2. Provide the necessary values in the respective fields as described in the table.
   For example, if you select **Local** in the drop-down menu, then the Create Local Gateway dialog box opens. Provide the necessary values in the respective fields as described in the table.

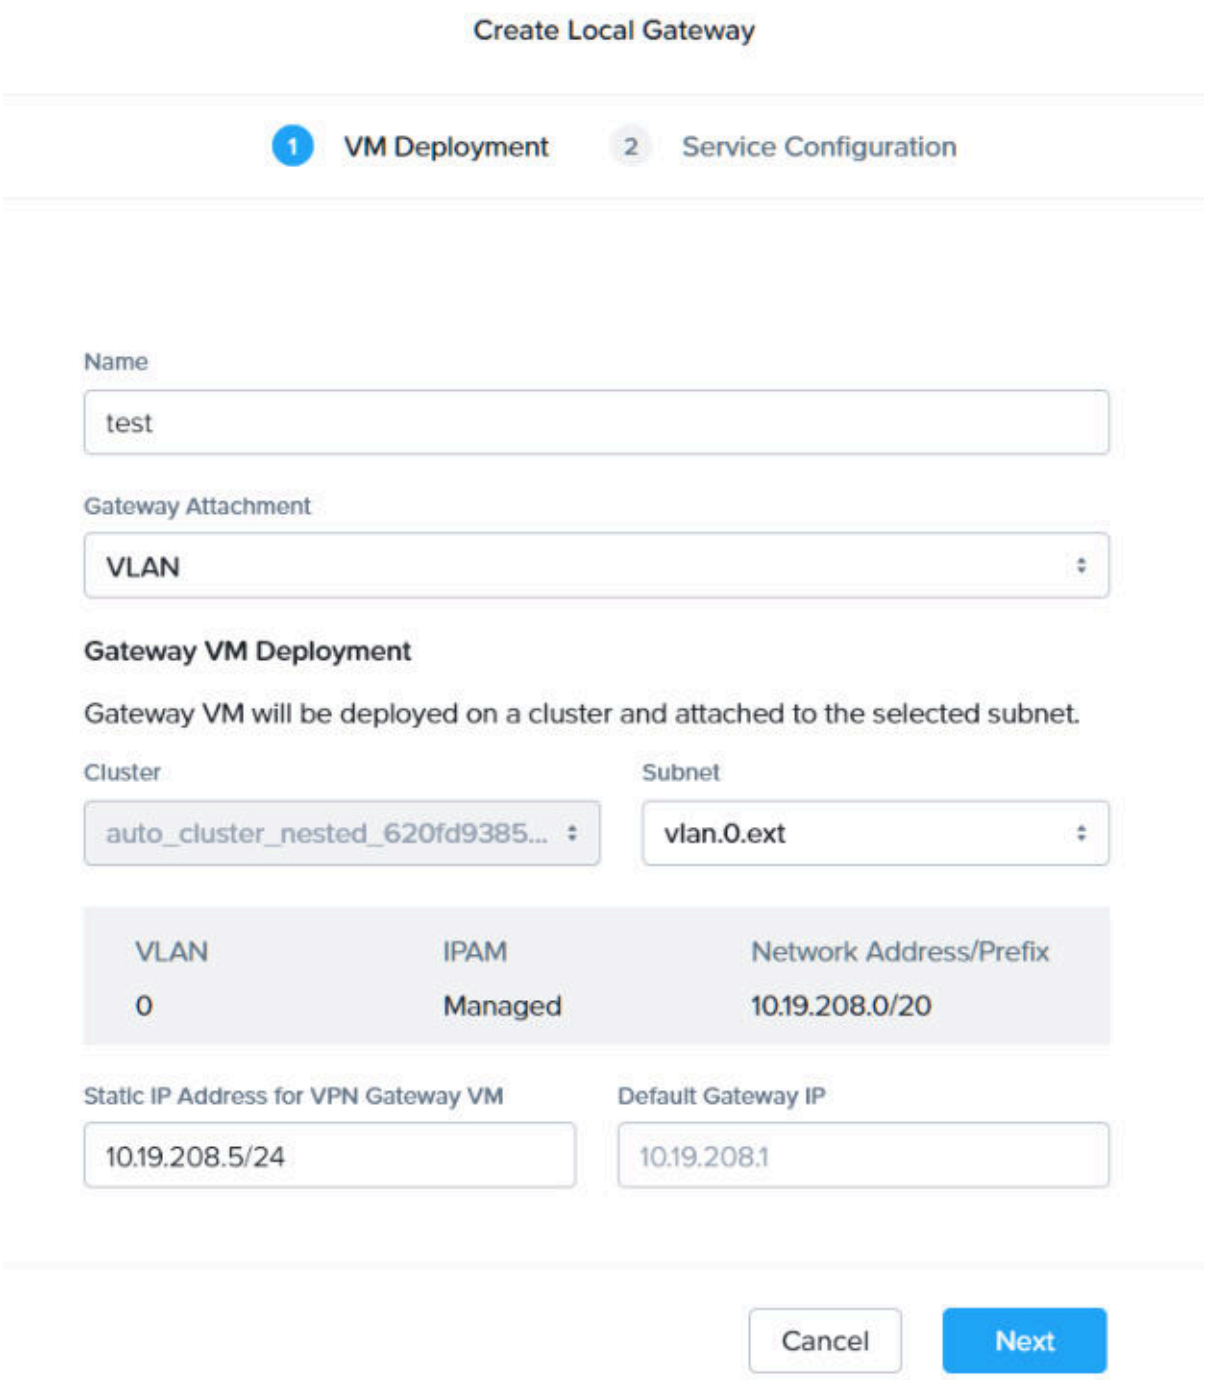

Figure 52: Sample Create Local Gateway - VM Deployment

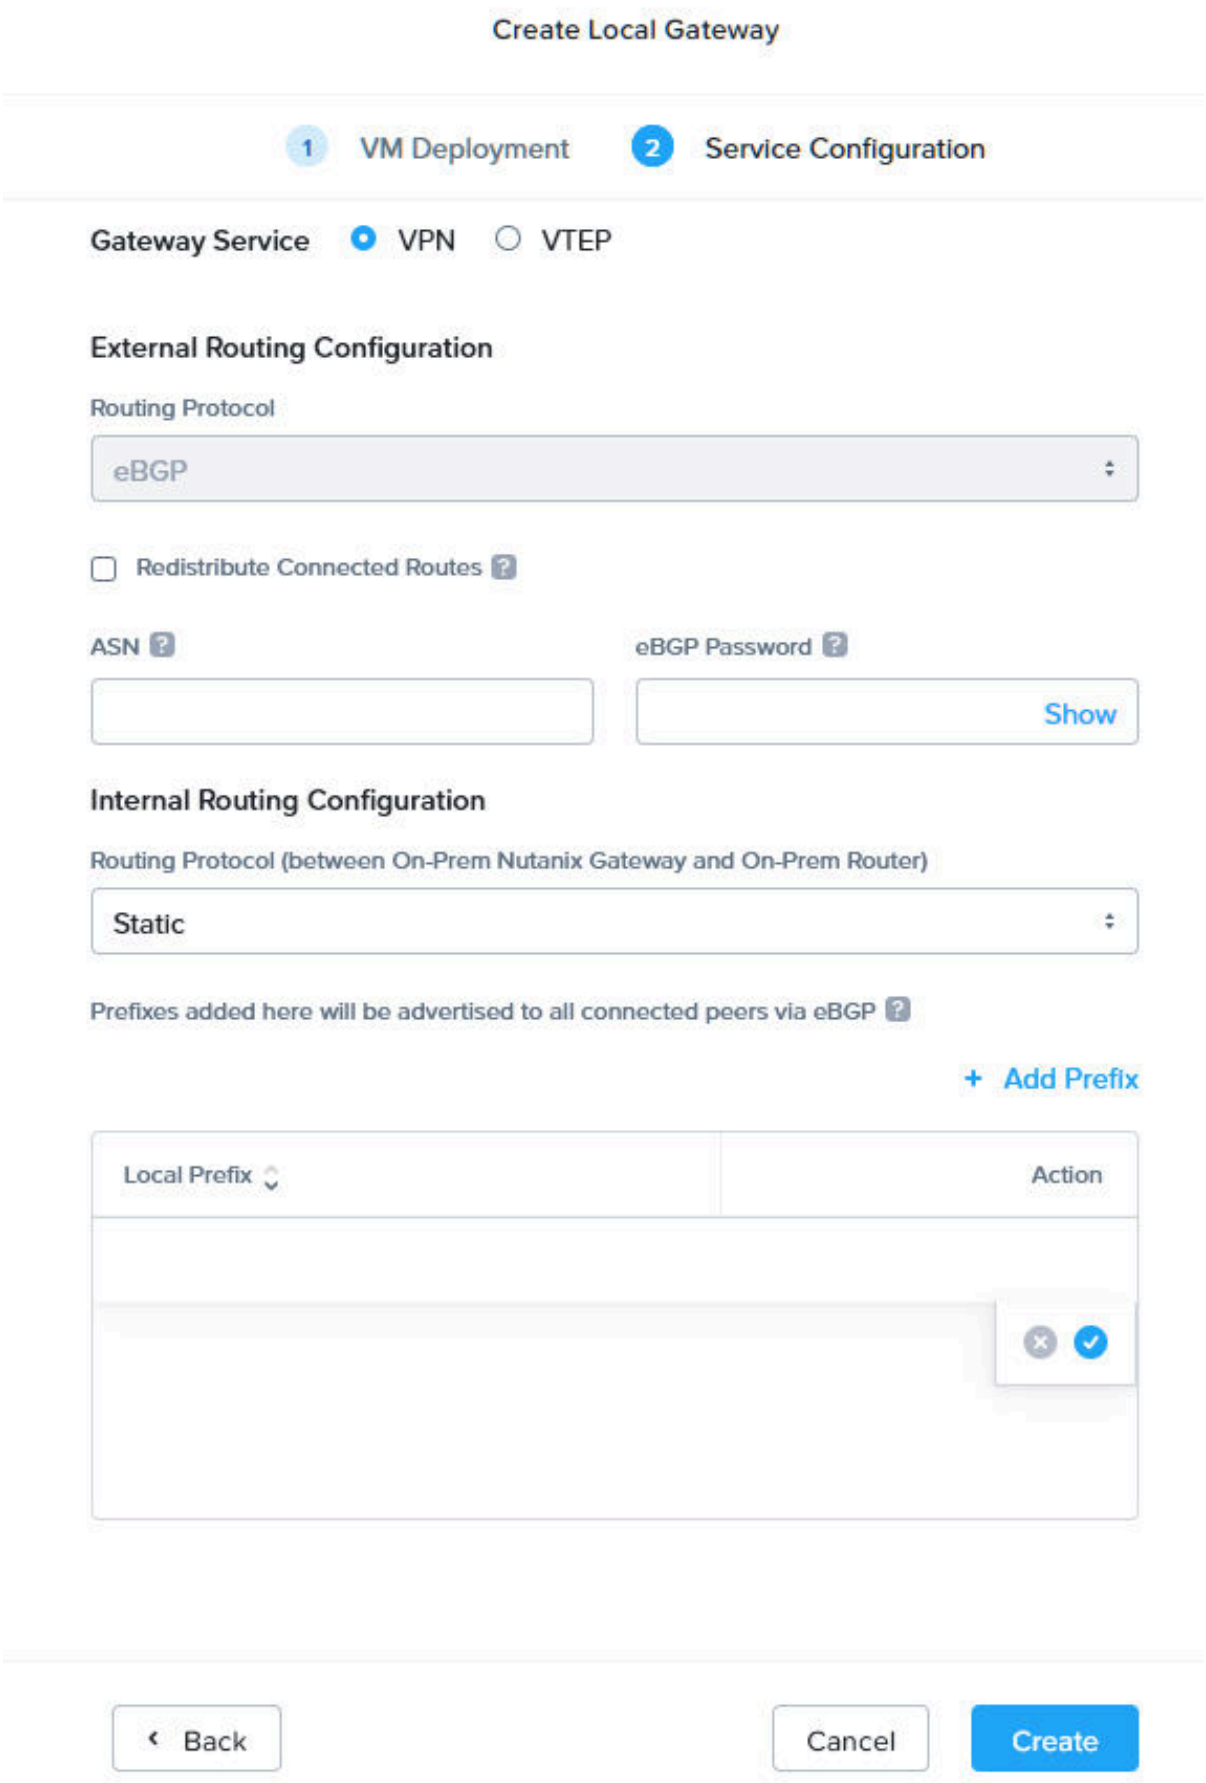

Figure 53: Sample Create Local Gateway - VPN Service Configuration

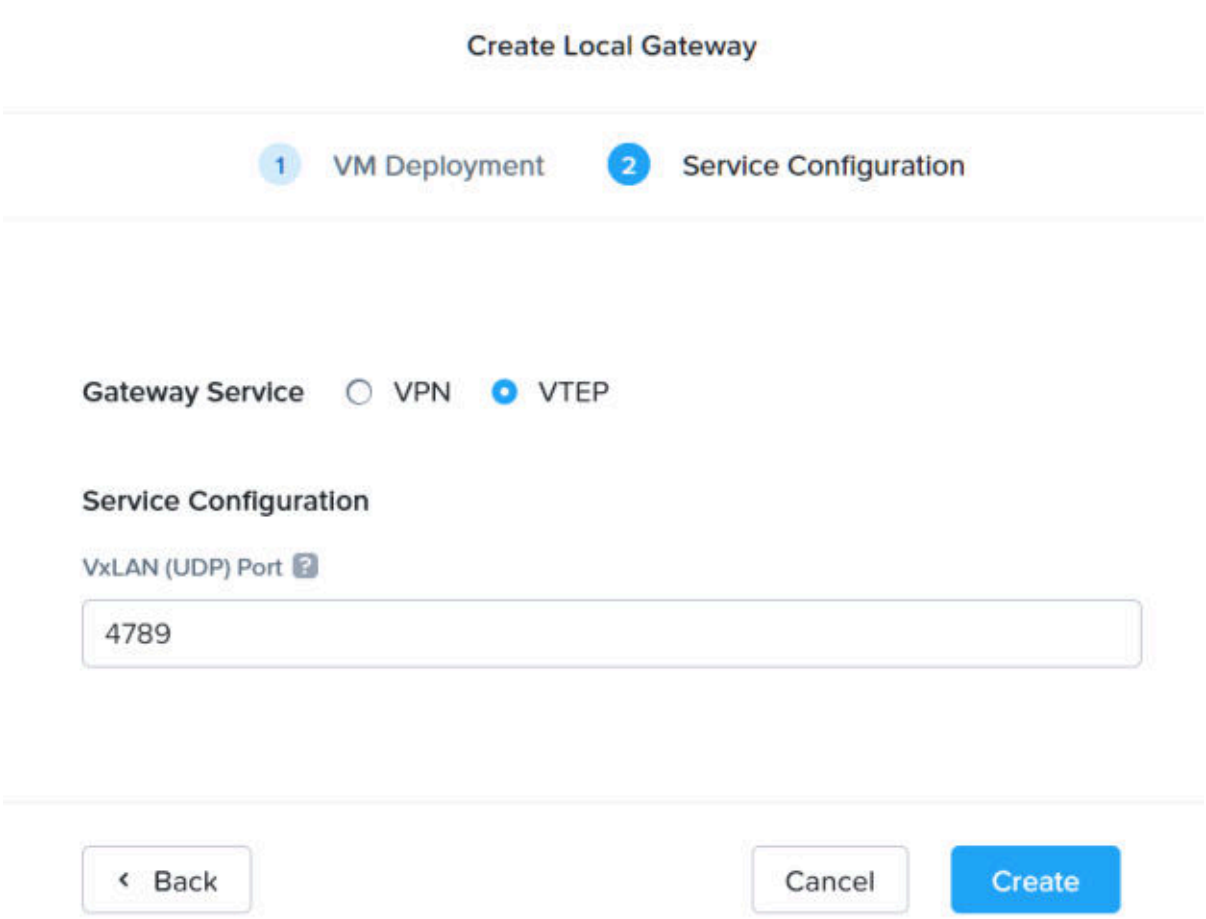

Figure 54: Sample Create Local Gateway - VTEP Service Configuration

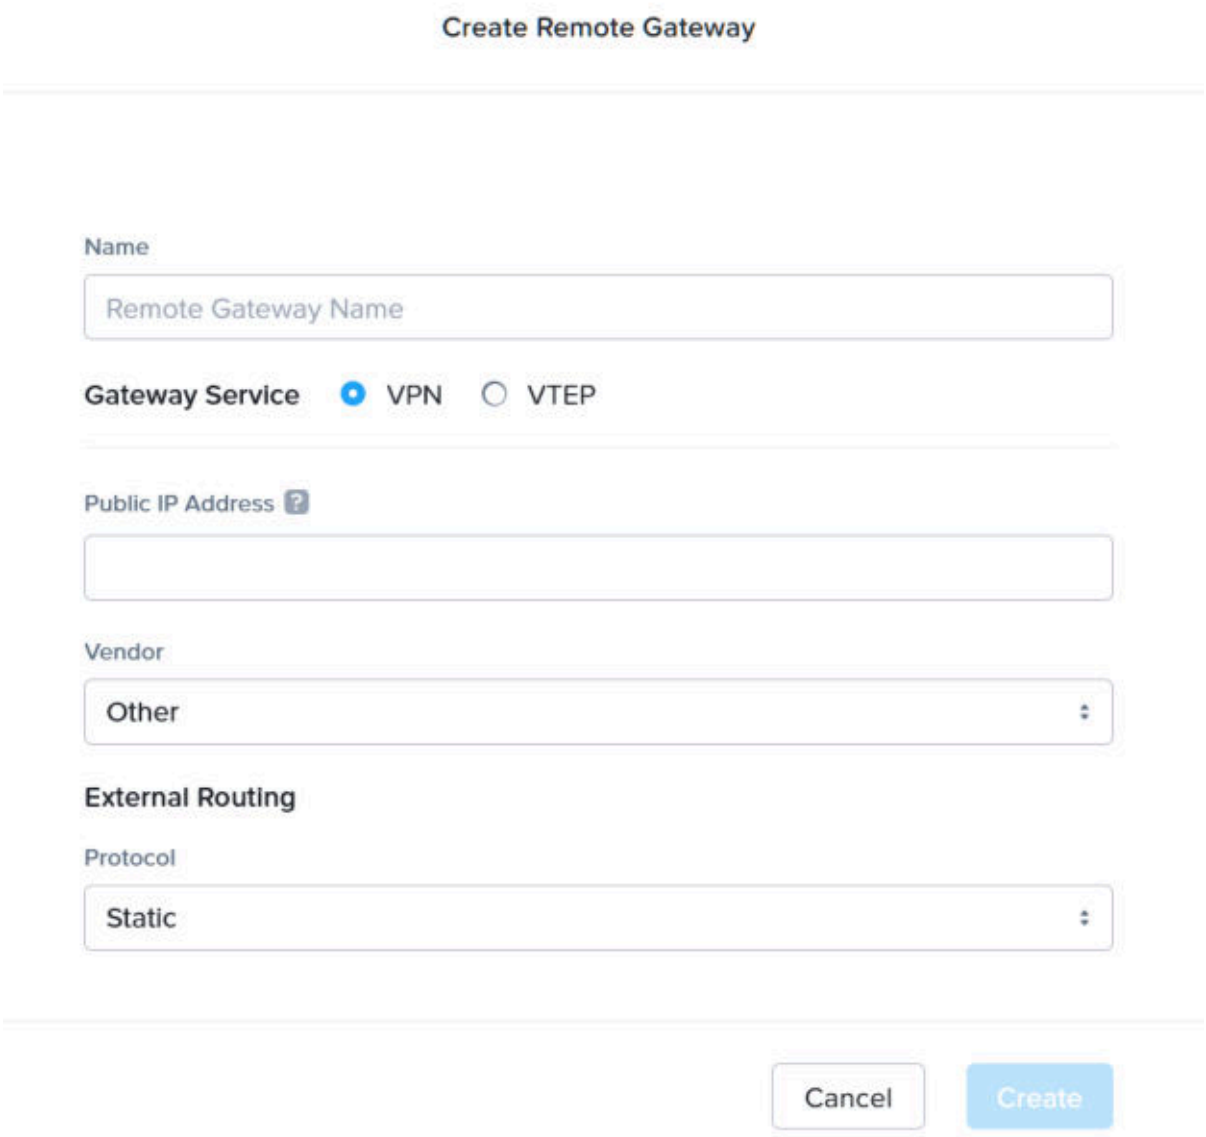

Figure 55: Sample Create Remote Gateway - VPN Gateway Service

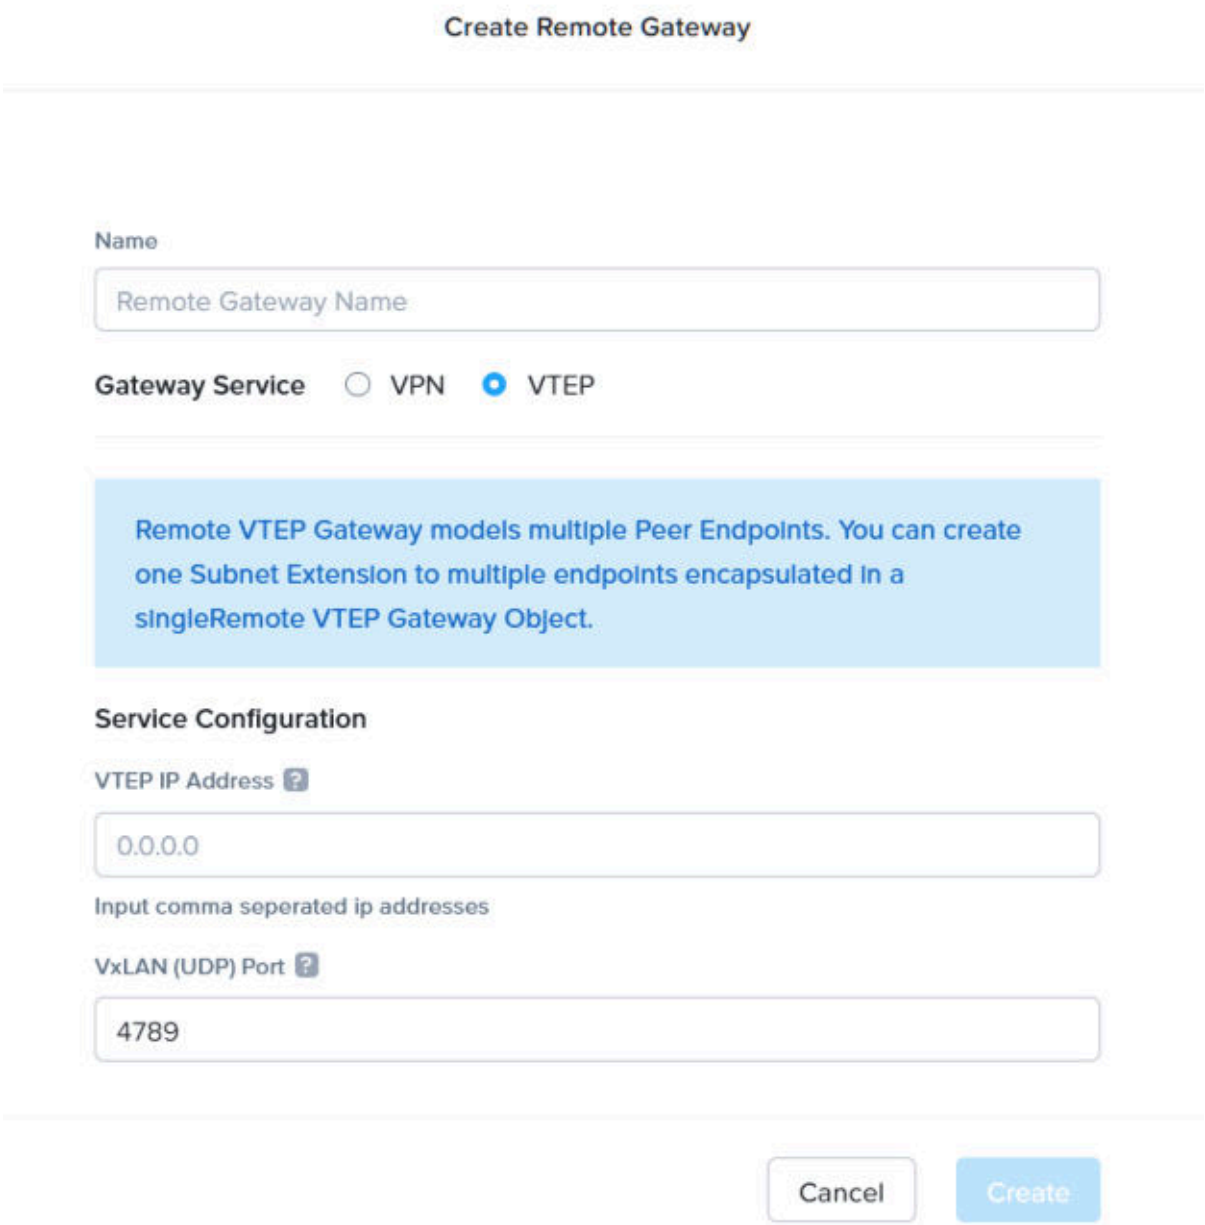

Figure 56: Sample Create Remote Gateway - VTEP Gateway Service

Table 12: Local Gateway Fields

Fields	Description	Values
VM Deployment		
Name	Enter a name for the network gateway.	(Name)

Fields	Description	Values
Gateway Attachments	(for Local gateway type only) Select the gateway attachment as **VPC** or **VLAN**. The VPN VM is deployed on a VPC VM or a cluster that has the selected VLAN respectively. 1. If you select **VPC**, then ***VPC Attachment*** is displayed. **VPC** is the default value for the ***Gateway Attachments*** field. The Gateway VM is deployed on the cluster and associated with the VPC selected in the ***VPC Attachment*** section. ***VPC attachment*** mode provides the options of eBGP and Static routing methods for external routing (configured in the **External Routing Configuration** section). 2. If you select **VLAN**, then the ***VLAN Attachment*** is displayed. The Gateway VM is deployed on the cluster that has the VLAN and the subnet specified in the ***VLAN Attachment*** section. ***VLAN attachment*** mode provides only the eBGP routing method for external routing.	(VLAN or VPC)
Gateway VM Deployment - VPC Attachment		
Cluster	Select the cluster on which you want to deploy the Gateway VM on.	(Name of the cluster)
VPC (If Gateway Attachment type is VPC)	Select the VPC configured on the selected cluster that you want to use for the Gateway VM deployment.	(Name of the VPC selected)
Floating IP (Optional)	Select a floating IP for the network gateway configuration. If you do not select a floating IP address then Prism Central allocates a floating IP automatically. This allocated floating IP is deleted when you delete the gateway. To request floating IPs and allocate them to subnets, see Requesting Floating IPs on page 66	(IP address)
Gateway VM Deployment - VLAN Attachment		
Cluster	Select the Cluster, from the drop down list, on which you want to deploy the Gateway VM on. Note: Only clusters with VLANs are available in the list.	(Name of the cluster)

Fields	Description	Values
Subnet	Select the subnet you want to attach the Gateway VM to, from the drop down list. Note: The list includes all the subnets you created on the selected cluster. After you select the subnet, the details of the subnet are displayed in a box below the **Subnet**field. The details include: VLAN ID, IPAM type being Managed or Unmanaged, and Network Address with Prefix.	(Name of the VLAN subnet)
Static IP Address for VPN Gateway VM	Enter the static IP address that the Gateway VM needs to use.	(IP Address with Prefix)
Default Gateway IP	Enter the default gateway IP of the subnet for the Gateway VM.	(IP Address)
Service Configuration		
Gateway Service	Select the gateway service you want to use for the gateway.	(VPN or VTEP)
VPN Service Configuration - External Routing Configuration (This section is available for VLAN and VPC attachment types)		
Routing Protocol	1. For **VPC** gateway attachments: Select **Static** for static routing. Note: You need to create static routes (see Creating Static Routes on page 83) for external routing and attach the route to the VPC selected in this configuration. 2. Select **eBGP** for eBGP based external routing. 3. For **VLAN** gateway attachments: External routing protocol is pre-set to **eBGP**. You cannot change the routing protocol.	(Static or eBGP)
Redistribute Connected Routes (Applicable only if VLAN type gateway attachment is selected)	(**VLAN** only) Select this checkbox to enable the redistribution of connected routes into the eBGP.	(Check mark or blank)

Fields	Description	Values
ASN (Only available if eBGP routing protocol is selected)	(For **eBGP** only) Enter the ASN for your on-prem gateway. If you do not have a BGP environment in your on-prem site, you can choose any number. For example, you can choose a number in the 65000 range. Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs. ASN must be distinct in case of eBGP.	(Number)

eBGP Password	(For **eBGP** in **Local** gateway type only) Enter the eBGP password for the eBGP route.	(Password: The password must be between 1 and 80 characters.
		
		• Characters allowed for Pre-Shared Key for IPSec
		  • a-z
		  • A-Z
		  • 0-9
		  • ~ ! @ # % ^ & * ( ) _ - + = : ; { } [ ] | < > , . / ? $
		  • Password length: Minimum 1 and maximum 64 characters.
		• Characters allowed for BGP passwords
		  • a-z
		  • A-Z
		  • 0-9
		  • ~ ! @ # % ^ & * ( ) _ - + = : ; { } [ ] | < > , . / ? $
		  • Password length: Minimum 1 and maximum 80 characters.
		
		)

VPN Service Configuration - Internal Routing Configuration (This section is available for VLAN attachment type only.)		

Fields	Description	Values
Routing Protocol (Between On-prem Gateway and On-prem Router)	Select the **Routing Protocol** to be used between on-premises Nutanix gateway and on-premises router. You can select: • `Static`: Select this protocol to provide a static route configuration for the VLAN gateway. • `OSPF`: Select this protocol to provide an OSPF routing configuration for the VLAN gateway. • `iBGP`: Select this protocol to provide a iBGP route configuration for the VLAN gateway. Note: For iBGP, the ASN must be the same between the Gateway appliance and the peer iBGP, when iBGP is selected as the internal routing protocol.	(Static or OSPF or iBGP)
+Add Prefix (Applicable to Static routing)	(For `Static` routing selected in **Routing Protocol**) Click this to enter a `Local Prefix` and click the check mark under **Actions** to add the prefix. If you click the X mark under **Actions**, the local prefix you entered is not added. The prefixes you add are advertised to all the connected peers via eBGP. The prefix must be a valid IP address with the host bits not set. You can add multiple local prefix IP addresses.	(prefix like /24)
Area ID (Applicable to OSPF protocol)	(OSPF only) Enter the OSPF area id in the IPv4 address format.	
Password Type	(OSPF only) Select the password type you want to set for the OSPF route. The options are: 1. `MD5`: Select this option to encrypt the packets with MD5 hash that can be decrypted with the MD5 password at the destination. 2. `Plain Text`: Select this option to set a clear-text password. 3. `None`: Select this if you do to set an open route without password protection	

Fields	Description	Values
Password	(OSPF only) Enter a password for the MD5 or Plain Text password type you select in the ***Password Type*** field. • For **MD5**: The password must be 1-16 characters long. Characters allowed for OSPF passwords (MD5) • a-z • A-Z • 0-9 • For **Plain Text**: The password must be 1-8 characters long. Characters allowed for OSPF passwords (Plain text): a-z.	
Peer IP (for iBGP)	Enter the IP Address of the On-prem router used to exchange routes with the network gateway.	(IP Address)
Password	Enter a password with 1-80 characters.	(Password)
VTEP Service Configurations		
VxLAN (UDP) Port	The default value provided is 4789. Do not change this.	(Number. Default value is 4789)

Table 13: Remote Gateway Fields

Fields	Description	Values
Name	Enter a name for the network gateway.	(Name)
Gateway Service	Select the gateway service you want to use for the gateway.	(VPN or VTEP)
VPN Service Configurations		
Public IP Address	Enter the public IP address of the remote endpoint. If a Floating IP is not selected, a new Floating IP is automatically allocated for the Gateway. These allocated IP addresses are deleted when the network gateway is deleted.	(IP Address)
Vendor	Select the vendor of the third party gateway appliance.	(Name of Vendor)
External Routing		

Fields	Description	Values
Protocol	1. Select **Static** for static routing. Note: You need to create static routes (see Creating Static Routes on page 83) for external routing and attach the route to the VPC selected in this configuration. 2. Select **eBGP** for eBGP based external routing.	(Static or eBGP)
eBGP ASN (Only available if eBGP routing protocol is selected)	(For **eBGP** only) Enter the ASN for your on-prem gateway. If you do not have a BGP environment in your on-prem site, you can choose any number. For example, you can choose a number in the 1-65000 range. Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs. ASN must be distinct in case of eBGP.	(Number)
VTEP Service Configurations		
VTEP IP Address	Enter VTEP IP Addresses of the remote endpoints that you want to create the gateway for. You can add IP addresses of multiple endpoints in one remote gateway.	(Comma separated list of IP Addresses)
VxLAN (UDP) Port	The default value provided is 4789. Do not change this.	(Number. Default value is 4789)

3. Click **Save**.

What to do next

The Gateway you create is displayed in the Gateways page.

Updating a Network Gateway

About this task

You can update a network gateway using the Update Gateway dialog box.

You can open the Update Gateway dialog box. The parameters in the Update Gateway dialog box are the same as those in the Create Local Gateway or Create Remote Gateway dialog box.

Procedure

1. Select the gateway you want to update on Gateways.
2. Click **Update** in the **Actions** menu.
3. Update the required details in the Update Gateway dialog box.
   You cannot modify some information. Such fields are greyed and in-actionable. If you need to modify such information, consider creating a new gateway with the updated parameters and deleting the current gateway.

4. Click **Save**.

Deleting a Network Gateway

About this task

If you want to delete a network gateway, you must first delete all the VPN connections associated with the gateway and only then you can delete the network gateway.

To delete a network gateway, do the following on the Gateway page.

Procedure

1. Do one of the following.
   - » Select the check box next to the name of the gateway and, in the **Actions** drop-down list, click **Delete**.
   - » Click the name of the gateway and, in the details page, click **Delete**.
2. In the confirmation dialog box, do the following.
   - » Click **Delete** to delete the entity.
   - » Click **Cancel** to exit without deleting the entity.

Creating a VPN Connection

About this task

Create a VPN connection to establish a VPN IPSec tunnel between VPN gateways in your on-prem site. Select the gateways between which you want to create the VPN connection.

To create a VPN connection, do the following on the **Networking** > **VPN Connections** page.

Procedure

1. Click the **Create VPN Connection** button on the VPN Connections page.
2. In the Create VPN Connection dialog box, provide the values in the respective fields.

Fields	Description and Values
Name	Enter a name for the connection.
VPN Connection	
IPSec Secret	Enter a secret password for the IPSec connection. To see the password, click **Show**. To hide the password, click **Hide**.
Local Gateway	Select the connection parameters on the local gateway as Initiator or Acceptor of VPN Tunnel connections.
VPN Gateway	Select the appropriate VPN Gateway as the local gateway for the VPN connection

Fields	Description and Values
VTI Prefix - Local Gateway	Enter a IPv4 Address with /<prefix>. Example: 10.25.25.2/30. This is the VPN Tunnel Interface IP address with prefix for the local gateway. The subnet for this IP address must be a /30 subnet with two usable IP addresses. One of the IP addresses is used for Local Gateway. Use the other IP address for the Remote Gateway.
Connection Handshake	This defines the type of handshake that the connection must use. There are two types of connection handshakes: 1. `Initiator`: The local VPN gateway acts as the initiator of the connection and thus initializes the VPN tunnel. 2. `Acceptor`: The local VPN gateway accepts or rejects incoming connection requests from other gateways. Note: In a VPN connection do not configure both the gateways (local gateway and remote gateway) in an endpoint as Initiators or as Acceptors. If you configure the local gateway as Initiator then configure the remote gateway as Acceptor in one endpoint and vice-versa in the (other) remote endpoint.
Remote Gateway	For a specific VPN connection, set the remote gateway as Initiator or Acceptor when you configure the VPN connection on the Remote Gateway.
VPN Gateway	Select the appropriate VPN Gateway as the remote gateway for the VPN connection.
VTI Prefix - Remote Gateway	The VPN Tunnel Interface IP address with prefix for the local gateway. Provide a IPv4 Address with /<prefix>. Example: 10.25.25.2/30. This is the VPN Tunnel Interface IP address with prefix for the local gateway. The subnet for this IP address must be a /30 subnet with two usable IP addresses. One of the IP addresses is used for Local Gateway. Use the other IP address for the Remote Gateway.
Advanced Settings	Set the traffic route priority for the VPN connection. The route priority uses Dynamic route priority because the priority is dependent on the routing protocol configured in the VPN gateway.
Route Priority - Dynamic Route Priority	Set the route priority as an integer number. The greater the number, higher is the priority.

3. Click **Save**.

What to do next

The VPN connection you create is displayed in the VPN Connections page.

Updating VPN Connection

About this task

You can update a VPN Connection using the Update VPN Connection dialog box.

You can open the Update VPN Connection dialog box. The parameters in the Update VPN Connectiondialog box are the same as those in the Create VPN Connection dialog box.

Procedure

1. Select the VPN Connection you want to update on the VPN Connection.
2. Click **Update** in the **Actions** menu.
3. Update the required details in the Update VPN Connection dialog box.
4. Click **Save**.

Deleting a VPN Connection

About this task

To delete a VPN connection, do the following on the VPN Connection page.

Procedure

1. Do one of the following.
   - » Select the check box next to the name of the VPN connection and, in the **Actions** drop-down list, click **Delete**.
   - » Click the name of the VPN connection and, in the details page, click **Delete**.
2. In the confirmation dialog box, do the following.
   - » Click **Delete** to delete the entity.
   - » Click **Cancel** to exit without deleting the entity.

Layer 2 Virtual Network Extension

You can extend a subnet between on-prem local and remote clusters or sites (Availability Zones or AZs) to support seamless application migration between these clusters or sites.

> Note: One or more on-prem cluster or sites managed by one Prism Central instance is defined as an Availability Zone or AZ. In this section, Availability Zone or AZ refers to and must be understood as one or more on-prem clusters or sites managed by one Prism Central. Local AZ refers to local on-prem clusters or sites managed by a Prism Central instance and remote AZ refers to another on-prem cluster or site managed by another Prism Central instance.

With Layer 2 subnet extension, you can migrate a set of applications to the remote AZ while retaining their network bindings such as IP address, MAC address, and default gateway. Since the subnet extension mechanism allows VMs to communicate over the same broadcast domain, it eliminates the need to re-architect the network topology, which could otherwise result in downtime.

Layer 2 extension assumes that there are underlying existing layer 3 connectivity already available between the Availability Zones. You can extend a subnet from a remote AZ to the primary (Local) AZ (and other remote AZs in case of VTEP-based subnet extensions)

- You can extend a Layer 2 subnet across two Nutanix AZs over either VPN or Virtual tunnel End Point (VTEP). SeeLayer 2 Virtual Subnet Extension Over VPN on page 113.
- You can extend a Layer 2 subnet between a Nutanix AZ and one or more non-Nutanix datacenters only over VTEP. See Layer 2 Virtual Subnet Extension Over VTEP on page 118.

You can extend subnets for the following configurations.

- IPAM Type. Managed and unmanaged networks.
- Subnet Type. On-prem VLAN subnets and VPC subnets.
- Traffic Type. IPv4 unicast traffic and ARP.
- On-prem Hypervisor. AHV and ESXi

> Note: If your cluster is ESXi, use vCenter Server to manually configure the port group attached to the subnet you want to extend. Set the security settings, **Promiscuous mode** and **Forged transmits** to **Accept** on the vSwitch as shown in the following image.

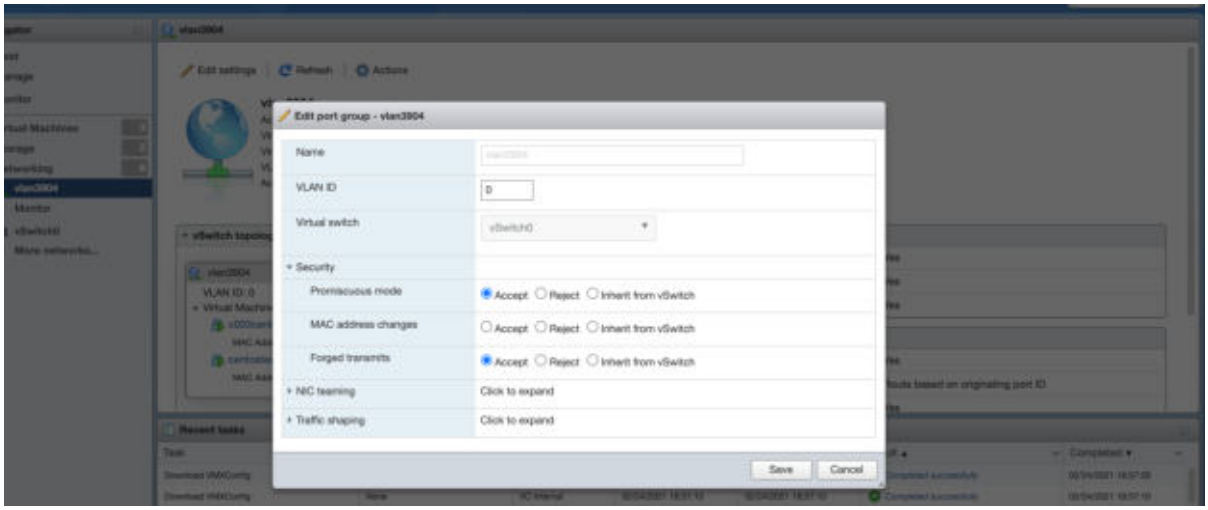

Figure 57: ESXi Host Port Group Configuration

Prerequisites for Setting Up Subnet Extension

Ensure the following before you configure Layer 2 subnet extension between your on-prem AZs.

- Ensure that the Prism Central versions support Layer 2 virtual subnet extension as specified in the Release Notes. See *AOS Family Release Notes* and *Release Notes | Prism Central* as applicable.

  See the *Prism Central Upgrade and Installation Guidelines and Requirements* section of the *Acropolis Upgrade Guide* for instructions about how to upgrade a Prism Central instance through the Prism Central web console.

- Ensure that you pair the Prism Central at the local AZ with the Prism Central at the remote AZ to use **Create Subnet Extension** wizard to extend a subnet across the AZs and facilitate bidirectional communication between these clusters or sites. Using paired availability zones it is possible to configure both VXLAN over VPN and VTEP based subnet extension. You can also extend subnets using the manual gateway and connection workflows instead of pairing the AZs.

  See the Pairing Availability Zones on page 114 for instructions about how to pair the local and remote AZs.

- Ensure that you set up a default static route with `0.0.0.0/0` prefix and the External Network next hop for the VPC you use for any subnet extension. This allows NTP and DNS access for the Network Gateway appliance.

Best Practices for Subnet Extension

Nutanix recommends the following configurations to allow IP address retention for VMs on extended subnets.

- When using Nutanix IPAM ensure the address ranges in the paired subnets are unique to avoid conflict between VM IP addresses across extended subnets.
- If the source and target sites use third-party IPAM, ensure that there are no conflicting IP address assignments across the two sites.

  > Note: If the source and target sites use Nutanix IPAM, the Prism Central web console displays a message that indicates an IP address conflict if one exists.

- If connectivity between sites already provides encryption, consider using VTEP only subnet extension to reduce encryption overhead.
- Use the **Subnet Extension to a Third Party Data-Center** workflow in the following scenarios
  - To extend a subnet to more than one other AZ. This is also known as point to multi-point.
  - To extend subnets between clusters managed by the same Prism Central.

Subnet Extension Workflow

You can manage Layer 2 subnet extension on the Subnet Extensions tab of the Connectivity page. Open the Subnet Extensions by clicking the hamburger icon in the top-left corner of the **Dashboard** and then clicking **Connectivity**.

- You can create point-to-point Layer 2 subnet extensions between two AZs over VPN or VTEP by opening the Create Subnet Extension Across Availability Zones dialog box. See Extending a Subnet Over VPN on page 116 for VPN-based extensions. See Extending a Subnet Across Availability Zones Over VTEP on page 119 for VTEP-based extensions.
- You can create point-to-point or point-to-multipoint Layer 2 subnet extensions to third party datacenters over VTEP by opening the Create Subnet Extension To A Third Party Data-Center dialog box. See Extending a Subnet to Third Party Datacenters Over VTEP on page 123.
- You can update a subnet extension that extends across AZs using the Update Subnet Extension Across Availability Zones dialog box. The Update Subnet Extension Across Availability Zones has the same parameters and fields as the Create Subnet Extension Across Availability Zones dialog box. You can open the Update Subnet Extension Across Availability Zones dialog box by:
  - Selecting the subnet extended across AZs in the Subnet Extensions and clicking the **Update** button.
  - Clicking the subnet extended across AZs in the Subnet Extensions and clicking the **Update** button on the **Summary** tab.

You can update a subnet extension that extends to multiple AZs or third party datacenters using the Update Subnet Extension To A Third Party Data-Center dialog box. Update Subnet Extension To A Third Party Data-Center dialog box has the same parameters and fields as the Create Subnet Extension To A Third Party Data-Center dialog box. You can open the Update Subnet Extension To A Third Party Data-Center dialog box by:

- Selecting the subnet extended to third datacenters in the Subnet Extensions and clicking the **Update** button.

- Clicking the subnet extended to third datacenters in the Subnet Extensions and clicking the **Update** button on the **Summary** tab.

See Updating an Extended Subnet on page 127.

Layer 2 Virtual Subnet Extension Over VPN

Subnet extension using VPN allows seamless, secure migration to a new datacenter or for disaster recovery. VPN based Layer 2 extension provides secure point to point connection to migrate workloads between Availability Zones. Consider VTEP-only Subnet Extension without VPN when encryption is not required.

Subnet extension using VPN is useful:

- When the two Availability Zones (where the subnets to be extended belong) do not have any underlying secure connectivity. For example, when connecting over the Internet, VPN (IPSec) provides the necessary connectivity and encryption (security).
- Sometimes when you need to move (lift-and-shift) workloads from a VLAN subnet to a VPC subnet retaining the same VM IP addresses . You need connectivity from other subnets to workloads that have already migrated to VPC. In such cases, VPN provides the Layer 3 connectivity and encryption between the VPC segment of extended subnet to other VLAN subnets.

Prerequisites for Setting Up Subnet Extension Over VPN

- See Layer 2 Virtual Network Extension on page 110 for general prerequisites to extend subnets.
- Set up VPN gateway services and a VPN connection between local AZ and the remote AZ. The subnet extension feature supports only the Nutanix VPN solution (not a third-party VPN solution) at the both the local and remote AZs. See the Virtual Network Connections on page 92 for instructions about how to upgrade the VPN gateway VM at the local and remote clusters or sites.

  > Note: Ensure that the VPN gateway version is 5.0 or higher. See the Updating a Network Gateway on page 107 section of the *Nutanix Flow Networking Guide* for instructions about how to upgrade the network gateway at the local and remote sites.

- Configure subnets with the same IP CIDR prefix at the source and target sites. For example, if the IP prefix at one site is 30.0.0.0/24, the IP prefix at the other site must also be 30.0.0.0/24. The network and mask must match at both AZs.
- Configure distinct DHCP pools for the source and target sites with no IP address overlap. Separate DHCP pools ensure no IP address conflicts occur for dynamically assigned IP addresses between the two AZs.
- Procure two free IP addresses, one from each subnet, for the Network Gateway in the subnets to be extended. These IP addresses are configured as local IP address and remote IP address for the subnet extension in the **Subnet Extension** wizard. These two free IP addresses are the externally accessible IP addresses for the local gateway, and the remote

gateway. Those two usable IP addresses are already contained inside the VPN connection and must not conflict with the following:

- DHCP pools on any of the Availability Zones.
- Gateway IP address on any of the Availability Zones.
- IP addresses allocated to existing user VMs on any of the Availability Zones.
- IP addresses used by Network Gateway Management NIC subnet (IP pool 100.64.1.0/24)
-

Limitation

To use subnet extension over a VPN, both sites must use the VPN service of the Nutanix Network Gateway. Consider VTEP-only subnet extension to connect to non-Nutanix third party sites.

Pairing Availability Zones

To extend the subnets between two on-prem Prism Central (availability zones) using VPN, pair the availability zones (AZs) with each other.

About this task

To pair one on-prem AZ with another on-prem AZ, perform the following procedure at both the AZs.

Procedure

1. Log on to Prism Central.
2. Click the hamburger icon at the top-left corner of the window. Go to **Administration** > **Availability Zones** in the left pane.

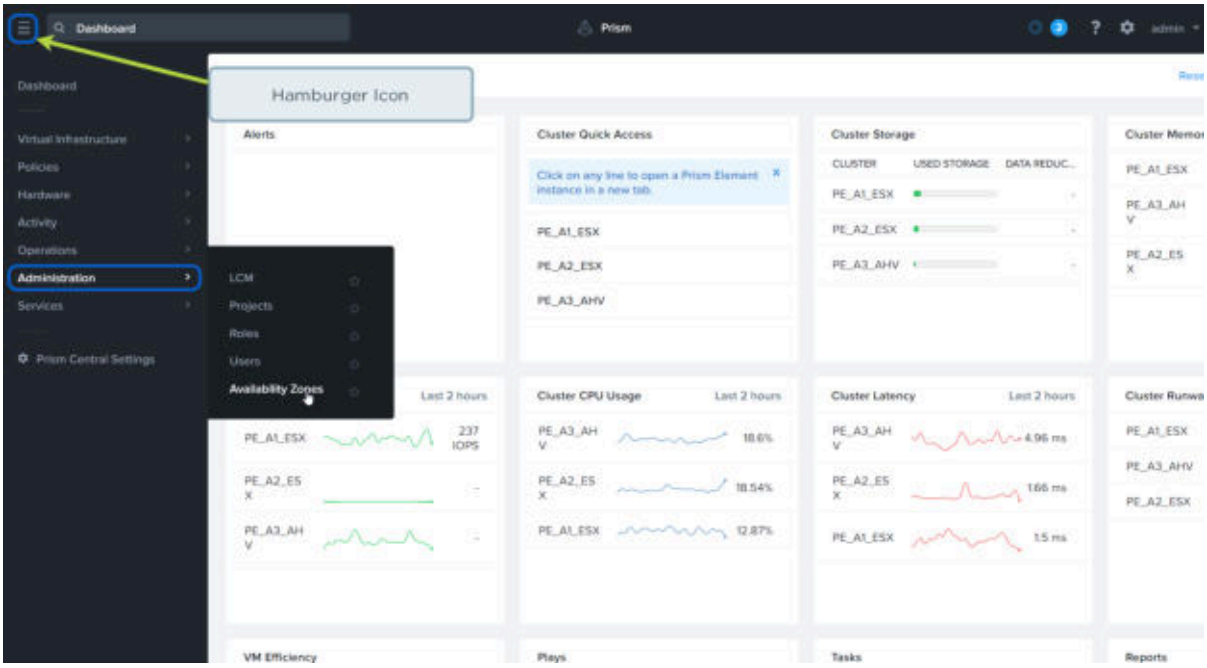

Figure 58: Pairing Availability Zone

3. Click **Connect to Availability Zone**.
   Specify the following information in the Connect to Availability Zone window.

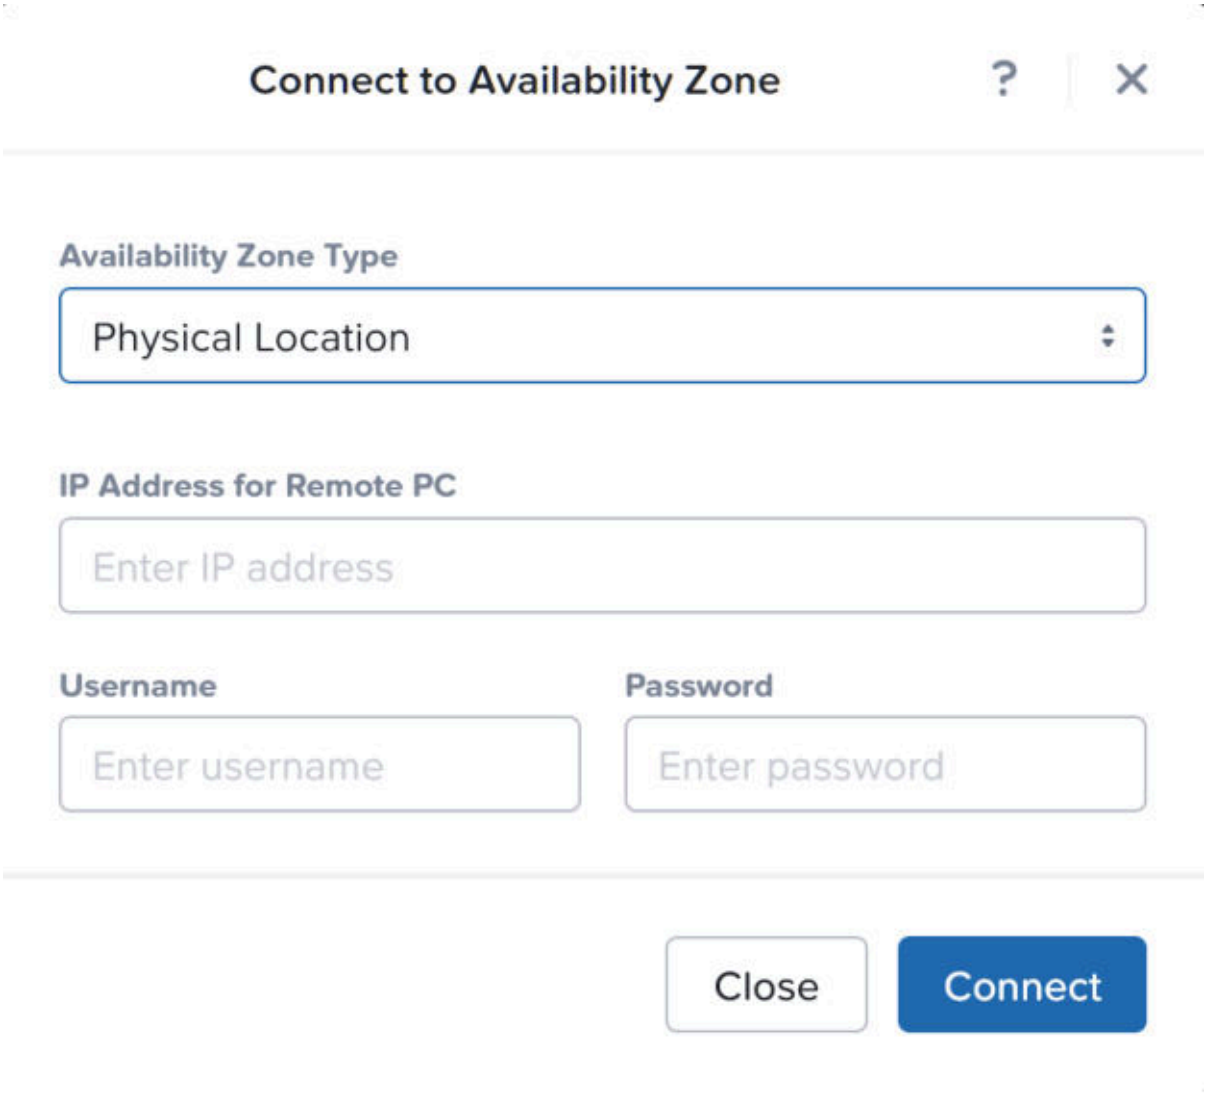

Figure 59: Connect to Availability Zone

   a. **Availability Zone Type**: Select **Physical Location** from the drop-down list.
      A physical location is an on-prem availability zone (site).
   b. **IP Address for Remote PC**: Enter the IP address of the recovery site Prism Central.
   c. **Username**: Enter the username of your recovery site Prism Central.
   d. **Password**: Enter the password of your recovery site Prism Central.
4. Click **Connect**.

What to do next

The connected AZs are listed in the Availability Zones dashboard with the name of the availability zone and connectivity status. Connectivity status is displayed as `Reachable` if the connection to the AZ is working.

Extending a Subnet Over VPN

The subnet extension allows VMs to communicate over the same broadcast domain to a remote site or Availability Zone (AZ).

Before you begin

See Layer 2 Virtual Network Extension on page 110 and Layer 2 Virtual Subnet Extension Over VPN on page 113 for information on prerequisites and best practices for extending a subnet.

About this task

Perform the following procedure to extend a subnet from the on-prem site.

Procedure

1. Click the hamburger icon in the top-left corner of the **Dashboard** > **Networking & Security** > **Connectivity** > **Subnet Extension**.
2. On the Subnet Extension page, select **Create Subnet Extension** > **Across Availability Zones** .

3. In the Create Subnet Extension Across Availability Zones dialog box, enter the necessary details as described in the table.

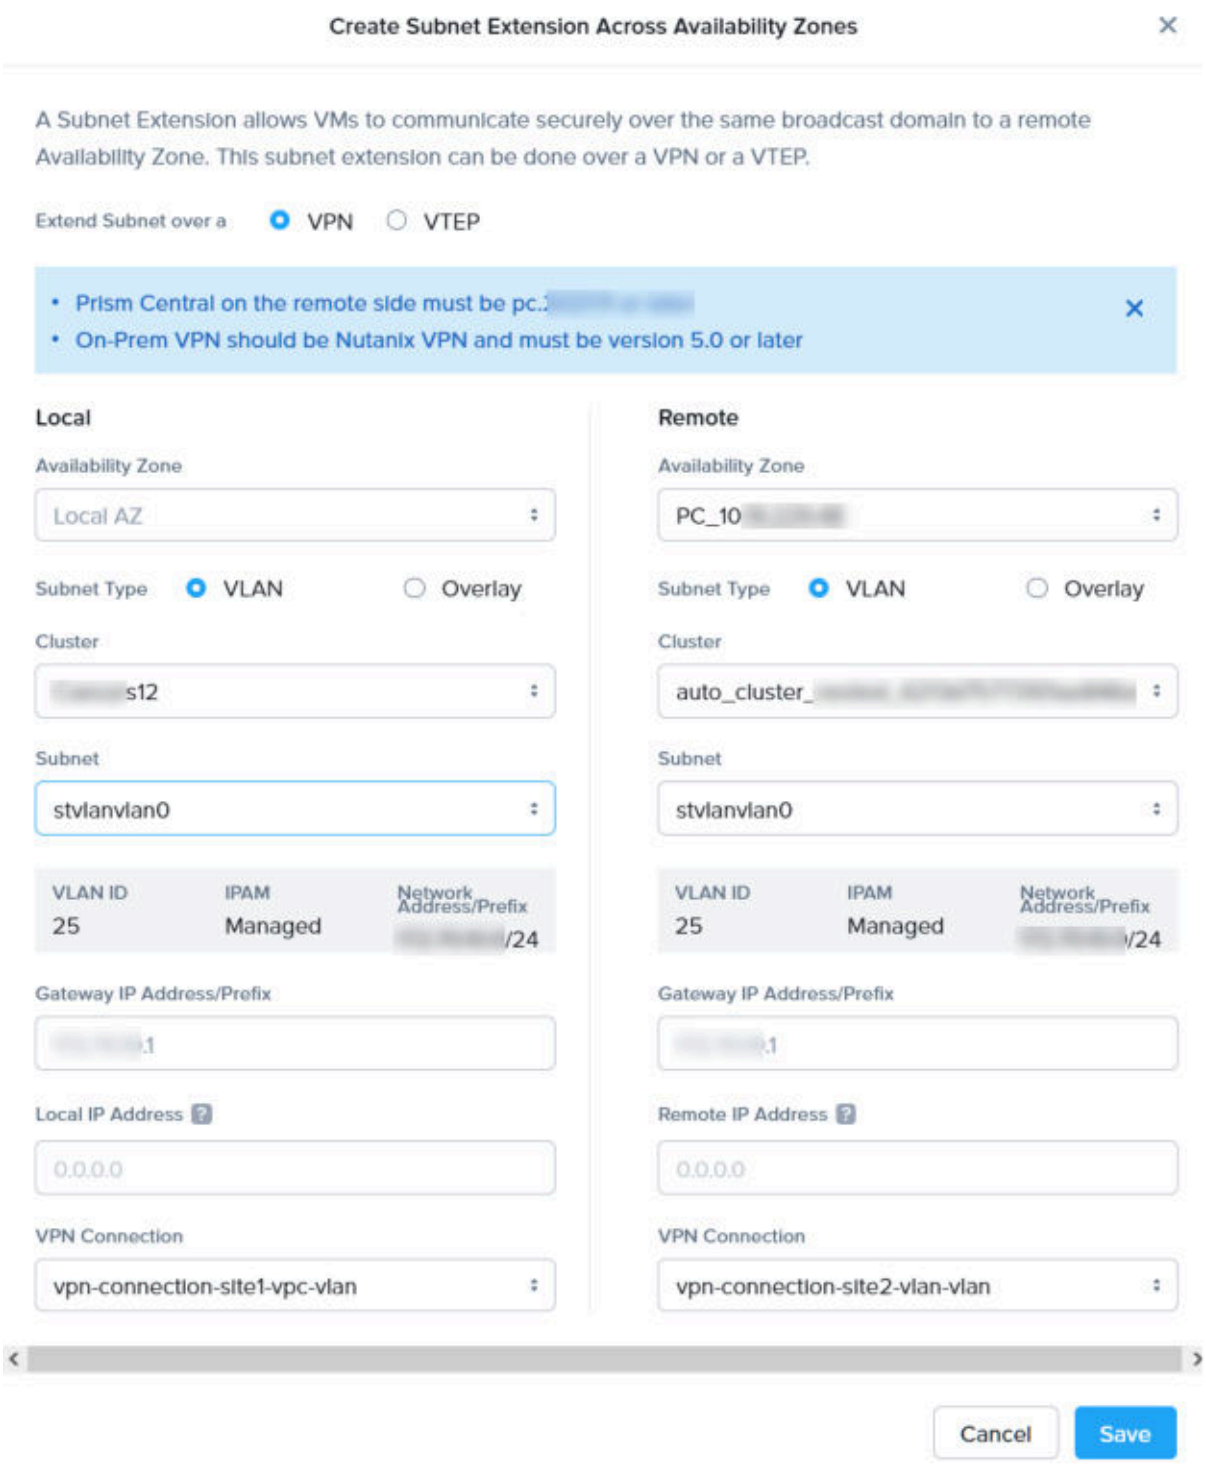

Figure 60: Create Subnet Extension Across Availability Zones

Fields	Description	Values
Extend Subnet over a	Select the gateway service you want to use for the subnet extension.	(VPN or VTEP)
Note: Configure the following fields for the **Local** and the **Remote** sides of the dialog box.		
Availability Zone	(For Local) Local AZ is pre-selected default. (For Remote) Select the appropriate AZ from the drop-down list of AZs.	(Local: Local AZ) (Remote: Dropdown list of AZs.)
Subnet Type	Select the type of subnet that you want to extend.	(VLAN or Overlay)
Cluster	Displayed if your selected VLAN subnet. Select the cluster from the dropdown list of clusters.	(Name of cluster selected from dropdown list)
VPC	Displayed if your selected Overlay subnet. Select the appropriate VPC from the dropdown list of VPCs.	(Name of VPC selected from dropdown list)
Subnet	Select the subnet that needs to be extended.	(Name of subnet selected from dropdown list)
(Network Information frame)	Displays the details of the VLAN or Overlay network that you selected in the preceding fields.	(Network information)
Gateway IP Address/ Prefix	Displays the gateway IP address for the subnet. This field is already populated based on the subnet selected.	(IP Address)
(Local or Remote) IP Address	Enter a unique and available IP address that are externally accessible IP addresses in **Local IP Address** and **Remote IP Address**.	(IP Address)
VPN Connection	Select the appropriate VPN Connection from the dropdown list that Flow networking must use for the subnet extension. See Creating a VPN Connection on page 108 for instructions to create VPN connection.	(Name of VPN connection selected from the dropdown list)

4. Click **Save**.

   A successful subnet extension is listed on the **Subnet Extension** dashboard. See .

Layer 2 Virtual Subnet Extension Over VTEP

Subnet extension using Virtual tunnel End Point (VTEP) allows seamless migration to new datacenters or for disaster recovery. VTEP based Layer 2 extension provides point-to-multipoint connections to migrate workloads from one Availability Zone to multiple Availability

Zones without encryption. If you need security and encryption, consider using Subnet Extension over VPN.

Subnet extension using VTEP is useful:

- When both subnets that need to be stretched are Nutanix subnets (managed or unmanaged). VTEP provides an optimized workflow to stretch the two subnets.
- When both subnets are connected over an existing private and secure link that does not need additional encryption.
- When one Nutanix subnet needs to be stretched across one or more non-Nutanix networks, sites, or datacenters. Subnet Extension with third-party VTEPs provides point-to-multipoint connectivity to third party datacenters assuming that there is underlying layer 3 connectivity between these VTEPs.

VTEP-based Layer 2 Subnet Extension provides the following advantages:

- Layer 2 subnet extension from one AZ to multiple AZs.
- Layer 2 subnet extension between Nutanix AZs and non-Nutanix third party VTEP-based AZs.
- The Remote VTEP Gateway is a set of endpoint IP addresses. You can add endpoint IP addresses to an existing operational Remote VTEP Gateway without stopping the subnet extension services. This on-the-fly addition enables you to extend the subnets to more AZs than originally planned, or perform maintenance, without disrupting the running services or configuring new remote VTEP gateways.

Prerequisite for Setting Up Subnet Extension Over VTEP

- See Layer 2 Virtual Network Extension on page 110 for general prerequisites to extend subnets.
- Set up VTEP local and remote gateway services on local and remote AZs. In case of point-to-multipoint extension, ensure that you create local and remote VTEP gateways on all the remote AZs that the subnet needs to be extended to.
- For each extended subnet within the same Network Gateway appliance ensure that you have unique VxLAN Network Identifiers (VNIs) that you can use for the VTEP subnet extensions. VNI may be any number between 0 and 16777215.

Extending a Subnet Across Availability Zones Over VTEP

The subnet extension over VTEP allows VMs to communicate two Availability Zones (AZ) without a VPN connection.

Before you begin

See Layer 2 Virtual Network Extension on page 110 and Layer 2 Virtual Subnet Extension Over VTEP on page 118 for information on prerequisites and best practices for extending a subnet.

About this task

To extend a subnet over VTEP across two availability zones (AZs), do the following.

Procedure

1. Open the Create Subnet Extension Across Availability Zones in one of the following ways:

   » On the Subnet Extensions tab, click > **Create Subnet Extension** > **Across Availability Zones** > **.**

   » In the Subnets dashboard, select the subnet you want to extend and click **Actions** > **Extend** > **Across Availability Zones**

   » In the Subnets dashboard, click the subnet you want to extend. On the subnet details page, click **Extend** > **Across Availability Zones**.

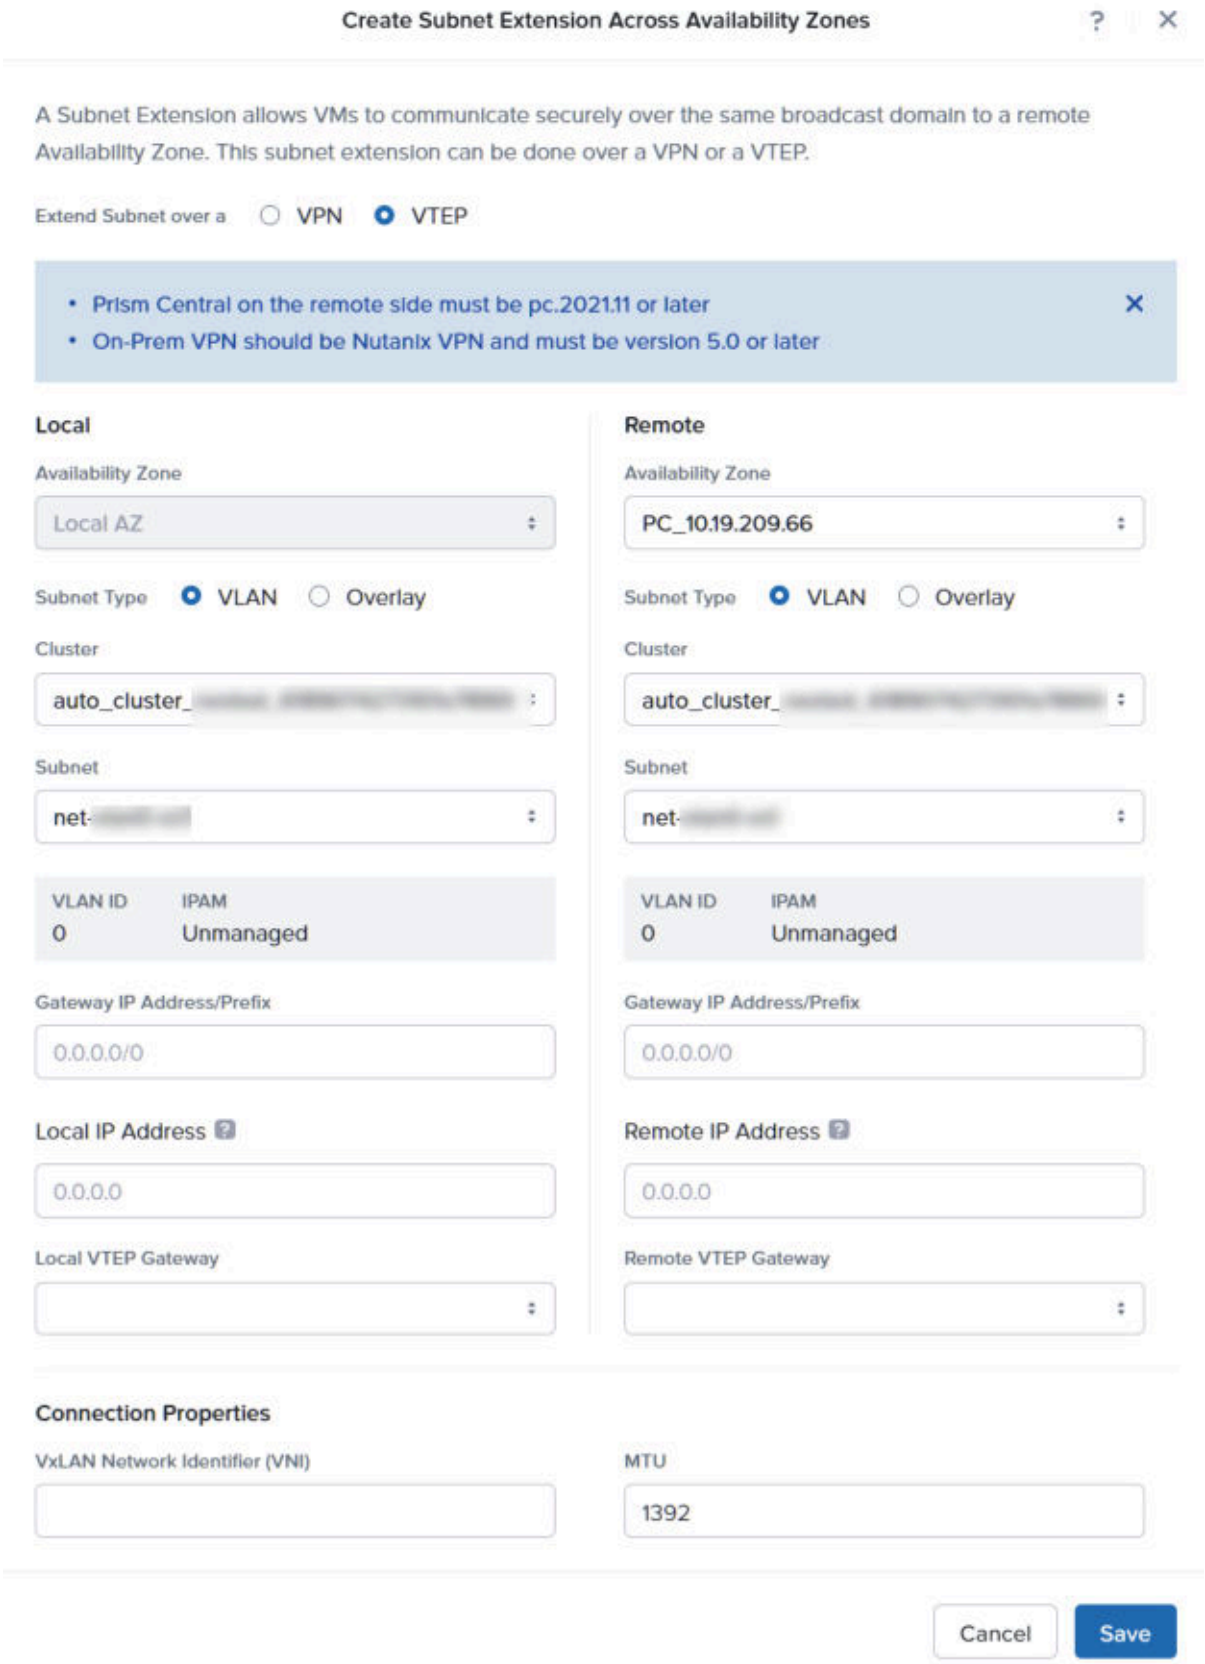

Figure 61: Example of Create VTEP Extension Across AZs with VLAN Subnet

2. For **Extend Subnet over a**, select **VTEP**.

3. Enter or select the necessary values for the parameters in the **Local** and **Remote** (AZ) sections as described in the table.

Parameters	Description and Value
Availability Zone	Displays the name of the paired availability zone at the local AZ.
Subnet Type	Select the type of the subnet - VLAN or Overlay that you are extending.
Cluster	Select the name of the cluster in the local AZ that the subnet is configured for.
Subnet	Select the name of the subnet at the local AZ for network. The VLAN ID and the IPAM - managed or unmanaged are displayed in the box below the **Subnet** field.
Gateway IP Address.	Enter the gateway IP address of the subnet you want to extend. Ensure that you provide the IP address in <IP-address/network-prefix> format. for example the gateway IP is 10.20.20.1 in a /24 subnet then provide the gateway IP address as `10.20.20.1/24`. Note: For an unmanaged network, enter the gateway IP address of the created subnet.
Local IP Address	Enter a unique and available (unused) IP address from the subnet provided in **Subnet** for the Network Gateway appliance.
Remote IP Address	Enter a unique and available (unused) IP address from the subnet provided in **Subnet** for the remote Network Gateway appliance.
Local VTEP Gateway	Select the local VTEP gateway you created on the local AZ. See Creating a Network Gateway on page 96 for information about creating VTEP gateways.
Remote VTEP Gateway	Select the VTEP gateway you created on the remote AZ. See Creating a Network Gateway on page 96 for information about creating VTEP gateways.
Connection Properties	
VxLAN Network Identifier (VNI)	Enter a unique number from the range 0-16777215 as VNI. Ensure that this number is not reused anywhere in the local or remote VTEP Gateways.
MTU	The default MTU is 1392 to account for 108 bytes of overhead and the standard physical MTU of 1500 bytes. VPC Geneve encapsulation requires 58 bytes and VXLAN encapsulation requires 50. However, you can enter any valid MTU value for the network, taking this overhead into account. For example, if the physical network MTU and vs0 MTU are 1600 bytes, the Network Gateway MTU can be set to 1492 to account for 108 bytes of overhead. Ensure that the MTU value does not exceed the MTU of the AHV Host interface and all the network interfaces between the local and remote AZs.

4. Click **Save**.
   After the subnet is extended, the extension appears in the Subnet Extensions list view.

Extending a Subnet to Third Party Datacenters Over VTEP

The subnet extension over VTEP allows VMs to communicate with multiple remote sites or Availability Zones (AZ) that may be third party (non-Nutanix) networks, or datacenters. It also provides the flexibility of adding more remote AZs to the same VTEP-based extended Layer 2 subnet. Examples of compatible VTEP gateways are switches from Cisco, Juniper, Arista, and others that support plain VXLAN VTEP termination.

About this task

To extend a subnet over VTEP across multiple availability zones (AZs) or third party datacenters, do the following.

Procedure

1. Open the Create Subnet Extension To A Third Party Data-Center in one of the following ways:

   » On the Subnet Extensions tab, click > **Create Subnet Extension** > **To A Third Party Data-Center**

   » In the Subnets dashboard, select the subnet you want to extend and click **Actions** > **Extend** > **To A Third Party Data-Center**

   » In the Subnets dashboard, click the subnet you want to extend. On the subnet details page, click **Extend** > **To A Third Party Data-Center**.

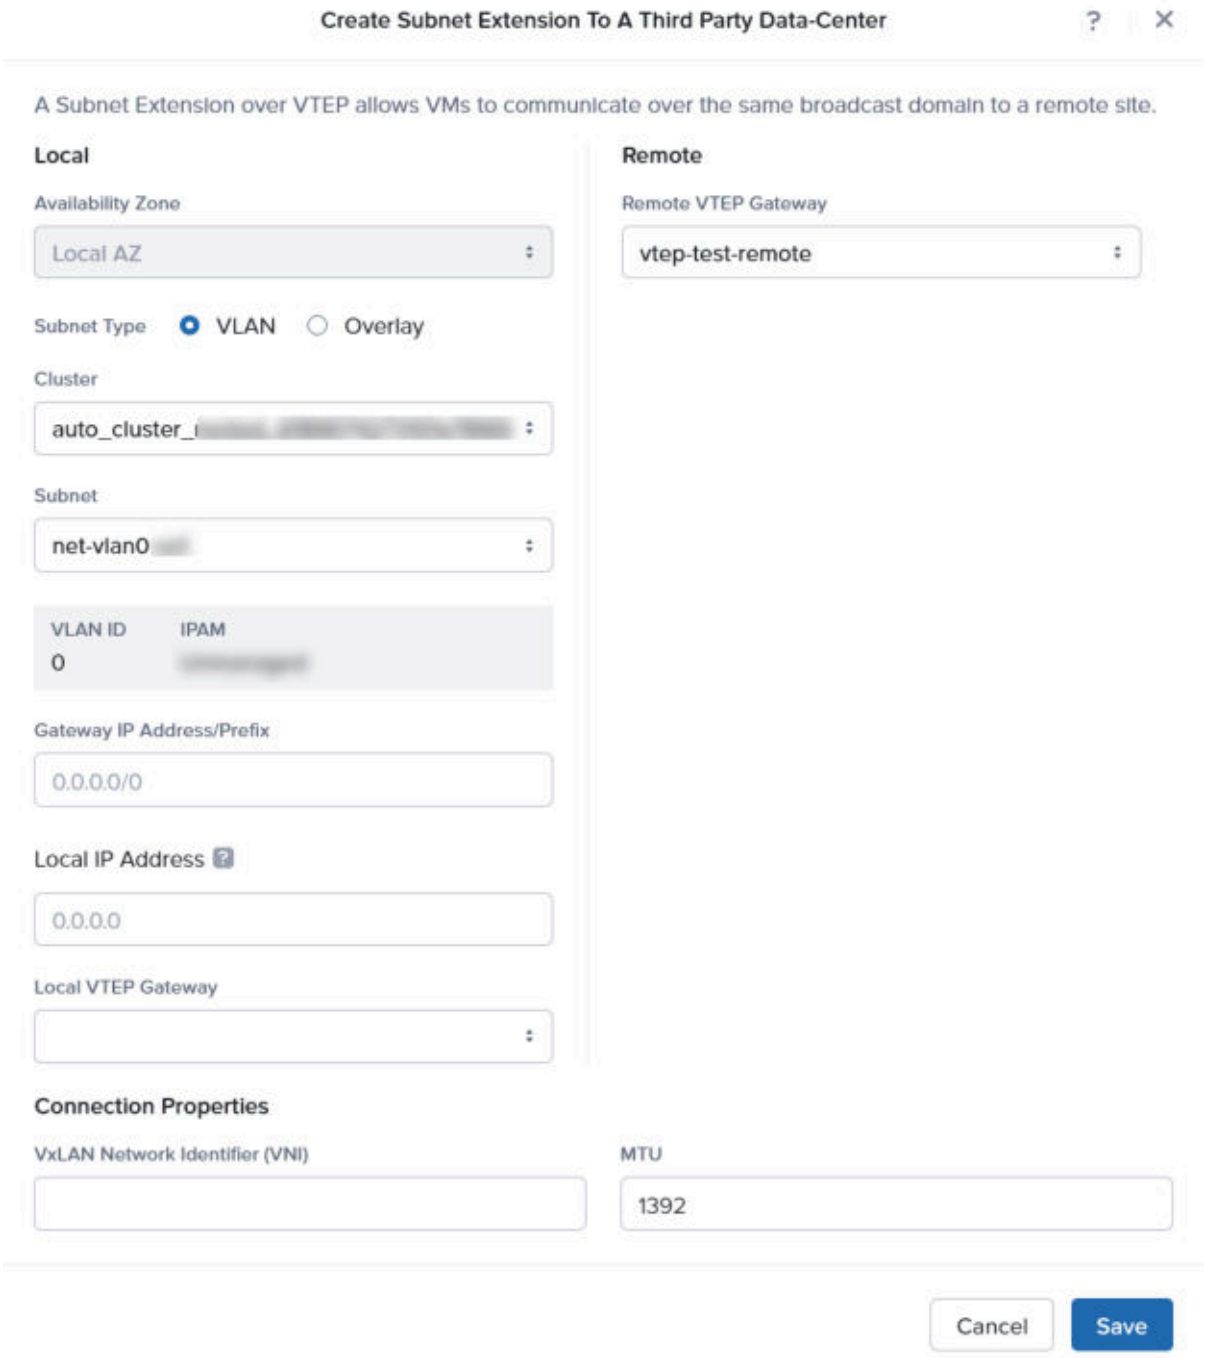

Figure 62: Example of Create VTEP Extension To A Third Party Data-Center with VLAN Subnet

2. Enter or select the necessary values for the parameters in the **Local**, **Remote** (AZ), and **Connection Properties** sections as described in the table.

Parameters	Description and Value
Local	

Parameters	Description and Value
Availability Zone	Displays the name of the paired availability zone at the local AZ.
Subnet Type	Select the type of the subnet - VLAN or Overlay that you are extending.
Cluster	Select the name of the cluster in the local AZ that the subnet is configured for.
Subnet	Select the name of the subnet at the local AZ for network. The VLAN ID and the IPAM - managed or unmanaged are displayed in the box below the **Subnet** field.
Gateway IP Address	Enter the gateway IP address of the subnet you want to extend. Ensure that you provide the IP address in <IP-address/network-prefix> format. for example the gateway IP is 10.20.20.1 in a .24 subnet then provide the gatewway IP address as `10.20.20.1/24`. Note: For unmanaged network, enter the gateway IP address of the created subnet.
Local IP Address	Enter a unique and available (unused) IP address from the subnet provided in **Subnet**.
Local VTEP Gateway	Select the local VTEP gateway you created on the local AZ. See Creating a Network Gateway on page 96 for more information about creating a local VTEP gateway.
Remote	
Remote VTEP Gateway	Select the remote VTEP gateway you created on the local AZ. See Creating a Network Gateway on page 96 for more information about creating a remote VTEP gateway.
Connection Properties	
VxLAN Network Identifier (VNI)	Enter a unique number from the range 0-16777215 as VNI. Ensure that this number is not reused anywhere in the networks that the Prism Central and Cluster are a part of.
MTU	The default MTU is 1392 to account for 108 bytes of overhead and the standard physical MTU of 1500 bytes. VPC Geneve encapsulation requires 58 bytes and VXLAN encapsulation requires 50. However, you can enter any valid MTU value for the network, taking this overhead into account. For example, if the physical network MTU and vs0 MTU are 1600 bytes, the Network Gateway MTU can be set to 1492 to account for 108 bytes of overhead. Ensure that the MTU value does not exceed the MTU of the AHV Host interface and all the network interfaces between the local and remote AZs.

3. Click **Save**.
   After the subnet is extended, the extension appears in the Subnet Extensions list view.

Updating an Extended Subnet

The *Update Subnet Extension Across Availability Zones* has the same parameters and fields as the *Create Subnet Extension Across Availability Zones* dialog box.

About this task

You can update a subnet extension that extends across AZs using the Update Subnet Extension Across Availability Zones or the Update Subnet Extension To A Third Party data center dialog box. The Update Subnet Extension Across Availability Zones or the Update Subnet Extension To A Third Party data center dialog box has the same parameters and fields as the Create Subnet Extension Across Availability Zones or the Create Subnet Extension To A Third Party data center dialog box, respectively.

Based on the type of the subnet extension that you want to modify, refer to the following:

Procedure

- Extending a Subnet Over VPN on page 116
- Extending a Subnet Across Availability Zones Over VTEP on page 119
- Extending a Subnet to Third Party Datacenters Over VTEP on page 123

Removing an Extended Subnet

About this task

Perform the following procedure to remove the subnet extension. This procedure deletes the extended subnet between the two Availability Zones (AZs) or between one Nutanix AZ and one or more third party subnets. Deleting the subnet extension does not automatically remove the network gateways or VPN connections that may have automatically been created by the Subnet Extension wizard. You need to separately delete these entities created automatically when the subnet was extended.

> Note: Removing an extended subnet from a cluster or AZ (either source or target AZs) automatically deletes the extended subnet from the corresponding source or target AZs.

Procedure

1. Click the hamburger icon in the top-left corner of the **Dashboard**.
   The main feature list appears.

2. Click **Network & Security** > **Connectivity** > > **Subnet Extensions**.
   The Subnet Extensions tab displays a list of the extended subnets.

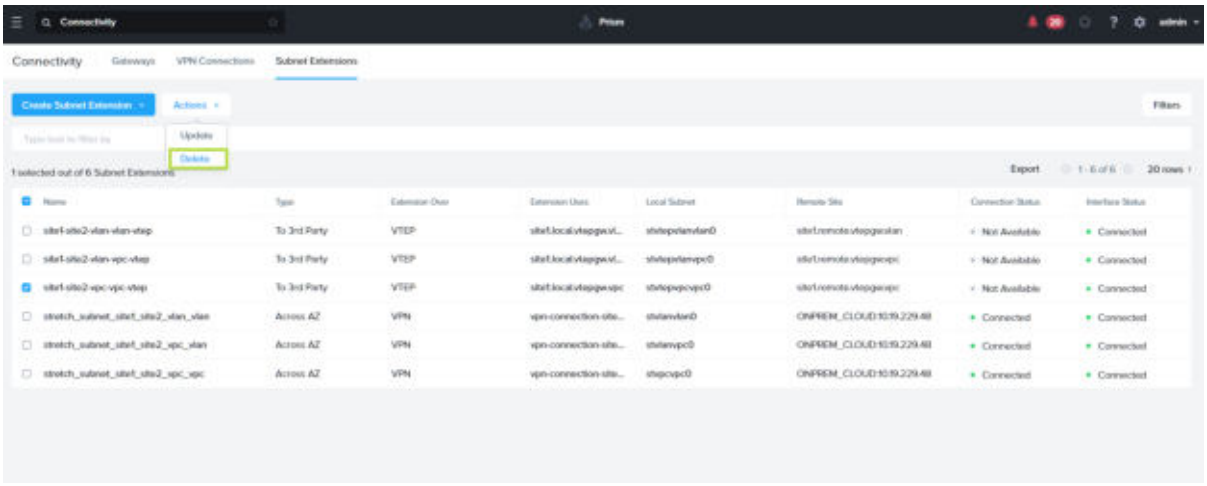

Figure 63: Sample Subnet Extensions dashboard

3. Select the subnet extension you want to remove.
4. Click **Actions** > **Delete**
   The confirmation dialog box is displayed.
5. Click **Remove**.
   Click **Cancel** to close the dialog box without removing the subnet extension.

What to do next

Check the list in the Subnet Extensions tab to confirm that the subnet extension is removed.

COPYRIGHT

Copyright 2022 Nutanix, Inc.

Nutanix, Inc.
1740 Technology Drive, Suite 150
San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

License

The provision of this software to you does not grant any licenses or other rights under any Microsoft patents with respect to anything other than the file server implementation portion of the binaries for this software, including no licenses or any other rights in any hardware or any devices or software that are used to communicate with or in connection with this software.

Conventions

Convention	Description
`variable_value`	The action depends on a value that is unique to your environment.
`ncli> command`	The commands are executed in the Nutanix nCLI.
`user@host$ command`	The commands are executed as a non-privileged user (such as nutanix) in the system shell.
`root@host# command`	The commands are executed as the root user in the vSphere or Acropolis host shell.
`> command`	The commands are executed in the Hyper-V host shell.
`output`	The information is displayed as output from a command or in a log file.

Version

Last modified: March 23, 2022 (2022-03-23T21:59:35+05:30)

NUTANIX

Welcome to the Nutanix family!

Thank you for choosing Nutanix Enterprise Cloud Platform! We look forward to bringing unprecedented simplicity and power to your enterprise datacenter as you embark on a journey to make IT infrastructure invisible. All of us at Nutanix are humbled by the trust that you, and thousands of other organizations in more than 150 countries, have placed in our people and solutions.

We want to ensure that you have an enriching experience as you run your business applications and services with Nutanix. You will have the opportunity to work with our world-class support organization, and see first-hand why we outpace every other IT vendor with a Net Promoter Score (NPS) of 90+ for over 6 years. And, you'll come to understand why the Association of Support Professionals has honored Nutanix as one of the top ten best support teams in recent years.

Nutanix Support Highlights

- 24x7x365 support with ten global support centers, including the United States (east and west coast), Mexico, Australia, Japan, China, India, and the Netherlands
- Leading technology experts consisting of VCPs, vExperts, CCIEs, MSFT, Linux, Kubernetes, Nutanix Certified Professional(NCP), Nutanix Certified Advanced Professional(NCAP), Nutanix Platform Experts(NPX) and Nutanix Support Specialists (NSS) with years of experience in supporting datacenter technologies and products
- Four-hour parts replacement for customers with mission-critical applications in most major metropolitan locations
- Full access to the Nutanix Customer Portal, containing validated and easy-to-consume information for all of our products, services and technology

Please visit Product Support for a comprehensive list of resources that will ensure great Nutanix Support experience.

I have listed a few of my favorite resources at https://www.nutanix.com/welcome. I invite you to explore these as you build and operate your Nutanix Enterprise Cloud.

With warm regards,

Rajiv Ramaswami
President & CEO, Nutanix

©2020 Nutanix, Inc. All Rights Reserved

SUPPORT QUICK REFERENCE GUIDE

SUPPORT CASE LIFECYCLE

All Support cases, whether opened through the web or the phone, go through a consistent lifecycle. The stages of the case lifecycle include:

1. **Create your Profile**
   If you are new to Nutanix Support, we need you to create a Customer Profile on our support portal.
   Visit http://portal.nutanix.com and click on "Sign up now"

   If you are having issues creating your account or profile, please e-mail portal-accounts@nutanix.com

2. **Collect Information to Expedite Resolution**
   Providing us with background information about the issue you are facing will help us understand your issue better and do some preliminary research before we engage with you to make our interaction more effective and efficient. We would therefore ask you to provide the following information when you open the case:

   - System Serial Number
   - Cluster ID
   - Software Versions of all relevant software (Nutanix OS and Hypervisor)
   - Time System became unavailable
   - Whether the issue is reproducible
   - Steps taken thus far in remediation
   - Results of Nutanix Cluster Check (NCC)
   - Workload characteristics – Applications running and number of VMs per block

3. **Create a Support Case**

   - Online: http://portal.nutanix.com
   - US Phone: 1-855-NUTANIX (688-2649) Ext 3
   - Other local country toll free numbers are available at: Support-phone-numbers

4. **Working on a case with SRE**
   An SRE is assigned to your case and owns your problem until we mutually agree the issue can be closed. He or she will contact you via email and/or phone as appropriate during the resolution process. Your responsibility is to have the appropriate people and resources available to work with the SRE during the service hours defined in your support agreement. We will inform you if we need to reassign your case for better skills or time zone match.

5. **Escalate a Support Case (Optional)**
   You can escalate a support case at any time via the support portal using the escalate button found in your case, by contacting your SRE, or by asking to speak with the Head of Support for your region. Support escalations may occur because your expectations of an issue resolution are not in line with the outlined action plan and timeframes that follow the normal resolution process.

   The issue is escalated to the Worldwide Support management team. The manager who takes on the escalation then creates a task force consisting of the necessary resources from Technical Support, Engineering, QA, and Product Management and formulates an action plan to address the issue.

   This action plan is then shared with you, the checkpoint schedule and the milestones determined, and agreement sought on the closure criteria of the escalation. Updates are provided on the progress as agreed upon, until the issue is resolved.

6. **Close a Support Case**
   A case is closed when you confirm that a resolution has been reached, or if we do not hear from you within two weeks of a request for information, and multiple attempts have been made to contact you during this period. A case may also be closed without final resolution, with acknowledgement and agreement from you.

7. **Customer Satisfaction Surveys**
   After a support case is closed, you will be invited by email to fill out a short survey about your experience. Your feedback is a valuable way of measuring how well the Nutanix Worldwide Support team is meeting your expectations. Customer satisfaction surveys give you the opportunity to provide us with valuable information to help improve our interactions with you, as well as any product improvement suggestions.

MORE HELPFUL LINKS AND INFORMATION

You can always find more information at: FAQ

TARGET INITIAL SUPPORT RESPONSE

To make sure your support request is prioritized correctly, Nutanix will use the following guidelines to assess issues and provide an initial response in a timely manner based on their priority level.

Priority	Description	Response Time
P1	**Emergency** – System is not available and productivity has been halted. Product is unusable in its current state. All data corruption issues are assigned this priority level.	**Within 1 Hour** (Production) **Within 30 Min** (Mission Critical)
P2	**Critical** – System is available but experiencing issues that have a direct impact on productivity. Major inconvenience.	**Within 4 Hours** (Production) **Within 2 Hours** (Mission Critical)
P3	**Normal** – System is having an occasional issue that has been identified as needing to be resolved, but the issue has not greatly affected productivity. Minor inconvenience.	**Within 8 Hours** (Production) **Within 4 Hours** (Mission Critical)

REPLACEMENT PARTS TARGET RESPONSE

- **Production** – Next Business Day (8AM-5PM)
- **Mission Critical** – 4 hours (available in most major metropolitan locations)

Should a customer choose to leverage our Field Engineer (FE) support offering an FE will be dispatched to the customer site to perform hardware replacement only. The FE does not perform any hardware or software troubleshooting and relies on the remote SRE to perform any software activities. The remote SRE will provide software subject matter expertise, prepare the environment for hardware replacement, initiate system restoration and perform the diagnostic health check to validate the environment.

NUTANIX™

©2020 Nutanix, Inc. All Rights Reserved

SUPPORT QUICK REFERENCE GUIDE

NUTANIX SUPPORT TEAMS AROUND THE GLOBE

Nutanix Support Teams are located in Australia, China, Japan, India(Bangalore and Pune), Mexico, The Netherlands, United States (East and West coasts) and are available 24 hours a day, 7 days a week, 365 days a year. We align our teams and your support case with your local business hours, and the rest of our global teams are there for you when need us.

BEST PRACTICES WHEN WORKING WITH NUTANIX SUPPORT

1. **Keep Your Profile and Location Up to Date**
   We encourage you to create and maintain your company information on our Support Portal. If you move your Nutanix equipment, it is important that you also update the install location.

2. **Educate Your Administrators**
   Please consider getting formally trained on the Nutanix suite of products. We have found that customers who invest in Nutanix education courses for their administrators are much more effective in avoiding problems, or speeding up issue resolution by defining the symptoms of problems and working with us to quickly resolve the underlying issue. You can find more information here: Training-Certification

3. **Plan Ahead**
   We recommend reviewing Nutanix Release Notes and other related Product documentation prior to deploying our products or performing an upgrade. These are available on: http://portal.nutanix.com.

4. **Enable Nutanix Pulse and Remote Tunnel Access**
   To the extent allowed by your organization's security policies, we encourage you to use our product's Nutanix Pulse feature, which regularly reports system health, and can automatically create a support case for certain issues it detects. The Remote Tunnel allows support staff to login remotely via secure SSH tunnel to proactively diagnose and resolve alerts which have been triggered from Nutanix Pulse or from the customer. Remote Tunnel access is an 'on demand' feature that is enabled or disabled by the customer in response to a request from Nutanix support to troubleshoot the system. Used together, they allow us to resolve support cases 30% faster.

5. **Regularly run Nutanix Cluster Check (NCC)**
   NCC is a framework of scripts that can help diagnose cluster health. To find additional information regarding NCC please go to: http://portal.nutanix.com and select 'Support Tools' from the Downloads section.

NUTANIX™
YOUR ENTERPRISE CLOUD

T. 855.NUTANIX (855.688.2649) | F. 408.916.4039
info@nutanix.com | www.nutanix.com | @nutanix

©2020 Nutanix, Inc. All rights reserved. Nutanix, the Nutanix logo and the other Nutanix products and features mentioned herein are registered trademarks or trademarks of Nutanix, Inc. in the United States and other countries. All other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).

Objects 3.4

Objects User Guide

March 15, 2022

Contents

Nutanix Objects Overview....6
Salient Features of Objects....6
Usage of Objects....7
Use Cases and Recommendations for NFS on Objects....8
NFS-S3 Interoperability....9
Advantages of Objects....10
Objects Architecture....11
Terminology Reference....13
Objects Workflow....13

Enabling Objects....15
Objects License Management....18
Open Source Software Usage....19

Finding Objects Version....20

Role-based Access Control for Objects....21
Prerequisites for Configuring Role-based Access Control for Objects....21
Role-Based Access Control Workflows for Objects....22
Objects Permissions....22
Creating Full Administrator Role for Objects....23
Creating View-only Administrator Role for Objects....25
Creating Bucket Administrator Role for Objects....25
Creating Infra Administrator Role for Objects....27
Creating a Custom Role for Objects....29

Deployment and Network Prerequisites....32
Deployment Prerequisites - AHV and ESXi....32
Prerequisites - ESXi....33
Network Configuration....34
AHV Configuration....35
ESXi Configuration....40
Shared Versus Single Network....47
URL and Port Requirements....48
Limitations....50
Limitations of NFS....50

Object Store Service Deployment....52
Object Store Instance Naming Conventions....52
Creating or Deploying an Object Store (Prism Central)....52
Viewing Object Store Deployments....60
Access Objects Endpoints....62
Deleting an Object Store....62
Object Store Expansion....63

Expanding Storage for an Object Store....64
Scaling Out an Object Store....68
Managing FQDN and SSL Certificate....73
Adding FQDNs....74
Setting up SSL Certificate for an Object Store....76
Object Store Deployment at a Dark Site (Offline Deployment)....78

Types of Objects Users....80

Directory Configuration and Access Key Generation....81
Configuring Directories....81
Generating Access Key for API Users....85
Viewing API Users....88
Managing API Keys....89
Deleting API Users....90

Bucket Creation, Operations and Bucket Policy Configuration....92
Bucket Naming Conventions....92
Creating and Configuring an S3 Bucket....92
Creating and Configuring an NFS Bucket....94
Managing NFS Allowlist....98
Bucket Policy Configuration....101
Object Versioning....101
Lifecycle Policies....102
Cloud Tiering....104
Legal Hold for Objects....112
WORM Bucket....113
Configuring a Bucket for Static Website Hosting....115
Cross-Origin Resource Sharing (CORS) Overview....118
Viewing Buckets....120
Buckets Filter Options....121
Bucket Summary....122
Updating a Bucket....122
Sharing a Bucket....123
Bucket Access Policies....125
Listing the Shared Buckets....129
Viewing Bucket Users....130
Deleting a Bucket....131

Objects Browser....133
Administrator Workflow: Objects Browser....134
Launching the Objects Browser....135
Supported Operations....136
Bucket Operations....136
Object CRUD Operations....148
Bucket Summary....157

Objects Streaming Replication....158
Replication Guarantees and Topologies....159
Replication Prerequisites....161
Adding Remote Prism Central as Availability Zone....162

Setting up IAM Synchronization with a Different PC 164
Creating Replication Rules between Buckets 166
Deleting a Replication Relation 169
Viewing Replication Statistics for a Bucket 170
Achieving Fault Tolerance for IAM 172

Baseline Replicator Tool 174
Accessing and Running the Baseline Replicator Tool 174

Monitoring and Alerts 177
Viewing Performance of Object Stores 177
Viewing Performance of Buckets 178
Viewing Object Store Usage 180
Assigning Quota Policy to a User 180
Viewing Buckets Usage 182
Viewing Alerts 183
Nutanix Objects Specific Alerts 185

Objects Notifications 190
Notification Types for Objects 190
Configuring Events Notification 195
Creating Notification Rules for Data Events 198

CRUD Operations by Using S3 APIs 202
Authentication 202
Supported and Unsupported APIs 202
Supported APIs 202
Unsupported APIs 206
Objects Tagging APIs Overview 207
API Operations Supported for Tagging 207
S3 Select API Overview 208
Supported SQL Functions 209

Error Responses 214
REST Error Responses 214
List of Error Codes 214

Integration with Backup Applications 217

Objects LCM Upgrades 218
Microservices Platform (MSP) 219
Finding the MSP Version 219
Finding the Primary and Secondary MSP Clusters 219
Upgrading MSP Controller 219
Upgrading Objects Manager 220
Upgrading Objects Service 221

Troubleshooting Objects 222

Shutting Down Objects VMs..... 222
Powering on Objects VMs..... 224
Detection of Slow Connections..... 225
Prism Element Clusters or Subnets Not Listed in the User Interface..... 225

Copyright..... 227

NUTANIX OBJECTS OVERVIEW

Nutanix Objects™ (Objects) is a software-defined Object Store Service. This service is designed with an Amazon Web Services Simple Storage Service (AWS S3) compatible REST API interface capable of handling petabytes of unstructured and machine-generated data. Objects addresses storage-related use cases for backup, and long-term retention and data storage for your cloud-native applications by using standard S3 APIs. You no longer have to introduce an external, separately managed storage solution.

Objects is deployed and managed as part of the Nutanix Enterprise Cloud OS. It enables users of the Nutanix platform to store and manage unstructured data on top of a highly scalable hyper-converged architecture. In comparison to cloud-hosted solutions, this on-premises model offers more consistent control over the costs associated with storing objects, along with more transparency about the location of those objects.

You can manage objects by using Prism Central or the S3-compatible REST APIs after an administrator has authorized the applications and users to access buckets accordingly.

For more information on Objects architecture, refer to *Nutanix Bible*.

Salient Features of Objects

The salient features of Objects are as follows:

Write-Once-Read-Many (WORM)

Nutanix Objects enables you to create WORM buckets to prevent anyone (including an administrator) from modifying or deleting data while the policy is active. WORM policies help you comply with strict regulations often mandated by the healthcare, financial, and government sectors. You can apply WORM policies to any bucket to prevent any updates to the data until the WORM policy expires.

Nutanix provides a 24-hour grace period for letting you test the WORM policy. The bucket creator can undo the applied policy within the allowed grace period of 24 hours.

After the grace period, no one can delete the policy or reduce the retention period. The option to extend the retention period is available. You must enable versioning on a WORM bucket so that you can update an object. The updated object gets stored as a new version. Versioning ensures that older versions of the object get retained, and the data never gets overwritten. For more information, see WORM Bucket on page 113.

Immutability

Retrieved data must be the same as originally written. This means that you can retrieve the stored object (and its associated metadata) in the future without any modification. If you enabled WORM, the object remains unchanged and is also undeletable.

Object Versioning

When you enable object versioning, uploading multiple copies of the same object creates multiple versions. You can use the versioning feature to protect data from accidental overwrite or deletion or as an option for reverting to a previous state.

You can also disable object versioning on a bucket at any time and set the objects in the bucket to expire after a specified time. By default, there is no limit on the number of versions for an object, considering the storage space is available. For more information, see Object Versioning on page 101.

Data Life Cycle Management

You can use age-based data retention policies to enforce compliance with strict data retention regulations that dictate the time to store specific data. For example, HIPAA regulation mandates you to retain medical data for six years from the creation date. You

can set a data retention policy to delete all data created in a specific bucket six years after its creation date.

> Note: The WORM policy supersedes any set retention policy.

You can set retention policies at the bucket level for non-WORM entities to specify the maximum number of maintained versions for each of the objects in the bucket.

The retention policies then delete older versions of the objects in first in, first out mode to create space for the most recent versions. You can also approach retention policies from a time perspective, where objects in a bucket expire after a certain amount of time you specify. If you do no set retention policies, the only limit to how many versions can be maintained depends on the storage space. For more information, see Lifecycle Policies on page 102.

Multipart Upload

Multipart upload allows you to reduce slow upload times by breaking large pieces of data into chunks. Then, the system can handle the data separately to increase upload speeds if you upload them simultaneously. You can also use multipart uploads to prevent losing progress during an upload. For example, if there is a connectivity loss, most applications retry only the unsuccessful chunk. Hence, you do not have to upload the entire object again. For more information, see Supported APIs on page 202.

Data-at-Rest Encryption with Native Key Management

Nutanix Objects provides a FIPS 140-2 compliant data-at-rest encryption solution. To deliver this capability, Objects uses the underlying AOS encryption capability. You can set encryption at an entire cluster level, always encrypting all data. With the native key management, the Nutanix cluster manages the keys, so the solution requires no additional device management or third-party costs.

Identity and Access Management

Native IAM functionality ensures that you have access only to the buckets and objects you created and granted access permissions. Each user gets a pair of access and secret keys that can be used by their applications to access Nutanix Objects. You can also generate access and secret keys for an Active Directory group. Administrators can revoke and regenerate keys at any time. For more information, see Directory Configuration and Access Key Generation on page 81.

Multi-protocol Access

Objects allows you to create buckets using both S3 and NFS protocols. NFS protocol support is natively implemented over Objects and build on the same foundation that powers the S3 protocol. For more information, see Use Cases and Recommendations for NFS on Objects on page 8, NFS-S3 Interoperability on page 9, Limitations of NFS on page 50, Creating and Configuring an NFS Bucket on page 94, and Creating and Configuring an S3 Bucket on page 92.

Usage of Objects

Following are examples of solutions you can implement by using Objects:

- Backup – You can integrate Objects with the backup applications such as Commvault, HYCU, Veeam, and Veritas. You can create backups to protect your data with a simple, scalable, and cost-effective active archive solution. You can start with small storage and scale to petabytes of storage to deliver great performance. Objects supports the multipart upload API with which you can reduce slow upload times by breaking data into chunks and upload documents, images, and videos to the global namespace.

- Long-term Retention – You can use Objects for long-term data retention. Built-in object versioning provides storage protection and searches your data without the problem of tape systems. Versioning maintains previous copies of the object to avoid data loss from overwrites or deletes.

  Some organizations, especially in the health, legal, and government sectors, require to comply with the strict regulations mandating requirements such as the minimum period for which data must be available or who can alter data. Nutanix Objects allows you to store data over a long period with features designed to comply with such strict regulations. For example, WORM buckets prevent anyone (including an administrator) from modifying or deleting data while the policy is active.

- DevOps – Easy application access to your data in a global namespace using simple PUT and GET commands makes Objects the perfect fit for your dev-ops data. Nutanix Objects offers broad support for most programming environments, and you can access it using standard internet networking infrastructures, such as HTTP and HTTPS protocols. With Nutanix Objects, you can easily access the objects with S3-compatible HTTP REST API requests from different locations. The bucket-level metrics allow you to track resource utilization with ease. Also, Identity and Access Management (IAM) support enable secure access to Objects resource and services.

  DevOps and IT ops can use the S3-compatible interface for cross-geo, cross-team collaboration, and agile development.

  Consider the following scenario. Joe and Adam work together on the same team as developers for an image processing application called XYZ. Joe, residing in New York, develops code for the data creation and storage of XYZ. Adam, residing in Seattle, performs the testing of XYZ. Several buckets dedicated to the application XYZ are stored in the Nutanix Objects store instance deployed in the New York office. After completing the programming for the latest feature, Joe stores the latest version of XYZ in a bucket. Adam can use GET request automation-scripts to pull the latest version of the code and store the latest test results into another bucket using PUT requests.

  Thus, the team works more efficiently and quickly. Also, it demonstrates the immutable capabilities of object storage, where Joe's data remains unchanged while Adam performs testing.

Use Cases and Recommendations for NFS on Objects

Use Cases

NFS protocol support provides multiprotocol access (NFS 3.0, S3) to data stored within Objects buckets. NFS protocol support is natively implemented over Objects and builds on the same foundation that powers the S3 protocol.

NFS access over Objects is ideally suited for large scale, read-heavy workloads with sequential accesses where data is ingested once and minimally or never modified later. Some ideal uses are:

- Backup applications – use Objects as a large scale NFS repository for backups while you migrate the underlying storage to Objects.
- Analytics – use Objects multiprotocol access for in place analytic as you ingest data through object or S3 interface and access through file or NFS interfaces that analytic systems require.

Recommendations

Nutanix recommends not to use NFS access over Objects for use cases that require modification to data including file edits and renames. For example, end user computing, home shares, app data that is frequently modified like presentations or CAD files.

NFS-S3 Interoperability

This section explains how objects in the S3 namespace get mapped to files and directories in the NFS namespace and vice-versa.

NFS to S3

- All the files and symbolic links created in the NFS namespace appears in the S3 namespace as objects.
- Any directory created from the NFS namespace will not appear in the S3 namespace as S3 protocol does not have any notion of directories.
- Any S3 operation like ObjectHead, ObjectGet or ObjectDelete on the directories fails with the error message `NfsDirectoryOperationNotAllowed (Operation prohibited on an NFS directory)`. Therefore, a file present inside directories and subdirectories appears as a single object in the S3 namespace as shown in the following table:

Table 1: NFS to S3

NFS	S3

• dir1/
  • dir2/
  • file
	dir1/dir2/file

S3 to NFS

- Objects created from the S3 protocol appears as files and directories in the NFS namespace based on the object name. If the object name contains a directory-like hierarchy, then the object is stored in a hierarchical namespace as identified its name.

  Following table shows an example where an object `a/b/c` created from S3 protocol appears in the NFS namespace as `a/`directory which contains the subdirectory `b/` which in turn contains the file `c`.

Table 2: S3 to NFS

S3	NFS

a/b/c	• a/
	  • b/
	  • c

Note:

- These implicit directories are created when an object containing directory hierarchy is accessed from the NFS namespace.
- The implicit directories continue to exist in the NFS namespace after the object containing directory hierarchy has been deleted.
- Folders created from the Objects Browser also appears as a directory in the NFS namespace.

For more information about Limitations of NFS on Objects, refer to Limitations of NFS on page 50.

Advantages of Objects

Following are the advantages of object storage with Nutanix Objects:

No Silos

Nutanix provides file services (Files) and block services (Volumes) as part of the Acropolis Distributed Storage Fabric (DSF). You can also add Nutanix Objects to the solution, thus allowing block, file, and object storage solutions to coexist with no silos. You can deploy and manage these features in a single environment.

Security-First Approach

Nutanix integrates security into every step of its solution stack from the early stages of development. For example, the stack conforms to Security Technical Implementation Guides (STIGs), which maintain a security baseline configuration based on common standards established by the National Institute of Standards and Technology (NIST). STIGs use machine-readable code to automate self-healing and compliance with the security standards for AOS and AHV. Nutanix also complies with *SEA 17a-4*, which specifies requirements for data retention and accessibility. Nutanix Objects also supports Data-at-Rest Encryption.

Capacity Optimization

Nutanix Objects leverages data-capacity optimization technologies such as compression and erasure coding (EC-X) in the background. In addition to compression savings, EC-X increases the usable storage capacity in a cluster with no overhead to the active path write.

In a cluster with the redundancy factor 2, two copies of the data get replicated among all nodes for resilience. Checksums of the data are stored with the metadata to ensure validity if corruption occurs. Therefore, the cluster (redundancy factor 2) uses half of its raw storage capacity to store copies of its data. EC-X performs the OR operation on these copies of data to compute a parity block. The original data blocks and the parity form an erasure code strip. This process reduces the number of actual data copies needed to protect the environment from a single node failure.

Due to the computational overhead of calculating parity during write operations, EC-X is suitable for workloads with infrequent write operations such as Objects with WORM policies enabled.

Suppose you have three data blocks (A, B, and C) stored in a four-node cluster, with each data block distributed across three nodes. Without EC-X, you may have the original three blocks of data plus three more copies of those blocks, which sum up to a total of six blocks of data. With EC-X enabled, you can reduce the storage footprint to four blocks of data that include the original three blocks of data plus their parity. These blocks are present in different nodes. If a node fails, the system can reconstruct the data from the data blocks and parity data present on the available nodes. Example: If the node hosting block B fails, the system can reconstruct its data from blocks A, C, and P (parity).

Additionally, when the disks are not getting used heavily, the system regularly performs background checksum validations and data restoration to maximize protection against data corruption. These techniques (combined with a high tolerance for disk and node failures in a self-healing and distributed system) forms a resilient and reliable enterprise object-storage platform.

Ease of Use and Simple Management

Nutanix Objects is compatible with HTTP S3 REST APIs for object and bucket storage. Objects supports all the necessary methods (including `GET`, `PUT`, `POST`, `DELETE`, and `LIST`) to perform basic object operations.

Nutanix Objects inherits the simplicity of use and management of Prism Central and provides the managing, monitoring, and reporting capabilities that are simple to administer.

For example, Prism Central automatically provides checks for interoperability and compatibility during a software update task. Alerts inform users about specific conditions, such as object storage components hitting capacity thresholds and reduced High Availability (HA). The Prism GUI reports on the size of the object store, number of buckets and objects, average throughput, latency, and the number of GET and PUT operations.

Cost

With Nutanix Objects, the cost to store a large amount of unstructured data is less compared to traditional storage solutions. Compared to public cloud providers, Nutanix Objects does not charge for data ingress or egress.

Objects Architecture

Nutanix Objects integrates with the solution stack through services that run inside the Prism Central VM and manage all the other components and services of object storage.

The Nutanix cluster deploys VMs to handle the multiple components that provide the object storage API and lookups for the objects. These components run as containerized services in a Kubernetes cluster. Objects follows a modular and scale-out design where each component focuses on a single core function, thus allowing you to scale out any component independently to match the workload demands.

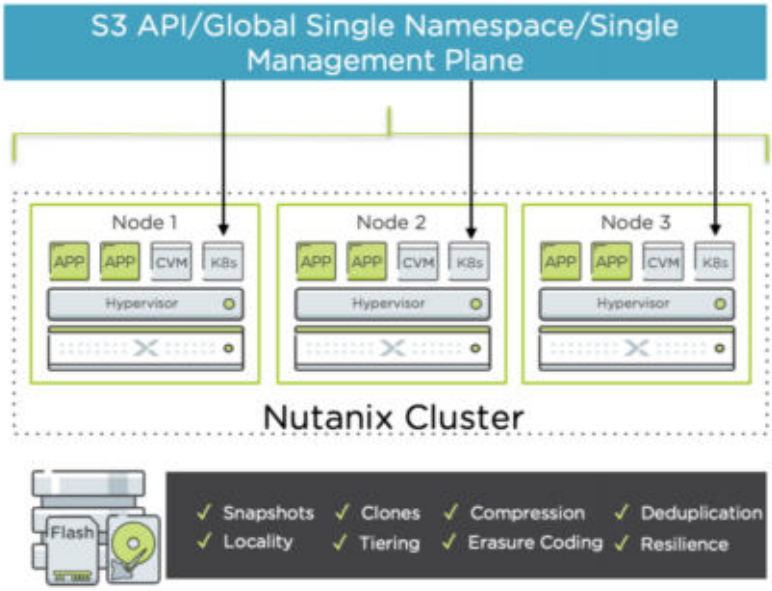

Figure 1: S3-Compatible Objects Storage Integrated with Nutanix

When you make `GET` and `PUT` requests to your object storage endpoint, you first hit the front-end adapter, a native, built-in load balancer that manages the S3 REST API calls and directs them to the right worker VM.

The worker VM runs different services that include the following:

- An object controller service, which supervises the data management layer that interfaces with AOS and coordinates with the metadata service.
- A metadata service that manages the metadata and serves as a general key-value store that also handles partitioning and region mapping.
- A life cycle management service that controls life cycle, audits, and background maintenance activities.
- An identity and access management service that handles user authentication for accessing buckets.

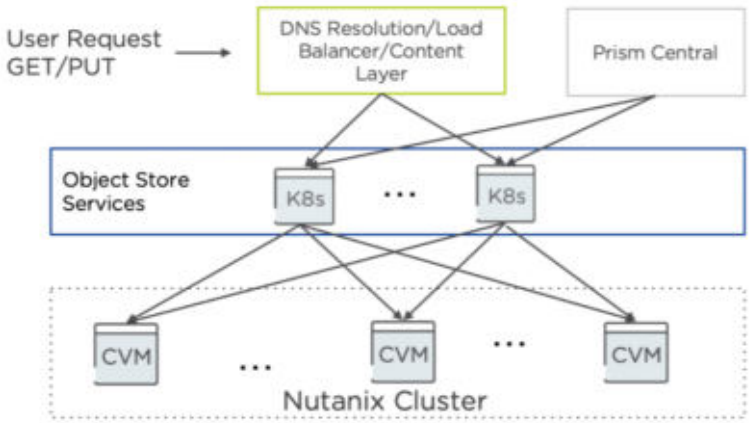

Figure 2: Objects Layered Architecture

Terminology Reference

Following are the terms that you frequently encounter when you are using Objects:

Table 3: Terminology Reference

Terminology	Description
Bucket	An organizational unit exposed to the users and contains the objects. A deployment may have one or more buckets.
Object	Object represents the data uploaded by the user or application. The actual unit (blob) of storage and the item interfaced by using the API (GET or PUT).
S3	The term used to describe the Amazon Web Services (AWS) interface. This term is now used synonymously for an object service. S3 is also used to describe the object API which you use to interact with an object store.
Storage Network	A VLAN required for communication between Objects services.
Public Network	A VLAN used to access the Object Store endpoints externally.
Microservices Platform (MSP)	A platform based on Kubernetes where all the Objects microservices run.
AHV IPAM	IP Address Management (IPAM) is a feature of AHV that allows it to assign IP addresses automatically to VMs by using DHCP. You can configure each virtual network with a specific IP address subnet, associated domain settings, and groups of IP address pools available for assignment to VMs.
Worker VMs	Virtual machines created during object store deployment that host various containerized Objects services. Also, Worker VMs are referred to as Objects nodes.
Objects Browser	A User Interface (UI) that allows the users to directly launch the object store instance in a web browser and perform bucket and object level operations.

Objects Workflow

This section describes the basic workflow of Objects. The subsequent sections provide detailed information about each step.

About this task

Following is the Objects workflow:

Procedure

1. Enable Objects from Prism Central. Refer to Enabling Objects on page 15.
2. Deploy object store on a desired cluster. Refer to Object Store Service Deployment on page 52.
3. Generate the access keys. Refer to Generating Access Key for API Users on page 85.
4. Set up the Secure Sockets Layer (SSL) certificates for the object store. Refer to Setting up SSL Certificate for an Object Store on page 76.
5. Access the object store endpoints from the third-party clients or Objects Browser. Refer to Access Objects Endpoints on page 62 and Objects Browser on page 133.
6. Create buckets using either S3 or NFS protocol. Refer to Bucket Creation, Operations and Bucket Policy Configuration on page 92.
7. Upload objects and perform object operations using Objects Browser or S3 APIs. Refer to Supported Operations on page 136 and Supported APIs on page 202.
8. Expand the object store if the storage is getting full. Refer to Expanding Storage for an Object Store on page 64.

ENABLING OBJECTS

To create an object store, you first need to enable Objects in Prism Central then add a license in Prism Element or Prism Central depending on the versions in use. Only users with Objects Admin, Prism Admin, Cluster Admin, or User Admin privilege can enable Objects.

Before you begin

A non-admin user can enable objects only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- Objects Admin
- Cluster Admin
- User Admin

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

To enable Objects in Prism Central, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.

2. To enable the Object Store Services, click **Enable**.

   Note: Enabling of Object Store Services is performed only once.

   Note: After you enable Objects, ensure that you perform LCM inventory and upgrade the MSP and Objects Manager to the latest versions before you start with deployment. For more information, see Objects LCM Upgrades on page 218.

   A welcome page appears with a prerequisites details to create an object store.

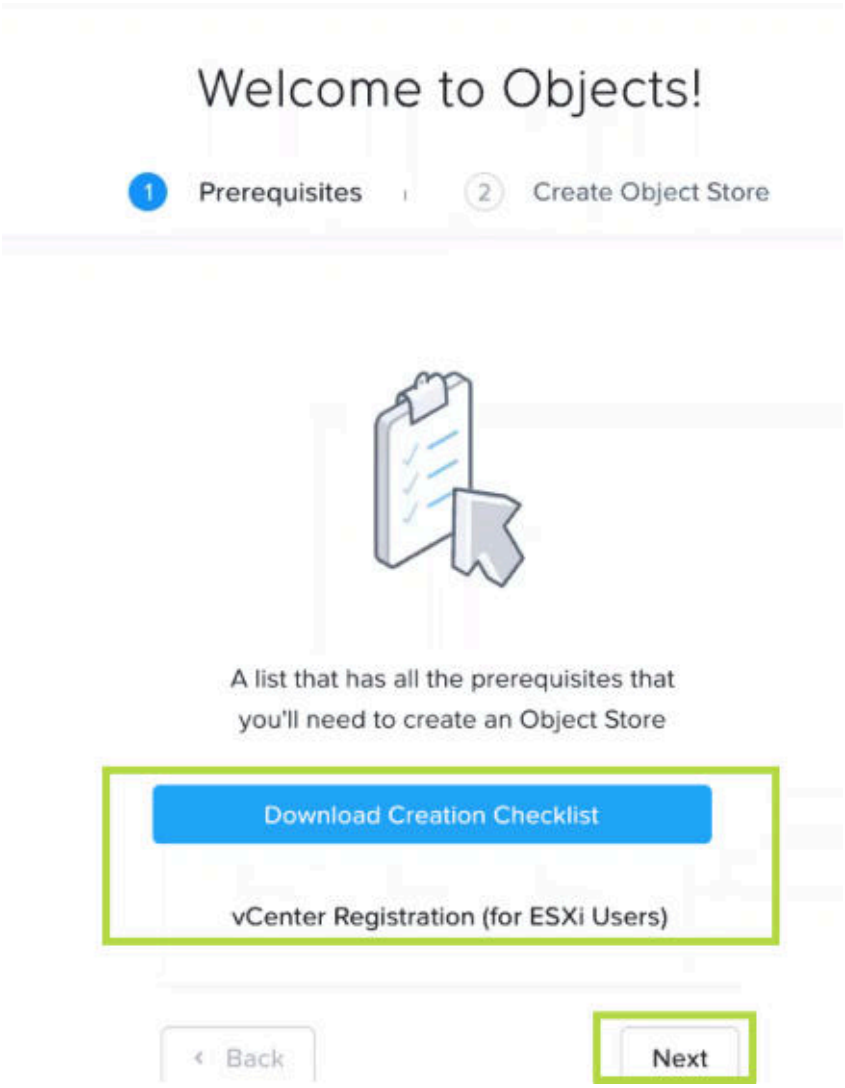

Figure 3: Objects Prerequisites

3. Click **Download Creation Checklist** to download the list of prerequisites for deploying an object store.

4. (Only for ESXi clusters) Click **vCenter Registration** before deploying object store on an ESXi cluster.

   To deploy object stores on the ESXi clusters, you need to provide the vCenter credentials and configure the IPAM for the ESXi networks. For more information, refer to Managing vCenter for Object Service on page 43.

   Note: A non-admin user can perform vCenter Registration only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

   - Create Object Store
   - View Object Store

   Administrators and non-admin users are Prism Central users with the following roles:

   - Administrator—A Super Admin or a Prism Admin in Prism Central.
   - Non-admin user—A Prism Central user without any administrator privileges.

   For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

5. Click **Next**.

6. Click **Create Object Store** to start creating the first object store.
   For more information on creating the object store, refer to Creating or Deploying an Object Store (Prism Central) on page 52.

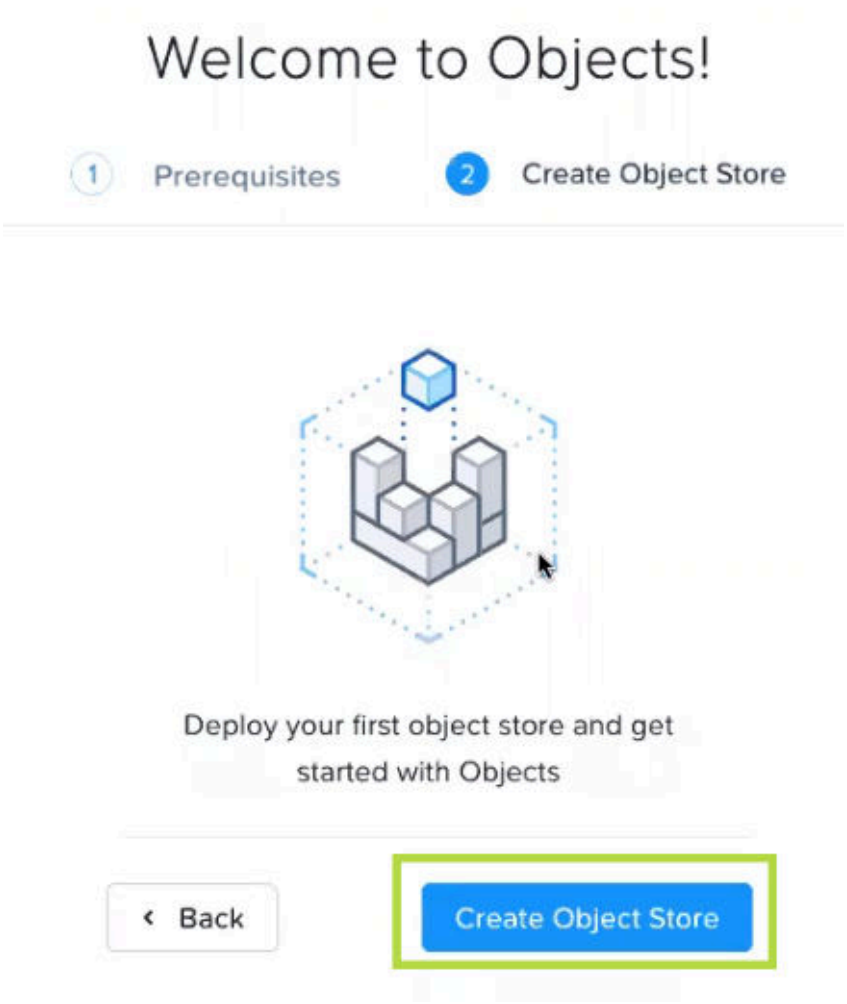

Figure 4: Objects Pre-configuration

Objects License Management

Nutanix provides licenses that you can apply manually to help ensure access to all the features. This feature enables you to administer your environment based on your current and future needs. With the introduction of multi-cluster feature in Objects, you can now manage (add, verify, and monitor) Objects license from the Prism Central, not from the Prism Element.

Refer to the following image to understand the licensing matrix.

Prism Central (PC) version → / Prism Element (PE) version ↓	5.11.x / 5.15.x	5.16.x	5.17.x or above
5.11.x / 5.15.x	Objects license in PE	Objects license in PE	Objects license in PE
5.16.x	Objects license in PE	Objects license in PE	Objects license in PE
5.17.x or above	Upgrade PC to version 5.17.x. Then, Objects licensing will be in PC.	Upgrade PC to version 5.17.x. Then, Objects licensing will be in PC.	Objects license in PC.

Figure 5: Licensing Matrix

> Note: Up to 2 TB Objects storage is free for all users.

Types of Licenses

Nutanix provides the following two types of Objects licenses.

- Objects (For AOS): This license allows you to deploy Objects on clusters with AOS licenses.
- Objects (Dedicated): This license is used when deploying an Objects-only cluster without AOS licenses.

Adding Licenses to Prism Central

You can add Objects license to the Prism Central manually. For more information, see Manage Licenses Manually section in the *Nutanix Licensing Guide*.

You can view the Objects usage information in the Prism Central. For more information, see Displaying License Features and Details section in the *Nutanix Licensing Guide*.

Open Source Software Usage

For more information about Objects open source licensing details, refer to Nutanix Support Portal.

FINDING OBJECTS VERSION

You can find the installed versions of Objects Manager and Objects Service from Prism Central.

About this task

To find the version of Objects, do the following:

Procedure

1. Log on to Prism Central.
2. On the sidebar, go to **Administration** > **LCM** > **Inventory**.
3. Click **Perform Inventory**.
   The Installed Versions list displays the version of Objects Manager and Objects Service.

ROLE-BASED ACCESS CONTROL FOR OBJECTS

Prism Central supports role-based access control (RBAC) that allows you to configure and provide customized access to the users based on their assigned roles for the objects entities.

The access policies and permissions are user-based policies and applies only to the actions performed on the Prism Central web console. These policies and permissions do not apply to the actions performed on Objects Browser or any S3 compatible client application even if the same users have Prism Central access and Objects Browser access.

From the Prism Central roles dashboard, you can define and assign the following additional roles to users or user groups:

- A set of predefined built-in roles (system roles)
- Additional custom roles

Objects support the following built-in roles (system roles) that are defined by default:

- Objects Admin—permissions required to enable, create, and manage object stores.
- Objects Viewer—permissions required to perform read-only operations on object stores.

> Note: For more information on Role Based Access Control, see:
>
> - *Controlling User Access (RBAC)* in the *Security Guide*.
> - *Built-in Role Management* in the *Security Guide*
> - *Operations Management* in the *Prism Central Guide*.

Prerequisites for Configuring Role-based Access Control for Objects

Make sure that your system meets the following requirements before you start to configure role-based access control for objects.

- Prism Central version pc.2021.9 or later
- Objects Manager version 3.4 or later
- Objects Service version 3.4 or later (recommended version)
- Microservices Infrastructure is enabled in Prism Central

> Note: For information on enabling Microservices Infrastructure, see *Enabling Microservices Infrastructure* in the *Prism Central Guide*.

Only administrators (Super Admin or a Prism Admin in Prism Central) can create roles and access control policies for objects. For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

> Note: To view the Prism Central version, log on to the Prism Central web console, click the **Entity** menu, and click **Settings** > **Prism Central Management**. The Prism Central version is listed in the **Prism Central Summary** section.

Role-Based Access Control Workflows for Objects

Role-based access control (RBAC) allows users to configure and provide different levels of access to the users based on the role of the user within an organization.

You can create the following roles and assign them to the users based on your requirements:

- Full administrator
- View-only administrator
- Buckets administrator
- Infra administrator
- Custom

> Note: For more information on role-based access control, see:
>
> - *Built-in Role Management* in the *Security Guide*
> - *Controlling User Access (RBAC)* in the *Security Guide*.
> - *Operations Management* in the *Prism Central Guide*.

Objects Permissions

Administrators can use the Prism Central RBAC feature to create access control policies and attach it to objects users. A non-admin user can perform object-related tasks only after an administrator creates an access control policy on a role in Prism Central for the non-admin user.

The administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

The following table lists the permissions that an administrator can grant to a non-admin user to perform various objects tasks. These permissions allow the users to perform the following tasks:

Table 4: Objects Permissions

Permission	Tasks Users Can Perform
Add Multicluster	Add Prism Element clusters to expand the storage.
Create Endpoint	Create remote endpoints for tiering.
Create Object Store	Create object stores.
Create Quota Policy	Create quota policies for objects.
Delete Object Store	Remove object stores.
Delete Quota Policy	Remove the quota policies created by the respective user for objects.

Permission	Tasks Users Can Perform
Download Certificate	Download the CA-signed certificates for an object store.
Edit Buckets	Edit the buckets in an object store.
Regenerate Self Signed Certificate	Regenerate self-signed certificates for an object store.
Replace Certificates	Replace certificates for an object store.
Scale Out	Add more resources to an existing object store cluster.
Set Capacity	Add additional storage capacity to an existing object store cluster.
Set Notification Endpoint	Add the notification endpoints.
Update Endpoint	Update the existing tiering endpoints.
Update Multicluster	Remove or modify the existing expanded storage (Prism Element cluster).
Update Quota Policy	Update the quota policies.
View Buckets	View the buckets in an object store.
View Endpoint	View the existing tiering endpoints.
View Multicluster	View the Prism Element clusters added for expanded storage.
View Notification Endpoint	View the notification endpoints.
View Object Store	View the object stores.
View Quota Policy	View the quota policies.

Creating Full Administrator Role for Objects

You can create full access administrator policy for users or user groups, and grant full access rights to all operations on Objects instances using role-based access control (RBAC) from the Prism Central web console.

About this task

To create a full access policy for Objects instances, do the following:

Note:

- Only users with administrator privileges create a full access policy for Objects instances.
- It might take some time (around five minutes, or so) for the changes made to the policy to be effective.

Procedure

1. Log on to the Prism Central web console as an Admin user, and click the **Entity** > **Administration** > **Roles**.
   The Roles page appears.

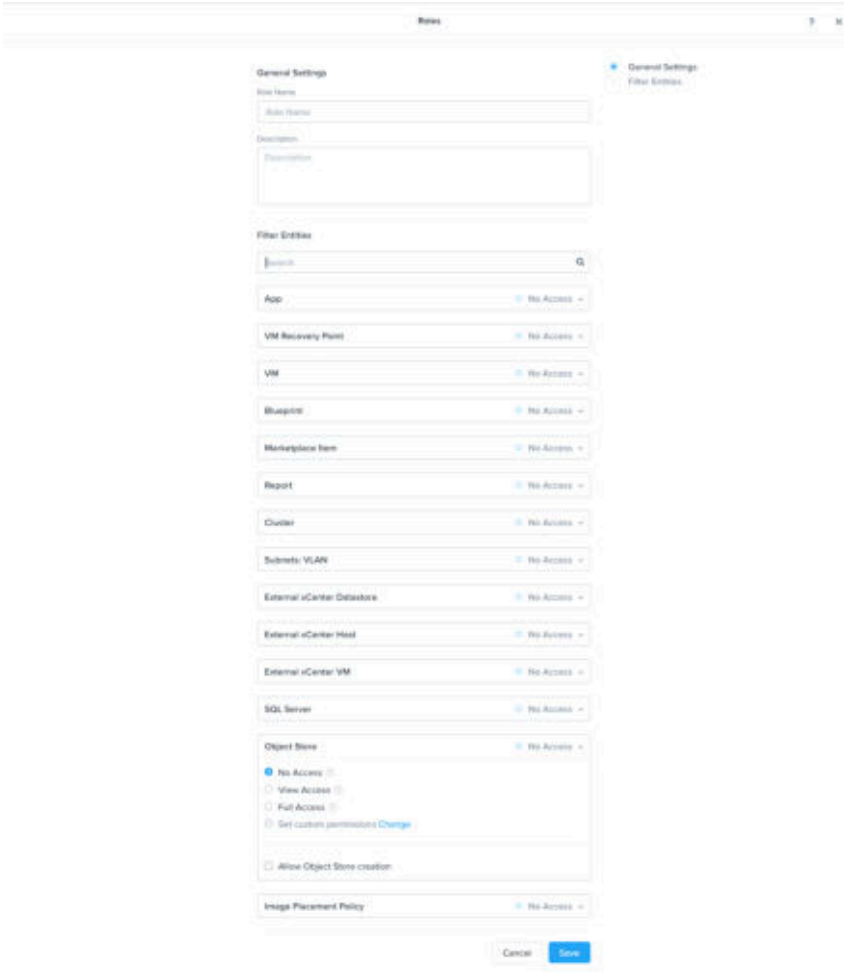

Figure 6: Roles Page

2. Select Objects Admin role and click **Actions** > **Manage Assignment**.
   The Role Assignment page appears.
3. Click **Add New** to create a new role.
4. From the **Select Users or Groups** section, select the required Active Directory group from the pull-down list, and then select one or more users, or user groups using the **Search User** field.

   Note: You can only add users who are part of an Active Directory group.

5. From the **Entities** section, select object store as the **Entity Type**, and then select the following:

   - Individual Entity
   - Select **All Objects** or enter the name of the Objects instance on which you need to grant custom permissions for the user.

6. Click **Save**.
   An administrator user with full access is created.

Creating View-only Administrator Role for Objects

You can create a view-only administrator role for users or user groups and grant complete access to read only operations on Objects instances using role-based access control (RBAC) from the Prism Central web console.

About this task

> Note: Only users with administrator privileges can do this task.

To create a view-only admin access policy for Objects instances, do the following:

Procedure

1. Log on to the Prism Central web console as an Admin user, and click the **Entity** > **Administration** > **Roles**.
   The Roles page appears.
2. Select Objects Viewer role and then click **Actions** > **Manage Assignment**.
   The Role Assignment page appears.
3. Click **Add New** to create a new role.
4. From the **Select Users or Groups** section, select the required Active Directory group from the pull-down list, and then select a user or user groups using the **Search User** field.

   > Note: You can add only those users who are part of a Windows Active Directory group.

5. From the **Entities** section, select object store as **Entity Type**, and then select the following:

   - **Individual Entity**
   - Select **All Objects** or enter the name of the objects instance on which you need to grant read-only access to the user.

6. Click **Save**.
   A view-only administrator access control policy for one or more users or user groups is created.

Creating Bucket Administrator Role for Objects

You can create a bucket administrator policy to manage all buckets activities on objects instances using role-based access control (RBAC) from the Prism Central web console.

About this task

> Note: Users with only administrator privileges can do this task.

To create a buckets admin policy, do the following:

Procedure

1. Log on to the Prism Central web console as an Admin user, and click the **Entity** > **Administration** > **Roles**.
   The Roles page appears.
2. Click **Create Role**.
   The **Roles** page appears.
3. Enter the following details in the **Roles** page:
   a. **Role Name**: Enter a name for the role.
   b. **Description**: Enter a brief description for the role.
4. From the **Object Store** section, click **Change** next to the **Set Custom Permissions** option.
   The Custom Object Store Permissions window appears.

> Note: You can also enter Object Store in **Filter Entities** to go to the object store section.

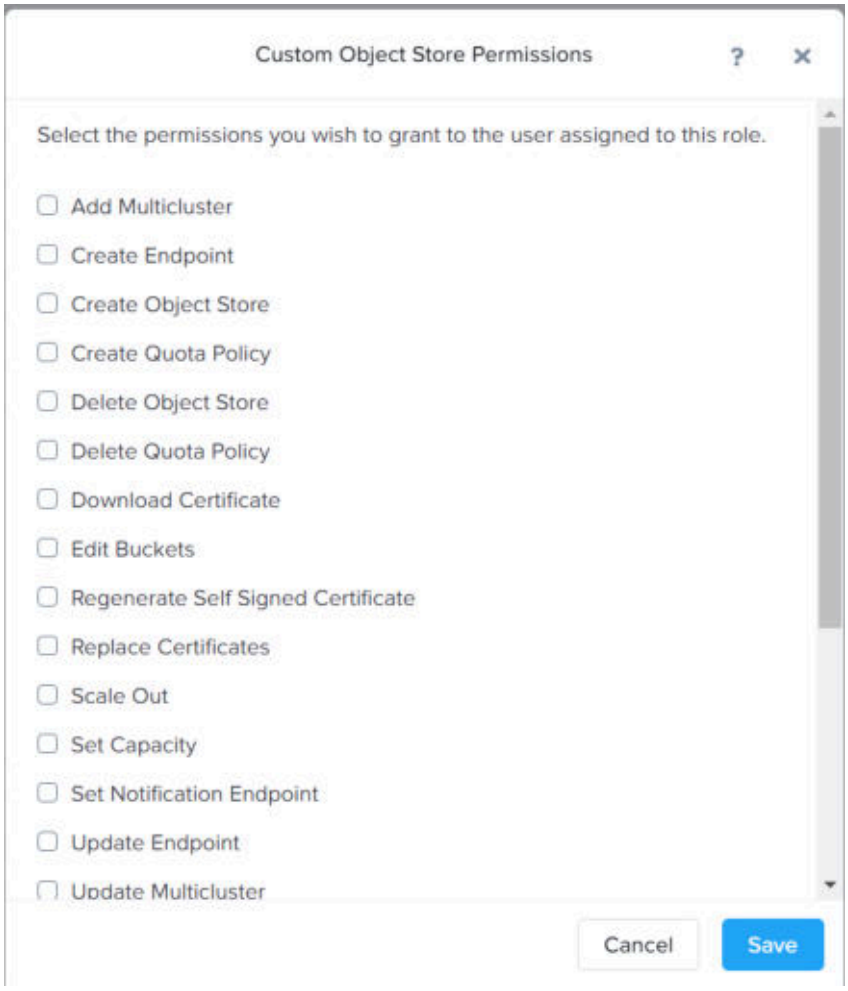

Figure 7: Custom Object Store Permissions

5. Select the following permissions from the list and click **Save**:

   - Edit Buckets
   - View Buckets
   - View Object Store

   Note:

   - When you select certain permissions, other additional permissions might automatically get added to the permission list. Read the on-screen instructions to understand the additional permissions that are automatically granted to the new custom role.
   - Users with Edit Bucket permission can perform bucket operations, such as create, update, and delete on the buckets available in the object stores that are assigned to the respective users.

6. Click the newly created role, and click **Actions** > **Manage Assignment**.
   The Role Assignment page appears.
7. Click **Add New** to assign the role to more users or user groups.
8. From the **Select Users or Groups** section, select the Windows AD group, and then select the user or user groups using the **Search User** field.

   Note:

   - You can add only those users who are part of an Active Directory group.
   - When a permission is granted to a role, the specific permission is not always inherited to the access control policy that is created with that role. This role needs to be assigned to a specific objects instance.

9. From the **Entities** section, select object store as the **Entity Type**, and then select the following:

   - Individual Entity
   - Select All Objects or enter the name of the Objects instance on which you need to grant custom permissions for the user.

10. Click **Save**.
    The Buckets Admin access control policy is saved and listed in the **Roles** page. The user now has access to perform bucket management tasks (create, update, and delete buckets) on all buckets in that objects instance.

    Note: Certain actions on buckets might need additional permissions to be assigned to the custom role. For example, to set a lifecycle policy based on tiering on a bucket, see Tiering section for the required permissions.

Creating Infra Administrator Role for Objects

An administrator can provide access to users or to user groups to manage all the existing objects instances without providing any write access to buckets on the objects instances, from the Prism Central web console.

About this task

> Note: Users with only administrator privileges can perform this task.

To create infra administrator access control policy, do the following:

Procedure

1. Log on to the Prism Central web console as an administrator user, and click **Entity** > **Administration** > **Roles**.
   The Roles page appears.
2. Click **Create Role**.
   The **Roles** page appears.
3. Enter the following details in the **Roles** page:
   a. **Role Name**: Enter a name for the role.
   b. **Description**: Enter a brief description for the role.
4. From the object Store section, select **Set Custom Permissions** and click **Change**.
   The Custom Object Store Permissions window appears.

   > Note: You can also enter Object Store in **Filter Entities** to go to the **Object Store** section.

5. Select the following permissions from the list and click **Save**:
   Add Multicluster, Create Quota Policy, Delete Quota Policy, Download Certificate, Regenerate Self Signed Certificate, Replace Certificates, Scale Out, Set Capacity, Set Notification Endpoint, Update Multicluster, Update Quota Policy, View Buckets, Edit Buckets, View Multicluster, View Notification Endpoint, View Object Store, and View Quota Policy.

   > Note: When you select certain permissions, other additional permissions might automatically get added to the permission list. Read the on-screen instructions to understand the additional permissions that are automatically granted to the new custom role.

6. Click the newly created role, and click **Actions** > **Manage Assignment**.
   The Role Assignment page appears.
7. Click **Add New** to assign the role to users or user groups.
8. From the **Select Users or Groups** section, select the Windows AD group, and then select the users or user groups using the **Search User** field.

   > Note:
   >
   > - You can add only those users who are part of an Active Directory group.
   > - When a permission is granted to a role, the specific permission is not always inherited to the access control policy that is created with that role. This role needs to be assigned to a specific objects instance.

9. From the **Entities** section, select object store as the **Entity Type**, and then select the following:

   - Individual Entity
   - Select **All Objects** or enter the name of the objects instance on which you need to grant custom permissions for the user.

     Note: When a permission is granted to a role, the specific permission is not always inherited to the access control policy that is created with that role. This role needs to be assigned to a specific objects instance.

10. Click **Save**.

    The Infra Administration access control policy is saved and listed in the **Roles** page. The user now has access to manage existing objects instances. However, the user cannot perform any bucket update operations, create any objects instances, or delete any objects instances.

Creating a Custom Role for Objects

An administrator can create custom roles and provide access to users or to user groups from the Prism Central web console.

About this task

Note: Only users with administrator privileges can do this task.

For information on creating a custom role, see *Creating a Custom Role* in the *Security Guide*.

To create custom access control policy (custom role), do the following:

Procedure

1. Log on to the Prism Central web console as an administrator user, and click **Entity** > **Administration** > **Roles**.

   The Roles page appears.

2. Click **Create Role**.
   The **Roles** page appears.

3. Enter the following details in the **Roles** page:

   a. **Role Name**: Enter a name for the role.

   b. **Description**: Enter a brief description for the role.

4. From the Object Store section, select **Set Custom Permissions** and click **Change**.
   The Custom Object Store Permissions window appears.

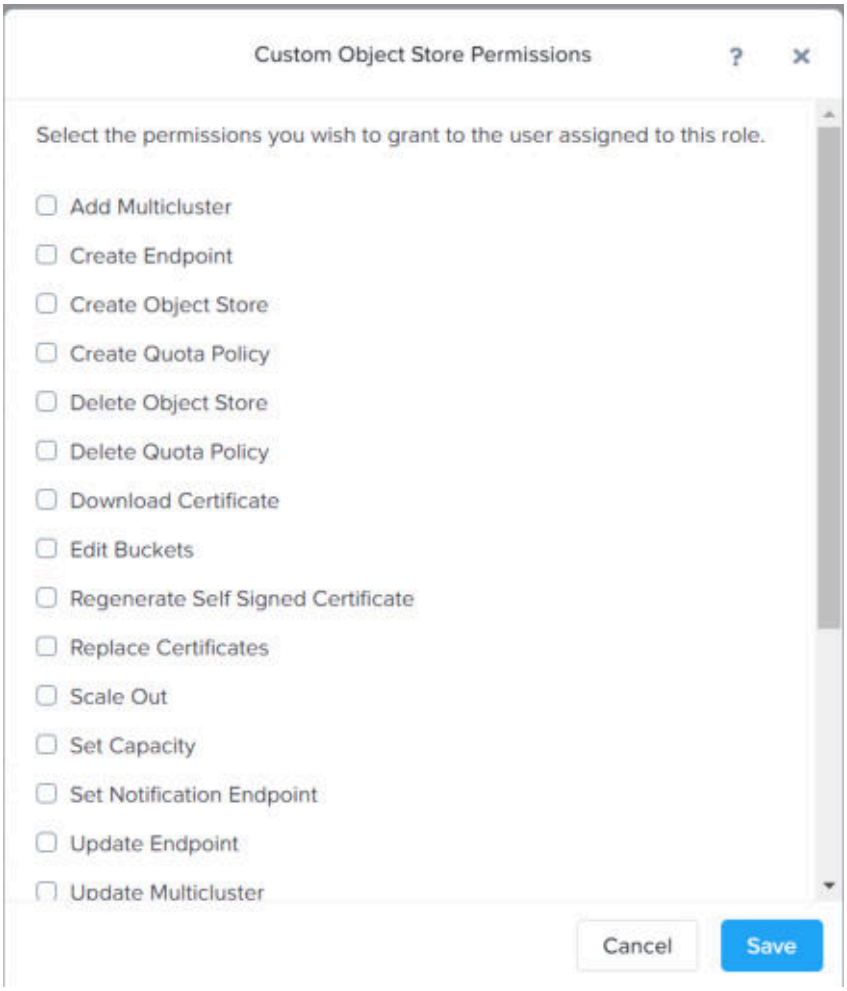

Figure 8: Custom Object Store Permissions Window

> Note: You can also enter Object Store in **Filter Entities** to go to the **Object Store** section.

5. Select the required permissions and click **Save**.

> Note: When you select certain permissions, other additional permissions might automatically get added to the permission list. Read the on-screen instructions to understand the additional permissions that are automatically granted to the new custom role.

6. Select the newly created role and click **Actions** > **Manage Assignment**.
   The Role Assignment page appears.

7. Click **Add New** to assign the role to users or user groups.

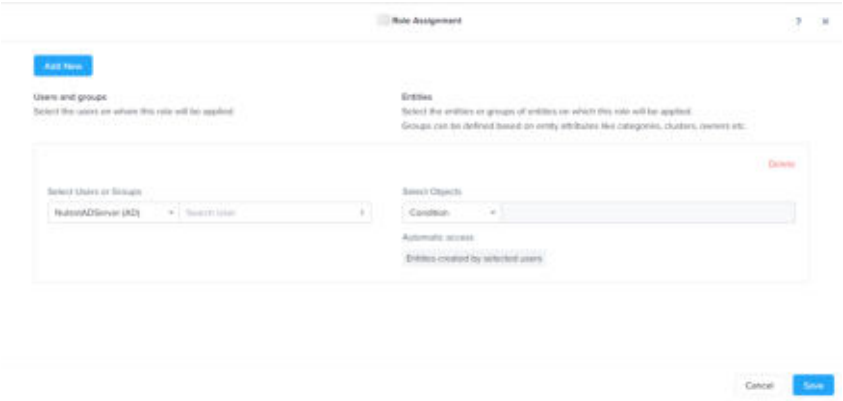

Figure 9: Role Assignment Window

8. From the **Select Users or Groups** section, select the required Active Directory group from the pull-down list, and then select one or more users or user groups using the **Search User** field.

   Note:

   - You can add only those users who are part of an Active Directory group.
   - When a permission is granted to a role, the specific permission is not always inherited to the access control policy that is created with that role. This role needs to be assigned to a specific Objects instance.

9. From the **Entities** section, select Object store as the **Entity Type**, and then select the following:

   - **Individual Entity**
   - Select **All Objects** or enter the name of the Objects instance on which you need to grant custom permissions for the user.

     Note: When a permission is granted to a role, the specific permission is not always inherited to the access control policy that is created with that role. This role needs to be assigned to a specific Objects instance.

10. Click **Save**.

    The custom access control policy is saved and listed in the **Roles** page. The user now has access to manage existing Objects instances granted to him in the policy.

DEPLOYMENT AND NETWORK PREREQUISITES

Ensure that you have met the deployment prerequisites and configured the network for **AHV** or **ESXi** before proceeding for an object store deployment.

Deployment Prerequisites - AHV and ESXi

Before deploying Object Store Services on AHV or ESXi, review this section carefully to ensure you have met the prerequisites. This section combines prerequisites for sites with internet access (online) and dark site (offline) deployments. Unless specified, the requirements are for both online and dark site (offline) deployments.

General Requirements

Ensure that your environment conforms to the following requirements before running Objects:

> Note: ESXi is supported only with Objects 3.0 or later versions.

- The hypervisor can be AHV or ESXi.
- Prism Element version 5.11.2 or later and Prism Central version 5.17.1 or later running in your environment.
- Recommended browser: Google Chrome
- Minimum of one node in a cluster running AHV or ESXi.

  > Note: Objects use no more than 12 vCPUs for each AHV or ESXi node.

- Ensure that no AHV or ESXi host or Prism Element or Prism Central upgrade is in progress while deploying Objects.
- Ensure that the Object Store domain should be dedicated for the object store deployment.

  For example, if the top level domain is `mycompany.com`, then the object store domain can be a sub domain such as `objectstore.mycompany.com`.
- Recommended that proxy should be able to reach the guest VM.
- Recommended to enable Pulse (only for site with Internet access).
- For an online installation, it is recommended to have a high speed and stable Internet connection.

  > Note: Image download times out after 90 minutes and the deployment fails.

- Recommended to upgrade to the latest version of the MSP controller for deployment in a dark site. Refer to Microservices Platform (MSP) on page 219.
- Ensure that the LCM web server is accessible from the Prism Element on which Objects is to be deployed in a dark site.
- Ensure that the LCM web server is accessible through the proxy, if set on Prism Central for the dark site deployment.
- Allow Prism Central and Prism Element to access the web server through port 80 for dark site deployment.

- Ensure to upgrade to MSP 2.0 or later versions and Objects 3.0 or later versions for expanding the object store cluster.
- Ensure to upgrade to Objects 2.1 or later versions for scaling out the object store cluster.

Network Requirements

Ensure to configure the following network requirements before running Objects:

- Configure Domain Name Servers (DNS) on both Prism Element and Prism Central.
- Configure Network Time Protocol (NTP) servers on both Prism Element and Prism Central.
- Set up the Virtual IP address and the data services IP address on the Prism Element where you plan to deploy Objects. Also, ensure that you set the DSIP on the PE cluster where the Prism Central is deployed.
- **AHV**: Ensure VLANs that are required internally for Object Store Services and externally for accessing the Object Store endpoints are configured on Prism Element correctly. Follow the guidelines provided in the Network Configuration on page 34 section.

  ESXi: Refer to ESXi Prerequisites.
- Ensure that you have an Internet connectivity for both Prism Element and Prism Central for online deployment. If you do not have Internet, refer to Object Store Deployment at a Dark Site (Offline Deployment) on page 78.

Prerequisites - ESXi

Before you start deploying Object Store Services on ESXi clusters, review this section carefully to ensure you have met the prerequisites. This section combines prerequisites for both sites with internet access (online) and dark site (offline) deployments.

Network Prerequisites

Ensure to configure the following network requirements before running Objects:

- Ensure that the ESXi network to be used for Objects deployment is available on all ESXi hosts.
  - If the network is configured using the standard vSwitch, ensure that all hosts have the network with the same name and VLAN.
  - If the network is configured using the Distributed Virtual Switch (DVS), ensure that all hosts are part of DVS.
- Ensure that you meet the Objects ESXi IPAM requirements. For more information, see ESXi Configuration on page 40.
- Ensure that only nodes belonging to a single Prism Element (PE) get added to an ESX cluster in vCenter.
- Ensure the VMware NSX network is present on all hosts of both the primary and secondary ESXi Metro Storage Clusters.

vCenter and ESXi Prerequisites

Ensure that the ESXi cluster complies with the following settings:

- ESXi clusters must be registered to a vCenter.

- Ensure that Distributed Resource Scheduler (DRS) and High Availability (HA) solutions are enabled on the cluster. For more information, see the *Create a Cluster* topic in the *VMware vSphere Product Documentation*.
- Ensure that the AOS cluster where you are deploying the object stores is mapped to a single ESXi cluster.
- Supported ESXi hypervisors versions are 6.5 or above.
- Objects version 3.0 and MSP 2.0 and later versions are supported.
- The minimum supported version of NSX-T is 2.4. Refer to KB article 000008545.

Prism Central VM (PCVM) Prerequisites

Ensure that the PCVM complies with the following:

- Ensure that PCVM is hosted and registered on the Prism Element (PE). It is required to provision Volume from the hosted PE to the MSP controller for image conversion.
- Ensure that the vCenter IP address is allowed in the proxy configuration and vCenter is connected to Prism Central.

  Worker VMs, Prism Central, and Prism Element must have a direct connection with the vCenter and not through a proxy server.

Network Configuration

The Nutanix Objects architecture uses two networks - **Objects Storage Network** (internal) and **Objects Public Network** (external).

Objects Storage Network is a virtual network used for internal communication between the components of an object store. **Objects Public Network** is a virtual network used by external clients to access the Object Store.

> Note:
>
> - You can have two virtual networks, each for Objects Storage Network and Objects Public Network, but it is not mandatory. You can have the Objects Storage Network and the Objects Public Network on the same virtual network. However, it is recommended to have the Objects Storage Network and the Objects Public Network on different virtual networks for production deployments.
> - It is recommended that the Objects Storage Network should be same as the CVM or the Hypervisor network. A single network enables the traffic between Objects and CVM to flow within the same network, hence avoiding cross-network hop which tends to be network constrained in some deployments. The traffic that flows between Objects and CVM is a function of the capability of the underlying AOS and is going to be significant in a dedicated objects deployment.
> - If you want to use different networks for Objects Storage Network and CVM network, ensure that the network bandwidth between the top-of-rack switch and the L3 device is fast enough to avoid network congestions. Alternatively, you can enable L3 functionality on the top-of-rack switch.

LACP and Link Aggregation

We recommend you to configure link aggregation with LACP and balance-tcp. Objects is a network-heavy workload and requires a high network bandwidth. You can achieve a faster

network by aggregating multiple physical links into one logical network interface that can be used for all traffics.

You require LACP and link aggregation to take full advantage of the bandwidth provided by multiple links. In Open vSwitch (OVS), it is accomplished through dynamic link aggregation with LACP and load balancing using balance-tcp. For more information, see *LACP and Link Aggregation* section in the *AHV Networking Guide*.

On ESXi based Objects deployments, you can also use the vSphere Distributed Switch with Route based on physical NIC load (LBT).

After entering maintenance mode for the desired host, configure link aggregation with LACP and balance-tcp using the following commands:

1. If upstream LACP negotiation fails, the default AHV host configuration disables the bond, thus blocking all traffic. The following command allows fallback to the active-backup bond mode in the AHV host in the event of LACP negotiation failure:

   ```
   nutanix@CVM$ ssh root@192.168.5.1 "ovs-vsctl set port br0-up other_config:lacp-fallback-ab=true"
   ```

2. In the AHV host and on most switches, the default OVS LACP timer configuration is slow, or 30 seconds. This value—which is independent of the switch timer setting—determines how frequently the AHV host requests LACPDUs from the connected physical switch. The fast setting (1 second) requests LACPDUs from the connected physical switch every second, thereby helping to detect interface failures more quickly. Failure to receive three LACPDUs —in other words, after 3 seconds with the fast setting—shuts down the link within the bond. Nutanix recommends setting lacp-time to fast on the AHV host and physical switch to decrease link failure detection time from 90 seconds to 3 seconds.

   ```
   nutanix@CVM$ ssh root@192.168.5.1 "ovs-vsctl set port br0-up other_config:lacp-time=fast"
   ```

3. Enable LACP negotiation and set the hash algorithm to balance-tcp.

   ```
   nutanix@CVM$ ssh root@192.168.5.1 "ovs-vsctl set port br0-up lacp=active"
   ```

   ```
   nutanix@CVM$ ssh root@192.168.5.1 "ovs-vsctl set port br0-up bond_mode=balance-tcp"
   ```

4. Enable LACP on the upstream physical switches for this AHV host with matching timer and load balancing settings. Confirm LACP negotiation using ovs-appctl commands, looking for the word "negotiated" in the status lines.

   ```
   nutanix@CVM$ ssh root@192.168.5.1 "ovs-appctl bond/show br0-up"
   ```

   ```
   nutanix@CVM$ ssh root@192.168.5.1 "ovs-appctl lacp/show br0-up"
   ```

5. Exit maintenance mode and repeat the preceding steps for each node and every connected switch port one node at a time, until you have configured the entire cluster and all connected switch ports.

AHV Configuration

You can configure and manage virtual networks through Prism Element and use these virtual networks to deploy an Object Store through Prism Central.

Once you have successfully added these virtual networks (with AHV IPAM enabled) in Prism Element for AHV, you can use these virtual networks to deploy an object store through Prism Central. For more information about network configurations and enabling AHV IPAM, refer to Network Configuration for VM Interfaces in *Prism Web Console Guide*.

The following section describes the Objects Storage (internal) and Objects Public (external) networks in more detail:

Objects Storage Network

An Object Store uses this private virtual network to communicate between services. The number of IP addresses required for deploying your Objects store instance varies according to the deployment size. For more information, see the *IP Address Consumption* section.

Requirements of a virtual network used for an Object Store:

- AHV IPAM must be enabled.
- Sufficient IP addresses must be available in the IPAM DHCP Pool. Refer to *IP Address Consumption - Based on Deployment Size*
- Two static IP addresses outside the DHCP pool are required for each object store. Later while deploying the object store from Prism Central, use the static IP addresses for the Objects Storage Network configuration based on the object store storage and resource requirements.

> Note: Objects internal services use the `10.100.0.0/16` and `10.200.0.0/16` subnets therefore the subnet used for internal interface IP addresses cannot conflict with either of these subnets.

Objects Public Network

An Object Store uses this virtual network (with AHV IPAM enabled) to allow access from external clients.

Requirements of a virtual network to access an Object Store externally:

- The virtual network must have AHV IPAM enabled.

  While configuring the virtual network, make sure you specify details such as VLAN ID, Network IP Address/Prefix Length, Gateway IP address, and DNS Servers.

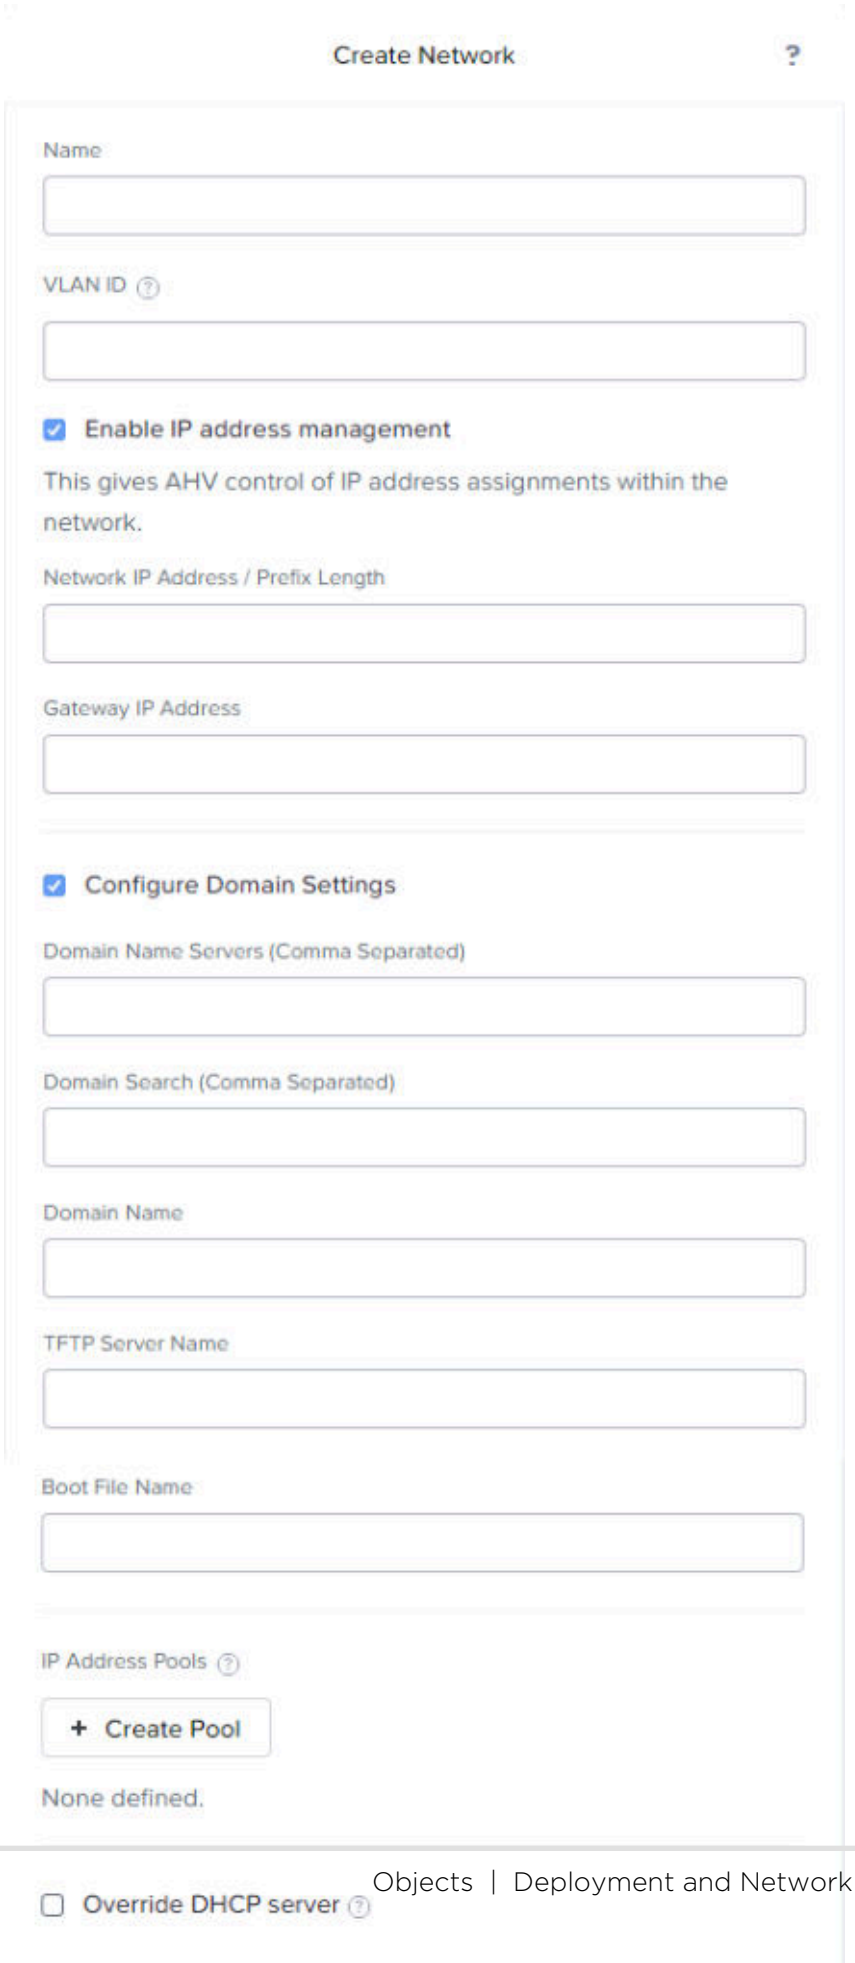
Create Network
?
Name
VLAN ID
Enable IP address management
This gives AHV control of IP address assignments within the network.
Network IP Address / Prefix Length
Gateway IP Address
Configure Domain Settings
Domain Name Servers (Comma Separated)
Domain Search (Comma Separated)
Domain Name
TFTP Server Name
Boot File Name
IP Address Pools
+ Create Pool
None defined.
Override DHCP server

- While deploying an object store from Prism Central, use up to four static IP addresses for the objects public network configuration based on the object store storage and resource requirements.

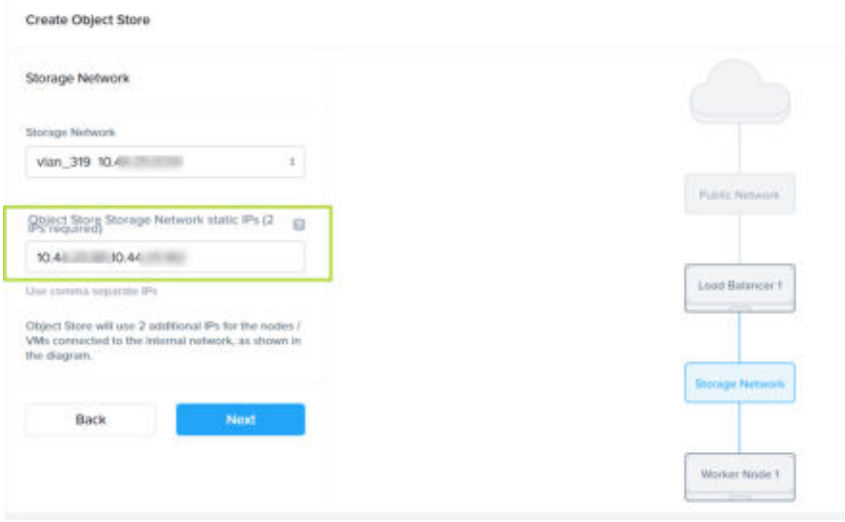

Figure 11: Network Configuration - Objects Storage IP addresses

Figure 12: Network Configuration - Objects Public IP addresses

Note: You can have two virtual networks, each for Objects Storage Network and Objects Public Network, but it is not mandatory. You can have the Objects Storage Network and the Objects Public Network on the same virtual network. However, it is recommended to have the Objects Storage Network and the Objects Public Network on different virtual networks for production deployments.

Note: All the IP addresses may not get used during the deployment. The number of IP addresses used depends on the size of your deployment. The unused IP addresses get reserved for future usage.

You can view a list of deployed object stores, and the general and networking details of the object stores. For example, you can view the Objects Public IP addresses for your deployment. For more information, see Viewing Object Store Deployments on page 60.

IP Address Consumption

This section describes the IP address consumption based on the deployment size.

Objects Public Network

Up to four static IP addresses are required for exposing the Objects S3 Endpoint based on the selected worker nodes. The client connects to these IP addresses for all the S3 requests. Object supports a maximum of four public IP addresses for each Object Store.

Objects Storage Network

The IP address consumption for the Objects Storage Network is as follows:

- From DHCP pool
  - Each Worker VM requires one IP address.
  - One IP address for each load balancer for communications with Objects workers.
  - Two IP addresses for High Availability of internal MSP services.
- Outside DHCP Pool - Two static IP addresses required for MSP DNS and API servers.

Number of Worker Nodes	AHV Managed Networks					
	Objects Storage Network					Objects Public Network
	From DHCP Pool				Outside DHCP Pool	Number of static IP addresses required for load balancers for communication with external S3 clients
	Number of IP addresses required for Worker VMs	Number of IP addresses required for the Load Balancer VMs to communicate with Worker VMs	Number of IP addresses for HA of internal MSP services	Number of IP addresses required for IPAM DHCP Server	Number of Static IP addresses required for MSP DNS and API Server	
1	1	1	2	1	2	1
2	2	1	2	1	2	1
3	3	2	2	1	2	2
4	4	3	2	1	2	3
5	5	4	2	1	2	4
N	N	4	2	1	2	4

Figure 13: IP Address Consumption - Based on Deployment Size

ESXi Configuration

For ESXi clusters, you can perform the Objects ESXi IPAM configuration.

You can do the Objects ESXi IPAM configuration from the **Objects service** > **vCenter Management** available in the Prism Central web console. For more information, see Managing vCenter for Object Service on page 43.

Objects requires you to add the IPAM range for the ESXi networks that you want to use for object store deployment.

Objects uses these ESXi networks for two purposes:

- To deploy Objects VMs that host the various Objects services (also referred to as **Objects Storage Network**).
- To deploy load balancer VMs that provide object store endpoint to the S3 clients (also referred to as **Objects Public Network**).

The following section describes the Objects Storage (internal) and Objects Public (external) networks in more detail:

Objects Storage Network

An Object Store uses this private virtual network to communicate between services. The number of static IP addresses required for deploying your Objects store instance varies according to the deployment size. For more information, see the *IP Address Consumption* section.

Requirements of a virtual network used for an Object Store:

- Add the IPAM range for the ESXi networks in the vCenter Management page.
- Sufficient IP addresses must be available. Refer to *IP Address Consumption - Based on Deployment Size*
- Subnet, Gateway, and DNS IP address to be used for the ESXi network. The provided values must be valid.

  Note: Objects internal services use the `10.100.0.0/16` and `10.200.0.0/16` subnets therefore the subnet used for internal interface IP addresses cannot conflict with either of these subnets.

Objects Public Network

An Object Store uses this virtual network to allow access from external clients.

Requirements of a virtual network to access an Object Store externally:

- Up to four static IP addresses that can either be part of or outside the IPAM range. Later while deploying the object store from Prism Central, use the static IP addresses

for the Objects Public Network configuration based on the object store storage and resource requirements.

Note: If the **Objects Public Network** is different from the **Objects Storage Network**, then only subnet and Gateway values are needed. IPAM range and DNS IP address values are optional.

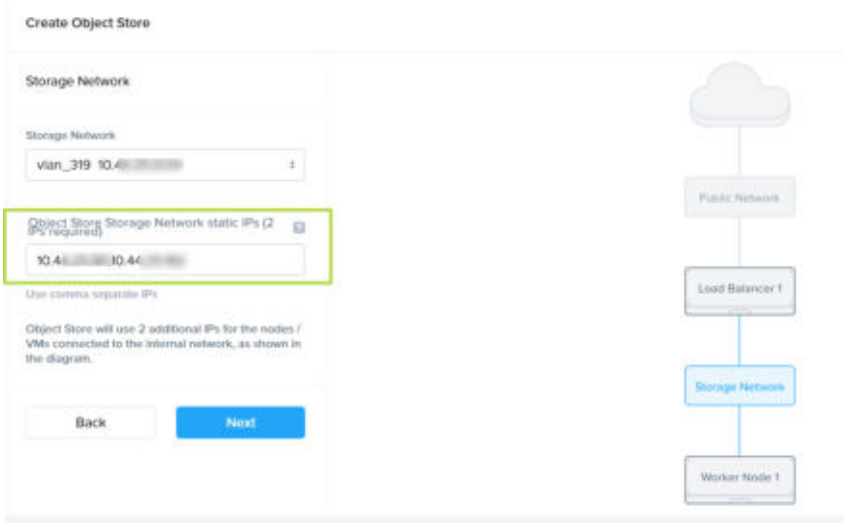

Figure 14: Network Configuration - Objects Storage IP addresses

Figure 15: Network Configuration - Objects Public IP addresses

Note: You can have two virtual networks, each for Objects Storage Network and Objects Public Network, but it is not mandatory. You can have the Objects Storage Network and the Objects Public Network on the same virtual network. However, it is recommended to have the Objects Storage Network and the Objects Public Network on different virtual networks for production deployments.

Note: All the IP addresses may not get used during the deployment. The number of IP addresses used depends on the size of your deployment. The unused IP addresses get reserved for future usage.

You can view a list of deployed object stores, and the general and networking details of the object stores. For example, you can view the Objects Public IP addresses for your deployment. For more information, see Viewing Object Store Deployments on page 60.

IP Address Consumption

This section describes the IP address consumption based on the deployment size.

Objects Public Network

Up to four static IP addresses are required for exposing the Objects S3 Endpoint based on the selected worker nodes. The client connects to these IP addresses for all S3 requests. Object supports a maximum of four public IP addresses for each Object Store.

> Note: The IP addresses can be within or outside the range of the IPAM network.

Objects Storage Network

The IP addresses get consumed from the IPAM range configured for the ESXi networks as follows:

- Each Worker VM requires one IP address.
- One IP address for each load balancer for communications with Objects workers.
- Two IP addresses for High Availability of internal MSP services.
- Two static IP addresses required for MSP DNS and API servers.

Number of Worker Nodes	Objects Storage Network					Objects Public Network
	Objects ESXi IPAM Range Configured for the ESXi Network					*Number of static IP addresses required for load balancers for communication with external S3 clients
	Number of IP addresses required for Worker VMs	Number of IP addresses required for the Load Balancer VMs to communicate with Worker VMs	Number of IP addresses required for HA of internal MSP services	Number of IP addresses required for IPAM DHCP Server	Number of Static IP addresses required for MSP DNS and API Server	
1	1	1	2	1	2	1
2	2	1	2	1	2	1
3	3	2	2	1	2	2
4	4	3	2	1	2	3
5	5	4	2	1	2	4
N	N	4	2	1	2	4

* Public access IP addresses can be within or outside the range of the IPAM Network

Figure 16: IP Address Consumption - Based on Deployment Size

Managing vCenter for Object Service

To deploy Object Stores on ESXi clusters, you need to provide the vCenter credentials and configure the IPAM for ESXi networks.

About this task

For AHV clusters, IPAM is configured in the Prism Element. For ESXi clusters, you can perform the IPAM configuration from the Objects > vCenter Management available in the Prism Central web console.

Managing vCenter for object service consists of two steps.

- Add the vCenter IP address and login credentials in the Object service within the Prism Central to create a trust relationship. Nutanix does not store the login credentials after the connection is established between the vCenter and Prism Central.

- Add IPAM for one or more ESXi networks.

  When deploying the Object Store instance, you need to define the Objects Storage Network and Object Public Network. You can use these pre-configured networks for Objects internal and public access.

To register the vCenter in the Prism Central, do the following:

Procedure

1. Log on to the source Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click **vCenter Management**.
3. Click **Add vCenter**.

Figure 17: vCenter Management Page

4. In the vCenter page, enter the IP address and login credentials of the vCenter.

   The details you enter are used to generate a certificate and build a trust relationship between the vCenter and Prism Central.

   Note: Nutanix does not store the login credentials after the connection is established between the vCenter and Prism Central.

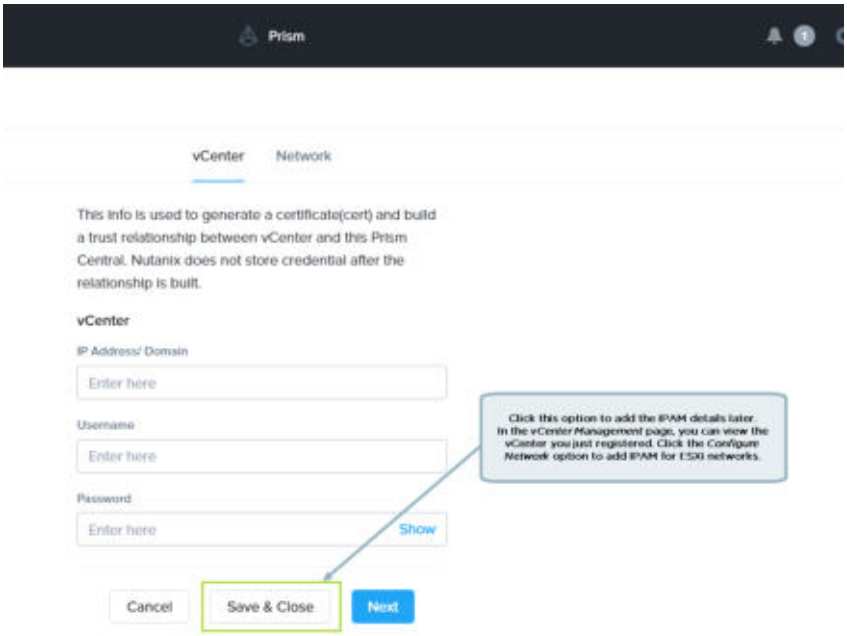

Figure 18: Register vCenter

   Note: For a user with non-administrator privileges to perform vCenter Registration, the administrator must create an access control policy for the non-administrator user on a role that must have the following minimum permissions:

   - Create Object Store
   - View Object Store

5. Click **Next** to add IPAM for the ESXi network.

   Else, click **Save & Close** to add the IPAM details later. In the vCenter Management page, you can click **Configure Network** in the **Actions** column to proceed with adding the IPAM for ESXi networks.

6. Click **Add Network**.
   The Add Network page appears.

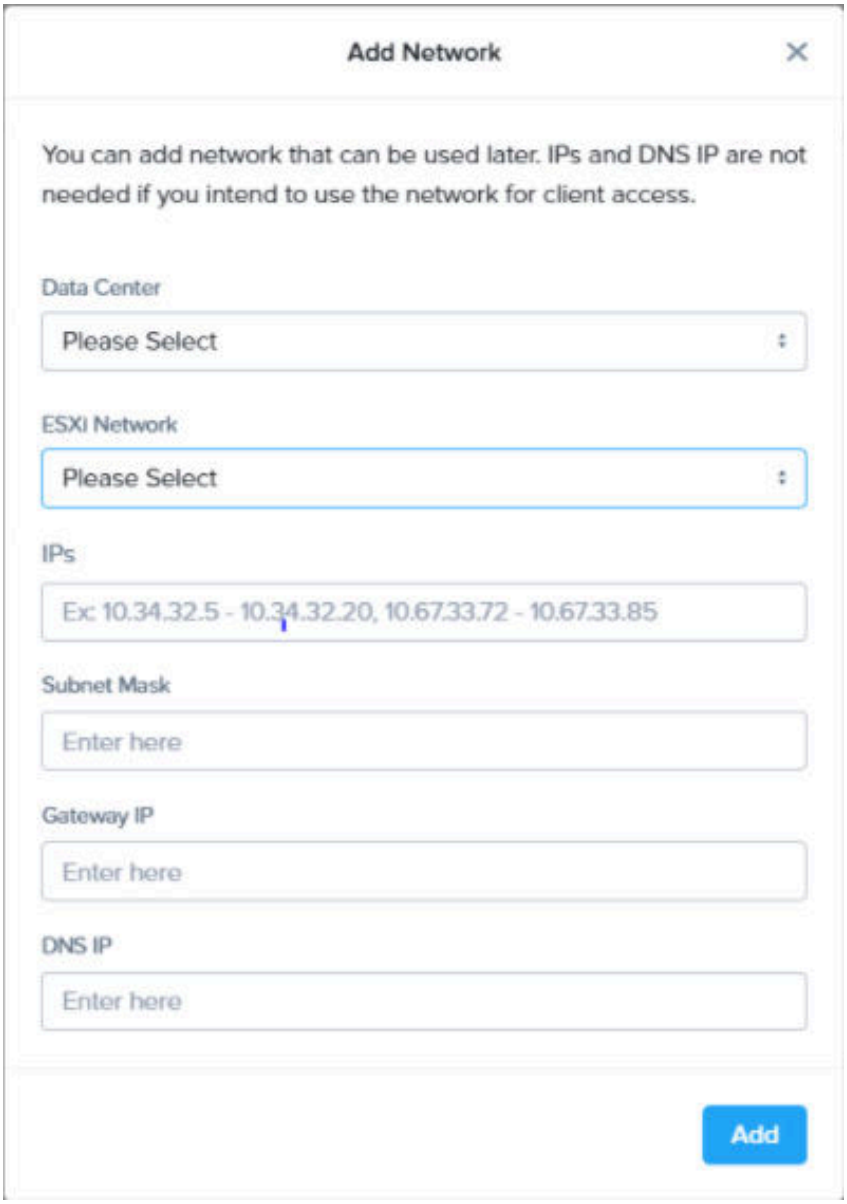

Figure 19: Add IPAM for ESXi Network

7. In the Add Network page, do the following:

   a. In the **Data Center**, select the data center where your ESXi cluster is located.

      The **ESXi Network** drop-down will get populated with a list of supported ESXi networks belonging to the data center.

   b. In the **ESXi Network**, select the ESXi network you want to use for deployment.

   c. Enter the IP address range, subnet mask, Gateway IP address, and DNS IP address.

      The IP address range you provide will be used for the Objects Storage Network. The IP address range and DNS IP address are optional if the IPAM will be used only for the Objects Pubic Network.

   d. Click **Add** to complete your IPAM configuration.

      If you want to add more networks, click **Add Network** in the Configure page and enter the details.

      > Note: It is recommended to use separate networks for the Objects Storage Network and Objects Public Network.

   e. Click **Update** to close the **Configure** page.

      In the vCenter Management page, in the **Network** column, you can view the count of IPAM.

      In the **Actions** column, you can click the **Delete** option to delete the vCenter.

      > Note: Even if you delete a vCenter, the IPAM details would still be available. If you add the deleted vCenter again, the IPAM details previously added gets recovered.

What to do next

You can start with the deployment of the Object Store on the ESXi clusters. For more information, see Object Store Service Deployment.

Shared Versus Single Network

You can configure Objects to use a single network for all services or separate the subnets used for internal communications and client communications.

You can have two virtual networks, each for Objects Storage Network and Objects Public Network, but it is not mandatory. You can have the Objects Storage Network and the Objects Public Network on the same virtual network. However, it is recommended to have the Objects Storage Network and the Objects Public Network on different virtual networks for the production deployments.

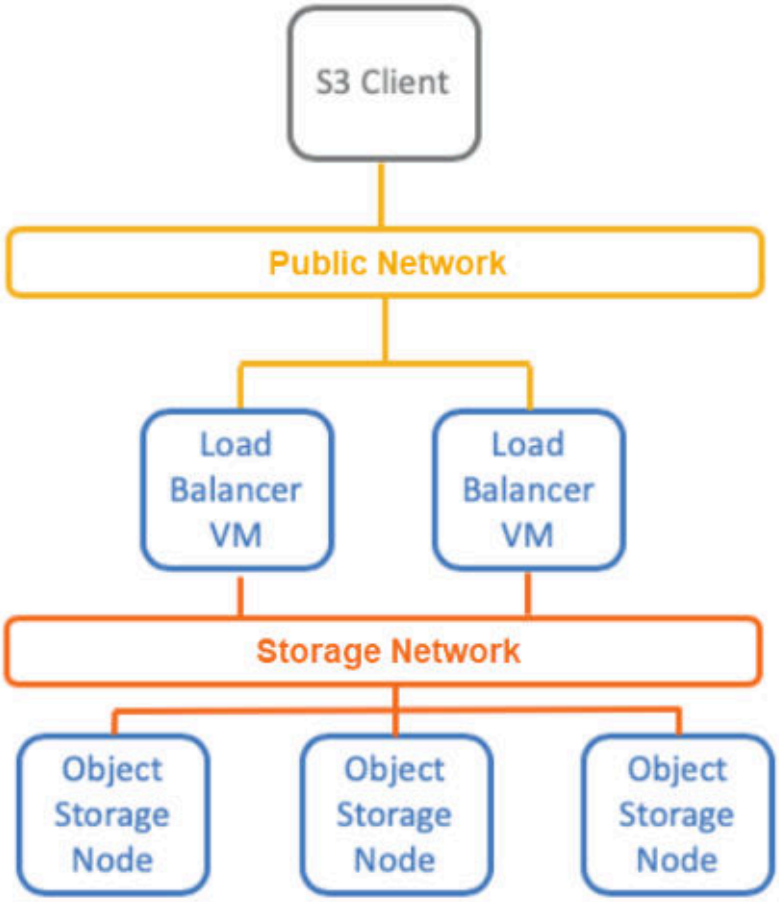

Figure 20: Objects Network Architecture - Using Two Separate Networks (Storage and Public)

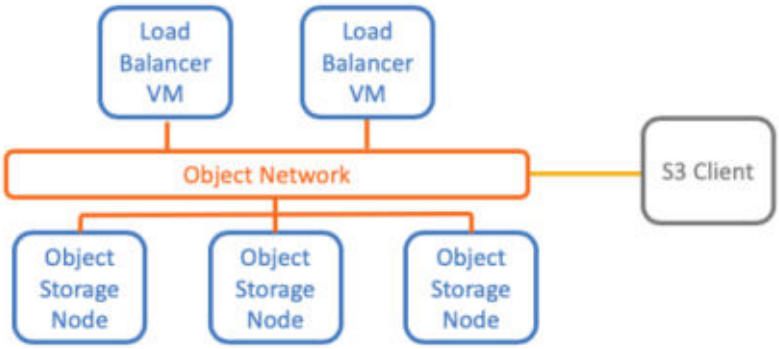

Figure 21: Objects Network Architecture - Using One Network (Storage and Public)

URL and Port Requirements

This section describes the URL and port requirements for Objects instance deployment.

URL Requirements

The following URLs are used by the Objects server:

- Prism Central must have access to `download.nutanix.com`, Amazon Elastic Container Registry (ECR) at `464585393164.dkr.ecr.us-west-2.amazonaws.com`, `api.ecr.us-west-2.amazonaws.com`, and `prod-us-west-2-starport-layer-bucket.s3.us-west-2.amazonaws.com`.
- Gold image is downloaded from *download.nutanix.com*.

> Note: URLs are not required for the dark site deployment.

Port Requirements

Refer to the Port Reference Guide for the required ports for your Objects deployment.

> Note: Internal network IP ranges such as 10.100.x.x and 10.200.x.x could be seen on the firewall if the inter MSP VM traffic passes through the firewall.

Refer to the following diagrams to understand Objects network architecture.

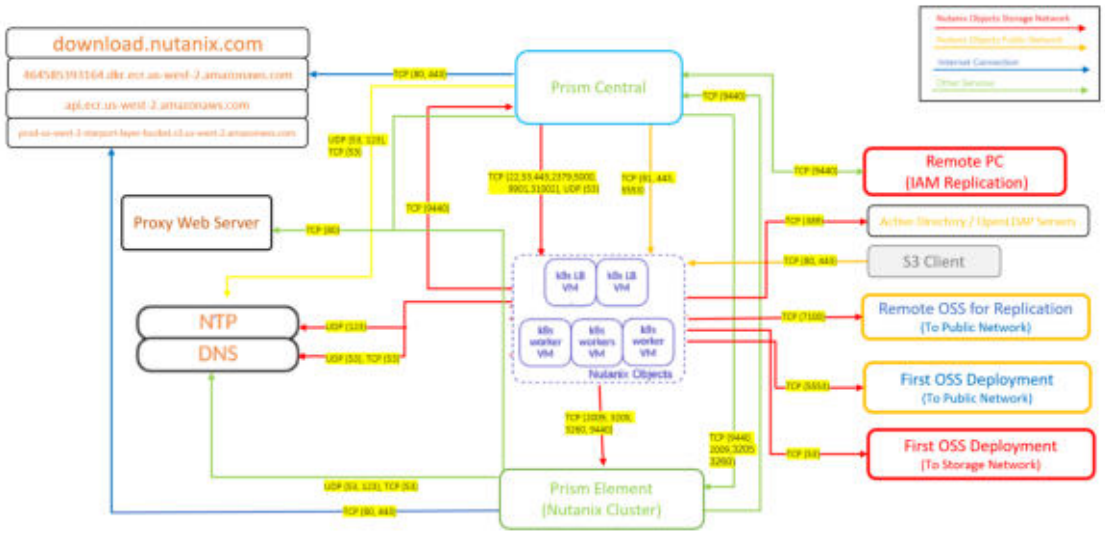

Figure 22: Objects Network Diagram with Ports (AHV)

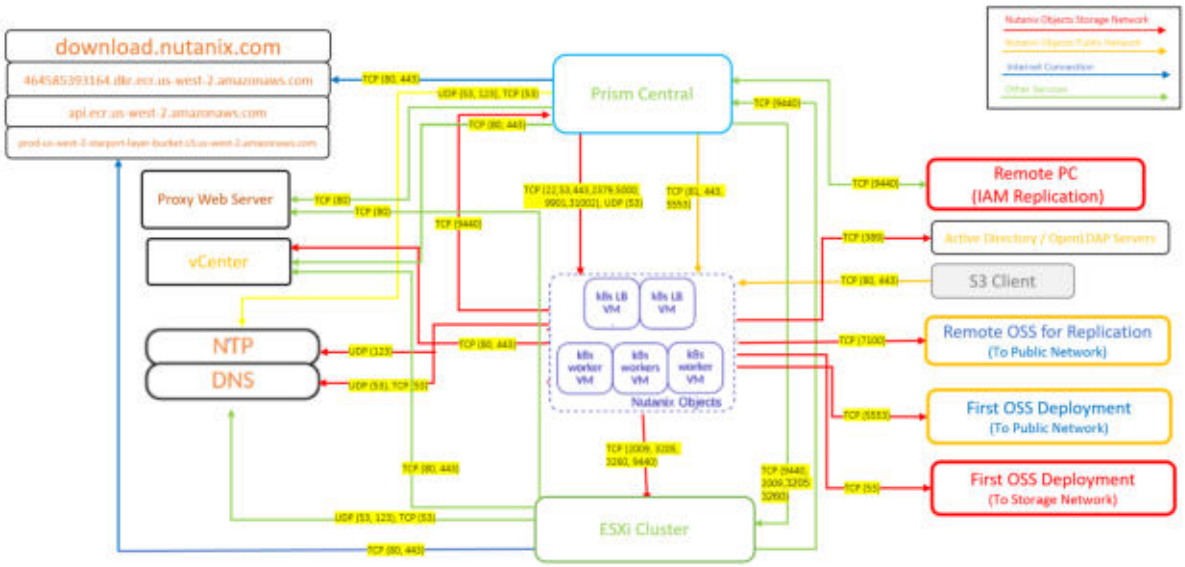

Figure 23: Objects Network Diagram with Ports (ESXi)

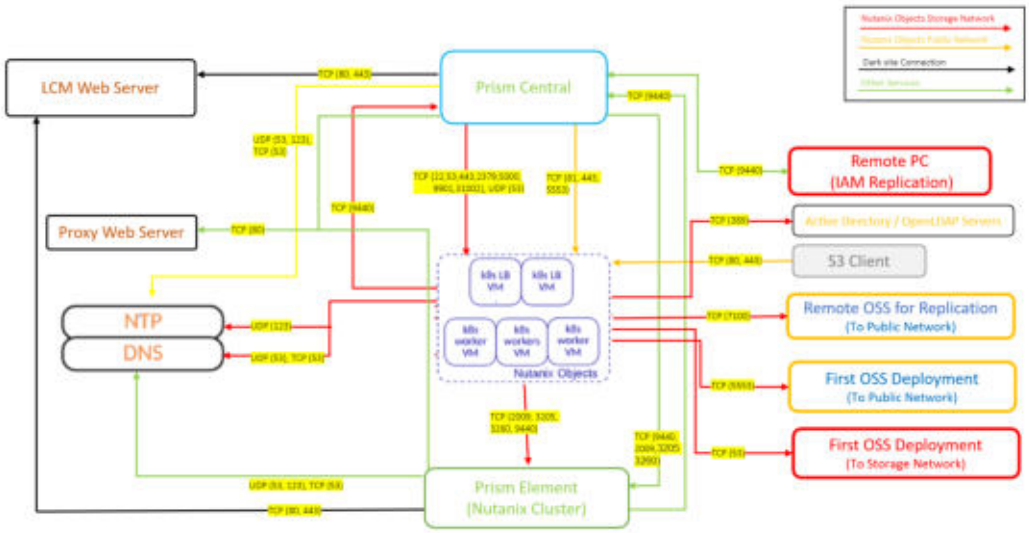

Figure 24: Objects Network Diagram with Ports (Dark Site)

Limitations

The following section lists the limitations for Objects.

System Limitations

- Once an object store gets deployed, you cannot change the Data Services, Controller VM, Microservices Platform (MSP), and Prism Central IP addresses.
- Prism Central and Prism Element de-registration and reregistration is not supported.

Limitations of NFS

This section lists the limitations of multiprotocol access.

- You can enable NFS access only at the time of bucket creation.
- NFS-enabled buckets are exposed as NFS shares, which can be mounted by the NFS clients.
- You can perform update (only multiprotocol access configurations), delete, and share actions on NFS buckets.

  > Note: You can delete NFS-enabled buckets only if it does not have any object and any directory explicitly created from the NFS protocol.

- You cannot enable other S3 bucket features, such as Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications on NFS-enabled buckets.
- You can only create symbolic links through NFS. Hard links are not supported.
- You can rename files and links; however, you cannot rename a directory. After rename, the file-handle of the renamed file changes.
- In addition to the S3 object naming convention, following are some cases, which are not allowed for an object name in the NFS-enabled bucket:
  - Object name cannot be . or ..
  - Object name cannot start with /,./ or ../
  - Object name cannot contain //, /./ or /../ pattern.
  - Object names cannot end with /. or /..
  - If an object name contains a directory hierarchy, then each component in the object name separated by / cannot be more than 255 bytes in length.
- Objects with / as suffix or objects-browser created folders, which will appear as directory in the NFS namespace cannot be deleted from the NFS protocol.
- If the objects are created from the S3 protocol in such a way that they end-up creating a conflict in the object and directory name, then only the file will be visible in the NFS namespace and the directory will get hidden. Only on renaming or removing the file to some other name, the directory will become visible again .

  For example, suppose we create two objects **`dir1/dir2/file`** and **`dir1/dir2`** from the S3 protocol. As we can see, **`dir2`** is a directory as well as file, so it creates a conflict. Now, while traversing this namespace from NFS protocol, once we are inside **`dir1/`** directory, only the object **`dir2`** would be visible and not the subdirectory **`dir2/`**. So, it is recommended to not use conflicting names for objects while creating them from S3 protocol.

OBJECT STORE SERVICE DEPLOYMENT

Objects is a highly available and distributed Object Store Service which can store petabytes of data.

To start using Objects, you need to deploy the Object Store Service. Refer to Creating or Deploying an Object Store (Prism Central) on page 52. You can also perform offline deployment (site without Internet access) of an object store. Refer to Object Store Deployment at a Dark Site (Offline Deployment) on page 78.

You can deploy multiple Objects instance on a single Prism Element registered to a Prism Central, or you can register multiple Prism Elements to one Prism Central and deploy as many object stores on each of these Prism Elements provided you have sufficient storage.

You can do the following operations after deploying an Objects instance:

1. Configure the directory and generate the access key. Refer to Directory Configuration and Access Key Generation on page 81.
2. Create buckets within the object store. Refer to Creating and Configuring an S3 Bucket on page 92
3. Upload objects and meta-data to the buckets by using the S3 APIs. Refer to Supported APIs on page 202.

Note:

- After registering Prism Element to Prism Central, wait for 10 minutes before starting the deployment.
- Parallel Object Store Service deployments are not supported.
- If your deployment fails due to precheck failures, you can resume the deployment after fixing the configuration.
- For objects containers, the EC delay will be reduced from 7 days to 3 days for old and new deployments.

Object Store Instance Naming Conventions

The name of an object store must conform to the following rules:

- Be unique across all existing object store names in Objects.
- Begin with a letter, and end with a letter or number.
- Can contain alphanumeric or hyphen characters.
- Not contain any special character other than hyphen.
- Minimum of 1 and a maximum of 16 characters long.

Note: You cannot change the name after creating the object store.

Creating or Deploying an Object Store (Prism Central)

Before creating a bucket and uploading objects, you must deploy an object store. You can deploy multiple Object Store Services on a single Prism Element registered to a Prism Central, or you can deploy multiple object stores on each Prism Element accessible from Prism Central provided you have sufficient storage.

Before you begin

Make sure you satisfy the deployment prerequisites before starting the deployment. Refer to Deployment Prerequisites - AHV and ESXi on page 32.

A non-admin user can create or deploy an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Create Object Store
- View Cluster
- View Subnet (only for AHV)

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

In addition to the above permissions, the user must also have:

- View Cluster access to the Prism Element on which the Objects instance will be deployed
- View Subnet access to the managed networks that will be used for deployment if the Prism Element runs on AHV

> Note: When you create an access control policy, select Cluster as the Entity type, and select Individual Entity, and then select All Clusters or the particular cluster on which you need to deploy the Object Store.

The user, who creates an Objects instance, gets full access to that Object instance automatically. This full access can be later revoked or modified by the administrator user.

About this task

> Note:
>
> - You cannot deploy a single Object Store Service across multiple Prism Elements.
> - If you exit by clicking **X** while creating an object store, the entered input and the pre-checks status are not saved.
> - Deployment can take a minimum of 30 minutes.
> - A new container with `msp-<uuid>` is created for Objects deployment on ESXi. This container will be used for downloading VM images for MSP workers.

To create an object store, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click **Create Object Store**.

3. In the Create Object Store: Prerequisites window, click **Continue** if you fulfil the prerequisites for AHV or ESXi.

4. On the **Object Store Details** section, do the following, and then click **Next**:

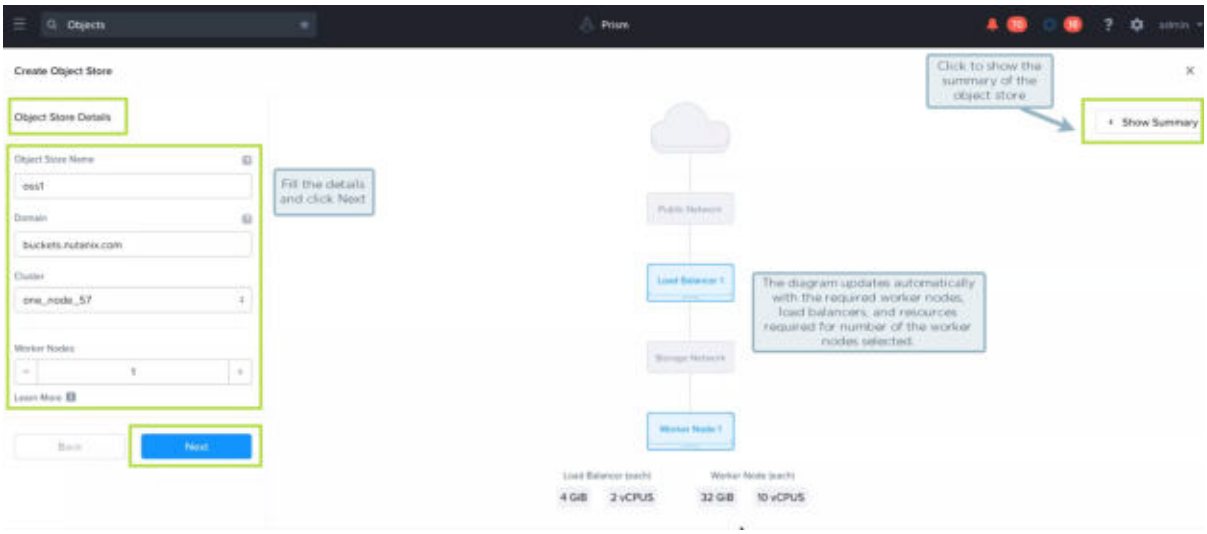

Figure 25: Object Store Details

- You can view the summary of the object store by clicking **Show Summary** on the right pane.
- You can hover over the help icon to get more information about the respective fields.
- The diagram updates automatically with the required worker nodes, load balancers, and resources required for number of the worker nodes selected.
- Active component will be highlighted in the diagram.

a. **Object Store Name**: Enter the name of the object store you want to create.
For guidelines on choosing a compliant name, refer to Object Store Naming Conventions.

b. **Domain**: Enter the domain name.

> Note: This domain name is the default domain name for all the object store in that cluster.

Domain naming conventions:

- Must contain at least one dot.
- Can contain either alpha numeric, hyphen, or underscore characters.
- Cannot start or end with a hyphen or underscore or dot.

c. **Cluster**: Select the cluster where the object store will be deployed.

d.

> Note:
>
> - The configured worker nodes cannot exceed the worker nodes of the cluster.
> - Resources cannot be reduced once added.
> - Actual performance depends on a variety of factors.

Worker Nodes: Add the number of worker nodes.

- Each VM is assigned 10 vCPUs and a DHCP IP address.
- A minimum of 10vCPUs and 32 GiB of memory is required.
- Each click on plus for Worker Nodes add 10vCPUs and 32 GiB of memory.
- vCPU and memory are linked. vCPU has to be in multiple of 10 and memory in multiple of 32.

5. The Storage Network section, do the following, and then click **Next**:

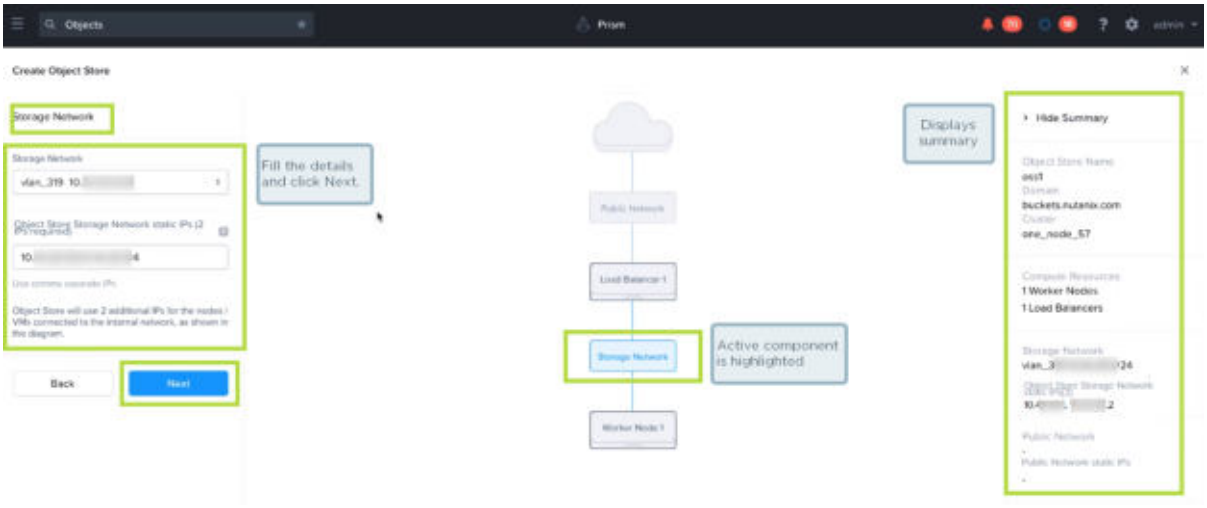

Figure 26: Storage Network Details

a. **Storage Network**: Select Objects storage network that is used for the internal communication between the components of an object store.

   For more information on the IP address requirements according to the deployment size, see Network Configuration.

b. (Only for AHV) **Object Store Storage Network Static IPs (2 IPs required)** : Enter two storage network IP addresses separated by a comma.

   - These storage network IP addresses are within the Objects Storage Network.
   - Object Store will use two additional IP addresses for the nodes or VMs connected to the internal network.

   > Note: For ESXi, these two internal IP addresses are not required and selected automatically from the IPAM range configured for ESXi networks.

6. On the Public Network section, do the following and then click **Save & Continue**:

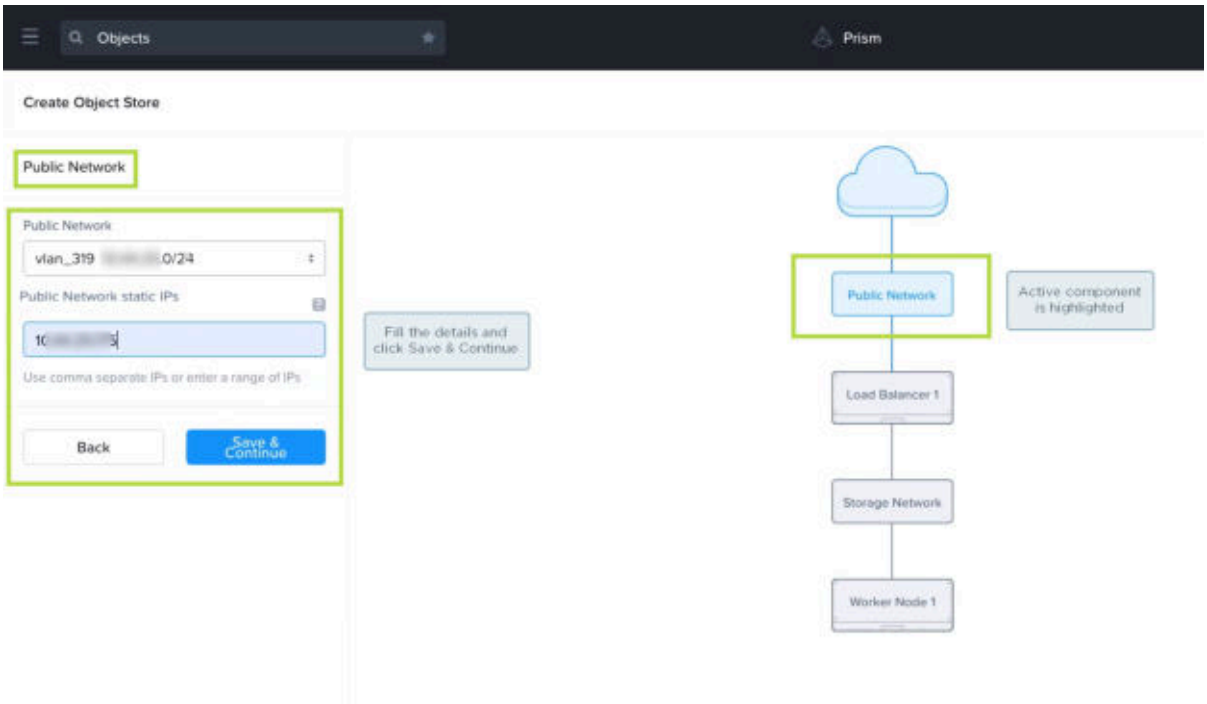

Figure 27: Public Network Details

a. **Public Network**: Select Objects public network that is used to allow access to the Object Store from the external clients.

   This VLAN should have up to four IP addresses in the usable IP address range. This network can be the same as the Storage Network. For more information, see Shared Versus Single Network on page 47.

b. In the **Public Network Static IPs**: Enter the public access IP addresses (one for each load balancer) separated by a comma or as an IP address range.

   For example, if one Load Balancer is used, then only one IP address is required. Also, you can enter the IP addresses in a range, `10.2.3.1-10.2.3.4`, or separated by a comma `10.2.3.1, 10.2.3.2, 10.2.3.3, 10.2.3.4`

   AHV - These IP addresses are within the Objects Public Network and used to access the object store.

   ESXi - The public access IP addresses can be within or outside the range of the IPAM network.

   For more information about network configurations, refer to Network Configuration on page 34.

You can click **Save for Later** if you wish to continue with the deployment later. The object store will be saved in the list of object stores. Select the object store, and then click **Actions** > **Complete Deployment** to complete the deployment.

Pre-checks starts before the deployment begins. A list of checks performed is displayed in the UI. Also, a VM image named predeployment_port_vm, and two VMs named predeployment_objects_public and predeployment_objects_storage are created.

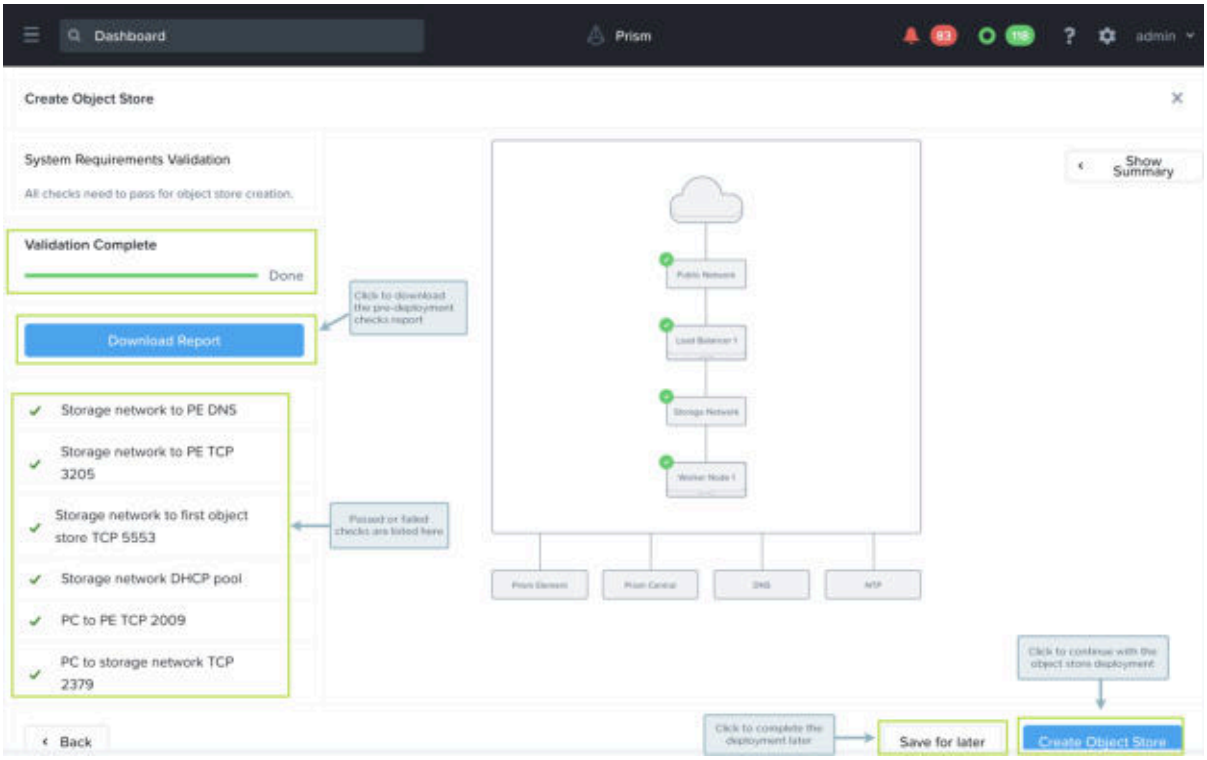

Figure 28: Pre-deployment checks status and report

7. Note: All pre-checks must pass to start the deployment.

Depending upon the pre-check result, do the following:

» If the pre-checks are passed, click **Download Report** to download the report, and then click **Create Object Store** to start with the object store deployment.

The report contains the name of the check, status and message.

» If the pre-checks fail, an error message is displayed in the UI. Click **Download Report** to download the report. A `Fail` status is displayed next to the check name with a message.

The object store is saved and you can complete the deployment after fixing the failed checks.

Once the pre-checks are passed, you can view the status of the object store deployment.

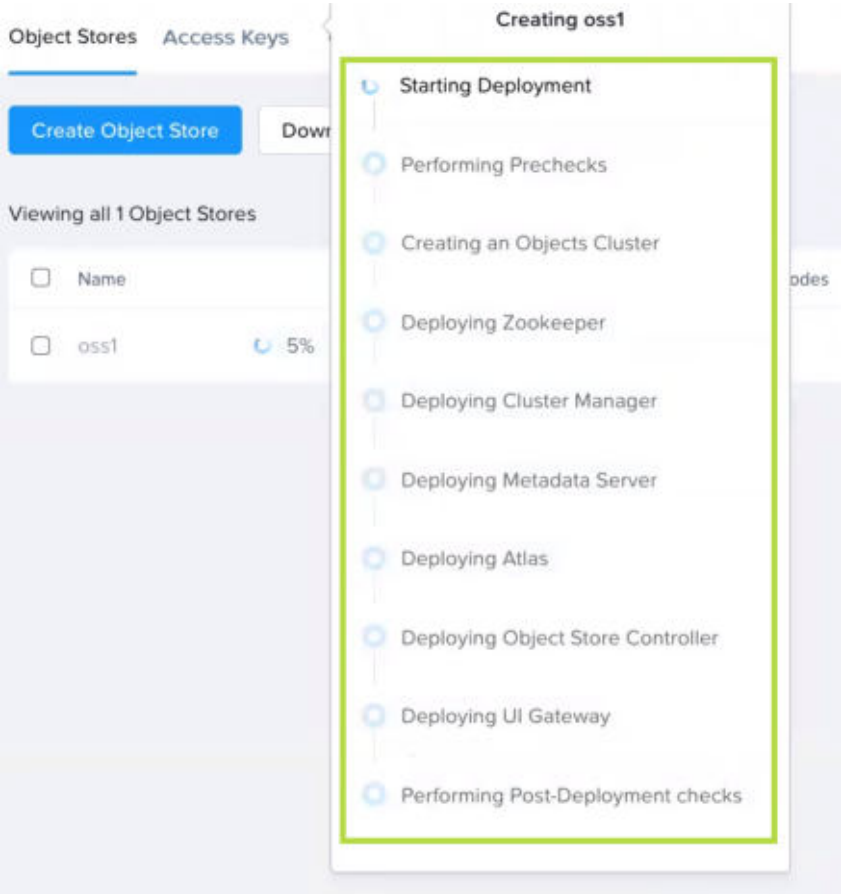

Figure 29: Deployment Steps

You can view the deployment progress in percentage and each step in the grid by hovering over the loading icon.

Note: Use an encrypted cluster for encrypting the bucket.

Warning: Do not delete MSP VMs (created with a prefix in Objects deployment) from the vCenter or Prism Central. There are no checks to identify the deletion of MSP VMs. You can identify the MSP VMs by their name. The naming conventions are as follows:

- MSP Worker VM: `<deployment_name>-XXXXXX-default-N`
- Load Balancer MSP VM: `<deployment_name>-XXXXXX-XXXXXXXXXX-envoy-N`

What to do next

After deploying the object store, you can perform the following:

- Configure directory and generate access key. Refer to Directory Configuration and Access Key Generation on page 81.

- Access the endpoint provided as part of Client Access Network (S3 Endpoints) by using HTTP or HTTPS. Refer to Access Objects Endpoints on page 62.
- Create and configure buckets. Refer to Creating and Configuring an S3 Bucket on page 92.
- Create objects using S3 APIs. Refer to Supported S3 APIs
- Expand the object store storage. Refer to Expanding Storage for an Object Store on page 64
- Scale out an object store. Refer to Scaling Out an Object Store on page 68
- Share buckets. Refer to Sharing a Bucket on page 123.

Viewing Object Store Deployments

You can view a list of deployed object stores, and the general and networking details of the object stores.

Before you begin

A non-admin user can view an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have minimum permission to view an object store and the non-admin user must be assigned to specific objects instances.

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

To view the object store, do the following:

Procedure

1. Log on to the Prism Central web console.
2. Click the **Entity** menu > **Services** > **Objects**.

Figure 30: Object Stores List View

A list of existing object stores appears. The following steps describe the fields that appear in the object store table. You can click on the name of an object store to open the object store

in a new window. You can click **View By** to view the **General**, **Usage** or **Networking** details of an object store. A dash (-) is displayed in the field when a value is not available or applicable.

General view:

- **Name**: Displays the name of the object store.
- **Version**: Displays the version of Objects in which the object store was created.
- **Domain**: Displays the domain of the object store.
- **Nodes**: Displays the number of nodes of the object store.
- **Usage (Logical)**: Displays the object store usage in GiB or TiB.
- **Buckets**: Displays the number of buckets in an object store.
- **Objects**: Displays the number of objects in an object store.
- **Alerts**: Displays the alerts for a particular object store.

  For more information, refer to Viewing Alerts on page 183.
- **Notifications**: Displays notifications if any.
- **Objects Public IPs**: Displays the endpoints or the IP addresses used by the client.

  You can access these endpoints by using HTTP and HTTPS protocols. For more information, refer to Access Objects Endpoints on page 62.

Networking view:

- **Name**: Displays the name of the object store.
- **Cluster**: Displays the cluster in which object store is deployed.
- **Objects Public Network**: Displays the VLAN required for accessing Object Store endpoints externally.
- **Objects Public IPs**: Displays the public IP address of Objects. This is used to access Objects Browser.
- **Objects Storage Network**: Displays the VLAN required internally for deploying Object Store Services on Prism Element.
- **Storage Network Static IPs**: Displays the internal configured IP addresses.

Usage view:

- **Name**: Displays the name of the object store.
- **Usage (Logical)**: Displays the cluster in which object store is deployed.
- **Local Usage**: Displays the amount of data stored locally in the Objects cluster.
- **Tiered Usage**: Displays the amount of data tiered out and stored remotely based on lifecycle policies.
- **Licensed Usage**: Displays the usage from a single object store accounted against the licensed capacity. It is calculated as (Local Usage + Tiered Usage).

  Note: Data tiered to another Objects endpoint is not included.

For information about Network Configurations, refer to Network Configuration on page 34.

Access Objects Endpoints

Objects endpoints are the entry point to Objects. You can access these endpoints by using the HTTP or HTTPS protocols through any third party clients which support S3 APIs.

Accessing Buckets and Objects within an Object Store Instance

Objects supports path-style and virtual hosted-style bucket access.

- Path-Style Access: The path-style syntax requires that you use the endpoint when attempting to access a bucket, and the request specifies a bucket by using the first slash-delimited component of the Request-URI path.

  For example, if you have a bucket with bucket name as **bucket-name** and the object name as **example.jpg**, and you want to use the path-style syntax. Following is the correct request:

  ```
  PUT /bucket-name/example.jpg HTTP/1.1
  Host: object-store-name.domain-name
  ```

- Virtual Hosted-Style Access: The virtual hosted-style syntax is used to address a bucket in a REST API call by using the HTTP Host header. This method requires the bucket name to be DNS-compliant.

  For example, if you have a bucket with bucket name as **bucket-name** and the object name as **example.jpg**, and you want to use the virtual hosted-style. Following is the correct request:

  ```
  PUT /example.jpg HTTP/1.1
  Host: bucket-name.object-store-name.domain-name
  ```

  For virtual hosted-style access, allow the Objects FQDN in the DNS server with the wild card allowlist. For example, following are the expected DNS entries for the Objects endpoint.

  ```
  objects.subdomain.example.com.        IN      A       192.168.5.101
  objects.subdomain.example.com.        IN      A       192.168.5.102
  objects.subdomain.example.com.        IN      A       192.168.5.103
  objects.subdomain.example.com.        IN      A       192.168.5.104
  *.objects.subdomain.example.com.      IN      A       192.168.5.101
  *.objects.subdomain.example.com.      IN      A       192.168.5.102
  *.objects.subdomain.example.com.      IN      A       192.168.5.103
  *.objects.subdomain.example.com.      IN      A       192.168.5.104
  ```

Deleting an Object Store

You can delete both successful and failed object store deployments. However, for deletion of successful deployment, ensure to first delete the objects and buckets within that object store.

Before you begin

A non-admin user can delete an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Delete Object Store
- View Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

Note:

- Deployments in progress cannot be deleted.
- You cannot delete the primary cluster without deleting the secondary cluster as the primary object store cluster hosts all the common services. To delete the primary cluster, ensure all the secondary clusters are deleted.
- Multi-cluster containers are not deleted on secondary Prism Elements if the Objects deployment is not successful.
- If a failure occurs while replacing an SSL certificate for an object store, you cannot delete that object store deployment; however, you can try replacing the SSL certificate again.
- Deletion of an object on the local object store may not immediately lead to space reclamation on the AWS S3. Space reclamation happens when all the source objects mapped to the corresponding objects in the AWS S3 gets deleted.

To delete an object store deployment, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. In the Object Stores table, select the object store which you want to delete, and then click **Actions** > **Delete**.

Figure 31: Delete Object Store Window

3. s
4. On the confirmation dialog box, click **Delete**.
   A message appears to confirm the object store deletion.

Object Store Expansion

This section describes the expansion options available with object store.

Storage Expansion

If your existing object store storage is less than 85% full or you want to use a different cluster for object store storage, you can expand the storage of your object store cluster by adding new nodes to the existing cluster or by adding additional clusters.

For example, if you have an object store deployed on a 10 TB cluster named Cluster1 and your storage is almost or less than 8.5 TB full, then you can either add a node to the existing Cluster1 or add a new cluster or existing clusters (such as Cluster2 and Cluster3) if they have sufficient storage capacity. In this example, Cluster1 is the existing primary cluster on which the object store was initially deployed and hosts the worker VMs for the object store. Cluster2 and Cluster3 are secondary clusters that are later added for capacity expansion.

Scale-Out (Compute and Memory Expansion)

Scaling out an object store enables you to add more resources to an existing object store cluster. You can scale out the CPU count and memory, and optionally, add more storage in addition to the current usage. Scale out uses an extra 10 vCPUs and 32 GiB memory for the newly added node.

Expanding Storage for an Object Store

Before you begin

Following are the requirements to add a new cluster to an existing object store.

- Ensure to upgrade to MSP 1.0.5 and Objects 2.0 before expanding the object store cluster.
- Ensure that the new clusters are registered to the same Prism Central cluster where the object store is deployed.
- Primary and secondary clusters should have versions greater than AOS 5.11.2.
- Up to 4 secondary clusters are supported.
- VMware ESXi and Nutanix AHV clusters are supported. Microsoft Hyper-V cluster is not supported.
- Data services IP address must be configured in each secondary cluster.
- Firewall must be running on the CVMs in the cluster.
- Maximum latency required for adding secondary clusters is less than 5 milliseconds.

> Note:
>
> - You cannot add secondary clusters if the primary cluster is full. You should add the secondary clusters before the primary storage reaches 80% of the total capacity.
> - When you add a secondary cluster to an empty object store (with no objects) and if suddenly the secondary cluster becomes unreachable, then an alert will not be displayed for that Prism Element and that cluster will not be shown in the user interface. However, once you start adding objects to that object store, an alert will be raised and the unreachable cluster will be marked in the user interface in a minute.

> CAUTION: Primary clusters cannot be removed once successfully added to an object store.

A non-admin user can expand storage for an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Multicluster
- Add Multicluster

- Update Multicluster
- View Cluster
- View Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

In addition to the above permissions, the user must have View Cluster permission on the Prism Element that needs to be added to the Object store.

Update Multicluster is needed only to update the capacity or to delete an added cluster.

About this task

To expand the object store storage, do the following:

Procedure

1. Log on to the Prism Central web console.
2. Click the **Entity** menu > **Services** > **Objects**.
3. Click the name of the object store for which you want to expand the storage.

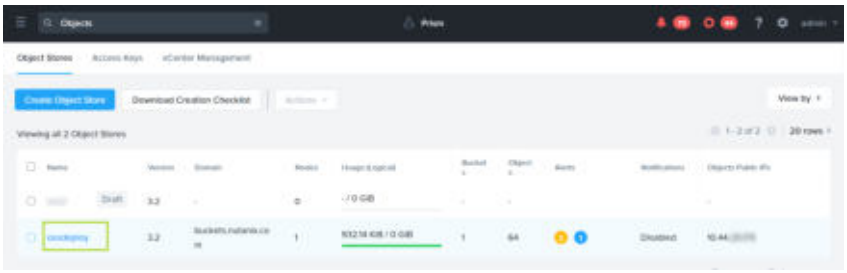

Figure 32: Object Stores List View

The object store opens on a new window.

4. Click **Clusters**.

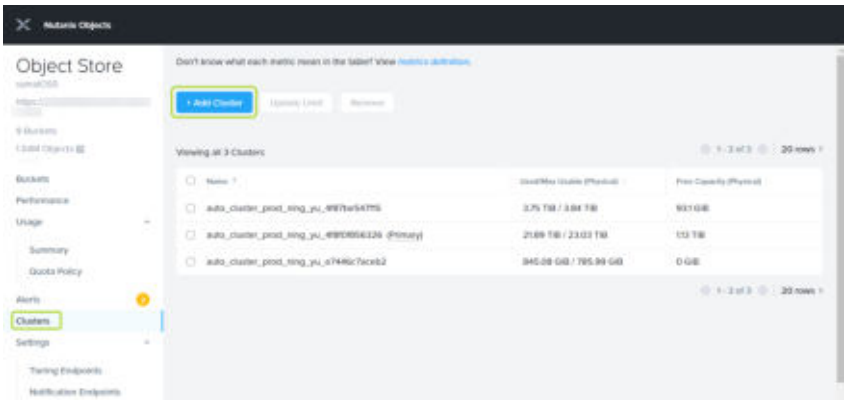

Figure 33: Object Store Clusters List

Note: You can expand your object store cluster by adding one secondary cluster at a time.

The table lists the Usage, Max Usable and Free Capacity of the cluster. All the capacities listed are physical capacities, not logical.

- Usage (Physical): Physical capacity used by this object store on the cluster.
- Max Usable (Physical): Maximum physical capacity on the cluster can be used by this object store. This is calculated as the total physical capacity of the cluster * any limit set for this cluster - capacity used by other workloads. Your available capacity might be less than you planned for as other workloads are taking up the space.
- Free Capacity (Physical): Additional capacity the object store can consume on this cluster within the specified limit. It is possible that there might not be free capacity available for the object store to consume even if current consumed capacity is less than max usable since there may be other workloads consuming capacity from this cluster.

For example, Cluster A is a primary cluster with 10 TB as the total capacity. No hard limit can be set on this cluster as it is a primary cluster. So, the Max Usable Physical Capacity of this object store cluster is 10 TB. However, the current object store usage is 5 TB and other workloads usage is 2 TB. So, the additional Physical Free Capacity available for the object store is 3 TB. However, in case of Cluster B which is the secondary cluster with total capacity as 15 TB, the hard limit is 50%. So, the Max Usable Physical Capacity of this object store cluster is 50 % * 15 TB = 7.5 TB. However, the current object store usage is 2 TB and other workloads usage is 10 TB. So, the additional Physical Free Capacity available for the object store should be 5.5 TB, but as the the other workloads consume 2.5 TB of the Max Usable Physical Capacity of the object store, so the remaining Physical Free Capacity available for the object store is only 3 TB. So, other workloads can consume the Max Usable Physical Capacity of the object store; however, an object store cannot go beyond the set limit.

Your primary cluster where the object store is deployed will be displayed in the table and this cluster cannot be removed. You can also view the free and used storage space.

5. Click **Add Clusters**.

   Note:

   - Once you have added 4 secondary clusters, the **Add Clusters** button will be disabled.
   - When adding a cluster, you cannot remove another cluster or reduce the limit of another cluster; however while removing the cluster or reducing the limit of the cluster, you can add another cluster.

   A list of clusters registered to the Prism Central will be displayed. Hypervisor type, total physical capacity and free physical capacity of the clusters will also be displayed. However, the clusters that are already added as secondary clusters to an object store will not be displayed.

6. Once you select the cluster, under **Set up hard limit** section, select the usage limit in percentage, or select **Custom** and enter a custom limit for an object store, and then click **Done**.

   This limits the object store to use a maximum capacity on the selected cluster. If you exceed the limit, an alert is generated.

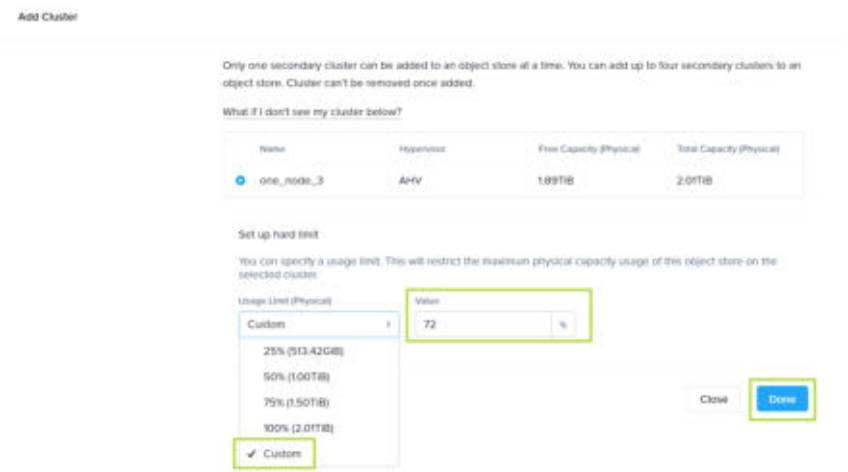

Figure 34: Setting Up Hard Limit

You can change the limit (increase or decrease) once a secondary cluster is added. Select the cluster, and then click **Update Limit**. You cannot update the limit for the primary cluster.

You can also remove the secondary clusters once added. Select the cluster, and then click **Remove**. The removed cluster can be added back to the multi-cluster.

Note: Points to note when updating the limit of the secondary clusters or removing the secondary clusters.

- If the secondary cluster is empty, then the removal of that cluster takes up to 3 hours.
- If the secondary cluster has some data, and the storage limit is reduced or the cluster is removed, then starting the data migration process may take up to 7 hours.
- When you reduce the limit of the secondary cluster and if the used capacity of the cluster is less than the updated limit capacity, then no data migration takes place and limit is changed without any delay.
- When adding a cluster, you cannot remove another cluster or reduce the limit of another cluster; however while removing the cluster or reducing the limit of the cluster, you can add another cluster or increase limit of another secondary cluster.
- If you cancel removing the cluster or decreasing the limit, then the last updated limit remains the same and any data migrated to other clusters will not be migrated back to this cluster.
- While reduction of limit for one cluster is ongoing, you can increase the limit of another cluster but cannot decrease.

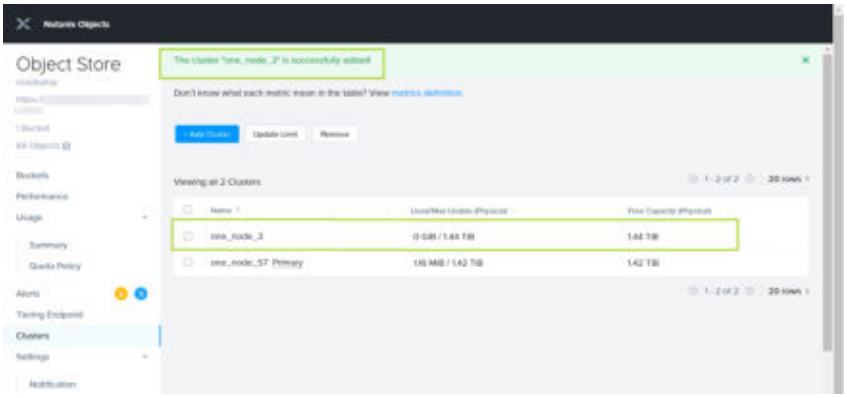

Figure 35: Setting Up Hard Limit

A new cluster is added to the object store. If any secondary cluster addition fails, you can remove that cluster. You can also see the usage of these clusters in the **Usage** tab. For more information, refer to Viewing Object Store Usage on page 180.

Note: It will take a minute to add a new cluster.

Scaling Out an Object Store

Scaling out an object store enables you to add more resources to an existing object store cluster. You can add a worker node, and optionally, add more storage in addition to the current usage. An additional 10 vCPUs and 32 GiB memory get added for each worker node.

Before you begin

- Make sure that you have at least a three-node cluster for performing scale out.
- Make sure that physical resources are available.

A non-admin user can scale out an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Scale Out
- View Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

> Note:
>
> - Objects version lower than Objects 2.1 does not support scale out. Upgrade Objects to the latest version to use the scale out feature.
> - You can scale out one node or VM at a time.
> - During scale out of an object store, no disruption happens. You can launch the object store during scale out.
> - If physical resources (VMs) are deployed, rollback is not supported. However, if deployment of physical resources fails or if deployment fails prior to deploying physical resources and your cluster is not scaling, you can roll back. For rolling back scale out of object store, contact Nutanix Support at http://portal.nutanix.com.

You can perform compute scale out (adding worker nodes) and storage scale out (adding additional storage capacity) for an object store.

To scale out an object store, do the following:

Procedure

1. Log on to the Prism Central web console.
2. Click the **Entity** menu > **Services** > **Objects**.

3. In the object store table, select the object store that you want to scale out.

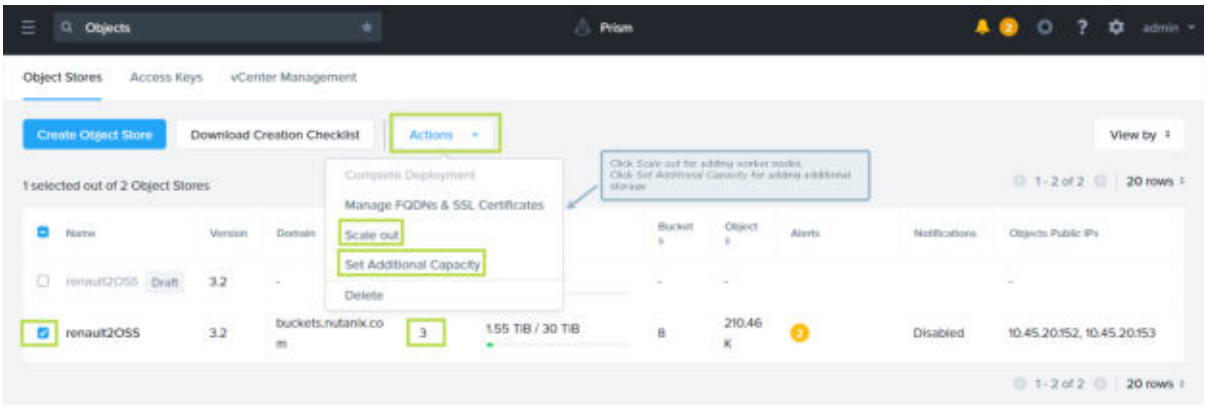

Figure 36: Object Store Selection for Scale Out

4. If you want to add a worker node (compute scale out), do the following:

   a. In the **Actions** list, click **Scale out**.
      The Compute Scale Out window appears.

   b. Click **Add Nodes**.

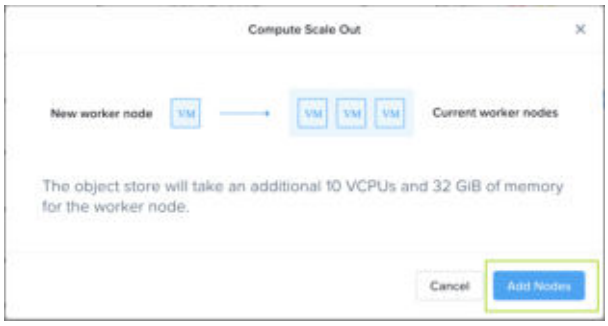

Figure 37: Compute Scale Out

The object store is now scaling out. This process takes about 5 to 10 minutes. You can track step by step deployment progress with scale out workflow progress for the object store.

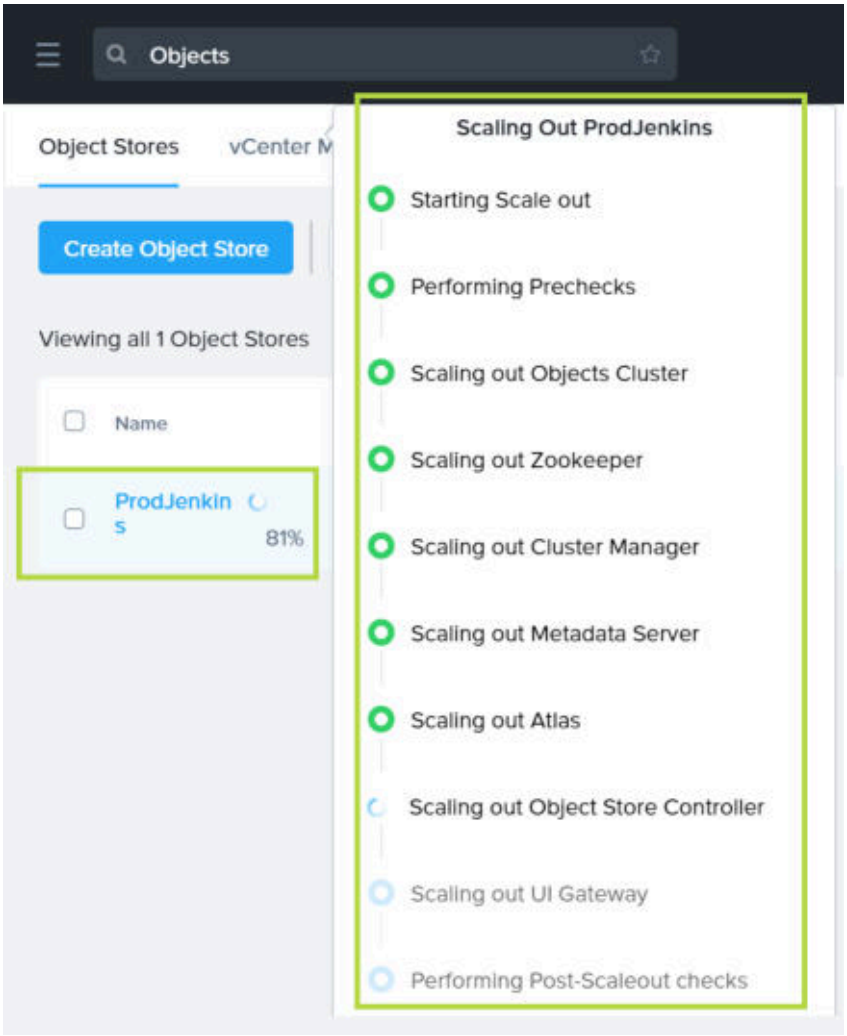

Figure 38: Object Store Scale Out Progress

Once the object store scale out is completed, new node is added to the object store. The object store takes an additional 10 vCPUs and 32 GiB of memory for the worker node.

5. If you want to add more storage to the current usage (storage scale out), do the following:

   a. In the **Actions** list, click **Set Additional Capacity**.
   The Set Additional Capacity page appears.

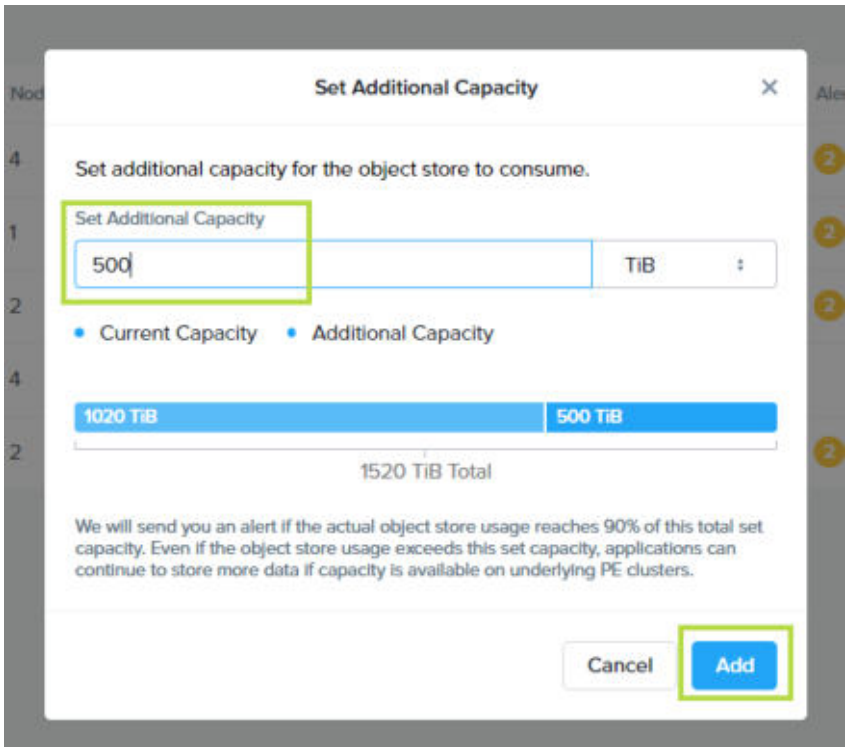

Figure 39: Storage Scale Out

   b. Enter the amount of additional storage you want to add to the current usage and click **Add** to complete.

> Note: Objects generates an alert if the logical usage of your object store reaches 90% of the specified value. For more information about alerts, refer to Viewing Alerts.

Managing FQDN and SSL Certificate

You can add multiple FQDNs, download the CA certificate, and set up or replace the SSL certificate for your object store.

A non-admin user can manage FQDN and SSL certificates only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Regenerate Self Signed Certificate
- Replace Certificates
- Download Certificate
- View Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

Adding FQDNs

You can now create multiple FQDNs for an object store. The FQDN used while creating an object store is the default FQDN and rest of the FQDNs are alternate FQDNs.

About this task

To add a FQDN, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Select the object store for which you want to create a FQDN.
3. Click **Actions** > **Manage FQDNs and SSL Certificates**.

4. In the New FQDN field, type the FQDN and then click **+FQDN**.
   Guidelines for naming FQDN:

   - Must contain at least two dots.
   - Can contain either alpha numeric, hyphen, or underscore characters.
   - Cannot start or end with a hyphen or underscore or dot.
   - Duplicate domain name is not allowed.
   - A FQDN cannot be at a sub-domain of another FQDN.

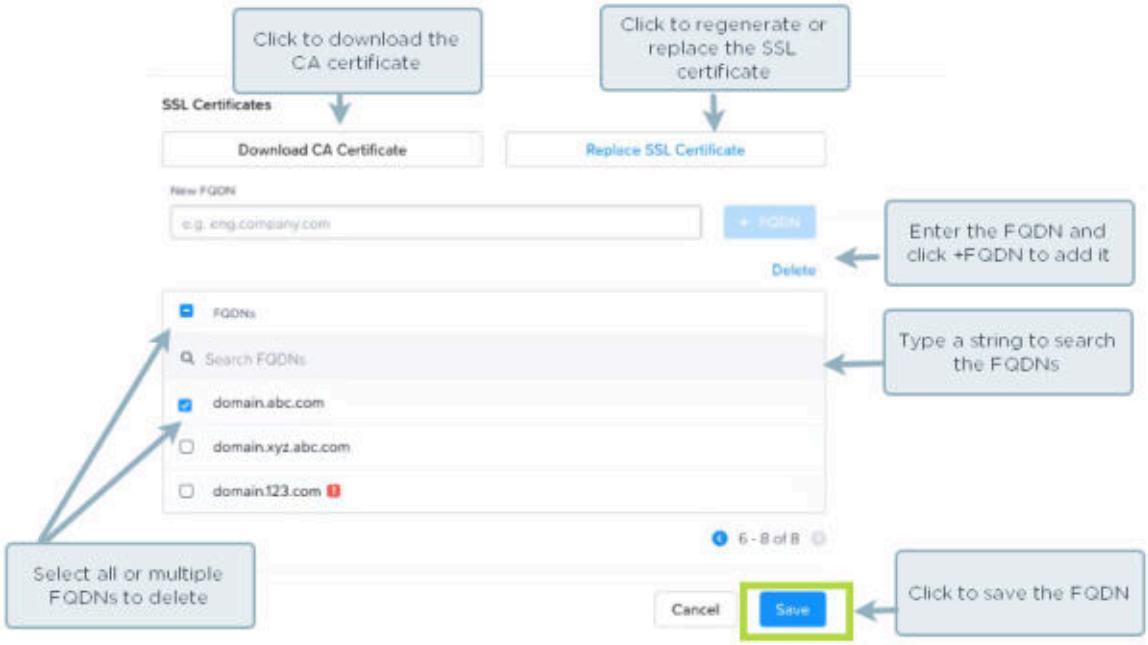

Figure 40: Adding FQDNs

You can select one or more FQDNs, and click **Delete** to delete the FQDNs. However, you cannot delete the default FQDN (the FQDN used while creating an object store).

The new FQDN is listed in the table.

> Note: A warning message is shown if the new FQDN is missing in the DNS configuration of the certificate. To add the FQDN to the certificate do either of the following:
>
> - Regenerate the SSL certificate to add all the newly added domains in the certificate.
> - Replace the SSL certificate by importing a CA signed certificate. Make sure to add all the domains in that certificate .

5. Click **Save**.
   A confirmation dialog box appears to replace the SSL certificate.

6. Click **Yes, Replace**.
   The new FQDN is added and the SSL certificate remains the same.

   > Note: The object store will be unreachable for 2-3 minutes and you will not be able to perform any operations on that object store.

What to do next

You can also regenerate or replace the SSL certificate. Refer to Setting up SSL Certificate for an Object Store on page 76.

Setting up SSL Certificate for an Object Store

By default, self-signed Secure Socket Layer (SSL) certificates are generated. If you have strong security requirements, you can replace the default certificate for the object store to securely connect to the object store while you use the HTTPS protocol. You can replace the certificates either by regenerating a self-signed certificate, or if you have a Certificate Authority (CA)-signed certificate, then you can import your private key and certificate files. When you replace the existing certificate, it removes the web browser certificate error warnings.

Before you begin

Ensure that you have met the following requirements:

- The private key should be RSA key type with key size 2048- or 4096-bit. Contents of the private key can be in PKCS#1 standard, unencrypted PKCS#8 standard, and PEM format.
- The provided public certificate must be signed by the provided CA.

  > Note: If you want the server to return the server certificate and the chain of intermediate certificates, upload the server certificate and the chain of intermediate certificates as a public certificate in a single file.

- The public certificate must have the FQDN of the Object Store Service along with the wild card in either CN or SAN. For example, if the object store name is `objects-2021` and domain is `companyname.com`, then the FQDN should be objects-2021.companyname.com. Then, the certificate must have *.objects-2021.companyname.com, objects-2021.companyname.com.

About this task

To set the SSL certificate, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Select the object store for which you want to set the certificate.
3. Click **Actions** > **Manage FQDNs and SSL Certificates**.
4. Under SSL Certificates, click **Replace SSL Certificate**.

5. In the Replace SSL Certificate window, select one of the following:

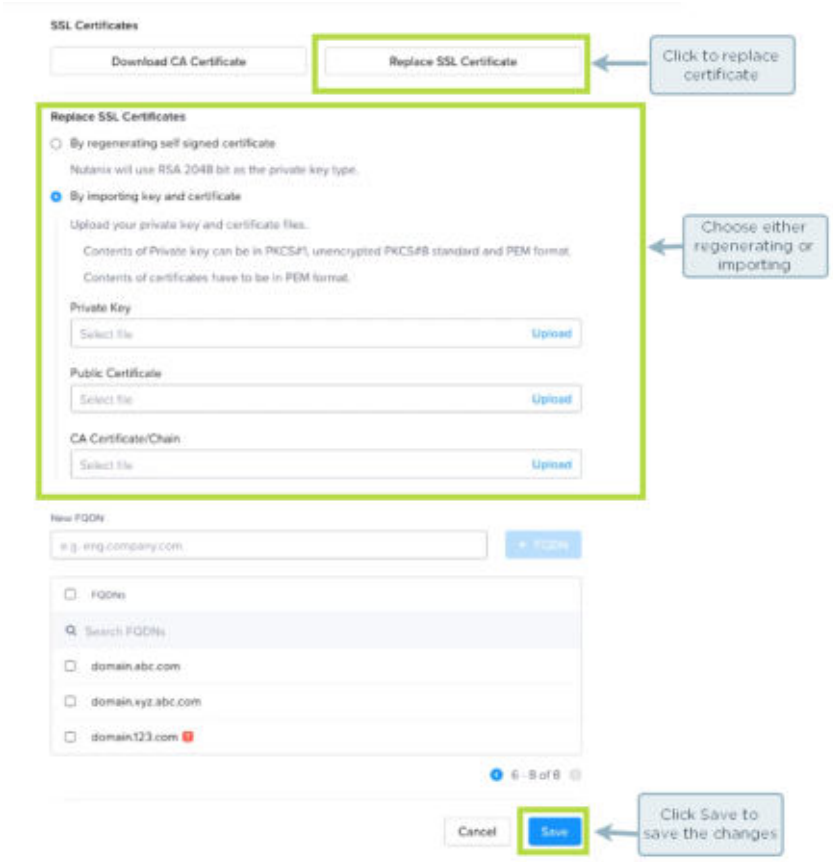

Figure 41: Replacing SSL Certificate

a. **By regenerating self signed certificate**: Uses RSA 2048 bit as the private key type.

A self signed certificate is a certificate signed by the same entity that verifies the certificate.

a. **By importing key and certificate**: Upload your private key and certificate files.

- **Private Key**: Click **Upload** to upload the private key.

  This key is used to decrypt the message.

- **Public Certificate**: Click **Upload** to upload the public certificate.

  A public key certificate is an electronic document used to prove the ownership of a public key. The certificate includes information about the key, information about the identity of its owner (called the subject), and the digital signature of an entity that has verified the contents of the certificate (called the issuer).

- **CA Certificate/Chain**: Upload the CA certificate.

  A certificate chain is the certificate of a particular CA, plus the certificates of any higher CAs up through the root CA.

You can also add or delete the FQDNs while regenerating or replacing the SSL certificate. Refer to Adding FQDNs on page 74.

6. Click **Save**.
   The default SSL certificate is replaced.

Object Store Deployment at a Dark Site (Offline Deployment)

Dark site deployment is a process for deploying Objects at a site without Internet access.

Before you begin

Make sure you satisfy the deployment prerequisites before starting the deployment. Refer to Deployment Prerequisites - AHV and ESXi on page 32.

About this task

> Note: MSP version needs to be MSP 1.0.8 before you can upgrade to MSP 2.2.1 and for Objects to be visible for upgrade in the dark-site mode.

To deploy an object store in a dark site, do the following:

Procedure

1. From a device that has public Internet access, go to Nutanix Portal and select **Entity Menu** > **Downloads** > **LCM**. Download the **LCM Dark Site Bundle** tar file.
2. Set up a web server, and upload **LCM Dark Site Bundle** to the server and extract the files in a directory in the base of a web server.
   For example, you can create a directory name `release` and extract the files in this directory.
   To view several examples of setting up a web server, see LCM Dark Site Guide.
3. From a device that has public Internet access, go to Nutanix Portal and select **Entity Menu** > **Downloads** > **Nutanix Objects**. Download the Nutanix Compatibility `nutanix_compatibility.tgz` and its signature `nutanix_compatibility.tgz.sign` files.
4. Transfer the compatibility and its signature tar files to your web server and replace the existing compatibility files with the new files.
5. From a device that has public Internet access, go to Nutanix Portal and select **Downloads** > **Nutanix Objects**. Download the required version of the `objects-x.x.tar.gz` tar file that you want to use or download the latest version of the `objects-x.x.tar.gz` tar file, and then upgrade Objects to that latest version before deployment.
   x.x represents the Nutanix Objects version.
6. Transfer the `objects-x.x.tar.gz` tar files to the web server and extract the files in a directory in the base of a web server.
   For example, you can create a directory name `release` and extract the files in this directory.
7. From a device that has public Internet access, go to Nutanix Portal and select **Downloads** > **Microservices Platform (MSP)**. Download the required version of the `msp-x.x.x.tar.gz` tar file that you want to use or download the latest version of the `msp-x.x.x.tar.gz` tar file, and then upgrade the MSP to that latest version before deployment.
   `x.x.x` represents the MSP version.
   Use command `mspctl controller version` to check the MSP version.
8. Transfer the `msp-x.x.x.tar.gz` tar files to the web server and extract the files in the same directory where `objects-x.x.tar.gz` is extracted.

9. Enable Objects to start the msp-controller on Prism Central. Refer to Enabling Objects on page 15.

   > Note: MSP Controller version should be 1.0.3 or later. If the MSP Controller version is below 1.0.3, upgrade the MSP Controller.

10. (Optional) Check the MSP Controller version and upgrade the MSP Controller to 1.0.3.

    For more information on checking the version and performing upgrade of MSP Controller, refer to Microservices Platform (MSP) on page 219.

11. To configure the dark site on MSP, SSH into the Prim Central VM as an admin user and run the following command.

    ```
    admin@pcvm$ mspctl controller airgap enable --url=http://x.x.x.x/directoryname
    ```

    x.x.x.x is the IP address of the LCM web server and *directoryname* is the name of the directory where the packages were extracted.

    For example, `admin@pcvm$ mspctl controller airgap enable --url=http://10.48.111.33/release`. Here, *10.48.111.33* is the IP address of the LCM web server and *release* is the name of the directory where the packages were extracted.

    To verify the configuration, run the following command.

    ```
    admin@pcvm$ mspctl controller airgap get
    ```

12. Deploy an object store through Prism Central. Refer to Creating or Deploying an Object Store (Prism Central) on page 52

TYPES OF OBJECTS USERS

There are two ways to manage Objects.

- Prism Central: Prism Central users can access Objects by using the Prism Central web console. These users can create user accounts and perform all the operations except any object-specific operations such as PUT, DELETE, copy, and list objects on any of the buckets (own or others). Prism Central users can also access a bucket by using the S3 APIs. These users can view buckets of all the API users, and can also share buckets of any API user with any other users.

  The following additional types of Prism Central users with varying access levels, as described below, can be defined using role-based access control (RBAC).

  - Objects only admin - Admin user having access to only Objects service within Prism Central. Within Objects, you can create the following three different users:

    - Full access to Objects service (Objects only—full access)
    - Read-only access to Objects service (Objects only—view access)
    - Restricted access to Objects service (user can customize the access)

  - Prism Central admin role with no access to Objects service at all within Prism Central (No Objects access)

- API: S3 API users (IAM users) cannot access Objects by using the Prism Central web console. They access buckets and perform operations only by using the S3 APIs. This includes S3-compatible applications. The API users have unconditional access to their own buckets, and limited or no access to buckets of other users based on the share policy. S3 API users are added using the Objects GUI. For more information, see Generating Access Key for API Users on page 85.

DIRECTORY CONFIGURATION AND ACCESS KEY GENERATION

You can configure the directory, add people, and generate access keys for the people. You can use these directories to search for people who can have access to the service. Only users with an access key can share buckets.

Configuring Directories

You can add directories that Objects can use to search for people who can have access to the service. You can also configure multiple Active Directory servers in the user interface and search across one or more Active Directories.

About this task

To configure the directories, do the following:

Before you begin

A non-admin user can configure directories only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- View Object Store
- Create Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click **Access Keys**.
3. Click **Configure Directories**.

4. Click **+ Add Directory**.

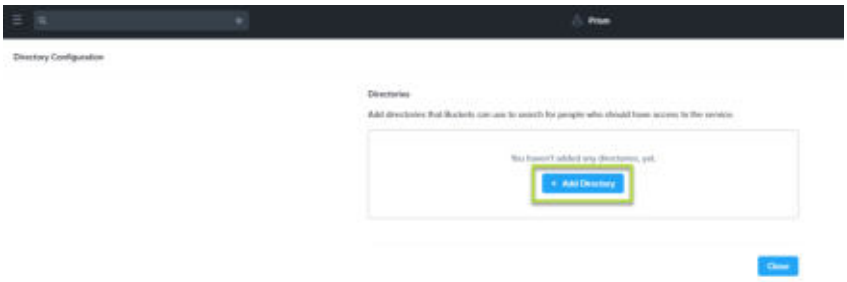

Figure 42: Directory Configuration

If you have already added the directory earlier, you will see a list of directories in this screen, and you can edit or remove the directory from the **Edit** or **Remove** link next to the directory name.

5. In the Add Directory window, select any one of the following:

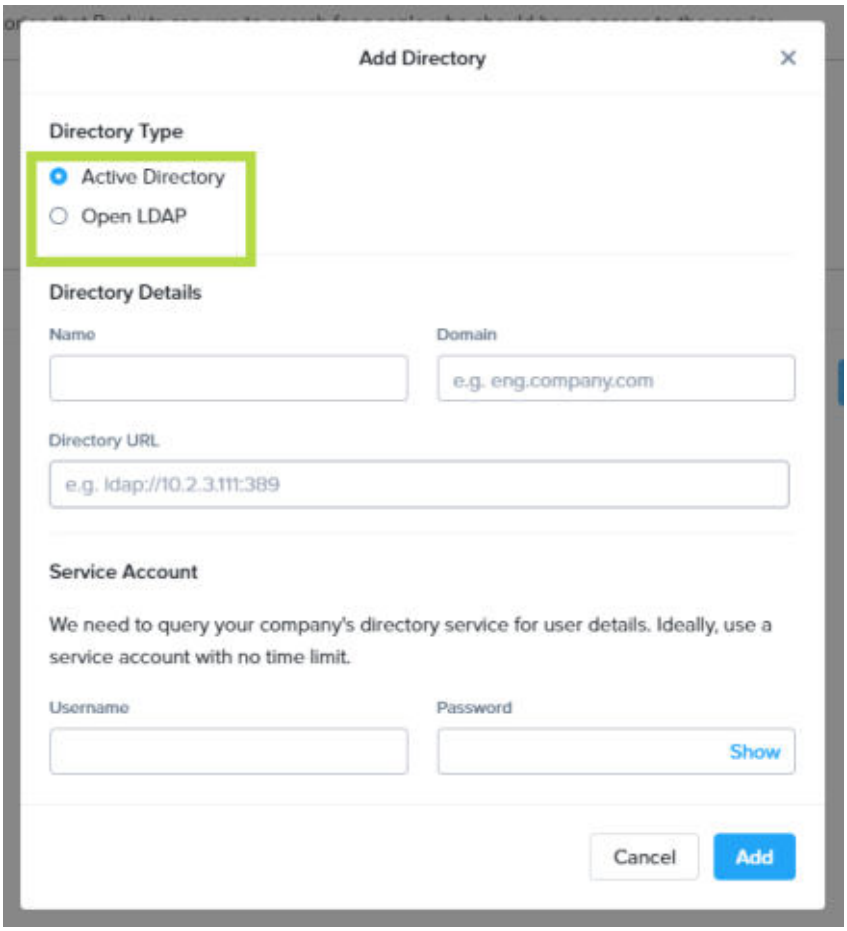

Figure 43: Add Directory Window

a. **Active Directory**: To add a directory through Active Directory, enter the following directory details and service account credentials:

- **Name**: Enter the name of the Active Directory server.
- **Domain**: Enter the domain that represents the top of the Active Directory tree and uses DNS to define its namespace.

  Usually, the domain is the DNS name of the company.
- **Directory URL**: Enter the Active Directory URL with port to access the Active Directory server.

  For example, `ldap://10.2.3.111:389`
- **Username**: Enter the username for accessing the Active Directory server to retrieve the user details.
- **Password**: Enter the password for accessing the Active Directory server to retrieve the user details.

> Note: To access Objects, no expiry should be set on the Active Directory account.

b. **LDAP**: To add a directory through Lightweight Directory Access Protocol (LDAP), enter the following directory details, user and group details hierarchy (or tree under which the user details for generating access keys have to be retrieved), and service account credentials:

- **Name**: Enter the name of the OpenLDAP server.
- **Domain**: Enter the domain that represents the top of the LDAP tree that uses DNS to define its namespace.

  Usually, the domain is the DNS name of the company.
- **Directory URL**: Enter the domain that represents the LDAP URL with port to access OpenLDAP server.

  For example, `ldap://10.2.3.111:389`
- **User Object Class**: Enter the LDAP object class value that defines users in the directory service. When the user is created, this list of user object classes is added to the attributes list of the user.
- **User Search Base**: Enter the location or the search starting point in the LDAP tree, which locates the users.

  For example, OU=people.
- **Username Attribute**: Enter the attribute names that are searched to retrieve users from the LDAP tree.
- **Group Object Class**: Enter the LDAP object class value that defines groups in the directory service to which the users belong.

  When the group is created, this list of group object classes is added to the attribute list of the user.
- **Group Search Base**: Represents the location or the search starting point in the LDAP tree under which groups are located.

  For example, OU=people
- **Group Member Attribute**: Enter the member attribute that specifies the group memberships.
- **Group Member Attribute Value**: Enter the group entries for which the memberships are specified by using the member attribute.

  These member attributes can have member attribute values specifying group membership in Distinguished Names (DNs). Member attribute values are used for group membership resolution.
- **Username**: Enter the username for accessing the openLDAP server to retrieve the user details.
- **Password**: Enter the password for accessing the openLDAP server to retrieve the user details.

  Note: To access Objects, no expiry should be set on the LDAP account.

What to do next

You can now generate access key for the API users. Refer to Generating Access Key for API Users on page 85.

Generating Access Key for API Users

You can add people (users) from the directory or through email address, and generate access key for the users. Only users with access keys can share buckets.

Before you begin

A non-admin user can generate access keys only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- View Object Store
- Create Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

To generate access keys for the users, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click **Access Keys**.
3. Click **+ Add People**.

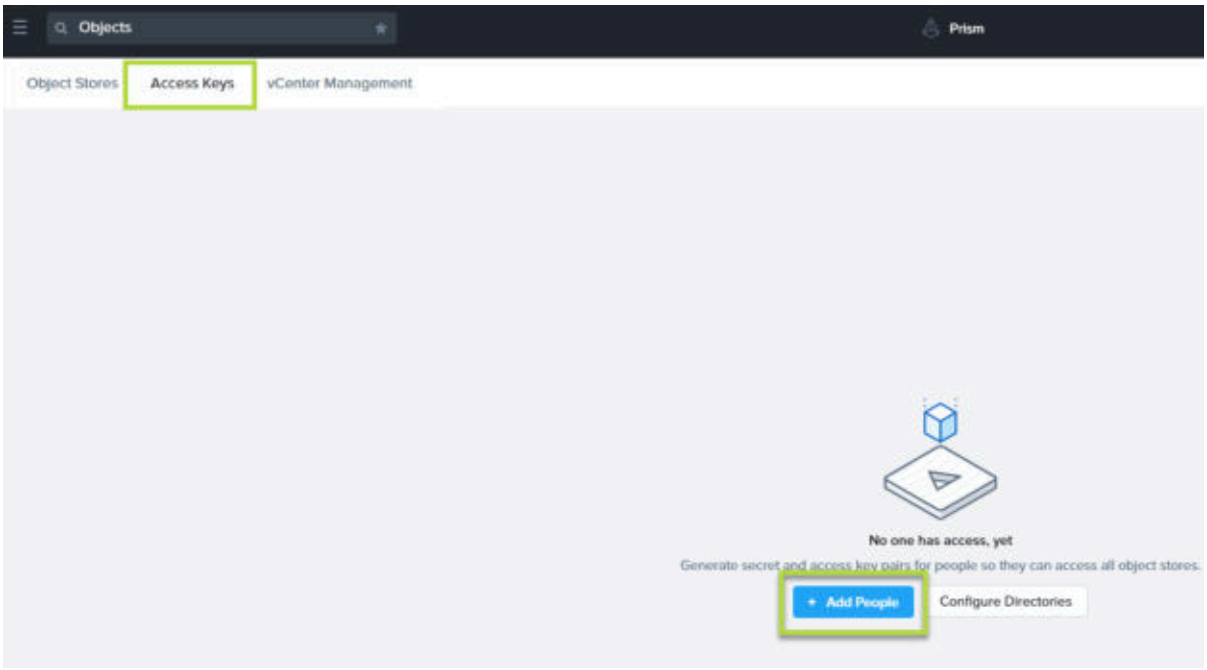

Figure 44: Adding People for Access Key Generation

4. In the Add People window, select any of the following options:

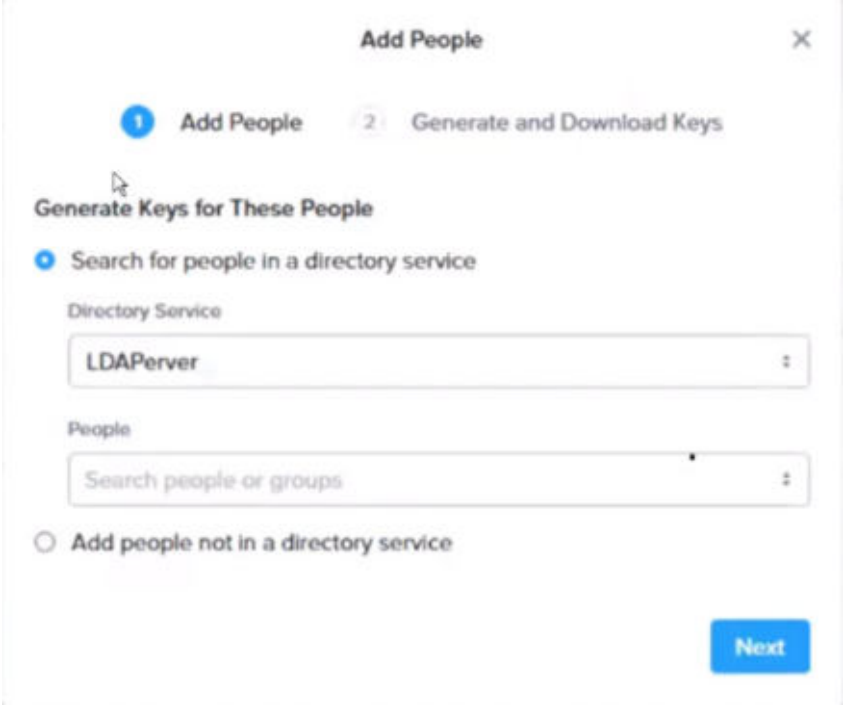

Figure 45: Add People Window: Search in Directory

a. **Search for people in a directory service**: Select to add people from the directory.
   For more information about adding a directory, refer to Configuring Directories on page 81.

   Note: You can use the Active Directory (AD) group to generate key pairs. Objects IAM generates keys pairs for each user as a separate file (inside a single zip file). The administrator can distribute these individual key pair files to the end-users.

b. **Add people not in a directory service**: Add email addresses of the people. Also, you can add a display name for the user.
   Adding the display name is optional and can contain up to 255 characters.
   Click **+Add** to add multiple users.

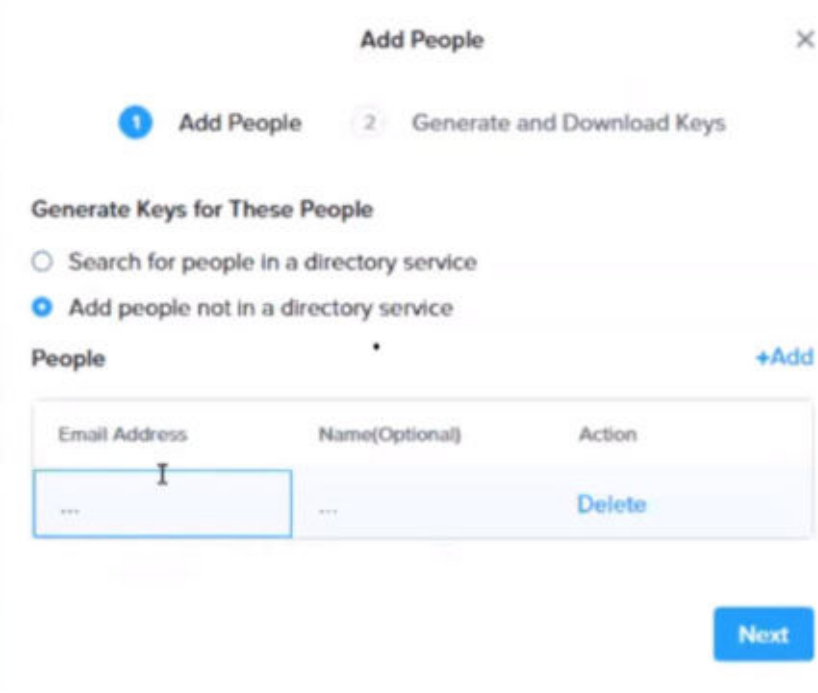

Figure 46: Add People Window: Enter Email Address

5. Click **Next** to open the Generate and Download Keys page.
   In this page, you generate and download the access keys. You can also add a tag to the access keys for key management.

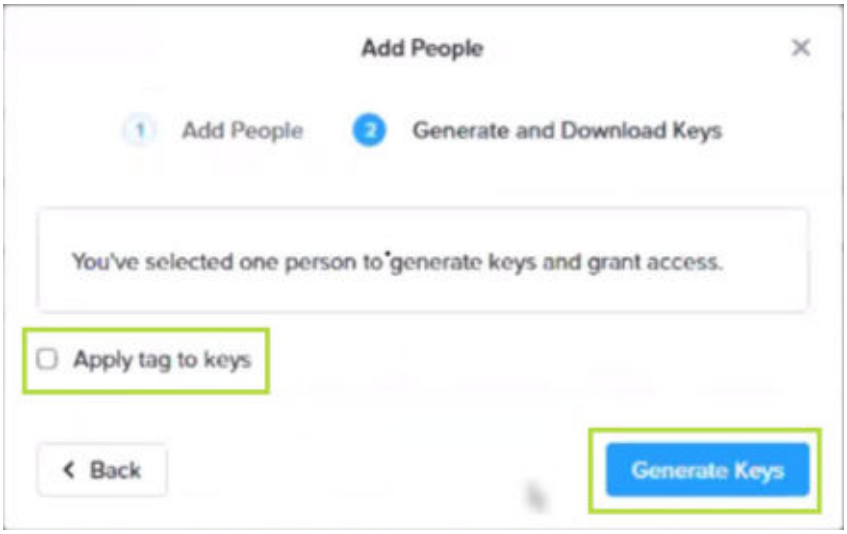

Figure 47: Generate and Download Keys Page

6. (Optional) Select **Apply tag to keys** check box and enter a tag name for the access keys.
   If you added multiple users, the same tag applies to all. The tags can contain up to 255 characters.

   Note: You cannot change the tag once added.

7. Click **Generate Keys**.
   The keys are generated for the selected people.

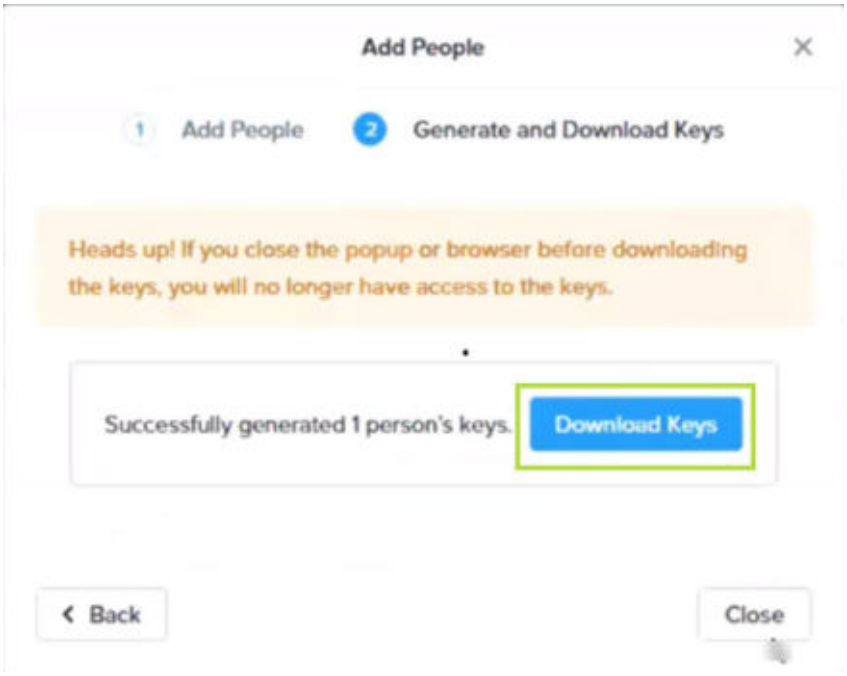

Figure 48: Generate and Download Keys Page

8. To download the keys, click **Download Keys**.
   The keys are downloaded.

> CAUTION: For Google Chrome, Microsoft Edge, and Internet Explorer, you can directly download the keys. For Safari and Firefox, after you click **Download keys**, a new tab opens listing the keys. You must copy and paste the keys at the desired location manually from the tab. You no longer have access to the keys after you close the tab.

Viewing API Users

You can view the list of people who can have access to Objects.

Before you begin

For a user with non-administrator privileges to view API users, the administrator must create an access control policy for the non-administrator user on a role that must have the minimum permission to View Object Store.

About this task

To view the list of API users of Objects, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.

2. Click **Access Keys**.

> Note: You can add keys for multiple users, but you cannot delete multiple users at the same time.

The table displays the list of API users.

Managing API Keys

If you lose or forget your access keys, you cannot retrieve the same key. However, you can add the key and get access. You can also delete a key.

Before you begin

A non-admin user can manage API keys only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- View Object Store
- Create Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

To add or delete the access keys, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click **Access Keys**.
   The table displays the list of API users.
3. Click **Manage** against a user.

Figure 49: Manage Keys Page

The Manage Keys page appears. This page provides you with the options to add or delete access keys.

4. To manage a key, you can do one of the following:

   » **Add Key**: Click this button to generate the access key for the user.

   » **Apply tag to keys**: Select this option if you want to associate a tag with the access key.

   > Note: You can add one key at a time and up to five keys for a user.

   » **Delete**: Select the access key and click this button to delete the access key.

   The access key of the user gets deleted and has no access to the object.

   > Note: You can delete one key at a time.

Deleting API Users

This section describes the procedure to delete an API user.

Before you begin

A non-admin user can delete API users only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- View Object Store
- Create Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

> Note: You cannot delete multiple API users at the same time.

To delete an API user, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click **Access Keys**.
   The table displays the list of API users.

3. Select the user you want to delete, and then click **Delete User**.
   The **Delete User** button appears at the top after you select a user.

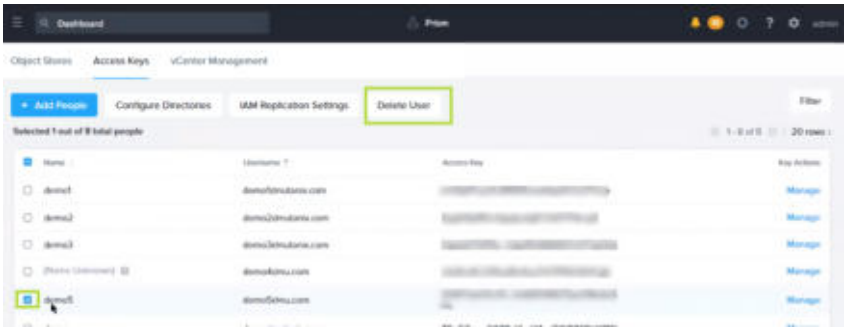

Figure 50: Access Keys Page - Deleting an API User

The user is deleted.

BUCKET CREATION, OPERATIONS AND BUCKET POLICY CONFIGURATION

Before you can upload objects to a bucket, you must create a bucket in an object store. You can create bucket using both S3 and NFS protocols. You can also configure a S3 bucket while creating it except applying the WORM policy.

A user with Edit Buckets permissions can perform all bucket operations, such as create, update, and delete on all the buckets on the assigned Object Store.

You can enable Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications to a S3 bucket. These S3 features cannot be enabled in a multi-protocol access bucket (NFS). However, you can also filter, update, share, and delete a bucket, and view all the existing buckets list created through both protocols.

Bucket Naming Conventions

The name of a bucket must conform to the following rules:

- Be a unique DNS-compliant name within a deployed bucket instance.
- Can contain alphanumeric, dot or hyphen characters.
- Begin with a lowercase letter or a number.
- Cannot contain uppercase and special characters.
- Minimum of 3 and a maximum of 64 characters long.

> Note: You cannot change the bucket name after creating bucket.

Creating and Configuring an S3 Bucket

You can create and configure the bucket settings. There are multiple options available for configuring a bucket and all are accessible at the time of creating a bucket. However, you cannot configure a WORM bucket while creating a bucket. You can edit WORM policies only after creating a bucket. You cannot enable multi-protocol access on the S3 bucket.

Before you begin

A non-admin user can create and configure S3 buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

> Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can perform a create bucket operation on the buckets available in the object stores that are assigned to the respective users.

About this task

Note: Ensure that the bucket names are unique for all the users.

To create and configure a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of object store in which you want to create a bucket.
   The object store opens in a new window.
3. Click **Create Bucket**.

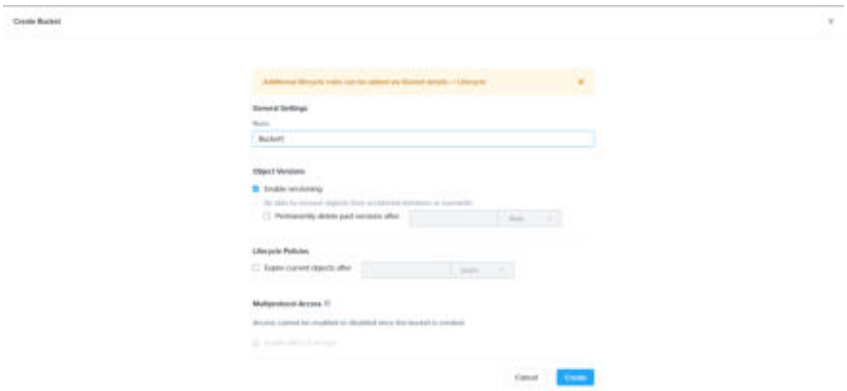

Figure 51: Create Bucket Window

Note: You can also create a bucket without enabling versioning and life cycle policies.

4. On the General Settings section, type a name for your bucket.
   For more information about naming buckets, refer to Buckets Naming Conventions.

5. (Optional) On the Object Versions section, configure the following:
   For more information about versions, refer to Object Versioning.

   a. **Enable Versioning**: Select this check box to enable versioning on objects and to keep all the versions on the same bucket.
      To apply life cycle policy with versioning enabled, refer to Rules for Lifecycle Policy When Object Versioning is Enabled on page 103.

   b. **Permanently delete past versions after**: Select this check box to enter a time period to delete all the older versions of the objects.
      You can specify the number in days, months, or years.

      Note: When you select versioning, you are able to recover objects from accidental deletion or overwrite.

6. (Optional) On the Lifecycle Policies section, select the following property:
   For more information, refer to Lifecycle Policies.

   » **Expire current objects after:** Select to type a time period after which the current version of the object expires.

   You can specify the number in days, months, or years.

   Note:

   - If versioning is not enabled, the current object deletes permanently. When you enable versioning, the current object becomes a past object.
   - Multi-protocol access cannot be enabled on the S3 bucket. If you want to create buckets with multi-protocol access, refer to Creating and Configuring an NFS Bucket on page 94.

7. Click **Save**.
   The bucket is created successfully.

What to do next

After creating a bucket, you can perform object operations from the Objects Browser or the S3 APIs. For more information, refer to Object CRUD Operations on page 148 and Supported S3 APIs section.

Creating and Configuring an NFS Bucket

You can create and configure the bucket with multi-protocol access. If you are enabling multi-protocol access on a bucket, S3 features such as Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications cannot be enabled on that bucket. You can perform update (only multi-protocol access configurations), delete, or share operations on these buckets.

Before you begin

Make sure you are aware of the uses cases, recommendations, NFS-S3 interoperability, and limitations of NFS on Objects. Refer to Use Cases and Recommendations for NFS on Objects on page 8, NFS-S3 Interoperability on page 9 and Limitations of NFS on page 50.

A non-admin user can create and configure NFS buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must

have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can perform a create bucket operation on the buckets available in the object stores that are assigned to the respective users.

About this task

Note:

- If you deployed object store prior to Objects 3.3 version, then you will not get the option to enable NFS for the bucket of that object store. Similarly, if you upgrade to later versions of Objects, then the option to enable NFS will not be available.
- Ensure that the bucket names are unique for all users.
- As the Objects-NFS does not support NLM (Network Lock Manager), no lock option is required while mounting the NFS bucket.
- The total and available bytes returned in the FSSTAT response denotes the logical capacity and logical available space and not the physical capacity of the cluster which also takes RF2 into consideration.

To create and configure a bucket with multi-protocol access, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of object store in which you want to create a bucket.
   The object store opens in a new window.

3. Click **Create Bucket**.

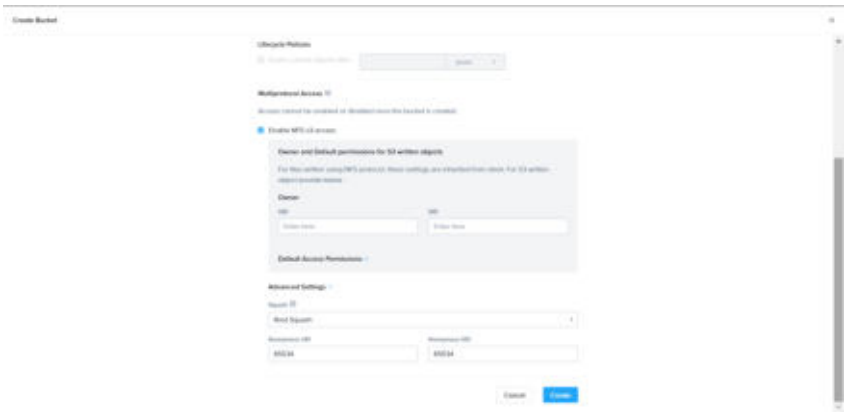

Figure 52: Create Bucket Window

4. On the General Settings section, type a name for your bucket.
   For more information about naming buckets, refer to Buckets Naming Conventions.

   Note: Versioning and lifecycle policies cannot be enabled on buckets with multi-protocol access.

   To create and configure buckets for S3 features, refer to Creating and Configuring an S3 Bucket on page 92

5. On the Multiprotocol Access section, select **Enable NFS v3 Access**.

   Note: Access cannot be enabled or disabled after the bucket is created.

6. For owner and default permissions for S3 written objects, do the following.

   Note: For files written using NFS protocol, these settings are inherited from the client.

   a. In the Owner section, enter the **UID** and **GID**.
      Any object (file or directory) created from S3 protocol have UserID (UID) and GroupID (GID) in the NFS namespace.

   b. In the Default Access Permissions section, by default read, write, and execute permissions are set for the files and directories for owner, group (other users in a group), and others (users that are not part of any group).

      Following are the default permissions for files:

      - **Owner**: Read, Write
      - **Group**: Read
      - **Others**: Read

      Following are the default permissions for directories:

      - **Owner**: Read, Write, Execute
      - **Group**: Read, Execute
      - **Others**: Read, Execute

      You can change these permissions as needed.

7. In the Advanced Settings section, select any one of the following squash options.

   » **None**: Select this option if you do not want to convert the UID and GID of the users on the server.

   » **Root Squash**: Select this option if the user has root privileges and you want to convert the UID and GID to an anonymous UID and GID on the server. The anonymous UID and GID are automatically generated, however, you can change it.

   » **All Squash**: Select this option if you want to map all users to a single identity. This will convert the UID and GID of all users to an anonymous UID and GID on the server.

8. Click **Create**.
   The bucket is created successfully.

9. Note: Before mounting a bucket, add the client to the NFS allowlist. Only the client present in the NFS allowlist will be given access to the NFS buckets.

   For more information on adding and managing clients to NFS allowlist, refer to Managing NFS Allowlist on page 98.

   Mount the bucket from the client VM.

   ```
   $ sudo mount -tnfs -o nfsvers=3,proto=tcp, soft -v objectstore-endpoint-ip-address:/bucketname path/to/mount
   ```

   For example,

   ```
   $ sudo mount -tnfs -o nfsvers=3,proto=tcp, soft -v 10.45.53.75:/test-nfs-bucket home/folder-1/mnt-point
   ```

   The bucket is mounted successfully.

What to do next

After creating a bucket, you create directories, files, and symbolic links from the NFS namespace. You can also perform object operations from the Objects Browser or the S3 APIs. For more information, refer to Object CRUD Operations on page 148 and Supported S3 APIs sections.

Managing NFS Allowlist

You need to add the IP addresses of the client VMs that are allowed to access an NFS bucket in an object store.

About this task

> Note: Make sure that you add the required client IP addresses to the allowed list before you mount a bucket.

To add clients to the NFS, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**. The **Object Stores** window appears.
2. Click the name of the object store in which the bucket exists.

3. In the left-side pane, click **NFS Clients Allowlist**, and then click **Add Client**.
   The Add NFS Clients window appears.

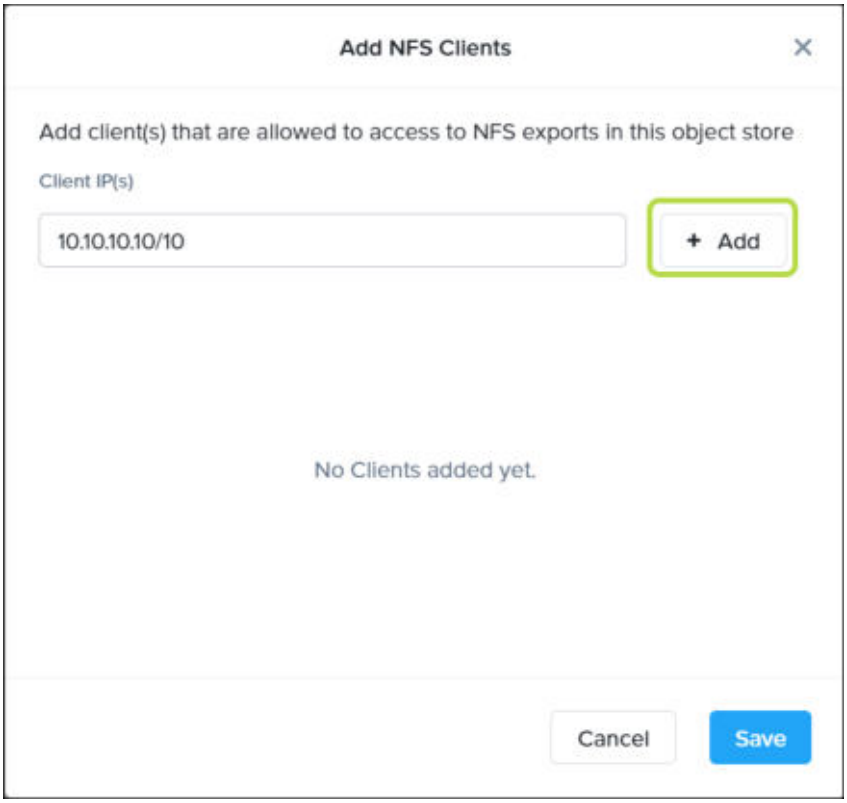

Figure 53: Add NFS Clients

   a. Enter the IP address of the client VM in Classless Inter-Domain Routing (CIDR) format and click **Add**.

   b. Click **Save**.
      The newly added IP address is listed in the **Added Clients** list.

      Note: The **Add** button appears only when you add the first IP address to the allowlist. To add more clients or to manage the existing clients, click **Manage Clients**.

4. Click **Manage Clients** to add or remove the clients.
   The Manage NFS Clients window appears.

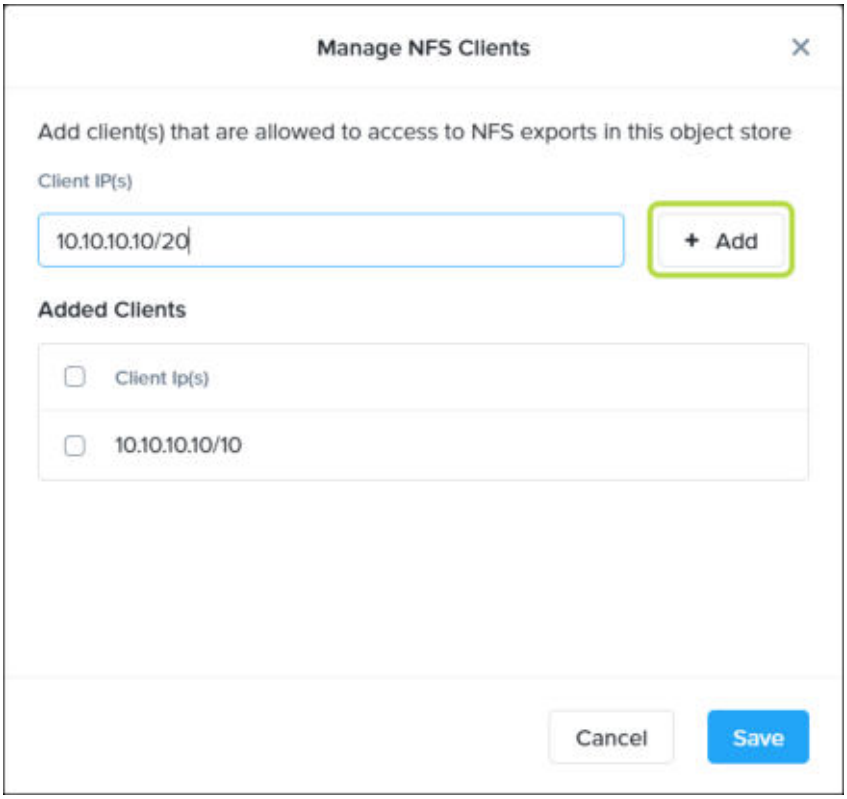

Figure 54: Manage NFS Clients

   a. (Optional) Enter the IP address of the client VM in Classless Inter-Domain Routing (CIDR) format and click **Add** to add client.

   b. (Optional) Select one or more IP addresses from **Client IP(s)** and click **Remove** to remove the clients.

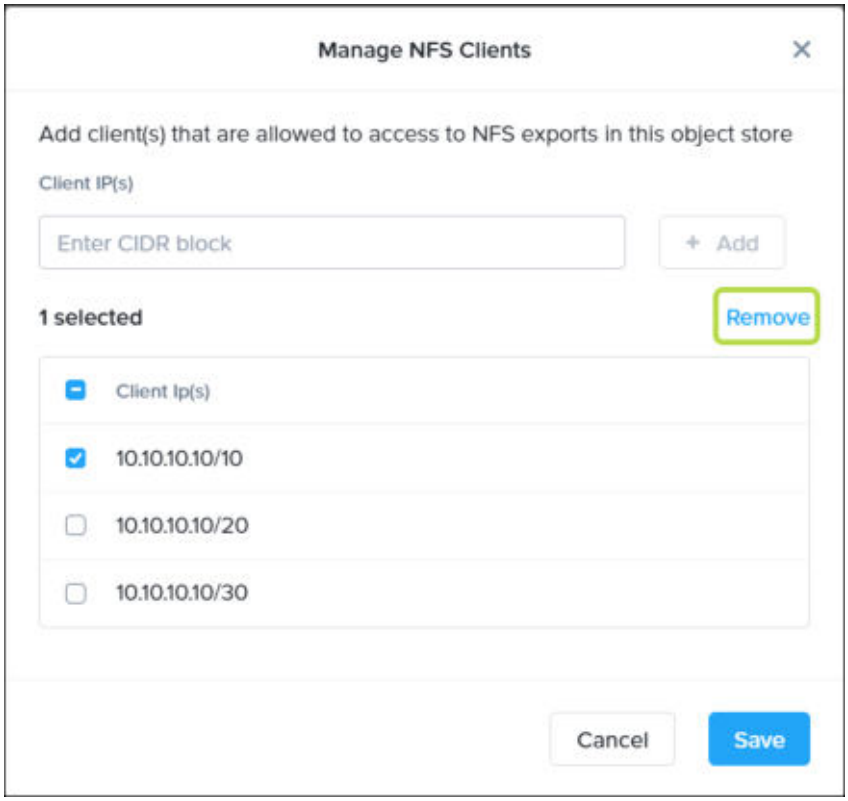

Figure 55: Manage NFS Clients

c. Click **Save**.
The IP address is removed from the **Added Clients** list.

What to do next

After adding the required clients to the NFS allowlist, you can mount the buckets and then perform tasks, such as creating files, directories, or symbolic links. For mounting the bucket, refer to Creating and Configuring an NFS Bucket on page 94.

Bucket Policy Configuration

You can configure multiple policies in a bucket such as object versioning, lifecycle policies, WORM, static website hosting and Cross-Origin Resource Sharing (CORS). Object versioning and lifecycle policies are accessible at the time of creating a bucket, however, you can apply WORM policy only after creating a bucket.

Note:

Note: Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications cannot be enabled for buckets created using NFS protocol.

Object Versioning

Object versioning enables you to keep multiple versions of an object in one bucket. By default, versioning is disabled for a new bucket. You can enable versioning while creating a bucket or editing a bucket. Refer to Creating and Configuring an S3 Bucket on page 92.

With object versioning option, you can enable versioning on objects of that particular bucket. Object versioning varies depending upon the lifecycle policies applied to the object.

For more information about the lifecycle rules for object versioning, refer to Rules for lifecycle policy.

Note:

- Versioning cannot be enabled for the buckets created using NFS protocol.
- For a versioned bucket, the number of objects shown for each bucket is indicative of the number of versions of all the objects present in the bucket.
- You cannot disable but can suspend the object versioning at any time.

When you suspend versioning, accumulation of the new object versions is stopped and previous object versions are retained.

A non-admin user can enable versions on objects only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can perform object versioning on the buckets available in the object stores that are assigned to the respective users.

Lifecycle Policies

Lifecycle policy enables you to create or update a set of rules that define actions that Nutanix Objects applies to a group of objects. With these policies, you can expire objects when no longer required or move them to a low-cost storage tier to preserve for a longer-term.

A non-admin user can create lifecycle policies only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can apply lifecycle policies on the buckets available in the object stores that are assigned to the respective users.
- Lifecycle policies cannot be applied for the buckets created using NFS protocol.
- Rules or any updates to the rules get applied to the new objects that you create and will not apply to the objects existing before the rule creation or update.

You can apply these policies while creating or editing a bucket. You can create multiple rules within a lifecycle policy. This means that different objects within the bucket can have different rules based on prefixes and tags.

Example: You can create rule 1 to expire the current versions of the objects with the tag value as a **dev**. Similarly, you can create multiple rules with different tiering and expiring configurations and apply them to other objects using the prefixes and tags.

With lifecycle policies, you can configure a lifecycle policy rule to:

- Automatically delete objects after a specified number of days or months, or years from the date of object creation.
- Tier objects to a S3-compatible object storage bucket after a specified number of days or months, or years from the date of tiering rule creation.
- Expire the current version and previous versions of an object independently. This means that you can set different expiration durations for the current version and the previous versions.
- Expire the incomplete multi-part uploads of an object.
- Apply to all objects or a subset of objects based on prefixes, tags, or both.

For example, if you want to store log files or business transaction records for a fixed period and after that period, you want to delete them.

Note: You cannot recover the objects once it is deleted.

Rules for Lifecycle Policy When Object Versioning is Enabled

Following are the rules for the lifecycle policy when you enable object versioning:

- If you apply lifecycle policy **Expire current objects after # days/months/years** on an object with versioning enabled, it deletes the current version of the objects after the specified time, and does not delete any past versions of the objects.
- If you apply lifecycle policy **Permanently delete past version after # days/months/years** on an object with versioning enabled, it deletes all the past versions after the specified time. This specified time gets calculated from the day the object version becomes non-current or past. This operation does not delete the current version.

  Note: This policy cannot be configured on a bucket with suspended versioning.

- If you apply both the lifecycle policies **Expire current objects after # days/months/years** and **Permanently delete past version after # days/months/years** on an object with versioning enabled, it deletes all the past versions based on the time specified in **Permanently delete past version after # days/months/years**, and the current version expires based on the time specified in **Expire current objects after # days/months/years**.
- If you apply lifecycle policy **Expire or Abort incomplete multipart uploads after # days/months/years** on an object, it deletes the parts associated with the multipart uploads after the specified time.
- If you apply any lifecycle policy to a WORM bucket with versioning enabled, the lifecycle policy is not applicable until the WORM retention period has elapsed.

Cloud Tiering

Cloud tiering enables you to move objects to another S3-compatible object store bucket for saving storage space in the Nutanix Objects cluster. Tiering can help you to save costs by sending the infrequently accessed objects to platforms such as AWS S3, Microsoft Azure Blob Storage, and Google Cloud Platform (GCP). The supported endpoints are AWS S3, Azure Blob Storage, Google Cloud Platform (GCP) and a different Objects instance.

A non-admin user can configure a tiering endpoint only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Create Endpoint
- Update Endpoint
- View Endpoint
- View Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

Cloud tiering is managed through lifecycle policies. You can configure multiple lifecycle rules for different objects within a bucket.

Cloud tiering configuration consists of the following steps:

- Step 1: Configure a remote endpoint in the Object Store

  You can configure remote endpoints using the Tiering Endpoints page in the Object store instance. The objects within the bucket get moved to the endpoint according to the rules you configure.
- Step 2: Create lifecycle rules for a bucket

  You can create a tiering rule for a bucket where you define the scope (all objects or a subset of objects), select the remote endpoint, and specify the details for the rule.

Points to Note

Before you start with object tiering, note the following points:

- It is strongly recommended that
  - Only encrypted data is stored on buckets for which tiering to the cloud is enabled.
  - Object Store admins enable audit trails for the S3 bucket or other storage endpoints to ensure data is not being accessed or tampered by external malicious entities and that all access is only coming from the Object Store instance.
  - Admins follow recommended security best practices of AWS or other storage end points while setting up buckets for tiering.
- Tiering lifecycle policies are non-retroactive. The policy gets applied to the new objects that you create. The policy will not apply to existing objects.
- Tiering lifecycle policies get applied to both versioned and non-versioned objects in the same way. A separate non-current version transition action in lifecycle policies is not supported.
- Do not perform write operations on the configured-endpoint bucket. Ensure that the endpoint bucket gets used only for the Object Store instance.
- Removing a configured endpoint in the Object Store instance is not supported.
- There is a N:N relationship between the source bucket and the configured endpoint bucket. You can create multiple tiering-lifecycle rules for different objects within a bucket and use a separate endpoint for each tiering rule. Also, an endpoint bucket can be a destination to many Object Store buckets.
- Objects within a WORM-enabled bucket will continue to adhere to their WORM property even after getting tiered out to the endpoint bucket.
- Only the object data gets tiered to the endpoint bucket. The metadata of the object is not tiered.
- The behavior to access tiered objects using the `Object Get` method remains the same. If you send a request to retrieve an object from the Object store and the object is already tiered out, the Object store fetches the data from the endpoint bucket and fulfills your request.

Configuring a Tiering Endpoint

The infrequently used objects are moved to the configured endpoints according to the configured lifecycle rule for tiering. The supported endpoints are AWS S3, Microsoft Azure Blob Storage, Google Cloud Platform (GCP), and a different Objects instance.

About this task

Note:

- Cloud tiering supports GCP endpoint as Other S3 compatible endpoint type.

- For tiering the objects to GCP, you need to create a bucket in GCP with the following default configurations. This bucket will be used as the tiering endpoint.
  - **Location type**: Multi-region
  - **Default storage class**: Standard
  - **Public access prevention**: Off
  - **Access control**: Uniform
  - **Protection tools**: None
- Ensure that the tiering endpoint has a Certificate Authority (CA) signed certificate. Note that self-signed certificates are not supported.

To configure a tiering endpoint, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store.
   The object store opens on a new window.
3. In the left-side pane, click **Tiering Endpoint**.
4. To open the Add Endpoint page, click **Create**.

   If no endpoints are configured, the **Add** button appears in the center of the page. Click **Add** to start an endpoint configuration.

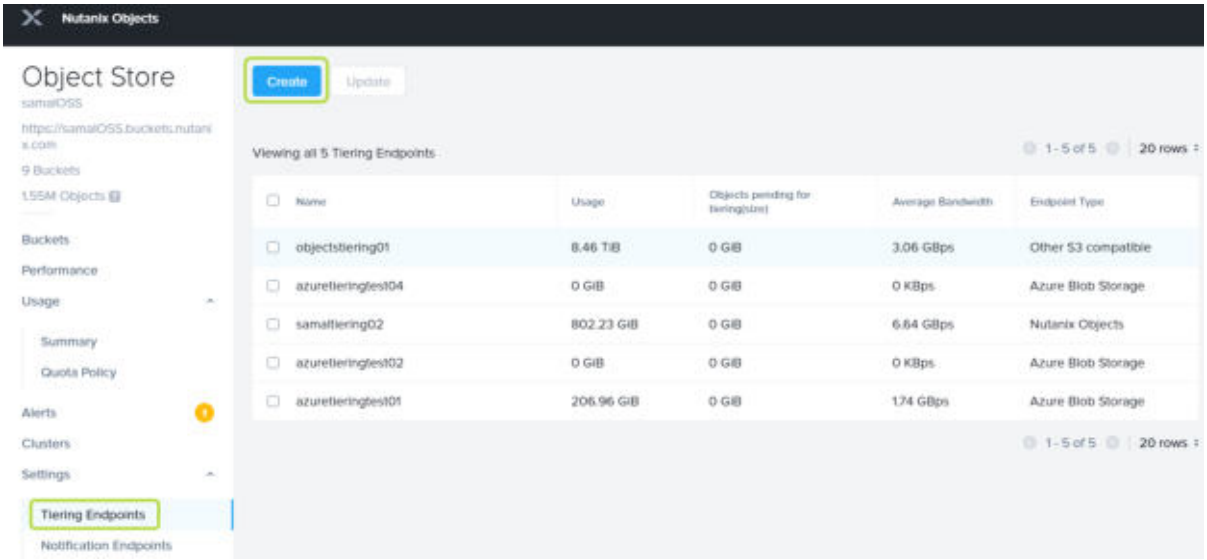

Figure 56: Tiering Endpoint Page

5. (Only for Nutanix Objects as the endpoint) Enter the following details:

   a. **Name**: Enter a name that identifies your tiering endpoint.

   b. **Endpoint Type**: Select Nutanix Objects as the endpoint type from the pull-down list.

   c. **Service Host**: Enter the complete URL of the endpoint that you want to use for tiering the objects.

      For example, the Objects public IP address or the domain name of an Object store instance.

      Example of a Nutanix Objects endpoint: `example.buckets.company.com`.

   d. **Bucket Name**: Enter the name of the bucket within the service host to which you want the objects to tier out.

   e. **Access Key**: Enter the access key of the bucket owner.

   f. **Secret Key**: Enter the secret key of the bucket owner.

   g. **Skip SSL Certificate Validation**: Select this check box to skip SSL certificate validation.

6. (Only for Other S3 Compatible as the endpoint) Enter the following details:

   a. **Name**: Enter a name that identifies your tiering endpoint.

   b. **Endpoint Type**: Select Other S3 Compatible as the endpoint type from the pull-down list.

   c. **Service Host**: Enter the complete URL of the endpoint that you want to use for tiering the objects.

      The AWS S3 URL format is `s3.<region>.amazonaws.com`

      Example of an AWS S3 endpoint: `s3.us-east-1.amazonaws.com`.

   d. **Bucket Name**: Enter the name of the bucket within the service host to which you want the objects to tier out.

   e. **Access Key**: Enter the access key of the bucket owner.

   f. **Secret Key**: Enter the secret key of the bucket owner.

7. (Only for Azure Blob Storage as the endpoint) Enter the following details:

   a. **Name**: Enter a name that identifies your tiering endpoint.

   b. **Endpoint Type**: Select Azure Blob Storage as the endpoint type from the pull-down list.

   c. **Container Name**: Enter the name of the container within the service host to which you want the objects to tier out.

   d. **Account Name**: Enter the Azure account name where the containers are located.

   e. **Secret Key**: Enter the secret key of the container owner.

8. Click **Save** to complete configuring the endpoint.

   You can also update the access and secret keys of the bucket for the configured endpoints. For example, if the bucket owner generates new access and secret keys, the owner can update the configured endpoint with the new keys using the **Update** option.

   Select the endpoint and click **Update** to open the Update Endpoint page. Update the required fields and click **Save** to complete the workflow.

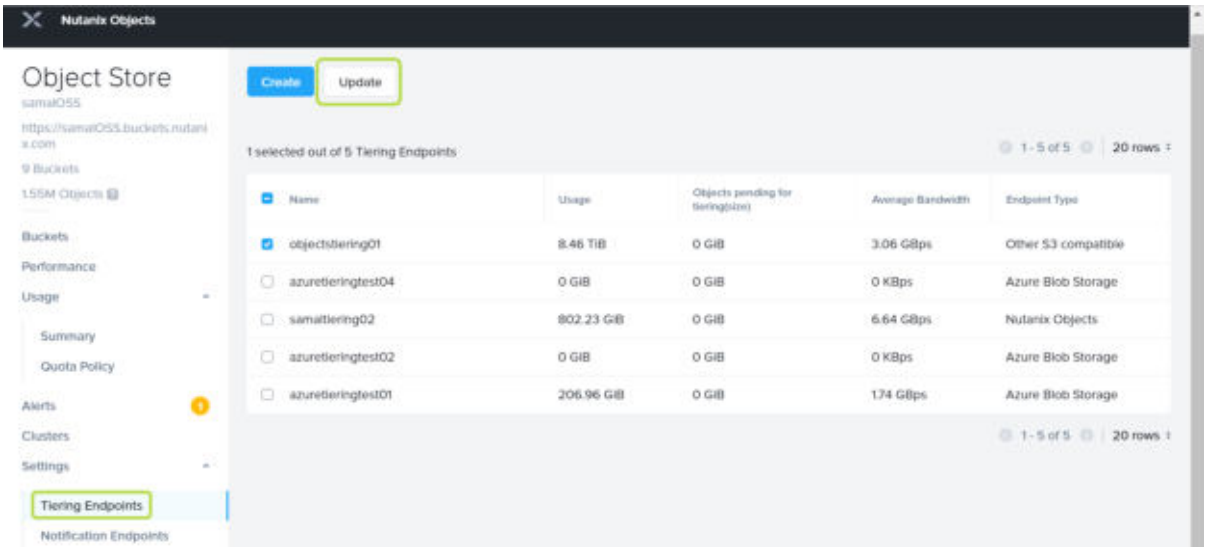

Figure 57: Updating a Tiering Endpoint

What to do next

After you configure an endpoint, create lifecycle rules for tiering objects within your bucket.

Creating a Lifecycle Rule

You can create a lifecycle rule for a bucket where you define the scope (all objects or a subset of objects), select the remote endpoint, and specify the details for the rule.

Before you begin

A non-admin user can create lifecycle rules for a bucket in cloud tiering only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Endpoint
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can create lifecycle rules bucket operation on the buckets available in the Object stores that are assigned to the respective users.

About this task

When creating a bucket, you can enable versioning, set a rule to delete past versions, and add a lifecycle policy to expire the current objects after the specified time. Once the bucket gets created, you can navigate to the bucket page and configure object-tiering lifecycle rules.

To create a lifecycle rule to tier out objects, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the Object Store that contains the bucket where you want to create tiering rules.
   The object store opens in a new window.
3. In the **Buckets** table, click on the bucket that you want to tier out.
   The Bucket page appears.
4. In the left-side pane, click **Lifecycle**.
   You can view the rules that you defined while creating the bucket.
5. Click **Create Rule** to open the Create Rule page.
6. In the Scope page, do the following, and then click **Next**:
   a. In the **Name** box, enter a name that identifies the rule you are creating.
   b. In the **Scope** list, select **All objects** or **Tags/Prefix**.
      - Select **All objects** to apply the tiering rule to all the objects within the bucket.
      - Select **Tags/Prefix** to apply the tiering rule to specific objects. You can filter objects by entering a prefix, tags, or both. Objects with the prefix and tags you specified get filtered and the tiering rule applies to those objects.

7. In the Configure Rule page, do the following, and then click **Next**.

   a. Select the **Tiering** check box.

      1. In the **Endpoint** list, select the endpoint where you want to tier out the objects.
      2. Enter a time period after which you want to move the objects to the selected endpoint. You can specify the number in days, months, or years.

      Note: The time for tiering objects must be less than the expiration time.

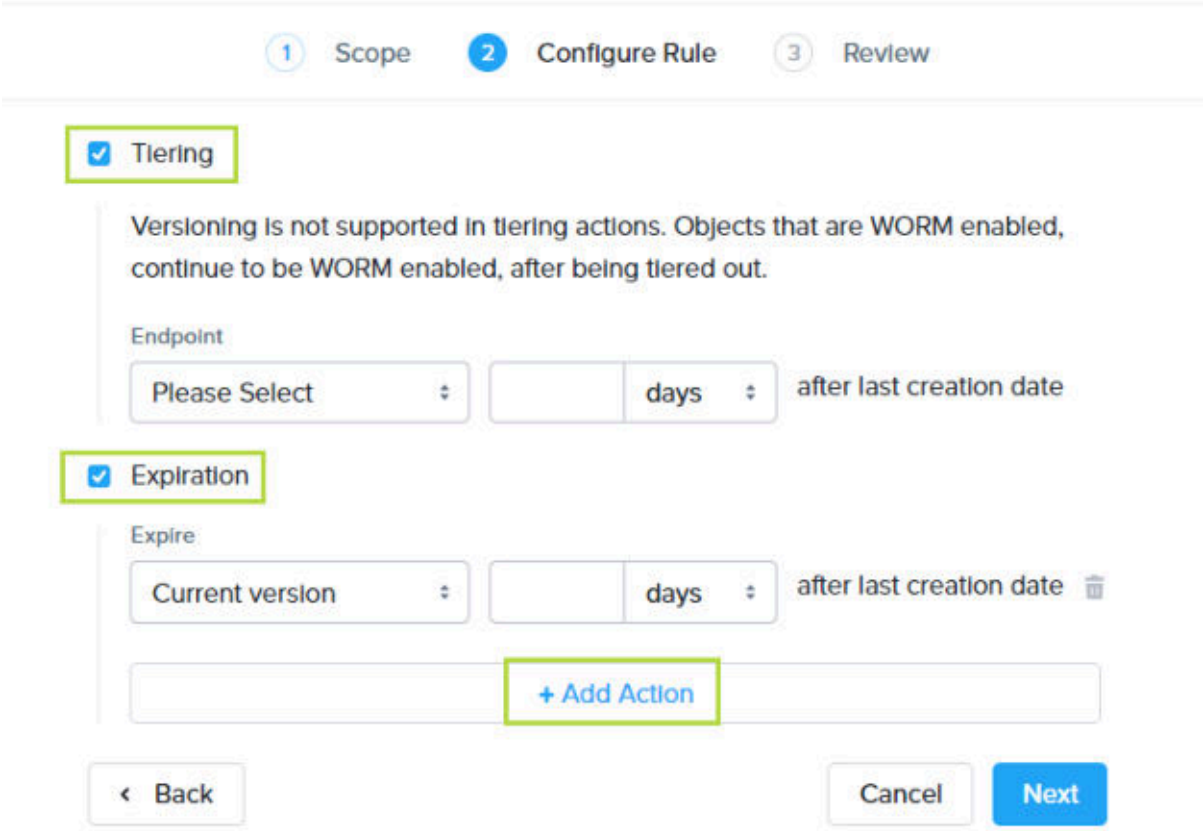

Figure 58: Creating Tiering Rule

   b. Select the **Expiration** check box.

      1. In **Expire**, select **Current version**, **Previous version**, or **Multipart uploads** according to your requirement.

         The **Previous version** option appears only for a version-enabled bucket.
      2. Enter a time period after which you want the objects to expire. You can specify the number in days, months, or years.

      You can click **Add Action** to add multiple expiration rules. You can create expiration rules for current version, previous version, and multipart uploads.

8. In the Review page, check your configurations. Then, click **Done** to create the rule.

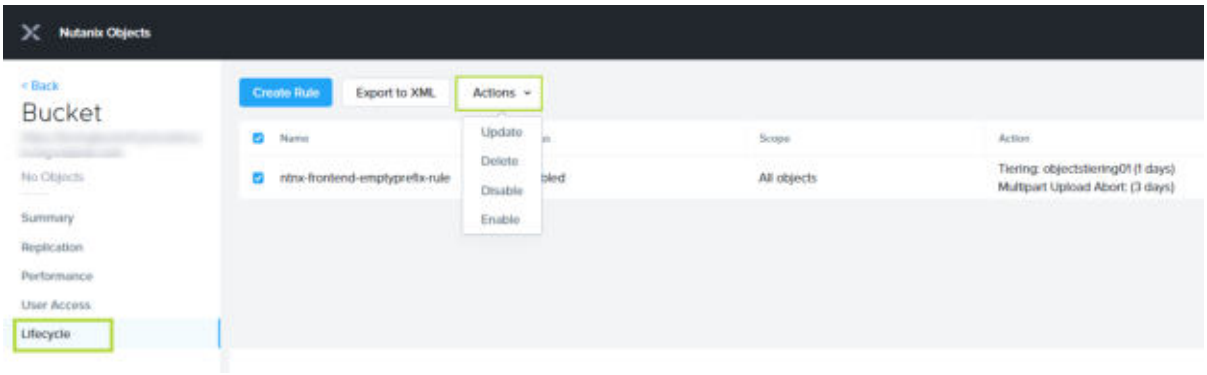

Figure 59: Tiering Rule - Actions

- You can select a bucket, and update, delete, disable, and enable the rule using the **Actions** drop-down.
- You can also export multiple rules to an XML file by clicking **Export to XML**.

The rule you just created gets enabled and appears in the Rules table.

Tiering Statistics

Cloud tiering supports endpoints such as, Nutanix Objects, Other S3 Compatible, and Azure Blob Storage. You can view the amount of object data moved to the endpoint bucket and the amount of pending data. You can also view the statistics for the source bucket (Object Store bucket) and the endpoint bucket.

Object Store Bucket - Statistics

In the Object Store instance page, click a bucket to view its tiering statistics.

You can view the following information:

- **Tier out object size**: The amount of object data that has been tiered from this source bucket to the endpoint buckets.
- **Space pending reclamation**: The amount of the deleted data and metadata left out due to incomplete operations.
- **Total size of objects marked for tiering (pending)**: The amount of the object data (eligible for tiering) that is in the process of tiering.

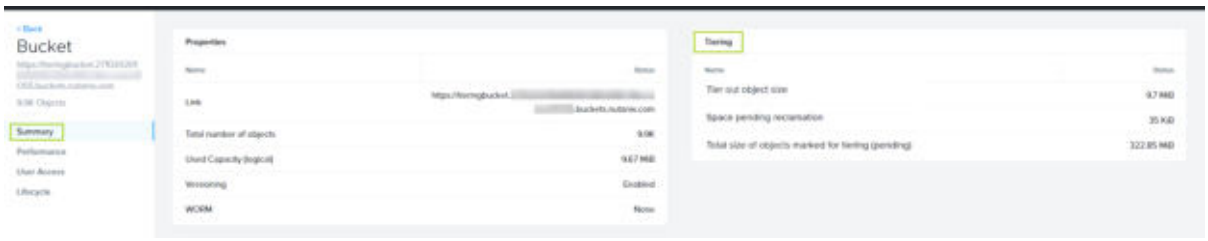

Figure 60: Source Bucket - Statistics

Endpoint Bucket - Statistics

You can view the tiering statistics for a configured endpoint-bucket. In the Object Store page, in the left-side pane, click **Tiering Endpoint**.

You can view the following information:

- **Name**: Name of the endpoint.
- **Usage**: The amount of object data moved to the endpoint bucket from the source buckets.
- **Objects pending for tiering (size)**: Amount of the object data that is in the process of tiering.
- **Average Bandwidth**: The rate at which the object data from all the source buckets is getting transferred to the endpoint bucket.
- **Endpoint Type**: Displays the type of endpoint.

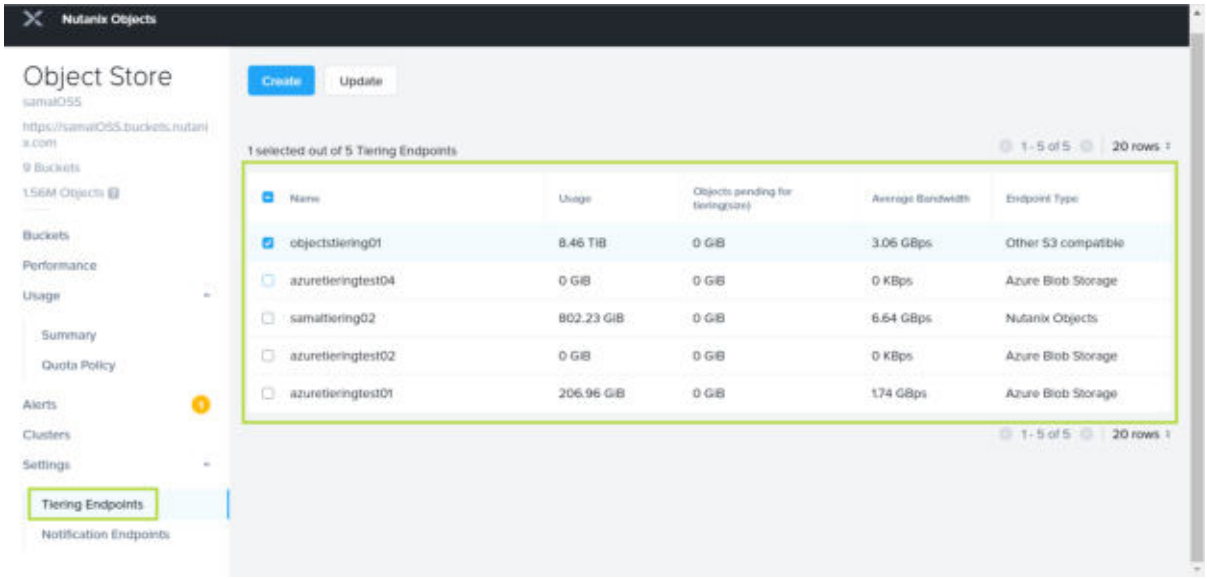

Figure 61: Endpoint Bucket - Statistics

Note: The tiering statistics are updated in the user interface after a tiering task is completed. However, when large amount of data is tiered to an endpoint, you might experience a delay in seeing the updated statistics in the user interface.

Legal Hold for Objects

Legal hold allows you to lock an object indefinitely. There is no expiration date. A legal hold remains active until an authorized user explicitly removes it. Implementing a legal hold prevents the object data from being modified or deleted.

There are various cases where you would want to implement a legal hold on your object data. For example, put an indefinite lock on objects that you want to preserve for legal cases, audits, or compliance purposes.

Note the following points about the legal hold functionality in Nutanix Objects:

- A user with the write permission on the bucket can apply a legal hold to the objects within the bucket.
- Only an administrator or bucket owner can remove the legal hold applied to the objects.

- A legal hold can be applied while writing an object or to an existing object also.
- If a bucket gets created using the Prism GUI and a user with write permissions applies a legal hold to the objects within the bucket, an administrator user needs to be set up to remove the legal hold. Contact Nutanix support to set up an administrator user with the necessary permissions. It is recommended that the user should create a bucket. In this case, the user will be the bucket owner and will have permission to remove the legal hold.
- If you are replicating a bucket created in Object v3.1 (with the legal hold applied) to a bucket created in Objects v3.0, the legal hold attributes will not get replicated.

A legal hold can be applied to the objects using the AWS S3 APIs. For more information, see Supported APIs on page 202.

WORM Bucket

Write-once-read-many (WORM) buckets protect your data and metadata. You can configure a WORM bucket to allow the creation of new objects and to prevent overwrites or deletion of the existing content for a particular retention period. By default, versioning is not enabled on a bucket. When you apply the WORM policy on a bucket, you can choose to enable versioning. In some industries, regulations or compliance rules mandate long-term records retention, sometimes for more than 7 years. For example, in the financial and health services industry, you must maintain the records in its original state, which cannot be overwritten or erased.

A non-admin user can perform WORM bucket operations only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can create WORM buckets on the buckets available in the Object stores that are assigned to the respective users.

When you increase the retention period of a bucket, the new retention period applies to the existing objects as well as the newly added objects.

When you apply any lifecycle policy to a WORM bucket with or without versions enabled, the policy is not applicable until the WORM retention period has passed.

Note:

- WORM cannot be enabled for the buckets created using NFS protocol.
- You cannot enable WORM policy while creating a bucket.

- You can set 1-day as minimum and 100 years as maximum limit for the retention period.

Operations on a WORM Bucket

This topic lists the valid and invalid operations on a WORM bucket. You can refer to this list before performing any operation on a WORM bucket.

Valid Operations

Following are the valid operations:

- Delete objects.

  Note: This operation does not delete the data, but create a delete marker on top of the existing versions of objects.

- Delete the delete marker.
- Delete the target version of object after the retention period.
- Create a new version of the object and retain old versions.
- Extend the retention period.

Invalid Operations

Following are the invalid operations:

- Enable WORM policy while creating a bucket.
- Enable WORM policy without specifying the retention period.
- Reduce the retention period.
- Delete the targeted version before the retention period.
- Change the version state.
- Does not support legal hold and governance mode.

Applying WORM Policy to Buckets from Prism Central

You can apply WORM policy only after the bucket is created. A WORM bucket allows you to create new objects and prevents you to overwrite or delete the existing content for a particular retention period.

About this task

Warning: You cannot modify or delete objects inside a WORM bucket for the specified time period.

To apply WORM policies to buckets.

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store in which the bucket exists.
3. Click **Buckets** in the left pane.

4. Select the bucket to apply the policy.
5. Click **Configure WORM**.
   The Configure WORM window appears.

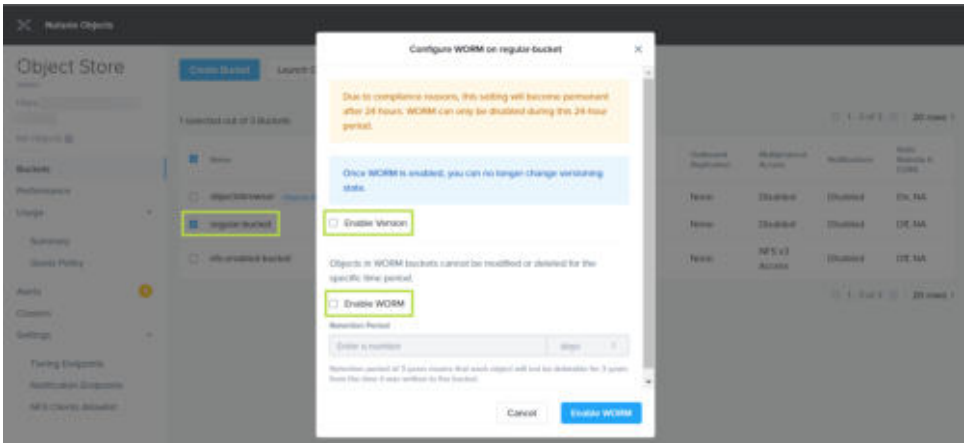

Figure 62: Configure WORM Window

Note:

- If versioning is disabled on a bucket, you can enable versioning while enabling WORM. However, once the WORM is enabled, you cannot change the versioning state.
- You cannot suspend versioning on WORM buckets.

CAUTION:

- For the legal compliance reasons, the setting becomes permanent after 24 hours. You can disable the WORM policy within the first 24 hrs of grace period.
- You can only edit the retention period to increase the length of retention. You cannot decrease the retention period.

6. (Optional) Click the **Enable Version** check box to enable versioning.
7. Click the **Enable WORM** checkbox.
8. Type the retention period (in years or months or days) in the **Retention Period** field by entering a number, and then selecting the time period from the drop-down menu.
9. Click **Enable WORM**.
   This procedure successfully applies the WORM policy to the bucket.

Configuring a Bucket for Static Website Hosting

You can use Objects to host a static website that has individual web pages with static content. To host a static website on Objects, you can configure a bucket for website hosting, and then upload your website files (objects) to the bucket. When you configure a bucket as a static website, you enable static website hosting, and optionally, add an index document and an error page. You can upload files (such as index document and error page) to the bucket using an S3 browser. The S3 browser uses the S3 protocols by providing the access and the secret key.

You can also choose to redirect to a website. Once you have configured your bucket as a static website, you can access the bucket through the object store endpoints for your bucket.

Before you begin

A non-admin user can configure static website hosting only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

> Note:
>
> - Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
> - Users with Edit Bucket permission can configure buckets for static website hosting on the buckets available in the Object stores that are assigned to the respective users.

About this task

> Note:
>
> - Static website hosting cannot be configured for the buckets created using NFS protocol.
> - Once you configure static website for a bucket, you cannot turn off this feature from the Objects user interface. To turn off static website for buckets, contact Nutanix Support.

To configure a bucket for static website hosting, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the Object Store, which contains the bucket.
   The object store opens in a new window.
3. In the **Buckets** table, select the bucket to configure it for static website hosting.

4. Click **Actions** > **Static Website**.
   The Configure Static Website window appears.

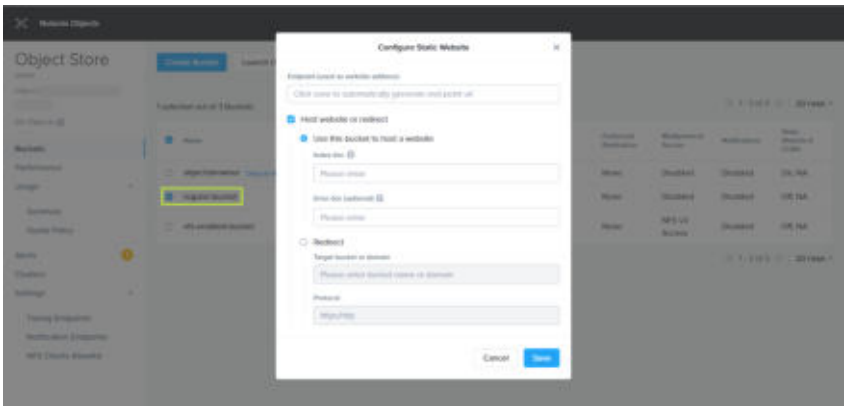

Figure 63: Configure Static Website Window

5. By default, the endpoint is auto-populated when you click **Save** at the last step.

   For example, when an endpoint auto populates, the URL will be in the format`<objectstorename>.<domain>/<bucketname>`. For example, `testobjectstore.nutanix.com/teamobjects`. However, if they have set up the DNS correctly, then you can also access the website with `<bucketname>.<objectstorename>.<domain>` endpoint using HTTP or HTTPS. For example, `https://teamobjects.testobjectstore.nutanix.com`.

6. Click the **Host Website or Redirect** check box.

   » **Use this bucket to host a website**: Select this option to use the bucket to host the website. Optionally, you can enter the name of the index document (for example, myindex.html) and an error page.

   An index document is a web page that Objects returns when you request the root of a website. It is the default page that loads when you are not requesting any specific page. After you enable static website hosting for your bucket, you can upload an HTML file with the index document name (for example, myindex.html) to your bucket. For example, if you specify no object in the URL, then the website loads the index page (myindex.html) that you have configured. If you have not configured an index document, then a website access to the root will return an access denied error message.

   A custom error page is a web page that Object returns when an error occurs. For example, if you are trying to load an object that does not exist, the website loads the error page that you have configured.

   » **Redirect**: Select this option to enter a website URL to redirect to that website. For example, when you try to access the bucket endpoint, you will be redirected to this website. Protocol used is either HTTP or HTTPS.

7. Click **Save**.

   An endpoint is auto-generated when you click **Save**. This endpoint will be the object store endpoint for your bucket and is used as the website address.

   You can now use your bucket as a static website. You can use the endpoint to test your static website.

Cross-Origin Resource Sharing (CORS) Overview

Cross-Origin Resource Sharing (CORS) allows a web application loaded in one domain to access the restricted resources that are requested from another domain. With CORS support in Objects, you can create rich web applications and allow cross-origin sharing of resources from Objects.

For example, if you upload an image (for example, **image1.png**) in a bucket (for example, first-bucket) in domain1 that contains some security-related information, and you are not allowed to make access to **image1.png** from a website (for example, **www.example.com**) on domain2. Then you can configure a CORS policy for that first-bucket to allow **www.example.com** to access the resources of the first-bucket.

> Note:
>
> - CORS cannot be configured for the buckets created using NFS protocol.
> - When you configure a bucket for static website hosting, public can have only read access to that bucket. POST, PUT, and DELETE requests on the bucket will be denied.

For configuring CORS for a bucket, create an XML document with the following. The document size is limited to 64 KB.

- Rules that identify the origins that you will allow to access your bucket.
- The operations (HTTP methods) that will support each origin.
- Other operation-specific information.

When Objects receives a cross-origin request for a bucket, it checks the CORS configuration on the bucket and uses the first `CORSRule` rule that matches the incoming browser request to enable a cross-origin request. You can add up to 100 rules to the configuration.

Following are the conditions to match the rules.

- The `Origin` header of the request must match `AllowedOrigin` elements.
- The request method (for example, GET, PUT, HEAD, and so on) or the `Access-Control-Request-Method` header for a pre-flight `OPTIONS` request must be one of the `AllowedMethod` elements.
- Every header specified in the `Access-Control-Request-Headers` request header of a pre-flight request must match an `AllowedHeader` element.

Following is an example of CORS configuration with two `CORSRule`:

```
<CORSConfiguration>
 <CORSRule>
   <AllowedOrigin>http://www.example.com</AllowedOrigin>
   <AllowedMethod>PUT</AllowedMethod>
   <AllowedMethod>DELETE</AllowedMethod>
   <AllowedHeader>*</AllowedHeader>
 </CORSRule>
 <CORSRule>
   <AllowedOrigin>*</AllowedOrigin>
   <AllowedMethod>GET</AllowedMethod>
 </CORSRule>
</CORSConfiguration>
```

- The first `CORSRule` allows cross-origin PUT and DELETE requests whose origin is `http://www.example.com` origin. The rule also allows all headers in a pre-flight `OPTIONS` request through the `Access-Control-Request-Headers` header. So, in response to any pre-flight `OPTIONS` request, Objects will return any requested headers.

  > Note: Other than pre-flight `OPTIONS` request, no other requests are denied that fail the CORS policy checks.

- The second rule allows cross-origin GET requests from all the origins.

  The `*` wild-card character refers to all the origins.

Configuring CORS on a Bucket

Cross-Origin Resource Sharing (CORS) allows a web application loaded in one domain to access the restricted resources that are requested from another domain.

Before you begin

A non-admin user can configure CORS only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

> Note:
>
> - Administrators and non-admin users are Prism Central users with the following roles:
>
>   - Administrator—A Super Admin or a Prism Admin in Prism Central.
>   - Non-admin user—A Prism Central user without any administrator privileges.
>
>   For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
> - Users with Edit Bucket permission can configure CORS on the buckets available in the object stores that are assigned to the respective users.

About this task

You set this configuration on a bucket so that the bucket can service cross-origin requests.

To configure CORS for a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the Object Store, which contains the bucket.
   The object store opens in a new window.
3. In the **Buckets** table, select the bucket to configure CORS.

4. Click **Actions** > **CORS**.

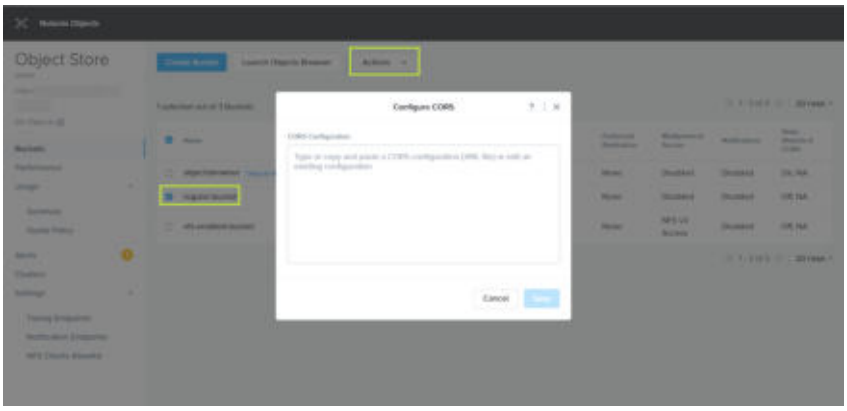

Figure 64: Configure CORS Window

The Configure CORS window appears.

5. Type or copy and paste a configuration file, or edit an existing configuration.
   The configuration file must be an XML file.

6. Click **Save**.
   The CORS configurations are saved for the bucket.

Viewing Buckets

The Buckets view allows you to view the list of buckets in the object store and access detailed information about each bucket.

Before you begin

A non-admin user can view buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with View Buckets permission can view all the buckets on an assigned object store.

About this task

To view the list of buckets, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store, and then click **Buckets**.

Figure 65: Buckets List View

A new window appears with the list of buckets in a tabular view. If you are on another page, you can click **Buckets** on the left pane to get back to this page.

The following list describes the fields that appear in the buckets list. A dash (-) is displayed in a field when a value is not available or applicable.

- **Name**: Displays the name of the bucket. Click the name to display the bucket summary. Refer to
- **Size**: Displays the size of the bucket.
- **Num. Objects**: Displays the number of objects in a bucket.
- **Versioning**: Displays if versioning is enabled or disabled in a bucket.
- **WORM**: Displays if WORM is enabled or disabled in a bucket.
- **Outbound Replication**: Displays the outbound replication status.
- **Notifications** Displays if notifications are enabled or disabled.
- **Static Website & CORS**: Displays if static website or CORS is configured for a bucket.

You can also launch the Objects Browser for this object store, click **Launch Objects Browser**.

You can identify the filters (if any) applied to the list of entities from the query field. This field displays all the filter options that are currently in use and also allows for basic filtering on the entity name. For more information on filter options, refer to Bucket Filter Options.

Buckets Filter Options

After selecting an entity type, the table displays the complete list of that entity across the registered clusters. You can filter this list by clicking the far right icon in the menu bar to display the Filters pane. This pane includes a set of fields that vary according to the type of entity. Select the desired field values to filter the list on those values. An entry appears in the current filters field for each selected value. You can remove a filter by clicking the X for that value in the current filters field.

Numeric filters have **To** and **From** fields to specify a range. These fields can take numeric values along with units. For example, the filter adjusts the scale accordingly when you type in `10 K` or `100 M`.

You can do the following in the current filters field:

- Remove a filter by clicking **X** for that filter.
- Remove all filters by clicking **Clear** (on the right).
- Use a saved filter list by selecting from the pull-down list.

Bucket Summary

This page displays the bucket properties and tiering details.

To view the Summary page of a bucket, click the name of the bucket in the buckets table. Refer to Viewing Buckets on page 120.

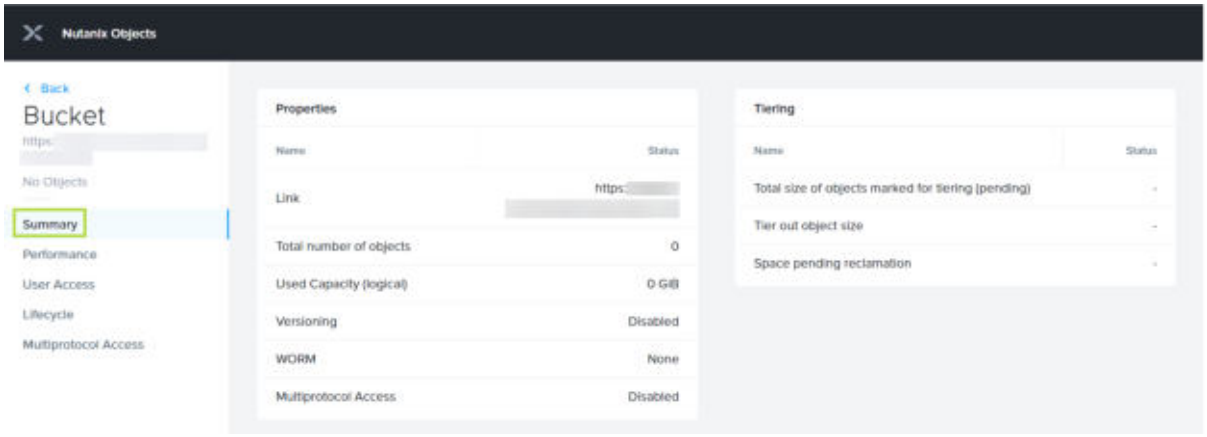

Figure 66: Bucket Summary

Updating a Bucket

You can update the bucket settings after creating the buckets and adding the objects to the buckets.

Before you begin

A non-admin user can update buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

- Users with Edit Bucket permission can perform an update bucket operation on the buckets available in the object stores that are assigned to the respective users.

About this task

Note:

- You cannot disable versioning (if enabled) but you can suspend it.
- You cannot edit multiple buckets at a time.

To edit a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the Object Store, which contains the bucket.
3. In the **Buckets** table, select the bucket for which you want to change the settings.
4. Click **Actions** > **Update**.
   The Update Bucket window appears.

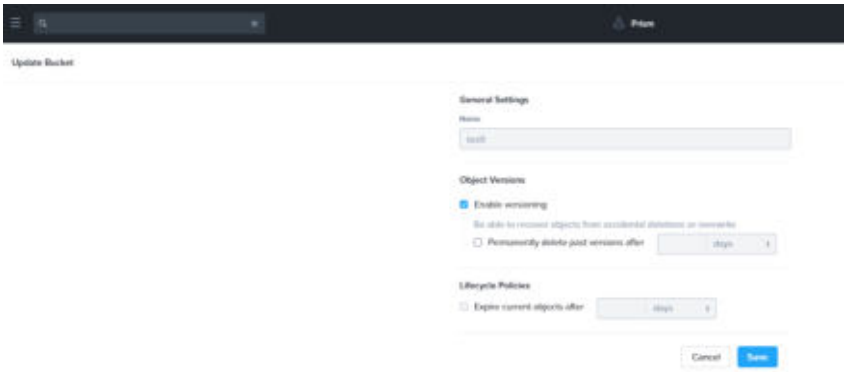

Figure 67: Update Bucket Window

5. Edit the settings that you want to change.
   For more information about the bucket settings, refer to Bucket Configuration.
6. Click **Save**.
   The changes are saved.

Sharing a Bucket

You can share a bucket with multiple users that have access keys.

Before you begin

A non-admin user can share a bucket only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can share a bucket available in the object stores that are assigned to the respective users.
- You can share only one bucket at a time.

About this task

To share a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the Object Store, which contains the bucket.

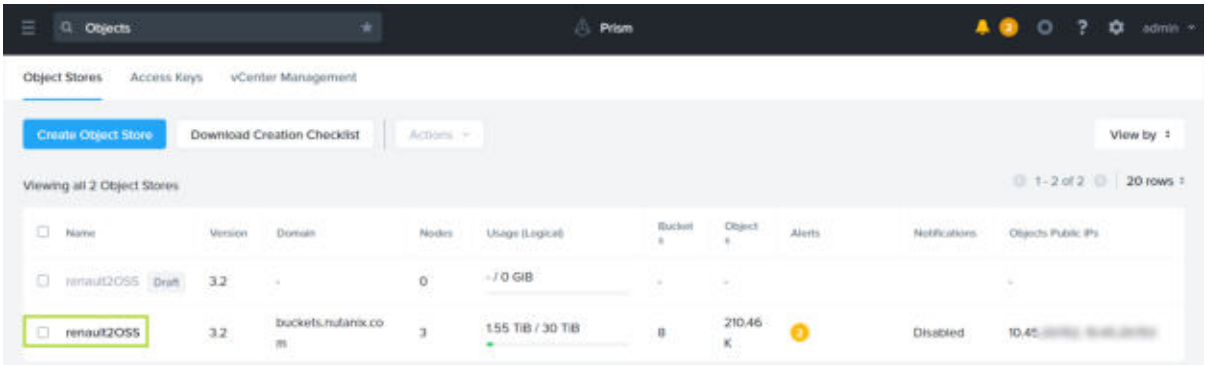

Figure 68: Object Store List

The object store opens in a new window.

3. Click **Buckets** from the left pane, and in the **Buckets** table, select the bucket that you want to share with other users.

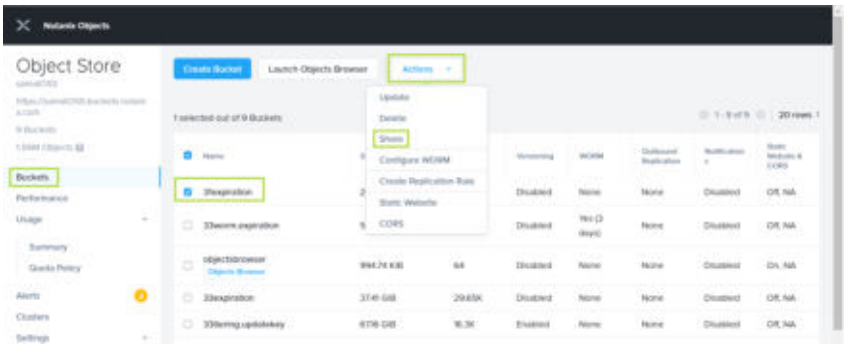

Figure 69: Sharing the Bucket

> Note: You can only add users who have access keys.

4. Click **Actions** > **Share**.
5. Type the email address of the user and set the permission for that user.

Figure 70: Share Bucket Window

To generate access keys, refer to Generating Access Key for API Users on page 85.
To know more about Bucket Permissions, refer to Bucket Access Policies.

6. To add more users, click **+ Add User**.
7. Click **Save**.
The bucket is now shared with the listed users with the allotted permissions.

What to do next

You can list the buckets that are shared with you.

For more information, refer to Listing the Shared Buckets on page 129.

Bucket Access Policies

This section describes the authorization and access policies implemented in Objects.

Nutanix supports the following list of policies on a bucket:

- **Read**: Provides read-only access to the user.
- **Write**: Provides write-only access to the user.

- **Read and write**: Provides both read and write access to the user.

> Note: To evaluate the policy for a user, the union of the user specific policy and the anonymous policy is computed. For example, if the user specific policy is `Read-only` but the anonymous policy for the bucket is `Write-only`, then the resulting policy for the user for that bucket will be `Read-Write`.

Table 5: Access Roles

Access Role	Description
User context with respect to buckets	
Owner	The user who creates the bucket. The owner can grant and revoke access to the non-admin users on the bucket.
Non-admin user	The user who gains access to the bucket based on the policy assigned by the owner of that bucket. The non-admin user can only perform operations on the bucket based on the access policy assigned by the owner.
User privilege levels (iam)	
Admin user	Admin users can read, list, and update the policy for the buckets that are owned by other users. They have certain access on the buckets that are not created by them.
Standard user	Any user who is not an admin. They have full control on the bucket which they create.

Permissions for Nutanix Objects Operations

This section describes the operations that can be performed by the owners and the admin users.

Owner Privileges

The owner can perform any S3 operations on a bucket and has read and write policy by default. They have full control over the buckets they create and can grant access to the buckets to other users (non-admin users).

Users can only list buckets owned by them. They cannot list the shared buckets.

Non-admin users can perform most operations like GET, PUT, DELETE; however, they cannot make policy changes to the buckets. Non-admin users can perform the following actions:

- Grant access to other users.
- Enable or disable versioning on the bucket.
- Enable or disable WORM policy.

Table 6: Operations Performed by an Owner

PUT Bucket
GET Bucket (List Objects), HEAD Bucket
DELETE Bucket
GET Bucket versioning
PUT Bucket versioning
GET Bucket acl
GET Bucket object lock configuration
PUT Bucket object lock configuration
DELETE Bucket lifecycle
GET Bucket lifecycle configuration
PUT Bucket lifecycle configuration

Admin Privileges

An admin user has special privileges in the system and can perform most of the operations on the buckets owned by another user.

An admin can perform the following actions:

- Create a bucket in the user interface.
- View (list) buckets from all the users.
- Modify the policy of a bucket.
- Enable or disable versioning.
- Enable or disable object locks (WORM) owned by another user.
- Delete the bucket that belongs to another user if the bucket is empty.

They cannot PUT or UPDATE objects (or initiate multipart upload), and DELETE objects.

Users can only view the buckets that they create, unless the user is an admin user.

API Accessible for the Corresponding Policies

The following table shows the list of operations allowed based on the access policy:

Table 7: API Accessible for the Corresponding Policies

S3 API	Read	Write	Read and Write
HEAD Bucket	Yes	Yes	Yes
HEAD Object	Yes	Yes	Yes
ListObjectsV2	Yes	No	Yes
ListObjects	Yes	No	Yes
GET Object	Yes	No	Yes
PUT Object	No	Yes	Yes

S3 API	Read	Write	Read and Write
COPY Object	No	Yes	Yes
DELETE Object	No	Yes	Yes
DELETE Objects	No	Yes	Yes
GET Bucket Location	Yes	No	Yes
GET Bucket versioning	Yes	No	Yes
ListObjectVersions	Yes	No	Yes
ListParts	Yes	No	Yes
MultipartUpload (CreateMultipartUpload, AbortMultipartUpload, CompleteMultipartUpload)	No	Yes	Yes
UPLOAD Part	No	Yes	Yes
Upload Part – Copy	No	Yes	Yes
PutObject	No	Yes	Yes
multipart upload parts	No	Yes	Yes
GET Bucket lifecycle configuration	Yes	No	Yes

Note:

- Only the bucket owner and admin user can perform PUT Bucket versioning and PUT Bucket lifecycle operations.
- To perform object copy operation of size more than 500 MB, do one of the following:
  - Use multipart copy part operation with part size as 500 MB.
  - Use the object copy API with a large read timeout for the client. For example, the read timeout should be larger than the default 60 seconds for Boto3 python client.

Access Policy Matrix

The following table defines access permissions between specific policies and access roles:

Access Policy Matrix

Policy	Owner	Non-admin user	Admin user (Non-owner)	Standard user with read or write access
Read and list objects, get policy, versioning config	Yes	No	Yes	Yes
Write and delete objects	Yes	No	No	Yes

Policy	Owner	Non-admin user	Admin user (Non-owner)	Standard user with read or write access
Delete bucket	Yes	No	Yes	No
Update Policy	Yes	No	Yes	No
Update Versioning config	Yes	No	Yes	No
Update worm config	Yes	No	Yes	No

Listing the Shared Buckets

Listing the Shared Buckets feature extends the S3 ListBuckets API and displays the buckets owned by or shared with the current user.

The current behavior of S3 API for ListBuckets is to list the buckets that the current owner (caller of the API) owns. If you have created a bucket, then you are the owner of that bucket. For example, when an admin creates a bucket using the Objects user interface, the owner of all such buckets is the admin user. You (owner of the buckets) can choose to share buckets with one or more IAM users with Read or Write or Read-Write permissions. IAM users with whom a bucket is shared can access the bucket and its objects subject to the permissions granted. However, the Vanilla S3 API for ListBuckets does not allow for discovery (listing) of buckets that were shared with them.

In earlier versions of Objects, if you share a bucket, you have to inform the user which bucket is shared with them. However, with Objects 2.1 and later versions, Listing the Shared Bucket feature extends the S3 ListBuckets API and displays the buckets owned by or shared with the current user. If you delete a bucket that is shared, the bucket will automatically be removed from the list.

> Note:
>
> - This feature is a departure from the S3 specifications and is enabled by default. To disable the feature, contact Nutanix Support.
> - This feature is only enabled for sharing relations created using Objects 2.1 or later versions. You cannot list the buckets created using earlier versions of Objects. A workaround to enable this feature retroactively is to have the owner of the bucket remove all the sharing-relations and recreate them using the latest Objects version.

For example, an Admin created three buckets B0, B1, and B2. B1 was shared with User1 with Read-Write access and B2 was shared with User2 with Write access. Following are the outputs of the ListBuckets API call before and after the introduction of Listing the Shared Buckets feature.

Before the introduction of Listing the Shared Bucket feature:

```
Admin B0 B1 B2
User1
User2
```

After the introduction of Listing the Shared Bucket feature:

```
Admin B0 B1 B2
User1 B1
User2 B2
```

Viewing Bucket Users

You can view all the users and their permissions for a particular bucket.

Before you begin

A non-admin user can view bucket users only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with View Buckets permission can view all the buckets on an assigned object store.

About this task

To view the usage of the bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store, which stores the bucket, and then click **Buckets**.
3. Click the name of the bucket.

Figure 71: Buckets List

4. Click **User Access**.

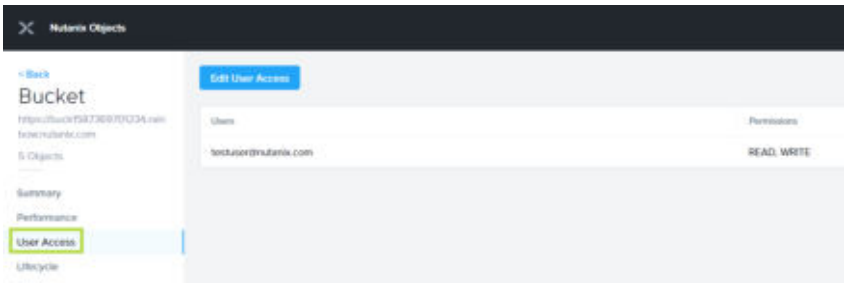

Figure 72: List of Users

The table displays the list of users and their permissions.

You can edit the user access by clicking **Edit User Access**.

For more information about editing user access, refer to the procedure from Step 4 in the Sharing a Bucket on page 123 topic.

Deleting a Bucket

If you no longer require a bucket, you can delete the bucket.

Before you begin

A non-admin user can delete a bucket only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Bucket permission can delete a bucket available in the object stores that are assigned to the respective users.

About this task

You can delete the following:

- Multiple empty buckets at a time.

- A bucket with versioning enabled; however, it permanently deletes all the versions of the buckets.

> Note: You cannot delete buckets containing objects and WORM enabled buckets

To delete a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store, which stores the bucket, and then click **Buckets**.
3. Click the check box next to the name of the bucket that you want to delete.
4. Click **Actions** > **Delete**.

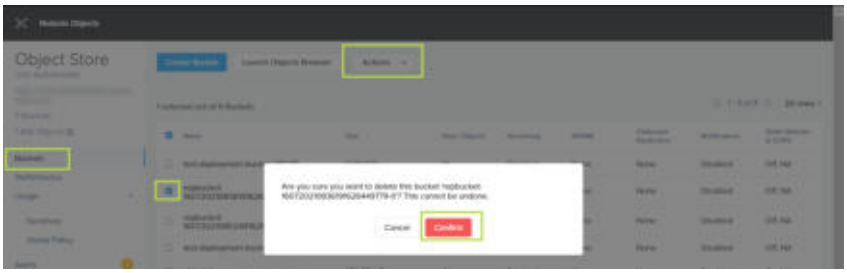

Figure 73: Delete Bucket Window

5. On the dialog box, click **Confirm**.
   A message appears to confirm the deletion of the bucket.

OBJECTS BROWSER

The Objects Browser is a User Interface (UI) that allows the users to directly launch the object store instance in a web browser and perform bucket and object level operations. This eliminates the steps of logging into Prism Central and launching the Objects service.

The administrators use Prism Central to manage the object stores. When an administrator creates an object store, a default bucket (objectsbrowser) gets created. The static website hosting capability is already enabled on this default bucket. The Objects Browser UI is hosted on this default bucket.

To launch the Objects Browser, select the object store and click **Actions** > **Launch Objects Browser**. Objects Browser will be opened in a new window. You can also launch the Objects Browser from the buckets list page (Refer to Viewing Buckets on page 120).

Alternatively, if the you have multiple load balancer IP addresses for the object store, you can click any of the IP addresses to launch the Objects Browser for that object store.

Figure 74: Prism Central - Launch Objects Browser

The administrator can also share the object store URL with the user to access the Objects Browser UI. The URL can be formed using the Objects Public IP address.

`http://objects-public-ipaddress/objectsbrowser`

> Note: You can use http or https to access the Objects Browser.

Refer to the following image to understand the URL formation.

Figure 75: Objects Browser URL

> Note:
>
> - Refreshing the page would log out the user and cancel the pending and in-progress uploads.
> - Any uploads taking more than 60 minutes gets terminated and retried three times before getting canceled.

Administrator Workflow: Objects Browser

The administrator needs to perform the following steps to grant access to the Objects Browser to a particular user.

About this task

Note:

- Objects Browser does not store the access and secret keys.
- If the user refreshes the page with the refresh button of the web browser, they will need to provide the credentials again, and any unsaved changes will be lost (including any pending or in-progress uploads).

To grant access to Objects Browser, do the following:

Procedure

1. Add the IAM user.
   For more information, refer to Generating Access Key for API Users on page 85.
   After the administrator adds the user, the access and secret keys are generated.
2. (Optional) Grant access permissions (read, write, or both) to the buckets.
   For more information, refer to Sharing a Bucket on page 123.

   Note: This step is optional and can be performed if the administrator wants to share bucket of another user with a new user. If the administrator does not perform this step, the buckets list will be empty for the IAM users. The users can still create buckets from the Objects Browser UI.
3. Share the Objects Browser URL with a user to launch the object store in a web browser.
4. Share the access and secret keys that you generated for the user.
   The user would require to enter the keys in the login page of the Objects Browser UI.

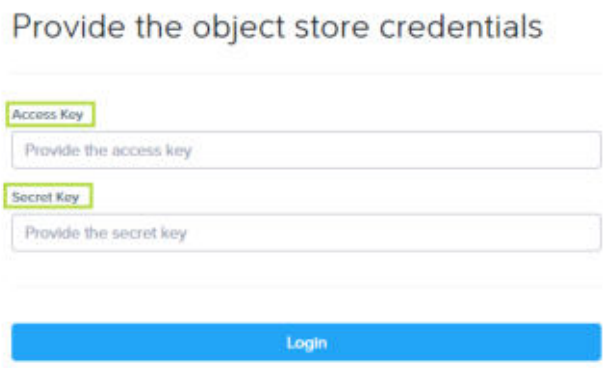

Figure 76: Objects Browser Login Page

What to do next

The user can log in to the Objects Browser UI and perform CRUD operations within the object store. For example, create buckets, upload objects, assign a tag to objects, and so on. Refer to Supported Operations on page 136.

Launching the Objects Browser

You can use the object store URL and keys shared by the administrator to launch the Objects Browser UI.

About this task

To launch the Objects Browser, do the following:

Procedure

1. Open the object store URL shared by the administrator in a web browser.
2. In the Objects Browser login page, enter the **Access** and **Secret** keys to access the object store instance.

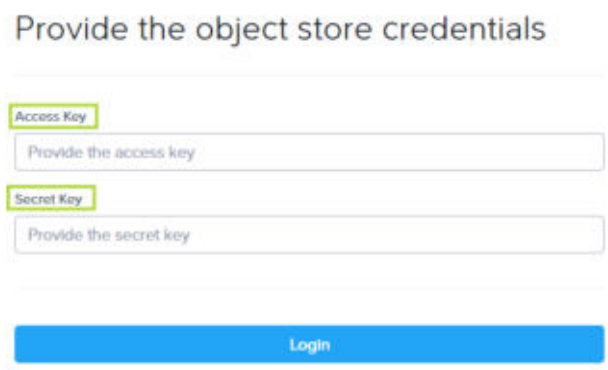

Figure 77: Objects Browser Login Page

Note: Access and Secret keys will be shared by the Administrator.

The object store will be opened in the Objects Browser.

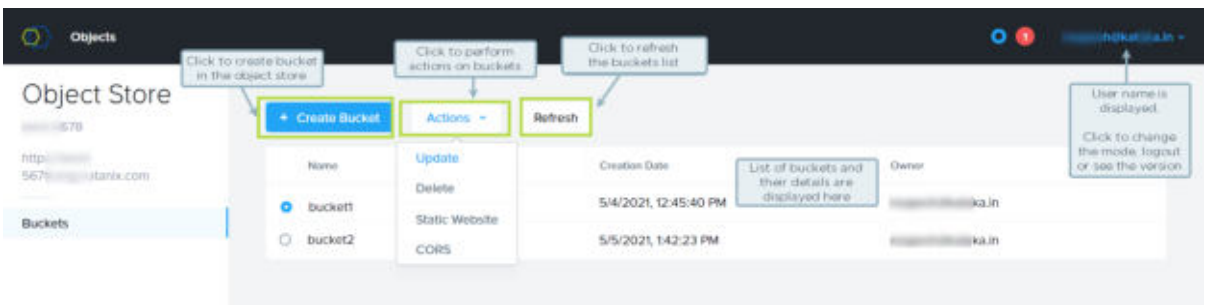

Figure 78: Object Browser Landing Page

What to do next

You can now perform CRUD operations on buckets and objects. Refer to Supported Operations on page 136. You can also change the mode from light to dark, logout, or see the version by clicking the user name at top-right corner.

Supported Operations

This section describes the various CRUD operations that a user can perform in buckets and objects of an object store from the Objects Browser UI.

Bucket Operations

After you log into the Objects Browser, all the buckets that the administrator shared with you are listed with creation date and owner information.

> Note: You must have read and write permissions to perform various CRUD operations on the buckets and objects.

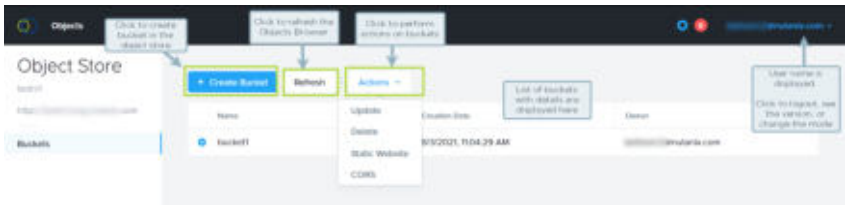

Figure 79: Objects Browser - Bucket CRUD Operations

You can perform the following operations at the bucket level:

- **Create Bucket**: Allows you to create buckets. For more information, see Creating an S3 Bucket Using Objects Browser on page 137.
- **Lifecycle**: Allows you to create lifecycle rules.

  Click the name of the bucket, and then click the **Lifecycle** option at the left pane. For information, refer to Creating Lifecycle Rules on page 143.
- You can use the **Actions** list to update bucket properties, host a static websites, configure CORS, and delete a bucket.
  - **Update**: Allows you to update the bucket properties.

    To update bucket properties, select the bucket, and then click **Actions** > **Update**.
  - **Delete**: Allows you to delete a bucket.

    To delete a bucket, select the bucket, and then click **Actions** > **Delete**.

    > Note: You can only delete an empty bucket. In the case of a version-enabled bucket, the delete operation performed on an object is not permanent. The object gets removed from

the list and moved to Recycle Bin and a Delete Marker gets created. For more information, see Understanding Object Versions on page 155.

- **Static Website**: Allows you to configure a bucket to host a static website.

  You can configure a bucket for website hosting, and then upload your website files (objects) to the bucket. For more information see, Configuring a Bucket for Static Website Hosting on page 146.

- **CORS**: Allows you to configure CORS on a bucket.

  This allows the bucket to service cross-origin requests. For more information see, Configuring CORS on a Bucket on page 147.

Creating an S3 Bucket Using Objects Browser

You can create, modify, and delete a bucket using Objects Browser.

About this task

Note:

- Make sure that the bucket names are unique for all users.
- You cannot configure a WORM bucket while creating a bucket. You can edit WORM policies only after creating a bucket.
- You cannot enable multi-protocol access on an S3 bucket.

To create and configure a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu >**Services** > **Objects**
2. Click the name of object store in which you want to create a bucket and launch the Objects Browser.

   For information on launching the Objects Browser, see Launching the Objects Browser on page 135.

   The object store opens in a new window.

3. Click **Create Bucket**.

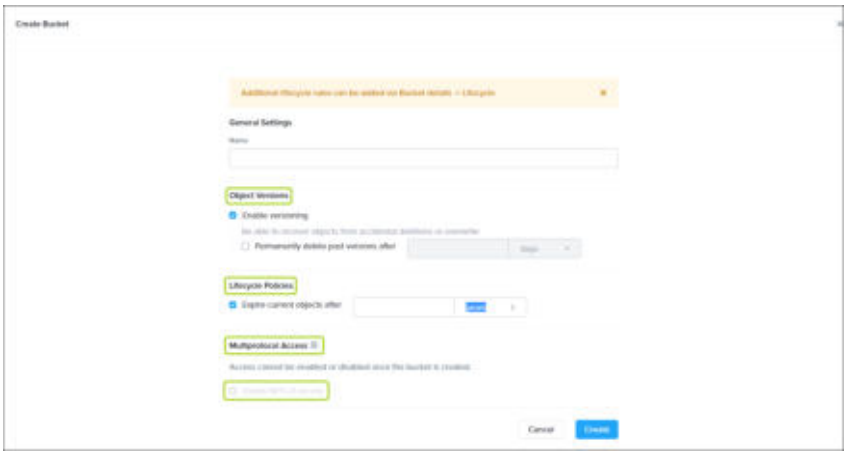

Figure 80: Create Bucket Window

The **Create Bucket** window appears.

4. On the General Settings section, type a name for your bucket.
   For more information about naming buckets, see Buckets Naming Conventions.
5. (Optional) On the Object Versions section, configure the following:
   For more information about versions, see Object Versioning.
   a. **Enable Versioning**: Select this check box to enable versioning of objects.
      This allows you to keep all versions of an object in the same bucket. For more information about versions, see Object Versioning.
      To apply life cycle policy with versioning enabled, see Rules for Lifecycle Policy When Object Versioning is Enabled on page 103.
   b. **Permanently delete past versions after**: Select this check box to enter a time period in days, months, or years. After the specified time, all the older versions of the objects are deleted.

      Note: Enabling versioning allows you to recover objects from accidental deletion and overwrite.
6. (Optional) On the **Lifecycle Policies** section, select the **Expire current objects after** option to enable lifecycle policies.
   Type a time period after which the current version of the object expires. You can specify the time in days, months, or years.

   Note:

   - For more information, see Lifecycle Policies.
   - If versioning is not enabled, the current object is deleted permanently. When you enable versioning, multiple versions of the same object is maintained in a bucket.
   - Multi-protocol access cannot be enabled on the S3 bucket. If you want to create buckets with multi-protocol access, see Creating and Configuring an NFS Bucket on page 94.

7. Click **Save**.
   The bucket is created successfully.

What to do next

After creating a bucket, you can create objects through S3 APIs and manage them. For more information, see Supported S3 APIs.

You can also perform various actions on a bucket, such as configuring static websites and configuring CORS on a bucket.

- For more information on creating lifecycle rules, see Creating Lifecycle Rules on page 143.
- For more information on configuring static websites, see Configuring a Bucket for Static Website Hosting on page 146.
- For more information on configuring CORS on a bucket, see Configuring CORS on a Bucket on page 147.

Updating a Bucket

After you create a bucket, you can update the bucket settings using Objects Browser.

About this task

To update a bucket settings, do the following:

Procedure

1. Log on to the Objects Browser.
   For more information about launching an Objects Browser, see Administrator Workflow: Objects Browser and Launching the Objects Browser.
2. From the object store, select the bucket that you need to update.

3. Click **Actions** menu > **Update**.

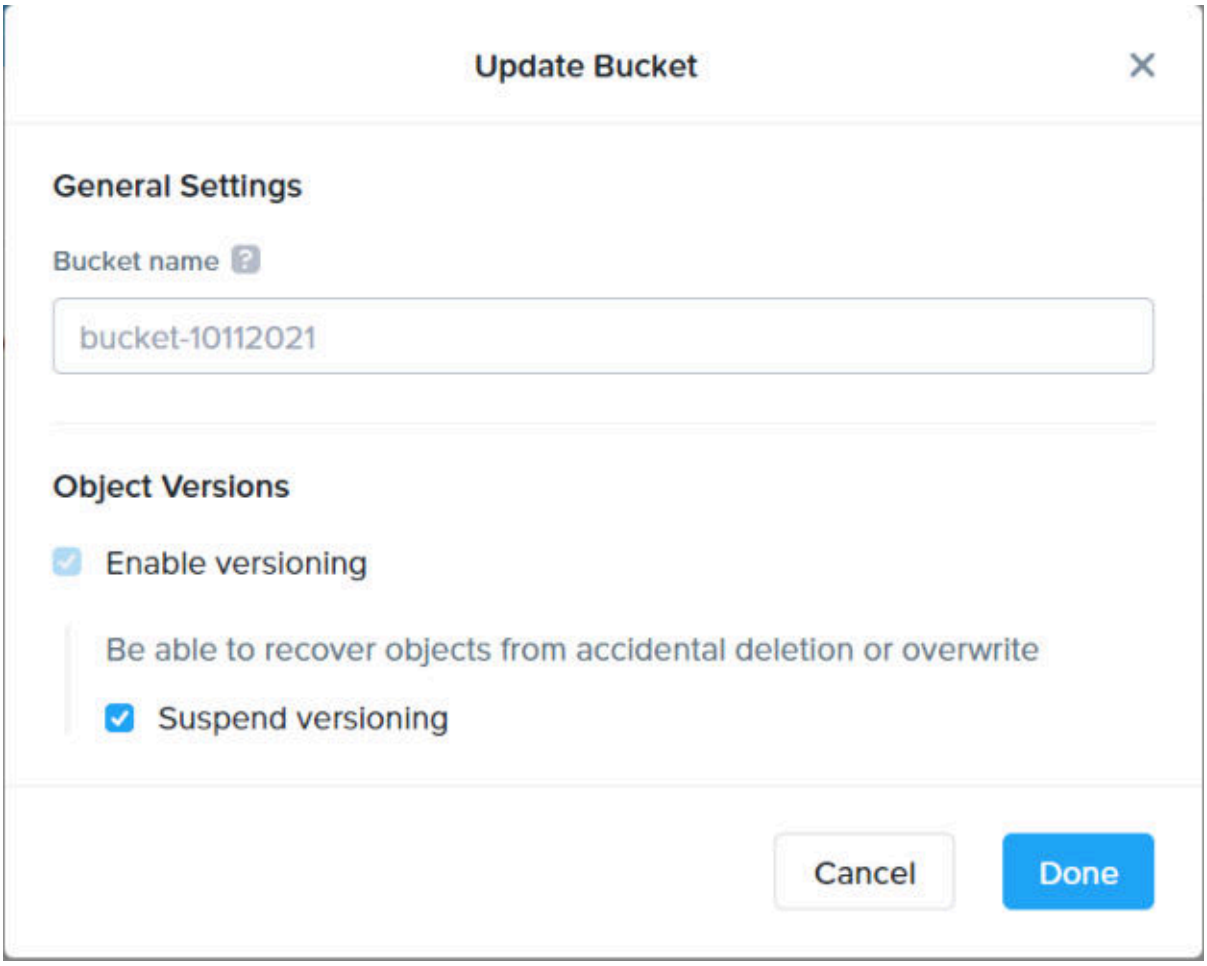

Figure 81: Update Bucket Window

4. On the Object Versions section, select one of the check boxes to enable or disable versioning of objects in the bucket.

   a. **Enable versioning**: Select this check box to enable versioning on objects and to keep all the versions of the object on the same bucket.

      Note: Select this option to recover objects from accidental deletion or overwrite.

   b. **Suspend versioning**: Select this check box to disable versioning of objects in a bucket.

      When you suspend versioning, accumulation of the new object versions is stopped. However, versions of objects already existing in the bucket are retained.

5. Click **Done**.
   The updated settings are applied to the bucket successfully.

Creating and Configuring an NFS Bucket Using Objects Browser

You can create and configure the bucket with multi-protocol access. If you are enabling multi-protocol access on a bucket, S3 features such as lifecycle policies, versioning, WORM, replication, static website, CORS, and notifications cannot be enabled on that bucket. You can perform update (only multi-protocol access configurations), delete, or share operations on these buckets.

Before you begin

Make sure you are aware of the uses cases, recommendations, NFS-S3 interoperability, and limitations of NFS on Objects. See Use Cases and Recommendations for NFS on Objects on page 8, NFS-S3 Interoperability on page 9 and Limitations of NFS on page 50.

About this task

Note:

- Ensure that the bucket names are unique for all users.
- As the Objects-NFS does not support NLM (Network Lock Manager), lock option is not required while mounting the NFS bucket.
- The total and available bytes returned in the FSSTAT response denotes the logical capacity and logical available space and not the physical capacity of the cluster which also takes RF2 into consideration.

To create and configure a bucket with multi-protocol access, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu >**Services** > **Objects**
2. Click the name of object store in which you want to create a bucket.
   For information on launching the Objects Browser, see Launching the Objects Browser on page 135.
   The object store opens in a new window.
3. Click **Create Bucket**.

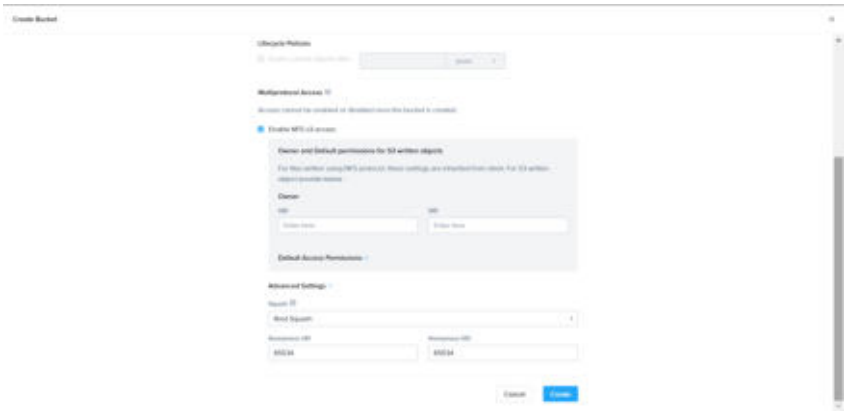

Figure 82: Create Bucket Window

4. On the General Settings section, type a name for your bucket.
   For more information about naming buckets, see Buckets Naming Conventions.

   Note: Versioning and lifecycle policies cannot be enabled on buckets with multi-protocol access.

   To create and configure buckets for S3 features, see Creating an S3 Bucket Using Objects Browser on page 137

5. On the Multiprotocol Access section, select **Enable NFS v3 Access**.

   Note: Access cannot be enabled or disabled after the bucket is created.

6. For owner and default permissions for S3 written objects, do the following.

   Note: For files written using NFS protocol, these settings are inherited from the client.

   a. In the Owner section, enter the **UID** and **GID**.
      Any object (file or directory) created from S3 protocol have UserID (UID) and GroupID (GID) in the NFS namespace.

   b. In the Default Access Permissions section, by default read, write, and execute permissions are set for the files and directories for owner, group (other users in a group), and others (users that are not part of any group).

      Following are the default permissions for files:

      - **Owner**: Read, Write
      - **Group**: Read
      - **Others**: Read

      Following are the default permissions for directories:

      - **Owner**: Read, Write, Execute
      - **Group**: Read, Execute
      - **Others**: Read, Execute

      You can change these permissions as needed.

7. In the Advanced Settings section, select any one of the following squash options.

   » **None**: Select this option if you do not want to convert the UID and GID of the users on the server.

   » **Root Squash**: Select this option if the user has root privileges and you want to convert the UID and GID to an anonymous UID and GID on the server. The anonymous UID and GID are automatically generated, however, you can change it.

   » **All Squash**: Select this option if you want to map all users to a single identity. This will convert the UID and GID of all users to an anonymous UID and GID on the server.

8. Click **Create**.
   The bucket is created successfully.

9. Note: Before mounting a bucket, add the client to the NFS allowlist. Only the client present in the NFS allowlist will be given access to the NFS buckets.

   For more information on adding and managing clients to NFS allowlist, see Managing NFS Allowlist on page 98.

   Mount the bucket from the client VM.

   ```
   $ sudo mount -tnfs -o nfsvers=3,proto=tcp, soft -v objectstore-endpoint-ip-address:/bucketname path/to/mount
   ```

   For example,

   ```
   $ sudo mount -tnfs -o nfsvers=3,proto=tcp, soft -v 10.45.53.75:/test-nfs-bucket home/folder-1/mnt-point
   ```

   The bucket is mounted successfully.

What to do next

After creating a bucket, you create directories, files, and symbolic links from the NFS namespace. You can also perform object operations from the Objects Browser or the S3 APIs. For more information, see Object CRUD Operations on page 148 and Supported S3 APIs sections.

Creating Lifecycle Rules

Lifecycle policy enables you to create or update a set of rules that define actions that Nutanix Objects applies to a group of objects within a bucket. With these policies, you can expire objects when no longer required.

Before you begin

Refer to Lifecycle Policies on page 102.

About this task

Note: Rules or any updates to the rules get applied to the new objects that you create and does not apply to the objects existing before the rule creation or update.

To create lifecycle rule, do the following:

Procedure

1. Log on to the Objects Browser.
2. In the list of buckets, click the name of the bucket for which you want to create rule.
3. Click **Lifecycle**.
4. Click **Create Rule**.

   When you have no rules created, you can directly import the XML. **Import XML** option is not available if you already have rules to avoid overwriting.

5. In the Scope section, do the following, and then click **Next**:

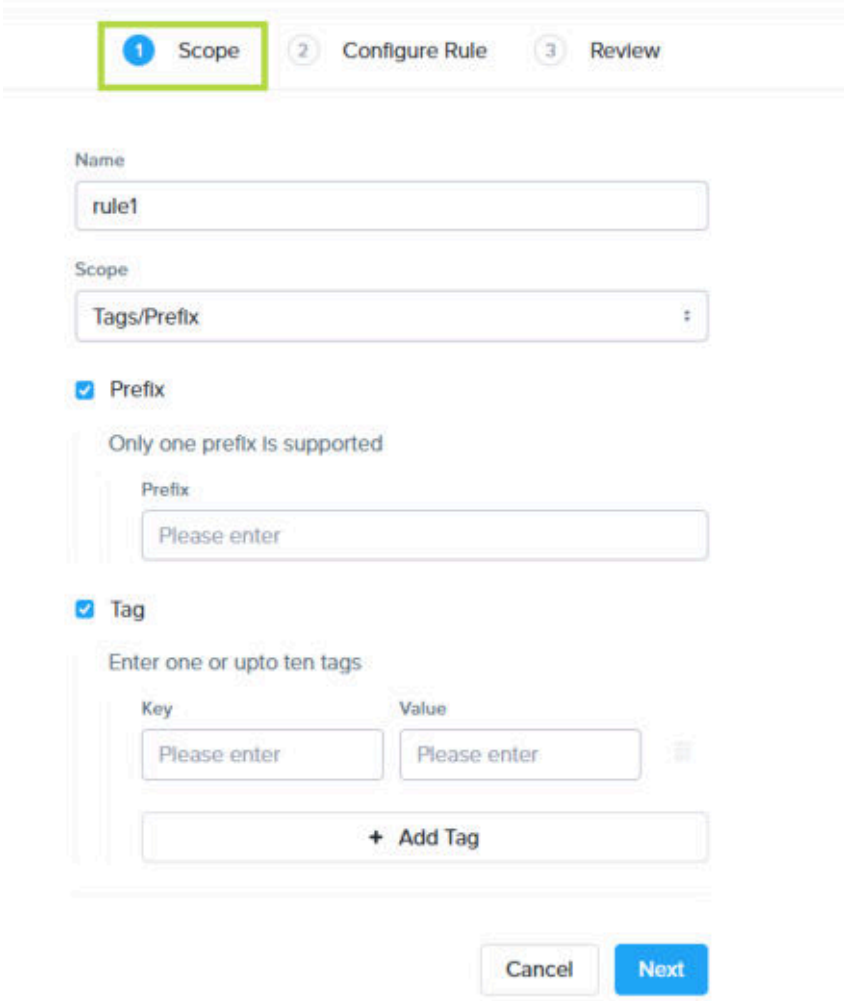

Figure 83: Defining Scope for a Rule

a. **Name**: Enter a name that identifies the rule you are creating.

b. **Scope**: Select the scope of the rule to either all objects of that bucket, or to tags and prefixes.

   » **Prefix**: You can enter only one prefix.

   » **Tag**: You can enter up to 10 tags in key value pair.

   Click **+ Add Tag** to add more tags.

6. In the Configure Rule section, do the following, and then click **Next**:

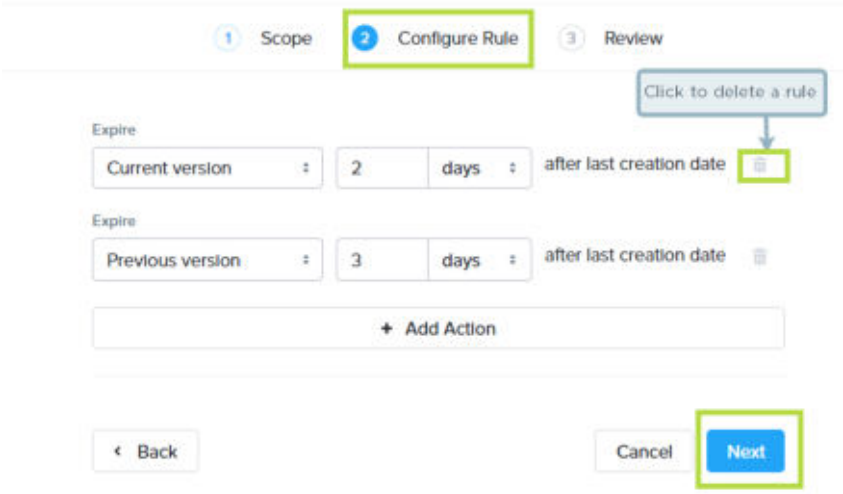

Figure 84: Configuring Rule

> Note: You cannot configure the tiering endpoints from the Objects Browser. You can configure it from the Objects UI and the same is visible as Read Only in the Objects Brower UI.

a. **Expire** Select which version to expire: **Current version**, **Previous version**, or **Multipart uploads** according to your requirement.

   - The **Previous version** option appears only for a version-enabled bucket.
   - Duplicate expiration field is not allowed. For example, you cannot create two rules for current version expiration.

   > Note: Multipart uploads expiration should not be specified with tag-based filters.

b. **after last creation date** Specify the number in days, months, or years after which that respective version of the object expires after last creation date of that object.

   You can click **Add Action** to add up to three expiration rules. You can create expiration rules for current version, previous version, and multipart uploads.

Click **Delete Icon** to delete the rule.

7. In the **Review** section, review the scope and actions, and then click **Done**.

   - You can select a bucket, and update, delete, disable, and enable the rule using the **Actions** drop-down.
   - You can also export the multiple rules to an XML file by clicking **Export to XML**.

   > Note: Actions listed are executed in sequence.

   The rule you just created gets enabled and appears in the Rules table.

Configuring a Bucket for Static Website Hosting

You can use Objects Browser to host a static website that has individual web pages with static content. To host a static website on Objects, you can configure a bucket for website hosting, and then upload your website files (objects) to the bucket. When you configure a bucket as a static website, you enable static website hosting, and optionally, add an index document and an error page. You can upload files (such as index document and error page) to the bucket using the Objects Browser or S3 browser. The S3 browser uses the S3 protocols by providing the access and the secret key. You can also choose to redirect to a website. Once you have configured your bucket as a static website, you can access the bucket through the object store endpoints for your bucket.

About this task

To configure a bucket for static website hosting, do the following:

Procedure

1. Log on to the Objects Browser.
2. In the **Buckets** table, select the bucket to configure it for static website hosting.
3. Click **Actions** > **Static Website**.
   The Configure Static Website window appears.

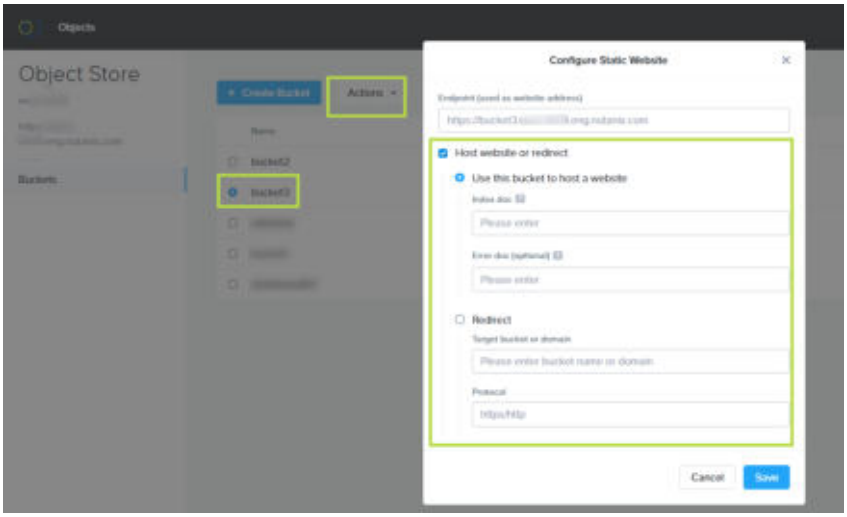

Figure 85: Configure Static Website Window

4. By default, the endpoint is auto-populated when you click **Save** at the last step.
   For example, when an endpoint auto populates, the URL will be in the format `objectstorename.domain/bucketname`. For example, `objectstore.nutanix.com/teamobjects`. However, if they have set up the DNS correctly, then you can also access the website with `bucketname.objectstorename.domain` endpoint using HTTP or HTTPS. For example, `https://teamobjects.objectstore.nutanix.com`.
5. Click the **Host Website or Redirect** check box.
   - **Use this bucket to host a website**: Select this option to use the bucket to host the website. Optionally, you can enter the name of the index document (for example, myindex.html) and an error page.

     An index document is a web page that Objects returns when you request the root of a website. It is the default page that loads when you are not requesting any specific page.

After you enable static website hosting for your bucket, you can upload an HTML file with the index document name (for example, myindex.html) to your bucket. For example, if you specify no object in the URL, then the website loads the index page (myindex.html) that you have configured. If you have not configured an index document, then a website access to the root will return an access denied error message.

A custom error page is a web page that Object returns when an error occurs. For example, if you are trying to load an object that does not exist, the website loads the error page that you have configured.

» **Redirect**: Select this option to enter a website URL to redirect to that website. For example, when you try to access the bucket endpoint, you will be redirected to this website. Protocol used is either HTTP or HTTPS.

To remove the static website configuration from the Objects Browser, uncheck **Host website or redirect**, and then click **Save**.

6. Click **Save**.

   An endpoint is auto-generated when you click **Save**. This endpoint will be the object store endpoint for your bucket and is used as the website address.

   You can now use your bucket as a static website. You can use the endpoint to test your static website.

Configuring CORS on a Bucket

Cross-Origin Resource Sharing (CORS) allows a web application loaded in one domain to access the restricted resources that are requested from another domain.

Before you begin

Refer to Cross-Origin Resource Sharing (CORS) Overview on page 118.

About this task

You set this configuration on a bucket so that the bucket can service cross-origin requests.

To configure CORS for a bucket, do the following:

Procedure

1. Log on to the Objects Browser.
2. In the **Buckets** table, select the bucket to configure CORS.

3. Click **Actions** > **CORS**.

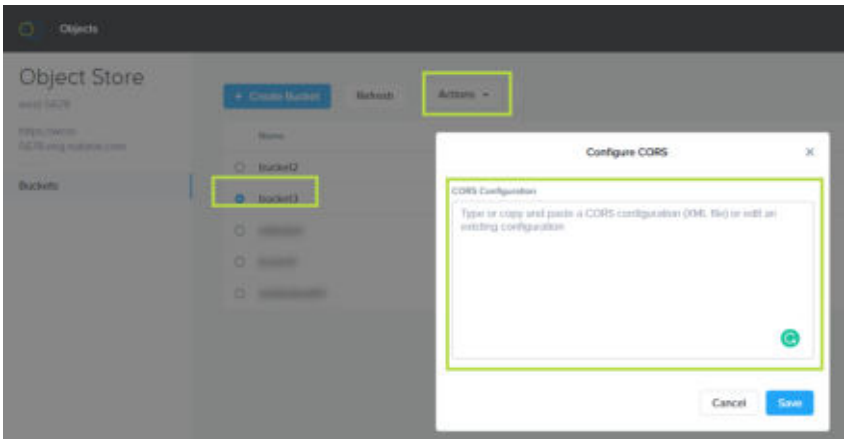

Figure 86: Configure CORS Window

The Configure CORS window appears.

4. Type or copy and paste a configuration file, or edit an existing configuration.
   The configuration file must be an XML file.
5. Click **Save**.
   The CORS configurations are saved for the bucket.

Object CRUD Operations

You can perform various CRUD operations on the objects within a bucket.

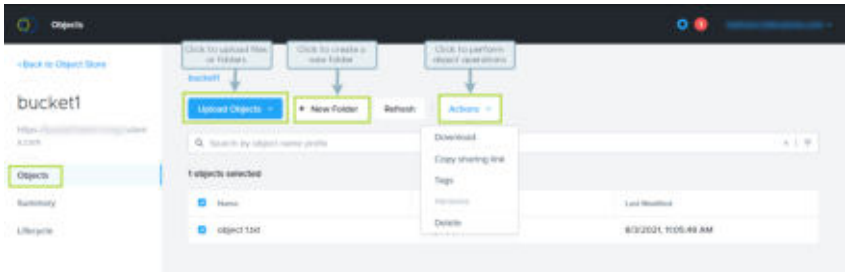

Figure 87: Objects Browser - Object CRUD Operations

The CRUD operations that you can perform on an object are as follows:

- **Upload Objects**: Allows you to upload files, or folders containing multiple files.

  For information on uploading objects and creating new folders, see Uploading an Object Using Objects Browser on page 150.

- **New Folder**: Allows you to create folders to organize and manage your data within a bucket.

  Note: You cannot create new folders in NFS-enabled buckets.

  For information on uploading objects and creating new folders, see Uploading an Object Using Objects Browser on page 150.

- You can use the **Actions** list to download an object, add tags for an object, manage object versions for version-enabled buckets, and delete an object.

  - **Download**: Allows you to download an object.

    To download an object, select the object the object, and then click **Actions** > **Download**.

    > Note: The object is downloaded to the default download folder of your browser. You can also download the object by clicking the name of the object in the objects list.

  - **Copy sharing link**: Allows you generate a link to the object that you can share with other users.

    To share an object with another user, select the object, and click **Actions** > **Copy sharing link**.

    Other users can directly open the shared object using the link and perform actions depending on the permissions. Users are prompted to log on, when they open the link.

  - **Tags**: Allows you to add tags to an object.

    > Note: You can add up to 10 tags for an object.

    For more information on adding tags, see Adding Tags to an Object Using Objects Browser on page 152.

  - **Versions**: Allows you to view and manage the different versions of an object. Versioning allows you to revert an earlier version of an object as its latest version.

    In the Versions page, you can revert, delete, and download the object versions. For more information, see Understanding Object Versions on page 155.

  - **Delete**: Allows you to delete any object.

    To delete an object, select the object the object, and then click **Actions** > **Delete**.

    > Note: For versioned buckets, the delete operation performed on an object is not permanent and a delete marker is created. You can view the deleted objects, and if needed, retrieve it from the **Recycle Bin**.
    >
    > For non-versioned buckets, the delete operation performed on an object is permanent and cannot be retrieved. Non-versioned buckets do not have a **Recycle Bin**.

- **Search**: You can perform prefix-based search in the search bar. You can perform the search in the objects list page and the Recycle Bin. You can enter any prefix as a search keyword and the objects starting with that keyword will be listed. For example, if you search for the prefix copy, all objects whose name start with the keyword copy are listed.

  You can also search for exact name matches. For example, if you search for the keyword copy.txt, the search result is an exact match. It will not search for the keyword starting with the keyword copy.text, but will find an exact match.

  > Note:
  >
  > - The search you do is limited to the folder that you are working on. For example, if you do a search from within the BigData folder, then the search results appear only from the objects in the BigData folder, and not from the entire list of objects.
  > - The search is case sensitive.

Object Naming Conventions

The object key name is a sequence of UTF-8 characters not exceeding 1024 bytes.

You can use the following characters to create an object key name or prefixes:

- Alphanumeric characters—lowercase letters (a-z), uppercase letters (A-Z), and numbers (0-9)
- Special characters—forward slash (/), exclamation mark (!), hyphen (-), underscore (_), period (.), asterisk (*), single quote ('), open parenthesis ((), and close parenthesis ())

Consider the following limitations before you create an object key name:

- Only UTF-8 character set is supported.
- The length of the object key name cannot exceed 1024 bytes.
- Special characters, such as pound (#) and percent (%) are not supported.

> Note: The following restrictions apply to the object key names:

Uploading an Object Using Objects Browser

After creating a bucket, you can use the Objects Browser to select and upload files or folders to a bucket.

About this task

To upload files or folders to a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu >**Services** > **Objects**
2. Click the name of object store in which you want to create a bucket and launch the Objects Browser.
   For information on launching the Objects Browser, see Launching the Objects Browser on page 135.
   The object store opens in a new window.
3. Click the bucket name to which you want to upload files or folders.
4. (Optional) Click **Upload Objects** > **Select Files** to select and upload a file.
   The **Upload Objects** window appears.

   > Note:
   >
   > - This page displays the upload progress, the upload status, the bucket name, the object name, the size of the object, and the actions that you can perform on the object. You can also search by object name, bucket name or prefix.
   > - Uploading objects is an asynchronous process. If the upload size is large, you can close the **Upload Objects** window and perform other operations in the Object

Browser UI. You can check the upload status by clicking the status icon in the top-right corner of the page.

- Multi-part upload is used to upload large files (more than 1 GB). This feature enables to upload objects of any size to the object store. Depending upon the size of the file, upload may take a few minutes.
- If you refresh the page using the refresh button of the web browser, logout, or navigate to any non-application URL, any unsaved changes will be lost. This includes any pending or in-progress uploads. A warning message appears if you try to log out or refresh the page while an upload is in progress. Keep the Objects Browser open until the uploads complete.

Figure 88: Upload Objects Window

a. Click the **Cancel All Updates** to cancel the updates that are in progress.

b. Click the **Close** button or the **X** icon, to close the **Upload Objects** window.

5. (Optional) Click **Upload Objects** > **Select Folder** to select and upload a folder and the files in the folder.

   The **Upload Objects** window appears.

6. Click **Create Folder** to create a new folder in the bucket.

   You can create multiple folders to organize the objects in a bucket.

   Note:

   - **Create New Folder** feature is not supported in NFS enabled buckets.
   - The objects within the folders is listed in the Recycle Bin with the folder name. For example, if object1.txt is in folder BigData, then the deleted object listed in the Recycle Bin will be displayed as BigData/object1.txt.
   - When you create a folder, Objects Browser creates a 0 byte object using the folder name you specified followed by a slash (/).

7. (Optional) Click **Summary** view the bucket summary.

What to do next

After you upload objects to the bucket, you can add a tag to an object, download an object, view and manage the versions of an object, and also delete an object. For more information, see:

- Object CRUD Operations on page 148
- Adding Tags to an Object Using Objects Browser on page 152
- Understanding Object Versions on page 155
- Deleting an Object Using Objects Browser on page 153

Adding Tags to an Object Using Objects Browser

You can use the Objects Browser to add tags to an object in a bucket.

About this task

This topics describes how to add tags to an object.

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu >**Services** > **Objects**
2. Click the name of object store and launch the Objects Browser.
   For information on launching the Objects Browser, see Launching the Objects Browser on page 135.
   The object store opens in a new window.
3. Click the bucket name that contains the object.

4. To tag an object, select the object that you need to tag, and click **Actions** > **Tags**.
   The Add Tags window appears.

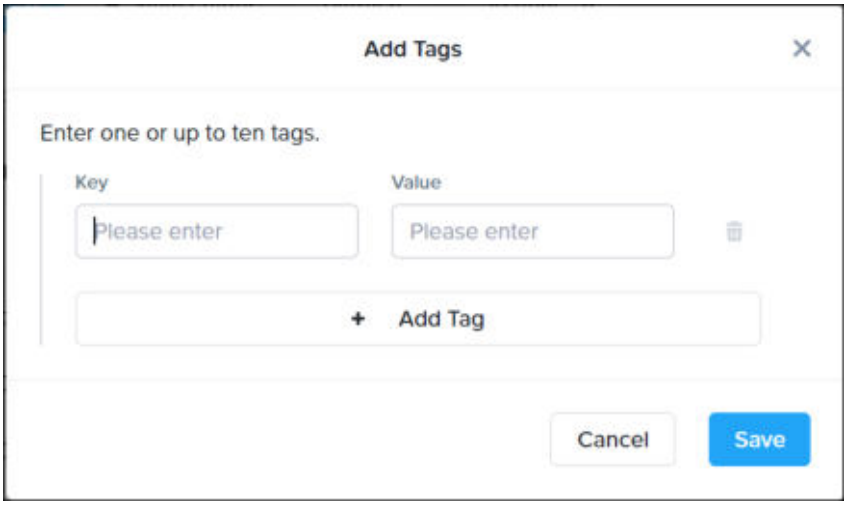

Figure 89: Adding Tags

   a. Enter the key and value pair.
   b. Click **Add Tag** to add more tags.

      Note: You can add up to 10 tags for an object.
      Click the delete icon against a key-value pair to delete it.

   c. Click **Save**.

What to do next

You can perform other object-related tasks, such as uploading objects, managing versions, and deleting an object. For more information, see:

- Object CRUD Operations on page 148
- Understanding Object Versions on page 155
- Deleting an Object Using Objects Browser on page 153

Deleting an Object Using Objects Browser

You can use the Objects Browser to delete objects in a bucket that you do not need.

About this task

To delete objects from a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu >**Services** > **Objects**

2. Click the name of object store and launch the Objects Browser.
   For information on launching the Objects Browser, see Launching the Objects Browser on page 135.
   The object store opens in a new window.
3. Click the bucket name that contains the object.
4. Select the objects that you need to delete, and click **Actions** > **Delete**.
5. Click **Confirm**, to confirm deletion of the objects.

   Note:

   - For non-versioned buckets, the delete operation performed on any object is permanent and cannot be reverted. The Recycle Bin is not available for non-versioned buckets.
   - For versioned buckets, the delete operation performed on any object is not permanent and a delete marker is created. The deleted objects can be viewed and recovered from the Recycle Bin.

6. Click **Recycle Bin**, select the objects to be deleted permanently, and then click **Delete Permanently**.

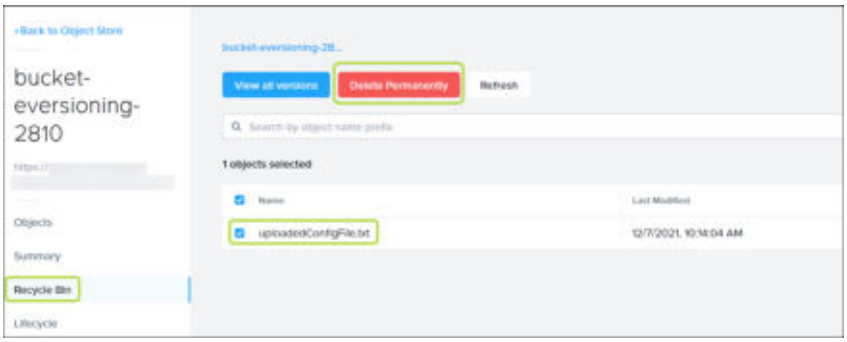

Figure 90: Delete Objects

Note:

- This action permanently deletes all versions of the selected objects. You cannot recover the deleted objects after you perform this operation.
- Any new version added to an object that is currently being deleted, gets automatically deleted.

What to do next

You can perform other object-related tasks, such as uploading objects and folders, managing versions, and adding tags to an object. For more information, see:

- Object CRUD Operations on page 148
- Understanding Object Versions on page 155
- Adding Tags to an Object Using Objects Browser on page 152

Understanding Object Versions

Versions allows you to view and manage the different versions of an object. From the **Versions** page, you can revert, delete, and download the different versions of an object. These options are available only for version-enabled buckets.

For more information about versions, see Object Versioning.

For more information on managing object versions, see Managing Object Versions Using Objects Browser on page 155.

When you delete a version-enabled object from the Bucket page, it is removed from the objects list and moved to the **Recycle Bin**.

To view the latest version of an object, click the object name, and then click **View all versions**. The latest version is listed with a **Delete Marker** banner in the **Recycle Bin**.

You can permanently delete any version by selecting the object version, and then by clicking the **Delete** button. You can also permanently delete the selected objects including all its versions from the **Recycle Bin**. Select the objects, and then click **Delete Permanently**.

Note:

- Selection of objects is limited to page-by-page basis. If you select all objects listed in the Recycle Bin, all the objects listed in the first page is deleted. For example, if the total objects in the Recycle Bin is 150, but the first page has 100 objects listed, then the 100 objects are deleted.
- For the selected objects, if any new version is added while the deletion is in progress, that version is also deleted.

To restore (revert) this object to any previous version, select the object, and then click **View all versions**. Select the version of the object, and then click **Revert**. Now, the reverted version is the latest version visible in the objects list. You can also select the latest delete marker and delete it to restore the object to last version.

Figure 91: Managing Object Versions - Using Objects Browser

Note: Versions with Delete Marker banner can only be deleted permanently. After deletion, they cannot be reverted or downloaded.

Managing Object Versions Using Objects Browser

You can use the Objects Browser to view and manage the versions of an object. Version allows you to revert an earlier version of an object as its latest version. You can also delete or download any version of the object.

About this task

To manage objects versions, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu >**Services** > **Objects**
2. Click the name of object store and launch the Objects Browser.
   For information on launching the Objects Browser, see Launching the Objects Browser on page 135.
   The object store opens in a new window.
3. Click the bucket name that contains the object.
4. Select the objects whose version you need to manage, and click **Actions** > **Versions**.

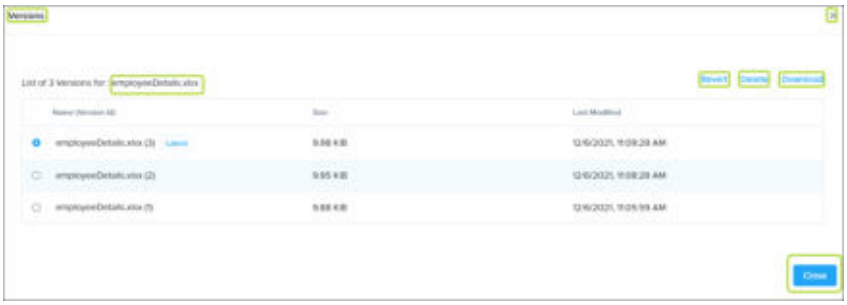

Figure 92: Objects Versions

Note: The latest version of an object is indicated by appending the banner `Latest` appended to the object name.

5. (Optional) Select a version that you need to revert as the latest version, and click **Revert**.

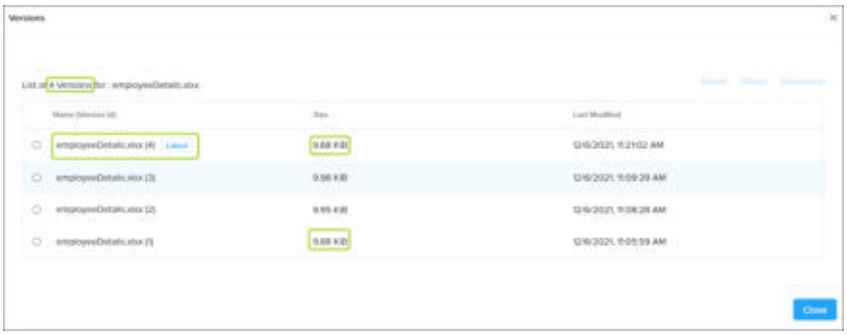

Figure 93: Revert Version

The above screen shot depicts an example, where `employeeDetails.xlsx (1)` is reverted to the latest version indicated by `employeeDetails.xlsx (4)` [Latest].

Note:

- Using the revert feature, you can designate any earlier version of an object as its latest version.
- If you deleted an object from the **Buckets** page, where all the objects within the bucket are listed. In the **Versions** page, you can select a version that does not have the Delete Marker banner, and then click **Revert**. The object gets restored in the **Bucket** page.

6. (Optional) Select a version of the object that you need to delete, and click **Delete**.
   This action permanently deletes the selected version of the object.
7. (Optional) Select a version of the object that you need to download, and click **Download**.
   You cannot download the object version that is marked with a `Delete Marker` banner. The object is downloaded to the default download folder of your browser.
8. Click **Close** or the **X** icon, to close the window.

Bucket Summary

The Summary page allows you to view the list of various policies applied to the bucket.

To view the Summary page of a bucket in an object store, click the name of the bucket in the buckets table, and then click **Summary**.

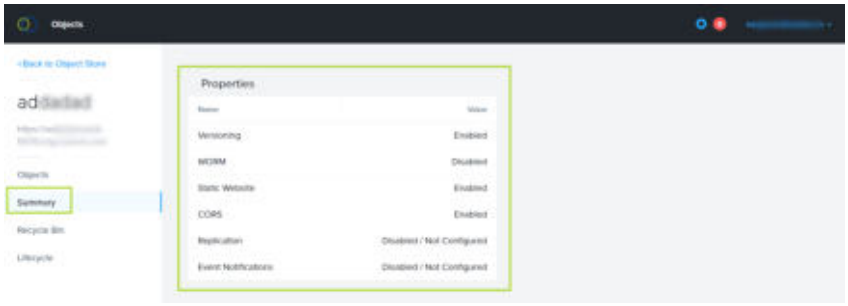

Figure 94: Objects Browser - Bucket Summary Page

OBJECTS STREAMING REPLICATION

Objects streaming replication enables automatic and asynchronous copying of source buckets to the target buckets in different Objects instances.

> Note: Replication is only supported for buckets created using S3 protocol.

Objects replication helps you with the following:

- Maintain another copy of data for disaster recovery.
- Maintain another copy of data at different sites for local access. This helps to minimize latency in accessing data.

Points To Remember

Note the following points about Objects replication.

- Replication can be set up between two buckets.

  You can replicate objects in a source bucket to a single destination bucket. For example, if you are replicating a Bucket A to Bucket B, then you cannot replicate Bucket A to Bucket C also.
- The source and target buckets must belong to different Objects instances.
- The Objects instances can be within the same Prism Central (PC) or different PCs.
- If the Objects instances belong to different PC clusters, a trust has to be set up between the PC clusters by adding each other as Availability Zones.

  Availability Zone - A physically isolated site where you can replicate the data that you want to protect. An instance of Prism Central represents an availability zone.
- For replication-enabled buckets, versioning and WORM modifications are prevented.
- Objects replication between buckets relies on the connection (for example, a secure VPN connection) for encrytion.
- Network Address Translation (NAT) performed by any device in between the two Objects instances is not currently supported.
- Proxy between Objects instances is not supported.

Types of Replication

There are two types of bucket replication.

Objects instances within the same PC

Replicating a bucket to an Objects instance within the same PC involves the single step of creating a replication rule.

Objects instances within different PC instances

Replicates buckets to an Objects instance within a different PC. This involves three steps.

- Step 1 - Adding the remote PC clusters as an availability zone.
- Step 2 - Performing IAM synchronization between the source PC cluster and remote PC clusters.
- Step 3 - Creating a replication rule.

After you set up the replication rule between the source and destination bucket, you can track the replication statistics for each bucket. For more information, see Viewing Replication Statistics for a Bucket.

Replication Guarantees and Topologies

Refer to this section to understand the replication guarantees, how replication works, and supported replication topologies.

Replication Guarantees

The following object attributes get replicated:

- Object operations - Object PUT, Object Copy, Updates (PutTags, PutObjectLock), and Deletes.
- Object metadata - ETag, create and modification time, and lock property.
- Version numbers, User metadata, and Tags.

> Note: Source and destination buckets are independently managed and not replicated. Also, any changes to the bucket policies (for example, access- or lifecycle-policies) does not get replicated.

How Replication Works

Note the following points about how the Objects replication works:

- Replication starts as soon as the object gets written on the source bucket.
- The replication completion time may vary depending on the object size and other factors such as available bandwidth, number of replications, and so on.
- Objects replications may not strictly follow the same time order in which they get written on the source bucket.
- Replication of the versions of an object may not happen in sequential order, but they get replicated eventually.

Replication Topology

> Note: Note the following points that are applicable to all the topologies:
>
> - Ensure that applications do not perform conflicting write operations (objects with the same name) on the remote bucket while replication is enabled.
> - There are no restrictions on the I/O operations performed on the remote bucket. You can perform read and write operations on a remote bucket.
> - A bucket can have a maximum of five inbound relationships. For example, Bucket A can be the destination bucket for a maximum of five source buckets.

The following are the topologies that you can use for your replication scenarios:

Single-Replication Relation

In this topology, you replicate objects one way from the source to the destination.

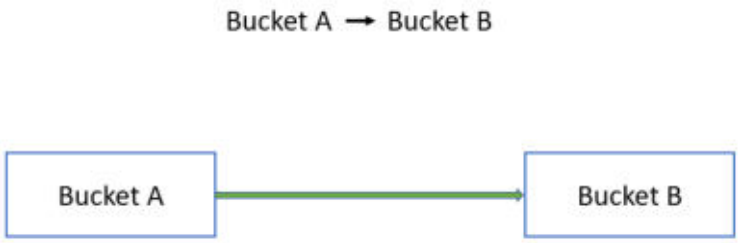

Figure 95: Single Replication - Unidirectional

Different buckets on the source Objects instance can replicate to buckets belonging to different Objects instances.

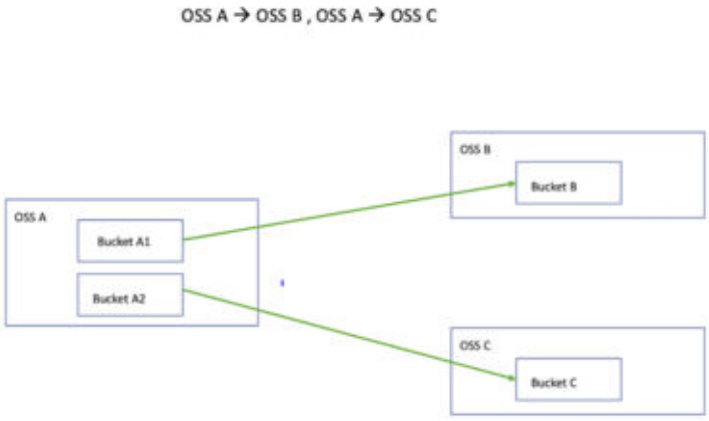

Figure 96: Single Replication - Different Objects Instances

Bidirectional Replication Relation

In this topology, you can set up a bidirectional-replication relation between a pair of buckets. Objects written on one side get replicated to the other side.

Independent replication relations betweens the buckets have to be created, that is, create a replication rule with Bucket A as the source and Bucket B as the destination and the other way around.

> Note: Ensure that the application does not write conflicting objects (objects with the same name) on both buckets while replication is enabled.

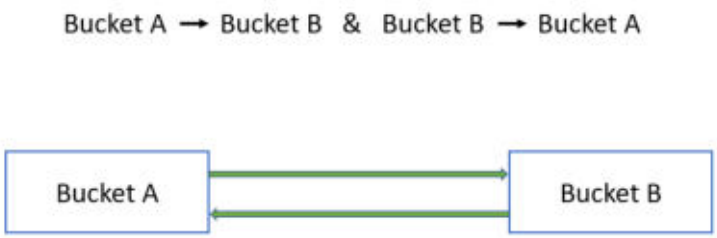

Figure 97: Bidirectional Replication

Chain Replication Relation

In this topology, objects from Bucket A can be replicated to Bucket B, and objects that originated in Bucket B can be replicated to Bucket C.

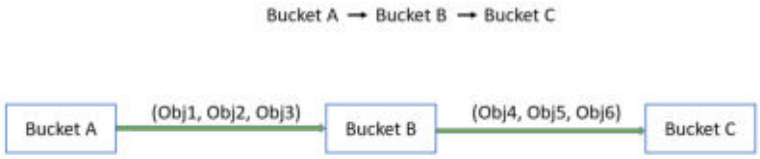

Figure 98: Chain Replication

N:1 Replication Relation

In this topology, a bucket can be a destination to many source buckets.

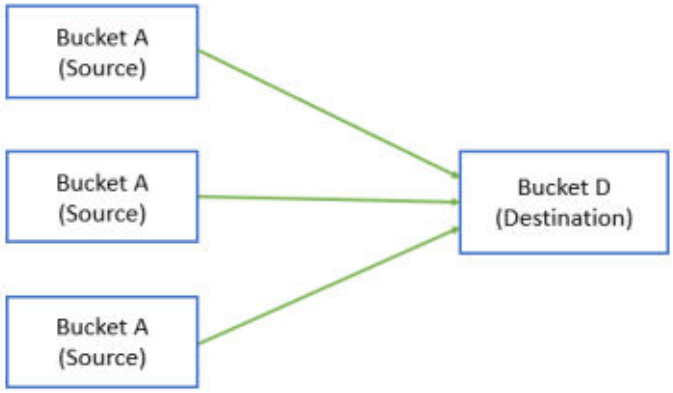

Figure 99: N:1 Replication

Replication Prerequisites

The prerequisites for replicating buckets are as follows:

- Object instances that contain the source and destination buckets must be deployed with Objects 3.0.
- For Objects instances within different PC instances replication, ensure to deploy Objects instances with IAM users if you want the same users to have access to the replicated objects.
- For Objects instances within different PC instances replication, you should establish a secure VPN connection between the PC clusters. For Objects instances within the same PC replication, you should establish a secure VPN connection between the clusters (PE).
- Ensure that the traffic to port 7100 on the load balancers of the Objects instances is allowed by the firewall for replication purposes.

Adding Remote Prism Central as Availability Zone

For ***Objects instances within different PC instances*** replication, you need to add the remote PC clusters as an Availability Zone (AZ) on the source PC cluster to establish a secure connection.

About this task

To add the remote PC as an availability zone, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Administration** > **Availability Zones**.

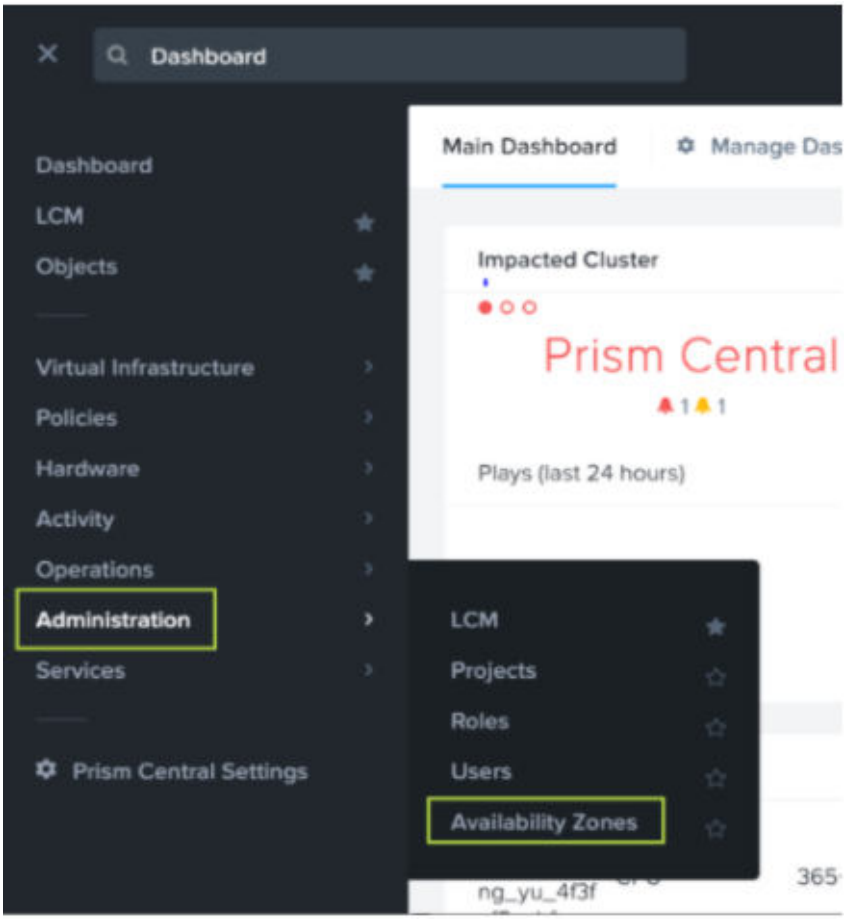

Figure 100: Administration - Availability Zones Page

2. Click **Connect to Availability Zone**.

3. In the **Availability Zone Type** list, click **Physical Location**.

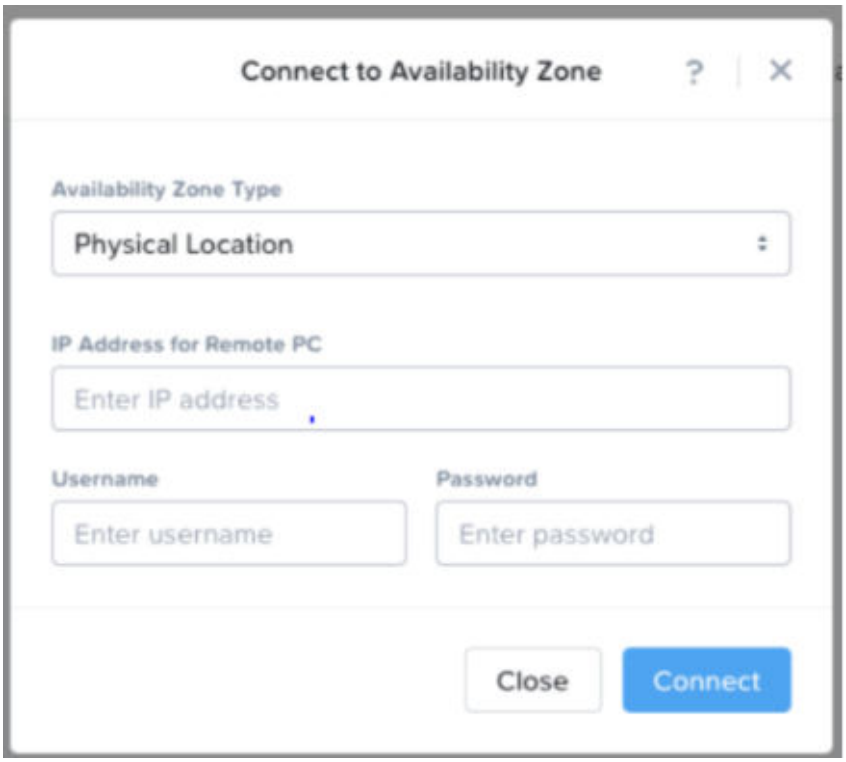

Figure 101: Add an Availability Zone

4. Enter the IP address and login credentials of the remote PC in the corresponding boxes.
5. Click **Connect**.
   The remote PC gets added as an AZ in the source PC and the connectivity status is shown as **Reachable**.

Setting up IAM Synchronization with a Different PC

To replicate Objects on a different Prism Central (PC), you must perform IAM synchronization of the source PC with the corresponding remote PC. The IAM users of the source Objects instance get replicated to the destination Objects instance belonging to a different PC with the same access key and secret key pair. The admin or bucket owner needs to provide permission to the users for the replicated buckets.

Before you begin

A non-admin user can set up IAM synchronization with a different Prism Central only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- Create Object Store
- View Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

To perform the IAM pairing, do the following:

Procedure

1. Log on to the source Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. In the **Objects** page, click the **Access Keys** tab.
3. Click **IAM Replication Settings**, and then click **Add IAM Pairing**.
4. In the **Target Prism Central** list, select the remote PC that you added as an Availability Zone (AZ) in the source PC.

   Note:

   - Only the PCs added in **Administration** > **Availability Zones** appear in this list.
   - The target Prism Central IAM must have all the Active Directory configured in this source Prism Central IAM.

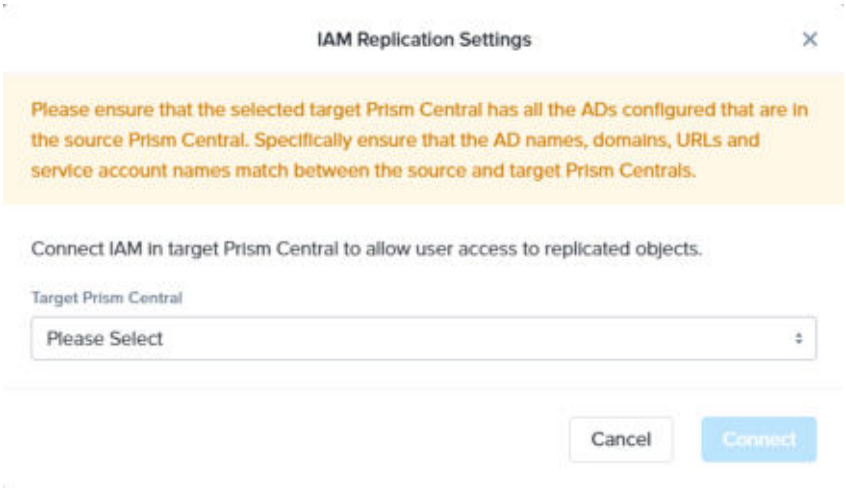

Figure 102: Add IAM Pairing

5. Click **Connect** to complete the pairing.

   Note: You cannot add more than five IAM replication targets.

   After you click **Connect**, all existing users get replicated to the target Prism Central IAM. You can monitor the progress of the replication in the IAM Replication Settings page. Once the IAM pairing is complete, all the future users and keys that are created and deleted get replicated to the target Prism Central IAM.

   If any of the replication fails, go to the IAM Replication Settings page and click **Sync** to replicate any unreplicated users to the target Prism Central. The administrators can view

the users for whom the replication failed and use the export option to download the list of errors.

> Note: The **Sync** option will not replicate a user or a key deletion. For a failed replication for delete operations, you must log in to the target Prism Central and delete the unwanted users and keys.

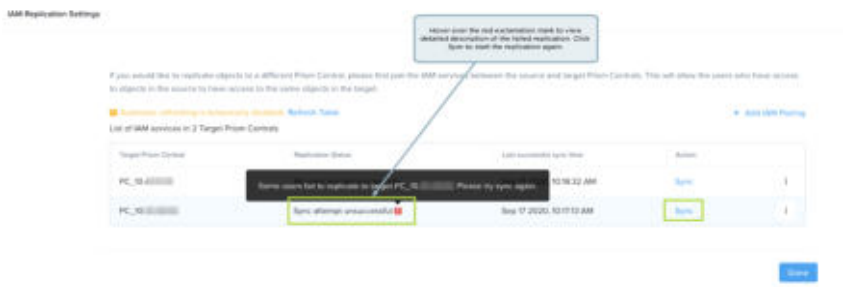

Figure 103: IAM Replication - Sync Option

Creating Replication Rules between Buckets

This section describes the steps to create a replication rule between buckets. Creating a replication rule involves selecting the replication type (same PC or remote PC) and defining the details of the destination bucket.

Before you begin

Make note of the following points.

- In the case of remote-PC cluster replication, ensure that the Fully Qualified Domain Name (FQDN) of the Object Store instances in two different PC clusters is different.
- If you plan to replicate data to a bucket on a different PC, ensure that you perform IAM synchronization of the source PC with the corresponding remote PC.
- Bucket access permissions are not automatically replicated to the destination bucket. You need to manually assign them to the destination bucket.
- Lifecycle policies are not replicated and can be assigned independently on the source and the destination buckets.
- If you are replicating a bucket created in Object v3.1 (with the legal hold applied) to a bucket created in Objects v3.0, the legal hold attributes will not get replicated.
- Versioning, WORM state, and WORM retention period must be the same between the source and destination buckets.
- Only objects added to the buckets or modified after the replication rule has been created will be transferred to the destination bucket.
- For replication-enabled buckets, versioning and WORM modifications are prevented. To make changes, you need to delete the replication rule, perform the required edits on the buckets, and again create the replication rule.
- When creating a replication rule for buckets in different Object Store instances, ensure that the first relationship to the Object store containing the destination bucket is created using the Prism GUI. For example, create a rule for Bucket 2 in Object Store A with Bucket 6 in Object Store B as the destination using the GUI. For successive replication rules for any

buckets in Object Store A to any other bucket in Object Store B, you can use the GUI or S3 API.

A non-admin user can create replication rules between buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with Edit Buckets permission can configure bucket replication on all the buckets on an assigned object store.

If the target Object store is on same Prism Central, then the user must have View Object Store access to the target Object store too. If the target Object store is on a different Prism Central, then the user will be able to select any Object store on the target Prism Central.

About this task

To create a replication rule between buckets, do the following:

Procedure

1. Log on to the source Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store.
   The object store opens on a new window. You can view the buckets in your Objects instance.
3. In the **Buckets** table, select your source bucket.

4. In the **Actions** list, click **Create Replication Rule**.

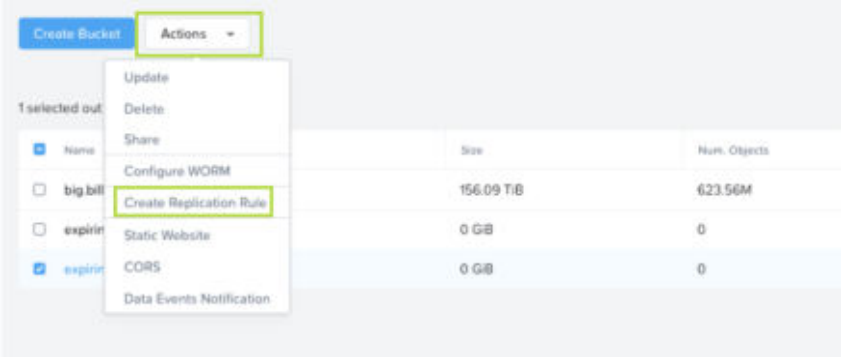

Figure 104: Create Replication Rule

The Create Replication Rule page appears.

5. If you want to replicate your bucket on a remote PC cluster, select the remote PC from the **Prism Central** list. All the remote PC clusters that you added as AZ on your source PC cluster appears in this list.

   Or

   If you want to replicate the source bucket on the same PC cluster, select **Local AZ** in the **Prism Central** list.

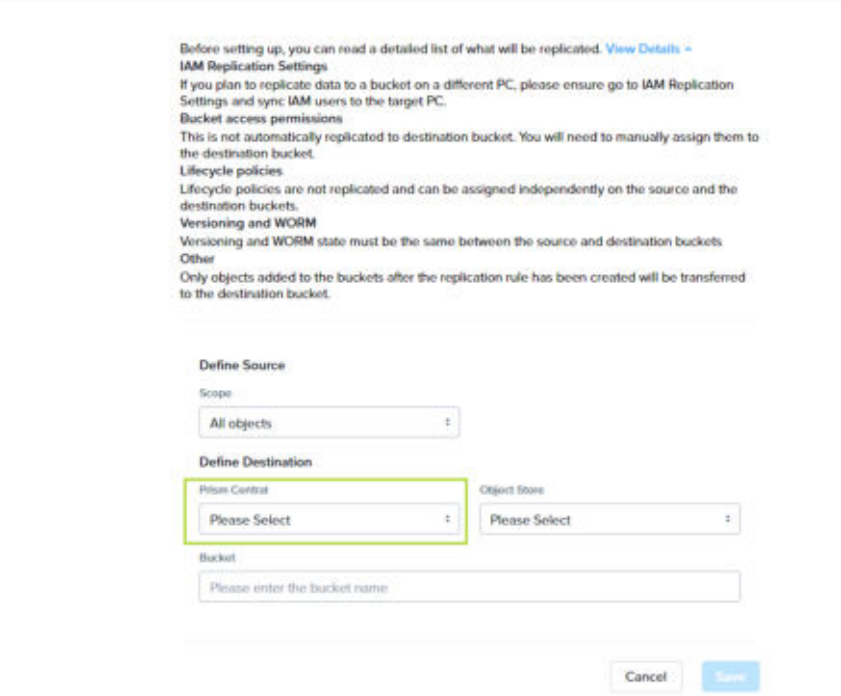

Figure 105: Create Replication Rule Page

6. In the **Object Store** list, select the Objects instance that contains the destination bucket.
7. In the **Bucket** box, enter the name of the destination bucket.
8. Click **Save** to complete.
   The replication rule between the buckets gets established. The replication time may vary depending upon various factors such as object size, workloads on the objects cluster, and so on.

   For more information about replication alerts, see Nutanix Objects Specific Alerts on page 185.

   Points to Note (Disconnected Availability Zone)

   Note the following points if you remove a Prism Central from the Availability Zones after setting up the IAM synchronization and bucket replication:

   - Successive IAM user additions and deletions will not get replicated.
   - Existing bucket replication will not be affected. But, you would not be able to create new replication rules from the GUI to the Objects Store deployed in the disconnected availability zone.

Deleting a Replication Relation

This section provides the steps to delete a replication relation.

About this task

A typical scenario where you may require to delete a replication rule is if you want to modify the versioning and WORM state of a replication-enabled bucket. Since versioning and WORM changes are prevented, you need to delete the replication rule, perform the required edits on the bucket, and again create the replication rule.

> Note: Deleting a replication relation would cause all pending replications for that relationship to be dropped immediately. It is recommended to wait for the pending replications to complete, then start making changes by deleting the relation.

To delete a replication rule for a bucket, do the following:

Procedure

1. Log on to the source Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store where the corresponding bucket is deployed.
   The object store opens on a new window. You can view the list of buckets in your Objects instance.
3. In the **Buckets** table, select the replication-enabled bucket.

4. In the **Actions** list, click **Delete Replication Rule**.

Figure 106: Delete Replication Rule

A warning message appears.

5. Click **Delete**.
   The replication is disabled for the bucket. No new objects will be replicated to the destination bucket.

Viewing Replication Statistics for a Bucket

This section describes the steps to view the replication statistics for a bucket.

Before you begin

A non-admin user can view replication statistics only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.
- Users with View Buckets permission can view the replication statistics on all the buckets on an assigned object store.

About this task

To view the replication statistics, do the following:

Procedure

1. Log on to the source Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.

2. Click the object store in which your bucket is created.

   The object store opens on a new window. You can view the buckets in your Objects instance.

3. Click on the bucket name.

Figure 107: Object Store - List of Buckets

4. Click **Replication** on the left pane to view the replication statistics.

   If the bucket has both inbound and outbound traffics, you can select either of the traffic types from the **Traffic Type** list, and then select the corresponding source or destination bucket.

   - **Outbound statistics** - Shows the statistics for replication of objects from the selected bucket (source) to the destination bucket.
   - **Inbound statistics** - The selected bucket can be a destination to many source buckets. Inbound statistics show the data for the replication of objects to the selected bucket. You can select the source bucket and view the inbound statistics from that bucket.

     Note: Statistics on the destination bucket can lag compare to the statistics on the source bucket as the update happens periodically.

   - **Last Replication Point** - Point-in-time up to which all the objects created on the source bucket have been replicated.
   - **Number of Objects pending replication** - Objects count pending to get replicated.
   - **Objects size pending replication** - Total amount of data pending to get replicated.
   - **Average Bandwidth** - The rate of the amount of data transferred from the source to destination. For inbound relationships, the average bandwidth is the cumulative value of all the incoming data from the source buckets. The bandwidth graph helps you to visualize the progress of the replication.

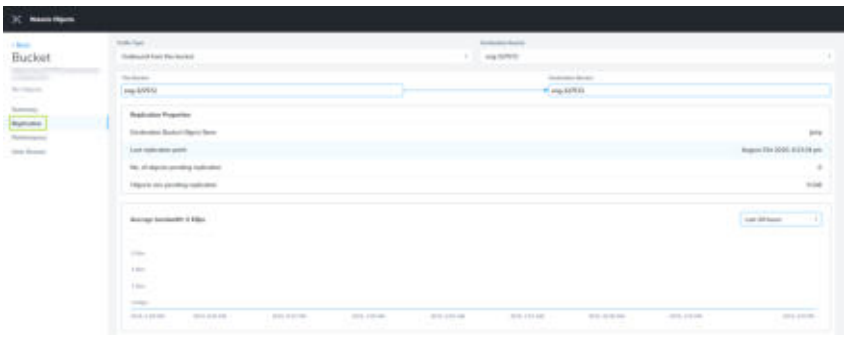

Figure 108: View Replication Statistics

Achieving Fault Tolerance for IAM

This section describes how you can use IAM replication and bucket replication to withstand data center failures.

IAM Service on Objects

At present, all the Object Stores (Objects instance) in the same PC share a single IAM. This IAM instance resides in the first object store (primary) that is deployed in the PC. If that primary object store becomes unavailable, the IAM service will not be available for all secondary object stores of that PC.

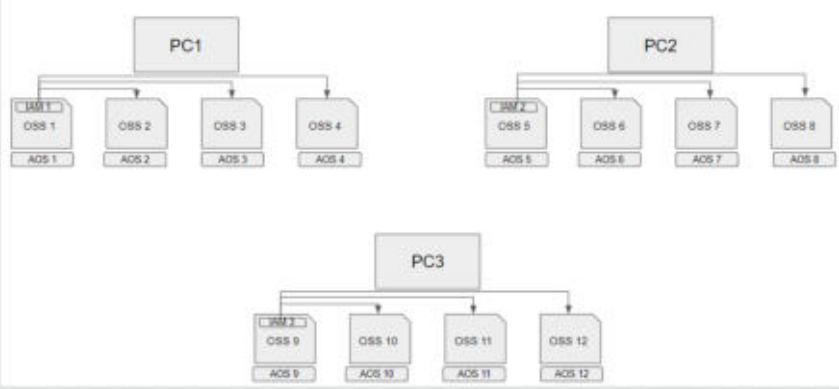

Figure 109: IAM Service on Objects

For example, you have one PC managing all your clusters that is PC1. This PC1 has four object stores (OSS 1, 2, 3 and 4). However, IAM 1 resides only in the first object store OSS 1 (primary object store). All other secondary object stores on that PC (OSS 2, 3 and 4) are for back up of OSS 1 and relies on IAM 1 for authentication. If OSS 1 goes down, then all other back ups also goes down. In this case, Nutanix recommends you the following IAM replication configuration.

> Note: The Objects cluster will work even when the PC fails.

Recommended IAM Replication Configurations

Nutanix recommends to transition to multiple PC, that is multiple IAM solution for achieving fault tolerance for IAM replication. You can have one or more PCs to manage IAM replication. The remote PC can have different object stores, and you can set up IAM replication between primary PC and remote PC.

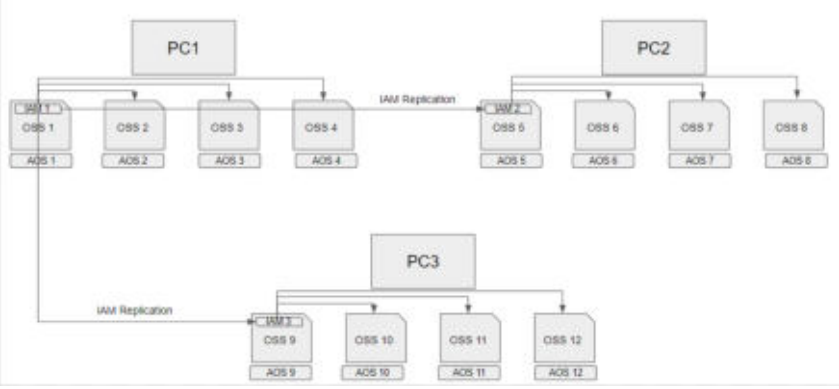

Figure 110: Objects IAM Service with Replication across Prism Centrals

For example, in the preceding image, you have three PCs managing all your clusters that is PC1, PC2, and PC3. PC1 has four object stores (OSS 1, 2, 3 and 4) with IAM 1 residing only in the first object store OSS 1 (primary object store). All other secondary object stores on that PC (OSS 2, 3 and 4) relies on IAM 1 for authentication.

Similary, PC2 and PC3 have different object stores respectively (OSS 5, 6, 7, 8 and OSS 9, 10, 11, 12) with IAM 2 and IAM 3 residing only in the first object stores OSS 5 and OSS 9 (primary object stores) of PC2 and PC3.

Now as recommended, you use PC2 and PC3 for replication of PC1. In this case, even if any object store in PC1 goes down, then replicated buckets will be available in PC2 and PC3 without any disruption.

BASELINE REPLICATOR TOOL

Baseline Replicator tool is a tool with which you can replicate existing objects in your bucket to a bucket at a remote site. You can use this tool only after a replication relationship is created with the destination bucket. You need to first identify the destination bucket, and then run this tool against the source bucket.

Baseline Replicator tool performs the following operations:

- Lists objects in the source bucket.
- Touches meta-data of each object on the bucket so that it gets setup for native replication.

> Note:
>
> - This tool does not replicate the objects to the destination bucket. Instead, it sets up the objects for replication and expects native replication protocol to replicate it over to the destination bucket.
> - You cannot replicate an object from the source bucket if an object with the same name already exists on the destination bucket.
> - Baseline replicator tool can be used to replicate upto 100 million objects in a single bucket.
> - When baseline replication is setup between two versioned buckets, all the versioned objects will be replicated over to the destination bucket. However, the delete markers from the source bucket are not replicated over to the destination bucket.

Accessing and Running the Baseline Replicator Tool

Prism Central admin users can access this tool inside the Service Manager pod (aoss_service_manager) in the Prism Central. Copy the tool from the Service Manager pod to a temporary folder on the Prism Central VM or the Linux VM. You can also download the Baseline Replicator tool from the Nutanix Support Portal in your Linux VM or the Prism Central VM.

Before you begin

Ensure that you conform to the following requirements before running this tool:

- Your machine should have connectivity with the object store (endpoint URL) hosting the source bucket.
- You have a valid secret key and access key.
- User must have read-write access to the bucket.

About this task

To run the Baseline Replicator tool, do the following:

Procedure

1. SSH into the Prism Central VM or the Linux VM where the tool is present.
2. Ensure the tool can be run.

```
ls -l path/to/the/tool
```

3. (Optional) To see the tools usage, run `nutanix@pcvm$ /tmp/baseline_replicator`.

   The Baseline Replicator tool has the following parameters:

   - source_bucket_name: Name of the source bucket.
   - source_endpoint_url: Endpoint URL of the object store hosting the bucket.
   - source_access_key: Access key of the source bucket.
   - source_secret_key: Secret key of the source bucket.
   - region_name: Specify the region name.
   - keys_prefix: Choose which objects to replicate by giving a prefix. For example, `ntnx` is the keys prefix, then only the objects with the prefix `ntnx` will be scheduled for replication.
   - start_key: If any interruption occurs while replicating, provide the start_key to resume the replication from that object.
   - log_path: Provide the path if you want to save the log in a specific location.
   - max_concurrency: Specify how many concurrent operations you want the tool to perform. Default concurrency is 50. You can increase the number based on the resources of the VM where the tool is running.
   - skip_prelim_test: When the tool is run, a preliminary test is run. If you want to skip the prechecks, use this --skip_prelim_test parameter.
   - log_level: Specify the levels of the logs you want to display, such as Debug, Info, Warn, Error, Fatal in order of decreasing verbosity.
   - resume: Use this option if the tool was stopped abruptly. You can run the tool again with --resume, and the tool takes care of identifying the last object that was setup for replication and provides that as the start_key.

   The following output appears with the usage details.

```
nutanix@NTNX-10-44-195-21-A-PCVM:~$ /tmp/baseline_replicator
usage: baseline_replicator [-h] --source_bucket_name SOURCE_BUCKET
                           --source_endpoint_url SOURCE_ENDPOINT_URL
                           --source_access_key <SOURCE_ACCESS_KEY>
                           --source_secret_key <SOURCE_SECRET_KEY>
                           [--region_name <REGION_NAME>]
                           [--keys_prefix <KEY_PREFIX>]
                           [--start_key <START_KEY>] [--log_path <LogFile>]
                           [--max_concurrency <Concurrency>]
                           [--skip_prelim_test [SKIP_PRELIM_TEST]]
                           [--log_level <DEBUG>] [--resume [RESUME]]
baseline_replicator: error: argument --source_bucket_name is required
```

   Figure 111: Baseline Replicator Usage

   Here `tmp` is the folder name where the tool is copied.

4. Copy the tool from the Service Manager pod and paste it in any folder, or download the tool from Nutanix Support Portal and paste it to any folder.

```
nutanix@pcvm$ docker cp aoss_service-manager:/svc/bin/baseline_replicator /tmp/
```

5. Run the Baseline Replicator tool.

```
nutanix@pcvm$ /tmp/baseline_replicator --source_endpoint_url=source-endpoint-
url --source_bucket_name=bucket-name --source_access_key=source-access-key --
source_secret_key=source-secret-key --max_concurrency=concurrency
```

For example,

```
nutanix@pcvm$ /tmp/baseline_replicator --source_endpoint_url=http://10.45.28.101 --
source_bucket_name=1million --source_access_key=v0cfi3kRMBv0cfi3kRMBv0cfi3kRMBz --
source_secret_key=NTY7_Xok_UEYOvz-Y9zdkVs_koV-YrMa --max_concurrency=200
```

If you lose connectivity or any interruption occurs, all objects may not get tagged for replication. In that case, you can check the logs to find the last object that was successfully setup for replication. You can then restart the replication by providing the start key parameter (start_key) to resume the replication from that object. Alternatively, you can restart the script by running the same command but with an additional --resume parameter.

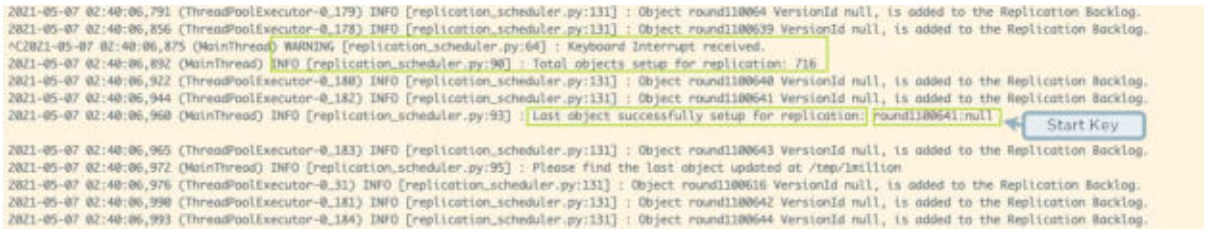

Figure 112: Checking Logs for Start Key

For example,

```
nutanix@pcvm$ /tmp/baseline_replicator --source_endpoint_url=http://10.45.28.101 --
source_bucket_name=1million --source_access_key=v0cfi3kRMBv0cfi3kRMBv0cfi3kRMBz --
source_secret_key=NTY7_Xok_UEYOvz-Y9zdkVs_koV-YrMa --start_key=round1100641:null --
max_concurrency=200 --resume
```

You can check the logs to find the last object that was set up for replication. Go to `/tmp/bucket-name` to find the logs.

MONITORING AND ALERTS

You can monitor the usage and the performance of object stores and buckets, and you can also view system generated alerts for Objects.

A non-admin user can view the usage and performance of object stores and buckets, and also view system generated alerts for objects only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have view object store permission and the non-admin user must be assigned to specific objects instances.

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

> Note: For the object counts that are shown at the object store instance level, Nutanix Objects counts each upload of a multipart upload as a separate object until the object is finalized. However, for the object count at the bucket level, Nutanix Objects does not include the upload counts in the objects count since an object is considered as uploaded in a bucket only after it is finalized. For example, if you have done a multipart upload of 10 objects (suppose 5 parts of each object), then the multipart count of the objects is shown as 50 in the object store instance level. However, at the bucket level, it is shown as 0 because the objects are not finalized or completely uploaded.

Viewing Performance of Object Stores

You can view and analyze the performance of an object store by placing the cursor anywhere on the horizontal axis to display the value at that time. You can also select the time interval (last 3 hours, last 6 hours, and last 12 hours) from the pull-down list above the graphs.

About this task

To view performance of an object store, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.

2. Click the name of the object store, and then click **Performance**.

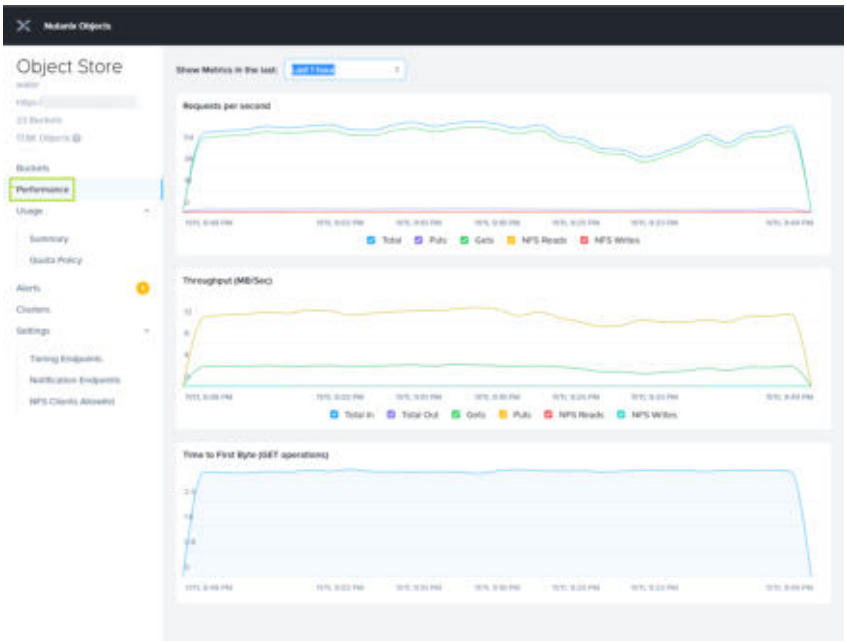

Figure 113: Object Store Performance Graph

The graph shows the following information:

a. **Requests Per Second**: Displays the following graphs:

- The **Total** graph displays the total input and output requests in each second in the object store or bucket.
- The **Puts** graph displays the total input in each second
- The **Gets** graph displays the total output in each second.
- The **NFS Reads** graph displays the total NFS reads in each second.
- The **NFS Writes** graph displays the total NFS writes in each second.

b. **Throughput (MB per sec)**: The Throughput graph displays granular read and write throughput in MB in each second.

You can see the total in, total out, gets, puts, NFS reads, and NFS writes in a bucket.

c. **Time to First Byte (GET Operations)**: This graph displays the time taken to read the first byte from the object in milliseconds.

Viewing Performance of Buckets

You can view and analyze the performance of a bucket by placing the cursor anywhere on the horizontal axis to display the value at that time. You can also select the time interval (last 3 hours, last 6 hours, and last 12 hours) from the pull-down list above the graphs.

About this task

Note: The Throughput chart shows data for each connection. This data is not the cumulative throughput across all the clients and connections.

To view performance of a bucket, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the Object Store, which contains the bucket.
   The object store opens in a new window.
3. Click **Bucket** in the left pane, select the name of the bucket, and then click **Performance**.

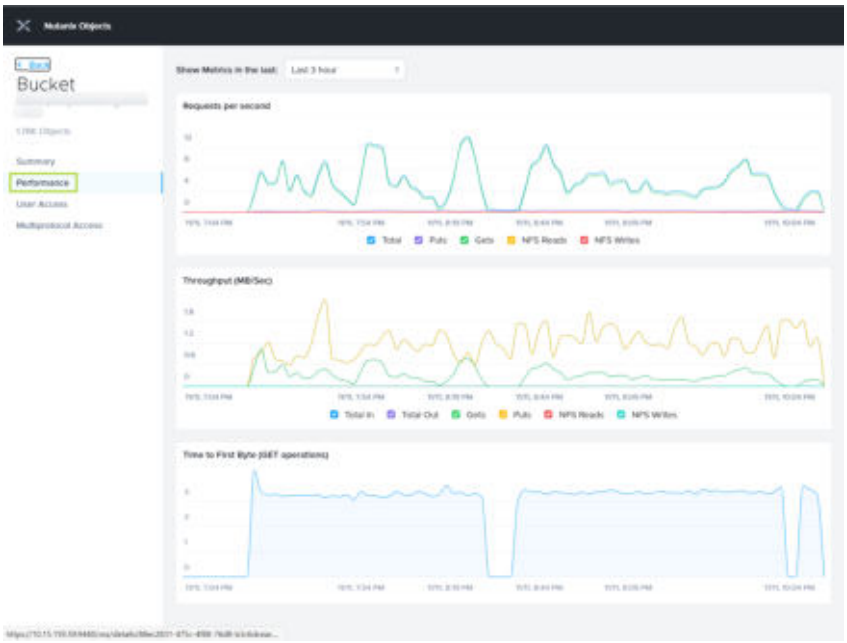

Figure 114: Bucket Performance Graph

The graph displays the performance of the bucket.

a. **Requests Per Second**: Displays the following graphs:

   - The **Total** graph displays the total input and output requests in each second in the object store or bucket.
   - The **Puts** graph displays the total input in each second
   - The **Gets** graph displays the total output in each second.
   - The **NFS Reads** graph displays the total NFS reads in each second.
   - The **NFS Writes** graph displays the total NFS writes in each second.

b. **Throughput (MB per sec)**: This graph displays granular read and write throughput in MB per second.

   You can see the total in, total out, gets, puts, NFS reads, and NFS writes in a bucket.

c. **Time to First Byte (GET Operations)**: This graph displays the time taken to read the first byte from the object in milliseconds.

Viewing Object Store Usage

The Usage tab displays the physical and logical storage usage. You can view the space used across all buckets in an object store.

About this task

To view the object store usage, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store to see its usage.
   The object store opens on a new window.
3. Click **Usage**.

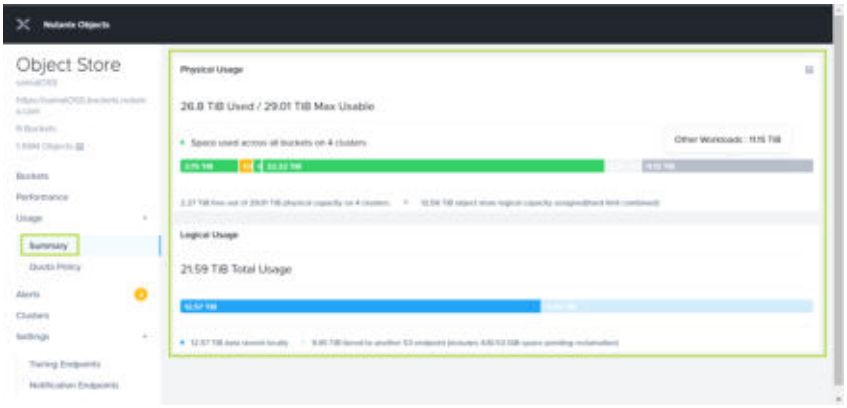

Figure 115: Object Store Usage

The storage usage section displays the following information:

Physical Usage: Displays the physical capacity used by the object store on the cluster.

- Total used and total usable storage for the cluster.
- Space used across all buckets on the existing cluster displayed in a bar.
- Total and free physical capacity of the cluster.
- Logical capacity assigned for the object store.

Logical Usage: Displays the logical capacity used by the object store on the cluster. As the Redundancy Factor is 2 (RF-2), the data is stored in two backups.

- Data stored locally
- Data tiered to another S3 endpoint

Assigning Quota Policy to a User

Assigning a quota policy to a user enables Objects to set soft thresholds on storage limits of all buckets owned by the user and number of buckets created by the user within an object store. You can assign quota policy to multiple users at the same time. If any user exceeds the limit, they can still create buckets and objects, however an alert is raised in the Alerts page. For example, if a user has 20 TB storage and has six buckets, then storage for all the six buckets will

be accounted. If you are an owner of a bucket and have shared your bucket with multiple users, then your storage quota will be used up even if other users use the shared bucket storage. If a user is deleted and recreated in the object store, the previous storage usage and quota policies will not apply to the user and will be treated as a new user. If some multi-part operations are pending, those pending parts do not count towards the quota.

Before you begin

A non-admin user can assign quota policy to a user only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- Create Quota Policy
- Delete Quota Policy
- View Quota Policy
- View Object Store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

> Note:
>
> You cannot list the buckets only owned by a particular user. So, if you have exceeded your storage quota and try to delete a bucket shared with you to release the storage space, the quota level remains as is. You have to delete your own bucket to release the storage space.
>
> For example, if user A has exceeded the storage quota and tries to list all the buckets, the bucket list will also contain the buckets that were shared with user A by user B. To release the storage space within user A quotas, if user A deletes objects in the buckets shared with user A by user B, the quota level remains as is.
>
> You can run `get-bucket-acl` command on the bucket to list the user owned buckets, and then delete the data from that bucket.

To assign a quota policy to a user, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. To see the object store usage and to assign quota policy to the users, click the name of the object store.
   The object store opens on a new window.
3. Click **Usage** > **Quota Policy**.
4. Select the user for which you want to create the quota policy.

5. Click **Create Policy**.

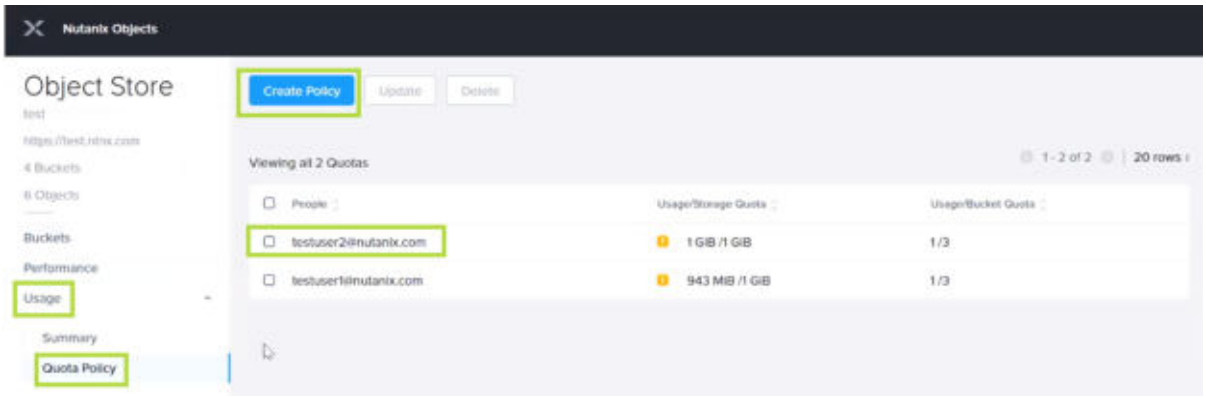

Figure 116: Create Quota Policy

A Configure Quota window appears.

6. Select a user or multiple users from the **People** drop-down.

> Note: You can create quote policy for multiple users simultaneously, but you cannot assign multiple quotas to the same user.

7. Click the **Storage Limit** check box and enter the soft limit for storage usage.
8. Click the **Bucket Limit** check box and enter the soft limit for bucket usage, and then click **Save**.
   A list of users with the configured storage- and bucket-quota, and the current usage is displayed.

   When the storage usage or bucket usage reaches 90% of the defined limit, you are notified with an exclamation mark.

   A background task checks for any violation of quota policies and generates this alert. So, it can take about an hour to generate the alert. On a heavily loaded object store, it can take up to a day to generate an alert.

   Also, if you delete a user from the object store, then the quota policy of the user is not deleted and can still generate alerts.

What to do next

You can view the alerts. For more information, refer to Viewing Alerts on page 183 section.

Viewing Buckets Usage

The Usage tab displays the storage usage in a graphical form. You can view the space used across the bucket.

About this task

To view the buckets usage, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.

2. Click the name of the object store, which stores the bucket, and then click **Buckets**
3. Click the name of the bucket.

Figure 117: Buckets List

Note: Only finalized objects of a multipart upload are shown in the buckets table.

4. Click **Usage**.

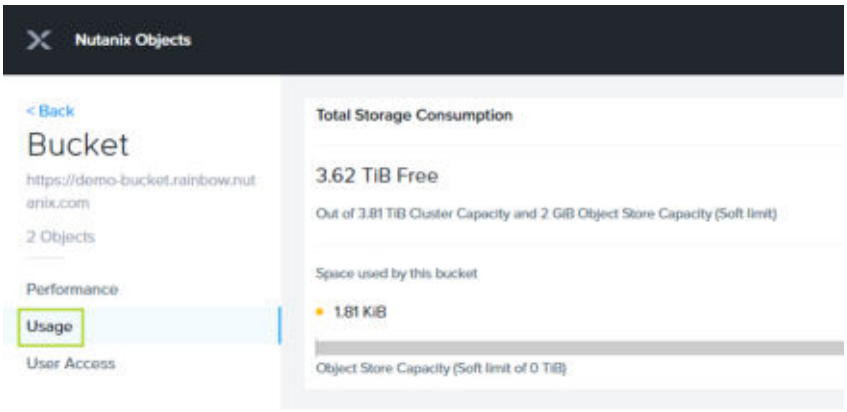

Figure 118: Bucket Usage

The storage usage section displays the following information:

- Space used across all buckets.
- Total physical capacity of the cluster.
- Logical capacity of the object store.

Note: For NFS enabled buckets with sparse files, the sparseness also contributes to the usage (logical).

Viewing Alerts

The Alerts tab apprise administrators of informational, warning, and critical alerts. The Alerts tab presents all notifications and events in an easy-to-consume table format. The table presents each alert with the color-coded severity level, a description of the alert, and the time-stamp.

About this task

To view the object store alerts, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store.
   The object store opens in a new window.
3. Click **Alerts**.

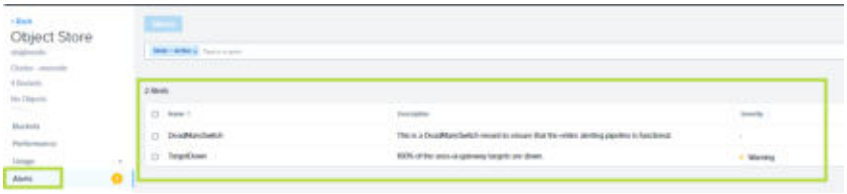

Figure 119: Alerts List

The alerts table displays the following information:

Table 8: Alerts Table

Field	Description
Name	This column displays the name of the alert.
Description	This column displays a description of the alert issue.
Severity	This column displays the severity level of the alert. Following are the three severity levels: • **Critical**: A critical alert is one that requires immediate attention. • **Warning**: A warning alert is one that might need attention soon. • **Info**: An info alert highlights a condition to be aware of.
Created	This column displays the number of days since the alert was created. Hover over the number of days to view the date and time when the alert was created.
Silenced	This column provides the option to suppress the alert.

You can also suppress multiple alerts by using **Silence** at the top-left of the screen.

> Note: Silencing of alerts is temporary.

You can use **Filters** at the top-right corner to filter the alerts you want to view. You can filter the alerts by severity levels (**Critical**, **Warning**, and **Info**) and state (**Active** and **Silenced**).

Nutanix Objects Specific Alerts

The following tables describe the Objects specific alerts. Apart from the Objects specific alerts, alerts that monitor the underlying infrastructure of the Objects clusters are also triggered. For more information about these alerts, you can contact Nutanix Support at https://portal.nutanix.com.

Table 9: Objects Alert - HighTimeToFirstByte

Name	`HighTimeToFirstByte`
Description	The `HighTimeToFirstByte` alert appears when the Time To First Byte (TTFB) for all the HTTP GET operations in the past 10 minutes exceeds 1 second.
Alert Message	`Get operations issued to the object store in the past 10 minutes have been showing TTFB of <value> msec`
Cause	High network latency, improper sizing, and component crashes may generate `HighTimeToFirstByte` alert.
Impact	The response for Object GET requests becomes slow.
Resolution	Do the following: • Check for the networking issues impacting the nodes in the Prism Element cluster. • Check if the GETs in each second and PUTs in each second shown on the **Performance** page of the object store exceeds the estimated workload at the time of deployment. If there is an indication of an overload, consider reducing the workload on the Object Store Service (OSS). • Check for crashing services on the Object Store. • Check the Prism Element cluster on which you deployed the object store for any alerts that may be relevant to `HighTimeToFirstByte` alert.

Table 10: Objects Alert - HighObjectStoreUsage

Name	`HighObjectStoreUsage`

Description	The **HighObjectStoreUsage** alert appears when the total object store space usage exceeds the estimated capacity specified at the time of deployment.
Alert Message	**Current object store usage <val> TB exceeds the provisioned capacity <val>TB**
Cause	The following may cause the **HighObjectStoreUsage** alert to generate. • Missing lifecycle management policies • Performing too many object operations.
Impact	Object store gets overloaded, thus causing slow performance.
Resolution	Do the following: • Delete nonessential objects from the object store to reduce the number of objects in the object store. • If the object store usage is unexpectedly high based on the workload and available resources, check for misconfigured Lifecycle Policies on the buckets in the object that may be retaining objects longer than intended. • If the storage space on the Prism Element cluster outside of the object store is high, consider deleting nonessential VMs, snapshots, data, or adding extra nodes to free the storage space in the Prism Element cluster.

Table 11: Objects Alert - HighErrorRate

Name	**HighErrorRate**
Description	The **HighErrorRate** alert appears when the object store returns one or more HTTP 4XX or HTTP 5XX errors in each second for the last 10 minutes.
Alert Message	**Operations issued to Nutanix Buckets have been failing with 5XX/4XX errors with observed error rate <val>/s over the past 10 minutes**
Cause	Improper credentials and component crashes may generate the **HighErrorRate** alert.
Impact	The object store operations fail.

Resolution	Do the following: • Check the client applications for the correct access and secret key combination. • Check for any crashing services on the object store. • Check the Prism Element cluster on which you deployed the object store for any alerts that may be relevant to `HighErrorRate` alert.

Table 12: Objects Replication Alert - ReplicationRPOTimeExceeded

Name	`ReplicationRPOTimeExceeded`
Description	The `ReplicationRPOTimeExceeded` alert appears when the replication of pending objects does not finish even after 12 hours of RPO.
Alert Message	`Last sync time for bucket <bucket name> exceeded RPO time by <time_period_secs>`
Cause	The cause can be a combination of many reasons, such as low network bandwidth, sizing, replication failures, and so on.
Impact	Replication of bucket to remote site lags.
Resolution	Do the following: • Check if any other replication alerts are generated. • Check the connectivity or if the available network bandwidth is sufficient. • Check for any alerts generated due to service crashes.

Table 13: Objects Replication Alert - RemoteEndpointStorageFull

Name	`RemoteEndpointStorageFull`
Description	The `RemoteEndpointStorageFull` alert appears when the object store instance where the remote bucket exists runs out of storage space.
Alert Message	`Storage full on replication endpoint <endpoint name> over the last 15 minutes.`
Cause	The remote objects instance storage becomes full.

Impact	Replication of bucket to remote site fails.
Resolution	Add capacity to the remote objects PE cluster.

Table 14: Objects Replication Alert - RemoteEndpointUnreachable

Name	`RemoteEndpointUnreachable`
Description	The `RemoteEndpointUnreachable` alert appears when the object store instance where the remote bucket exists is not reachable.
Alert Message	`Replication to endpoint <endpoint name> has lost connectivity for the last 15 minutes`
Cause	The following can be the reasons for this alert: • The network connectivity to the remote object store instance is lost. • The remote object store instance is down.
Impact	Replication of bucket to remote site fails.
Resolution	Do the following: • Check the network connectivity to the remote objects instance. • Check if the remote objects instance is up and running.

Table 15: Objects Nfs Alert - HighNfsOpsDropRate

Name	`HighNfsOpsDropRate`
Description	The `HighNfsOpsDropRate` alert is triggered when some of the NFS ops are not being executed for past 10 minutes.
Alert Message	`Due to high payload, few operations submitted to NFS are not executed for the past 10 minutes. Operations are dropped when the outstanding NFS operations have exceeded the threshold value of 1000, or when the QoS queue has reached its maximum capacity of 128 ops, or when the operation wasn't admitted to the queue within 10 seconds.`

Cause	The following can be the reasons for this alert: • NFS payload is causing the number of ops to exceed the threshold value of 1000. • Overall high payload is causing QoS queue to reach its maximum capacity.
Impact	NFS ops are being dropped. Latencies may increase due to client retries.
Resolution	Do the following: • Reduce the overall payload on the impacted bucket. • Modify maximum number of concurrent NFS requests. For example, in CentOS, this can be done by changing `/proc/sys/sunrpc/tcp_slot_table_entries` value to 128.

OBJECTS NOTIFICATIONS

Notifications for Objects enables you to send the completed events logs to the configured endpoints in your Objects instance. This helps you with centralized events log management, thus enabling you to monitor and analyze the logs and identify performance or configuration issues. TCP is the supported protocol for Objects notifications.

The following endpoints are supported for logging the notification events:

- **Syslog**—System Logging Protocol is a standard protocol for sending events log to the Syslog server. Enter the host name or host IP address and port number of your Syslog server when configuring the endpoints. Syslog server should be up and running when performing the endpoint configuration in your Objects instance.
- **Nats-streaming**—A lightweight, reliable streaming platform built on the top of the core NATS platform that provides persistent logs. Enter the host name or host IP address and port number of your NATS streaming server when configuring the endpoints. NATS streaming server should be up and running when performing the endpoint configuration in your Objects instance.

  Note: The default topic used to create the NATS queue is *OSSEvents*. You can use this as the subject while using the NATS client to connect to the NATS streaming server.

Notification Types for Objects

There are two types of notifications for Objects that can be enabled for either successful, or failed, or both successful and failed instance-level and data events.

- **Instance-level events**—All operations performed on a bucket. For example, create a bucket, update a bucket, delete a bucket, enable versioning, and so on. The following instance-level events are supported for successful, or failed, or both successful and failed events:

  Note: All the instance-level events notifications are enabled by default once you configure the endpoints. You cannot enable or disable the notifications for a single instance-level event.

Table 16: S3 APIs for Only Successful Instance-level Events

Supported Instance-level Events	S3 APIs for Only Successful Instance-level Events
Create bucket	S3:BucketCreated:Put
Delete bucket	S3:BucketRemoved:Delete
Add bucket policy	S3:BucketPolicy:Put
Enable static website	S3:BucketWebsite:Put
Disable static website	S3:BucketWebsite:Delete
Create bucket CORS	S3:BucketCORS:Put
Delete bucket CORS	S3:BucketCORS:Delete
Update bucket lifecycle policy	S3:BucketLifecyclePolicy:Update
Delete bucket lifecycle policy	S3:BucketLifecyclePolicy:Delete

Supported Instance-level Events	S3 APIs for Only Successful Instance-level Events
Enable bucket lock (WORM)	S3:BucketLock:Enabled
Disable bucket lock (WORM)	S3:BucketLock:Disabled
Add bucket replication rule	S3:BucketReplication:Put
Delete bucket replication rule	S3:BucketReplication:Delete
Enable bucket versioning	S3:BucketVersioning:Enabled
Suspend bucket versioning	S3:BucketVersioning:Disabled
Enable bucket notification	S3:BucketNotify

Table 17: S3 APIs for Only Failed Instance-level Events

Supported Instance-level Events	S3 APIs for Only Failed Instance-level Events
Create bucket	S3:BucketCreatedError:Put
Delete bucket	S3:BucketRemovedError:Delete
Add bucket policy	S3:BucketPolicyError:Put
Enable static website	S3:BucketWebsiteError:Put
Disable static website	S3:BucketWebsiteError:Delete
Create bucket CORS	S3:BucketCORSError:Put
Delete bucket CORS	S3:BucketCORSError:Delete
Update bucket lifecycle policy	S3:BucketLifecyclePolicyError:Update
Delete bucket lifecycle policy	S3:BucketLifecyclePolicyError:Delete
Enable bucket lock (WORM)	S3:BucketLockError:Enabled
Disable bucket lock (WORM)	S3:BucketLockError:Disabled
Add bucket replication rule	S3:BucketReplicationError:Put
Delete bucket replication rule	S3:BucketReplicationError:Delete
Enable bucket versioning	S3:BucketVersioningError:Enabled
Suspend bucket versioning	S3:BucketVersioningError:Disabled
Enable bucket notification	S3:BucketNotifyError

Table 18: S3 APIs for Both Successful and Failed Instance-level Events

Supported Instance-level Events	S3 APIs for Both Successful and Failed Instance-level Events
Create bucket	S3:BucketCreated:Put and S3:BucketCreatedError:Put
Delete bucket	S3:BucketRemoved:Delete and S3:BucketRemovedError:Delete

Supported Instance-level Events	S3 APIs for Both Successful and Failed Instance-level Events
Add bucket policy	S3:BucketPolicy:Put and S3:BucketPolicyError:Put
Enable static website	S3:BucketWebsite:Put and S3:BucketWebsiteError:Put
Disable static website	S3:BucketWebsite:Delete and S3:BucketWebsiteError:Delete
Create bucket CORS	S3:BucketCORS:Put and S3:BucketCORSError:Put
Delete bucket CORS	S3:BucketCORS:Delete and S3:BucketCORSError:Delete
Update bucket lifecycle policy	S3:BucketLifecyclePolicy:Update and S3:BucketLifecyclePolicyError:Update
Delete bucket lifecycle policy	S3:BucketLifecyclePolicy:Delete and S3:BucketLifecyclePolicyError:Delete
Enable bucket lock (WORM)	S3:BucketLock:Enabled and S3:BucketLockError:Enabled
Disable bucket lock (WORM)	S3:BucketLock:Disabled and S3:BucketLockError:Disabled
Add bucket replication rule	S3:BucketReplication:Put and S3:BucketReplicationError:Put
Delete bucket replication rule	S3:BucketReplication:Delete and S3:BucketReplicationError:Delete
Enable bucket versioning	S3:BucketVersioning:Enabled and S3:BucketVersioningError:Enabled
Suspend bucket versioning	S3:BucketVersioning:Disabled and S3:BucketVersioningError:Disabled
Enable bucket notification	S3:BucketNotify and S3:BucketNotifyError

For more information, see Configuring Events Notification on page 195.

The following is the structure of the notification output for instance-level events that get published to the endpoints.

```
{"EventType":"s3:BucketCreated:Put","Key":"frst","Records":
[{"eventError":"","eventVersion":"2.0","eventSource":"aws:s3","awsRegion":"us-
east-1","eventTime":"2022-02-09T07:40:13Z","eventName":"s3:BucketCreated:Put","userIdentity":
{"principalId":"upgrade@ntnx.com"},"requestParameters":
{"sourceIPAddress":"10.44.91.64"},"responseElements":{"x-amz-
id-2":"110000+1229+2129+240","x-amz-request-id":"110000+1229+2129+240","x-ntnx-origin-
endpoint":"0.0.0.0"},"s3":
{"s3SchemaVersion":"1.0","configurationId":"SuccessfulWriteManagementEvent","OSSDomain":"rc4.eng.nutanix.
{"name":"frst","ownerIdentity":
{"principalId":"upgrade@ntnx.com"},"arn":"arn:aws:s3:::frst"}}}],"level":"info","msg":"","time":"2022-02-
```

- **Data events**—Data events are specific to data operations. These events can be enabled for each bucket. To enable notifications for successful, or failed, or both data events for a bucket, you need to create notification rules. You can define the scope of a notification

rule by selecting **All Objects** or **Subset of objects**. For more information, see Creating Notification Rules for Data Events on page 198.

The following is the structure of the notification output for data events that get published to the endpoints.


```
{"EventType":"s3:ObjectCreated:Put","Key":"put","Records":
[{"eventError":"","eventVersion":"2.0","eventSource":"aws:s3","awsRegion":"us-
east-1","eventTime":"2022-02-09T08:04:48Z","eventName":"s3:ObjectCreated:Put","userIdentity":
{"principalId":"upgrade@ntnx.com"},"requestParameters":
{"sourceIPAddress":"10.44.91.64"},"responseElements":{"x-amz-
id-2":"110000+1229+2791+297","x-amz-request-id":"110000+1229+2791+297","x-ntnx-origin-
endpoint":"0.0.0.0"},"s3":
{"s3SchemaVersion":"1.0","configurationId":"5836c513-51ef-45a9-9c84-6a0c01996e40","OSSDomain":"rc4.eng.nu
{"name":"put","ownerIdentity":
{"principalId":"upgrade@ntnx.com"},"arn":"arn:aws:s3:::put"},"object":
{"key":"csv.docx","size":"31","eTag":"c7a74379a77360a6b85926adb8b2aafb","versionId":"1","sequencer":"1100
```


The following data events are supported for successful, or failed, or both successful and failed events:

> Note: You can enable notifications for all Objects Create Events, Objects Access Events, or Objects Delete Events, or you can choose to enable notifications for one or more sub-categories of Objects Create Events (such as Put, Post, Copy, Multipart Upload Complete), Objects Access Events (such as Get, Head), or Objects Delete Events (such as Delete, Delete Markers).

Table 19: S3 APIs for Only Successful Data Events

Supported Data Events	S3 APIs for Only Successful Data Events
All Create Events	s3:ObjectCreated:*
Put	s3:ObjectCreated:Put
Post	s3:ObjectCreated:Post
Copy	s3:ObjectCreated:Copy
Multipart Upload	s3:ObjectCreated:CompleteMultipartUpload
All Access Events	s3:ObjectAccessed:*
Get	s3:ObjectAccessed:Get
Head	s3:ObjectAccessed:Head
All Delete Events	s3:ObjectRemoved:*
Delete	s3:ObjectRemoved:Delete
Delete Markers	s3:ObjectRemoved:DeleteMarkerCreated

Table 20: S3 APIs for Only Failed Data Events

Supported Data Events	S3 APIs for Only Failed Data Events
All Create Events	s3:ObjectCreatedError:*

Supported Data Events	S3 APIs for Only Failed Data Events
Put	s3:ObjectCreatedError:Put
Post	s3:ObjectCreatedError:Post
Copy	s3:ObjectCreatedError:Copy
Multipart Upload	s3:ObjectCreatedError:CompleteMultipartUpload
All Access Events	s3:ObjectAccessedError:*
Get	s3:ObjectAccessedError:Get
Head	s3:ObjectAccessedError:Head
All Delete Events	s3:ObjectRemovedError:*
Delete	s3:ObjectRemovedError:Delete
Delete Markers	s3:ObjectRemovedError:DeleteMarkerCreated

Table 21: S3 APIs for Both Successful and Failed Data Events

Supported Data Events	S3 APIs for Both Successful and Failed Data Events
All Create Events	s3:ObjectCreated:* and s3:ObjectCreatedError:*
Put	s3:ObjectCreated:Put and s3:ObjectCreatedError:Put
Post	s3:ObjectCreated:Post and s3:ObjectCreatedError:Post
Copy	s3:ObjectCreated:Copy and s3:ObjectCreatedError:Copy
Multipart Upload	s3:ObjectCreated:CompleteMultipartUpload and s3:ObjectCreatedError:CompleteMultipartUpload
All Access Events	s3:ObjectAccessed:* and s3:ObjectAccessedError:*
Get	s3:ObjectAccessed:Get and s3:ObjectAccessedError:Get
Head	s3:ObjectAccessed:Head and s3:ObjectAccessedError:Head
All Delete Events	s3:ObjectRemoved:* and s3:ObjectRemovedError:*
Delete	s3:ObjectRemoved:Delete and s3:ObjectRemovedError:Delete
Delete Markers	s3:ObjectRemoved:DeleteMarkerCreated and s3:ObjectRemovedError:DeleteMarkerCreated

Configuring Events Notification

While configuring events notification, you specify the endpoints. Once you have configured events notification, by default, only successful instance-level events get logged to the endpoints. You can choose to log only successful, only failed, or both successful and failed events, or none.

Before you begin

A non-admin user can configure event notifications only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Set Notification Endpoint
- View Notification Endpoint
- View Object store

> Note: Administrators and non-admin users are Prism Central users with the following roles:
>
> - Administrator—A Super Admin or a Prism Admin in Prism Central.
> - Non-admin user—A Prism Central user without any administrator privileges.
>
> For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

About this task

To configure events notification, do the following:

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store.
   The object store opens on a new window.
3. Click **Settings** > **Notification Endpoints**.
   The Notification Endpoints page appears.

   This page allows you to configure the endpoints by entering the endpoint server details. The available endpoints are **Syslog** and **Nats-streaming**. You can configure any one of these endpoints or both.

4. In the **Syslog** tab, select the **Enable** check box to configure Syslog as the endpoint.

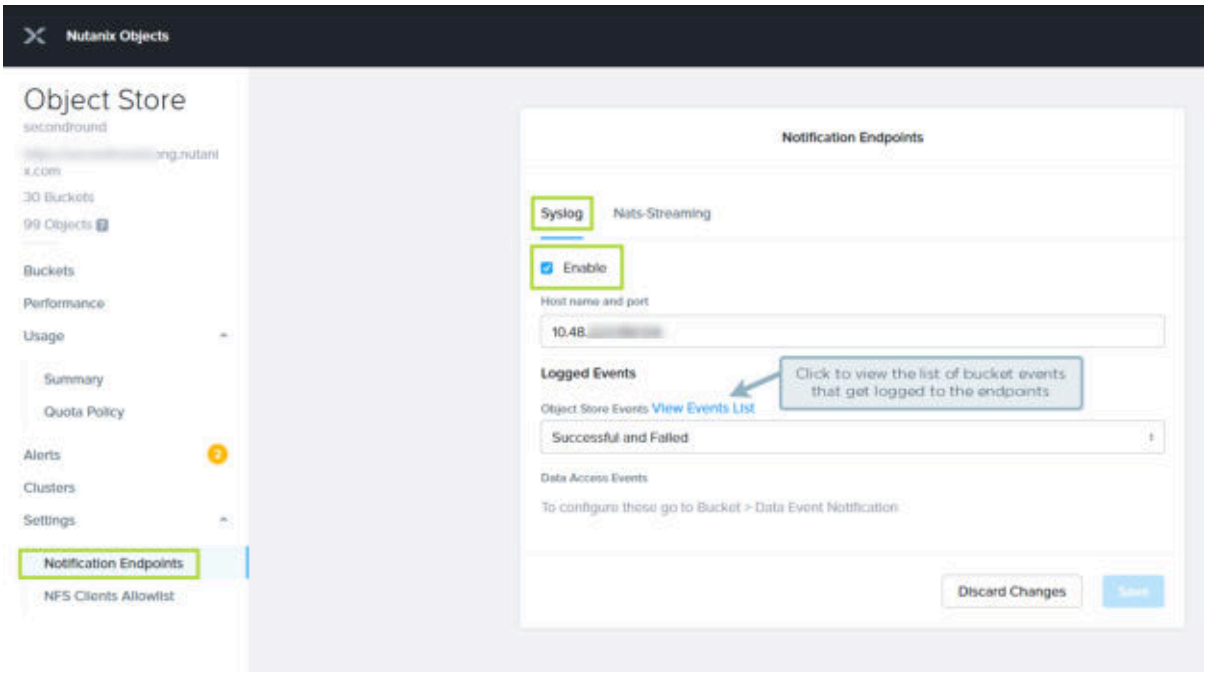

Figure 120: Notification Endpoints: Syslog

a. In the **Host name and port** box, enter the host name and port number of your Syslog server in the *Host name:Port* format.

b. In the **Logged Events** section, select one of the following Object Store events for which you want to enable notifications. By default, **Successful only** option is selected.

- **Successful and Failed**: Logs both successful and failed object store events.
- **Successful only**: Logs only successful object store events.
- **Failed only**: Logs only failed object store events.
- **None**: Logs none of the object store events.

> Note: You can click **View Events List** link to view the list of bucket events that get logged to the endpoints.

5. In the **Nats-streaming** tab, select the **Enable** check box to configure Nats-streaming as the endpoint.

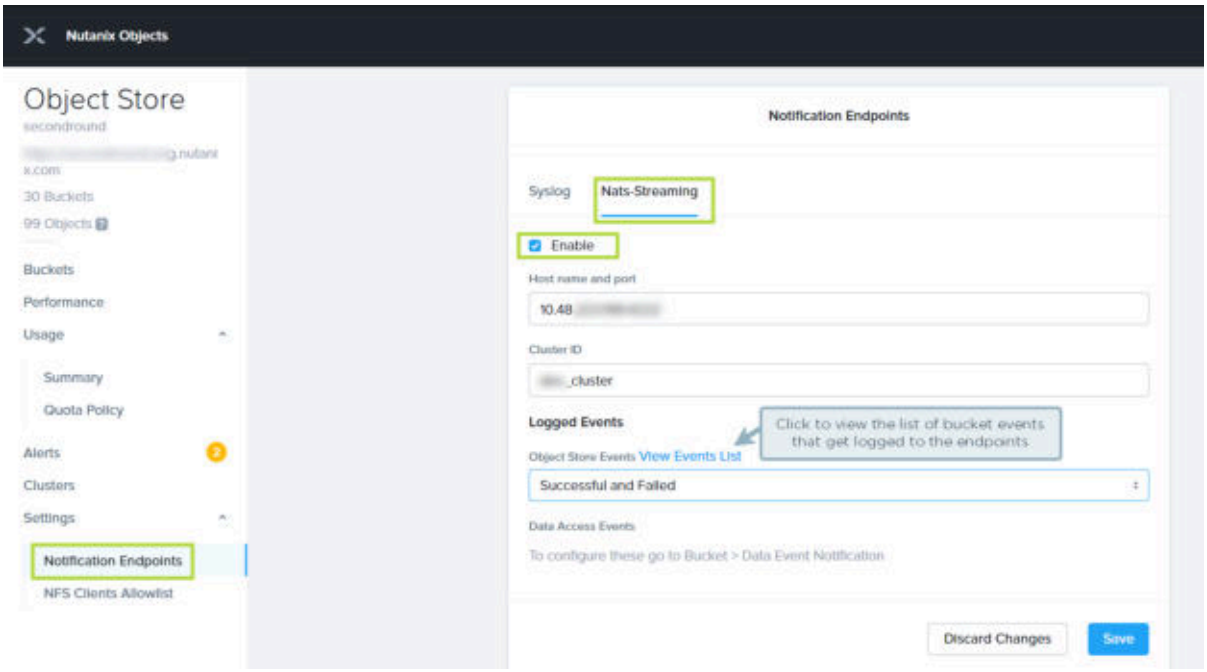

Figure 121: Notification Endpoints: Nats-Streaming

   a. In the **Host name and port** box, enter the host name and port number of your NATS streaming server in the *Host name:Port* format.

   b. In the **Cluster ID** box, enter the cluster ID of the server that you used to start the NATS server.

   c. In the **Logged Events** section, select one of the following Object Store events for which you want to enable notifications. By default, **Successful only** option is selected.

      - **Successful and Failed**: Logs both successful and failed object store events.
      - **Successful only**: Logs only successful object store events.
      - **Failed only**: Logs only failed object store events.
      - **None**: Logs none of the object store events.

      Note: You can click **View Events List** link to view the list of bucket events that get logged to the endpoints.

6. Click **Save** to complete the events notification configuration.

   Notification is enabled for your Objects instance. Also, the status of the notification endpoint is shown in the **Notifications** column. For example, if the NATS endpoint gets disconnected

due to network issues, a red icon appears in the **Notifications** column. Hovering over the red icon displays `Cannot connect to nats endpoint` message.

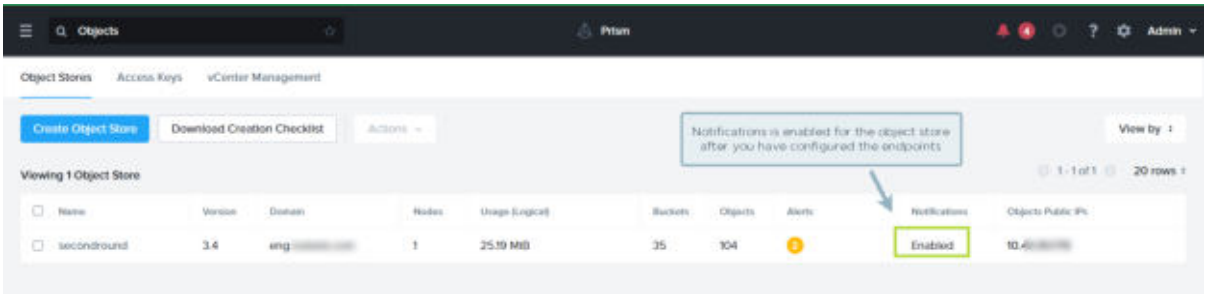

Figure 122: Object Store List: Enabled Notifications for Object Store

7. (Optional) To undo an endpoint configuration, do the following.

   a. Click **Settings** > **Notification Endpoints** to open the Notification Endpoints page.

   b. Clear the check box for the endpoint that you want to remove, and then click **Save**.

   c. In the Data Event Notification page, delete all the notification rules corresponding to the endpoint that you removed before adding a new notification rule to the bucket. For more information, see Creating Notification Rules for Data Events on page 198.

What to do next

You can now create notification rules for data events of the bucket, refer to Creating Notification Rules for Data Events on page 198.

Creating Notification Rules for Data Events

To keep track of the data events that get completed successfully, you need to create notification rules for the buckets.

Before you begin

A non-admin user can create notification rules for data events only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Notification Endpoint
- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
  - Administrator—A Super Admin or a Prism Admin in Prism Central.
  - Non-admin user—A Prism Central user without any administrator privileges.

  For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

- Users with Edit Bucket permission can create notification rules for data events on all buckets available in the Object stores that are assigned to the respective users.

About this task

To create a notification rule for a bucket, do the following.

Procedure

1. Log on to the Prism Central web console, and click the **Entity** menu > **Services** > **Objects**.
2. Click the name of the object store.
   The object store opens on a new window. You can view the buckets in your Objects instance.
3. In the **Buckets** table, select the bucket for which you want to create a notification rule.
4. Click **Actions** > **Data Events Notification**.
   The Data Events Notification page appears.
5. Click **Create Rule** to open the Create Rule page.
   If a rule is already created for that bucket, click **+ Add Rule** to add a new rule in the Create Rule page.

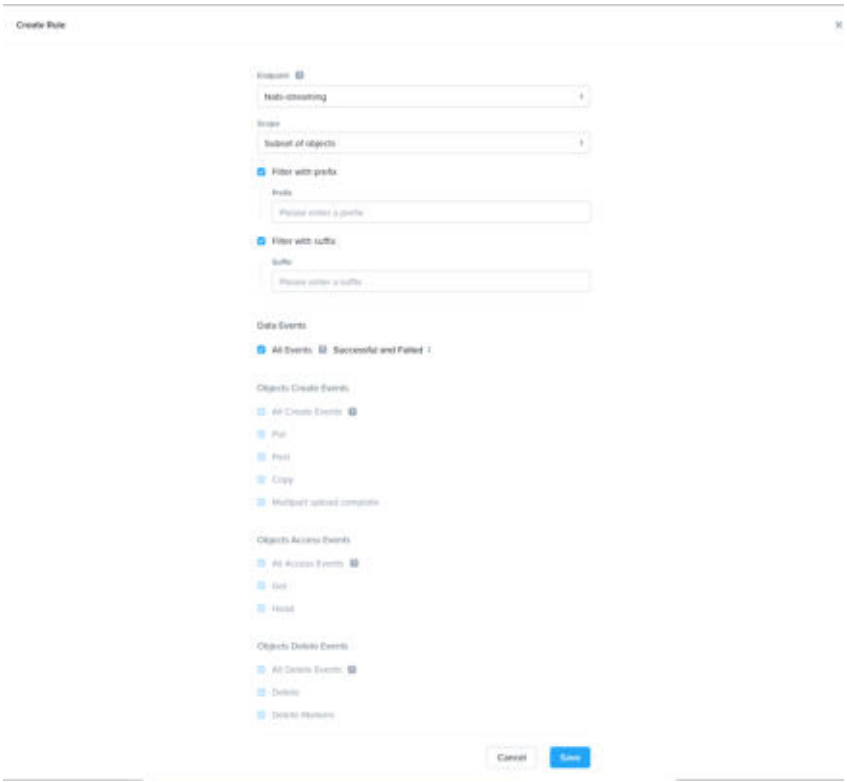

Figure 123: Create Rule

6. In the Create Rule page, do the following.

   a. In the **Endpoint** list, select the endpoint where you want the data events for this bucket to get logged.

      To update endpoints for your Objects instance, see Configuring Events Notification.

   b. In the **Scope** list, select **All Objects** or **Subset of objects**.

      - If you select **All Objects**, the notification rule you are creating applies to all the objects.
      - Select **Subset of objects** to apply the notification rule to specific objects. You can filter Objects by entering a prefix, suffix, or both. Objects with the prefix and suffix you specified get filtered and the rule applies to those objects.

        Note:

        - You cannot proceed without providing suffix or prefix for the scope Subset of objects.
        - You can enter only one prefix and suffix. You can create another rule with a different suffix and prefix.
        - The suffix-prefix pair in two rules corresponding to the same event should not be overlapping. For example, any given Object name *xyz.jpg* should be selected by a maximum one rule corresponding to the given event type.

   c. In the **Data Events** section, click **All Events** check box to select all data events, or select a check box next to the name of the events listed under Objects Create Events, Objects Access Events, and Objects Delete Events, and choose from the following that you want to get logged to the endpoint.

      - **Successful and Failed**: Logs both successful and failed object store events.
      - **Successful only**: Logs only successful object store events.
      - **Failed only**: Logs only failed object store events.

        Note: When you select **All Events** check box, all the Objects Create Events, Objects Access Events, and Objects Delete Events will be selected.

7. Click **Save** to complete, and then click **Done** to close the Create Rule page.

   The notification rule you created appears in the list of data events notification rules in the Data Event Notification page. You can click **All Events** or ***Number*** **Events** in the Events column

to view the configured data events for that bucket. For example, click **8 Events** to view the 8 data events configured.

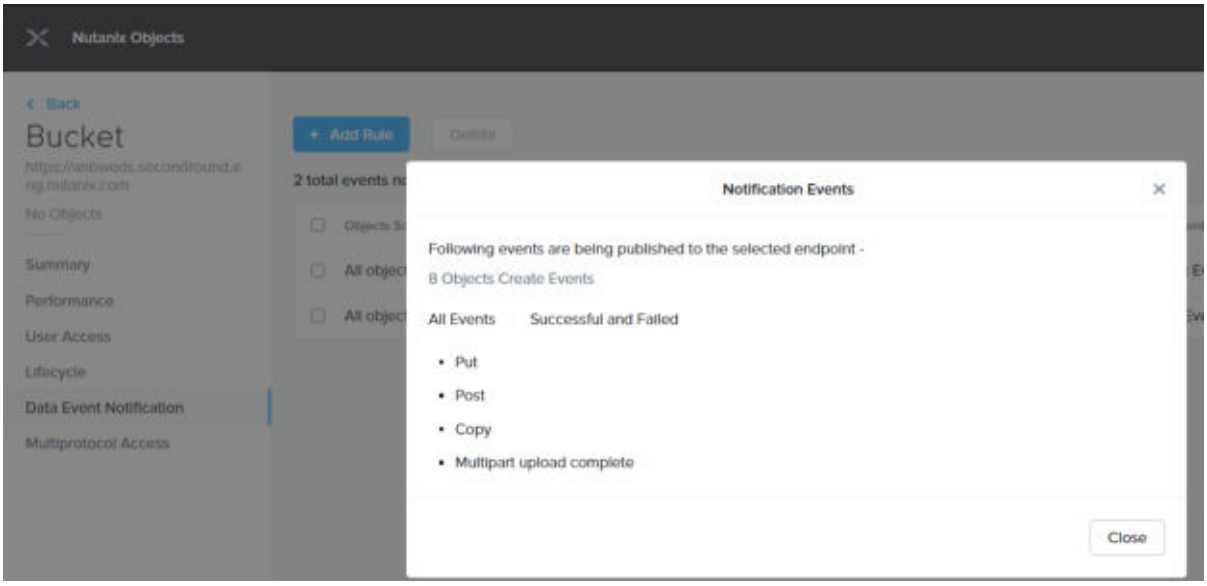

Figure 124: Configured Notification Events

Also, you can delete a notification rule. In the Data Event Notification page, select a notification rule from the list, and then click **Delete**.

After you create a notification rule for a bucket, the status in the **Notifications** column changes to **Enabled**.

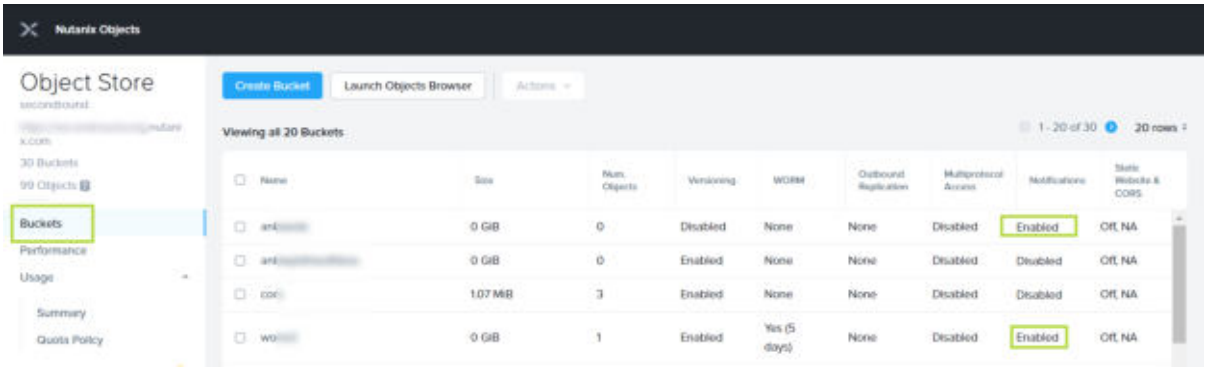

Figure 125: Notification Rule - Enabled for Data Events

CRUD OPERATIONS BY USING S3 APIS

You can create and manage buckets and objects by using the S3 APIs, which are then visible in Prism Central.

Authentication

You can send requests to Objects by using the REST API or the Amazon Web Services Software Development Kit (AWS SDK) wrapper libraries that wrap the underlying S3 REST API. Every interaction with Objects is authenticated. In this authentication process, the identity of the requester who is trying to access Objects is verified with a signature value. The signature value is generated from the AWS access keys (access key ID and secret access key) of the requester. This AWS access keys and endpoint URL is provided by the administrator to the user.

If you are using the AWS SDK, the libraries compute the signature from the keys you provide. However, if you make direct REST API calls, the signature is computed from the request.

For creating the buckets and objects, you need the following information from the administrator:

- Endpoint URL (Static IP address and the Port number)
- Access key ID
- Secret access key

Once you get this information, you can import the SDK libraries, create a session and a client, and then you can start making the requests (for example, creating buckets and objects).

> Note: Objects supports the following:
>
> - Signature Version 2 and Signature Version 4 (regular and pre-signed)
> - Streaming signed payloads for PUT requests with Signature Version 4

Supported and Unsupported APIs

This section describes Objects support for Amazon S3 API. The object store service is available on the following ports:

- HTTP: 80
- HTTPS: 443

> Note: Transport Layer Security (TLS) 1.2 is supported on Objects.

Supported APIs

The following table lists the supported S3 API methods.

> Note:
>
> - User-provided meta-data header names get stored in lower case.
> - User-defined object meta-data is supported for `PUT Object` and `PUT Object – Copy` APIs and limited to 2 KB in size.

Table 22: Supported S3 APIs and Parameters

Supported S3 APIs	Request Parameters/Request Headers/ Request Elements/Request Body
PUT Bucket	Common Headers on page 205 + CreateBucketConfiguration, LocationConstraint (By default, LocationConstraint “us-east-1” is only supported.)
PUT Bucket Lifecycle	Common Headers on page 205 + Content-MD5 + AbortIncompleteMultipartUpload, And, Date, Days, DaysAfterInitiation, Expiration, Filter, ID, Key, LifecycleConfiguration, NoncurrentDays, NoncurrentVersionExpiration, NoncurrentVersionTransition, Prefix, Rule, Status
PUT Bucket Policy	Common Headers on page 205 + JSON string containing the policy contents containing the policy statements
PUT Bucket versioning	Status, VersioningConfiguration
Complete Multipart Upload	Common Headers on page 205 + CompleteMultipartUpload, Part, PartNumber, ETag
PUT Object - Copy	Common Headers on page 205 + x-amz-copy-source, x-amz-metadata-directive, x-amz-copy-source-if-match, x-amz-copy-source-if-none-match, x-amz-copy-source-if-unmodified-since, x-amz-copy-source-if-modified-since, x-amz-tagging, x-amz-tagging-directive, x-amz-meta-, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
DELETE Objects	Single query string parameter "delete" + Common Headers + Content-MD5, Content-Length, Delete, Quiet, Object, Key, VersionId
GET Object	Common Headers on page 205 + Range, If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
HEAD Object	Common Headers on page 205 + Range, If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
List Multipart Uploads	Common Headers on page 205 + delimiter,max-uploads, key-marker, prefix, upload-id-marker

Supported S3 APIs	Request Parameters/Request Headers/ Request Elements/Request Body
List Object Versions	Common Headers on page 205 + delimiter, key-marker, max-keys, prefix, version-id-marker
GET Bucket (List Objects) Version 1	Common Headers on page 205 + delimiter, marker, max-keys, prefix
GET Bucket (List Objects) Version 2	Common Headers on page 205 + delimiter, max-keys, prefix, list-type, continuation-token, fetch-owner, start-after
List Parts	Common Headers on page 205 + uploadId, max-parts, part-number-marker
PUT Object	Common Headers on page 205 + Content-Length, Content-MD5, Expect, x-amz-tagging, x-amz-meta-, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date, x-amz-object-lock-legal-hold
INITIATE Multipart Upload	Common Headers on page 205 + x-amz-meta-, x-amz-tagging, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
UPLOAD Part	Content-Length, Content-MD5, Expect
Upload Part - Copy	x-amz-copy-source, x-amz-copy-source-range, x-amz-copy-source-if-match, x-amz-copy-source-if-none-match, x-amz-copy-source-if-unmodified-since, x-amz-copy-source-if-modified-since
PUT Bucket Object Lock Configuration	Common Headers on page 205 + ObjectLockEnabled Rule
DELETE Bucket	Common Headers on page 205 + Bucket
Delete Multiple Objects	Common Headers on page 205 + Bucket, Delete, Quiet, Objects, Key, VersionId
GET Object Retention	Common Headers on page 205 + Bucket, Key, VersionId
PUT Object Retention	Common Headers on page 205 + Bucket, Key, VersionId, Retention, Mode, RetainUntilDate
PUT Object Legal Hold	Common Headers on page 205 + Bucket, Key, versionId, LegalHold, Status
GET Object Legal Hold	Common Headers on page 205 + Bucket, Key, versionId
Select Object Content	Expression, ExpressionType, InputSerialization, CSV, AllowQuotedRecordDelimiter, Comments, FieldDelimiter, FileHeaderInfo, QuoteCharacter, QuoteEscapeCharacter, RecordDelimiter, OutputSerialization, CSV

> Note: The maximum size limit of a multipart object is 16 TiB. This limit is different from the standard S3 API limit of 5 TiB.

For more information about objects tagging API, refer to Objects Tagging APIs Overview on page 207.

For more information about S3 Select API, refer to S3 Select API Overview on page 208.

Pre-signed URLs can be generated for Objects. For more information about generating the pre-signed URLs, see *Signing and authenticating REST requests* section in the *Amazon Simple Storage Service Developer Guide*.

> Note: You can limit a presigned request by specifying an expiration time. You can set an expiration time of more than seven days.

The following table lists the supported APIs that use only Common Headers on page 205:

Supported APIs
GET Bucket lifecycle configuration
PUT Bucket lifecycle configuration
GET Bucket Object Lock Configuration
GET Bucket Location
GET Bucket Policy
GET Bucket versioning
GET Bucket ACL
GET Bucket notification configuration
GET Bucket cors
GET Bucket website
GET Bucket replication
PUT Bucket cors
PUT Bucket notification configuration
PUT Bucket website
PUT Bucket replication
DELETE Bucket cors
DELETE Bucket website
DELETE Bucket Lifecycle
DELETE Bucket Policy
DELETE Bucket replication
DELETE Object
HEAD Bucket
LIST Bucket
ABORT Multipart Upload

Common Headers

You can use the following headers while making requests:

Table 23: List of Common Headers

Authorization
Content-Length
Content-Type
Content-MD5
Date
Expect
Host
x-amz-content-sha256
x-amz-date
x-amz-security-token

For more information on common headers, refer to *Common Request Headers* section in the *Amazon Simple Storage Service API Reference Guide*.

Unsupported APIs

The following table lists the unsupported S3 API methods:

Table 24: Unsupported S3 APIs

Unsupported S3 APIs
GET Bucket accelerate configuration
GET Bucket analytics configuration
GET Bucket encryption
GET Bucket inventory configuration
GET Bucket logging
GET Bucket metrics configuration
GET Bucket requestPayment
List Bucket Analytics Configurations
List Bucket Inventory Configurations
List Bucket Metrics Configurations
PUT Bucket accelerate configuration
PUT Bucket acl
PUT Bucket analytics configuration
PUT Bucket encryption
PUT Bucket inventory configuration
PUT Bucket logging
PUT Bucket metrics configuration
DELETE Bucket analytics configuration

Unsupported S3 APIs
DELETE Bucket encryption
DELETE Bucket inventory configuration
DELETE Bucket metrics configuration
DELETE Bucket encryption
Copy
GET Object ACL
PUT Object ACL
Restore Object
Select Object Content

Objects Tagging APIs Overview

A tag is a label that you assign to an object, and it consists of a key and a value pair that you can define.

The tagging APIs feature allows you to add, retrieve, and remove tags for your objects.

- Retrieving objects metadata also returns the number of tags associated with the object, if any.
- Tagging is also supported with a few other Object APIs.
- The maximum number of tags allowed per object is 10.
- Tag keys can be up to 128 Unicode characters in length, and tag values can be up to 256 Unicode characters in length.
- Tag keys and values are case-sensitive.

API Operations Supported for Tagging

Currently, Objects only supports APIs for tagging objects. Tag-based bucket lifecycle policy management and object listing are not supported.

Nutanix Objects supports the following object tagging API operations.

- **PUT Object tagging**: Replace the tags associated with an object. You can add the tags in the request body.

  The following two scenarios are involved:

  - You can add tags to an object with no tags associated with it.
  - You can replace the existing tags associated with an object.
- **GET Object tagging**: Retrieves the tags associated with an object.
- **DELETE Object tagging**: Deletes the tags associated with an object.

The following Object APIs also support tagging:

- GET Object (returns tag count, if any)
- PUT Object
- PUT Object-Copy

- Initiate Multipart Upload

S3 Select API Overview

S3 Select is a feature to select partial content of an object and returns the result.

With S3 Select feature, you can only fetch the content of interest. This allows applications to speed up their queries drastically as the size of the sub-set of object is smaller compared to the actual size of the entire object. S3 Select is supported on objects stored in CSV format and returns the result in CSV format.

> Note: S3 select does not support compression.

The `SelectObjectContent` API filters the contents of an object located in Nutanix Objects using an SQL statement. In the request, you must specify the SQL expression and data serialization format (CSV) of the object. Objects uses this format to parse object data into records, and returns only records that match the specified SQL expression. You must also specify the data serialization format (CSV) for the response.

Basic Command Support

Objects S3 Select feature support only the `SELECT` SQL command.

The S3 Select query operates on a single object. The primary commands involves data selection and filtering.

The following clauses are supported for the SELECT command.

- SELECT list: The `SELECT` list names the columns, functions, and expressions that you want the query to return.

  The following are supported:

  - Aggregate operators such as AVG, COUNT, MAX, MIN
  - Arithmetic operations (For example, +, /)
  - _1, _2 positional operators for CSV.
- FROM clause: Using `FROM` clause, you can select from arrays or objects within a JSON object. You can use `S3Object` as a basic reference.
- WHERE clause: The `WHERE` clause filters rows based on the condition. A condition is an expression that has a Boolean result.

  The following conditions are supported:

  - AND, NOT, OR operators
  - Comparision (For example, <, >=)
  - Pattern matching (For example, LIKE, %)
- LIMIT clause: The `LIMIT` clause limits the number of records that you want the query to return based on number.

The following is the general syntax of a `SELECT` query.

```
SELECT fields FROM S3Object WHERE condition LIMIT num
```

> Note: The specific parameters differ based on the format and schema of the object.

For more information on the supported SQL functions, refer to Supported SQL Functions on page 209.

Requirements

The following is the requirement for using S3 SELECT:

- You must have `s3:GetObject` permission for the object you are querying.

Limitations

The following are limitations for using S3 SELECT:

- The maximum length of a SQL expression is 256 KB.
- The maximum length of a record in the input or result is 1 MB.
- Complex operations like sub-queries or joins are not supported.

Supported SQL Functions

S3 Select supports the following SQL functions.

Table 25: Type of Supported Functions

Type of Functions	Functions with Parameters
Aggregate	**avg (X)**: Returns the average value of all X within a group. String and BLOB values that are not numbers are interpreted as 0. **count (X) or count(*)**: Returns a count of the number of times X is in a group. The count(*) function (with no arguments) returns the total number of rows in the group. **max(X)**: Returns the maximum value of all values in the group. It is the value that is returned last in an ORDER BY on the same column. **min(X)**: Returns the minimum value of all values in the group. This is the first value that appears in an ORDER BY of the column. **sum(X)**: Returns sum of all values in the group.

Type of Functions	Functions with Parameters
Conditional	**CASE**: A CASE expression is similar to IF-THEN-ELSE. The optional expression that occurs in between the CASE keyword and the first WHEN keyword is called the base expression. There are two forms of the CASE expression: • **A CASE with a base expression**: The base expression is evaluated just once and the result is compared against the evaluation of each WHEN expression from left to right. The result of the CASE expression is the evaluation of the THEN expression that corresponds to the first WHEN expression for which the comparison is true. Or, if none of the WHEN expressions evaluate to a value equal to the base expression, the result of evaluating the ELSE expression, if any. If there is no ELSE expression and none of the WHEN expressions produce a result equal to the base expression, the overall result is NULL. • **A CASE without a base expression**: Each WHEN expression is evaluated and the result treated as a boolean, starting with the leftmost and continuing to the right. The result of the CASE expression is the evaluation of the THEN expression that corresponds to the first WHEN expression that evaluates to true. Or, if none of the WHEN expressions evaluate to true, the result of evaluating the ELSE expression, if any. If there is no ELSE expression and none of the WHEN expressions are true, then the overall result is NULL. A NULL result is considered untrue when evaluating. **coalesce(X,Y,...)**: Returns a copy of the first argument, or NULL if all arguments are NULL. Coalesce() must have at least 2 arguments. **nullif(X,Y)**: Returns its first argument if the arguments are different and NULL if the arguments are the same. This function searches its arguments from left to right for an argument that defines a collating function and uses that collating function for all string comparisons.

Type of Functions	Functions with Parameters
Conversions	**CAST**: A CAST expression of the form CAST(expr AS type-name) is used to convert the value of expr to a different storage class specified by type-name. The following are the Conversion Processing ways for the different type-name: • Casting a value with no affinity converts into a BLOB. Casting to a BLOB first casts the value to TEXT, then interprets the resulting byte sequence as a BLOB instead of TEXT. • Casting a BLOB value to TEXT, the sequence of bytes that make up the BLOB is interpreted as text encoded using the database encoding. • Casting a BLOB value to a REAL, the value is first converted to TEXT. • Casting a BLOB value to INTEGER, the value is first converted to TEXT. • Casting a TEXT or BLOB value into NUMERIC yields either an INTEGER or a REAL result.

Type of Functions	Functions with Parameters
Date	**date(time-value, modifier, modifier, ...)** **time(time-value, modifier, modifier, ...)** **datetime(time-value, modifier, modifier, ...)** All the date and time functions take a time value as an argument. The time value is followed by zero or more modifiers. The date() function returns the date in the `YYYY-MM-DD` format. The time() function returns the time in the `HH:MM:SS` format. The datetime() function returns in the `YYYY-MM-DD HH:MM:SS` format. **strftime(format, time-value, modifier, modifier, ...)** The strftime() function takes a format string as its first argument. It returns the formatted date according to the format string specified as the first argument. The following is a list of valid strftime() substitutions: • %d - day of month: 00 • %f - fractional seconds: SS.SSS • %H - hour: 00-24 • %j - day of year: 001-366 • %J - Julian day number • %m - month: 01-12 • %M - minute: 00-59 • %s - seconds since 1970-01-01 • %S - seconds: 00-59 • %w - day of week 0-6 with Sunday==0 • %W - week of year: 00-53 • %Y - year: 0000-9999 • %% - %

Type of Functions	Functions with Parameters
String	**length(X)**: Returns the following based on the value of X: • For a string value, returns the number of characters (not bytes) in X prior to the first NULL character, • For a blob value, returns the number of bytes in the blob. • For a NULL value, returns NULL. • For a numeric value, returns the length of a string representation of X. **lower(X)**: Returns a copy of string X with all ASCII characters converted to lower case. Note: The default lower() function works for ASCII characters only. **substring(X,Y,Z)**: Returns a substring of input string X that begins with the Y-th character and which is Z characters long. Returns the following based on different scenarios: • If Z is omitted, returns all characters through the end of the string X beginning with the Y-th. • If Y is negative, then the first character of the substring is found by counting from the right rather than the left. • If Z is negative, then the Z characters preceding the Y-th character are returned. • If X is a string, then the indices refer to actual UTF-8 characters. • If X is a BLOB, then the indices refer to bytes. **ltrim(X,Y)**: Returns a string formed by removing any and all characters that appear in Y from the left side of X. If the Y argument is omitted, ltrim(X) removes spaces from the left side of X. **rtrim(X,Y)**: Returns a string formed by removing any and all characters that appear in Y from the right side of X. If the Y argument is omitted, rtrim(X) removes spaces from the right side of X. **trim(X,Y)**: Returns a string formed by removing any and all characters that appear in Y from both ends of X. If the Y argument is omitted, trim(X) removes spaces from both ends of X. **upper(X)**: Returns a copy of input string X in which all lower-case ASCII characters are

ERROR RESPONSES

This section provides reference information about Objects errors responses and codes. When Objects request returns an error, the client receives an error response. The format of the error response is API specific; however, all the error responses have common elements.

REST Error Responses

When an error occurs, the header information contains the following: Content-Type: application/xml; An appropriate status code.

Table 26: REST Error Responses

Name
Code
Error
Message
RequestId
Resource

For more information on Error Responses, refer to the *REST Error Responses* section in the *Amazon Simple Storage Service API Reference Guide*.

List of Error Codes

The following table lists the error codes:

Table 27: Error Codes

Error Code	HTTP Status Code
AccessDenied	403 Forbidden
AuthorizationHeaderMalformed	400 Bad Request
BadDigest	400 Bad Request
BucketAlreadyExists	409 Conflict
BucketAlreadyOwnedByYou	409 Conflict
BucketNotEmpty	409 Conflict
EntityTooSmall	400 Bad Request
EntityTooLarge	400 Bad Request
IncompleteBody	400 Bad Request
InlineDataTooLarge	400 Bad Request
InternalError	500 Internal Server Error
InvalidAccessKeyId	403 Forbidden

Error Code	HTTP Status Code
InvalidArgument	400 Bad Request
InvalidBucketName	400 Bad Request
InvalidBucketState	409 Conflict
InvalidDigest	400 Bad Request
InvalidLocationConstraint	400 Bad Request
InvalidObjectState	403 Forbidden
InvalidPart	400 Bad Request
InvalidPartOrder	400 Bad Request
InvalidPolicyDocument	400 Bad Request
InvalidRange	416 Requested Range Not Satisfied
InvalidRequest	400 Bad Request
InvalidURI	400 Bad Request
KeyTooLongError	400 Bad Request
MalformedACLError	400 Bad Request
MalformedPOSTRequest	400 Bad Request
MalformedXML	400 Bad Request
MaxMessageLengthExceeded	400 Bad Request
MaxPostPreDataLengthExceededError	400 Bad Request
MetadataTooLarge	400 Bad Request
MethodNotAllowed	405 Method Not Allowed
MissingContentLength	411 Length Required
MissingRequestBodyError	400 Bad Request
NoSuchBucket	404 Not Found
NoSuchBucketPolicy	404 Not Found
NoSuchKey	404 Not Found
NoSuchLifecycleConfiguration	404 Not Found
NoSuchUpload	404 Not Found
InvalidVersion	404 Not Found
NotImplemented	501 Not Implemented
OperationAborted	409 Conflict
PermanentRedirect	301 Moved Permanent
PreconditionFailed	412 Precondition Failed
Redirect	307 Moved Temporarily
RequestIsNotMultiPartContent	400 Bad Request
RequestTimeout	400 Bad Request
RequestTimeTooSkewed	403 Forbidden

Error Code	HTTP Status Code
SignatureDoesNotMatch	403 Forbidden
ServiceUnavailable	503 Service Unavailable
SlowDown	503 Slow Down
TemporaryRedirect	307 Moved Temporarily
UnexpectedContent	400 Bad Request

For more information on List of Error Codes, refer to the *List of Error Codes* section in the *Amazon Simple Storage Service API Reference Guide*.

INTEGRATION WITH BACKUP APPLICATIONS

Objects is ideal for cost effective, scale-out storage. It provides a fully distributed, API-accessible storage platform that integrates directly into applications or used for backup, archiving, and data retention. Objects offers you a seamless way to switch from your traditional backup to object store backup. You can also perform multipart upload. Objects supports integration with back up applications such as Commvault, HYCU, Veeam and Veritas.

> Note: Some backup vendors have configurable size limit if larger than 5 TiB VM images need to be backed up (For example, HYCU). The backup appliance configuration needs to be changed in order to take advantage of larger size limit in Nutanix Objects.

For more information about Commvault Integration, refer to *Commvault with Nutanix* guide on the Nutanix Support Portal.

OBJECTS LCM UPGRADES

Objects version upgrades by using the Life Cycle Management (LCM) feature. You can perform LCM upgrades through Prism Central (PC). Objects is a part of the PC upgrades module in LCM. LCM upgrades the following components of Objects: ***Objects Manager*** and ***Objects Services***.

Objects Manager

Objects Manager is a containerized service running on PC VM. The Objects Manager is primarily responsible for taking input from a user for deploying the object store, validating the user inputs, managing certificates, deploying the Objects Services, and serves as an interface between PC and backing object store. A single Objects Manager can manage one or more object stores. In case of scale-out PC, the Objects Manager service runs on each of the PC nodes and provides high availability.

> Note: During upgrade of the Objects Manager, no disruption happens to the Objects IO. However, the user interface will not be available for a short period of time for statistics and management.

For more information on updating Objects Manager, refer to Upgrading Objects Manager on page 220.

Objects Service

Objects Service provides the object store interface and is responsible for storing and retrieving objects. The objects and metadata for the Objects are stored on the selected Prism Element clusters. The various services that perform the task are containerized and run on the Kubernetes platform. Each Objects Service instance provides a single global namespace. During deployment of the Objects Service, the required number of VMs are created for running the Kubernetes pods and the load balancer.

> Note: During upgrade of the Objects Services, disruption to the IO is expected as each of the internal services get upgraded. The upgrade process can take 15 to 30 minutes.

For more information on updating Objects Services, refer to Upgrading Objects Service on page 221.

> Note:
>
> - Upgrades should be in the following order:
>
> 1. Prism Central
> 2. Prism Element
> 3. MSP Controller. For more information refer to MSP.
> 4. Objects Manager
> 5. Objects Service
>
> - If Objects is not enabled and you upgrade the Prism Central, then the Objects Manager will be upgraded automatically. However, if Objects is enabled, then you have to upgrade the Objects Manager manually.

First Time Objects Users

You must perform the following tasks to install and update Objects in your environment.

1. Perform an inventory and upgrade to a compatible version of Prism Central by using the Life Cycle Management (LCM) feature in Prism Central, refer to the *Life Cycle Management Guide*.
2. Enable Objects, refer to Enabling Objects on page 15.
3. In LCM, perform inventory and update Objects to an available GA version.

Microservices Platform (MSP)

Microservices Platform (MSP) is a platform based on Kubernetes where all the Objects microservices run.

The lifecycle of MSP is controlled by a service running on Prism Central called MSP Controller. You need to upgrade the MSP Controller before upgrading the Objects Manager and Objects Service. The first object store cluster deployed on the Prism Central is the primary MSP or the primary object store. Rest all the service MSPs are the secondary MSPs or the secondary object stores.

> Note: The primary object store cluster hosts all the common services such as IAM, and you cannot delete the primary cluster without deleting the secondary cluster.

Finding the MSP Version

About this task

Run the following command on Prism Central VM to find the MSP version:

Procedure

1. Log on to Prism Central VM.
2. Run the following command:
   ```
   admin@pcvm$ mspctl controller version
   ```

Finding the Primary and Secondary MSP Clusters

About this task

Run the following command on Prism Central VM to find the primary and secondary MSP clusters:

Procedure

1. Log on to Prism Central VM.
2. Run the following command:
   ```
   admin@pcvm$ mspctl cluster list
   ```

   > Note: If the cluster_type value is **primary_msp**, then the cluster is a primary MSP cluster. Other clusters are secondary.

Upgrading MSP Controller

You can upgrade the MSP Controller from the LCM page in Prism Central.

About this task

> Note:

- For Objects deployment, it is recommended to upgrade to the latest MSP controller version.
- Ensure that you have MSP version 1.0.8 to:
  - upgrade to MSP 2.2.1.
  - allow the objects to be visible for upgrade in the dark-side mode.

To upgrade the MSP controller, do the following:

Procedure

1. Log on to the Prism Central web console, and then click the **Entity** menu **Administration** > **LCM** to open the LCM page.
2. Click **Inventory**, and then click **Perform Inventory**.
   The table lists all the updated modules.
3. Click **Updates** > **Software**.
4. Select the **MSP** check box, and then click **Update**.
5. Click **Apply Updates** to finish the update process.
   The MSP Controller is upgraded.

Upgrading Objects Manager

You need to upgrade MSP before upgrading the Objects Manager.

About this task

Note: Objects Manager upgrade will be disabled when there is at least one unresolved Objects Service upgrade failure.

To upgrade Objects Manager, do the following:

Procedure

1. Log on to Prism Central,
2. Click **Enable** Objects if Objects is not already enabled.
   For more information, refer to Enabling Objects on page 15.
3. Click the **Entity menu** > **Administration** > **LCM** on the Prism Central dashboard to open the LCM page.
4. Click **Inventory**.
5. Click **Perform Inventory**.

   Note: It takes 10 to15 minutes to list all the updated modules.

   The table lists all the updated modules.

6. If a new Objects Manager version is available, then select the **Objects Manager** check box, and then click **Update**.

   If multiple versions of Objects Manager are available, select the latest version. You can also select MSP, Objects Manager and Objects services, and upgrade them together. When multiple options for upgrade are selected, the upgrade happens serially and not in parallel.

7. Click **Apply Updates** to finish the update process.
   Objects Manager is updated to the latest version.

   > Note: The update process takes about 10 to 15 minutes to complete.

8. Click **Inventory** > **Perform Inventory** to verify the latest version.

Upgrading Objects Service

You can upgrade the Objects Service after the Objects Manager is upgraded.

About this task

> Note: Objects Service is available for upgrade only after deployment.

To upgrade Objects Service, do the following:

Procedure

1. Log on to Prism Central.
2. Click the **Entity menu** > **Administration** > **LCM** on the Prism Central dashboard to open the LCM page.
3. Click **Inventory**.
4. Click **Perform Inventory**.

   > Note: It takes 10 to 15 minutes to list all the updated modules.

   The table lists all the updated modules.

5. Click **Updates** > **Software**.
6. Select the **Objects Service** instance check box, and then click **Update**.

   Multiple Object Service instances will be listed separately. You can either select all of the instances and upgrade them together, or select and upgrade the instances individually. You can also select MSP, Objects Manager and Objects services, and upgrade them together. When multiple options for upgrade are selected, the upgrade happens serially and not in parallel.

7. Click **Apply Updates** to finish the update process.
   Objects Service is updated to the latest version.

   > Note: The update process takes about 15 to 30 minutes for each Object Service instance to complete.

8. Click **Inventory** > **Perform Inventory** to verify the latest version.

TROUBLESHOOTING OBJECTS

This section explains how to troubleshoot issues that you might encounter while using Objects.

Handling Deployment Failure

In case you encounter any deployment failure, contact Nutanix Support at https://portal.nutanix.com/.

Shutting Down Objects VMs

This section describes the steps to perform a graceful shutdown of Objects VMs on AHV or ESXi.

Before you begin

Make sure all read and write operations are stopped from the application before shutting down the VM.

About this task

> Warning: Ensure that you only perform shutdown operations as described in this procedure. Do not perform any destructive actions such as deleting a VM.

To shut down Objects VMs, do the following:

Procedure

1. Log on to the Prism Central web console, and click **Entity Menu** > **Services** > **Objects**.
2. Copy the Object Stores cluster name.

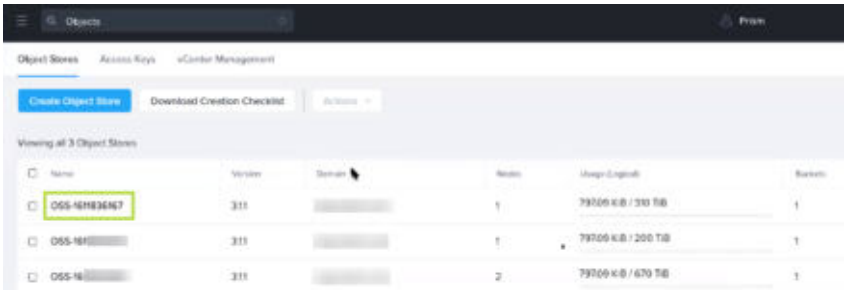

Figure 126: Object Store Cluster

3. SSH into any CVM on the Prism Element cluster where the Object Store is deployed.

4. To view the list of VMs, do any one of the following:

   » **AHV**: To view the list of VMs, run the following command.

   ```
   nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep '<objectstore-name>' -i
   ```

   For example, `nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep 'OSS-1611836167' -i`, where `OSS-1611836167` is the Object Store name.

   ```
   nutanix@NTNX-18FM6F370127-A-CVM:XX.XX.XXX.XX:~$ acli vm.list |grep 'OSS-1611836167' -i

   oss-1611836167-898db8-default-0 6596b630-cacd-48bd-90b5-b4f1698e260a

   oss-1611836167-898db8-ijpsvgbict-envoy-0 c25b5fd8-2787-46b6-b3b4-c982f6ddf1b9
   ```

   » **ESXi**: To view the list of VMs in vCenter, click **Hosts and Clusters** tab, expand the ESXi cluster.

   The list of VMs appear.

   To find the primary and secondary MSP clusters, run the command `mspctl cluster list`. If the cluster_type value is `primary_msp`, then the cluster is a primary MSP cluster. Rest clusters are secondary.

   > Note: It is recommended to shut down the envoy VMs first, and then proceed to shut down the worker VMs. Also, ensure that the primary cluster is shutdown at last as it hosts all the common services such as IAM.

5. To shut down the VMs, do any one of the following:

   » **AHV**: Run the following command:

   ```
   acli vm.shutdown <vm name>
   ```

   » **ESXi**: Right click the VM, and then click **Power** > **Power Off**.

   > Note: First shut down the envoy VMs.

   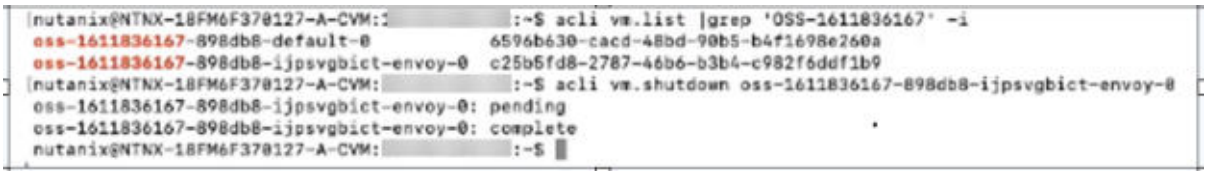

   Figure 127: Shutting Down the VMs: AHV

   For AHV, you can check the status of the VM from the Prism Central web console. Click the **Hamburger icon** > **Virtual Infrastructure** > **VMs** and confirm that the status of the VM is shown as `Off` in the **Power State** column.

   For ESXi, you can check the status of the VM from the vCenter. Click **Hosts and Clusters** tab, and expand the cluster where the VM is listed, and confirm that the status of the VM is shown as `Off`.

6. Perform Step 5 for all the VMs within the cluster.

What to do next

You can also power on Objects VMs. Refer to Powering on Objects VMs on page 224.

Powering on Objects VMs

This section describes the steps to power on the Objects VM after performing a graceful shutdown of Objects VMs on AHV or ESXi.

About this task

> Warning: Ensure that you only perform power on operations as described in this procedure. Do not perform any destructive actions such as deleting a VM.

To power on Objects VMs, do the following:

Procedure

1. Log on to the Prism Central web console, and click **Entity Menu** > **Services** > **Objects**.
2. Copy the Object Stores cluster name.

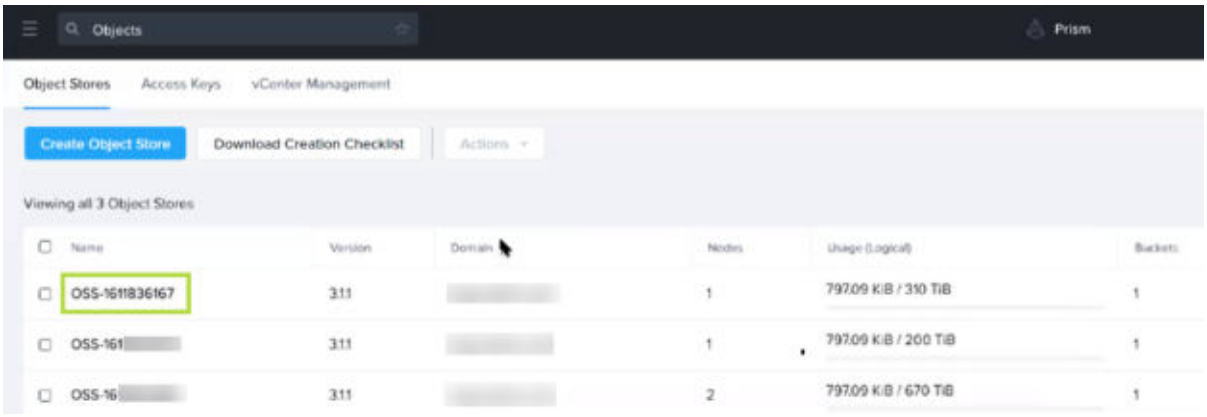

Figure 128: Object Store Cluster

3. SSH into any CVM on the Prism Element cluster where the Object Store is deployed.
4. To view the list of VMs, do any one of the following:

   » **AHV**: To view the list of VMs, run the following command.

   ```
   nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep '<objectstore-name>' -i
   ```

   For example, `nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep 'OSS-1611836167' -i`, where `OSS-1611836167` is the Object Store name.

   ```
   nutanix@NTNX-18FM6F370127-A-CVM:XX.XX.XXX.XX:~$ acli vm.list |grep 'OSS-1611836167' -i
   oss-1611836167-898db8-default-0 6596b630-cacd-48bd-90b5-b4f1698e260a
   oss-1611836167-898db8-ijpsvgbict-envoy-0 c25b5fd8-2787-46b6-b3b4-c982f6ddf1b9
   ```

   » **ESXi**: To view the list of VMs in vCenter, click **Hosts and Clusters** tab, expand the ESXi cluster.

   The list of VMs appear.

   To find the primary and secondary MSP clusters, run the command **`mspctl cluster list`**. If the cluster_type value is **`primary_msp`**, then the cluster is a primary MSP cluster. Rest clusters are secondary.

5. To power on the VMs, do any one of the following:

   » **AHV**: Run the following command:

   ```
   acli vm.on <vm name>
   ```

   » **ESXi**: Right click the VM, and then click **Power** > **Power On**.

   Note: First power on the worker VMs.

6. Check the following to ensure that the VMs are powered on.

   - For AHV, in the Prism Central web console, go to the **VMs** page. Select the VM that you powered on using the acli and perform the **Launch Console** action. You will view the login prompt if the VM is powered on.

     For ESXi, you can check the status of the VM from the vCenter. Click **Hosts and Clusters** tab, and expand the cluster where the VM is listed, and confirm that the status of the VM is shown as **On**.

   - In the Prism Central web console, go to **Entity Menu** > **Services** > **Objects**. Check for the following points for your Objects cluster:

     - Statistics are visible in the Buckets and Objects columns.
     - Click on the corresponding Objects cluster to check if it is reachable. After the Object Store page opens, check if the various statistics are visible. For example, Performance, Usage Summary, and so on.
     - In the Alerts page, ensure that there are no active alerts. If there are active alerts, wait for a few minutes. If the alerts persist, contact Nutanix Support.
     - Perform read and write operations to ensure that the object cluster is running.

Detection of Slow Connections

Nutanix Objects can detect slow performing client connections to Objects.

Any client connection must be able to read or write data at the rate of 2 MiB in every 30 second window (translates to approximately 69 KiB/s throughput). A client connection which is not able to transfer 2 MiB data in 30 second window will be treated as a slow connection.

Any slow connection will be terminated immediately to avoid Denial of Service (DoS) attacks and better manage the resources on the server.

Note: Up to 1000 active connections are accepted by each Objects endpoint. The slow performing client connections can potentially consume all slots and cause Denial of Service (DoS).

Prism Element Clusters or Subnets Not Listed in the User Interface

You might see that the **Cluster** drop down or the **Subnets** drop-down shows empty list or an error message when no Prism Element clusters or subnets are available for deployment respectively. Or you might see that no Prism Element clusters are listed during the addition of multi-cluster to the Object Store.

The following are the possible causes for the Prism Element clusters or subnets to not appear in the user interface:

- No other Prism Elements are registered to the Prism Central.

- Prism Element was just registered but the sync between the Prism Element and the Prism Central is not complete yet.
- The logged in user does not have access to any Prism Element.
- The admin has just created an access policy granting user access to a Prism Element.
- The logged in user does not have access to any subnets on the allowed Prism Element.
- The admin has just created an access policy granting user access to a Subnet.

Note: If you are just registered to Prism Element, or you have been just granted access, try again in 5 minutes.

COPYRIGHT

Copyright 2022 Nutanix, Inc.

Nutanix, Inc.
1740 Technology Drive, Suite 150
San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

AOS Version	OS Vendor	OS Name	OS Status	OS Bits	Platform
AOS 6.1	Canonical	Ubuntu 20.04.x LTS	-	64	x86
AOS 6.1	Canonical	Ubuntu 19.04.x	-	64	x86
AOS 6.1	Canonical	Ubuntu 18.04.x LTS	-	64	x86
AOS 6.1	Canonical	Ubuntu 16.04.x LTS	-	64	x86
AOS 6.1	Canonical	Ubuntu 16.04.x LTS	-	32	x86
AOS 6.1	Canonical	Ubuntu 14.04.x LTS	-	64	x86
AOS 6.1	Canonical	Ubuntu 14.04.x LTS	-	32	x86
AOS 6.1	Canonical	Ubuntu 12.04.x LTS	Legacy OS	64	x86
AOS 6.1	Canonical	Ubuntu 12.04.x LTS	Legacy OS	32	x86
AOS 6.1	CentOS	CentOS 8.3	Community	64	x86
AOS 6.1	CentOS	CentOS 8.2	Community	64	x86
AOS 6.1	CentOS	CentOS 8.1	Community	64	x86
AOS 6.1	CentOS	CentOS 8.0	Community	64	x86
AOS 6.1	CentOS	CentOS 7.9	Community	64	x86
AOS 6.1	CentOS	CentOS 7.8	Community	64	x86
AOS 6.1	CentOS	CentOS 7.7	Community	64	x86
AOS 6.1	CentOS	CentOS 7.6	Community	64	x86
AOS 6.1	CentOS	CentOS 7.5	Community	64	x86
AOS 6.1	CentOS	CentOS 7.4	Community	64	x86
AOS 6.1	CentOS	CentOS 7.3	Community	64	x86
AOS 6.1	CentOS	CentOS 7.2	Community	64	x86
AOS 6.1	CentOS	CentOS 7.1	Community	64	x86
AOS 6.1	CentOS	CentOS 7.0	Community	64	x86
AOS 6.1	CentOS	CentOS 6.10	Community	64	x86
AOS 6.1	CentOS	CentOS 6.10	Community	32	x86
AOS 6.1	CentOS	CentOS 6.9	Community	64	x86
AOS 6.1	CentOS	CentOS 6.9	Community	32	x86
AOS 6.1	CentOS	CentOS 6.8	Community	64	x86
AOS 6.1	CentOS	CentOS 6.8	Community	32	x86
AOS 6.1	CentOS	CentOS 6.7	Community	64	x86
AOS 6.1	CentOS	CentOS 6.7	Community	32	x86
AOS 6.1	CentOS	CentOS 6.6	Community	64	x86
AOS 6.1	CentOS	CentOS 6.6	Community	32	x86
AOS 6.1	CentOS	CentOS 6.5	Community	64	x86
AOS 6.1	CentOS	CentOS 6.5	Community	32	x86
AOS 6.1	CentOS	CentOS 6.4	Community	64	x86
AOS 6.1	CentOS	CentOS 6.4	Community	32	x86
AOS 6.1	CentOS	CentOS 6.3	Community	64	x86
AOS 6.1	CentOS	CentOS 6.3	Community	32	x86
AOS 6.1	CentOS	CentOS 6.2	Community	64	x86
AOS 6.1	CentOS	CentOS 6.2	Community	32	x86
AOS 6.1	CentOS	CentOS 6.1	Community	64	x86
AOS 6.1	CentOS	CentOS 6.1	Community	32	x86
AOS 6.1	CentOS	CentOS 6.0	Community	64	x86
AOS 6.1	CentOS	CentOS 6.0	Community	32	x86
AOS 6.1	CentOS	CentOS 5.11	Legacy OS	64	x86
AOS 6.1	CentOS	CentOS 5.11	Legacy OS	32	x86
AOS 6.1	Debian	Debian 8.5	Community	64	x86

AOS 6.1	Debian	Debian 9.4	Community	64	x86
AOS 6.1	Debian	Debian 10.1	Community	64	x86
AOS 6.1	Debian	Debian 10.5	Community	64	x86
AOS 6.1	Debian	Debian 10.6	Community	64	x86
AOS 6.1	Debian	Debian 10.7	Community	64	x86
AOS 6.1	Debian	Debian 10.8	Community	64	x86
AOS 6.1	Debian	Debian 10.9	Community	64	x86
AOS 6.1	FreeBSD	FreeBSD 12.1	Community	64	x86
AOS 6.1	FreeBSD	FreeBSD 11.2	Community	64	x86
AOS 6.1	FreeBSD	FreeBSD 11.1	Community	32	x86
AOS 6.1	FreeBSD	FreeBSD 11.1	Community	64	x86
AOS 6.1	FreeBSD	FreeBSD 11.0	Legacy OS	32	x86
AOS 6.1	FreeBSD	FreeBSD 11.0	Legacy OS	64	x86
AOS 6.1	FreeBSD	FreeBSD 10.4	Community	32	x86
AOS 6.1	FreeBSD	FreeBSD 10.4	Community	64	x86
AOS 6.1	FreeBSD	FreeBSD 10.3	Legacy OS	32	x86
AOS 6.1	FreeBSD	FreeBSD 10.3	Legacy OS	64	x86
AOS 6.1	FreeBSD	FreeBSD 10.2	Legacy OS	32	x86
AOS 6.1	FreeBSD	FreeBSD 10.2	Legacy OS	64	x86
AOS 6.1	FreeBSD	FreeBSD 10.1	Legacy OS	32	x86
AOS 6.1	FreeBSD	FreeBSD 10.1	Legacy OS	64	x86
AOS 6.1	FreeBSD	FreeBSD 9.3	Legacy OS	32	x86
AOS 6.1	FreeBSD	FreeBSD 9.3	Legacy OS	64	x86
AOS 6.1	Microsoft	Windows Server 2022	-	64	x86
AOS 6.1	Microsoft	Windows Server 1809	-	64	x86
AOS 6.1	Microsoft	Windows Server 2019	-	64	x86
AOS 6.1	Microsoft	Windows Server 2016	-	64	x86
AOS 6.1	Microsoft	Windows Server 2012 R2	-	64	x86
AOS 6.1	Microsoft	Windows Server 2012	-	64	x86
AOS 6.1	Microsoft	Windows Server 2008 R2	Legacy OS	64	x86
AOS 6.1	Microsoft	Windows Server 2008 R2 SP1	Legacy OS	64	x86
AOS 6.1	Microsoft	Windows Server 2008	Legacy OS	64	x86
AOS 6.1	Microsoft	Windows Server 2008	Legacy OS	32	x86
AOS 6.1	Microsoft	Windows 10	-	32	x86
AOS 6.1	Microsoft	Windows 10	-	64	x86
AOS 6.1	Microsoft	Windows 8.1 Update	-	64	x86
AOS 6.1	Microsoft	Windows 7	Legacy OS	32	x86
AOS 6.1	Microsoft	Windows 7	Legacy OS	64	x86
AOS 6.1	Microsoft	Windows Server 2003	Legacy OS	64	x86
AOS 6.1	NeoKylin	NeoKylin OS 7.0	-	64	x86
AOS 6.1	Oracle	Oracle Linux 8.4	-	64	x86
AOS 6.1	Oracle	Oracle Linux 8.3	-	64	x86
AOS 6.1	Oracle	Oracle Linux 8.2	-	64	x86
AOS 6.1	Oracle	Oracle Linux 8.1	-	64	x86
AOS 6.1	Oracle	Oracle Linux 8.0	-	64	x86
AOS 6.1	Oracle	Oracle Linux 7.9	-	64	x86
AOS 6.1	Oracle	Oracle Linux 7.8	-	64	x86
AOS 6.1	Oracle	Oracle Linux 7.7	-	64	x86
AOS 6.1	Oracle	Oracle Linux 7.6	-	64	x86
AOS 6.1	Oracle	Oracle Linux 7.5	-	64	x86

AOS 6.1	Oracle	Oracle Linux 7.4	-	64	x86
AOS 6.1	Oracle	Oracle Linux 7.3	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.10	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.10	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.9	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.9	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.8	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.8	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.7	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.7	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.6	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.6	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.5	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.5	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.4	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.4	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.3	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.3	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.2	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.2	-	32	x86
AOS 6.1	Oracle	Oracle Linux 6.1	-	64	x86
AOS 6.1	Oracle	Oracle Linux 6.1	-	32	x86
AOS 6.1	Rancher	RancherOS v1.5.5	Community	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 8.5	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 8.4	-	64	x86
AOS 6.1	Red Hat	Fedora CoreOS 32.20201104.3.0	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 8.3	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 8.2	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 8.1	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 8.0	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.9	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.8	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.7	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.6	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.5	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.4	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.3	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.2	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.1	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 7.0	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.10	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.10	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.9	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.9	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.8	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.8	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.7	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.7	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.6	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.6	-	32	x86

AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.5	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.5	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.4	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.4	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.3	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.3	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.2	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.2	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.1	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.1	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.0	-	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 6.0	-	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 5.11	Legacy OS	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 5.11	Legacy OS	32	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 4.8	Legacy OS	64	x86
AOS 6.1	Red Hat	Red Hat Enterprise Linux 4.8	Legacy OS	32	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 15 SP 3	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 15 SP 2	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 15 SP 1	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 15	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 12 SP 4	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 12 SP 3	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 12 SP 2	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 12 SP 1	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 12	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 11 SP 4	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 11 SP 3	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 11 SP 2	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 11	-	32	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 11	-	64	x86
AOS 6.1	SUSE	SUSE Linux Enterprise Server 11 SP 3	-	64	x86
AOS 6.1	AlmaLinux	AlmaLinux 8.3	-	64	x86
AOS 6.1	RockyLinux	RockyLinux 8.4	-	64	x86

STORAGE	STORAGE	STORAGE	GRAPHICS	GRAPHICS	SUPPORTED FEATURES
SCSI	IDE	PCI	GPU	vGPU	Hot Add Memory
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	Supported	-
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	-
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	-
Recommended	Supported	Supported	-	-	Supported

Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	-	-	Supported
Recommended	Supported	-	-	-	-
Recommended	Supported	-	Supported	Supported	-
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	-
Recommended	Supported	-	Supported	Supported	Supported
Recommended	Supported	-	Supported	Supported	Supported
-	Recommended	-	-	-	-
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported

Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
-	Supported	Recommended	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
-	Supported	Recommended	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
-	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	-
-	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	-
-	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	-
-	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	Supported	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-

Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	-
-	Supported	Recommended	Supported	-	Supported
-	Supported	Recommended	Supported	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	-	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	Supported
Recommended	Supported	Supported	Supported	-	-
Recommended	Supported	Supported	Supported	-	-
-	Supported	Recommended	-	-	-
-	Supported	Recommended	-	-	-
-	Supported	Recommended	-	-	-
Recommended	Supported	Supported	-	-	-
Recommended	Supported	Supported	Supported	Supported	Supported
Recommended	Supported	Supported	Supported	Supported	Supported

SUPPORTED FEATURES	SUPPORTED FEATURES	SUPPORTED FEATURES
Hot Add vCPU	UEFI	Secure Boot
Supported	Supported	Supported
Supported	-	-
Supported	Supported	-
-	Supported	-
-	-	-
-	-	-
-	-	-
-	Supported	-
-	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	Supported	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	Supported	-
Supported	Supported	-
Supported	Supported	Supported
Supported	-	-
Supported	-	-
-	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	-	-
-	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
Supported	-	-

Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
Supported	Supported	-
Supported	Supported	Supported
-	Supported	Supported
Supported	Supported	Supported
Supported	Supported	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	Supported	Supported
-	-	-
-	-	-
-	-	-
-	-	-
Supported	-	-
Supported	Supported	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	-	-
-	-	-
Supported	-	-

Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	Supported	Supported
Supported	Supported	Supported
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	Supported	-
Supported	Supported	Supported
Supported	Supported	Supported
Supported	Supported	Supported
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	Supported	-
-	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-

Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	-	-
Supported	-	-
Supported	-	-
Supported	Supported	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	-	-
Supported	Supported	-
Supported	-	-
Supported	-	-
Supported	-	-
-	-	-
-	-	-
-	-	-
Supported	-	-
Supported	-	-
-	-	-
Supported	-	-
Supported	-	-

Prism 6.0

Prism Central Guide

March 9, 2022

Contents

Welcome viii

Help Organization ix

1. Prism Central Overview 10
- Main Menu 10
- Entities Menu 13
- Settings Menu 17
- Searching for Information 20
- Understanding Displayed Statistics 29

2. Prism Central Administration 31
- Installing or Upgrading 31
- Licensing 31
- Prism Central Management 31
  - Logging Into Prism Central 32
  - Logging Out of Prism Central 35
  - Managing Prism Central 35
  - Finding the Prism Central Version 42
  - Finding the AHV Version on Prism Central 43
  - Finding the AOS Version Using Prism Central 44
  - Register (Unregister) Cluster with Prism Central 45
  - Expanding (Scale Out) Prism Central 52
  - Shutdown (Start Up) Prism Central VM 56
  - IP Address Reconfiguration 56

3. Main Dashboard 60
- Main Dashboard Display 62
- Widgets on a Dashboard 64
- Creating a New Dashboard 69
- Adding Dashboard Widgets 70
- Modifying a Dashboard 75

4. Entity Exploring 77
- Virtual Infrastructure Entities 82
  - VMs Summary View 82
  - Storage Containers Summary View 97
  - Volume Groups Summary View 108
  - vCenter Datastores Summary View 117
  - Catalog Items View 124
  - Images Summary View 125
  - Categories Summary View 128
  - OVAs View 131
- Network and Security Entities 135

Subnets View....135
Policies Entities....136
Security Policies Summary View....137
NGT Policies View....139
Image Placement Policies Summary View....140
Data Protection Entities....142
Protection Policies View....142
Recovery Plans View....143
VM Recovery Points View....144
Hardware Entities....145
Clusters Summary View....146
Hosts Summary View....165
Disks Summary View....175
GPUs Summary View....183
Activity Entities....186
Audits Summary View....186
Tasks View....190
Operations Entities....193
Reports View....194
Administration Entities....195
Users Summary View....195
Roles Summary View....199
Projects Summary View....203
Availability Zones View....208
LCM View....209
Services Entities....209

5. Alert and Event Monitoring....210
Alerts Summary View....210
Alert Details....218
Events Summary View....221
Event Details....222
Configuring Alert Emails....223
Configuring Alert Policies....226
Modifying System Alert Policies....227
Adding Custom Alert Policies....229
Viewing External Alert Policies....244
Stopping Alerts During a Maintenance Window....245
Prism Central Logs....247
Audit Log Events....250
Alerts/Health checks....252
Cluster....252
Controller VM....264
DR....270
License....277
Guest VM....277
Node....278
Network....279
Other....281
Prism Central VM....286
Storage....286
System Indicator....287

6. System Management....288

Configuring Prism Central When a Cluster Uses Proxy Servers ... 288
Who Needs to Use the Whitelist Method ... 289
Configuring an HTTP Proxy ... 290
Whitelisting Prism Central and Its Managed Clusters (nCLI) ... 295
Configuring Name Servers ... 296
Configuring NTP Servers ... 296
Configuring an SMTP Server ... 297
Configuring SNMP ... 298
Configuring Syslog Monitoring ... 304
Syslog Modules ... 305
Modifying UI Settings ... 308
Configuring a Banner Page ... 311
Internationalization (i18n) (Prism Central) ... 311
Localization (L10n) (Prism Central) ... 312
Changing the Language Settings (Prism Central) ... 312

7. Security and User Management ... 315

8. Virtual Infrastructure (Cluster) Administration ... 316
VM Management ... 316
Creating a VM (AHV) ... 317
Creating a VM (Self Service) ... 327
Managing a VM (AHV) ... 332
Managing a VM (Self Service) ... 342
Adding Multiple vGPUs to the Same VM ... 350
Migrating Live a VM with Virtual GPUs ... 354
Creating a VM (ESXi) ... 357
Managing a VM (ESXi) ... 362
Exporting a VM as an OVA ... 368
NGT Management in Prism Central ... 369
Storage Quality of Service (QoS) ... 376
Storage Management ... 379
Creating a Volume Group ... 380
Modifying a Volume Group ... 384
Managing Volume Group Connections ... 385
Managing Volume Group Virtual Disks ... 389
Catalog Management ... 390
Adding a Catalog Item ... 391
Deleting a Catalog Item ... 393
Image Management ... 394
Requirements ... 395
Limitations ... 395
Creating an Image ... 395
Modifying an Image ... 411
Importing Images to Prism Central ... 412
Category Management ... 413
Creating a Category ... 414
Modifying a Category ... 415
Assigning a Category ... 415
Affinity Policies for AHV ... 416
Operations Policy Management ... 417
Operations Policies View ... 417
Modifying an Operations Policy ... 418
Enabling an Operations Policy ... 420

Network and Security Management 421
Configuring Network Connections 423
Creating or Updating Virtual Switch 426
Deleting a Virtual Switch 431
Enabling LAG and LACP on the ToR Switch (AHV Only) 431
Modifying Network Connections 432
OVA Management 433
Exporting a VM as an OVA 433
Uploading an OVA 434
Concatenating Upload using APIs 440
Deploying an OVA as VM 442
Downloading an OVA 448
Renaming an OVA 448
Deleting an OVAs 449
Affinity Policies for AHV 450
Performing Other Administrative Tasks 451
vCenter Server Integration 452
Registering vCenter Server (Prism Central) 453
Managing vCenter Server Registration Changes (Prism Central) 456

9. Prism Self Service Administration 457
Prism Self Service Overview 457
Configuring Prism Self Service 458
Migrating Prism Self Service 461
Project Management 462
Creating a Project 462
Modifying a Project 465

10. Network and Security Management 466
Configuring Network Connections 468
Creating or Updating Virtual Switch 470
..... 471
Deleting a Virtual Switch 475
Enabling LAG and LACP on the ToR Switch (AHV Only) 475
Modifying Network Connections 476

11. Policies Management 478
NGT Policies 478
Creating a New Restart Policy 478
Modifying a Restart Policy 480
Image Placement Policies 480
Configuration Workflow 481
Associating a Registered Nutanix Cluster with Categories 481
Associating an Image with Categories 481
Configuring an Image Placement Policy 482
Updating an Image Placement Policy 483
How Prism Central Handles Conflicting Policies 483
Sample Scenarios and Configurations 484

12. Operations Management 486
Performance Monitoring 486
Analysis Dashboard 486

Sessions Dashboard....512
Closing the Session....518
Switching Between Sessions....518
Application Discovery....518
Application Discovery View....519
Enabling Application Discovery....526
Connecting to the Cloud....532
Adding (Removing) Clusters to Monitor....533
Configuring vCenter Authentication....534
Discovering Applications....535
Publishing Applications....536
Creating Discovery Policies....536
Cost Management (Xi Beam)....539
Application Monitoring....540
Application Monitoring Summary View....540
Configuring Application Monitoring....553
Modifying Application Monitoring....556
Operations Policy Management....556
Operations Policies View....557
Modifying an Operations Policy....558
Enabling an Operations Policy....560
Resource Planning....560
Capacity Runway View....561
Scenarios View....563
Updating Capacity Configurations....581
Behavioral Learning Tools....582
Task Automation....585
Overview....586
Action Gallery....586
Plays....586
Playbook Actions....587
Creating Playbooks using Triggers....596
Running a Playbook (Manual Trigger)....606
Configuring Manual Parameters....606
Exporting or Importing Playbooks....610
Integrating Slack....611
ServiceNow Integration with Prism Central (X-Play Support)....611
Reports Management....617
Creating a New Report....619
Managing a Report....620
Viewing Report Instances....623
Views in the Reports....626
Scheduling a Report....643
Downloading Reports....645
Configuring Report Settings....645
Role Based Access Control....649

13. Services Enablement....651
Enabling Calm....651
Enabling Files....652
Enabling Foundation Central....652
Enabling Karbon....652
Enabling Objects....653

14. Customer Support Services 654
Pulse Health Monitoring 654
Configuring Pulse 656
Prism Central Proxy for Pulse Data 659
Creating a Support Case 659
Viewing Case Status 665
Accessing the Nutanix Support Portal 666
Accessing the REST API Explorer 669

15. Help Resources 671
Accessing Online Help 671
Accessing the Nutanix Next Community 673
Glossary 674

Copyright 680
License 680
Conventions 680
Version 680

Booklist 681

WELCOME

You can monitor and manage entities across Nutanix clusters through a web console called Prism Central. This documentation describes how to use Prism Central.

- For a list and description of major topics, see Help Organization on page ix.
- For help with viewing this documentation, see Accessing Online Help on page 671.
- To access other Nutanix documents, see Nutanix support portal.

HELP ORGANIZATION

This documentation is organized as follows:

- Prism Central Overview on page 10 provides an overview of Prism Central including menu options and navigation tools.
- Prism Central Administration on page 31 describes how to install and manage Prism Central; how to upgrade Prism Central, NCC, and individual clusters through Prism Central; and how to manage licenses.
- Main Dashboard on page 60 describes the main (home) dashboard and how to customize it.
- Entity Exploring on page 77 describes how to use the entity browser to view detailed information about VMs, containers, disks, and other entities across the registered clusters.
- Resource Planning on page 560 describes how to review and analyze current and potential resource needs in a Nutanix cluster.
- Performance Monitoring on page 486 describes how to monitor and analyze performance in a cluster.
- Services Enablement on page 651 describes how to enable certain applications (Calm, Karbon, and Objects) through Prism Central.
- Alert and Event Monitoring on page 210 describes how to monitor alerts and events across the registered clusters and how to configure alert policies and notification.
- System Management on page 288 describes how to configure various system settings (NTP, SMTP, and so on) for Prism Central.
- Security and User Management on page 315 describes how to configure user authentication, manage local user accounts, control user access using roles, install SSL certificates, and control remote (SSH) access.
- Virtual Infrastructure (Cluster) Administration on page 316 describes how to administer individual clusters through Prism Central such as how to create and manage VMs.
- Prism Self Service Administration on page 457 describes how to implement self-service features through Prism Central.
- Policies Management on page 478 describes how to create and apply policies to manage and control access and operations in a cluster.
- Reports Management on page 617 describes how to create and run reports about infrastructure resources.
- Task Automation on page 585 describes how to automate routine administrative tasks through Prism Central by using the X-play feature.
- Customer Support Services on page 654 describes how to configure Pulse (system diagnostics reporting feature), access the Nutanix customer support portal, and create a support ticket through Prism Central when you need help.
- Help Resources on page 671 describes how to use the online help to find what you need and how to access the Nutanix Next Community portal.

1

PRISM CENTRAL OVERVIEW

Nutanix provides an option to monitor and manage multiple clusters through a single web console. This multi-cluster view, known as Prism Central, is a centralized management tool that runs as a separate instance comprised of either a single VM or a set of VMs. Prism Central provides the following features:

- Single sign on for all registered clusters.
- Customizable main dashboard that displays summary information across the registered clusters (see Main Dashboard on page 60).
- Virtual infrastructure dashboards with drill-down options to manage VMs, storage containers, catalog items, images, categories, recoverable entities, and subnets across the registered clusters (Virtual Infrastructure Entities on page 82).
- Dashboards to manage security policies, protection policies, recovery plans, Nutanix Guest Tools (NGT) policies, and image placement (see Policies Entities on page 136).
- Hardware component dashboards with drill-down options to view detailed information about individual clusters, hosts, disks, and GPUs across the registered clusters (see Hardware Entities on page 145).
- Activity monitors for alerts, events, audits, and tasks (see Activity Entities on page 186).
- Tools to analyze system activity, plan for resource needs, create usage reports, and automate routine administrative tasks (see Operations Entities on page 193).
- Dashboards to manage inventory (LCM), projects, roles, users, and availability zones (see Administration Entities on page 195).
- Service connection pages to enable Calm, Karbon, and Buckets (see Services Entities on page 209).
- Settings menu from which you can configure Prism Central functions (see Settings Menu on page 17).

Main Menu

The main menu at the top of every screen provides access to the features of Prism Central. This section describes each of the main menu options.

Viewing Choices

Clicking the collapse menu button (also know as a "hamburger" button) in the far left of the main menu displays the entities menu (see Entities Menu on page 13). The Main dashboard is the first screen that appears after logging into Prism Central (see Main Dashboard on page 60).

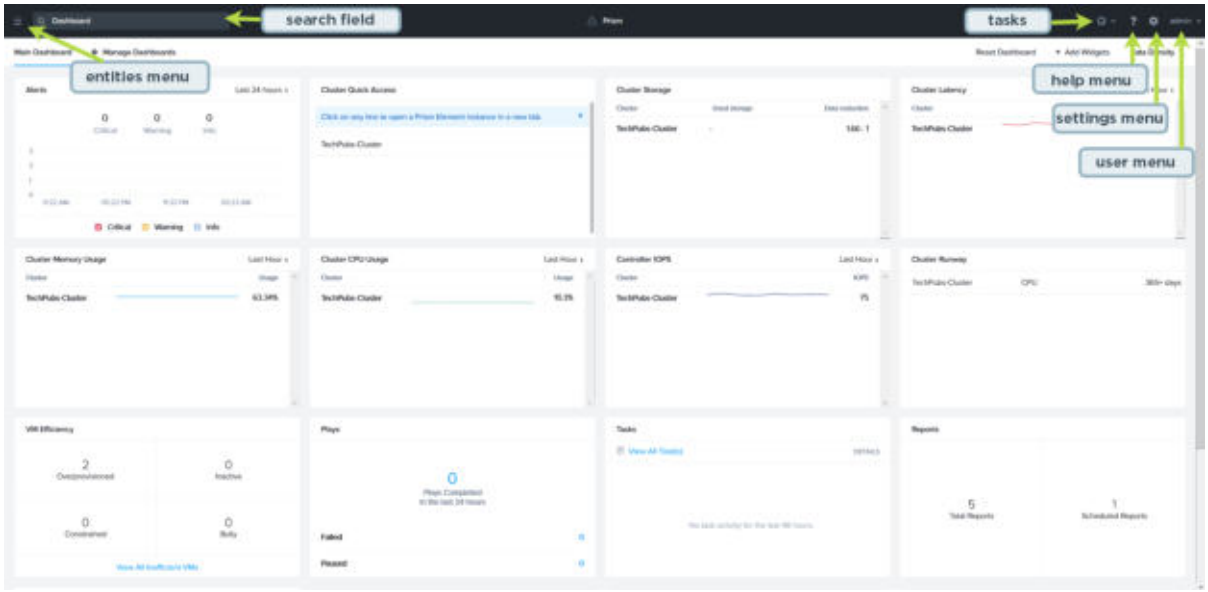

Figure 1: Prism Central Main Menu

Search, Alert, and Task Options

The main menu provides status information about active alerts and tasks plus access to advanced search capability:

- A field with a magnifying glass icon appears on the left of the main menu. Enter a string in this field to search for relevant content in Prism Central (see Searching for Information on page 20).
- An alerts icon appears on the right of the main menu when there are current critical (red), warning (yellow), or informational (gray) alert messages. The number of active alerts is displayed in the icon. Click the alert icon to display information about those alerts in the Alerts dashboard (see Alerts Summary View on page 210).
- A tasks icon appears to the right of the alerts when there are current tasks (running or completed within the last 48 hours). The icon is blue when all the tasks are running properly or have completed successfully. If one or more of those tasks do not complete successfully, the icon turns yellow (warning) or red (failure) . Click the tasks icon to see a drop-down list of the current tasks; click **View All Tasks** at the bottom of the list to display the tasks dashboard (see Tasks View on page 190).

Help Menu ("?" Icon)

A question mark icon appears on the right side of the main menu. Clicking the question mark displays a list of help resource options that you can select. The following table describes each option in the pull-down menu.

Table 1: Help Menu Options

Name	Description
Learn about search	Displays search guidelines (see Searching for Information on page 20).

Name	Description
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help on page 671).
Online Documentation	Opens the online help at the introduction page (see Accessing Online Help on page 671).
Support Portal	Opens a new browser tab (or window) at the Nutanix Support portal logon page (see Accessing the Nutanix Support Portal on page 666).
Nutanix Next Community	Opens a new browser tab (or window) at the Nutanix Next Community entry page (see Accessing the Nutanix Next Community on page 673). The portal is an online community site for customers and partners to exchange ideas, tips, and information about Nutanix technologies and related data center topics.
Create Support Case	Opens the **Create a new support case** page to view or create support cases with Nutanix customer support (see Creating a Support Case on page 659).

Settings Menu (Gear Icon)

A gear icon appears on the right side of the main menu. Clicking the gear icon launches the **Settings** page and displays a menu of tasks you can perform (see Settings Menu on page 17).

User Menu (<user_name> Icon)

A user icon appears on the far right side of the main menu with the current user logon name. Clicking the user icon displays a list of options to update your user account, log off from Prism Central, and other miscellaneous tasks. The following table describes each option in the pull-down menu.

Table 2: User Menu Options

Name	Description
Change Password	Opens the **Change Password** window to update your password (see “” ()).
Update Profile	Opens the **Update Profile** window to update your user name and email address (see “” ()).
Download Cmdlets Installer	Downloads the PowerShell installer for the Nutanix cmdlets. See the Nutanix *Command Reference* for information about installing the cmdlets locally and for cmdlet descriptions.
Download nCLI	Downloads the Nutanix command-line interface (nCLI) as a zip file to your local system. The download occurs immediately after clicking this option (no additional prompts). See the Nutanix *Command Reference* for information about installing the nCLI locally and for nCLI command descriptions.

Name	Description
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page (see Accessing the REST API Explorer on page 669).
About Nutanix	Opens the **About Nutanix** window that displays Nutanix operating system (AOS) and other version information (see Finding the Prism Central Version on page 42).
Nothing To Do?	Opens a game that is strictly for entertainment. To quit the game, click the "X" at the upper right of the screen.
Sign Out	Logs you out of Prism Central (see Logging Out of Prism Central on page 35).
Adjust Contrast (Chrome only)	Displays a contrast setting box at the bottom of the screen where you can set the Prism Central display to **Normal** (default) or **High** contrast. This option is available only when using the Chrome browser.

Entities Menu

Clicking the collapse menu ("hamburger") button in the main menu (see Main Menu on page 10) displays a menu of entities on the left. The following table describes each entity in the menu.

- The primary menu options are **Dashboard**, seven category names (**Virtual Infrastructure**, **Network & Security**, **Policies**, **Hardware**, **Activity**, **Operations**, **Administration**, and **Services**), and **Global Settings**. Clicking a category name displays a submenu of entities in that category.
- Clicking the entity name displays a dashboard or other appropriate page for that entity.
- For quicker access to one or more submenu entities, you can promote a submenu option to the primary menu by bookmarking that entity (clicking the star icon to the right of the name). This displays that entity name below the **Dashboard** entry. You can remove the bookmarked entry by clicking the star again.
- The entity menu pane closes automatically when you click in the display area to the right. To keep the entity menu pane open all the time, click the **Lock Navigation Bar** button at the bottom. The button is a toggle; to unlock the navigation bar, click the button again.
- When the entities menu is displayed, the collapse menu button changes to an X button. Click the X button to hide the menu.

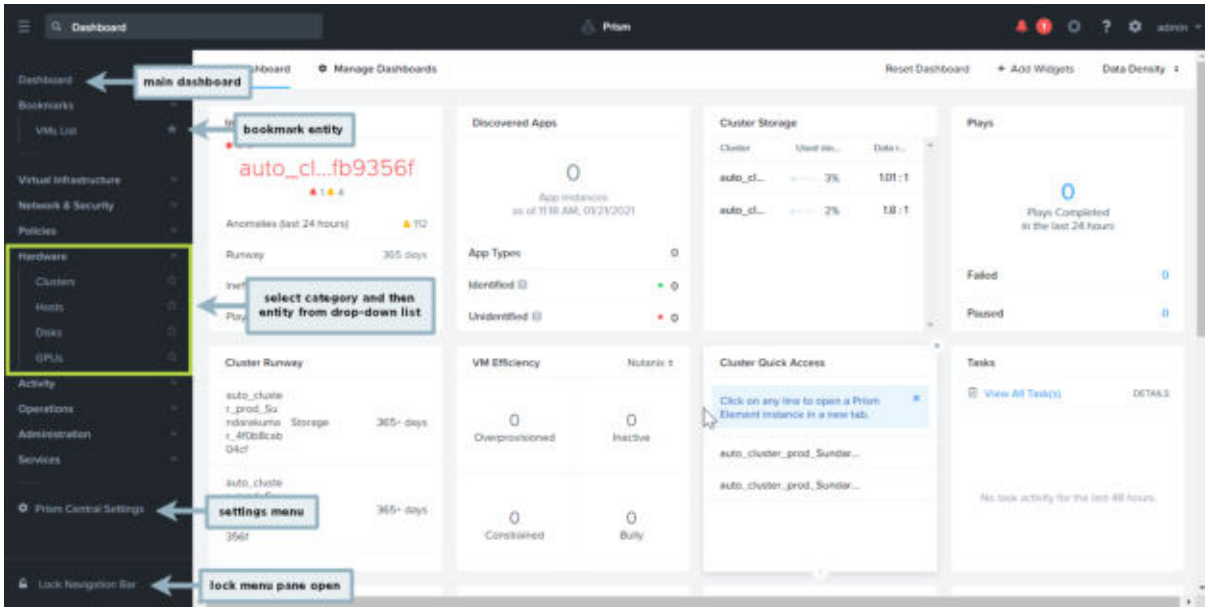

Figure 2: Entities Menu

Table 3: Entities Menu Options

Category	Name	Description
Dashboard		Displays the main dashboard (see Main Dashboard on page 60).
Virtual Infrastructure	VMs	Displays the VMs dashboard (see VMs Summary View on page 82).
	Storage Containers	Displays the Storage Containers dashboard (see Storage Containers Summary View on page 97).
	Volume Groups	Displays the Volume Groups dashboard (see Volume Groups Summary View on page 108).
	vCenter Datastores	Displays the vCenter Datastores dashboard (see vCenter Datastores Summary View on page 117). This option appears when there is at least one registered vCenter instance.
	Catalog Items	Displays the Catalog dashboard (see Catalog Items View on page 124).
	Images	Displays the Images dashboard (see Images Summary View on page 125).
	Categories	Displays the Categories dashboard (see Categories Summary View on page 128).
	OVAs	Displays the OVAs dashboard (see OVAs View on page 131).
Network & Security	Subnets	Displays the subnets dashboard (see Subnets View on page 135).

Category	Name	Description
	Virtual Private Clouds	Displays the virtual private clouds dashboard (see Network and Security Entities on page 135).
	Floating IPs	Displays the floating IPs dashboard (see Network and Security Entities on page 135).
	VPNs	Displays the VPNs dashboard (see Network and Security Entities on page 135).
Policies	Security	Displays the Security Policies dashboard (see Security Policies Summary View on page 137).
	NGT Policies	Displays the NGT Policies dashboard (see NGT Policies View on page 139).
	Image Placement	Displays the Image Placement Policies dashboard (see Image Placement Policies on page 480).
Data Protection	Protection Summary	Displays the protection summary dashboard.
	Protection Policies	Displays the Protection Policies dashboard (see Protection Policies View on page 142).
	Recovery Plans	Displays the Recovery Plans dashboard (see Recovery Plans View on page 143).
	VM Recovery Points	Displays the VM Recovery Points dashboard (see VM Recovery Points View on page 144).
Hardware	Clusters	Displays the Clusters dashboard (see Clusters Summary View on page 146).
	Hosts	Displays the Hosts dashboard (see Hosts Summary View on page 165).
	Disks	Displays the Disks dashboard (see Disks Summary View on page 175).
	GPUs	Displays the GPUs dashboard (see GPUs Summary View on page 183).
Activity	Alerts	Displays the Alerts dashboard (see Alerts Summary View on page 210)
	Events	Displays the Events dashboard (see Events Summary View on page 221).
	Audits	Displays the Audits dashboard (see Audits Summary View on page 186).
	Tasks	Displays the Tasks dashboard (see Tasks View on page 190).
Operations	Analysis	Displays the Analysis dashboard (see Analysis Dashboard on page 486).
	App Discovery	Displays the App Discovery dashboard or the enable wizard if application discovery is not enabled (see Application Discovery on page 518).

Category	Name	Description
	Cost Management	Provides access to the Xi Beam site (see Cost Management (Xi Beam) on page 539).
	Integrations	Displays the Monitoring Integrations dashboard or the enable wizard if monitoring integrations is not enabled (see Application Monitoring on page 540).
	Ops Policies	Displays the Operations Policies dashboard (see Operations Policy Management on page 417).
	Planning	Displays the capacity Planning dashboard (see Resource Planning on page 560).
	Playbooks	Displays the list of playbooks (see Task Automation on page 585).
	Reports	Displays the Reports dashboard (see Reports View on page 194).
Administration	LCM	Displays the life cycle manage (LCM) dashboard (see LCM View on page 209).
	Projects	Displays the Projects dashboard (see Projects Summary View on page 203).
	Roles	Displays the Roles dashboard (see Roles Summary View on page 199).
	Users	Displays the Users dashboard (see Users Summary View on page 195).
	Availability Zones	Displays the Availability Zones dashboard (see Availability Zones View on page 208).
Services	Calm	Displays the Calm feature page or the enable wizard if Calm is not enabled(see Enabling Calm on page 651).
	Files	Displays the Files dashboard or the enable wizard if Files is not enabled (see Enabling Files on page 652).
	Foundation Central	Displays the Foundation Central dashboard or the enable wizard if Foundation Central is not enabled (Enabling Foundation Central on page 652).
	Karbon	Displays the Karbon feature page or the enable wizard if Karbon is not enabled (see Enabling Karbon on page 652).
	Objects	Displays the Objects feature page or the enable wizard if Objects is not enabled (see Enabling Objects on page 653).
Prism Central Settings		Displays the settings menu (see Settings Menu on page 17).

Settings Menu

Prism Central includes a Settings page from which you can configure a variety of system services. You can access the Settings page by doing either of the following:

- Click the gear icon in the main menu (see Main Menu on page 10).
- Click the collapse menu ("hamburger") button and then select **Prism Central Settings** from the entities menu (see Entities Menu on page 13).

The Settings page displays a menu of tasks (on the left) you can perform. Click the task to open the window or page for that option in the pane to the right. The following table describes each menu option.

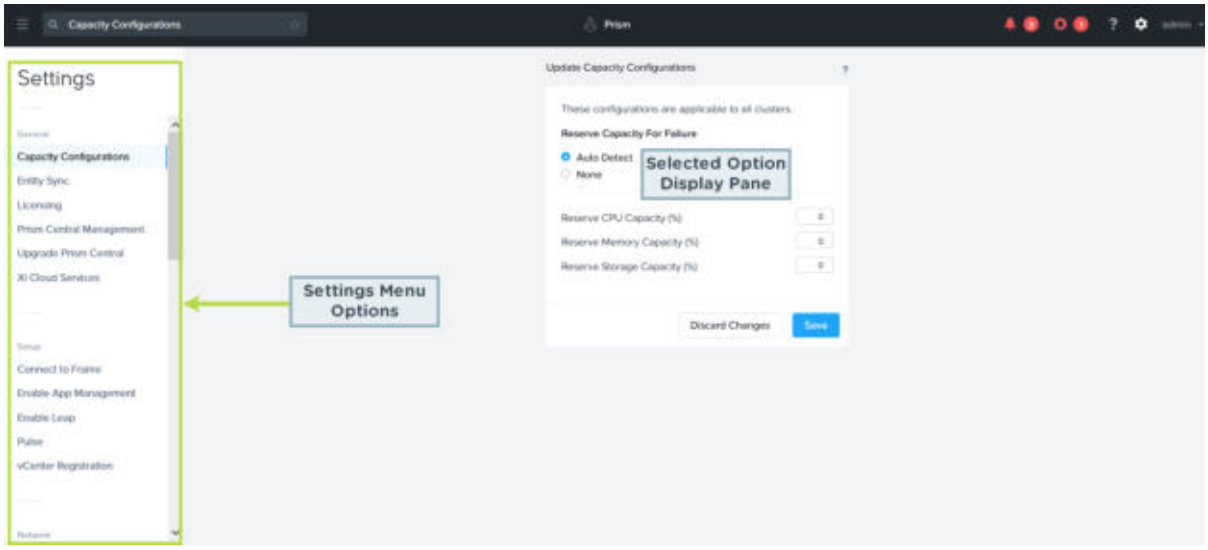

Figure 3: Settings Page

Table 4: Settings Menu

Category	Name	Description
General	Capacity Configurations	Opens the **Update Capacity Configurations** window, which allows you to configure sizing rules for cluster runway analysis and alerting (see Updating Capacity Configurations on page 581). This menu option does not appear if Prism Pro is disabled.
	Entity Sync	Opens the **Force Entity Sync** window, which synchronizes entities with the specified availability zones. See the *Leap Administration Guide* for more information.
	Licensing	Opens the **Licensing** page to install or update the Prism Central license that enables entitled Nutanix features (see Licensing on page 31).

Category	Name	Description
	Migrate SSP	Opens the **Migrate Prism Self Service to Prism Central** window to migrate a pre-AOS 5.5 Prism Self Service (hosted in Prism Element) to Prism Central (see Migrating Prism Self Service on page 461).
	Prism Central Management	Opens the **Manage Prism Central** page for viewing information about Prism Central and optionally expanding (scale out) Prism Central across multiple VMs (see Managing Prism Central on page 35).
	Upgrade Prism Central	Opens the **Upgrade Prism Central** window to upgrade the Prism Central VM to a newer version (see Installing or Upgrading on page 31).
	Xi Cloud Services	Opens the **Xi Cloud Services** windows, which allows you to log in to the Xi portal. See the *Xi Cloud Services Administration Guide* for more information.
Setup	Connect to Frame	Opens the **Connect to Frame** page to connect to Xi Frame. See the Xi Frame documentation for more information.
	Enable App Management	Opens the **Enable App Management** page to enable the Nutanix Calm features (see Enabling Calm on page 651).
	Enable Leap	Opens the **Leap** page to enable the Disaster Recovery as a Service (DRaaS) capability. See the *Leap Administration Guide* for more information.
	Pulse	Opens the **Pulse** page to enable (disable) the Nutanix automated support system (see Configuring Pulse on page 656).
	Self-Service Admin Management	Opens the Self-Service Admin Management page to modify the self-service configuration (see Configuring Prism Self Service on page 458). This menu option does not appear until a directory service is configured.
	vCenter Registration	Opens the **vCenter Registration** window to register (or unregister) clusters with vCenter (see Registering vCenter Server (Prism Central) on page 453).
Network	HTTP Proxy	Opens the **HTTP Proxies** window to configure an HTTP proxy to which the Nutanix software can connect (see Configuring an HTTP Proxy on page 290).
	Name Servers	Opens the **Name Servers** window to configure name servers (see Configuring Name Servers on page 296).
	NTP Servers	Opens the **NTP Servers** window to specify which NTP servers to access (see Configuring NTP Servers on page 296).

Category	Name	Description
	SNMP	Opens the **SNMP Configuration** window to configure SNMP (see Configuring SNMP on page 298).
Flow	ID Based Security	Opens the **ID Based Security** page to add Active Directory domain services configurations and to import user groups for identity-based security policies. See the *Nutanix Flow Guide* for more information.
	Microsegmentation	Opens the **Microsegmentation** page to enable this feature, which is disabled by default. Before you can configure and use application security policies, isolation environment policies, and quarantine policies, you must enable the feature. See the *Nutanix Flow Guide* for more information.
Security	Cluster Lockdown	Opens the **Cluster Lockdown** window, which allows you to delete (or add) public authentication keys used for SSH access into Prism Central. Removing all public keys locks down Prism Central from external access. For details, see “” () .
	SSL Certificate	Opens the **SSL Certificates** window to create a self-signed certificate. For details, see “” () .
User and Roles	Authentication	Opens the **Authentication Configuration** window to configure Prism Central authentication. For details, see “” () .
	Local User Management	Opens the **Local User Management** window. This window lists current users and allows you to add, update, and delete local user accounts. For details, see “” () .
	Role Mapping	Opens the **Role Mapping** window to configure role mappings that apply in the user authentication process. For details, see “” () .
Alerts and Notifications	Alert Email Configuration	Opens the **Alert Email Configuration** window, which allows you to configure rules and templates that apply to the alert messages from selected clusters (see Configuring Alert Emails on page 223).
	SMTP Server	Opens the **SMTP Server Settings** window to configure an SMTP server (see Configuring an SMTP Server on page 297).
	Syslog Server	Opens the **Syslog Servers** window to configure a syslog server (see Configuring Syslog Monitoring on page 304).
Appearance	Language Settings	Opens the **Language Settings** window to localize the Prism Central screen text (see Changing the Language Settings (Prism Central) on page 312).

Category	Name	Description
	UI Settings	Opens the **UI Settings** window to disable (or re-enable) the login screen background video and configure the automatic logout timeout (see Modifying UI Settings on page 308).
	Welcome Banner	Opens the **Edit Welcome Banner** window to create a welcome banner message that appears before users login to Prism Central (see Configuring a Banner Page on page 311).

Searching for Information

The main menu includes a search field on the left (see Main Menu on page 10) that allows you to find information about selected entities in various ways. If you need help at any time to navigate Prism Central or apply a search filter, you can click the question mark icon in the main menu and select **Learn about search**. This displays the Search Guidelines page that explains the search rules and options.

Search Basics

An entity is an object type such as a VM, cluster, node, security policy, project, report, event, alert, or audit. The search field is context-sensitive, which means it is populated automatically based on where you are in Prism Central. The first screen after logging on is the main dashboard, so "Dashboard" appears in the search field initially.

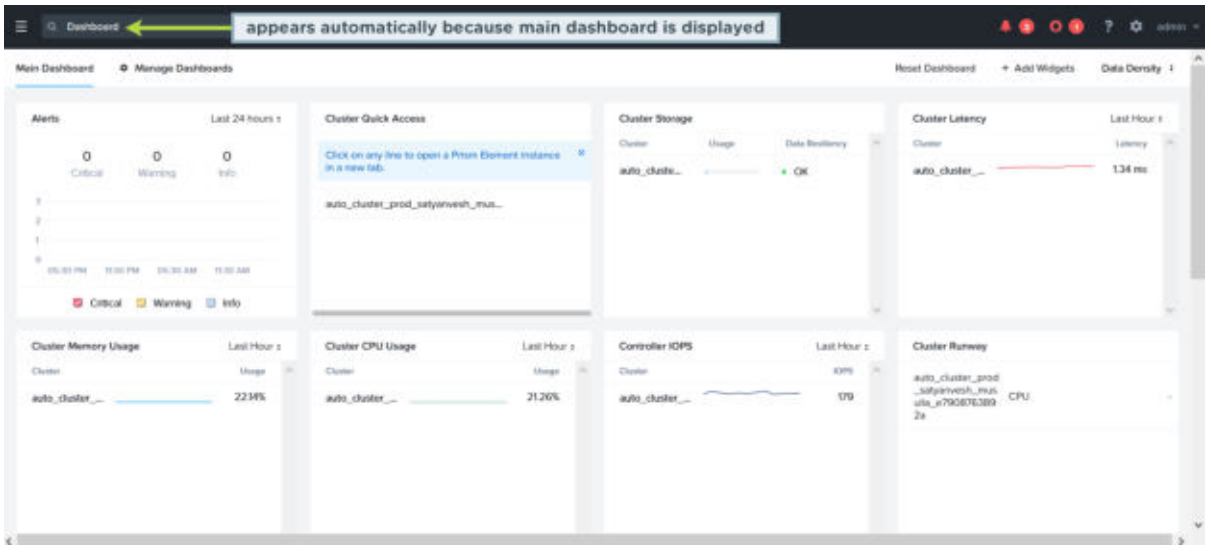

Figure 4: Search Field (initial screen)

As you go to different Prism Central screens, the search string changes automatically to match the current location, as illustrated in the following example from the clusters page.

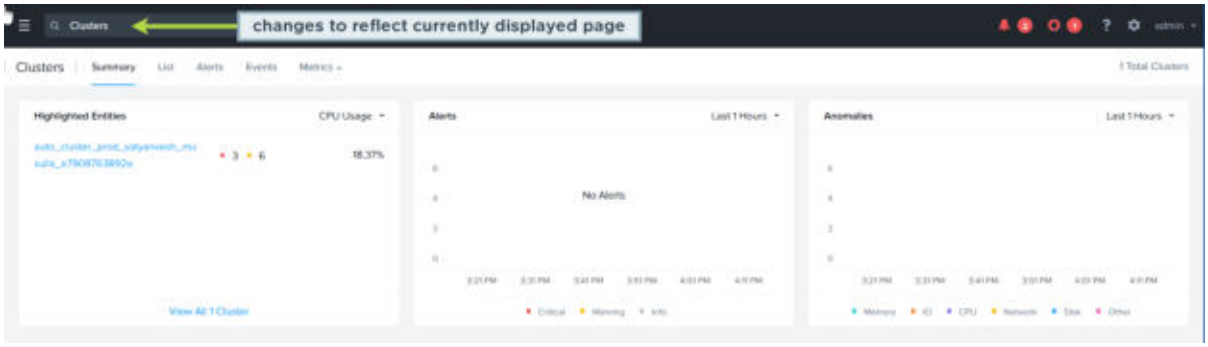

Figure 5: Search Field (from a different screen)

When you enter a string in the search field, a drop-down list appears with relevant matches across Prism Central. (Search strings are case insensitive.) For example, entering the string "VM" displays the following results:

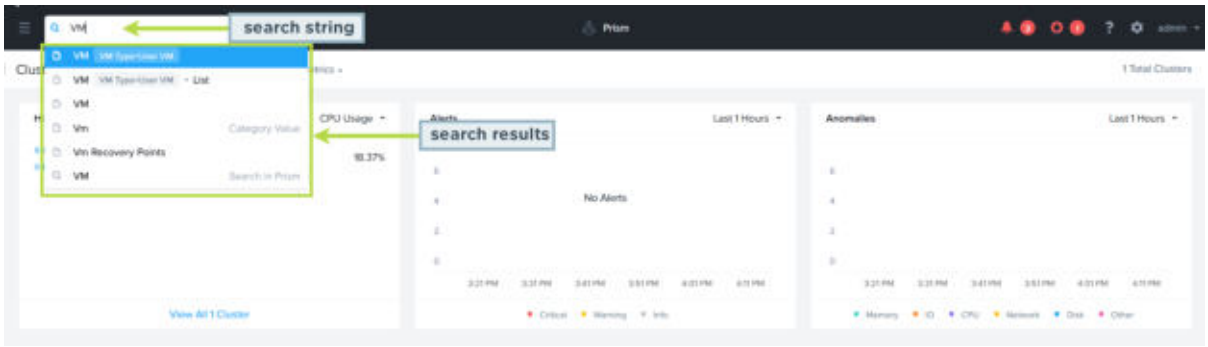

Figure 6: Search Results (for an entity type)

- "Category Value". Click this option to display the Category page with any VM-related entries.
- "VM Type=User VM". Click this option to display the VMs summary page filtered for user VMs.
- [blank]. Click this option to display the VMs summary page (unfiltered).
- "VM Type=User VM > List". Click this option to display the List tab of the user VMs page.
- "VM Type=User VM > Alerts". Click this option to display the Alerts tab of the user VMs page.
- "Search in Prism". Click this option to search across Prism Central for any information about the target entity. In the following example, two tables appear in the results, a list of the top VMs and a list of the top VM-related alerts. The top VMs list includes a link to the full list of

VMs, and the top alerts list includes a link to the full list of alerts (which in this case is the same). Clicking the VM or alert name takes you to the details page for that VM or alert.

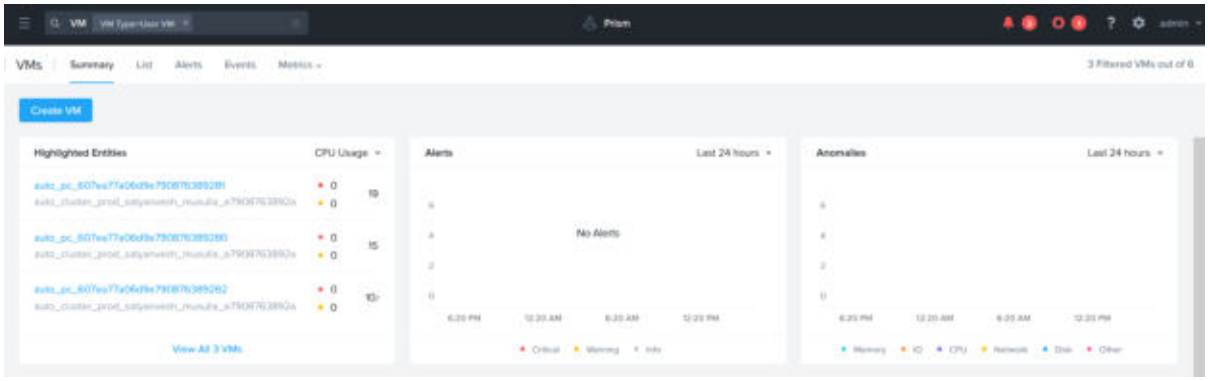

Figure 7: "Search in Prism" Results

Because search is context-sensitive, you can do a quick search from your current location without entering an explicit string. For example, if you are on the user VMs summary page, clicking in the search field (or clicking the forward slash [/] character) displays the search results for user VMs. In the following example, clicking the List, Alerts, or Events entry is the same as clicking those tabs on the page. Clicking "Memory Swap" displays the Memory Swap metrics page, and "Search in Prism" displays a results page for user VMs.

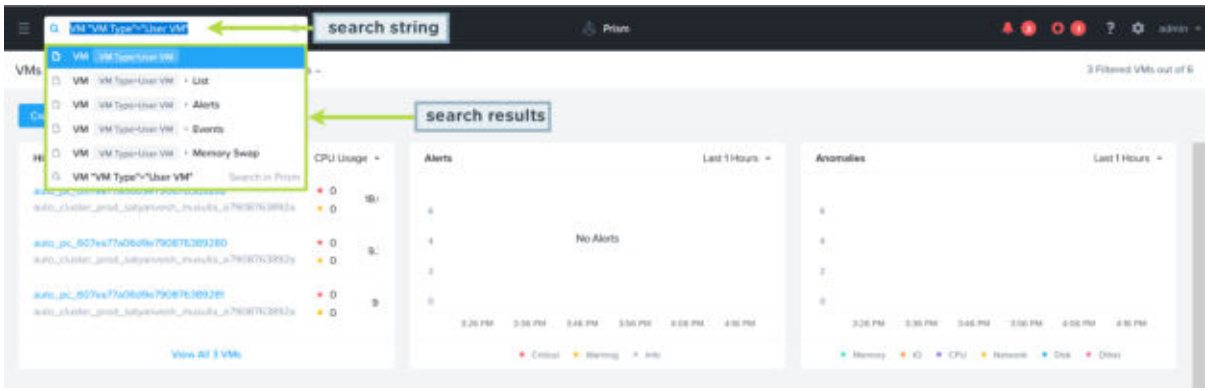

Figure 8: Search Results (for the current location)

Search provides intelligent responses based on whatever you enter in the field. For example, entering "scale out" returns a link to the Manage Prism Central page from which you can scale out Prism Central. If you enter a more generic string such as "version", the search engine returns results with links to pages across entities with version information.

Figure 9: Search Results (for generic string)

You can bookmark a filtered page (see Entity Exploring on page 77) by clicking the star icon in the search field. This adds a bookmark for that page (including any applied filters) to the entities menu (see Entities Menu on page 13).

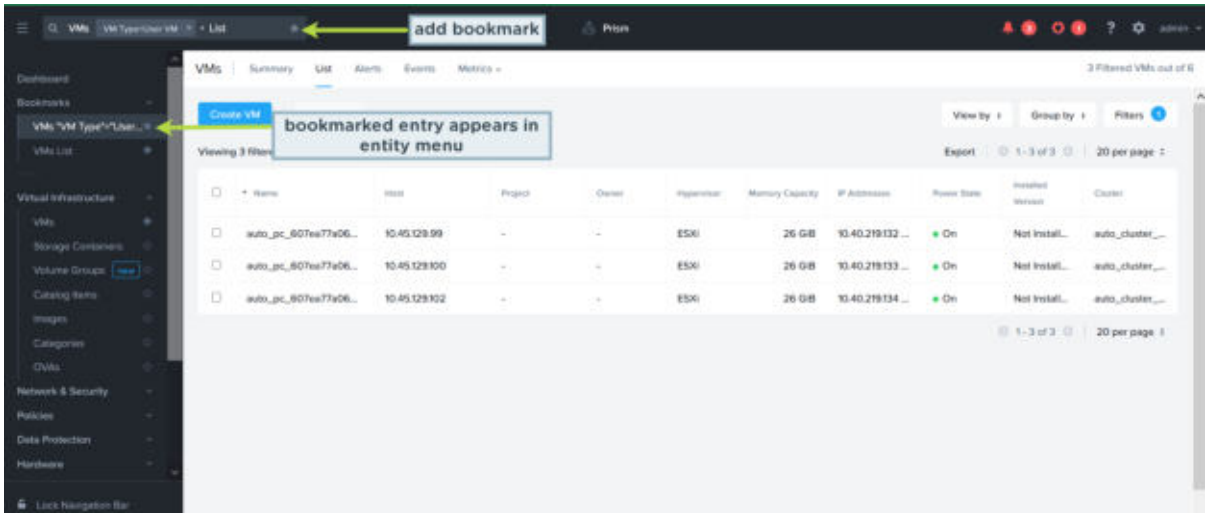

Figure 10: Bookmark Entry

On some pages, applied filters are not reflected in the main menu search field. For those pages, a local filters field appears. When you select one or more filters on such a page, the applied filters appear in the local filters field. You can do the following from this field:

- Remove a filter by clicking the X for that filter.
- Remove all filters by clicking **Clear** (on the right).
- Save the filter list by clicking the star icon. You can save a maximum of 20 filter lists per entity type.
- Use a saved filter list by selecting from the pull-down list.

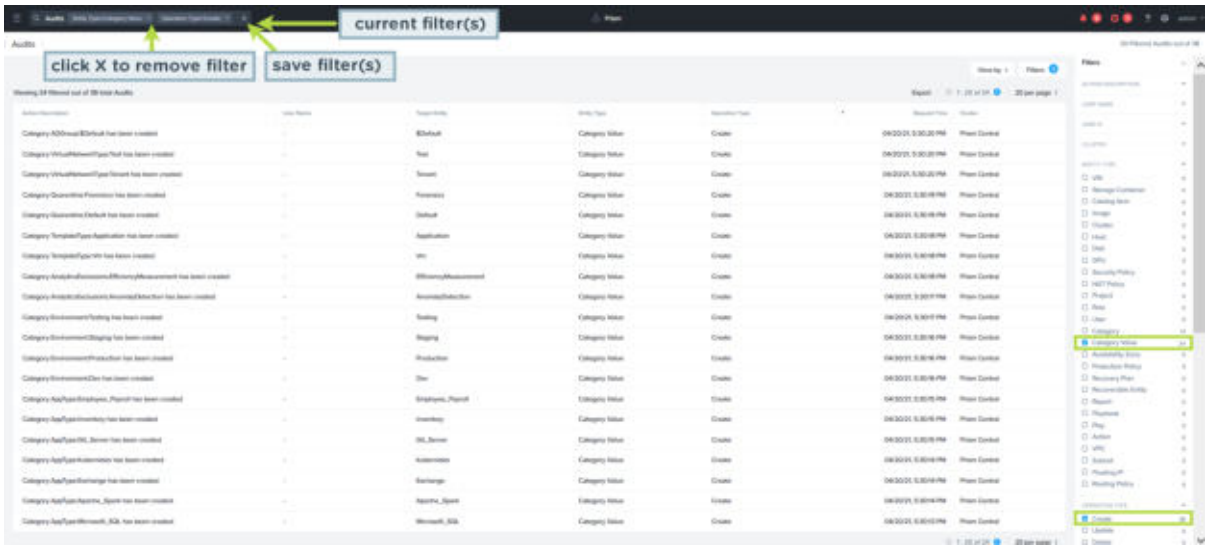

Figure 11: Local Filters Field

Query Syntax

The search field supports a range of simple and more complex query options in the following syntax forms:

- [entity|metric]: Enter an entity or metric type such as "cluster", "vm", or "alerts" to return information about the instances of that entity or metric across the registered clusters.
- <entity> <metric>: To refine an entity search, add a metric to the query. For example, "vm iops" returns IOPS information for the VMs.
- <entity> <attribute> <metric>: To further refine the search, add an attribute for the metric. For example, "node failure alerts" returns any failure alerts about the nodes.
- [metric|attribute] operator <value>: Enter an expression for a metric or attribute to return instances that match the expression. For example, ""block model"=1050" returns information about the NX-1050 nodes. See the following "Filter Expressions" section for the supported expression operators.
- <complex expression>: Combine the syntax forms to create complex queries. For example, "clusters hypervisor = AHV "cpu usage" < 30" returns information about clusters running the AHV hypervisor with CPU usage below 30%. Complex expressions have an implied AND so only instances that satisfy all conditions are returned.
- <action expression>: In addition to search queries, you can initiate certain actions from the query bar. For example, "<vm_name> launch console" launches a console window for the specified VM (if you are allowed to do so).

The following table describes the syntax rules for search queries.

Table 5: Syntax Rules

Rule	Example
RULES THAT GO TO WELL DEFINED PAGES	
entity type	vms
entity type + metric perspective (io, cpu, memory)	vms io
entity type + alerts	vm alerts
entity type + alerts + alert filters	vm alerts severity=critical
entity type + events	vm events
entity type + events + event filters	vm events classification=anomaly
entity type + filters (both metric and attribute)	vm “power state”=on
entity type + filters + metric perspective (io, cpu, memory)	vm “power state”=on io
entity type + filters + alerts	vm “power state”=on alerts
entity type + filters + alerts + (alert filters)	vm “power state”=on alerts severity=critical
entity type + filters + events	vm “power state”=on events

Rule	Example
entity type + filters + events + event filters	vm “power state”=on events classification=anomaly
entity instance (name, ip address, disk serial, ...)	vm1, 10.1.3.4, BHTXSPWRM
entity instance + metric perspective (io, cpu, memory)	vm1 io
entity instance + alerts	vm1 alerts
entity instance + alerts + alert filters	vm1 alerts severity=critical
entity instance + events	vm1 events
entity instance + events + event filters	vm1 events classification=anomaly
entity instance + pages	vm1 nics, c1 capacity
parent instance + entity type	c1 vms
alert title search	disk bad alerts
page name search	analysis, tasks
RULES THAT ONLY GO TO SEARCH RESULT PAGES (exclusively)	
entity type + metric	vm iops
entity type + attribute	vm power state
entity instance + metric	vm1 iops
entity instance + attribute	vm1 power state
entity type + filters + metric	vm “power state”=on iops
entity type + filters + attribute	vm “power state”=on hypervisor
help text search	upgrade cluster

Keywords

The following table lists the keywords or phrases you can use when formulating a query.

- The Object column identifies the type of object.
  - Entities: Lists the entities you can specify. The list is limited. For example, you cannot search for information about remote sites or network switches.
  - Fields: Lists the parameters (fields) you can specify.
  - Alerts: Lists the alert conditions you can specify.
  - Events: Lists the event conditions you can specify.
  - Actions: Lists the actions you can specify.
- The Entity column specifies the entities for which you can use that keyword in a query. For example, Alert queries can apply to any entity, but Fields has multiple rows broken down by entity type (such as cluster, VM, and disk) because there are select keywords that apply to each entity type.

- The Keywords column is a comma-separated list of the keywords or phrases you can use in a query for the specified object/entity type.

Table 6: Search Keywords

Object	Entity	Keywords
Entities	(n/a)	vm, cluster, node, container, disk
Fields	(not specified)	cpu usage, memory usage, disk usage, free physical storage, storage logical usage, saving ratio, savings, iops, read iops, write iops, io bandwidth, read io bandwidth, write io bandwidth, io latency, read io latency, write io latency, memory capacity, hypervisor
	Cluster	cluster name, ip address, version, number of hosts, cpu count, memory capacity, runway, storage runway, cpu runway, memory runway
	VM	vm name, ip address, host ip, virtual cpus count, power state, reserved memory, os, virtual hardware version
	Disk	serial, tier, mode, iops
	Container	container name, rf, on disk deduplication, perf-tier deduplication, compression, erasure coding
	Node	host name, ip address, service vm, hypervisor name, cpu capacity, cpu model, cpu sockets count, cpu cores count, cpu thread count, serial number, block model
Alerts	(any)	alert, alert title • Severity levels: critical, warning, info • Categories (impact types): capacity, performance, configuration, availability, system indicator • Dispositions types: resolved, unresolved, acknowledged, unacknowledged
Events	(any)	event type, title
Actions	VM	clone, migrate, delete, power on, power off, suspend, create vm, launch console, create network config, resume, snapshot, update, configure vm host affinity
	Cluster	launch prism element, unregister

Filter Expressions

You can use any of the following operators in an expression.

Table 7: Query Operators

Description	Operator
contains	~
does not contain	!~
starts with (“starts with” x)	=x*
ends with (“ends with” x)	=*x
equal to	=
not equal to	!=
greater than	>
greater than or equal to	>=
less than	<
less than or equal to	<=
a metric range, for example CPU usage between 10 and 30	=[10 to 30]

The following implicit operator rules apply to expressions:

- In most cases, an AND is applied to all filters in an expression. For example, "hypervisor=AHV iops>100" means the hypervisor is AHV and IOPS are over 100.
- If the filters are on the same attribute, an OR is applied. For example, “hypervisor=AHV hypervisor=ESXi” means the hyperviosr is either AHV or ESXi.
- If the filters are on the same metric, an AND is applied to create a range. For example, "iops>100 iops<500" means IOPS in the range 100-500.
- If multiple range filters are defined in a single query, an OR is applied across the ranges. For example, "iops>100 iops<500 cpu>20 cpu<40" means IOPS in the 100-500 range or CPU in the 20-40 range.

Search Units

The data for search queries comes from the following parameters.

Table 8: Data Parameters

Data Type	Unit	Parameter
bytes usage related properties	GiBiBytes	memory_capacity_bytes, storage.usage_bytes, storage.free_bytes, storage.logical_usage_bytes, data_reduction.saved_bytes, storage.usage_bytes, storage.user_container_own_usage_bytes, data_reduction.saved_bytes, memory_size_bytes

Data Type	Unit	Parameter
runway properties	days	capacity.runway, capacity.storage_runway, capacity.cpu_runway, capacity.memory_runway
disk capacity related properties	TebiBytes	storage.capacity_bytes, storage.user_capacity_bytes
CPU/memory usage properties	percent	hypervisor_cpu_usage_ppm, hypervisor_memory_usage

Example Queries

Here are examples of various query types.

- Entity queries:

```
<cluster_name>
<cluster_ip_address>
<disk_serial#>
VMs "Power State"=On List
VMs Hypervisor=AHV List
powered on vms "memory capacity" > 32
```

- Performance queries:

```
clusters running out of cpu
clusters hypervisor = AHV "cpu usage" < 30
vm iops
```

- Alert queries:

```
node failure alerts
<cluster_name> alerts
<cluster_name> critical availability alerts
<alert title>
Alerts "Create Time"="08-Nov-2018 9:46 AM to 08-Nov-2018 10:46 AM" Severity=Critical
```

- Action queries:

```
<vm_name> launch console
<cluster_name> launch prism element
create vm
```

- Exploration queries:

```
clusters hypervisor=AHV
vm os=Linux
<cluster_name> vms
"block model"=1050
"cpu model"=Intel
containers Rf > 2
clusters version=4.6.2
hosts iops < 1000
<cluster_name> powered off vms
disks tier=ssd
vms "cpu usage"
vms “power state”
```

Understanding Displayed Statistics

Various statistics are displayed in the web console and Prism Central screens. There are three possible sources for a statistic:

> Note: Most displayed statistics appear in 30 second intervals. The values in the tables represent the most recent data point within the last 30 seconds. Prism Central collects the statistical data from each registered cluster, so the process of collecting that data could result in a longer lag time for some statistics displayed in Prism Central.

1. Hypervisor. When the hypervisor can provide usage statistics, those figures appear where appropriate. ESXi provides such statistics, but Hyper-V and AHV do not. Getting the statistics from ESXi means numbers displayed in the Prism display screens (both Prism Central and Prism Element) should match the corresponding ones in vCenter.
2. Controller (Stargate). When hypervisor statistics are unavailable or inappropriate, the Controller VM provides statistics from Stargate. Controller-reported statistics might differ from those reported by the hypervisor for the following reasons:
   - An NFS client might break up large I/O requests into smaller I/O units before issuing them to the NFS server, thus increasing the number of operations reported by the controller.
   - The hypervisor might read I/O operations from the cache in the hypervisor, operations not counted by the controller.
3. Disk (Stargate). Stargate can provide statistics from both the controller and disk perspective. The difference is that the controller perspective includes read I/O operations from memory and disk I/O operations, while the disk perspective includes just the disk I/O operations.

> Note: The distinction between hypervisor, controller, and disk statistics applies only to storage-related statistics such as IOPS, latency, and bandwidth.

Field labels in the Prism display screens help identify the information source:

- A field name that does not include either "Controller" or "Disk" indicates that statistic derives from the hypervisor (for example "IOPS").
- A field name that includes the word "Controller" indicates that statistic derives from the controller (for example "Controller IOPS").
- A field name that includes the word "Disk" indicates that statistic derives from the disk (for example "Disk IOPS").

The following table identifies the information source for various statistics. Overview, VM, and storage statistics come from either the hypervisor or controller. IFor VM statistics in a mixed ESXi/AHV cluster, the source depends on the hypervisor hosting that VM (hypervisor for ESXi-based VMs and controller for AHV-based VMs). Hardware statistics come from the disk. Metrics in the analysis page can come from any of the sources (hypervisor, controller, or disk) depending on the type of metric.

Table 9: Source for Displayed Statistics

Hypervisor Type	Overview, VM, and Storage	Hardware	Analysis
ESXi	hypervisor (controller for some storage stats)	disk	metric dependent
Hyper-V	controller	disk	metric dependent

Hypervisor Type	Overview, VM, and Storage	Hardware	Analysis
AHV	controller	disk	metric dependent
Citrix Hypervisor	controller	disk	metric dependent
Mixed (ESXi + AHV)	hypervisor	disk	metric dependent
Prism Central	cluster dependent (hypervisor or controller)	disk	metric dependent

2

PRISM CENTRAL ADMINISTRATION

Administering Prism Central includes the following:

- Installing Prism Central and then upgrading Prism Central and registered clusters (see Installing or Upgrading on page 31)
- Managing licenses (see Licensing on page 31)
- Accessing Prism Central, managing the Prism Central instance (VM or set of VMs), and registering clusters with Prism Central (see Prism Central Management on page 31)

Installing or Upgrading

A Prism Central instance consists of either a single VM or a set of VMs that functions as a multi-cluster monitoring and control interface. You can install a Prism Central instance in either a Nutanix cluster or a non-Nutanix environment, although there are some limitations when you choose a non-Nutanix environment.

- To install or upgrade a Prism Central instance, see the "Prism Central Upgrade and Installation" topics in the *Acropolis Upgrade Guide*.

  > Note: Do not add any additional vDisks or NICs to a Prism Central VM above what is specified in the installation or upgrade instructions.

- To scale out an existing single VM Prism Central to multiple VMs, see Expanding (Scale Out) Prism Central on page 52.
- To upgrade NCC through Prism Central, see the "Upgrading NCC on Prism Central" topic the *Acropolis Upgrade Guide*.
- To upgrade one or more managed clusters through Prism Central, see the "Prism Central-Managed Clusters Upgrade" topic the *Acropolis Upgrade Guide*.
- If you use a proxy server, see Configuring Prism Central When a Cluster Uses Proxy Servers.

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the Nutanix License Manager Guide for information about how to license your cluster, apply add-on licenses, and manage your licenses in a variety of situations.

Prism Central Management

After installing Prism Central (see Installing or Upgrading on page 31), you can

- Log in and out through any supported browser (see Logging Into Prism Central on page 32 and Logging Out of Prism Central on page 35).

- Monitor the status of this Prism Central instance (see Managing Prism Central on page 35).
- Check version information (see Finding the Prism Central Version on page 42).
- Modify login page and automatic log off settings (see Modifying UI Settings on page 308).
- Register clusters with this Prism Central instance (see Register (Unregister) Cluster with Prism Central on page 45).
- Expand Prism Central from a single VM to a multi-VM instance (see Expanding (Scale Out) Prism Central on page 52).
- Shutdown or Start Up a Prism Central VM (see Shutdown (Start Up) Prism Central VM on page 56).
- Protect Prism Central to upto three AHV or ESXi clusters registered to it (see *Protecting Prism Central* in the *Prism Central Supplement* for pc.2021.7).
- Recover Prism Central from Prism Element running on the registered AHV or ESXi cluster (see *Recovering Prism Central (1-Click Disaster Recovery)* in the *Prism Central Supplement* for pc.2021.7).

Logging Into Prism Central

About this task

To log in to Prism Central, do the following. Knowledge base article KB 1661 lists default cluster credentials.

Procedure

1. Open a web browser, enter http://*management_ip_addr* in the address field, and press **Enter**. Replace *management_ip_addr* with the Prism Central VM IP address.

   > Note: The supported browsers are the current version and two major versions back of Firefox, Chrome, and Safari, plus Internet Explorer version 11 and Microsoft Edge.

   The browser redirects to the encrypted port (9440) and may display an SSL certificate warning. Acknowledge the warning and proceed to the site. If user authentication is enabled and the browser does not have the correct certificate, a denied access message may appear. For the complete list of required ports, see Port Reference.

2. If a welcome screen appears, read the message and then click the "Accept terms and conditions" bar at the bottom.

3. In the login screen, do one of the following: enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

   » Enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

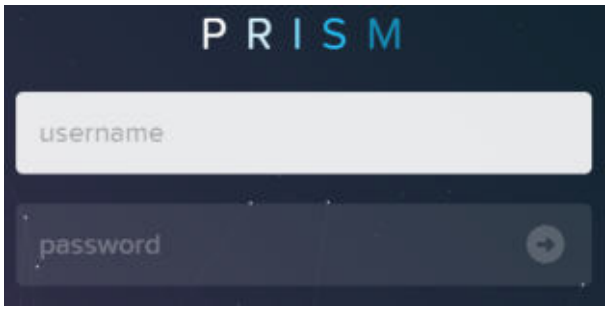

Figure 12: Login Screen

Note: If LDAP authentication is used, enter the user name in *username@domain* format; the *domain\username* format is not supported. The user principle name [UPN] attribute is used to find the user account in Active Directory, so your user name must be in that format.

   » If SAML authentication is enabled (see "" ()), enter your user name according to the configured identify provider in the indicated field. (This field appears on the login page only when SAML authentication is enabled.)

Figure 13: Login Screen (SAML)

Note: The login page includes background animation that is enabled by default. Click the **Freeze space time continuum!** link at the bottom right of the login screen to disable the animation (or the **Engage the warp drive!** link to enable the animation). To permanently disable (or enable) the animation, see Modifying UI Settings on page 308).

4. If you are logging in as an administrator for the first time (user name admin and default password Nutanix/4u), which requires that the default password be changed, enter a new

password in the **password** and **re-type password** fields and then press **Enter** or click the right arrow icon.

Figure 14: Login Screen (first admin login)

The password must meet the following complexity requirements:

- At least 8 characters long
- At least 1 lowercase letter
- At least 1 uppercase letter
- At least 1 number
- At least 1 special character
- At least 4 characters difference from the old password
- Should not be among the last 10 passwords

After you have successfully changed the password, the new password is synchronized across all Controller VMs and interfaces (Prism web console, nCLI, and SSH).

Note:

- The default password expiration age for the admin user is 60 days. You can configure the minimum and maximum password expiration days based on your security requirement.
  - nutanix@cvm$ sudo chage -M <MAX-DAYS> admin
  - nutanix@cvm$ sudo chage -m <MIN-DAYS> admin
- After you upgrade from an AOS earlier version and then attempt to log in to the Prism web console as the admin user, you are prompted to create a new admin user password.
- When you change the admin user password, update any applications and scripts using the admin user credentials for authentication. Nutanix recommends that you create a user assigned with the admin role instead of using the admin user for authentication. The *Prism Web Console Guide* describes authentication and roles.

5. If a license agreement screen appears (typically on the first login or if the EULA changed since the last login), which indicates the current EULA has not been acknowledged yet, do the following:

   a. Read the license agreement (on the left).

   b. Enter appropriate information in the **Name**, **Company**, and **Job Title** fields (on the right).

   c. Check the "I have read and agree to the terms ..." box.

   d. Click the **Accept** button.

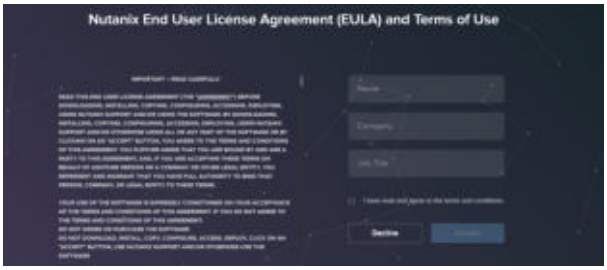

Figure 15: EULA Screen

Logging Out of Prism Central

Procedure

To log off from Prism Central, click the user icon in the main menu and then select the **Sign Out** option from the pull-down list. You are logged out immediately after selecting the option (no prompt or message).

Managing Prism Central

Prism Central consists of one or more VMs treated as a single instance that monitors and manages registered clusters. You can view information about the Prism Central VMs by clicking the gear icon in the main menu and then selecting **Prism Central Management** from the Settings menu (see Settings Menu on page 17). This displays the Manage Prism Central page.

> Note: This page is not available if Prism Central is hosted in a non-Nutanix environment, that is not hosted in a Nutanix cluster.

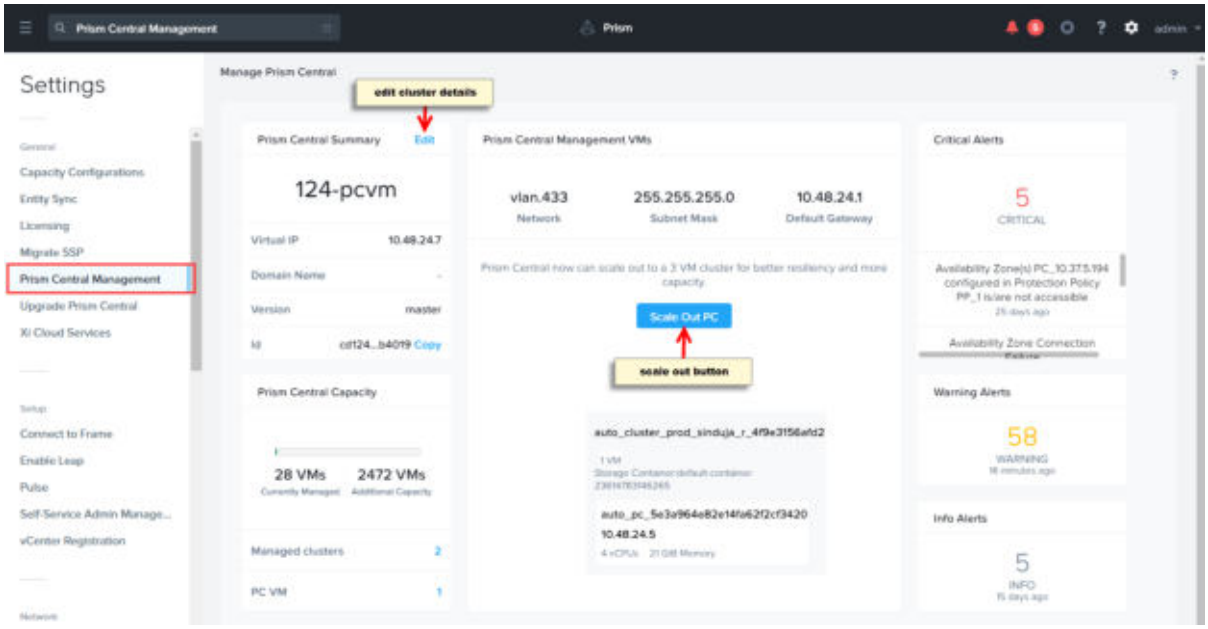

Figure 16: Manage Prism Central Page (1-VM instance)

The Manage Prism Central page contains the following:

- A "Prism Central Summary" pane on the upper left that displays fields for
  - Cluster name, which is "Unnamed" by default.
  - Domain name. A dash appears if a domain name is not specified.
  - Version. This is the Prism Central version number.
  - ID number. Click the copy link to copy the ID number so you can paste it elsewhere.
- A "Prism Central Capacity" pane on the lower left that displays fields for
  - Number of VMs monitored currently by this Prism Central instance.
  - Number of additional VMs (remaining capacity) this Prism Central instance can monitor.
  - Number of clusters registered to this Prism Central instance; click the number to display the clusters page (see Clusters Summary View on page 146).
  - Number of VMs that comprise this Prism Central instance; click the number to display the VMs page (see VMs Summary View on page 82).
- A "Prism Central VMs" pane in the middle that displays network address information (network name, subnet mask, and default gateway address) at the top and information below about each VM in the Prism Central instance (VM name and name of storage container in which it is located, IP address, number of vCPUs, and memory size). In addition, single VM instances include a **Scale Out PC** button to scale out this Prism Central instance (see Expanding (Scale Out) Prism Central on page 52).
- An "Alerts" pane on the right that displays a list of Prism Central-related alerts broken into Critical, Warning, and Info sections.

When the Prism Central instance consists of three VMs, the display changes slightly. The **Scale Out PC** button disappears, and a new **Add PCVM** button appears (top right). Prism Central instances are limited to a maximum of three VMs, but if you lose one of those VMs for any

reason, you can add a replacement by clicking the **Add PCVM** button. This displays the Add PC VMs page (similar to the **Scale Out PC** page) from which you can add the new VM.

> Note: Verify that the prerequisites in Expanding (Scale Out) Prism Central on page 52 are satisfied before adding a Prism Central VM.

To add or change the Prism Central domain name or virtual IP, click the **Edit** link in the "Prism Central Summary" pane. The Cluster Details window appears. Do the following:

1. In the **Cluster Name** field, enter a name for the cluster (default "Unnamed"). After naming the cluster, the field no longer appears when subsequently opening the window. Naming the cluster is optional.
2. In the **FQDN** field, enter the fully qualified domain name (FQDN) for the Prism Central cluster. For a multi-VM Prism Central, specifying a domain name allows balancing of requests across the Prism Central VMs to improve performance and resiliency. (This requires an administrator configure the domain name in the DNS server as the set of IP addresses associated with the Prism Central VMs.)

   When an FQDN is specified, it is the first choice for IP address resolution. Make sure the virtual IP (next field) is not mapped to multiple DNS entries as this can lead to resolution problems. If the FQDN fails to resolve in DNS, the existing Prism Central VM IP addresses are used.
3. In the **Virtual IP** field, enter an IP address that will be used as a virtual IP for the cluster. This is relevant if you have a multi-VM Prism Central.
4. When the fields are correct, click the **Update** button to save the changes and close the window.

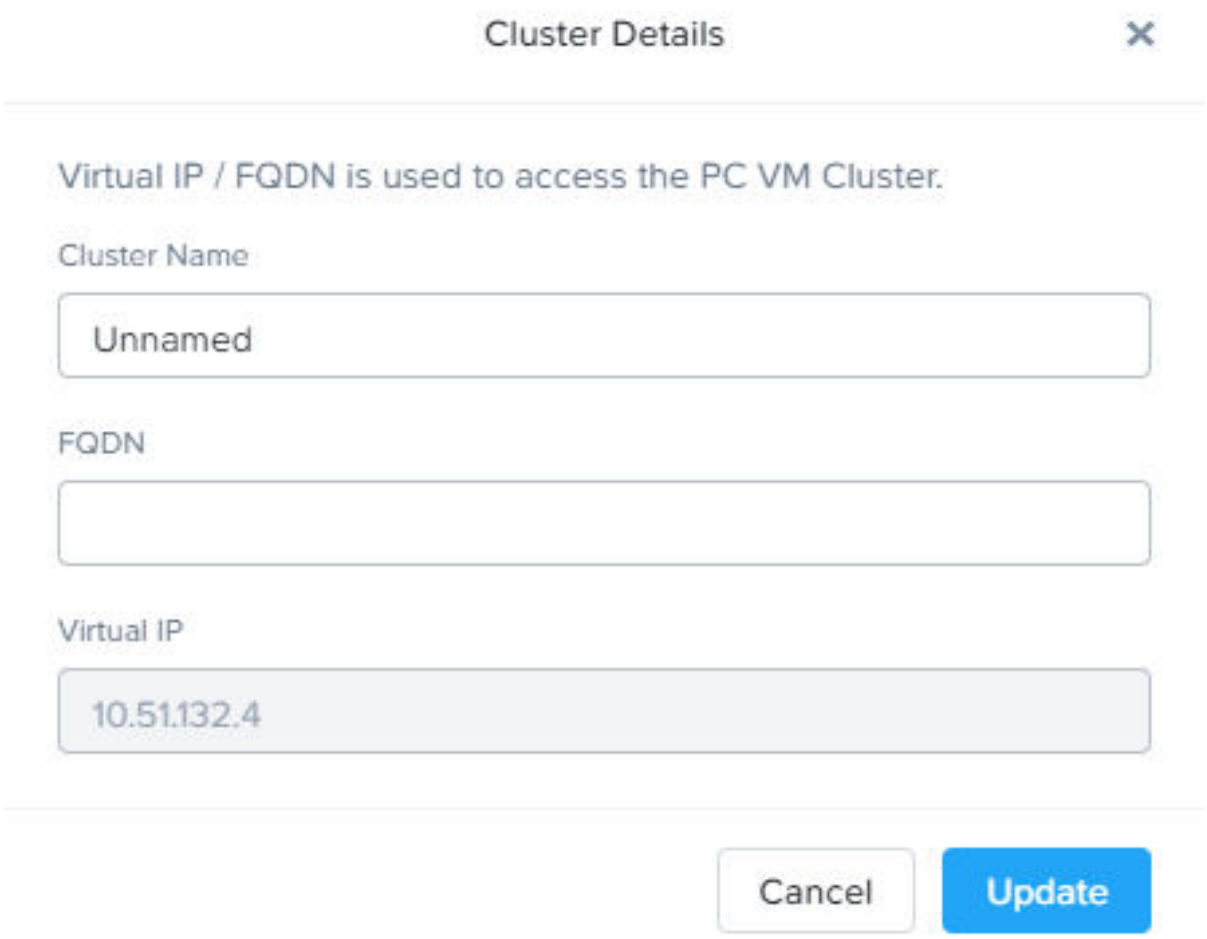

Figure 17: Cluster Details Window

Enabling Microservices Infrastructure

Some Prism Central services are packaged in containers as micro services, and the control plane for micro services platform (sometimes referred to as CMSP) enables the Microservices Infrastructure. The Microservices Infrastructure is disabled by default, but you can enable it at any time.

Before you begin

Address the following before enabling the Microservices Infrastructure:

- The Microservices Infrastructure is supported on clusters running AHV or ESXi only. For ESXi clusters you may need to enter your vCenter credentials (user name and password) and a network for deployment.
- The host cluster must be registered with this Prism Central instance (see Register (Unregister) Cluster with Prism Central on page 45).
- The host cluster virtual IP address and ISCSI data services IP address must be configured (see the "Modifying Cluster Details" section in the *Prism Web Console Guide*).
- Name servers and NTP servers must be configured in both the host cluster (Prism Element) and Prism Central.

- You may need to increase the Prism Central VM memory size before enabling the Microservices Infrastructure. The requirements are listed in KB 8932. If your Prism Central VM resources are below the recommended amount, manually increase the memory to the recommended amount.

About this task

To enable the Microservices Infrastructure, do the following:

CAUTION: Once enabled, the Microservices Infrastructure cannot be disabled. The following limitations apply after enabling the Microservices Infrastructure:

- Scaleout three-node PC VMs deployment: Small, large, or x-large PC VMs are supported. Reverting this deployment to a single PC VM deployment is not supported.
- Single PC VM deployment: Small, large, or x-large PC VM is supported. However, you cannot expand the single VM deployment to a scaleout three-node deployment.
- Prism Central backup, restore, and migration are not supported.

  Note: With a minimum version of pc.2021.7, Prism Central Disaster Recovery (PCDR) supports microservices infrastructure. After recovery, enable microservices infrastructure. For information about PCDR, see the *Prism Central Release Notes* for the minimum version.

- Changes to any of the following IP addresses are not supported: Prism Central virtual IP address, host cluster virtual IP address, host cluster iSCSI data services IP address
- Unregistering the host cluster from Prism Central is not supported. However, unregistering other clusters is still supported.
- Common Access Card (CAC) is not supported.
- Ensure that the IP addresses in subnet that you configure for microservices infrastructure do not conflict with the IP addresses in the management subnet.

  Note: Do not use the IP addresses in the subnet 10.100.0.0/24 for operational purposes like DNS.

Procedure

1. Go to the Manage Prism Central page (see Managing Prism Central on page 35).

2. In the **Prism Central on Microservices Infrastructure** box, click the **Enable Now** button.

Figure 18: Enable box (in Manage Prism Central page)

3. In the Enable Microservices Infrastructure window, do the following in the indicated fields:
   a. [first screen, no fields] Read the information and then click the **Continue** button.

      The window redisplays with two fields.

   b. **Prism Central Service Domain Name**: Enter a domain name for the service.

      > Note: Ensure that the domain is not longer than 16 characters. Prism Central supports a maximum of 16 characters for domain name for enabling microservices infrastructure. If the length of the domain name exceeds 16 characters, enabling microservices infrastructure fails.

      The domain name is used for exporting services under the Prism Central platform. The name must be unique and should be at least three level, for example test.nutanix.com. Do not use the name of an existing domain. (The parent domain's DNS server could redirect the DNS requests to the Prism Central service domain's DNS server to resolve the service names globally.) The IP address of the DNS server appears in the Prism Central management page.

      > Note: While running a minimum version of AOS 6.0.1, when you configure microservices infrastructure and Flow networking, ensure that the DNS name you configure for CMSP

does not end with test. Flow networking does not support test as a top level domain. For example, the following are valid domain configurations:

- my.cluster.domain
- my.test.cluster.test.domain

However, the following are examples of domains that Flow networking does not support:

- my.cluster.test
- my.cluster.domain.test

c. **Internal Network**: Select the network to use for Prism Central micro services communication from the pull-down list.

   If the desired network does not appear in the list, you must first define it (see Network and Security Management on page 421).

d. **Subnet Mask**: Enter the desired subnet mask.

   For a managed network, subnet mask and gateway IP are already configured, but you must enter those values if the selected network is unmanaged.

e. **Gateway IP Address**: Enter the IP address for the gateway.

f. **IP Address Range**: Enter a range of IP addresses the network can use.

   Enter a range of at least 10 available (unreserved) addresses for a 3-VM (scale out) Prism Central instance or at least 5 addresses for a 1-VM Prism Central instance. The addresses must be consecutive. For a managed network, the range of addresses for microservices must be outside the range of reserved addresses for the selected network.

g. When all the fields are correct, click the **Validate** button.

   The system checks the settings. When the settings are confirmed, a "Validation was successful" message appears and the **Validate** button changes to an **Enable** button.

h. Click the **Enable** button.

   The enablement process begins, which takes some time (around 45 minutes). You can monitor progress from the tasks dashboard (see Tasks View on page 190). When the microservices infrastructure is enabled, the contents of the enable box in the Manage Prism Central page changes to display configuration information such as the domain name, the used IP addresses, and the service domain's DNS IP address.

4. After the enablement tasks are completed, including the **IAM Migration & Bootstrap** task, log out of Prism Central. Wait at least 15 minutes before logging on to Prism Central.

Enabling Microservices Infrastructure (dark site)

Enabling the Microservices Infrastructure when there is no Internet access (dark site) requires some extra steps from the standard procedure.

About this task

To enable the Microservices Infrastructure at a dark site, do the following:

Procedure

1. From a device that has public Internet access, log on to the Nutanix support portal, select **Downloads** > **LCM** from the entity menu (hamburger icon), and then download the LCM Dark Site Bundle tar file.
2. Set up a local web server, upload LCM Dark Site Bundle to the server, and extract the files in a directory in the base of the web server.

   See the *Setting Up a Local Web Server* section in the *Life Cycle Manager Dark Site Guide* for more information. Set the following mime types in the LCM web server when you are using Windows Web Server: .json, .xz, .qcow2, .sign, .gz
3. In the Nutanix support portal **Downloads** page, select **Microservices Platform (MSP)**, download the **Upgrade (Darksite)** bundle to the web server base directory you used in step 2, and extract the files.
4. Now follow the standard procedure for enabling the Microservices Infrastructure through Prism Central (see Enabling Microservices Infrastructure on page 38).

Finding the Prism Central Version

About this task

You can check the version of your Prism Central instance using the Prism Central user interface or by running the following command.

Figure 19: Finding the Prism Central Version

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275778971001 on Brightcove.

```
nutanix@cvm$ ncli cluster info
```

Do the following to check the Prism Central version in the UI.

Procedure

1. Click the user icon in the main menu and then select the **About Nutanix** option from the pull-down list.

   An About Nutanix window appears that includes the AOS version number. It also includes a link to Nutanix patent information.

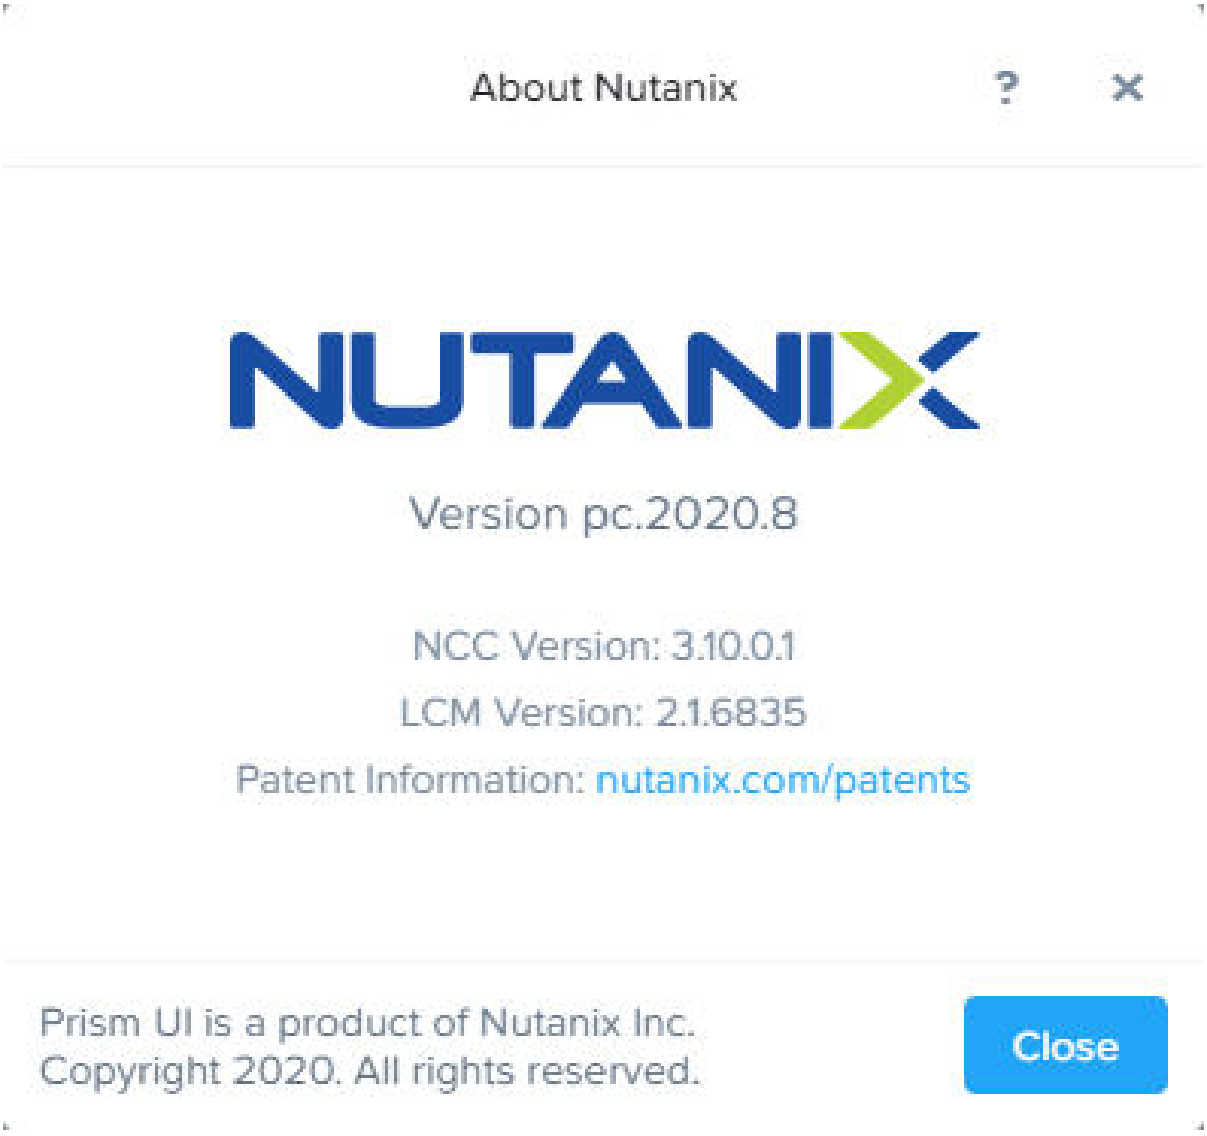

Figure 20: Example About Nutanix Window

2. Click **Close** to close the window.

Finding the AHV Version on Prism Central

You can see the installed AHV version in the Prism Central console.

About this task

To view the AHV version installed on any host in the clusters managed by the Prism Central, do the following.

Figure 21: Finding the AHV Version on Prism Central

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275777446001 on Brightcove.

Procedure

1. Log on to Prism Central.
2. In side bar, select **Hardware** > **Hosts** > **Summary** tab.
3. Click the host you want to see the hypervisor version for.
4. The Host detail view page displays the **Properties** widget that lists the Hypervisor Version.

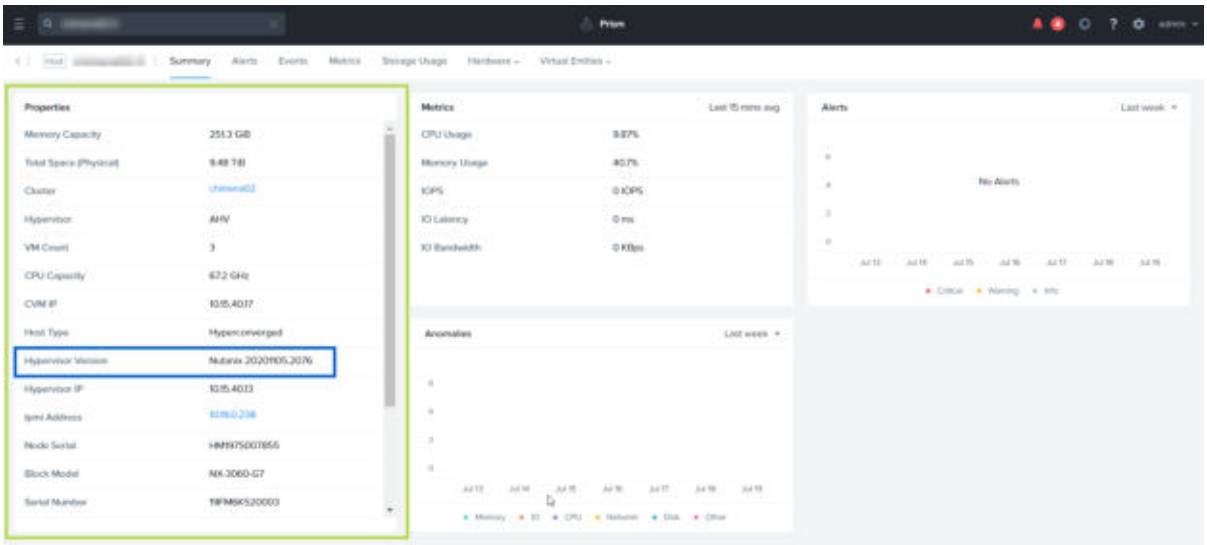

Figure 22: Hypervisor Version in Host Detail View

Finding the AOS Version Using Prism Central

To view the Nutanix AOS version running in the Prism Central, do the following:

Procedure

1. Log on to Prism Central.
2. Click the entities menu (hamburger menu) near the search field and then select **Hardware > Clusters** to view the cluster dashboard.

3. In the **List** tab (default), the AOS version of the cluster is listed against the cluster name.

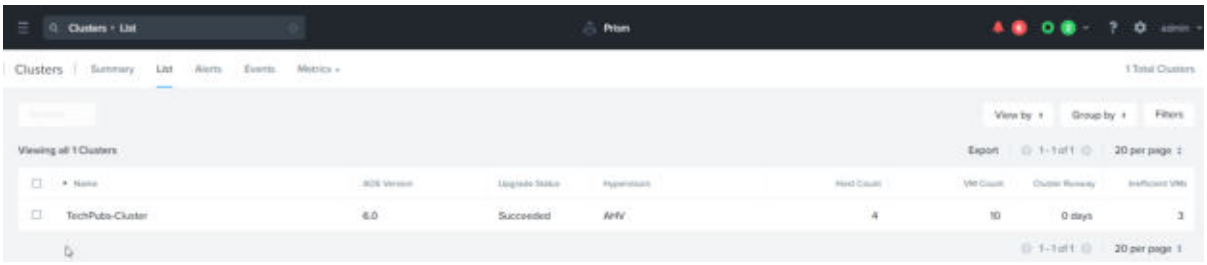

Figure 23: AOS Version in a Cluster

Register (Unregister) Cluster with Prism Central

Before you begin

- If you have never logged into Prism Central as the user admin, you need to log in and change the password before attempting to register a cluster with Prism Central.
- Do not enable client authentication in combination with ECDSA certificates on a registered cluster since it causes interference when communicating with Prism Central.
- Ports 9440 and 80 need to be open in both directions between the Prism Central VM and all the Controller VMs (and the cluster virtual IP address if configured) in each registered cluster. For the complete list of required ports, see Port Reference.
- A cluster can register with just one Prism Central instance at a time. To register with a different Prism Central instance, first unregister the cluster.

About this task

To register a cluster with Prism Central, do the following:

Procedure

1. Log in to the Prism web console on the target cluster as the user admin.
2. In the Health dashboard, from the **Actions** drop-down menu, select **Run Checks**. to run Nutanix Cluster Checks.

3. Do one of the following:

   » On the **Home** dashboard, click **Register or create new** from the **Prism Central** widget.

   » Click the gear icon and select **Prism Central Registration** from the **Settings** menu.

Figure 24: Home Dashboard

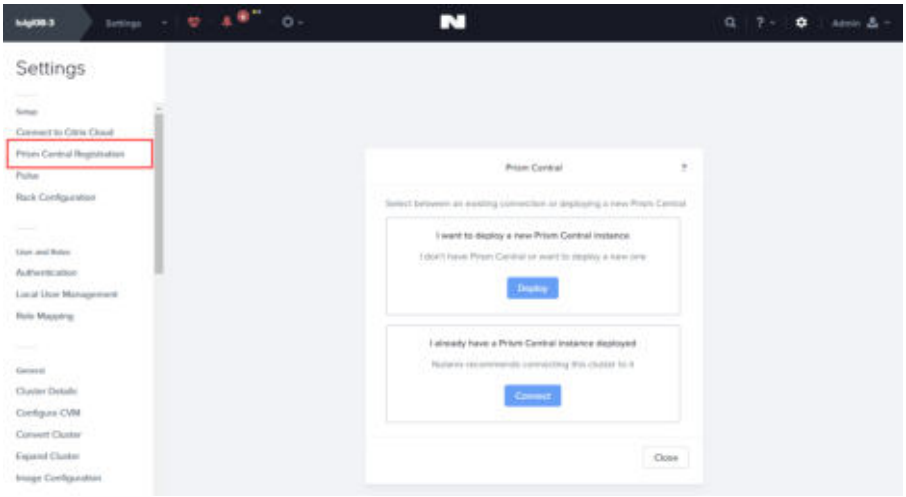

Figure 25: Settings Menu

4. In the first screen of the Prism Central dialog box, click the **Connect** button.

   This screen includes two options, **Deploy** and **Connect**. This procedure describes how to connect to an existing Prism Central instance. See *Installing Prism Central (1-Click Method)*

section in Acropolis Upgrade Guide for instructions on how to deploy a new Prism Central instance.

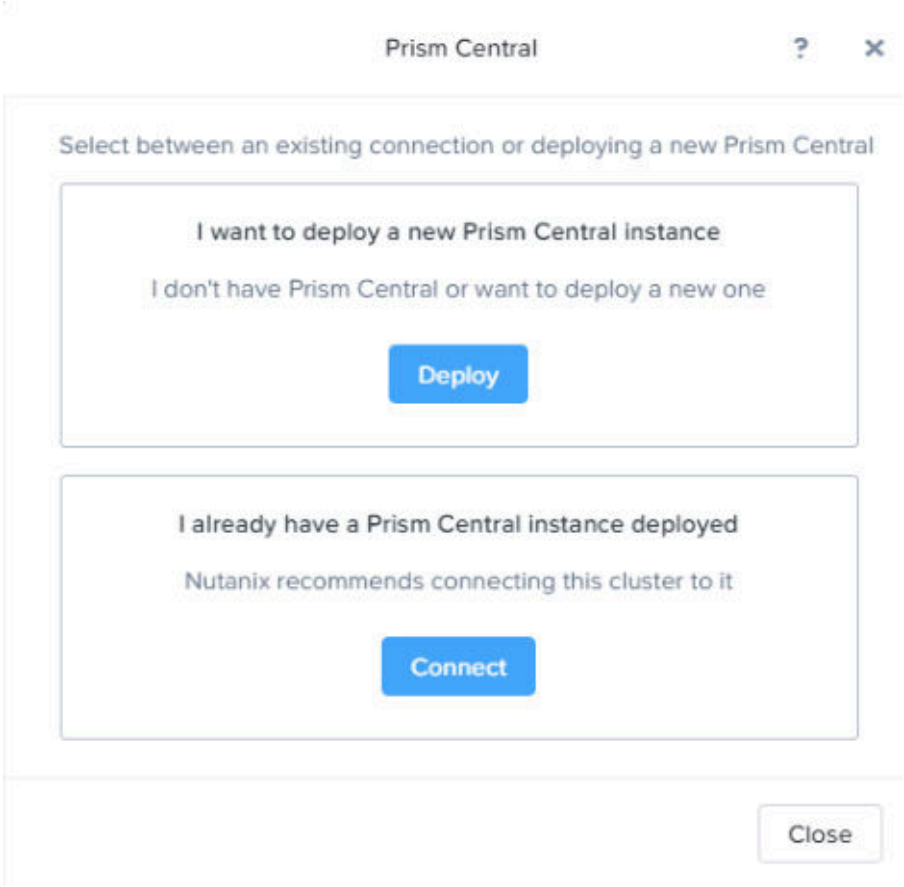

Figure 26: Prism Central Dialog Box (deploy/connect screen)

5. A screen appears outlining what services are available through Prism Element and Prism Central. After reviewing the message, click the **Next** button.

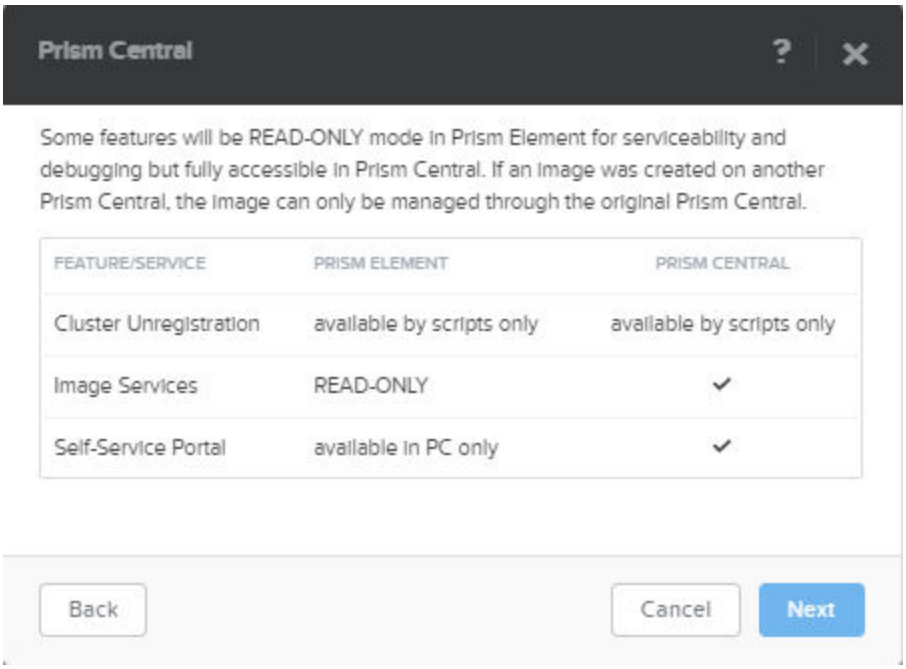

Figure 27: Prism Central Dialog Box (services screen)

6. In the registration screen, do the following in the indicated fields:
   a. **Prism Central IP**: Enter the IP address of the Prism Central VM.
   b. **Port**: The default port number is 80. This is an optional field. For the complete list of required ports, see Port Reference.
   c. **Username**: Enter admin as the Prism Central user name.
   d. **Password**: Enter the Prism Central admin user password.

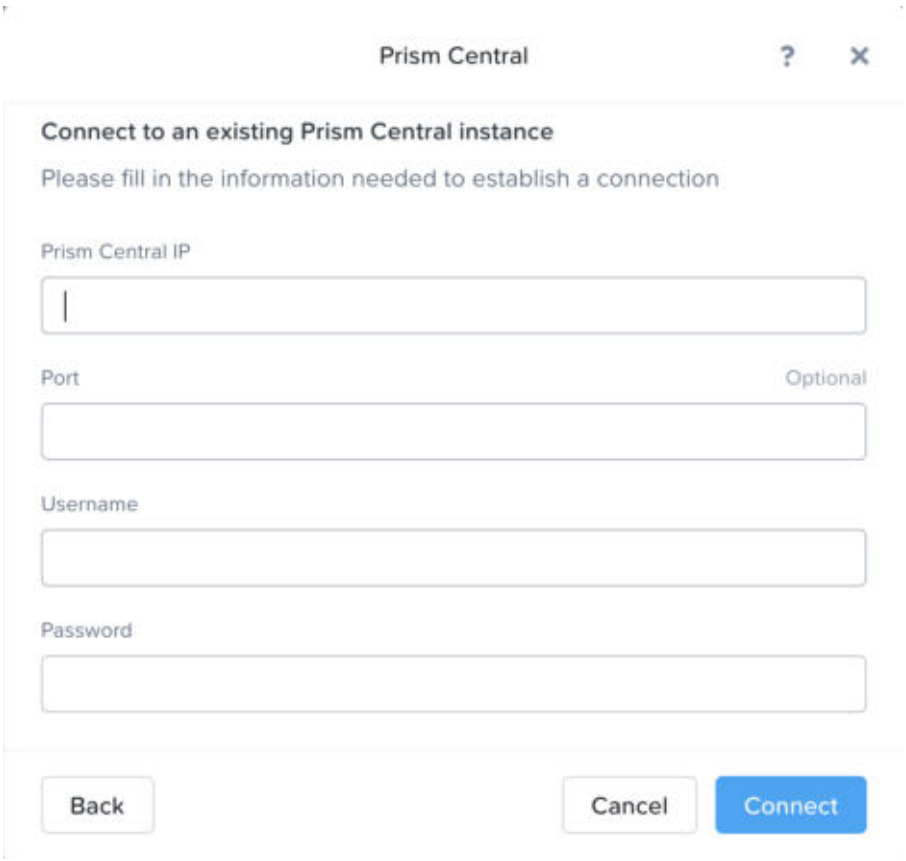

Figure 28: Prism Central Dialog Box (registration screen)

7. When all the fields are correct, click the **Connect** button to save the values and close the window.
   This registers the cluster on the specified Prism Central VM and allows the passing of information between the cluster and Prism Central.

   Note: The user credentials provided when registering a cluster (Prism Element) with Prism Central are only used once. After registration, modifying the admin password would not impact any communication between Prism Element and Prism Central.

Unregistering from Prism Central

About this task

To unregister a cluster from an existing Prism Central instance, do the following:

Note: Unregistering a cluster through the Prism GUI is no longer available. This option was removed to reduce the risk of accidentally unregistering a cluster because several features

(including role-based access control, application management, micro-segmentation policies, and self-service capability) require Prism Central to run your clusters. If a cluster is unregistered from Prism Central, not only will these features not be available but the configuration for them may also be erased. Therefore, only the following procedure is available to unregister a cluster. See KB 4944 for additional details if you have enabled Self Service, Calm, or other special features in Prism Central.

Procedure

1. Log on to any Controller VM of the registered cluster through an SSH session.
2. Run the **cluster status** command and verify that all services are in a healthy state.

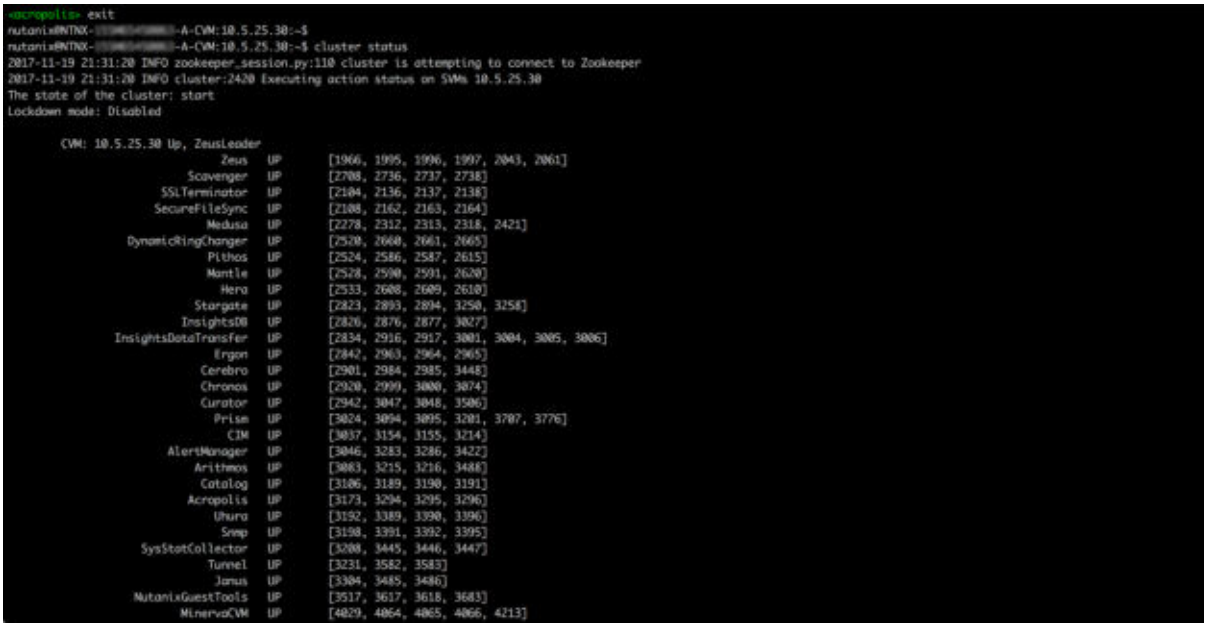

Figure 29: Cluster Status Command

3. Unregister the cluster from Prism Element.

   ```
   nutanix@cvm$ ncli multicluster remove-from-multicluster external-ip-address-or-svm-ips=pc-name-or-ip username=pc-username password=pc-password force=true
   ```

   Replace *pc-name-or-ip* with the Prism Central name or IP address and *pc-username* and *pc-password* with the login credentials for your Prism Central administrator account. This step can take some time (though typically just a few seconds). To check if the unregistration has completed, run the following command.

   ```
   nutanix@cvm$ ncli multicluster get-cluster-state
   ```

4. Get the UUID for the cluster.

   ```
   nutanix@cvm$ ncli cluster info
   ```

   Find the "Cluster UUID" value in the displayed information.

   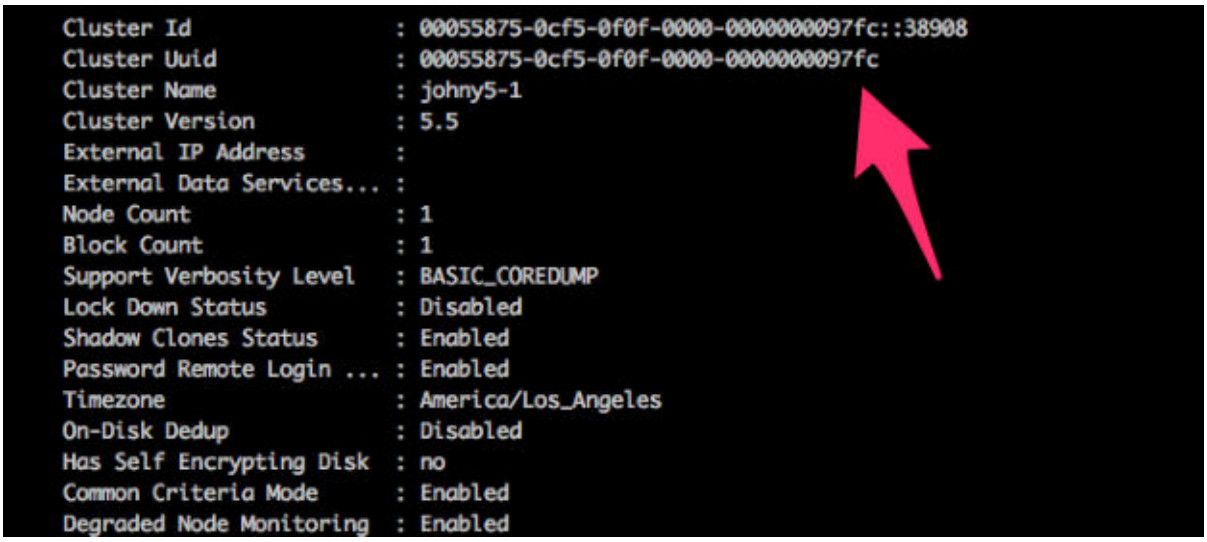

   Figure 30: Cluster Info Output

5. Log on to the Prism Central VM through an SSH session (as the "nutanix" user) and do the following:

   a. Run the unregistration clean-up script.

      ```
      [pcvm]$ python /home/nutanix/bin/unregistration_cleanup.py uuid
      ```

      Replace *uuid* with the value you obtained in step 4. This script removes all remaining registration information about that cluster and completes the unregistration process for the Prism Central VM.

   b. Get the UUID for Prism Central.

      ```
      [pcvm]$ ncli cluster info
      ```

      Find the "Cluster UUID" value in the displayed information (see step 4), which in this case is the UUID for Prism Central.

6. Go back to the Controller VM and run the **unregistration_cleanup.py** script to complete the unregistration process on the cluster.

   ```
   nutanix@cvm$ python /home/nutanix/bin/unregistration_cleanup.py uuid
   ```

   In this case the *uuid* is the Prism Central UUID obtained in step 5b.

What to do next

After performing these steps you can re-register the cluster with a new or re-created Prism Central instance.

If the clean up process does not complete successfully, try the following:

- Check the logs to indicate if there are any input errors when calling the script. The logs for the unregistration cleanup script can be found under **~/data/logs/unregistration_cleanup.log**.
- If errors occur during script execution, run the **cluster status** command and check that the cluster services are up and running. Rerun the script and check if it succeeds.

Note: If Prism Self Service is enabled and the unregistered cluster is set as the default cluster, you must update the default project to select a new (registered) cluster. See the topic *Modifying A Project* in the Prism Central Guide.

Expanding (Scale Out) Prism Central

Before you begin

The following requirements must be met before you can expand Prism Central or add a Prism Central VM:

- The specified gateway must be reachable.
- No duplicate IP addresses can be used.
- The container used for deployment is mounted on the hypervisor hosts.
- When installing on an ESXi cluster:
  - vCenter and the ESXi cluster must be configured properly. See the *vSphere Administration Guide for Acropolis (using vSphere Web Client)* for details.
  - vCenter must be registered in Prism.
  - DRS must be enabled in vCenter.
  - vCenter is up and reachable during the deployment.

About this task

If Prism Central is just a single VM currently, you can expand it to three VMs. This increases both the capacity and resiliency of Prism Central (at the cost of maintaining two additional VMs). To expand this Prism Central instance across multiple VMs, do the following:

Note:

- Scale out Prism Central is supported on AHV and ESXi clusters only. If you are expanding a pre-AOS 5.17 Prism Central "large" VM that contains just a single vDisk, the upgrade process adds three new vDisks to each VM. This allows Prism Central to shard Cassandra metadata to improve performance. While the listed disk capacity of a Prism Central VM includes all the vDisks, the additional vDisks do not impact overall capacity. During the scale out you can monitor progress ("VM disk attach" tasks) from the Tasks dashboard.
- All scale out Prism Central VMs must run on the same cluster. For example, running two VMs in cluster_1 and one VM in cluster_2 is not supported.

Procedure

1. Click the gear icon and then select **Prism Central Management** from the Settings menu (see Settings Menu on page 17).
   The Manage Prism Central page appears. This page provides information about this Prism Central instance.

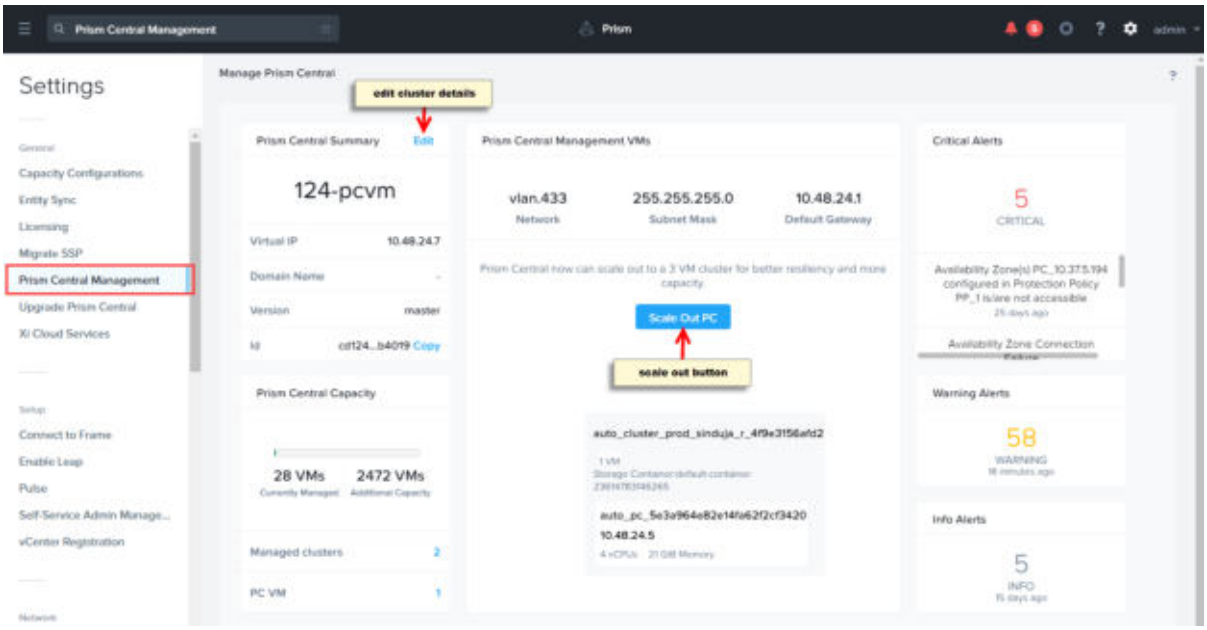

Figure 31: Manage Prism Central Page

2. To expand this Prism Central instance from one to three VMs, click the **Scale Out PC** button to display the Scale Out PC page and do the following:

   Note: A pop-up window appears explaining that scale out is a one-way process. Click the **Continue** button to display the Scale Out PC page. Once you scale out a Prism Central

instance from a single VM to multiple VMs, you cannot revert back. Deleting any of the Prism Central VMs may result in data loss.

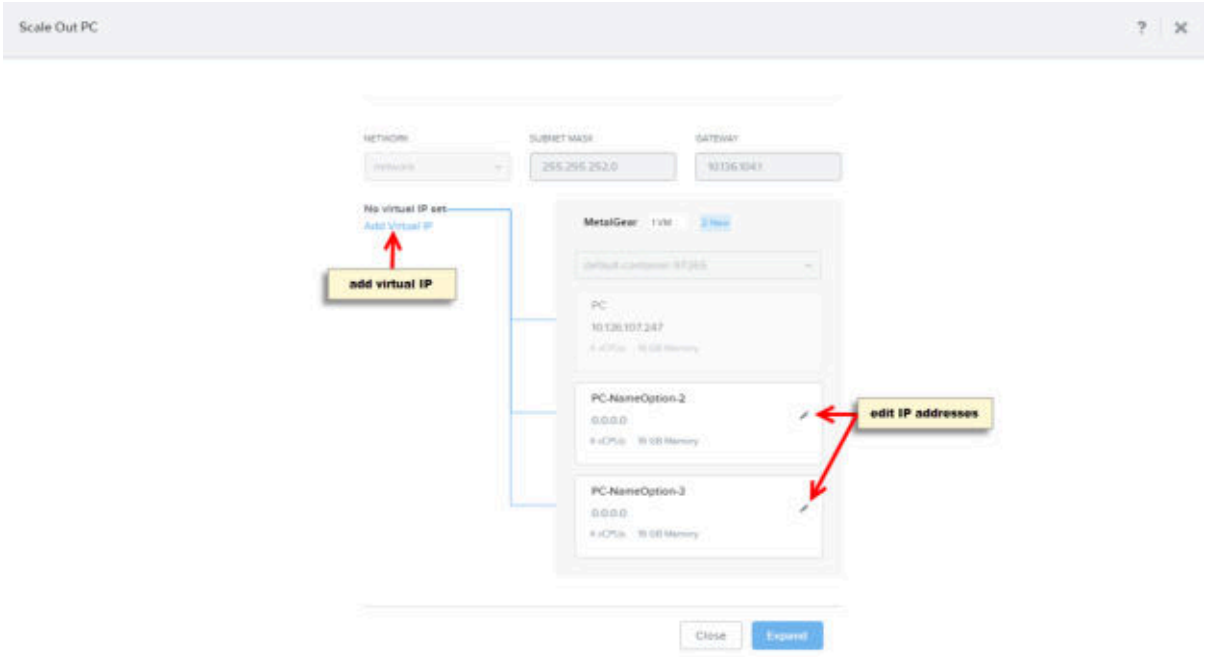

Figure 32: Scale Out PC Page

Note: The Prism Central VM image is deployed from the target cluster. If the required image cannot be accessed, typically because there is no Internet access (such as at a dark site), a

pop-up error message appears. In this case you have the option to manually download and deploy the image as follows:

1. Log on to the Nutanix customer support portal, click **Downloads** > **Prism Central**, and click the download link for the target version to save the Prism Central binary .TAR and metadata .JSON files on your local media.
2. Log on (using SSH) to any Controller VM in the cluster specified in the error message and copy the Prism Central binary .TAR and metadata .JSON files to the Controller VM.
3. Run the following command to deploy the Prism Central image:

   nutanix@cvm$ ncli software upload software-type=PRISM_CENTRAL_DEPLOY file-path=*file_path* meta-file-path=*metadata_file_path*

   The *file_path* is the full (absolute) path to the .TAR file, and the *metadata_file_path* is the full path to the .JSON file. After this step completes, you can continue the scale out procedure.

a. Review the **Network**, **Subnet Mask**, and **Gateway** fields, which display the network name, subnet mask value, and gateway IP address applied to this Prism Central instance.

   The values in these three fields are read-only and cannot be changed.

b. [optional] To specify a virtual IP address for Prism Central, click the **Add Virtual IP** link.

   This displays the Add PC Virtual IP window. A virtual IP can be used as a single point of access for Prism Central. Enter the IP address in the Virtual IP field and then click the **Update** button.

c. Specify IP addresses for the two new Prism Central VMs.

   1. Click the pencil icon for one of the new VMs. (The VM names are set automatically.) The **IP field** opens for editing. Enter the IP address and then click the green check mark icon to save that address.

Figure 33: Edit IP Field

   2. Repeat this step for the second new VM.

   The vCPU count and memory size in the new VMs are fixed and match the current Prism Central VM values; you cannot change these values.

d. When all the parameters are correct, click the **Expand** button.

   This starts the process of creating the new VMs and deploying this Prism Central instance. You can monitor progress from the Tasks page (see Tasks View on page 190).

What to do next

In some cases, especially when using Calm, a version mismatch with a linked product could occur after expanding Prism Central. To fix this situation, perform a life cycle manager (LCM)

inventory after completing the Prism Central expansion procedure. See the LCM documentation for instructions on how to perform an LCM inventory.

Shutdown (Start Up) Prism Central VM

A Prism Central VM is managed like any other VM in a cluster through the Prism web console (Prism Element) for the cluster in which the Prism Central VM is running.

> CAUTION: Prism Central supports features that could be damaged by shutting down the Prism Central VM at the wrong time. Do not shutdown a Prism Central VM until you are certain it can be done safely. If in doubt check with Nutanix customer service before attempting a shutdown.

- To shutdown (or start up) a Prism Central VM from the Prism web console, see the "Managing a VM (AHV)" or "Managing a VM (ESXi)" section in the *Prism Web Console Guide*.
- To shutdown (or start up) a Prism Central VM running on Hyper-V, see the appropriate Hyper-V documentation.

IP Address Reconfiguration

In certain cases, you might need to reconfigure the IP address and gateway of a Prism Central VM. You can use this procedure to reconfigure the IP address and gateway of single Prism Central instances as well as Prism Central VMs in a Scale out Prism Central (clustered PC VMs or PC cluster).

Preparing to Reconfigure the IP Address and Gateway of PC VMs

Perform the tasks described in this topic before you start the IP address reconfiguration procedure.

About this task

You must perform the following tasks while the Prism Central VMs are still on the existing IP addresses.

Procedure

1. Coordinate Prism Central downtime, because the features and functionality of Prism Central will be unavailable for the entire duration of the IP address reconfiguration procedure.
2. Create a table to map the existing IP addresses with the new IP addresses of the Prism Central VMs for your reference.
3. Use SSH to log on to any running Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.
4. Verify if the PC VM or PC cluster is in a stable state.

   ```
   nutanix@pcvm$ cluster status
   ```

   This command is valid even if you only have a single PC VM and do not have a PC cluster.
5. Run the NCC health checks to make sure that the PC cluster is in a healthy state.
6. If you want to reconfigure the IP addresses of the entities in the AOS clusters (such as IP addresses of CVMs and hypervisor hosts), on which the PC is hosted, reconfigure those IP addresses first before you reconfigure the IP addresses of the PC VMs.

7. Stop the PC VM or PC cluster.

```
nutanix@pcvm$ cluster stop
```

This command is valid even if you only have a single PC VM and do not have a PC cluster.

Wait to proceed until an output showing all the services as DOWN is displayed, except the Zeus and Scavenger services.

Reconfiguring the IP Address and Gateway of Prism Central VMs

Before you begin

Ensure that you have completed the tasks described in Preparing to Reconfigure the IP Address and Gateway of PC VMs on page 56.

About this task

The procedure to reconfigure the IP addresses of PC VMs differs depending on whether or not IP Address Management (IPAM) is enabled in an AHV cluster. The procedure to reconfigure the IP addresses of PC VMs in an ESXi cluster is the same as the procedure of an AHV cluster in which IPAM is disabled.

IP Address Management (IPAM) is a feature of AHV that allows it to assign IP addresses automatically to VMs by using DHCP. See *IP Address Management* section in the *AHV Administration Guide*.

> Note:
>
> In this document, "managed network" refers to an AHV network that has IPAM enabled, and "unmanaged network" refers to an AHV network that does not use IPAM.

Perform the following procedure to reconfigure the IP address of a single PC VM or PC VMs in a PC cluster.

> Note:
>
> Perform steps 3, 4, and 5 only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

Procedure

1. Run the external_ip_reconfig script.

```
nutanix@pcvm$ external_ip_reconfig
```

2. Follow the prompts to type the new netmask, gateway, and external IP addresses.

   A message similar to the following is displayed when the procedure begins:

   ```
   External IP reconfig started
   ```

   A message similar to the following is displayed if the procedure is completed successfully:

   ```
   External IP reconfig finished successfully. Restart all the CVMs and start the cluster.
   ```

   The message indicates that you must restart the CVMs. In this case, you are reconfiguring the IP addresses of the PC VMs, so you must restart the PC VMs and not the CVMs.

   > Note:
   >
   > If your PC VMs are running in a managed network of an AHV cluster, do not start the PC VMs yet but proceed to step 3.
   >
   > If your PC VMs are in an unmanaged network of an AHV cluster or are in an ESXi cluster, skip steps 3, 4, and 5, and proceed to step 6 directly.

   A message similar to the following is displayed if the procedure fails:

   ```
   External IP reconfig Failed
   ```

   If the procedure fails, run the following command to check the log for troubleshooting purposes:

   ```
   nutanix@pcvm$ cat ~/data/logs/ip_reconfig.log
   ```

3. (Optional) Use SSH to log onto any CVM of the AOS cluster that is hosting the PC cluster or PC VM.

   > Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

4. (Optional) Update the NIC attached to the PC VM with the new IP address of the PC VM.

   > Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   ```
   nutanix@cvm$ acli vm.nic_update pc-vm-name pc-vm-mac request_ip=true ip=pc-vm-ip
   ```

   Replace the variables with their appropriate values as follows:

   - *pc-vm-name*: Name of the Prism Central VM whose IP address you modified.
   - *pc-vm-mac*: MAC address of the Prism Central VM whose IP address you modified (press **Tab** after you type the name of the PC VM to automatically populate the MAC address).
   - *pc-vm-ip*: New IP address of the Prism Central VM.

5. (Optional) Log on to the Prism Element web console and modify the IP address range in the IP pool for the PC network to work with the modified IP address of Prism Central VMs. In the Prism Element web console, perform the following:

   > Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   a. In the **VM** dashboard, click the **Table** view.

   b. Select a **PC VM** and click **Update**.

   c. In the **Update VM** dialog box, scroll down to the **Network Adaptors (NIC)** section.

   d. In the **VLAN Name** column, note the names of the networks attached to your PC VM.

   e. Click **Close**.

   f. Click the gear icon in the top-right corner, and under **Settings**, click **Network Configuration**.

   g. In the **Network Configuration** dialog box, identify the network names you noted in step d.

   h. Select each network and click the pencil icon next to the network.

   i. Under **IP Address Pools**, modify the IP address range if the new IP address of the PC VM is in a different IP address range than the previous IP address range.

6. Restart each Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.

7. Once the PC VMs are powered on, use SSH to log on to each PC VM or the single PC VM if you do not have a PC cluster and verify if the output of the following verification script displays the new IP addresses:

   ```
   nutanix@pcvm$ external_ip_reconfig_verify.py
   ```

8. Start the PC cluster.

   ```
   nutanix@cvm$ cluster start
   ```

9. Perform the following verification steps:

   a. Log on to the Prism Element web console of every cluster attached to the PC VM whose IP address you changed.

   b. On the **Home** page, verify if the **PC Registration** status is displayed as **OK**. This might take a few minutes after the PC VM or PC cluster is powered on.

3

MAIN DASHBOARD

The main dashboard is the opening screen that appears after logging into Prism Central. It provides a dynamic summary view across the registered clusters. To view the main dashboard at any time, select **Dashboard** from the Entities menu (see *Entities Menu* section).

> Important: Nutanix recommends using Google Chrome web browser to access Prism Central. The display of dashboard and widgets is optimized for Google Chrome browser.

Main Screen Layout

In addition to the main menu options (see *Main Menu* section), the main dashboard includes a menu bar and a display section for a collection of widgets that appear as tiles with targeted information about the registered clusters in each tile. The menu bar includes the following options:

- **Main Dashboard** tab. Click this tab to display the main dashboard. Additional tabs appear for any custom dashboards you create.
- **Manage Dashboard** button. Click this button to create a custom dashboard, edit the name of a custom dashboard, or delete a custom dashboard (see Creating a New Dashboard on page 69).
- **Reset Dashboard** button. Click this button to reset the main dashboard to the default set of widgets.
- **Add Widget** button. Click this button to add a widget to the displayed dashboard (see Adding Dashboard Widgets on page 70).

- **Data Density** button. Click this button to select the density of data representation in the widgets. Select one of the following three options:
  - Light—The information is lightly packed in the widget with more spacing between the elements.
  - Default—The information is moderately packed with sufficient spacing between the elements. This scheme of data density is based on user research by Nutanix.
  - Dense—The information is densely packed with minimal spacing between the elements for essential legibility.

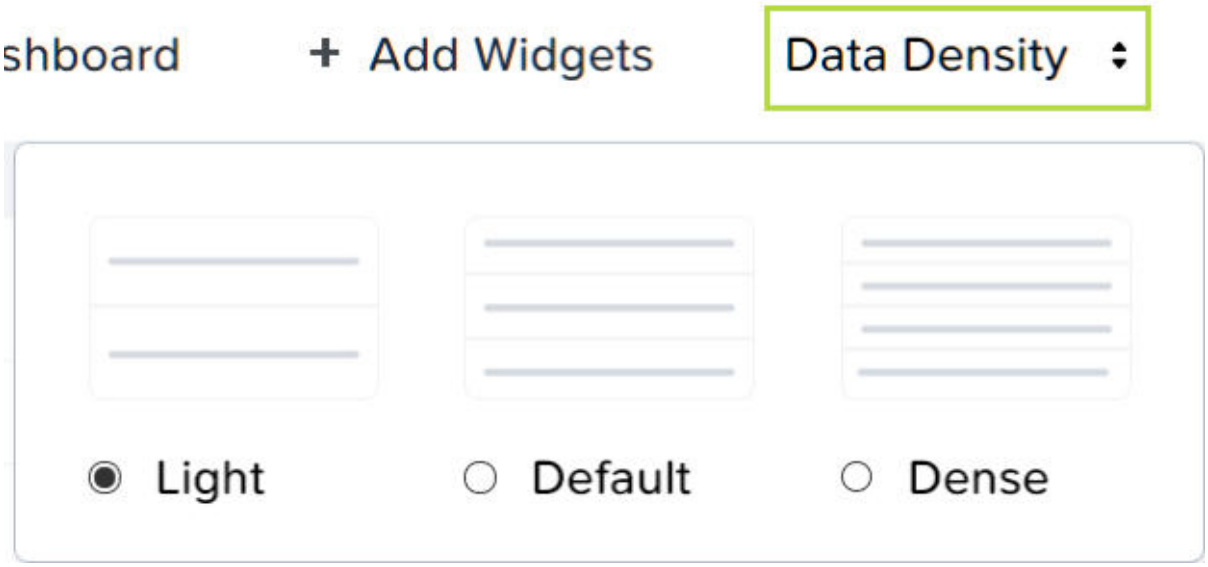

Figure 34: Data Density

Note: **Manage Dashboards** and **Add Widget** are Prism Pro license features. If Prism Pro is disabled, these buttons disappear, and you cannot customize or add dashboards. Any custom dashboards created before Prism Pro was disabled remain, but clicking the **Reset All** button (which replaces the **Reset Dashboard** button) deletes all custom dashboards and returns the main dashboard to the default state.

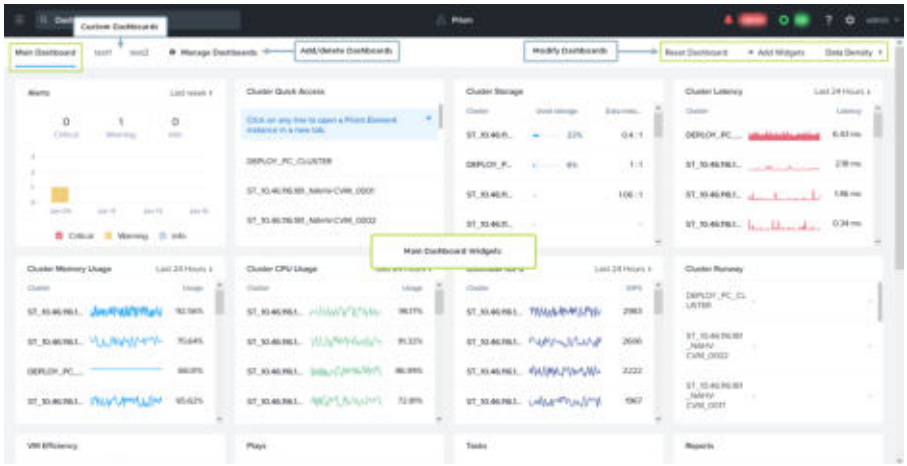

Figure 35: Main Dashboard

Main Dashboard Display

The main dashboard displays thirteen information tiles (widgets) by default. You can customize this view by adding or deleting widgets as desired (see Modifying a Dashboard on page 75).

The widgets are lined up in rows of four widgets per row. The default widgets are described in this section as they appear from left to right, top to bottom on the Main Dashboard page. A sample view of the Main Dashboard with the default widgets is displayed in the figure.

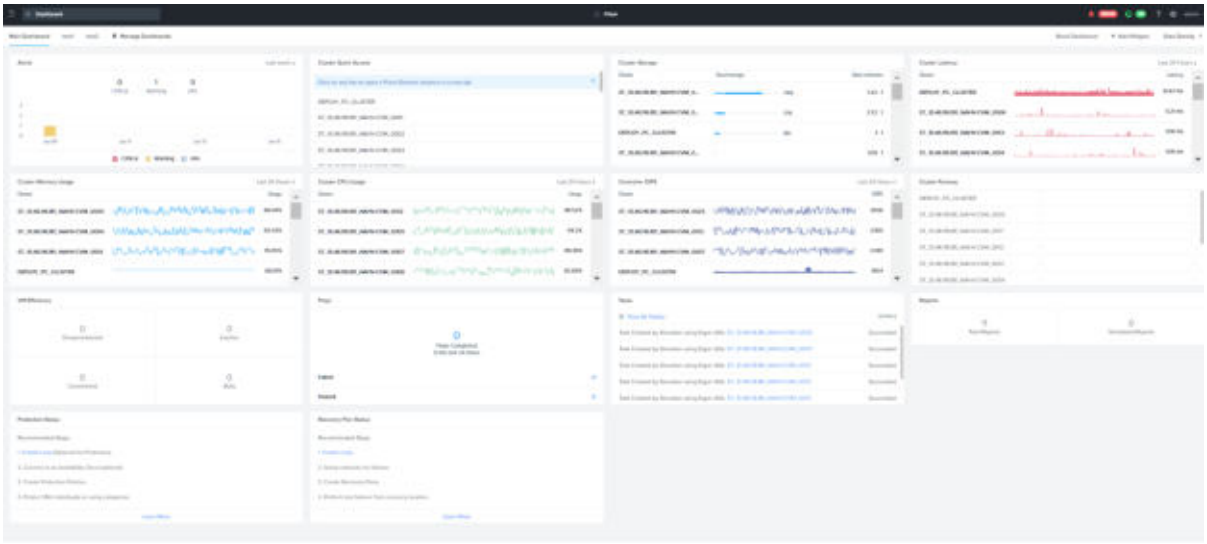

Figure 36: Main Dashboard with Default Widgets

The following figure displays the sample view of a widget with a period (time) selection:

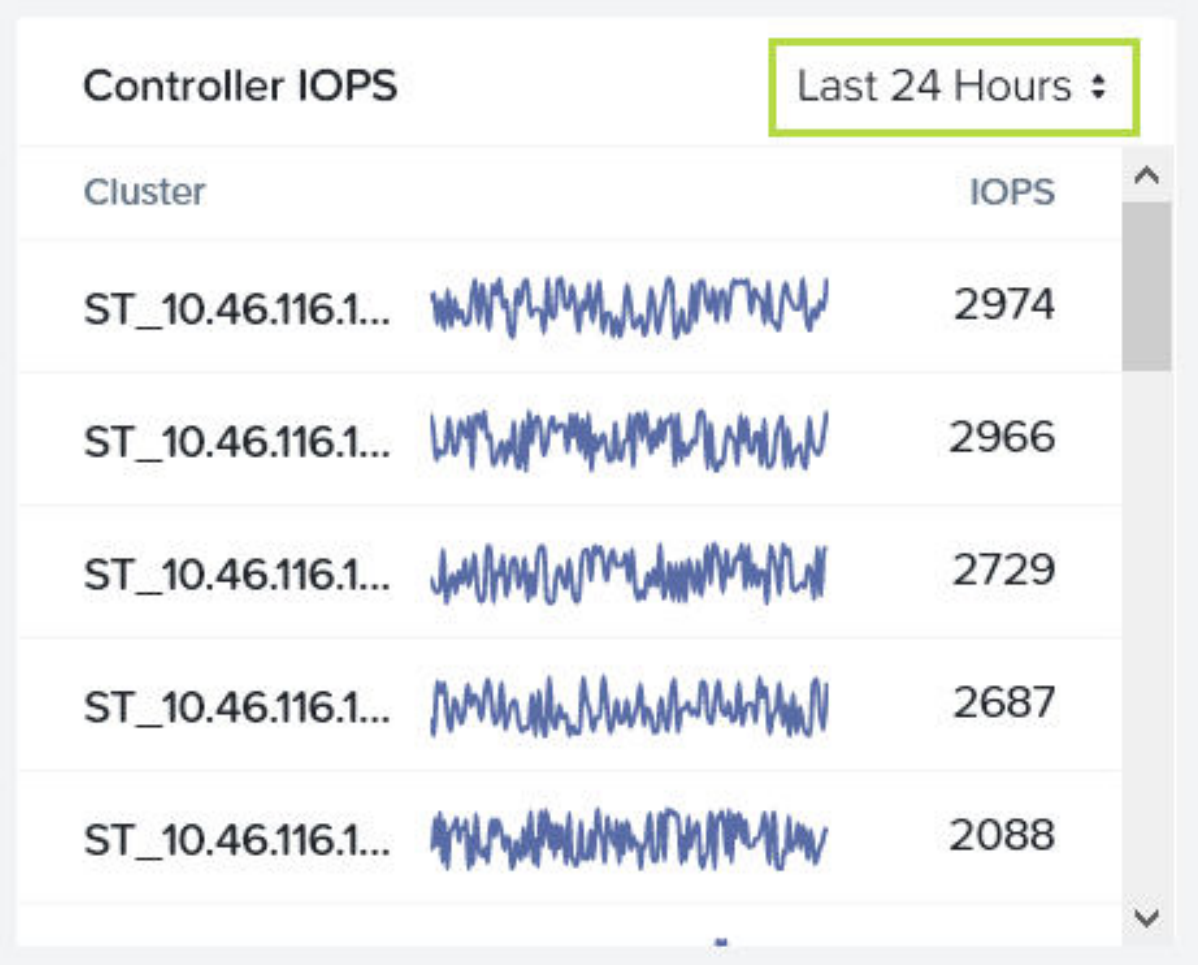

Figure 37: Sample Widget

- **Alerts**. This widget displays colored bar graph representing the alerts raised in the period selected. You can select the period on the drop down list at the top right corner of the widget. Each category of alerts is represented as a colored bar. Click any bar on the graph to see a list of the alerts represented in the bar. Click the **View All Alerts** link below the list to see all the alerts on the Alerts page.

  The Alerts categorized as **Critical**, **Warning** and **Info** are represented as red, yellow and blue bars respectively in the graphs.

- **Cluster Quick Access**. This widget displays a list of the clusters managed by the Prism Central. It allows you to open the Prism Element for a registered cluster in a new tab or window by simply clicking the line for that cluster. The cluster name appears on the left of each line while current health and alert icons (as appropriate) appear on the right.

- **Cluster Storage**. This widget displays storage and resiliency information for the highest usage clusters. Each line includes the cluster name, a usage column with a bar that visually indicates how much capacity is used currently, and a data resiliency column that displays the current status (critical, warning, OK, or unknown). The cluster list is ordered by data resiliency status with critical clusters at the top of the list. Click the cluster name to see the cluster details page. Hovering the cursor over the usage bar displays a box with usage details including data reduction ratio, usage percentage, used space, total space, fault domain (disk, node, or rack), and fault tolerance level (0, 1, or 2). The data reduction ratio indicates the data savings due to data reduction techniques such as deduplication,

compression, and erasure coding. (A 1:1 ratio indicates none of these data reduction techniques are in use currently.)

- **Cluster Latency**. This widget displays the total (read and write) IO latency average for the highest latency clusters. Clicking on a cluster name displays the summary page for that cluster.
- **Cluster Memory Usage**. This widget displays the percentage of total memory in use currently for the highest usage clusters. Clicking on a cluster name displays the summary page for that cluster.
- **Cluster CPU Usage**. This widget displays the percentage of total CPU in use currently for the highest usage clusters (or all clusters if there are fewer than five). Clicking on a cluster name displays the summary page for that cluster.
- **Controller IOPS**. This widget displays the total (read and write) controller IOPS for the highest volume clusters. Clicking on a cluster name displays the summary page for that cluster (see). The IOPS number comes from the controller when the hypervisor is AHV or Hyper-V and from the hypervisor when the hypervisor is ESXi.
- **Cluster Runway**. This widget alerts you to potential storage, CPU, or memory resource constraints across the clusters and provides an estimated runway (time remaining) before the resources are maxed out based on current usage. Clicking on a cluster name displays the capacity page for that cluster.
- **VM Efficiency**. This widget displays the number of VMs that are considered inefficient broken down by category: overprovisioned, inactive, constrained, and bully. It includes a link to the VMs dashboard for details about these VMs (see *VM Summary View*). See *Behavioral Learning Tools* in the *Prism Central Guide* for more information about VM efficiency.
- **Plays**. This widget displays a list of the Plays running on the cluster. The list is categorised to completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the Plays dashboard page.
- **Tasks**. This widget displays a list of recent tasks with the current status of each task. Click on the View All Tasks(s) link to view the Tasks page.
- **Reports**. This widget displays a table that lists the number of total and scheduled reports with a link to the reports dashboard (see *Reports Management* in the *Prism Central Guide*).
- **Protection Status**. This widget displays the protection plan status if one has been created. If not, it displays a set of recommended steps to create a protection plan.
- **Recovery Plan Status**. This widget displays the recovery plan status if one has been created. If not, it displays a set of recommended steps to create a recovery plan.

Widgets on a Dashboard

Resize the Widgets

You can vertically resize of a widget by clicking and dragging down the resize handle at the bottom of the widget. You cannot resize a widget horizontally.

The following figure provides an example of the resize handle.

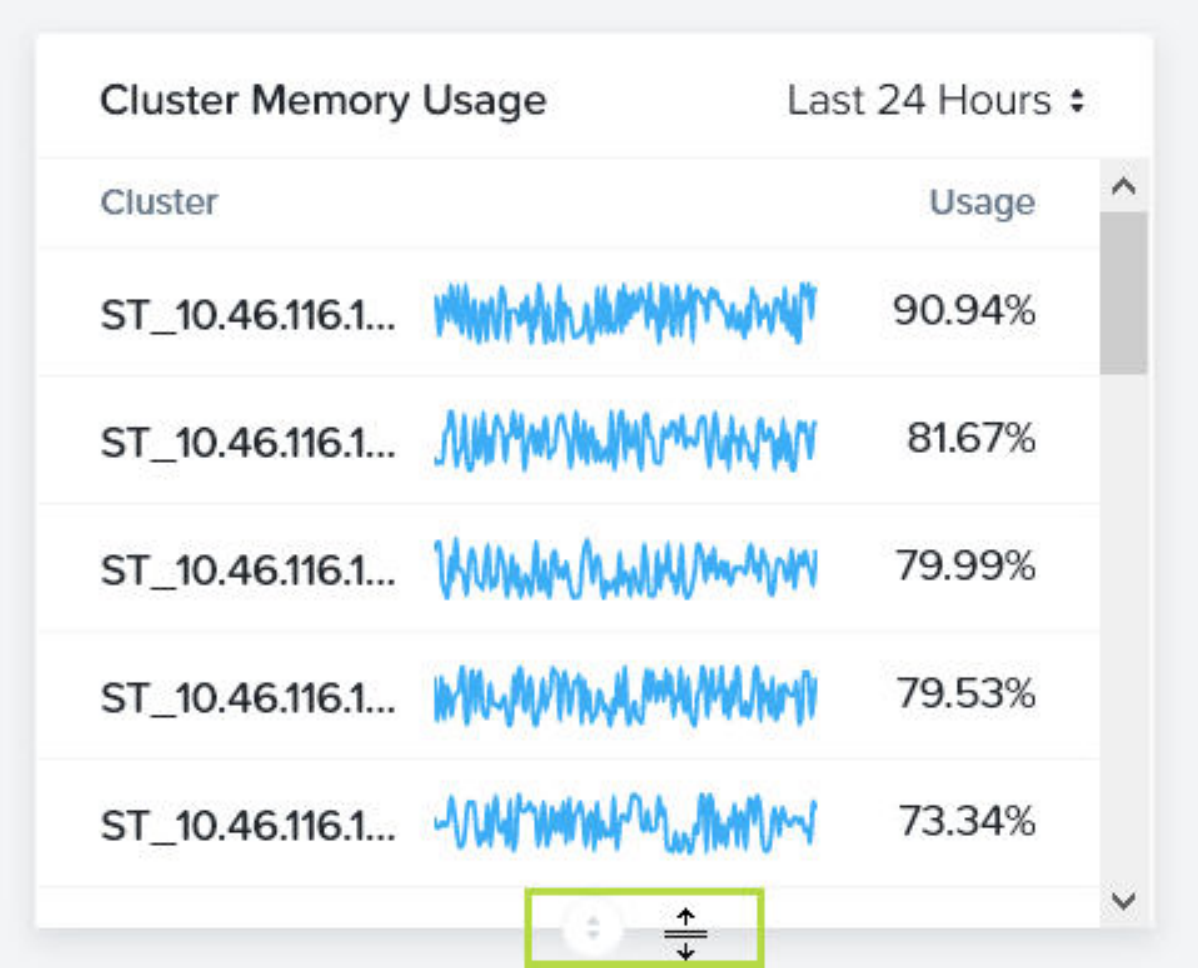

Figure 38: Vertically Resizing the Widget

When you click and drag up or down the resize handle to resize a widget, the widget placed below the resized widget shifts upward or downward.

You can resize a widget in multiples of half the default length of the widget, such as half, one-and-half, twice, two-and-half times the default length. Therefore, you can reduce the length to half or expand it to one-and-half times or twice the default length. Even if you attempt to resize it in any other variation of the default length, the widget auto-sizes itself to the nearest half length.

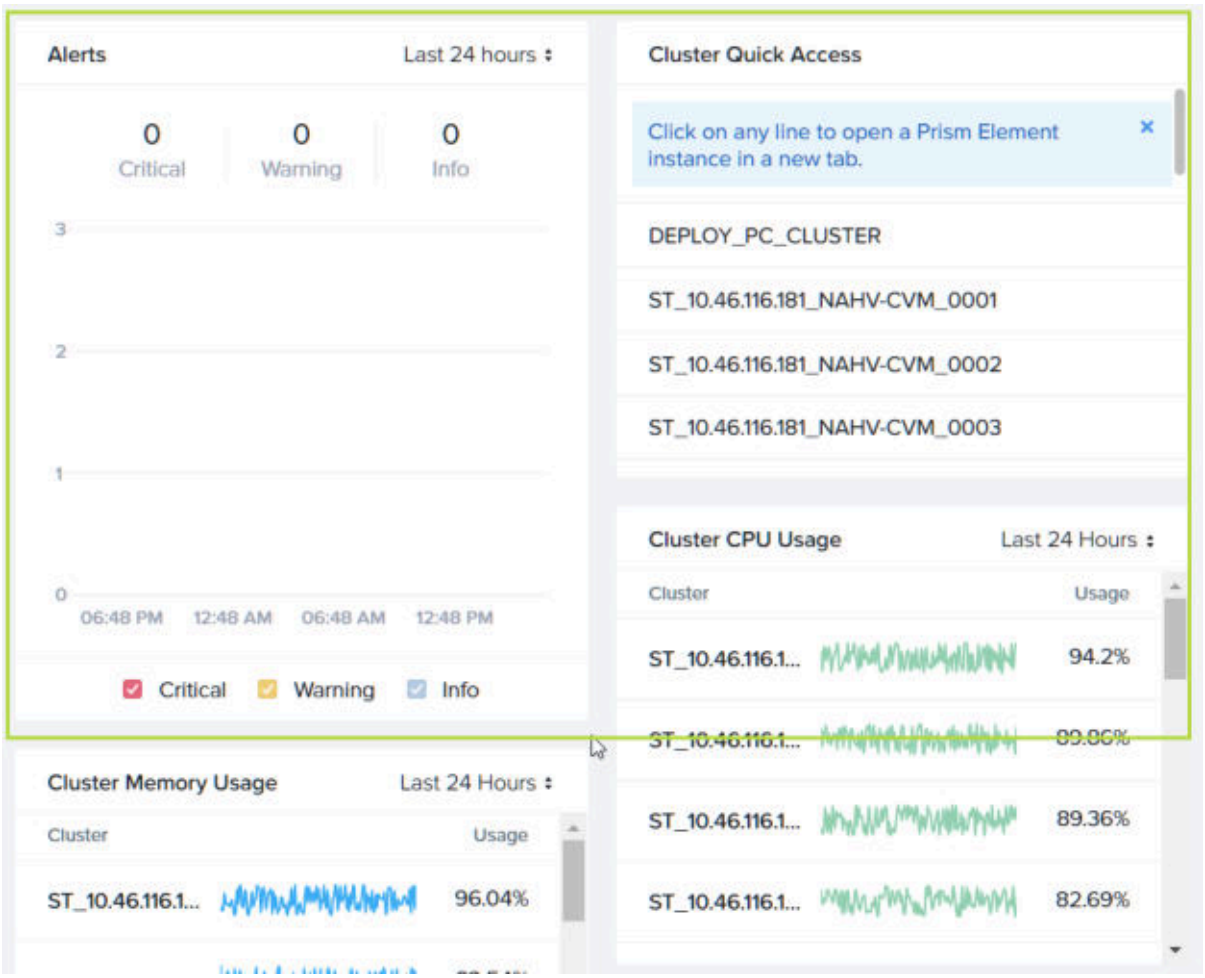

Figure 39: Widget Resize to One-and-half Times the Default Length

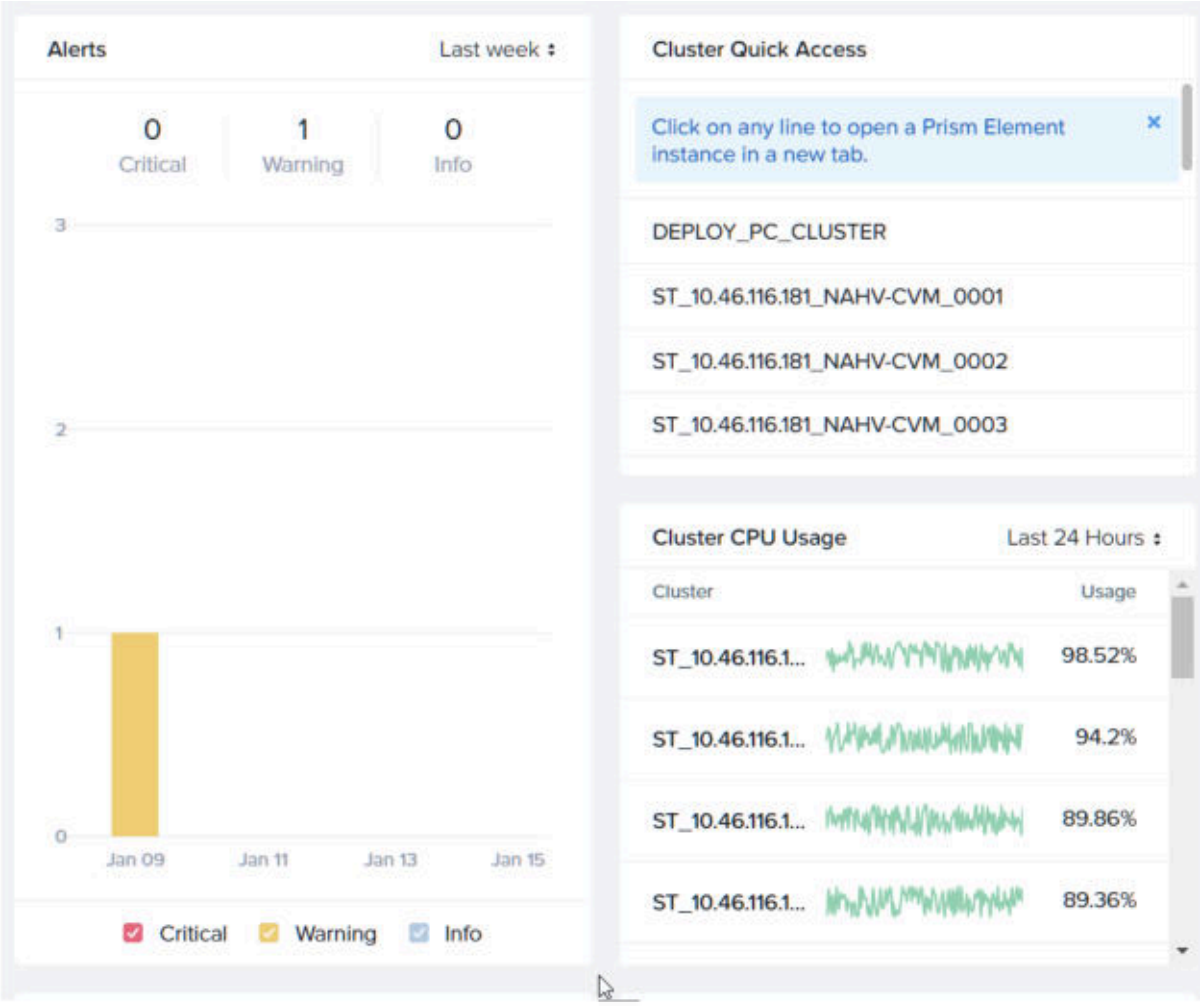

Figure 40: Widget Resize to Twice the Default Length

Drill-down or Expanding Data Elements in a Widget

You can expand or drill-down the data presented in the widgets by clicking on the data element. The following conditions are applicable:

- Expanding a data element in a widget is not possible in all widgets.
- In widgets that provide a list, clicking the list item or row opens an entity page in Prism Central or another application such as Prism Web Console in a new browser tab. For example, in the Cluster Quick Access widget, clicking the cluster name in the list opens the Prism Element Web Console for that cluster in a new browser tab. In the Cluster Storage widget, clicking a cluster name opens the cluster details view for that cluster.
- In widgets that provide numbers, clicking a number opens the respective entity page in Prism Central. For example, in the Reports widget, clicking the **Total Reports** or the **Scheduled Reports** number opens the Reports page. Similarly, clicking

- If the widget presents charts, click on the chart to display an expanded or drill-down view of the chart. Do not click on the entity (such as cluster name). If you do so, instead of displaying the drill-down view of the chart, the entity details page is displayed.

  The figure displays an example of expanding a chart in a widget.

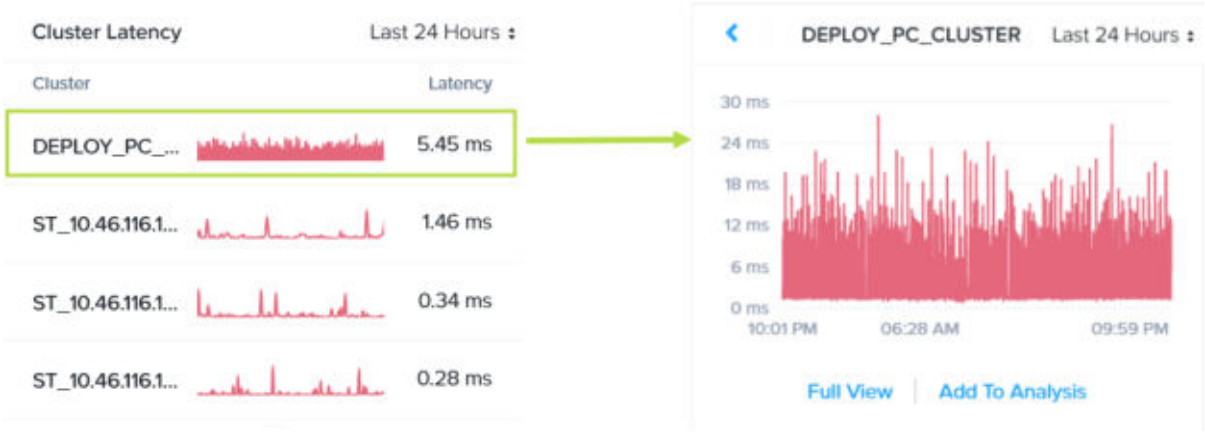

Figure 41: Widget Data Drill-down

At the bottom of the drill-down view, you can click **Full View** to open the details page of the entity(cluster or VM). Click **Add To Analysis** to add the chart to the Analysis dashboard (See *Analysis Dashboard* in *Prism Central Guide*.

Hover Details in Drill-down Views

After you click on a chart in a widget, the drill-down view of the chart is displayed. You can hover on any poin on the time chart to display the parameters at that point.

In the same example as presented in the Drill-down or Expanding Data Elements in a Widget on page 67, if you hover over the chart, the following sample details are displayed:

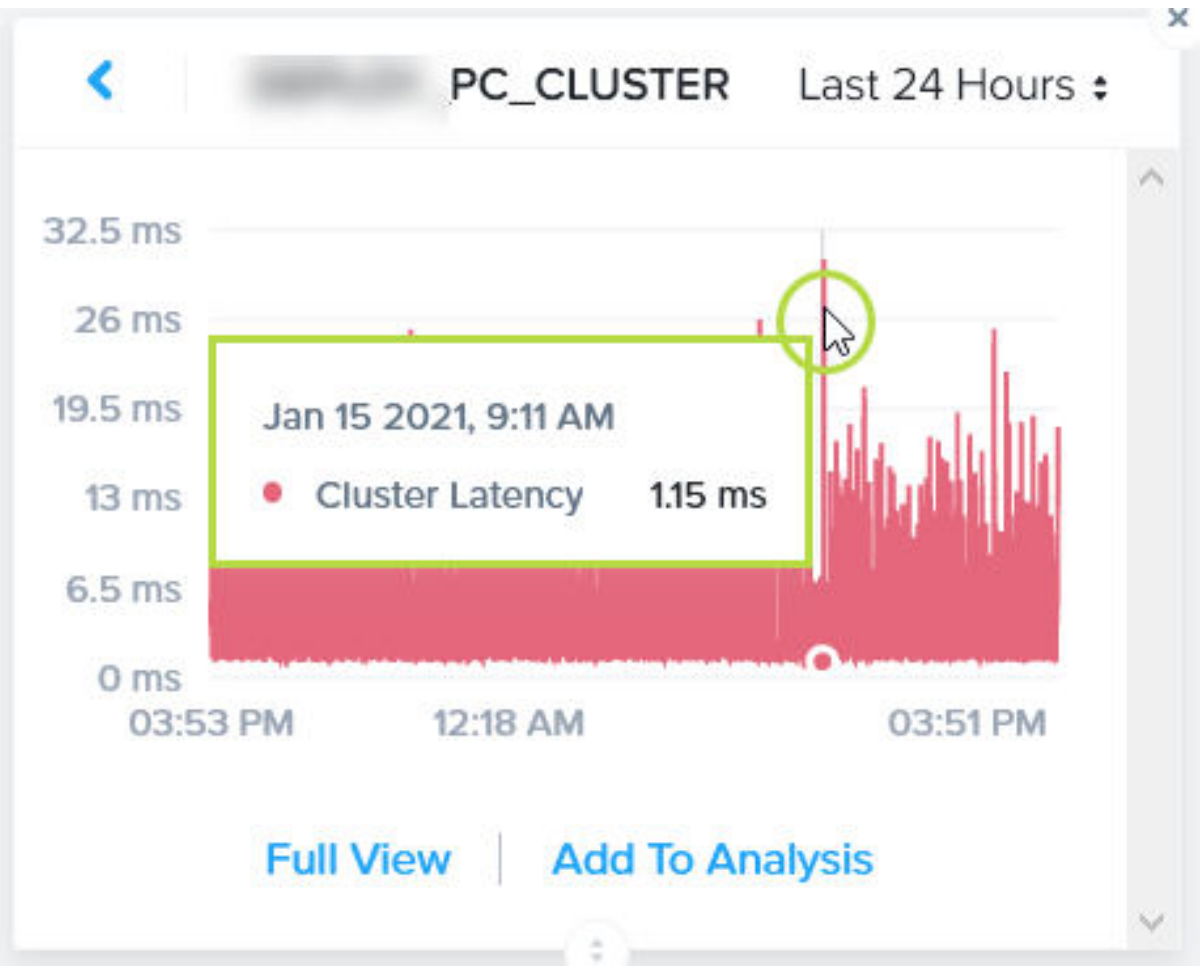

Figure 42: Hover Details

Creating a New Dashboard

About this task

The Main dashboard provides a default view into the clusters, but you can add custom views by creating custom dashboards. To create a custom dashboard, do the following:

> Note: Creating a dashboard requires a Prism Pro license. The **Manage Dashboards** button does not appear if Prism Pro is disabled.

Procedure

1. In the Main dashboard (see Main Dashboard on page 60), click the **Manage Dashboards** button.

2. In the Manage Dashboards window (see Modifying a Dashboard on page 75), do the following:

   a. Click the **New Dashboard** button.

   b. Enter a name for the dashboard in the displayed field.

   c. Click the **Save** button.

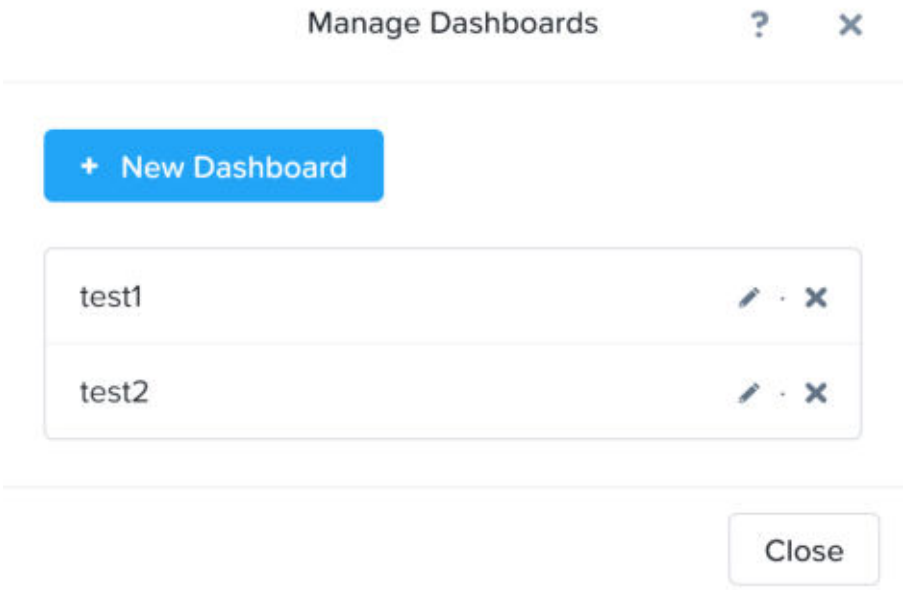

Figure 43: Manage Dashboards Windows (new dashboard)

3. Click the **Close** button to close the Manage Dashboards window.
   A tab for the new dashboard appears next to the **Main Dashboard** (or other custom dashboard) tab.

4. Click the tab for the new dashboard (to display that dashboard) and then click the **Add Widgets** button.
   The new dashboard is empty (no widgets) initially.

5. In the Add Widgets screen, select a widget and add it to the dashboard (see Adding Dashboard Widgets on page 70).
   Repeat this step until all the desired widgets have been added to the dashboard.

Adding Dashboard Widgets

Prism Central provides a selection of widgets that can be added to the home or a custom dashboard. Each widget provides information about a specific resource or usage topic that is displayed in a tile on the dashboard.

To add a widget to a dashboard, do the following:

> Note: Adding a widget to a dashboard requires a Prism Pro license. The **Add Widgets** button does not appear if Prism Pro is disabled.

1. Go to the target (main or custom) dashboard and click the **Add Widgets** button on the far right.

2. In the **Add Widgets** page, select the desired widget in the left column. (To find a particular widget, enter the name in the search field.) A preview of the selected widget appears in the middle column.
3. Enter values for the configurable parameters in the right column. The parameters vary by widget; some widgets do not have any configurable parameters.
4. Click the **Add to Dashboard** (or **Add & Return to Dashboard**) button to add the widget to the dashboard.

Figure 44: Add Widget Screen

The following table describes the widgets that you can add to a dashboard.

Table 10: Widget Parameters

Widget	Description	Parameter	Values
Custom Widgets			
Custom Alerts Widget	Displays a custom list of alerts.	Widget Name - Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster - Select the cluster(s) to monitor from the pull-down list. The default is all clusters.	All Clusters, <cluster name>
		Choose Size - Select the size (width/height dimensions) of the widget. The default is 1 x 1.	1 x 1, 2 x 1, 2 x 2, 2 x 3, 2 x 4, 4 x 4
		Severity. The default is critical.	Critical, Warning, Info

Widget	Description	Parameter	Values
		Entity. Select the entity of interest. The default is cluster.	Cluster, Storage, VM, Hardware, DR
Top Lists Widget	Displays a list of the top performers for a selected metric.	Widget Name - Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster - Select the cluster(s) to monitor from the pull-down list. The default is all clusters.	All Clusters, <cluster name>
		Choose Size - Select the size (width/height dimensions) of the widget. The default is 2 x 1.	2 x 1, 2 x 2, 2 x 3
		Entity. Select the entity of interest. The default is host.	host, VM
		Metric. Select the metric of interest. The default is IOPS.	IOPS, Memory Usage, CPU Usage, Bandwidth, Latency
Custom Chart Widget	Displays a performance graph for a selected entity and metric pair.	Widget Name - Enter a name for the widget. The default is "New Chart Widget".	(name)
		Entity Type. The default is host.	Host, Disk, Storage Pool, Storage Container, Virtual Machine, Virtual Disk
		Entity. (There is no default value.)	(entity name)
		Metric. (There is no default value.)	(metric name)
Cluster Info Widget	Displays cluster summary information about alerts, anomalies, runway, and inefficient VMs.	Widget Name - Enter a name for the widget. The default is "New Widget".	(n/a)
		Select a Cluster. Select the target cluster from the pull-down list.	(registered cluster names)
Virtual Infrastructure Widgets - This widget appears on the Main dashboard by default (see Main Dashboard on page 60).			

Widget	Description	Parameter	Values
VM Efficiency	Displays a four-cell table that lists the number of overprovisioned, inactive, constrained, and bully VMs in the cluster with links to the details.	(no customizable parameters)	(n/a)
Hardware Widgets - All these widgets (except Performance) appear on the Main dashboard by default (see Main Dashboard on page 60).			
Cluster CPU Usage	Displays the percentage of total CPU in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Latency	Displays the total (read and write) IO latency average for the highest latency clusters.	(no customizable parameters)	(n/a)
Cluster Memory Usage	Displays the percentage of total memory in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Quick Access	Displays a list of registered clusters. Health and alert icons appear for each cluster. Clicking the cluster line opens Prism (element) for that cluster in a separate tab or window.	(no customizable parameters)	(n/a)
Cluster Runway	Displays storage, CPU, and memory runway estimates (time remaining before the resource reaches capacity). See the *Capacity Runway Tab* section in *Resource Planning* for more runway information.	(no customizable parameters)	(n/a)
Cluster Storage	Displays storage statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Controller IOPS	Displays the total (read and write) controller IOPS for the highest volume clusters.	(no customizable parameters)	(n/a)
Impacted Cluster	Displays information about any clusters that are impacted (performance, capacity, or other potential issues) and may need attention.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Performance	Displays latency, bandwidth, and IOPS statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Alerts	Displays colored bar graph representing the alerts raised in the period selected. Clicking the **View All Alerts** link below the list to see all the alerts on the Alerts page.	(no customizable parameters)	(n/a)
Activity Widgets - The Tasks widget appears on the Main dashboard by default (see Main Dashboard on page 60).			
Tasks	Displays a list of recent tasks with the current status of each task and a link to the Tasks dashboard.	(no customizable parameters)	(n/a)
Discovered Apps	Displays the number of applications discovered in a specified set of monitored clusters. Click the **App instances** number to display the application discovery dashboard. Click the **Identified** or **Unidentified** number to display the application discovery dashboard filtered for that condition. This widget displays the discovered apps information only if you enable App Discovery.	(no customizable parameters)	(n/a)
Operations Widgets - These widgets appear on the Main dashboard by default (see Main Dashboard on page 60).			
Plays	Displays a list of the Plays running on the cluster. The list is categorised to completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the Plays dashboard page.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Reports	Displays a table that lists the number of total and scheduled reports with a link to the Reports dashboard (see).	(no customizable parameters)	(n/a)
Disaster Recovery Widgets - These widgets appear on the Main dashboard by default (see Main Dashboard on page 60).			
Protection Status	Displays the protection plan status if one has been created. If not, it displays a set of recommended steps to create a protection plan.	(no customizable parameters)	(n/a)
Recovery Plan Status	Displays the recovery plan status if one has been created. If not, it displays a set of recommended steps to create a recovery plan.	(no customizable parameters)	(n/a)

Modifying a Dashboard

About this task

The Main dashboard provides a default view into the registered clusters, but you can customize that view at any time. To modify the Main dashboard or any other dashboard you create, do the following:

> Note: Customizing a dashboard requires a Prism Pro license. The **Add Widgets** and **Manage Dashboards** buttons do not appear if Prism Pro is disabled.

Procedure

1. Go to the Main dashboard (see Main Dashboard on page 60) and click the tab of the dashboard to modify (main or previously added custom dashboard).
2. To add a widget to the displayed (main or custom) dashboard, click the **Add Widgets** button, select a widget from the Add Widgets screen, and add it to the dashboard (see Adding Dashboard Widgets on page 70).
3. To delete a widget, click the X icon in the upper right of the displayed tile.

   A prompt appears to verify the delete; click the **OK** button. The widget then disappears from the dashboard.
4. To reset the Main dashboard to the default set of widgets (after you have previously added or deleted widgets), click the **Reset Dashboard** button.

   A prompt appears to verify the reset; click the **OK** button. The main dashboard returns to its default view.

5. To rename a custom dashboard, do the following:
   a. Click the **Manage Dashboards** button.
   b. Click the pencil icon for that dashboard.
   c. Enter a new name in the displayed field.
   d. Click the **Save** button.

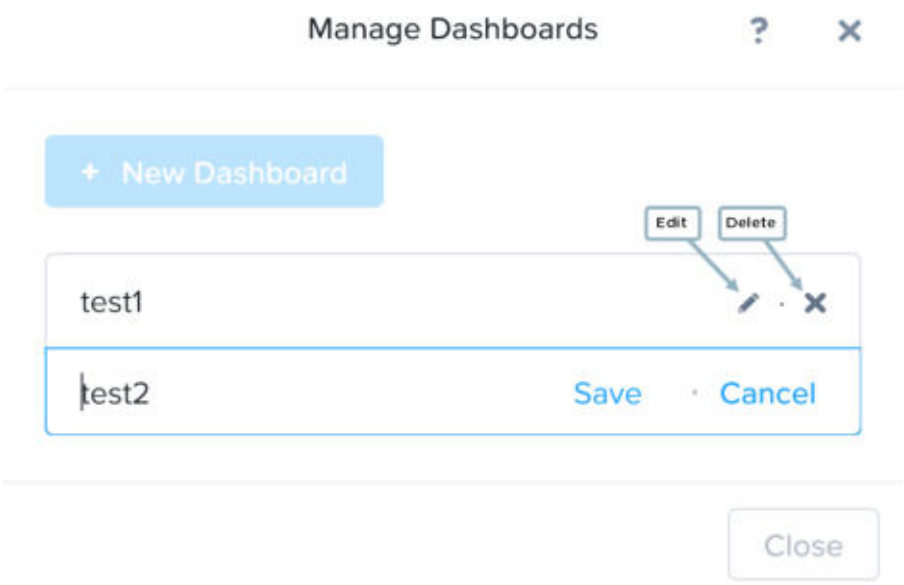

Figure 45: Manage Dashboards Window

6. To delete a custom dashboard, click the **Manage Dashboards** button and then click the X icon for that dashboard.
   A prompt appears to verify the delete; click the **OK** button. The tab for that dashboard disappears from the screen.

4

ENTITY EXPLORING

The entities menu (see Entities Menu on page 13) provides access to dashboards that allow you to view statistics about a range of entities across clusters and to organize that information in a variety of ways. An entity is an object type such as a VM, cluster, security policy, project, or report. Many of the entity dashboards include common elements in a summary view with drill-down options to view detailed information about individual entities of that type.

Note: This section describes common dashboard elements and how you can customize the view. See the appropriate "<entity> Summary View" section for a description of the specific fields and options applicable to each entity view. The Entities menu is organized into group categories, and the "<entity> Summary View" sections are located under one of the following group category headings:

- Virtual Infrastructure Entities on page 82 (VMs, storage containers, volume groups, vCenter datastores, catalog items, images, categories, OVAs)
- Network and Security Entities on page 135 (subnets, virtual private clouds, floating IPs, VPNs)
- Policies Entities on page 136 (security policies, NGT policies, image placement policies)
- Data Protection Entities on page 142 (protection summary, protection policies, recovery plans, VM recovery points)
- Hardware Entities on page 145 (clusters, hosts, disks, GPUs)
- Activity Entities on page 186 (alerts, events, audits, tasks)
- Operations Entities on page 193 (analysis, app discovery, cost management, integrations, ops policies, planning, playbooks, reports)
- Administration Entities on page 195 (LCM, projects, roles, users, availability zones)
- Services Entities on page 209 (Calm, Files, Foundation Central, Karbon, Objects)

These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

Screen Layout

Common display elements in many dashboards include the following:

- A selection pane on the left that lists the entity type and number at the top.
  - For hardware and virtual infrastructure entity types like clusters and VMs, the pane includes five tabs: Summary, List, Alerts, Events, Metrics. Click the desired tab to display that content in the main section of the screen.
  - For most other entity types, the pane is blank (no tabs), and the relevant content appears automatically in the main section of the screen.
  - For some entity types, the pane does not appear.
- When a list appears in the main section, the display may include:
  - A menu bar with (left to right) selection options, action options (if any), grouping options, and display options for the list, and a **Filters** button on the far right that when clicked displays the **Filters** pane (hidden by default).
  - A query field that identifies the filters (if any) being applied to the list. This field displays all filter options that are currently in use. It also allows for basic filtering on the entity name.
  - A table (list) of entities. What is included in the table is determined by the grouping, display, and filter options. Entities are displayed 20 per page by default (use the scroll bar to scroll a page), but you can specify a different number per page (10 to 60) from the pull-down list above the table (click "*X - XX of XXX*" to display list).
  - A download icon (just above the table on the right); click the icon to download the table contents in CSV format. You can download a maximum of 1000 rows.

Figure 46: Dashboard Layout (example)

Filter Options

When the main section displays a list of entities, you can filter the list by clicking the far right icon in the menu bar to display the **Filters** pane. This pane includes a set of fields that vary according to the type of entity. Select the desired field values to filter the list on those values. An entry appears in the search field (see Searching for Information on page 20) for each value you select. You can save a filter by clicking the start icon in the search field, or you can remove a filter by clicking the X for that value.

Numeric filters have **To/From** fields to specify a range. These fields can take numeric values along with units. For example, the filter adjusts the scale accordingly when you type in "10 K" or "100 M".

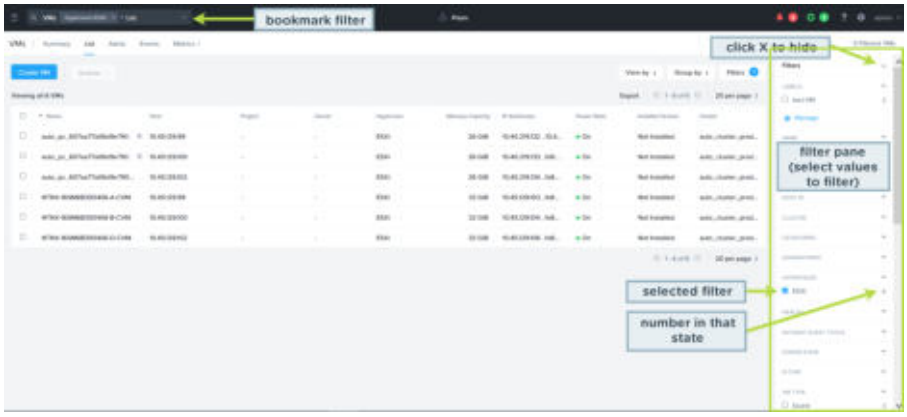

Figure 47: Filter Pane

Grouping Options

There are three (or four) drop-down menus on the right of the menu bar that allow you to organize the entity information in several ways.

- The **View by** menu allows you to specify the type of displayed information.

  Note: A general focus is available for all entities, but the other options are available only when appropriate for that entity.

  - Select **General** to display a set of general information parameters, as illustrated in the Explore Dashboard figure.
  - Select **Performance** to display a set of performance-specific parameters.
  - Select **Efficient** to display a set of efficiency-related parameters. "Efficiency" is determined through the VM behavioral learning engine (see Behavioral Learning Tools on page 582).
  - Select **GPU** to display a set of GPU-specific parameters.
  - Select **Data Protection** to display a set of data protection-related parameters.

  In addition to the preset **View by** displays, you can create one or more custom displays as follows:

  1. Select the **Add Custom** button at the bottom of the **Focus** menu.
  2. In the **<entity> Column** window, do the following:
     1. Enter a name for the custom display in the first (top) field.
     2. Select (click the blue + circle for) each entity property in the left column you want in your display. The selected entity properties appear in the right column. You can filter the entity-property list by entering a string in the search field above the left column.

You can select a maximum of 10 columns in one custom display including the Name column.

3. When the entity-property list in the right column is correct (you can remove an entity property from this list by clicking the red - circle for that entity), click the **Save** button.

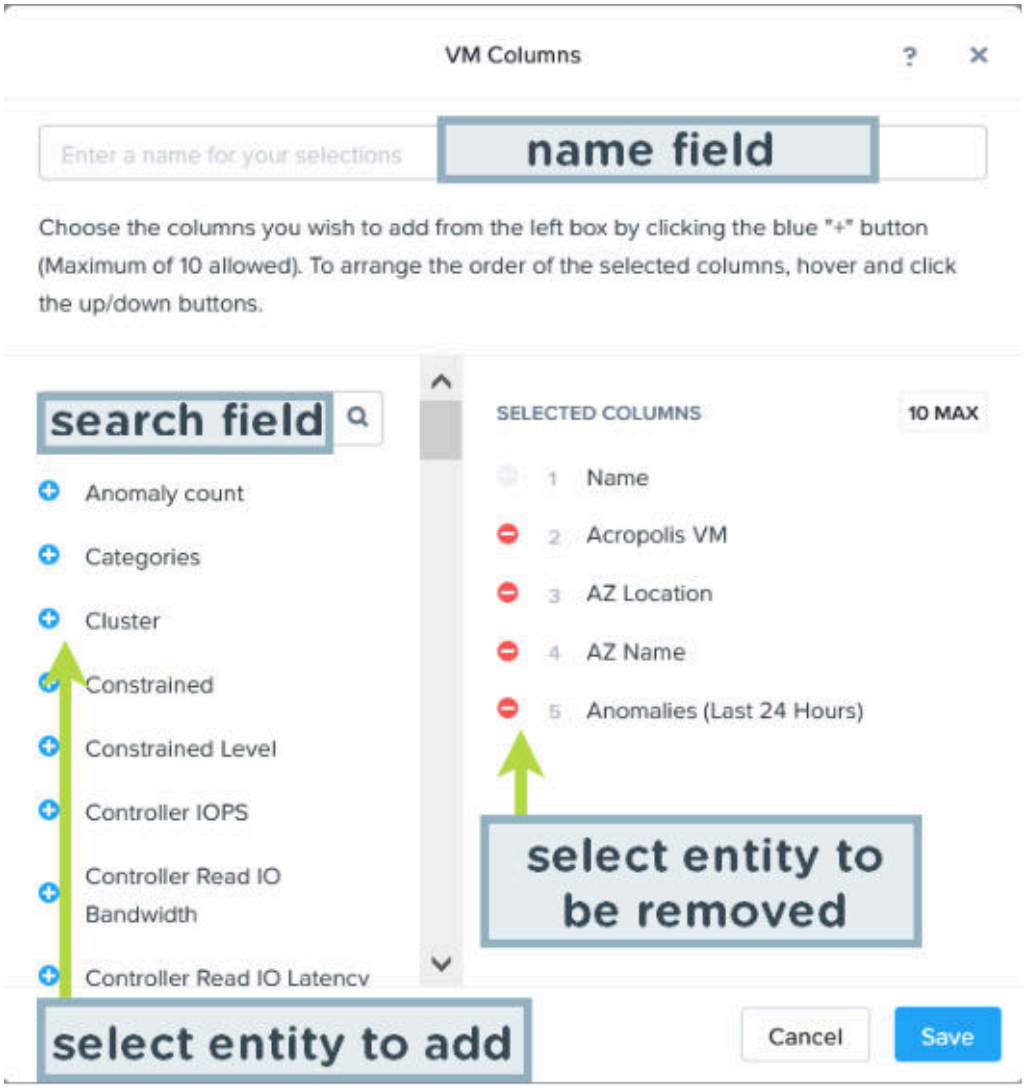

Figure 48: VM Column Window (example)

The **<entity> Column** window disappears and the new custom display appears (by name) in the **Focus** menu. To view that display at any time, select it from the menu.

> Note: Only the user that created a custom display can view that custom display. Other users and the administrative user cannot view that custom display.

- The **Color** menu allows you to color code the entries based on a selected parameter, as illustrated in the Circles View figure.
- The **Group** menu allows you to group the entries based on a selected parameter, as illustrated in the Circles View figure.

- The **Sort** menu allows you to sort the entries based on a selected parameter. This menu appears for the Tiles and Circles views only; ordering in the Grid (tabular) view is accomplished by clicking on a column header.

Action (and Selection) Options

The dashboard provides options to administer and categorize the entities. Actions are applied to selected entities that are currently visible. You can select a set of entities (rows) by selecting the first one and then using <Shift> select to select the last one.

- While some administrative tasks for a cluster must be done through Prism Element (see the *Prism Web Console Guide*), you can perform other administrative tasks directly from Prism Central. Action buttons for such tasks appear in the menu bar when an entity type is selected. For example, buttons for creating a VM and configuring the network appear in the VMs view. (No buttons appear when there are no relevant actions available for the selected entity type.)
- When one or more entities are selected from the list, an **Actions** menu appears in the menu bar. The menu includes actions that you can do to the selected entities. For example, the VMs menu includes actions such as power on or clone which will be executed on all the selected VMs. (Only currently valid actions are available; other menu options are inactive and appear gray.)
- When one or more entities are selected from the list, a label icon appears (next to the **Actions** menu). Labels allow you to create custom groupings for the entities. To add a label, click the icon and enter the label name in the field. The label is applied to all the selected entities in the list.

  > Note: Labels are available for VMs and clusters only.

- Use the selection menu to
  - Select all entries in the list by checking the box or choosing **Select all** from the pull-down list.
  - Unselect all by unchecking the box or choosing **Clear selections** from the pull-down list.
  - Show just the selected entries by choosing **Show selected entities** from the pull-down list.

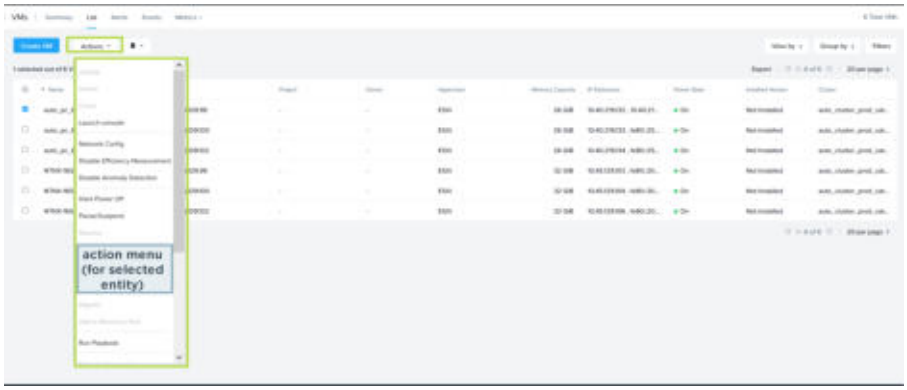

Figure 49: Action Options

Virtual Infrastructure Entities

You can access dashboards for the following virtual infrastructure entities from the **Virtual Infrastructure** category of the entities menu (see Entities Menu on page 13):

- VMs (see VMs Summary View on page 82)
- Storage Containers (see Storage Containers Summary View on page 97)
- Volume Groups (see Volume Groups Summary View on page 108)
- vCenter Datastores (see vCenter Datastores Summary View on page 117)
- Catalog Items (see Catalog Items View on page 124)
- Images (see Images Summary View on page 125)
- Categories (see Categories Summary View on page 128)
- OVAs (see OVAs View on page 131)

> Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

VMs Summary View

To access the VMs dashboard, select **Virtual Infrastructure** > **VMs** from the entities menu (see Entities Menu on page 13). The VMs dashboard summary view displays information about VMs across the registered clusters and allows you to access detailed information about each VM. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> Note: This section describes the information and options that appear in the VMs dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See VM Management on page 316 for information about creating and managing VMs.

Summary Tab

Clicking the **Summary** tab displays the following four widgets:

- **Suggested**: Displays a list of the VMs with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **Memory Usage**, and **IOPS**. Click the **View All XX VMs** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of VM-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

- **VM Efficiency**: Displays the number of VMs that are considered inefficient broken down by category (overprovisioned, inactive, constrained, and bully). This is the same widget that appears in the main dashboard (see Main Dashboard on page 60). See Behavioral Learning Tools on page 582 for more information about VM efficiency.

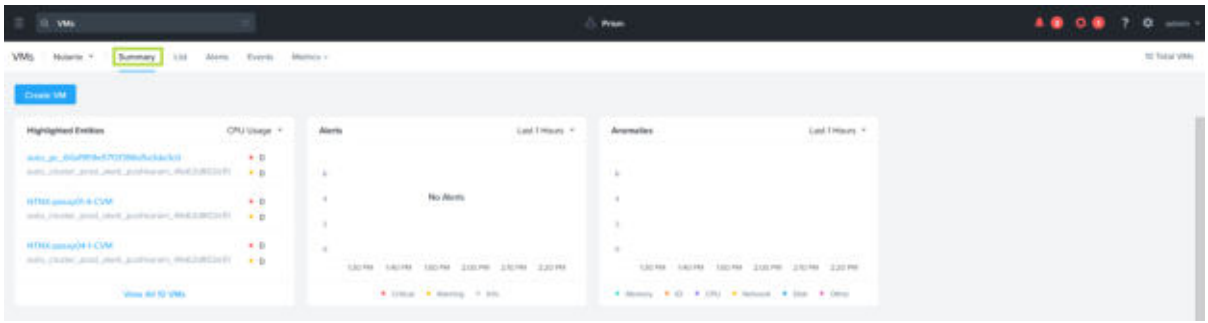

Figure 50: VMs Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the VMs across the registered clusters. The following table describes the fields that appear in the VMs list. The fields vary based on the **Focus** menu selection, which is **General**, **Performance**, **Efficiency**, or **GPU**. A dash (-) is displayed in a field when a value is not available or applicable.

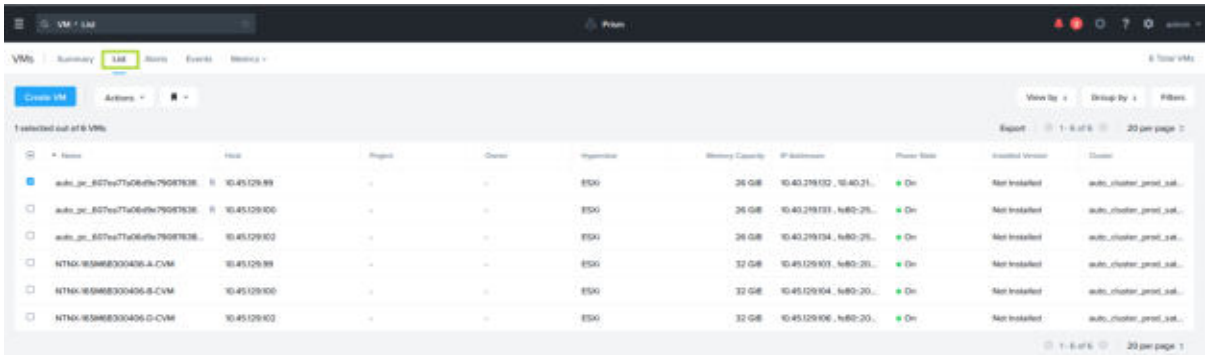

Figure 51: VMs List Tab

Table 11: VMs List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the VM name. Clicking on the name displays the details page for that VM (see VM Details View on page 90).	(VM name)
Host	Displays the host name. This field may be blank if the VM is powered off and a host is not assigned.	(host name)
Project	Displays the name of the project to which this VM belongs.	(project name)

Parameter	Description	Values
Owner	Displays the owner (user name) of this VM.	(user name)
Hypervisor	Displays the hypervisor type on which the VM is running.	AHV, ESX, or Hyper-V
Memory Capacity	Displays the total amount of memory available to this VM.	xxx [MB|GB]
IP Addresses	Displays the VM IP address(es).	(IP address)
Power State	Displays whether the VM is powered on or powered off	On (green), Off (red)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"Performance" Focus Fields		
Name	Displays the VM name.	(VM name)
Memory Usage	Displays the percentage of allocated memory capacity currently being used by this VM.	0 -100%
Controller Read IOPS	Displays read I/O operations per second (IOPS) for this VM.	(number)
Controller Write IOPS	Displays write I/O operations per second for this VM.	(number)
Controller I/O Bandwidth	Displays I/O bandwidth used per second for this VM.	xxx [MBps|KBps]
Controller I/O Latency	Displays the average I/O latency for this VM.	xxx [ms]
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"Efficiency" Focus Fields		
Name	Displays the VM name.	(VM name)
Efficiency	Displays the assessed efficiency of the VM. If a VM is performing in an expected range, the efficiency is listed as "Good". If not, the type of inefficiency is displayed as determined by the VM behavioral learning engine (see Behavioral Learning Tools on page 582).	Good, Constrained, Overprovisioned, Inactive, Bully; ("NA" if insufficient baseline data for categorization, which is typically 21 days)
Efficiency Detail	Displays why a VM is considered inefficient. A dash (-) appears for a "Good" VM.	(text message)
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"GPU" Focus Fields		

Parameter	Description	Values
Name	Displays the VM name.	(VM name)
GPU Configuration	Displays the GPU board (in case of Passthrough mode), vGPU software, type and number of vGPU instances in brackets (in case you configured the vGPU mode).	(GPU board/software-type name) Example: Passthrough - Nvidia Tesla M60 Example: vGPU - Nvidia GRID M60-8Q (2)
GPU Type	Displays the GPU configuration made as Passthrough or vGPU or if you did not configure GPU then "-"	[Passthrough|vGPU|-]
GPU Usage	Displays the percentage of GPU capacity being used by the VM.	(percentage)
GPU Framebuffer Usage	Displays the percentage of GPU framebuffer (RAM) capacity being used by the VM.	(percentage)

You can filter the VMs list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range. Click the plus (+) sign to the right of the parameter name to expand the fields for that parameter; click the minus (-) sign to contract those fields.

Table 12: Filter Pane Fields

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the pull-down list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)
Name	Filters on the VM name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the VM name condition/string. Note: In this and the following two fields, the condition menu options are **Contains**, **Doesn't contain**, **Starts with**, **Ends with**, and **Equal to**.	(VM name string)
Host	Filters on the host name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the host name condition/string.	(host name string)

Parameter	Description	Values
Cluster	Filters on the cluster name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the cluster name condition/string.	(cluster name string)
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a pull-down list appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of VMs tagged to each selected category is displayed on the right of the line.	(category name)
Hypervisor	Filters on the hypervisor type. Check one or more of the boxes to filter on those hypervisors. The number of VMs currently on each hypervisor type is displayed on the right of the line.	AHV, ESX, HyperV
Health	Filters on the VM health state (good, warning, or critical). Select one or more states to return a list of VMs in that state(s). The number of VMs currently in each state is displayed on the right of the line.	Critical, Warning, Good
Power State	Filters on the VM power state. Select one or more states to return a list of VMs in that state(s). The number of VMs currently in each state is displayed on the right of the line.	On, Off, Suspended, Paused, Unknown
VM Type	Filters on the type of VM. Select either user VM or Controller VM (or both). The number of VMs of each type is displayed on the right of the line.	User VM, AHV Controller VM
Memory Usage	Filters on the amount of memory capacity being used. Check the box for the desired percentage range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of VMs utilizing memory in that range.	([xx] to [yy]% range)
Read IOPS	Filters on the read IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of VMs with read IOPS in that range.	([xx] to [yy] range)
Write IOPS	Filters on the write IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of VMs with write IOPS in that range.	([xx] to [yy] range)

Parameter	Description	Values
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of VMs with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of VMs with average I/O latency in that range.	([xx] to [yy] range)
Over Provisioned	Filters for over-provisioned VMs. Check the box(es) for the desired type (high and moderate). Note: This and the following two fields filter based on a VM efficiency algorithm, which is part of the VM behavioral learning capabilities (see Behavioral Learning Tools on page 582).	High, Moderate
Constrained	Filters for constrained VMs. Check the box(es) for the desired type (high and moderate).	High, Moderate
Efficiency	Filters for certain VM profiles. Check the boxes for the desired profile types. There is one for efficient VMs (good) and four for inefficient VMs (bully, over-provisioned, constrained, inactive).	Bully, Over Provisioned, Constrained, Inactive VM, Good
GPU Configuration	Filters for GPU configuration information such as model name. Enter the GPU configuration information in the field and then check the box. As you type a pull-down list appears to help you select the correct configuration information.	(configuration info)
GPU Type	Filters for GPU operational mode. Check the box for one or more of the GPU types.	vGPU, Passthrough, Passthrough(Compute)
GPU Usage	Filters on the amount of GPU capacity being used. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of GPUs in that range.	([xx] to [yy]% range)
GPU Framebuffer Usage	Filters on the amount of GPU framebuffer (RAM) capacity being used. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of GPUs in that range.	([xx] to [yy]% range)
vGPU Guest Driver Version	Filters on the guest driver version. Enter the guest driver version number in the field.	(guest driver version number)

Export VM information. You can export the table containing the list of VMs and their information to a file in a CSV format by clicking the export icon on the right.

You can group the VM list in the following ways:

- The **Group** pull-down menu allows you to group the VM entries by cluster, hypervisor, power state, virtual CPU count, or health state. (You can only choose one.)
- When you select one or more VMs, the **Label** icon appears (to the left of the **Actions** menu). To assign the selected VMs a group label, click the **Label** icon and do one of the following:

Figure 52: Label Icon

  - To assign an existing label (if there are existing labels), select the desired label and then click **Apply changes**.
  - To assign a new label, enter a meaningful label in the blank text field and then click **Create new label**.

  You can now perform any actions available on this labeled group from the **Actions** menu (see following section). You can edit or delete the label by clicking the **Label** icon and selecting **Manage Labels**.

The VMs view includes two action buttons:

- To create a VM, click the **Create VM** button. You can create a VM on a cluster running AHV or ESXi (see Creating a VM (AHV) on page 317 or Creating a VM (ESXi) on page 357).
- To configure the network, click the **Network Config** button (see Configuring Network Connections on page 423).

The **Actions** menu appears when one or more VMs are selected. It includes the following actions:

- These actions can be applied to multiple VMs: **Delete**, **Power on**, **Power off**, **Pause/Suspend**, **Resume**, **Protect**, **UnprotectManage Categories**, **Quarantine VMs**, **Unquarantine VMs**, **Enable NGT**, **Disable NGT**, **Manage Ownership**
- These actions can be applied to only one VM at a time: **Update**, **Clone**, **Launch console**, **Snapshot,**, **Migrate**, **Configure VM Host Affinity**, **Add to Catalog**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the current state of the selected VM(s). See Managing a VM (AHV) on page 332 or Managing a VM (ESXi) on page 362 for instructions on how to perform these actions.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just VM-related alerts across the registered clusters (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just VM-related events across the registered clusters (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the VMs. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant performance information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

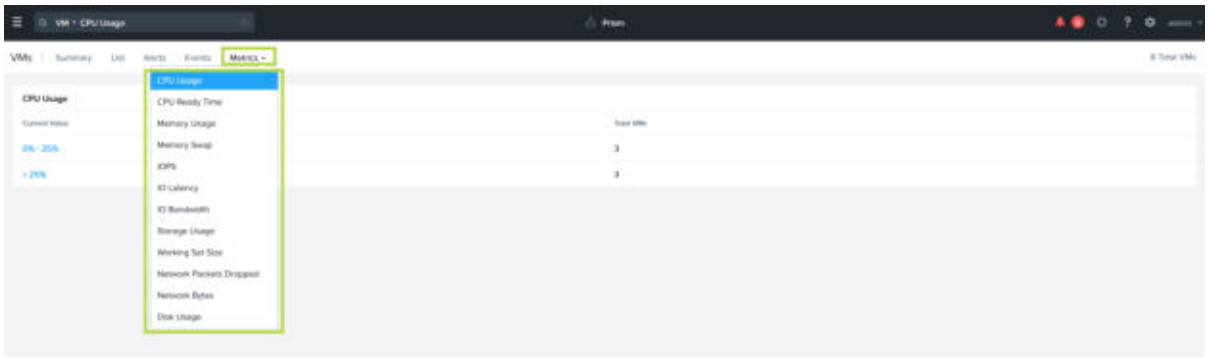

Figure 53: VMs Metrics Tab

Table 13: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total VMs (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those VMs. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
CPU Ready Time	Displays a CPU ready time percentage usage table.
Memory Usage	Displays a memory percentage usage table.
Memory Swap	Displays memory swap-out and swap-in rate tables.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.
Usage	Displays total, snapshot, and shared storage size tables.
Working Set Size	Displays total, snapshot, and shared working set size tables. Working set is the amount of memory that a VM requires in a given time interval.
Network Packets Dropped	Displays tables for the number of transmitted and received packets dropped.
Network Bytes	Displays tables for the amount of transmitted and received bytes (in GiBs).

VM Details View

To access the details page for a VM, go to the VMs **List** tab (see VMs Summary View on page 82) and click the VM name. You can also access the details page by clicking the VM name wherever that name appears, such as in a dashboard widget or search result.

The VM name and the following set of tabs appear on the left: Summary, Console, Data Protection, Alerts, Events, Metrics, NICs, Disks, Snapshots, and Categories. Click a tab to display that information on the right. (Click the **Back to VMs** link to return to the VMs summary view.)

Note: VirtIO must be installed in a VM for AHV to display correct VM memory statistics.

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the VM (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or

Last week from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

- Action button (above the widgets). Click the appropriate button to execute that administrative action on the VM. The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM. See Managing a VM (AHV) on page 332 for instructions on how to perform each action.

  Note: You can perform administrative actions on VMs in Acropolis managed clusters only. Unavailable actions are grayed out.

- 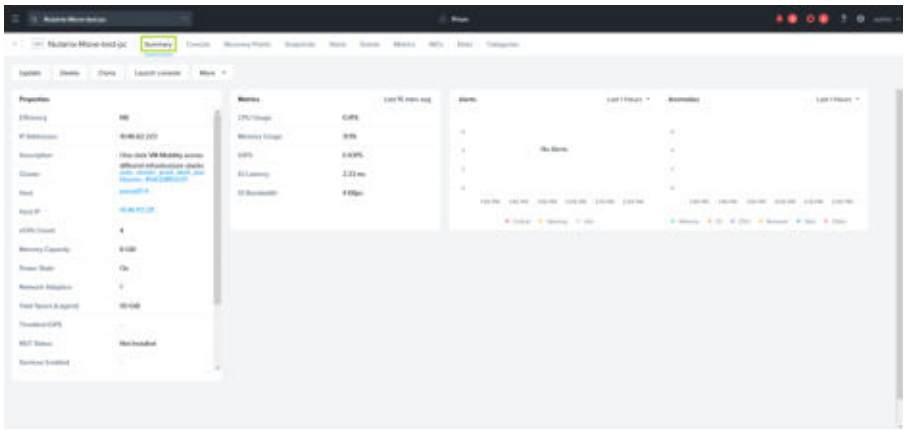

Figure 54: VM Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 14: VM Properties Fields

Parameter	Description	Values
Efficiency	Displays the efficiency state for this VM. If the efficiency is not good, an additional field may appear that specifies the problem. For example, if the VM is constrained, a Constrained field appears that identifies the constrained resource such as the CPU or memory. Note: The Anomalies, Efficiency, Overprovisioned, and Constrained parameters relate to the VM behavioral learning feature. See Behavioral Learning Tools on page 582 for a description of each parameter.	Bully, Over Provisioned, Constrained, Inactive VM, Good
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Host	Displays the host name. This field may be blank if the VM is powered off and a host is not assigned.	(host name)

Parameter	Description	Values
Host IP	Displays the host IP address.	(IP address)
Virtual CPU Count	Displays the number of virtual CPUs assigned to this VM.	(number)
Memory Capacity	Displays the amount of memory available to this VM.	xxx [MB|GB]
IP Addresses	Displays the IP address(es) assigned to the VM.	(IP address)
Power State	Displays whether the VM is powered on or powered off	On, Off
Network Adapters	Displays the number of network adapters available to this VM.	(# of adapter ports)
Disk Capacity	Displays the total disk capacity available to this VM.	xxx [GB|TB]
NGT Status		
Services Enabled	Displays the services enabled for this VM.	Calm, Karbon
NGT Installed Version		
NGT Cluster Version		
(the following fields appear when the VM is allocated to a GPU)		
GPU Type	Displays the GPU operational mode. If it is vGPU, the following fields also appear.	vGPU, Passthrough, None
GPU Configuration	Displays the vGPU profile used.	(vGPU profile name)
Framebuffer	Displays the size of the GPU framebuffer (RAM).	xxx GiB
Virtual Slice	Displays the virtual slice applied. The "virtual slice" reflects the approximate amount of physical GPU resources that the vGPU can receive. Note: The Virtual Slice and vGPU Guest Driver Version fields do not appear for passthrough GPUs.	(slice amount)
vGPU Guest Driver Version	Displays the version number of the vGPU guest driver.	(version number)

Console Tab

The **Console** tab displays the VM console screen. There are three icons above the console display (on the right).

- Click the left (three small boxes) icon to send a **Control-Alt-Delete** command to the console.
- Click the middle (camera) icon to take a screen shot of the console display.
- Click the right (box with an arrow in it) icon to open the console in a new window.

Recovery Points

The **Recovery Points** tab displays a list of recovery points (backup snapshots) when backups have been enabled. (The list is blank if there are no snapshots available.) The total number of recovery points and the latest and oldest recovery points are listed on the left. A list of all recovery points appears in a table on the right with the create time, location, expiry time, and recovery point type provided for each recovery point.

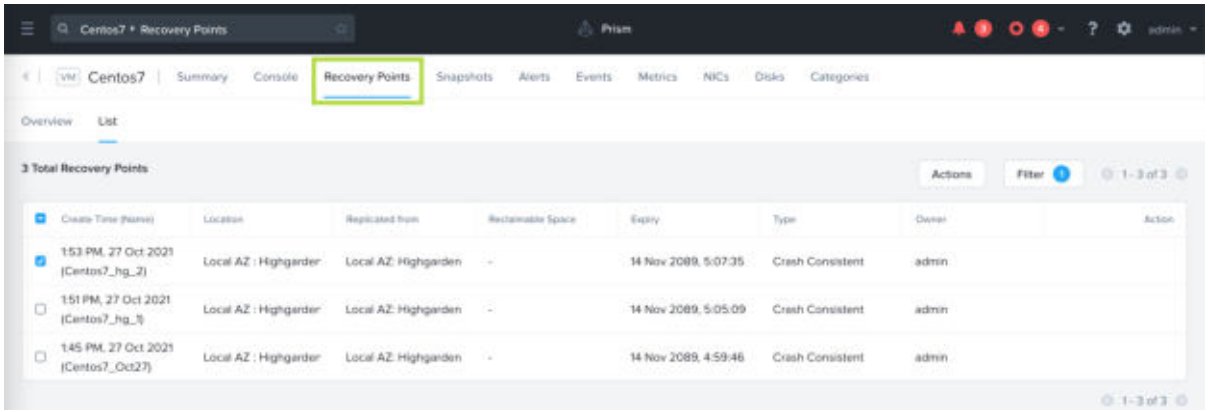

Figure 55: Recovery Points Tab

You can restore or replicate a VM from a recovery point. Select the recovery point and click the required action from the **Actions** menu:

- Click **Replicate** to replicate a VM from the selected recovery point either locally or remotely in a state of a chosen recovery point.
- Click **Restore** to restore the VM from the selected recovery point.
- Click **Delete** to delete the recovery point.

Snapshots Tab

The **Snapshots** tab displays information in tabular form about backup snapshots of the VM. Each line represent a snapshot, and the following information is displayed for each snapshot:

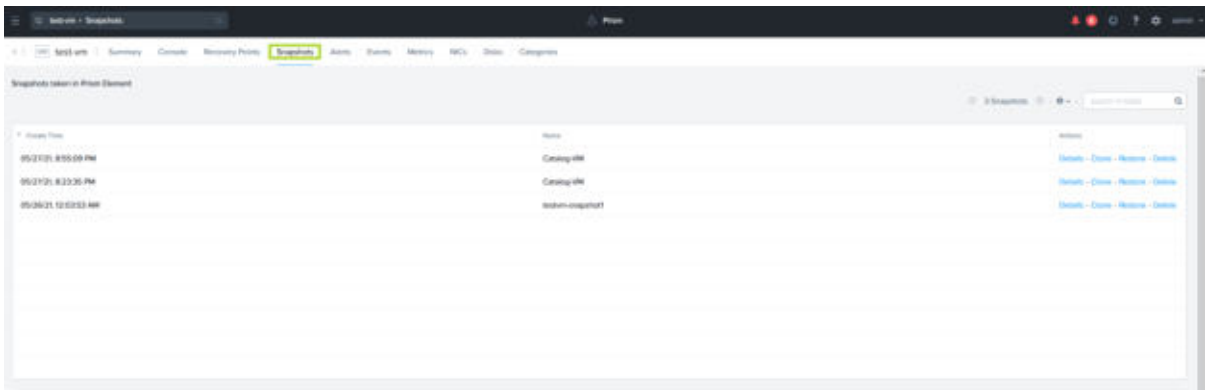

Figure 56: Snapshot Tab

- **Create Time**. Displays the time the backup snapshot was created (completed).
- **Name**. Displays a name for the backup if one was created.

- **Action**. Displays four action links:
  - Click the **Details** link to open a window that displays the snapshot details.

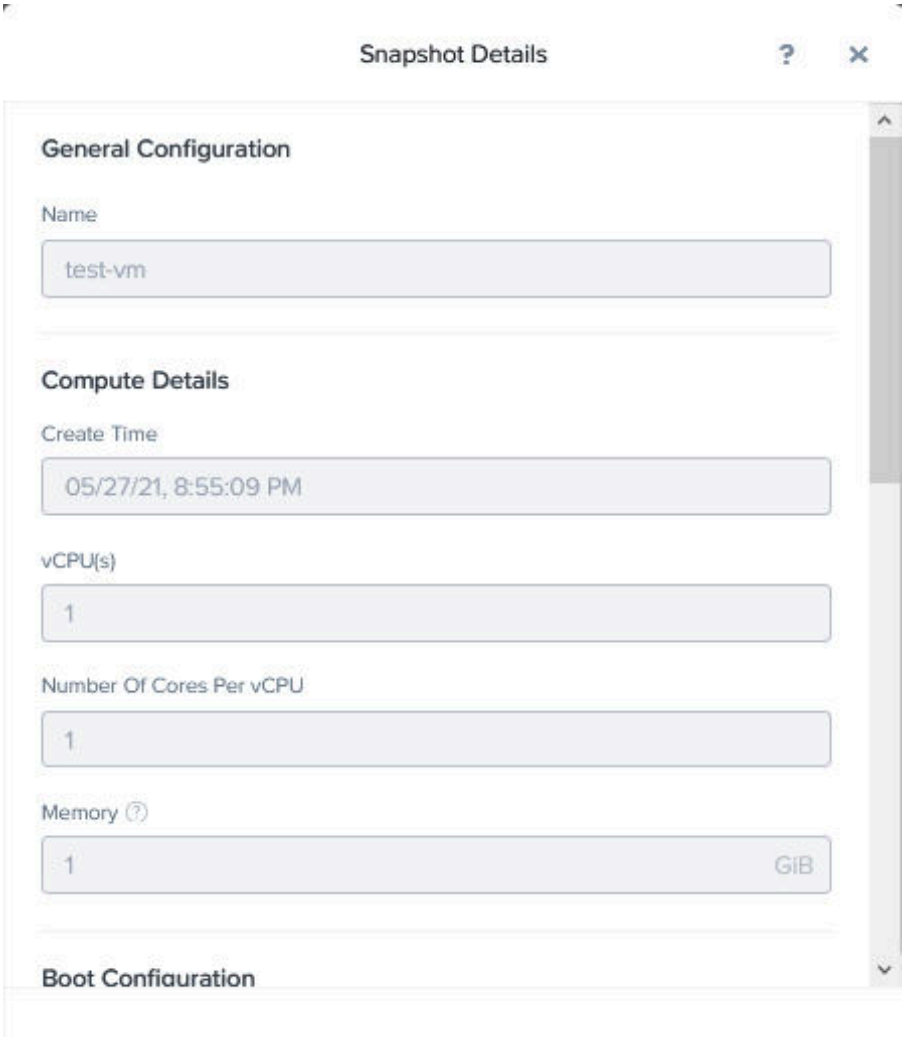

Figure 57: Snapshot Details Window

  - Click the **Clone** link to clone a VM from the snapshot.
  - Click the **Restore** link to restore the VM from the snapshot. This restores the VM back to the state of the selected snapshot.
  - Click the **Delete** link to delete the snapshot.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this VM (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this VM (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the VM. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 582). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 230).

The following table describes the available metrics. (Some of these metrics are not available on all hypervisors.)

Table 15: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the VM (0 - 100%).
CPU Ready Time	Displays the current, high, and low percentage of CPU wait time (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the VM (0 - 100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
IO Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
Usage	Displays separate graphs for current, snapshot, and shared storage usage (in GiBs) by the VM.
Working Set Size	Displays separate graphs for total, write, and read storage usage (in GiBs) for the VM working set size.
Network Packets Dropped	Displays separate graphs for the number of transmitted and received packets dropped.
Network Bytes	Displays separate graphs for the amount of transmitted and received bytes (in GiBs).

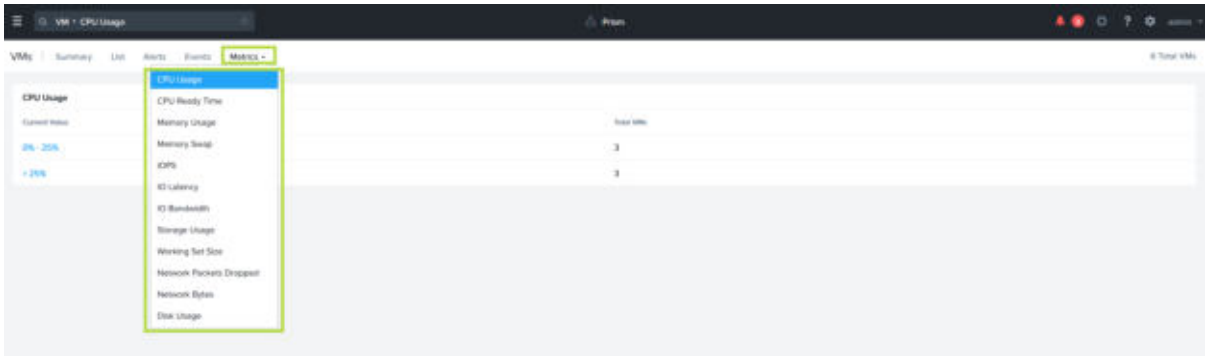

Figure 58: Metrics Tab: CPU Usage

NICs Tab

The **NICs** tab displays information in tabular form about the virtual NICs in the VM. Each line represent a virtual NIC, and the following table describes the fields.

Table 16: NIC Fields

Parameter	Description	Values
VLAN ID	Displays the VLAN name for this NIC.	(VLAN ID
MAC Address	Displays the virtual NIC MAC address.	(MAC address)
Network Connection State	Displays whether the NIC is connected to the network currently.	Connected, Disconnected
Requested IP Address	Displays the virtual NIC IP address	(IP address)
Action	Displays the available actions you can execute on this NIC.	Delete

Disks Tab

The **Disks** tab displays information in tabular form about the virtual disks in the VM. Each line represents a virtual disk, and includes the following fields.

- **Disk Address**: Displays the disk address (such as ide.0 or scsi.1).
- **Capacity**: Displays the disk capacity (in MiB or GiB).

Categories Tab

The **Categories** tab displays the categories and image placement policies associated with the VM. Each VM can have a one-to-many relationship with categories and the categories can have a many-to-one relationship with image placement policies.

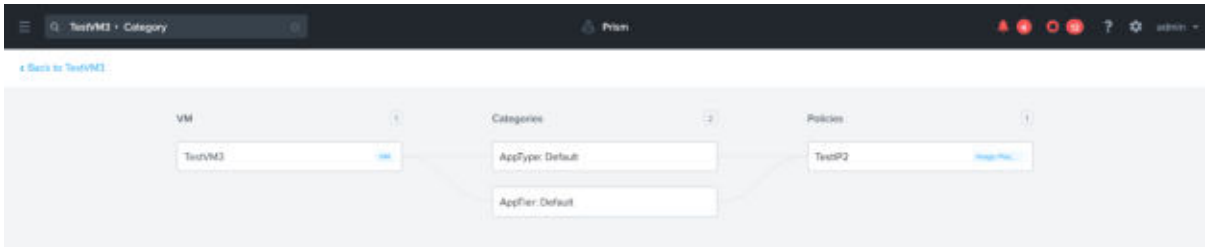

Figure 59: Categories Tab

For more information about categories management and image placement policies, see Category Management on page 413 and Image Placement Policies on page 480.

Storage Containers Summary View

To access the storage containers dashboard, select **Virtual Infrastructure** > **Storage Containers** from the entities menu (see Entities Menu on page 13). The storage containers dashboard summary view displays information about storage containers across the registered clusters and allows you to access detailed information about each storage container. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> Note: This section describes the information and options that appear in the storage containers dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in various ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities**: Displays a list of the storage containers with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that storage container. Click the **View all XX Storage Containers** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of storage container-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

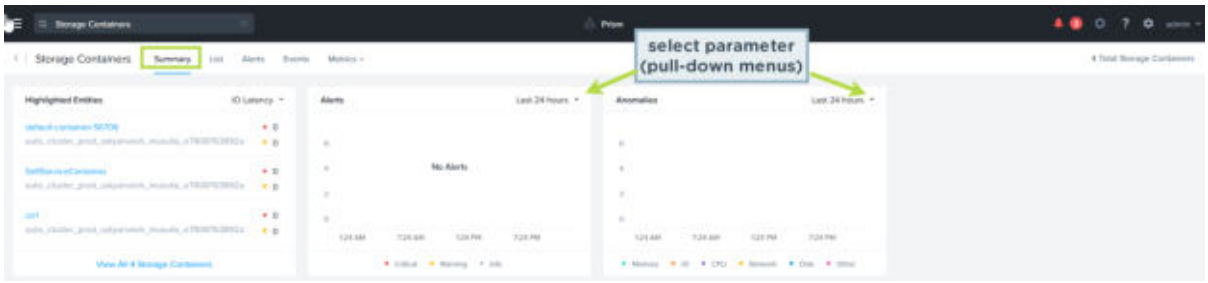

Figure 60: Storage Containers Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the storage containers across the registered clusters. The following table describes the fields that appear in the storage containers list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

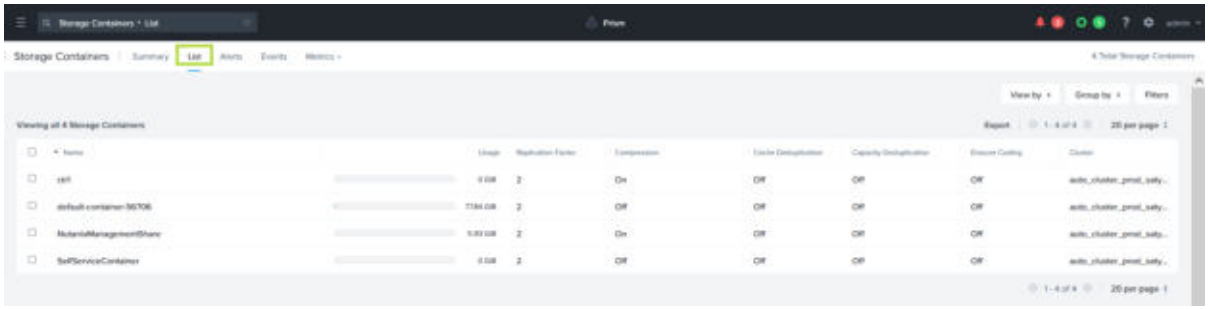

Figure 61: Storage Containers List Tab

Table 17: Storage Containers List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 103).	(name)
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the storage container. Hovering the cursor over the bar displays a window listing the used space, free space, and total space in the storage container.	xxx [GB|TB]
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]

Parameter	Description	Values
Capacity Deduplication	Displays whether capacity deduplication is enabled, that is deduplication compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Cluster	Displays the name of the cluster in which the storage container resides. Click the name to display the cluster details (see Cluster Details View on page 153).	(cluster name)
"Performance" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 103).	(name)
Free Space	Displays the amount of free space available in the storage container.	xxx [GB|TB]
Used Space	Displays the amount of used space in the storage container.	xxx [GB|TB]
Total Space	Displays the total amount of storage space in the storage container.	xxx [TB]
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
"Optimization" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 103).	(name)
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	*xx*:1

Parameter	Description	Values
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GB|TB]
Compression Delay	Displays the time delay to perform compression on the data.	xx [m]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GB|TB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	*xx*:1

You can filter the storage containers list based on various parameter values. The following table describes the filter options available when you open the storage containers view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 18: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of storage containers that satisfy the name condition/string.	(storage container name string)
Cluster	Filters on the cluster name. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of storage containers in the clusters that satisfy the name condition/string.	(cluster name string)
Cache Deduplication	Filters on the cache deduplication setting. Select one or more settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
Replication Factor	Filters on the replication factor. Check the box for the desired replication factor (1, 2, 3). Only existing RF values appear. For example, if all containers are at RF 2, just a single box for RF 2 would appear. The number of storage containers currently in each setting is displayed on the right of the line.	1, 2, 3

Parameter	Description	Values
Capacity Deduplication	Filters on the capacity deduplication setting. Select one or more settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
Erasure Coding	Filters on the erasure coding setting. Select one or both settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off
Free Space	Filters on the available storage space for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with available capacity in that range.	([xx] to [yy] GiB range)
Health	Filters on the storage container health state. Select one or more states to return a list of storage containers in those states. The number of storage containers currently in each state is displayed on the right of the line.	Critical, Warning, Good
Used Space	Filters on the used storage space for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with used storage in that range.	([xx] to [yy] GiB range)
Total Space	Filters on the total available capacity for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with maximum capacity in that range.	([xx] to [yy] GiB range)
IOPS	Filters on the current IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of storage containers with IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of storage containers with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of storage containers with average I/O latency in that range.	([xx] to [yy] range)

You can group the storage containers list in the following ways:

- The **Group** pull-down menu allows you to group the host entries by replication factor, compression, erasure coding, cluster, or health setting. (You can only choose one.)
- The **View by** pull-down menu allows you to group the host entries by the information parameters (fields), which vary depending on whether you selected the **General**, **Performance**, or **Optimization** focus. (You can only choose one parameter.) You can also use the **+ Add custom** entry to create a custom view.

There are no action options available from the storage containers view (no action buttons and no **Actions** menu options when a storage container is selected).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just storage container-related alerts across the registered clusters (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just storage container-related events across the registered clusters (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the storage containers. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

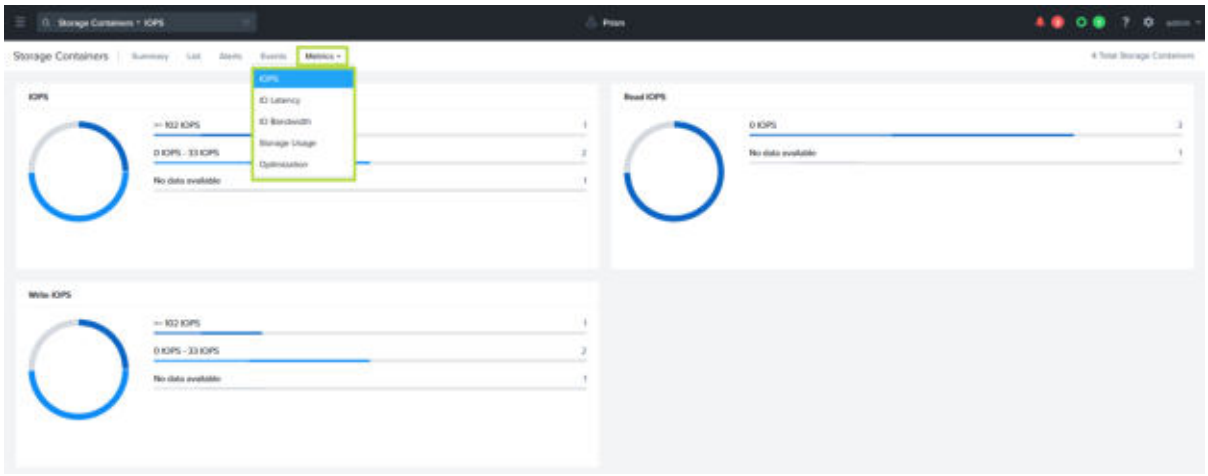

Figure 62: Storage Containers Metrics Tab

Table 19: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000). Note: The same format also applies to the other metrics in this table.
I/O Latency	Displays total, read, and write I/O latency rate graphs.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.
Storage Usage	Displays storage usage graph.
Optimization	Displays replication factor, data reduction ratio, and overall efficiency graphs.

Storage Container Details View

To access the details page for a storage container, go to the storage containers dashboard **List** tab (see Storage Containers Summary View on page 97) and click the storage container name. You can also access the details page by clicking the storage container name wherever that name appears, such as in a dashboard widget or search result.

The storage container name and the following set of tabs appear at the top: **Summary**, **Alerts**, **Events**, **Metrics**, and **Usage**. Click a tab to display that information. (Click the arrow at the upper left to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the storage container (see following table).
- A **Usage** widget that displays usage data for the storage container (see following table).
- An **Optimization & Performance** widget that displays optimization and performance data for the storage container (see following table).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last week** (default), **Last 1 Hours**, or **Last 24 hours** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

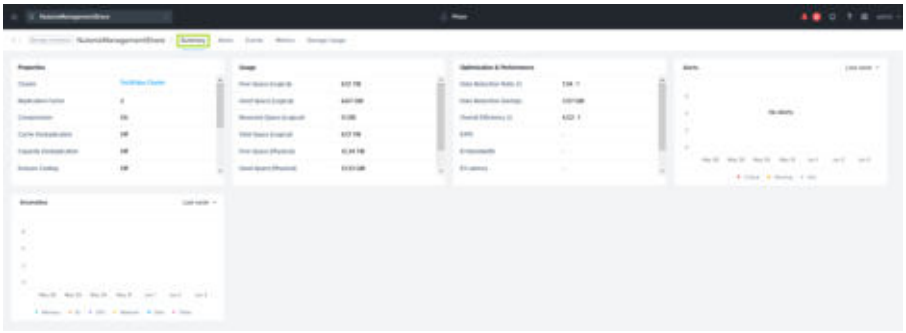

Figure 63: Storage Container Summary Tab

The following table describes the fields in the **Properties**, **Usage**, and **Optimization & Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 20: Storage Container Widget Fields

Parameter	Description	Values
"Properties" Widget		
Cluster	Displays the name of the cluster in which the storage container resides. Click the name to display the cluster details (see Cluster Details View on page 153).	(cluster name)
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether capacity deduplication is enabled, that is deduplication compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Filesystem Whitelists	Displays the IP addresses for file systems that are whitelisted.	[IP addresses]

Parameter	Description	Values
Compression Delay	Displays the delay (number of minutes) before data changes are compressed. A zero value indicates that compression is immediate (not delayed).	xx min
Thick Provision	Displays the reserved storage capacity for thick provisioned VMs.	xxx [GB|TB]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GB|TB]
"Usage" Widget		
Free Space (Logical)	Displays the amount of logical (effective) free space available in the storage container. Note: Logical space accounts for the storage container replication factor. RF1 means the logical and physical spaces are the same, RF2 means the logical space is half the physical space, and RF3 means the logical space is a third of the physical space.	xxx [GB|TB]
Used Space (Logical)	Displays the amount of logical used space in the storage container.	xxx [GB|TB]
Reserved Space (Logical)	Displays the amount of logical reserved space in the storage container.	xxx [GB|TB]
Total Space (Logical)	Displays the amount of logical total space in the storage container.	xxx [GB|TB]
Free Space (Physical)	Displays the amount of physical free space available in the storage container.	xxx [GB|TB]
Used Space (Physical)	Displays the amount of physical used space in the storage container.	xxx [GB|TB]
Total Space (Physical)	Displays the amount of physical total space in the storage container.	xxx [GB|TB]
"Optimization & Performance" Widget		
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	*xx*:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GB|TB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	*xx*:1

Parameter	Description	Values
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
Erasure Coding	Displays whether erasure coding is enabled for the storage container.	[On|Off]

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this storage container (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this storage container (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the storage container. Click the **Metrics** tab to display graphs for nine metrics: **Storage Controller IOPS**, **Storage Controller Read IOPS**, **Storage Controller Write IOPS Storage Controller Latency**, **Storage Controller Read Latency**, **Storage Controller Write Latency**, **Storage Controller IO Bandwidth**, **Storage Controller Read IO Bandwidth**, **Storage Controller Write IO Bandwidth**. Each graph is a rolling time interval performance or usage monitor. The baseline range (when established) appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 582). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week).
- Click the **Actions** pull-down menu (upper right of any graph) to monitor the metric performance in greater detail (see Analysis Dashboard on page 486):
  - Click **Add to Analysis** to add an active monitor for the metric on the Analysis dashboard.
  - Click **Select Analysis Session** to assign the metric to a target session.

- Click the filter button to display the filter panel; check (uncheck) the desired boxes to display (hide) those metrics.
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 230).

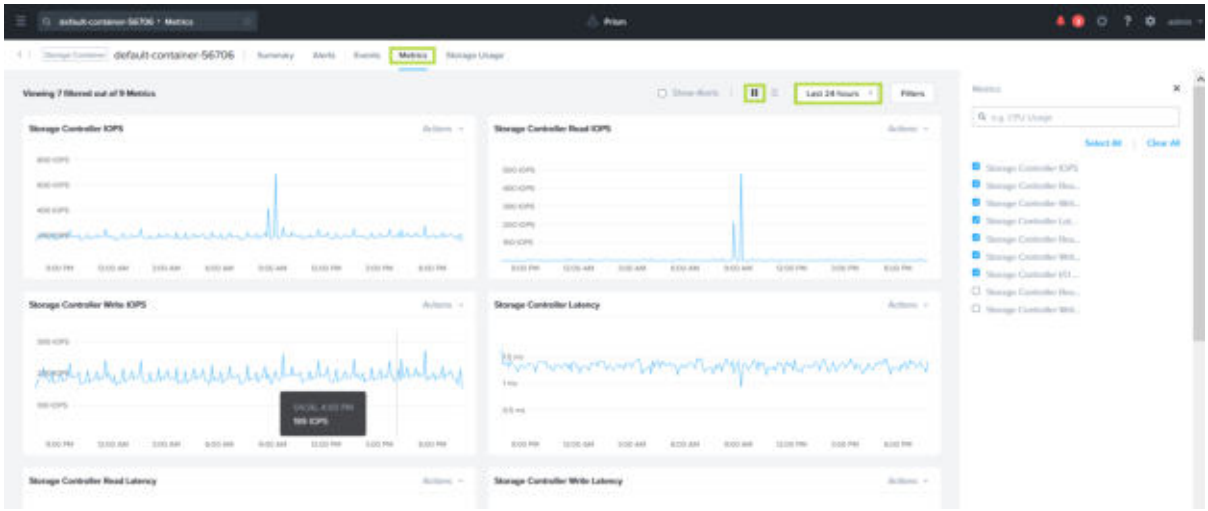

Figure 64: Storage Container Metrics Tab

Storage Usage Tab

The **Storage Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of storage container storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of container storage space used by each disk tier (SSD and DAS-SATA).

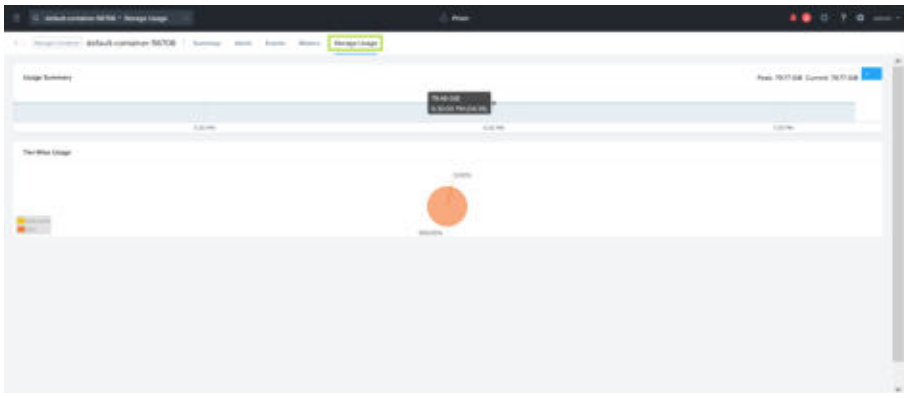

Figure 65: Storage Container Usage Tab

Volume Groups Summary View

To access the volume groups dashboard, select **Virtual Infrastructure** > **Volume Groups** from the entities menu (see Entities Menu on page 13). The volume groups dashboard summary view displays information about volume groups across the registered clusters and allows you to access detailed information about each volume group. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, and **Metrics**) with a display area below the selected tab.

Note:

- This section describes the information and options that appear in the volume groups dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in various ways.
- See Storage Management on page 379 for information about creating and managing volume groups.
- Some volume group details and actions are unavailable for clusters running a pre-6.0 AOS version.

Summary Tab

Clicking the **Summary** tab displays the following two widgets:

- **Highlighted Entities**: Displays a list of the volume groups with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that volume group. Click the **View all XX Volume Groups** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of storage container-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.

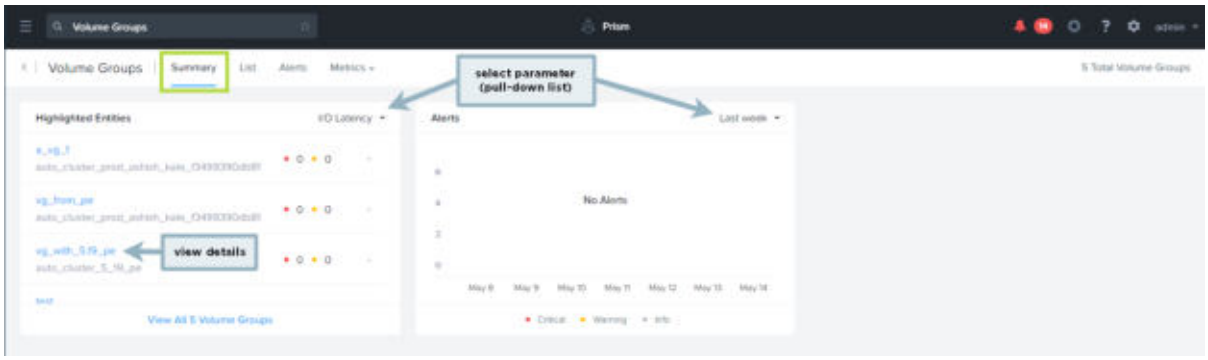

Figure 66: Volume Groups Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of volume groups across the registered clusters. The following table describes the fields that appear in the volume groups list. A dash (-) is displayed in a field when a value is not available or applicable.

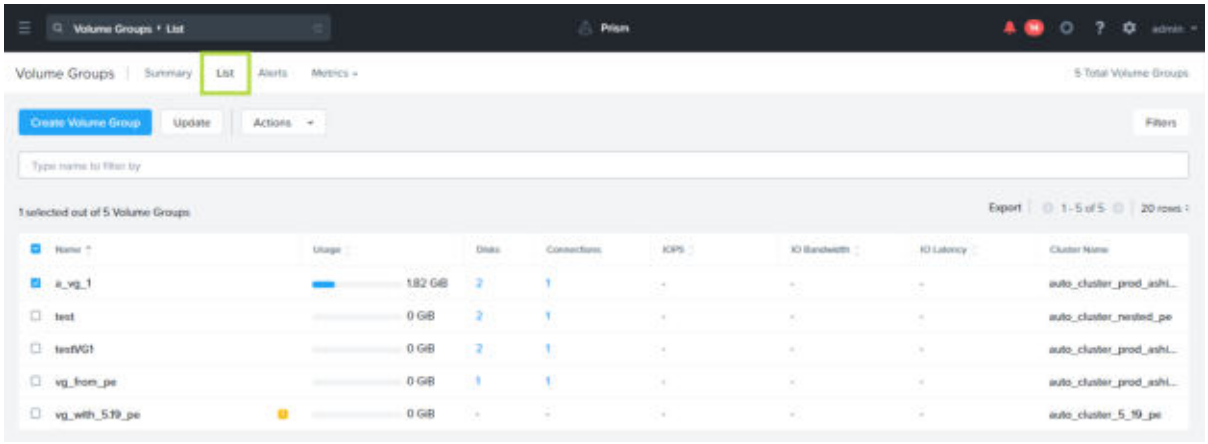

Figure 67: Volume Groups List Tab

Table 21: Volume Groups List Fields

Parameter	Description	Values
Name	Displays the name of the volume group. Click the name to display the details page for that volume group (see Volume Group Details View on page 112). An exclamation point icon appears if the cluster is running a pre-6.0 AOS version, which indicates the available information and actions are limited.	(name)
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the volume group. Hovering the cursor over the bar displays a window listing the used space, free space, and total space in the volume group.	xxx [GB|TB]
Disks	Displays the number of virtual disks in the volume group. Clicking the number displays the **Virtual Disks** tab of the details view (see Volume Group Details View on page 112).	(integer)
Connections	Displays the number of external client connections to the volume group. Clicking the number displays the **Connections** tab of the details view (see Volume Group Details View on page 112).	(integer)
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]

Parameter	Description	Values
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Cluster Name	Displays the name of the cluster in which the volume group resides. Click the name to display the cluster details (see Cluster Details View on page 153).	(cluster name)

You can filter the volume groups list based on various parameter values. The following table describes the filter options available when you open the volume groups view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 22: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the volume group name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of volume groups that satisfy the name condition/string.	(volume group name string)
Usage	Filters on the used storage space for a volume group. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of volume groups with used storage in that range.	([xx] to [yy] GiB range)
Cluster Name	Filters on the cluster name. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of volume groups in the clusters that satisfy the name condition/string.	(cluster name string)
IOPS	Filters on the current IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It returns a list of volume groups with IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It returns a list of volume groups with I/O bandwidth usage in that range.	([xx] to [yy] range)

Parameter	Description	Values
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of volume groups with average I/O latency in that range.	([xx] to [yy] range)

The volume groups view includes two action buttons and an actions menu:

- To create a volume group, click the **Create Volume Group** button (see Creating a Volume Group on page 380).
- To update or delete a volume group, click the **Update** or **Delete** button (see Modifying a Volume Group on page 384).
- To manage external client connections for a volume group, select the volume group and then select **Manage Connections** from the **Actions** menu (see Managing Volume Group Connections on page 385).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just volume group-related alerts across the registered clusters (see Alerts Summary View on page 210).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the volume groups. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

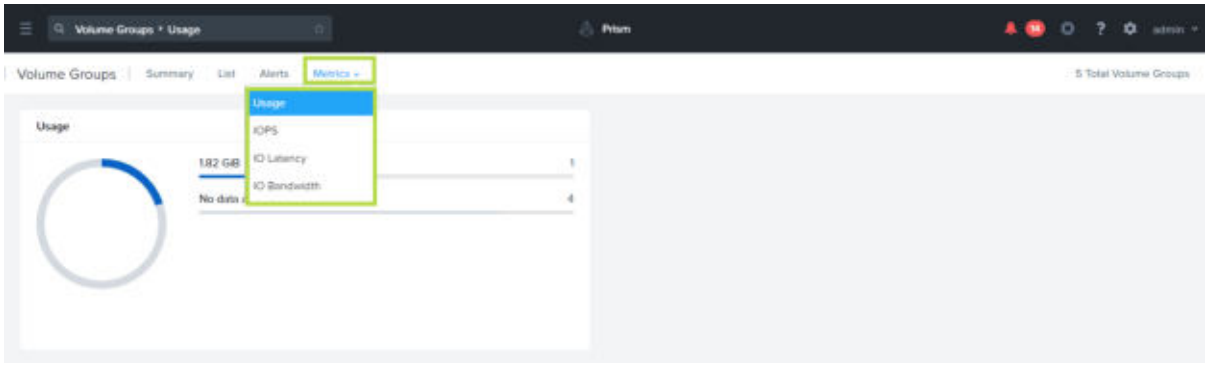

Figure 68: Volume Groups Metrics Tab

Table 23: Metrics Tab Fields

Metric	Description
Usage	Displays volume group usage graph.

Metric	Description
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000). Note: The same format also applies to the other metrics in this table.
I/O Latency	Displays total, read, and write I/O latency rate graphs.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.

Volume Group Details View

To access the details page for a volume group, go to the volume groups dashboard **List** tab (see Volume Groups Summary View on page 108) and click the volume group name. You can also access the details page by clicking the volume group name wherever that name appears, such as in a dashboard widget or search result.

The volume group name and the following set of tabs appear at the top: **Summary**, **Virtual Disks**, **Connections**, **Alerts**, and **Metrics**. Click a tab to display that information. (Click the arrow at the upper left to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the storage container (see following table).
- A **Usage and Performance** widget that displays usage data for the storage container (see following table).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last week** (default), **Last 1 Hours**, or **Last 24 hours** from the pull-down menu.

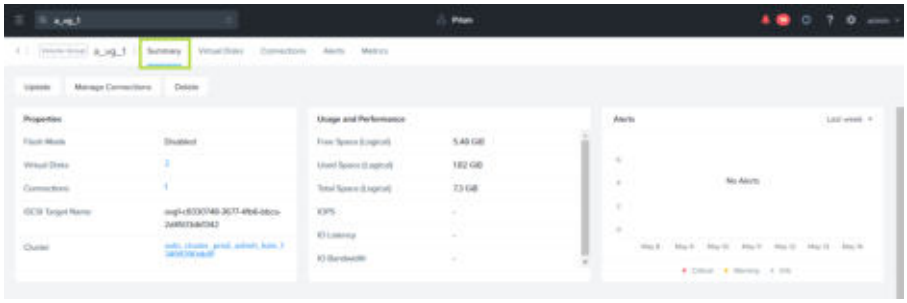

Figure 69: Volume Group Summary Tab

The following table describes the fields in the **Properties** and **Usage and Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 24: Volume Group Widget Fields

Parameter	Description	Values
"Properties" Widget		
Flash Mode	Displays whether Flash Mode is enabled for the volume group.	[Enabled, Disabled]
Virtual Disks	Displays the number of virtual disks in the volume group. Clicking the number displays the **Virtual Disks** tab.	(integer)
Connections	Displays the number of external client connections to the volume group. Clicking the number displays the **Connections** tab.	(integer)
iSCSI Target Name	Displays the name of the iSCSI target.	(ISCSI name)
Cluster	Displays the name of the cluster in which the volume group resides. Click the name to display the cluster details (see Cluster Details View on page 153).	(cluster name)
"Usage and Performance" Widget		
Free Space (Logical)	Displays the amount of logical (effective) free space available in the volume group. Note: Logical space accounts for the replication factor. RF1 means that the logical and physical spaces are the same, RF2 means that the logical space is half the physical space, and RF3 means that the logical space is a third of the physical space.	xxx [GB|TB]
Used Space (Logical)	Displays the amount of logical used space in the volume group.	xxx [GB|TB]
Total Space (Logical)	Displays the amount of logical total space in the volume group.	xxx [GB|TB]
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]

The **Summary** tab includes three action buttons:

- To update or delete the volume group, click the **Update** or **Delete** button (see Modifying a Volume Group on page 384).
- To manage external client connections for the volume group, click the **Manage Connections** button (see Managing Volume Group Connections on page 385).

Virtual Disks Tab

Clicking the **Virtual Disks** tab displays a list of virtual disks for the volume group. The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable.

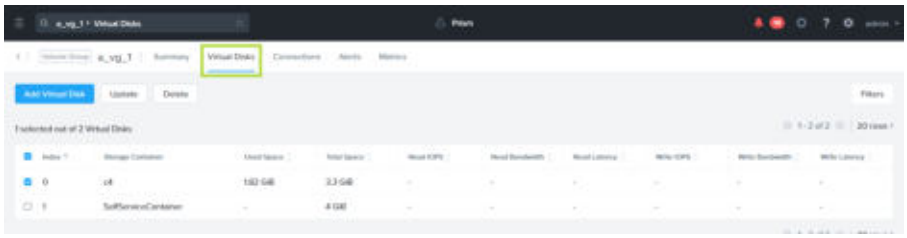

Figure 70: Volume Group Virtual Disks Tab

Table 25: Virtual Disks Fields

Parameter	Description	Values
Index	Displays the index number assigned to the virtual disk.	(integer)
Storage Container	Displays the name of the storage container in which the virtual disk is located.	(container name)
Used Space	Displays the amount of used space in the virtual disk.	xxx [GB|TB]
Total Space	Displays the total amount of storage space in the virtual disk.	xxx [GB|TB]
Read IOPS	Displays the current read I/O operations per second (IOPS) for the virtual disk.	[0 - unlimited]
Read Latency	Displays the average read I/O latency for the virtual disk.	xxx [ms]
Read Bandwidth	Displays read I/O bandwidth used per second for the virtual disk.	xxx [MBps|KBps]
Write IOPS	Displays the current write IOPS for the virtual disk.	[0 - unlimited]
Write Latency	Displays the average write I/O latency for the virtual disk.	xxx [ms]
Write Bandwidth	Displays write I/O bandwidth used per second for the virtual disk.	xxx [MBps|KBps]

You can filter the virtual disks list based on various parameter values. The following table describes the filter options available when you open the virtual disks view **Filter** pane. To apply

a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 26: Filter Pane Fields

Parameter	Description	Values
storage container	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of volume groups that satisfy the name condition/string.	(storage container name string)
[Total|Used] Space	Filters on the total or used storage space for a virtual disk. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of virtual disks with total or used storage in that range.	([xx] to [yy] GiB range)
[Read|Write] IOPS	Filters on the current read or write IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It returns a list of virtual disks with IOPS in that range.	([xx] to [yy] range)
[Read|Write] Bandwidth	Filters on the read or write I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It returns a list of virtual disks with I/O bandwidth usage in that range.	([xx] to [yy] range)
[Read|Write] Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of virtual disks with average I/O latency in that range.	([xx] to [yy] range)

The **Virtual Disks** tab includes three action buttons:

- To add a virtual disk to the volume group, click the **Add Virtual Disk** button (see Managing Volume Group Virtual Disks on page 389).
- To update or delete a virtual disk, select the target disk and then click the **Update** or **Delete** button (see Managing Volume Group Virtual Disks on page 389).

Connections Tab

Clicking the **Connections** tab displays a list of external client connections to the volume group. The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable.

Figure 71: Volume Group Connections Tab

Table 27: Connections Fields

Parameter	Description	Values
Connection	Displays the IP address or IQN of the external client.	(IP or IQN address)
Type	Displays the type of connection.	external client

The **Connections** tab includes three action buttons:

- To add an external client connection, click the **Add Connection** button (see Managing Volume Group Connections on page 385).
- To update or remove a connection, click the **Update** or **Remove** button (see Managing Volume Group Connections on page 385).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this volume group (see Alerts Summary View on page 210).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the volume group. Click the **Metrics** tab to display graphs for 10 metrics: **Usage**, **Controller IOPS**, **Controller Read IOPS**, **Controller Write IOPS Controller Average I/O Latency**, **Controller Average Read I/O Latency**, **Controller Average Write I/O Latency**, **Controller I/O Bandwidth**, **Controller Read I/O Bandwidth**, **Controller Write I/O Bandwidth**. Each graph is a rolling time interval performance or usage monitor. The baseline range (when established) appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 582). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- Click the **Actions** pull-down menu (upper right of any graph) to monitor the metric performance in greater detail (see Analysis Dashboard on page 486):
  - Click **Add to Analysis** to add an active monitor for the metric on the Analysis dashboard.
  - Click **Select Analysis Session** to assign the metric to a target session.

- Click the filter button to display the filter panel; check (uncheck) the desired boxes to display (hide) those metrics.
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 230).

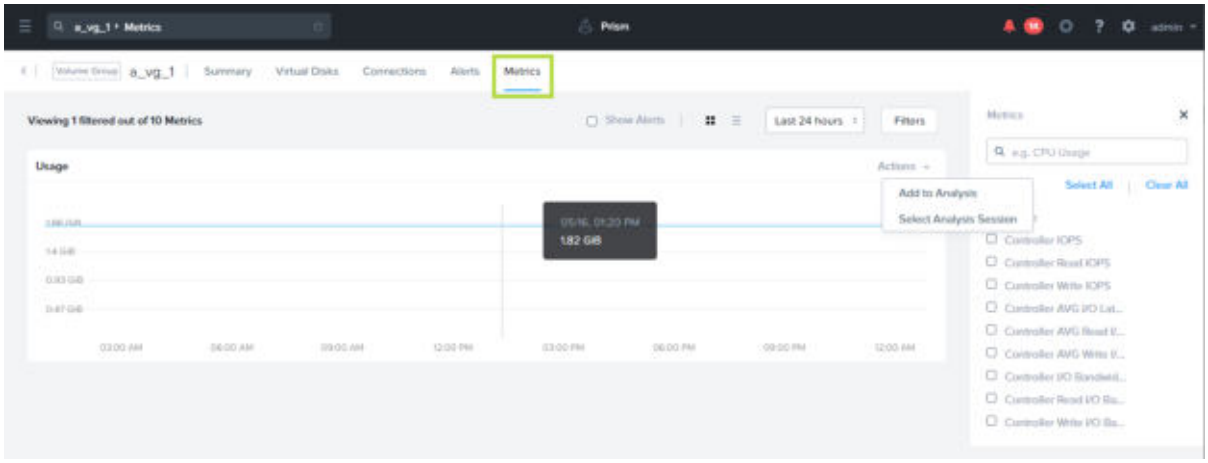

Figure 72: Volume Group Metrics Tab

vCenter Datastores Summary View

To access the vCenter datastores dashboard, select **Virtual Infrastructure** > **vCenter Datastores** from the entities menu (see Entities Menu on page 13). The vCenter datastores dashboard summary view displays information about datastores across the registered vCenter instances and allows you to access detailed information about each datastore. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below for the selected tab.

Note:

- This section describes the information and options that appear in the vCenter datastores dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- The vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 540 for more information.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities**: Displays a list of the datastores with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The only option is **Vsphere Datastore Numberreadaveraged Average**. Click the **View All XX vCenter Datastores** link at the bottom to display the **List** tab (following section).
- **Alerts**: Displays a list of datastore-related alerts that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.

- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

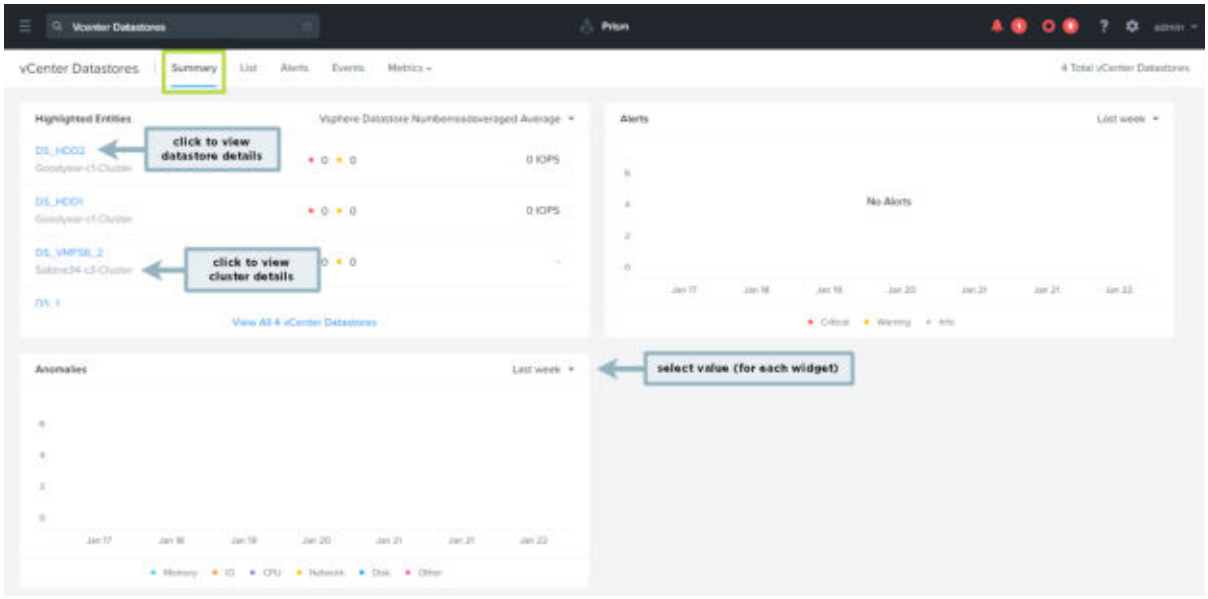

Figure 73: Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the registered vCenter datastores. The following table describes the fields that appear in the datastores list. A dash (-) is displayed in a field when a value is not available or applicable.

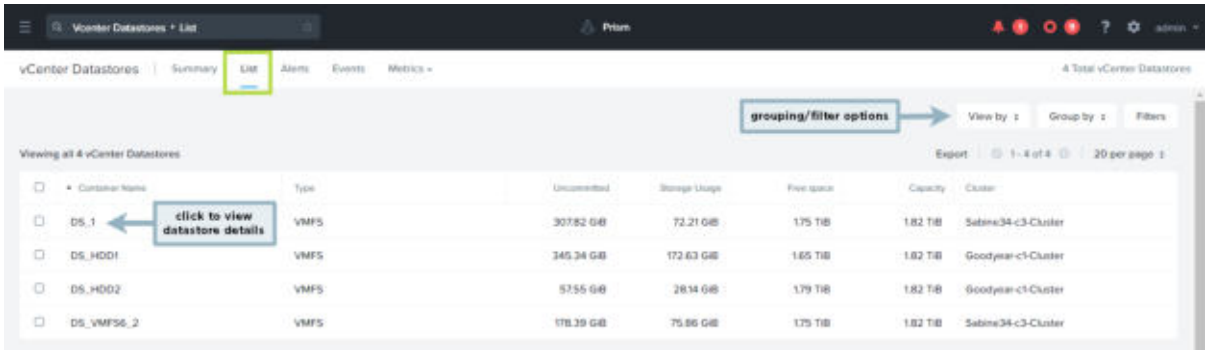

Figure 74: List Tab

Table 28: Datastores Fields

Parameter	Description	Values
Container Name	Displays the name of the datastore container. Click the name to display the details for that container (see vCenter Datastore Details View on page 121).	(container name)
Type	Displays the file system type.	VMFS, NFS
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB|TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB|TiB]
Free Space	Displays the amount of free storage space.	xx [GiB|TiB]
Capacity	Displays the total amount of storage capacity.	xx [GiB|TiB]
Cluster	Displays the name of the cluster in which the datastore resides.	(cluster name)

You can filter the datastores list based on a variety of parameter values. The following table describes the filter options available when you open the datastores view **Filters** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 29: Filter Pane Fields

Parameter	Description	Values
Container Name	Filters on the container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of containers that satisfy the name condition/string.	(name string)
Cluster	Filters on the cluster name. Enter a string in the field to display a list of clusters that satisfy the name string.	(cluster name)
Storage Usage	Filters on the amount of used storage space. Check the box(es) for the desired range or enter a range in the "**from <low> to <high> TiB**" field. The number of containers in each range is displayed on the right of the line.	([xx] to [yy] range)
Capacity	Filters on the amount of total storage capacity.	([xx] to [yy] range)
Free Space	Filters on the amount of free storage space.	([xx] to [yy] range)
Uncommitted	Filters on the amount of uncommitted storage space.	([xx] to [yy] range)
Type	Filters on file system type.	VMFS, NFS

To download the table contents in CSV format, click the **Export** link.

You can group the datastore containers list in the following ways:

- The **View by** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.
- The **Group by** pull-down menu allows you to group the datastore entries by cluster or type. (You can only choose one.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just datastore-related alerts across the registered clusters (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just datastore-related events across the registered clusters (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the datastores. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant performance information. The following table describes the available metrics.

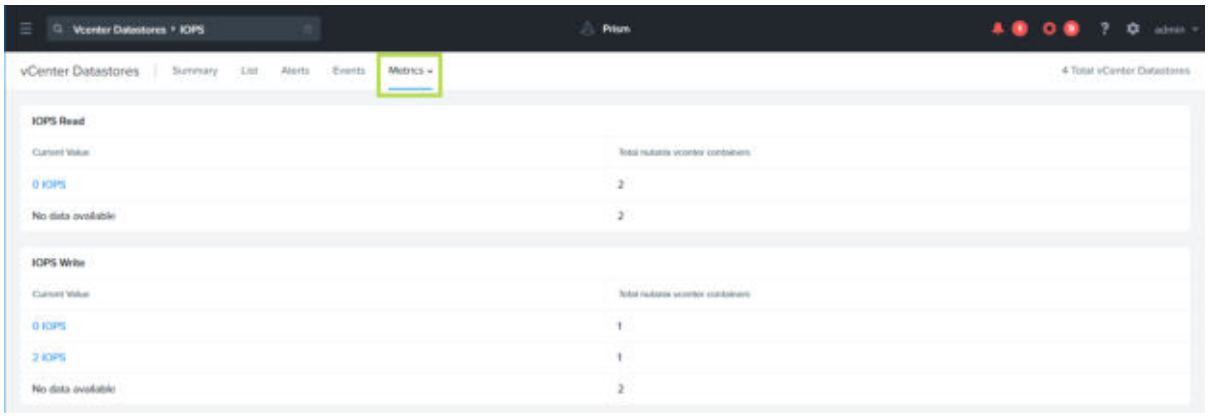

Figure 75: Metrics Tab (IOPS)

Table 30: Metrics Tab Fields

Metric	Description
IOPS	Displays the number of read and write IOPS for the listed number of containers. (Clicking an IOPS value redisplays the Summary tab filtered for those containers.)
Disk Usage	Displays the disk usage amount for the listed number of containers. (Clicking the usage value redisplays the Summary tab filtered for those containers.)

Metric	Description
Disk Capacity	Displays the disk capacity for the listed number of containers. (Clicking the capacity value redisplays the Summary tab filtered for those containers.)
Disk Provisioned	Displays the provisioned disk space for the listed number of containers. (Clicking the provisioned value redisplays the Summary tab filtered for those containers.)

vCenter Datastore Details View

To access the details page for a vCenter datastore, go to the vCenter datastores **Summary** tab (see vCenter Datastores Summary View on page 117) and click the datastore name. You can also access the details page by clicking the datastore name wherever that name appears, such as in a dashboard widget or search result.

The datastore name and the following set of tabs appear at the top: Summary, Alerts, Events, and Metrics. Click a tab to display that information below the tabs. (Click the **Back to vCenter Datastores** link to return to the datastores summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the VM (see following table).
- A **Metrics** widget that displays average metrics for the past 15 minutes (IOPS reads and writes, and disk usage, disk capacity, and disk provisioned amounts).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
-

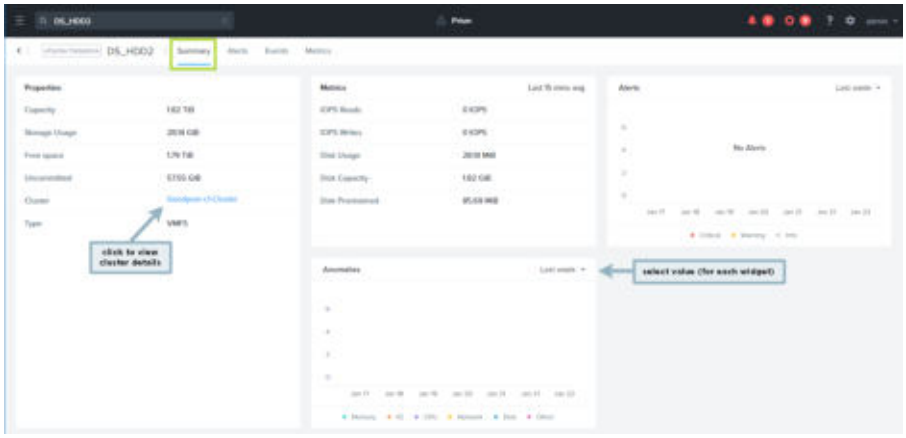

Figure 76: Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 31: Datastore Properties Fields

Parameter	Description	Values
Capacity	Displays the total amount of storage capacity.	xx [GiB|TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB|TiB]
Free Space	Displays the amount of free storage space.	xx [GiB|TiB]
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB|TiB]
Cluster	Displays the name of the cluster in which the datastore resides. Click on the cluster name to displays details about that cluster.	(cluster name)
Type	Displays the file system type.	VMFS, NFS

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this vCenter datastore (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this vCenter datastore (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab displays a set of graphs for various metrics. Each graph is a rolling time interval performance or usage monitor for that metric.

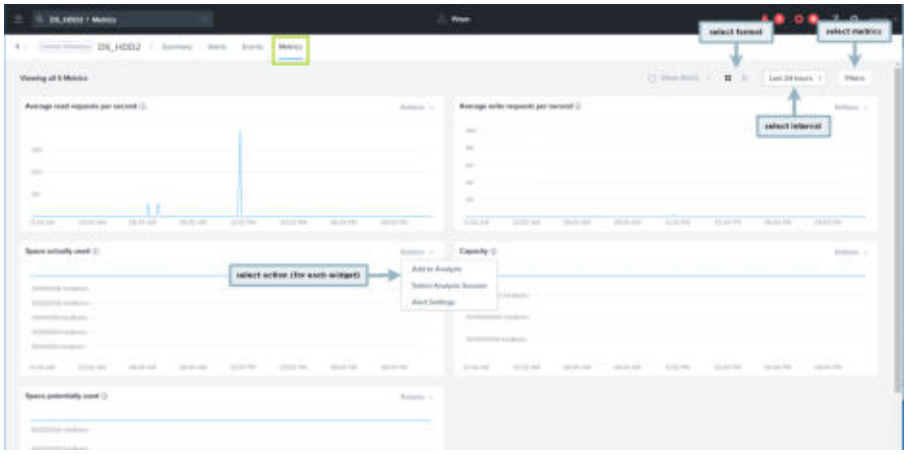

Figure 77: Metrics Tab

- Each graph includes a pull-down actions menu with some or all of the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard on page 486).
  - Select **Select Analysis Session**, which opens the **Select A Session** window, to add the metric to an existing or new analysis session. Choose an option from the pull-down list and then click the Add to Session button.

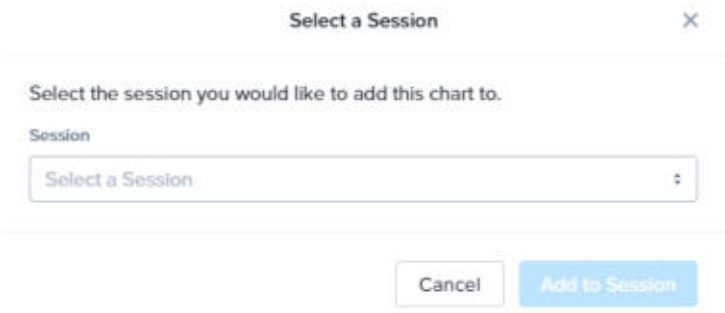

Figure 78: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 230).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired metrics.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to vCenter Datastores** link to return to the vCenter datastores dashboard.

The following table describes the metric graphs.

Table 32: Metric Performance Graphs

Metric	Description
Average read requests per second	Average number of read commands issued per second to the datastore during the collection interval.
Average write requests per second	Average number of write commands issued per second to the datastore during the collection interval.
Space actually used	Amount of space actually used by the virtual machine or the datastore. May be less than the amount provisioned at any given time, depending on whether the virtual machine is powered-off, whether snapshots have been created or not, and other such factors. Available from datastore and virtual machine target entities.
Capacity	Configured size of the datastore. Available from a datastore entity only.

Metric	Description
Space potentially used.	Amount of storage set aside for use by a datastore or a virtual machine. Files on the datastore and the virtual machine can expand to this size, but not beyond it. Available from datastore and virtual machine target entities.

Catalog Items View

To access the catalog items dashboard, select **Virtual Infrastructure** > **Catalog Items** from the entities menu (see Entities Menu on page 13). The catalog items dashboard allows you to view information about the images and VMs in the catalog.

> Note: This section describes the information and options that appear in the catalog items dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See Catalog Management on page 390 for information about the catalog service and how to add (or delete) items.

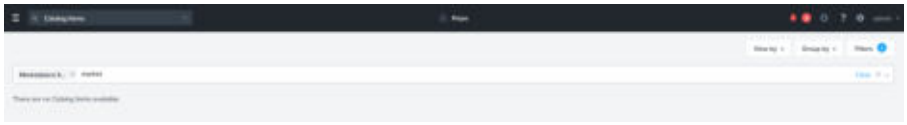

Figure 79: Catalog Items Dashboard

The following table describes the fields that appear in the catalog items list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 33: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the item name.	(item name)
Type	Displays the type of item (image or VM).	Image, Vm
Description	Indicates which user added this item to the catalog.	(text string)

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 34: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of catalog items that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It will return a list of catalog items whose description field satisfy the condition/string.	(description string)
Type	Filters on the item type. Check the box(es) for the desired item types.	Vm, Image

You can group the catalog items list in the following ways:

- The **Group** pull-down menu allows you to group by type (VM and image).
- The **View by** pull-down menu allows you to view by name, description, message, percent, reason, state, and type.

The **Actions** menu appears when one or more catalog items are selected. It allows you to delete the selected catalog items (see Deleting a Catalog Item on page 393).

Images Summary View

To access the images dashboard, select **Virtual Infrastructure** > **Images** from the entities menu (see Entities Menu on page 13). The images dashboard allows you to view information about images available from this Prism Central instance.

Note: This section describes the information and options that appear in the images dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See Image Management on page 394 for information about adding and managing images through Prism Central.

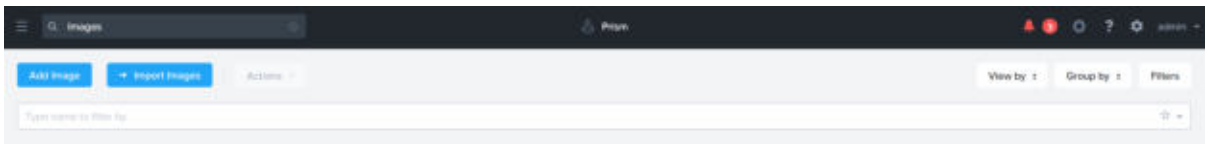

Figure 80: Images View

The following table describes the fields that appear in the images list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 35: Image List Fields

Parameter	Description	Values
Name	Displays the image name. Clicking the name displays the details page for that image (see Image Details View on page 127).	(name)
Description	Indicates which user uploaded this image.	(text string)
Type	Displays the image type.	ISO, Disk
Size	Displays the image size.	xxx [MB|GB]
Creator	Displays who created the image.	(name)

You can filter the images list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 36: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of storage containers that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It will return a list of images whose description field satisfy the condition/string.	(description string)
Type	Filters on the image type. Check the box(es) for the desired image types.	Disk, ISO

You can group the images list in the following ways:

- The **Group** pull-down menu allows you to group by type (disk and ISO).
- The **View by** pull-down menu allows you to view the entries by categories, creator, description, size, and type.

The images dashboard includes the following action buttons:

- **Add Image** Click to add an image (see Creating an Image on page 395).
- **Import Images** Click to import images from registered clusters to Prism Central (see Importing Images to Prism Central on page 412.

The **Actions** menu appears when one or more images are selected. It includes delete, update, and add image to catalog options (see Modifying an Image on page 411). The available

actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Image Details View

You can view detailed information about an image by going to the details page for that image.

To access the details page for an image, go to the images dashboard (see Images Summary View on page 125) and click the image name. An image details page includes the following:

- Image name (upper left). You can switch from one image to another by selecting a different image name from the pull-down list.
- Image information (middle left). The information includes the image description, type, size, and creator name (see Creating an Image on page 395).
- Tab options (lower left). Select the **Location** (default) or **Policies** tab to display that tab content to the right.
- Action buttons (upper right). When the window size is not large enough, some actions appear under a **More** menu button.
  - To delete the image, lick the **Delete** button
  - To modify the name, description, or type, click the **Update Image** button (see Modifying an Image on page 411).
  - To add the image to the catalog, click the **Add Image to Catalog** button (see Adding a Catalog Item on page 391).
  - To assign the image to a category, click the **Manage Categories** button (see Assigning a Category on page 415).
  - To open a help page in a separate tab or window, click the question mark icon.
  - To close the details page, click the X icon.

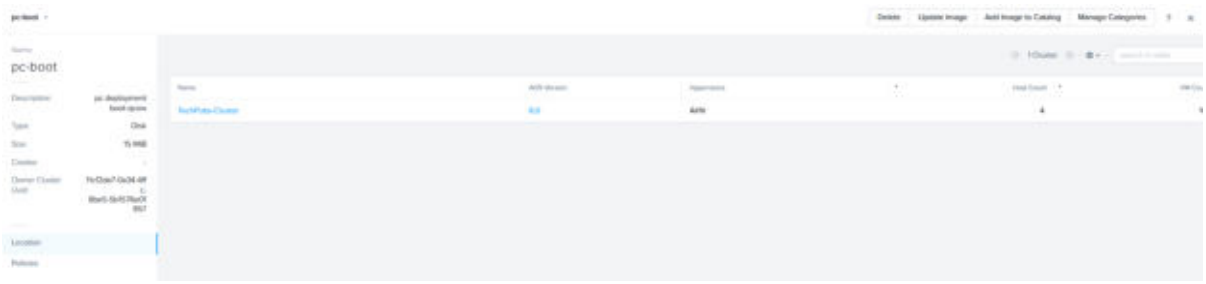

Figure 81: Image Details View (Location Tab)

Location Tab

The following table describes the information displayed in the **Location** tab, which appears by default when you first open the page.

Table 37: Location Fields

Parameter	Description	Values
Name	Displays the name of the cluster in which the image is stored. Each row represents another cluster that contains the image.	(cluster name)
AOS Version	Displays the AOS version running in the cluster.	(AOS version number)
Hypervisor	Displays the name of the hypervisor running in the cluster.	AHV, ESXi, Hyper-V
Host Count	Displays the number of hosts in the cluster.	(integer)
VM Count	Displays the number of VMs in the cluster	(integer)

Policies Tab

The **Policies** tab lists the policies that apply to the image (see Image Placement Policies on page 480). The following table describes the information displayed in the **Policies** tab.

Table 38: Policies Fields

Parameter	Description	Values
Policy Name	Displays the policy name.	(cluster name)
Status	Indicates whether the image is compliant per the policy.	compliant, noncompliant
Details	Lists the clusters in which the policy is enforced.	(cluster names)

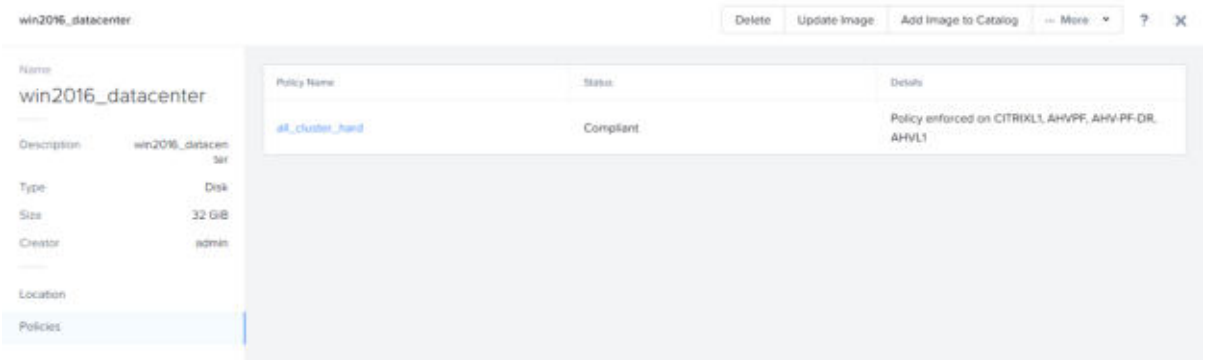

Figure 82: Image Details View (Policies Tab)

Categories Summary View

To access the categories dashboard, select **Virtual Infrastructure** > **Categories** from the entities menu (see Entities Menu on page 13). The categories dashboard allows you to view summary information about existing categories and access detailed information about each category.

Note: This section describes the information and options that appear in the categories dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See Category Management on page 413 for information on how to create, modify, and apply categories.

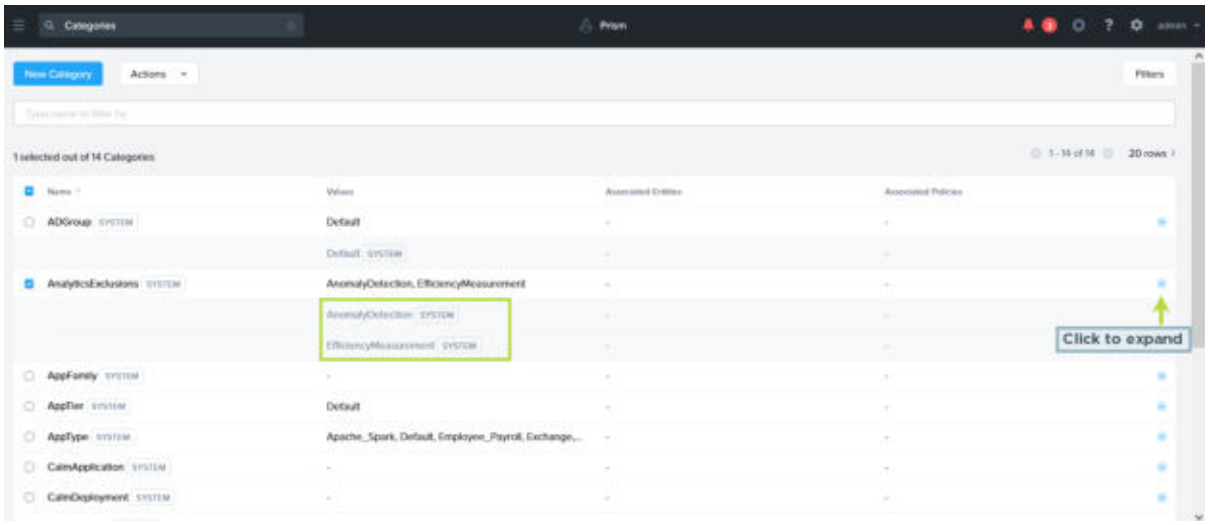

Figure 83: Categories Dashboard

The following table describes the fields that appear in the categories list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 39: Categories List Fields

Parameter	Description	Values
Name	Displays the name of the category.	(category name)
Value	Displays the values defined for the category. Click **Show more** (right of line) to see a line for each value. (Click**Show fewer** to collapse the list.) You may see an icon with a "showing X of Y" message to indicate there are additional values; click the icon to display the full list.	(value names)
Assigned Entities	Lists the number of entities assigned to this category.	(number)
Assigned Policies	Lists the number of policies assigned to this category.	(number)

You can filter the category list based on several parameter values. The following table describes the options available when you click the **Filter** button, which displays the Categories view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 40: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the category name. Select a condition from the pull-down list and enter a string in the field. It will return a list of categories that satisfy the name condition/ string. Note: In this and the Value field, the condition menu options are **Contains**, **Does not contain**, **Starts with**, **Ends with**, and **Equal to**.	(name string)
Entities	Filters on the entity type. Check the box for one or more entity types.	VM, Host,
Policies	Filters on the policy type. Check the box for one or more entity types.	Security Policy, Affinity Policy

There is a **New Category** action button to create a new category (see Creating a Category on page 414). The **Actions** menu appears when one or more categories are selected and includes the following options:

- This action can be applied to multiple categories: **Delete**
- This action can be applied to only one category at a time: **Update**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the selected category. See Modifying a Category on page 415 for instructions on how to perform these actions.

Category Details View

You can view detailed policy, value, and entity associations for a category by going to the details page for that category. To access the details page for a category, go to the categories dashboard (see Categories Summary View on page 128) and click the category name.

A category details page includes the following:

- Category name (upper left). You can switch from one category to another by selecting a different category name from the pull-down list in the upper left of the screen.
- Action buttons (upper right).
  - Click the **Update** button to update the category definition (see Modifying a Category on page 415); click the **Delete** button to delete the category. The button is grayed out if that action is not allowed. For example, you cannot delete system categories.
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.
- **Policies** (left), **Values** (middle), and **Entities** (right) columns that list the values defined for the category plus the policies and entities associated with that category. Placing the cursor over a policy, value, or entity displays lines that graphically indicate the association among the three parameters. For example, in the following figure placing the cursor over the

Forensics value shows that it is associated with the **Quarantine-Forensics** security policy, but it is not associated with any entities currently.

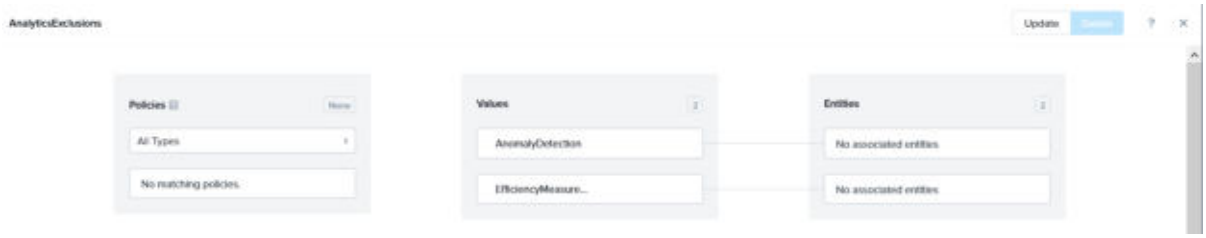

Figure 84: Category Details View

OVAs View

To access the OVAs dashboard, select **Virtual Infrastructure** > **OVAs** from the entities menu (see Entities Menu on page 13). The OVAs dashboard allows you to view information about the OVAs available in the

Note: This section describes the information and options that appear in the OVAs dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See OVA Management on page 433 for information about the OVA operations.

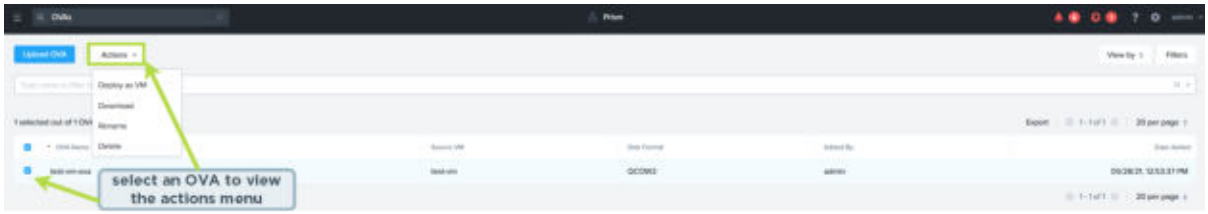

Figure 85: OVAs Dashboard

The following table describes the fields that appear in the catalog items list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 41: OVA List Fields

Parameter	Description	Values
OVA Name	Displays the name of the OVA file.	(file name)
Source VM	Displays the name of source VM that was exported as an OVA. The alert indicators for the task status are displayed beside the source VM for which the task is run. See the Alert indicator table for more information.	(file name)

Parameter	Description	Values
Disk Format	Displays the format of the disks in the OVA.	[QCOW2|VMDK]
Added by	Indicates which user added this OVA.	(text string)
Date added	Displays the date and time when the OVA was added.	(date in mm/dd/yyyy format, time in hr:min:sec [AM|PM] format) Example: 07/30/20, 3:00:20 PM

Table 42: OVA Alerts

OVA Alerts	Status Message sample	Indicates...
	Incorrect file uploaded. Multiple OVA files present in OVA.	That the upload process failed during validation. You must fix the issue in the OVA file or upload the correct, compatible OVA file.

Figure 86:

	Upload for OVA file has not started. Please start OVA file upload using Resume Upload option. Uploaded OVA file has not been verified. Please start verification using Resume Upload option.	That there is an interruption in one of the three sub-tasks in the upload process. You can resume the upload. See the *Resume Upload* action procedure.

A sample of the error hover message is as follows:

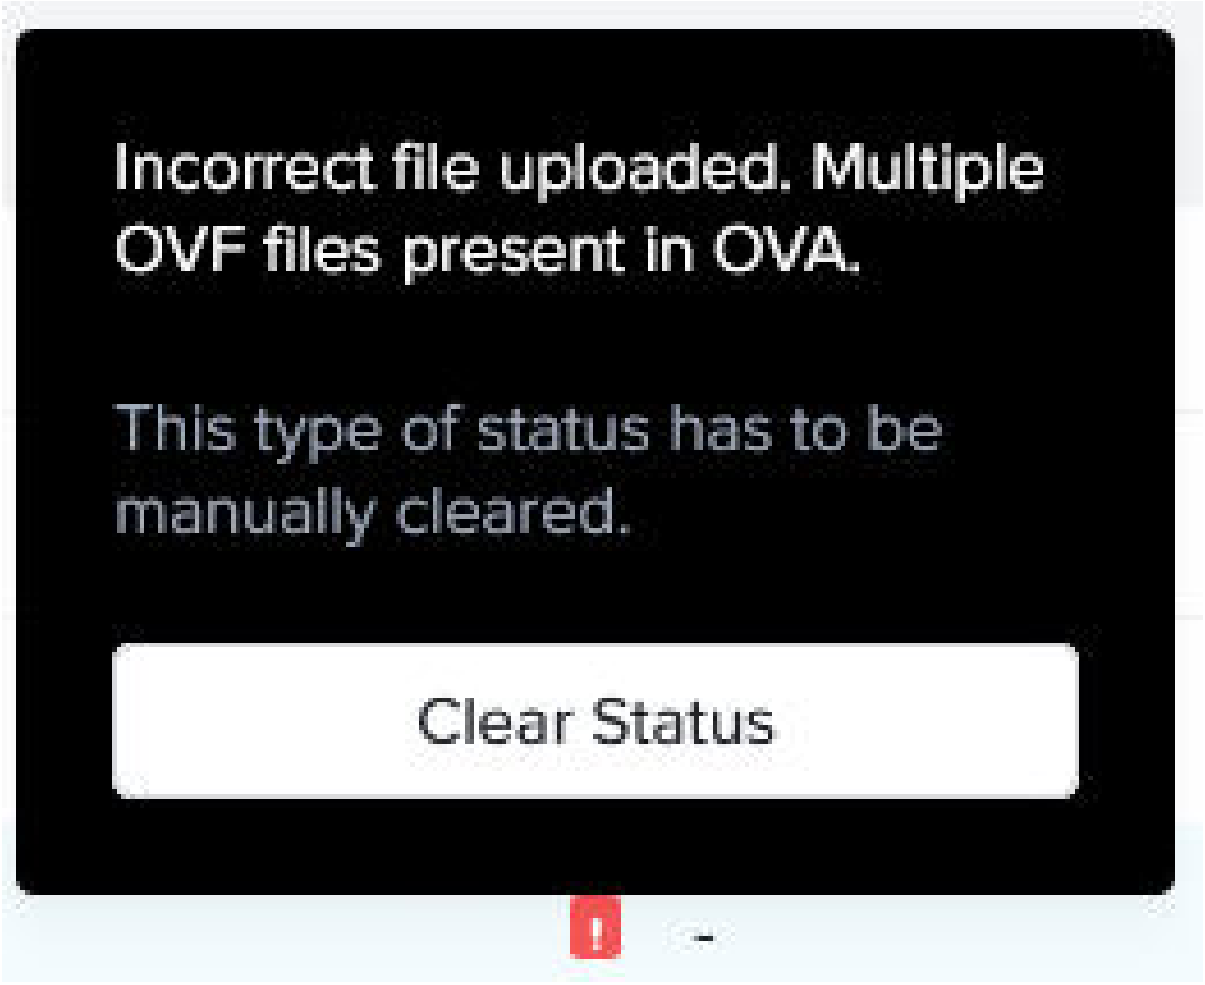

Figure 89: Sample error hover message

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

For each filter, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.

Table 43: Filter Pane Fields

Parameter	Description	Values
OVA Name	Filters on the OVA file name. It returns a list of OVAs that satisfy the name condition/string.	(OVA file name string)
Source VM	Filters on the Source VM name. It returns a list of OVAs for which the Source VM field satisfies the condition/string.	(Source VM name string)

Parameter	Description	Values
Disk Format	Filters on the disk format. Check the box(es) for the OVAs of the required disk format.	VMDK, QCOW2

The **Actions** menu appears when a OVA is selected. It allows you to:

- Deploy an OVA file as a VM.
- Download an OVA file to your local machine.
- Rename an OVA file.
- Delete an OVA file.
- Resume upload when an upload is interrupted and shows an error.

> Note: The resume upload action is not available in case the concatenate process of the upload is interrupted. In such a case, you need to run the Concatenate API from the REST API Explorer to resume the upload.

Network and Security Entities

You can access dashboards for the following network-related entities from the **Network & Security** category of the entities menu (see Entities Menu on page 13):

- Subnets (see Subnets View on page 135)
- Virtual private clouds
- Floating IPs
- VPNs
- See Network and Security Management on page 421 for information on how to configure network connections.

Virtual private clouds, floating IPs, and VPNs are advanced Flow Networking features. These features support flexible app-driven networking that focuses on VMs and applications instead of virtual LANs and network addresses. Flow Networking powers network virtualization to offer a seamless network experience with enhanced security. It is disabled by default. It is a software-defined network virtualization solution providing overlay capabilities for the on-premises AHV clusters.

> Important: The Flow Networking features are currently in technical preview. You need to enter a secret passphrase when you are enabling it. Contact your Nutanix account team to enable Flow Networking. Refer to the support FAQs for more details on support during technical preview.

Subnets View

To access the subnets dashboard, select **Subnets** from the entities menu (see Entities Menu on page 13). The subnets dashboard allows you to view information about the subnets configured for the registered clusters.

> Note: This section describes the information and options that appear in the Network and Security dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

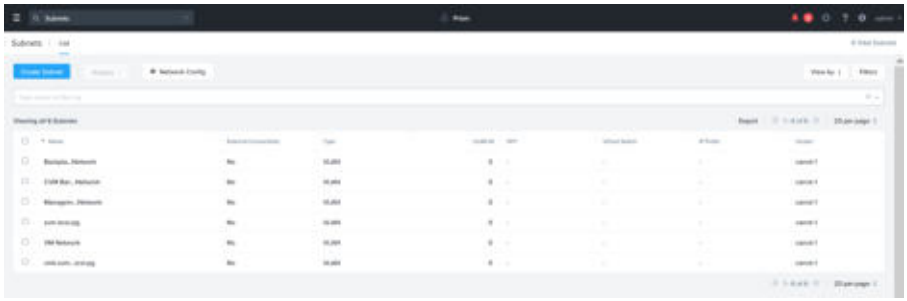

Figure 90: Subnets Dashboard

The following table describes the fields that appear in the subnets list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 44: Subnets Fields

Parameter	Description	Values
Name	Displays the subnet name.	(subnet name)
Type	Displays the subnet type.	VLAN
VLAN ID	Displays the VLAN identification number.	(ID number)
Virtual Switch	Displays the virtual switch that is configured for the VLAN you selected. The default value is the default virtual switch *vs0*. Note: The virtual switch name is displayed only if you add a VLAN ID in the VLAN ID field.	(virtual switch name)
IP Address Management	Displays the IPv4 Address of the network with the prefix .	(IPv4 Address/prefix)
Cluster Name	Displays the name of the cluster for which this subnet is configured.	(cluster name)

To filter the list by network name, enter a string in the filter field. (Ignore the **Filters** pane as it is blank.)

To view focused fields in the List, select the focus parameter from the **Focus** drop down list. You can create your own customised focus parameters by selecting **Add custom** from the drop down list and selecting the necessary fields after providing a **Name**, in the Subnet Columns.

There is a **Network Config** action button to configure a new network (see Configuring Network Connections on page 423). The **Actions** menu appears when one or more networks are selected and includes a **Manage Categories** option (see Assigning a Category on page 415).

Policies Entities

You can access dashboards for the following policy types from the **Policies** category of the entities menu (see Entities Menu on page 13):

- Security Policies (see Security Policies Summary View on page 137)

- NGT Policies (see NGT Policies View on page 139)
- Image Placement Policies (see Image Placement Policies on page 480)

Security Policies Summary View

To access the security policies dashboard, select **Policies** > **Security Policies** from the entities menu (see Entities Menu on page 13). The security policies dashboard allows you to view summary information about defined security policies.

> Note: This section describes the information and options that appear in the security policies dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See *Flow Microsegmentation Guide* for information about how to create and apply security policies.

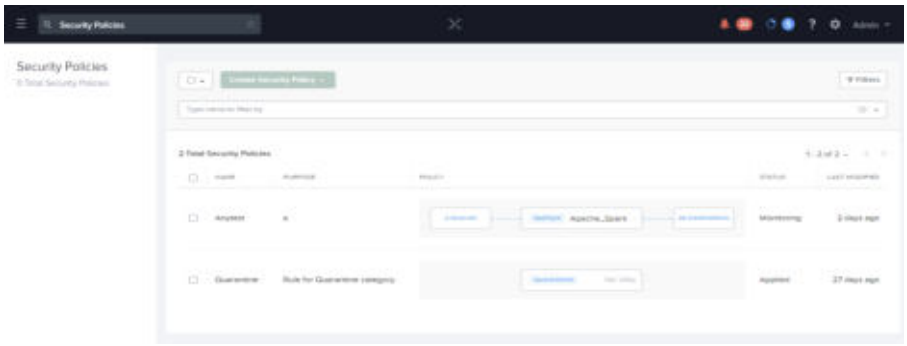

Figure 91: Security Policies Dashboard

The following table describes the fields that appear in the security policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 45: Security Policies List Fields

Parameter	Description	Values
Name	Displays the policy name. The policy is one of three types: application, quarantine, or isolation.	(name), Application, Quarantine, Isolation
Purpose	Describes (briefly) the policy's purpose.	(text string)
Policy	Displays (high level) what the policy does.	(boxed text)
Status	Displays the current status of the policy (either applied currently or in monitoring mode).	Applied, Monitoring
Last Modified	Displays the date the policy was last modified (or the creation date if the policy has never been modified).	(date)

You can filter the security polices list based on several parameter values. The following table describes the filter options available when you open the Security Policies view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 46: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of security policies that satisfy the name condition/string.	(policy name string)
Type	Filters on the policy type. Check the box for one or more of the policy types (application, quarantine, isolation). It will limit the list to just those policy types.	Application, Quarantine, Isolation
Status	Filters on the policy status. Check the box for applied or monitoring.	Applied, Monitoring

The security policies dashboard includes a **Create Security Policy** action button with a drop-down list to **Secure an Application** or **Isolation Environments**.

The **Actions** menu appears when one or more policies are selected. It includes options to update, apply, monitor, and delete. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Security Policy Details View

To access the details page for a security policy, click on the desired security policy name in the list (see Security Policies Summary View on page 137). The Security Policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The rule status appears below the name and indicates whether the policy is being applied currently or is in monitoring mode.
- Three columns appear that specify the Inbound policy (on the left), the affected entities (in the middle), and the Outbound policy (on the right).
- There are three action buttons (upper right).
    - Click the appropriate button to update, apply, monitor, or delete the policy (see “” () for details). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
    - Click the question mark icon to open a help page in a separate tab or window.
    - Click the X icon to close the details page.

Figure 92: Security Policy Details View: Monitoring Rule Example

Figure 93: Security Policy Details View: Applied Rule Example

NGT Policies View

To access the Nutanix guest tools (NGT) policies dashboard, select **Policies** > **NGT Policies** from the entities menu (see Entities Menu on page 13). The NGT policies dashboard allows you to view information about existing NGT policies and create new policies.

Note: This section describes the information and options that appear in the NGT policies dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See NGT Policies on page 478 for information on how to create or modify NGT policies.

The NGT dashboard displays a list of current policies that includes the name and type for each policy.

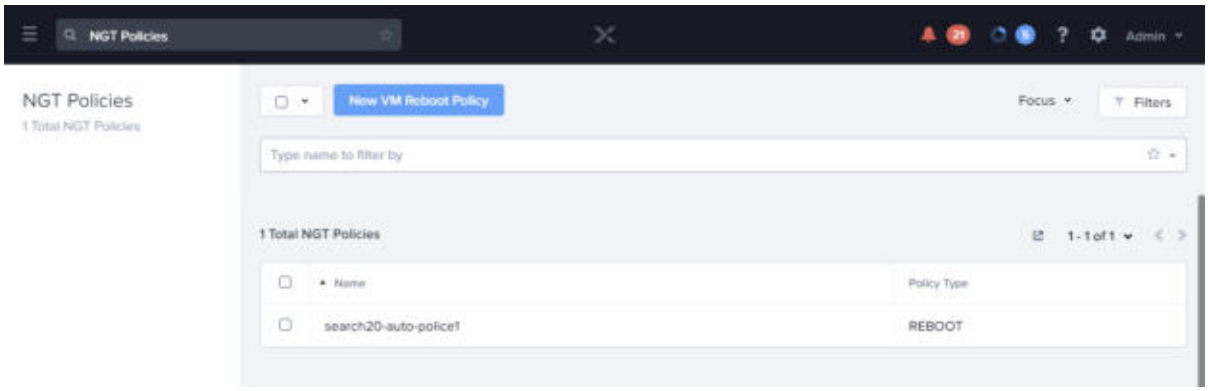

Figure 94: NGT Policies Dashboard

To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following parameters and values.

- **Name**: Enter name in search field.
- **Created By**: System, User

There is a **New VM Reboot Policy** action button to create a new reboot policy (see Creating a New Restart Policy on page 478). The **Actions** menu appears when one or more policies are selected and includes the following options:

- **Update** (see Modifying a Restart Policy on page 480)
- **Delete**

Image Placement Policies Summary View

The image placement policies dashboard allows you to view summary information about defined image placement policies.

To access the image placement policies dashboard, select **Policies** > **Image Placement** from the entities menu (see Entities Menu on page 13).

> Note: This section describes the information and options that appear in the image placement policies dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See Image Placement Policies on page 480 for information about how to create and apply image placement policies.

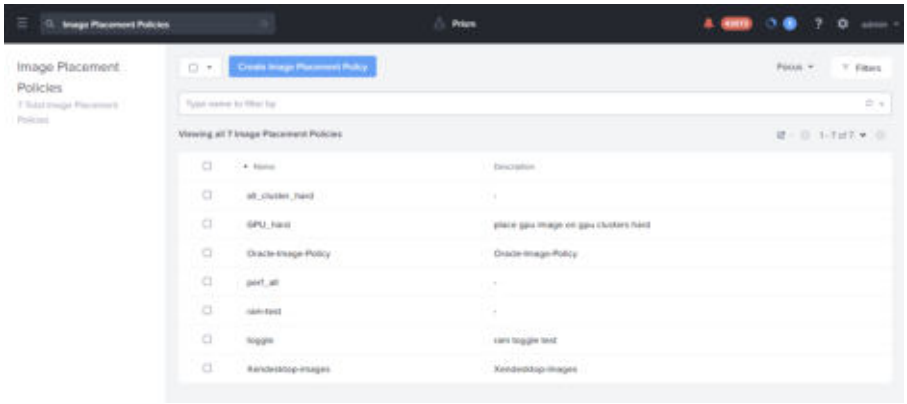

Figure 95: Image Placement Policies Dashboard

The following table describes the fields that appear in the image placement policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 47: Image Placement Policies List Fields

Parameter	Description	Values
Name	Displays the policy name.	(name)
Description	Describes (briefly) the policy's purpose.	(text string)

You can filter the image placement polices list by opening the **Filter** pane. To apply a filter, check the box and enter a text string in the **Contains** field for the Name or Description parameter. You can apply filters across both parameters.

The image placement policies dashboard includes a **Create Image Placement Policy** action button (see Configuring an Image Placement Policy on page 482).

The **Actions** menu appears when one or more policies are selected. It includes options to update and delete. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Image Placement Policy Details View

To access the details page for an image placement policy, click on the desired policy name in the list (see Image Placement Policies Summary View on page 140). The image placement policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The policy enforcement status appears below the name and indicates whether the enforcement is hard or soft.
- Two columns appear that specify the images in question (on the left) and the target clusters (on the right). Click **Expand all** (or **Collapse all**) to display (or hide) the list of images or clusters.
- There are four action buttons (upper right).
  - Click the **Update** or **Delete** button to update or delete the policy (see Updating an Image Placement Policy on page 483). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

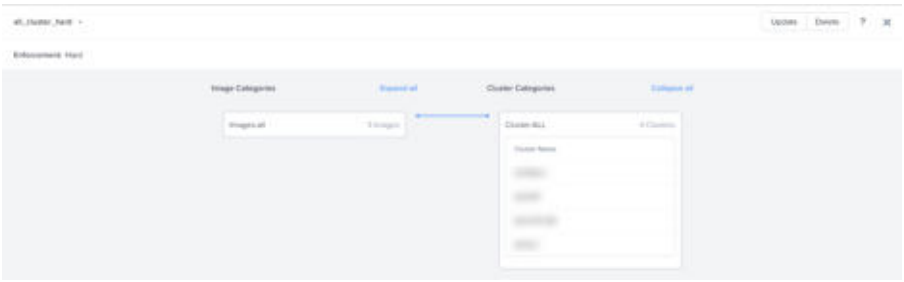

Figure 96: Image Placement Policy Details View

Data Protection Entities

You can access dashboards for the following data protection types from the **Data Protection** category of the entities menu (see Entities Menu on page 13):

- Protection Summary
- Protection Policies (see Protection Policies View on page 142)
- Recovery Plans (see Recovery Plans View on page 143)
- VM Recovery Points (see VM Recovery Points View on page 144)

Protection Policies View

To access the protection policies dashboard, select **Policies** > **Protection Policies** from the entities menu (see Entities Menu on page 13). The protection policies dashboard allows you to view information about the current policies and create new policies.

Note: This section describes the information and options that appear in the protection policies dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See the *Leap Administration Guide* for information about how to create and use protection policies.

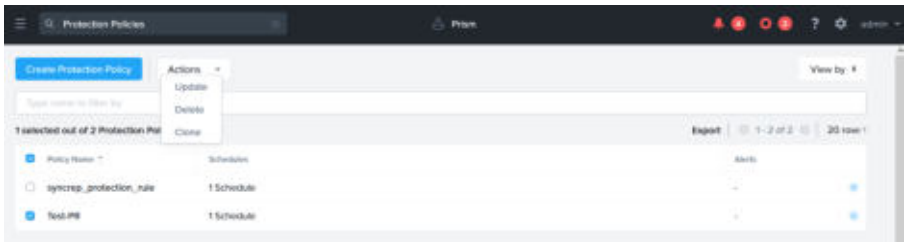

Figure 97: Protection Policies Dashboard

The following table describes the fields that appear in the protection policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 48: Protection Policies Fields

Parameter	Description	Values
Name	Displays the protection policy name.	(item name)
Source	Displays the replication source name for the protection policy	(source name)
Destination	Displays the replication destination name for the protection policy.	(destination name)
RPO	Displays the recovery point objective (RPO) for the protection policy	(time interval)

Parameter	Description	Values
Remote Retention	Displays the number of retention points at the remote availability zone.	*xx* Recovery Points
Local Retention	Displays the number of retention points at the local availability zone.	*xx* Recovery Points

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

Click the **Create Protection Policy** button to create a new protection policy. The **Actions** menu appears when one or more protection policies are selected and includes the following options:

- **Update**: Update the protection policy.
- **Clone**: Clone the protection policy.
- **Delete**: Delete the protection policy.

Recovery Plans View

To access the recovery plans dashboard, select **Policies** > **Recovery Plans** from the entities menu (see Entities Menu on page 13). The recovery plans dashboard allows you to view information about the current plans and create new plans.

> Note: This section describes the information and options that appear in the recovery plans dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See the *Leap Administration Guide* for information about how to create and use recovery plans.

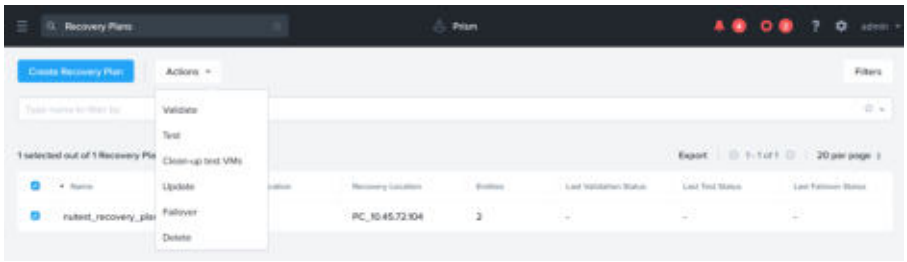

Figure 98: Recovery Plans Dashboard

The following table describes the fields that appear in the recovery plans list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 49: Recovery Plans Fields

Parameter	Description	Values
Name	Displays the recovery plan name.	(name)

Parameter	Description	Values
Source	Displays the primary availability zone for the recovery plan.	(zone name)
Destination	Displays the recovery availability zone for the recovery plan.	(zone name)
Entities	Displays the number of VMs and snapshots associated with the recovery plan in the availability zone in which you are viewing the recovery plan.	(integer)
Last Validation Status	Displays the status of the most recent validation of the recovery plan.	
Last Test Status	Displays the status of the most recent test performed on the recovery plan.	

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

Click the **Create Recovery Plan** button to create a new recovery plan. The **Actions** menu appears when one or more recovery plans are selected and includes the following options:

- **Validate**: Clone a recovery plan.
- **Test**: Test the recovery plan.
- **Update**: Update the recovery plan.
- **Failover**: Perform a failover operation.
- **Delete**: Delete the recovery plan.

VM Recovery Points View

To access the recoverable entities dashboard, select **Virtual Infrastructure** > **VM Recovery Points** from the entities menu (see Entities Menu on page 13). A recoverable entity is a VM that you can recover from a snapshot. The recoverable entities dashboard allows you to view information about the entities for which snapshots exist in the availability zone.

> Note: This section describes the information and options that appear in the recoverable entities dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See the *Xi Leap Administration Guide* for information about recoverable entities and how to use them.

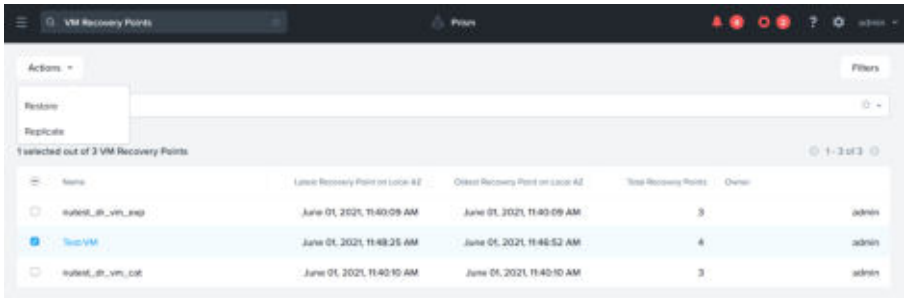

Figure 99: VM Recovery Points Dashboard

The following table describes the fields that appear in the recoverable entities list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 50: VM Recovery Points List Fields

Parameter	Description	Values
Name	Displays the recoverable entity name.	(item name)
Latest Recovery Point on Local AZ	Displays the time at which the most recent snapshot was taken. This is specific to the local availability zone.	(time and date)
Oldest Recovery Point on Local AZ	Displays the time at which the oldest available snapshot was taken. This is specific to the local availability zone.	(time and date)
Total Recovery Points.	Displays the total number of recovery points available locally.	(integer)
Owner	Displays the owner who created the recovery points	(user name)

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

The **Actions** menu appears when one or more recoverable entities are selected and includes the following options:

- **Restore**: Restore a VM entity from the snapshot taken at the recovery point.
- **Replicate**: Replicate the snapshot to a paired availability zone.

Hardware Entities

You can access dashboards for the following hardware components from the **Hardware** category of the entities menu (see Entities Menu on page 13):

- Clusters (see Clusters Summary View on page 146)
- Hosts (see Hosts Summary View on page 165)
- Disks (see Disks Summary View on page 175)
- GPUs (see GPUs Summary View on page 183)

Clusters Summary View

To access the clusters dashboard, select **Hardware** > **Clusters** from the entities menu (see Entities Menu on page 13). The clusters dashboard allows you to view summary information about registered clusters and access detailed information about each cluster. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> Note: This section describes the information and options that appear in the clusters dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities**: Displays a list of the clusters with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **IOPS**, and **Aggregate Hypervisor Memory Usage Pp**. Click the **View all XX Clusters** link at the bottom to display the **List** tab (following section).

- **Alerts**: Displays a list of cluster-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.

- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) , **Last 1 Hours**, or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

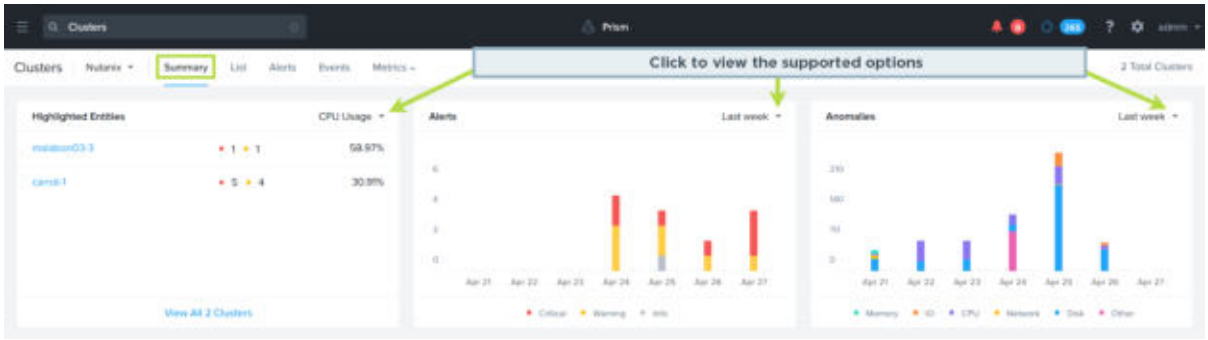

Figure 100: Clusters Summary Tab (Nutanix)

Figure 101: Clusters Summary Tab (Non-Nutanix)

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the registered clusters. The following table describes the fields that appear in the clusters list. The fields vary based on the **View by** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

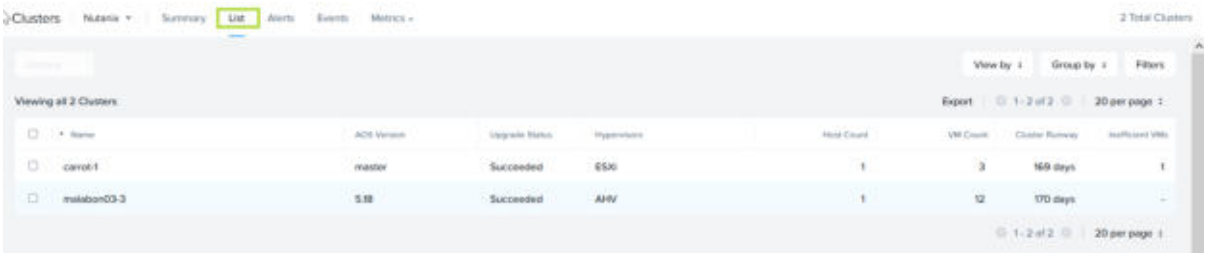

Figure 102: Clusters List Tab (Nutanix)

Figure 103: Clusters List Tab (Non-Nutanix)

Table 51: Clusters List Fields

Parameter	Description	Values
View by "General" Fields		
Name	Displays the cluster name. Clicking on the name displays the details page for that cluster (see Cluster Details View on page 153).	(cluster name)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Upgrade Status	Displays the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, Hyper-V
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)

Parameter	Description	Values
VM Count	Displays the total number of VMs in the cluster (in any state).	(number of VMs)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 153 for more information).	(number of days)
Inefficient VMs	Displays the number of inefficient VMs in the cluster (see Behavioral Learning Tools on page 582 for more information).	(number)
View by "Performance" Fields		
Name	Displays the cluster name.	(cluster name)
CPU Usage	Displays the percentage of CPU capacity in the cluster currently being used.	0 -100%
Memory Usage	Displays the percentage of memory capacity in the cluster currently being used.	0 -100%
IOPS	Displays total (both read and write) I/O operations per second (IOPS) for this cluster.	(number)
IO Bandwidth	Displays total I/O bandwidth used per second in this cluster.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency in this cluster.	xxx [ms]
View by "Capacity" Fields		
Name	Displays the cluster name.	(cluster name)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 153 for more information).	(number of days)
CPU Runway	Displays the predicted CPU runway (time period) before the cluster requires additional resources.	(number of days)
Memory Runway	Displays the predicted memory runway (time period) before the cluster requires additional resources.	(number of days)
Storage Runway	Displays the predicted storage memory runway (time period) before the cluster requires additional resources.	(number of days)
View by "Encryption" Fields		
Name	Displays the cluster name.	(cluster name)

Parameter	Description	Values
Encryption Type	Displays the cluster encryption type.	Not Encrypted, Software, Hardware, Software and Hardware
Encryption Scope	Displays the encryption scope.	Cluster, Entity
KMS Type	Displays the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)

You can filter the cluster list based on a variety of parameter values. The following table describes the filter options available when you open the Clusters view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 52: Filter Pane Fields

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the pull-down list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)
Name	Filters on the cluster name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of clusters that satisfy the name condition/ string.	(cluster name string)
AOS Version	Filters on AOS version. Select one or more versions to return a list of clusters running those version(s). The number of clusters currently running each version is displayed on the right of the line.	(Acropolis version numbers across clusters currently)
Hypervisors	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, HyperV
Health	Filters on the cluster health state (good, warning, or critical). Select one or more states to return a list of clusters in that state(s). The number of clusters currently in each state is displayed on the right of the line.	Critical, Warning, Good

Parameter	Description	Values
CPU Usage	Filters on the amount of total CPU being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
IOPS	Filters on the total (both read and write) IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of clusters with total IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the total I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of clusters with total I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of clusters with average I/O latency in that range.	([xx] to [yy] range)
Upgrade Status	Filters on the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Categories	Filters on the category type. Search for the category name. For example, ADGroup: $Default	(category name string)
Encryption Type	Filters on the encryption type. Check the box for the desired encryption type. It will return a list of clusters with the corresponding encryption type applied.	Not Encrypted, Software, Hardware, Software and Hardware
Encryption Scope	Filters on the encryption scope. Check the box for the entities where encryption is applied on. It will return a list of clusters with the corresponding encryption scope applied.	Cluster, Entity
KMS Type	Filters on the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External

You can group the clusters list in the following ways:

- The **Group by** pull-down menu allows you to group the cluster entries by AOS version, host count, or health state. (You can only choose one.)
- When you select one or more clusters, the **Label** icon appears (to the left of the **Actions** menu). To assign the selected clusters a group label, click the **Label** icon and do one of the following:

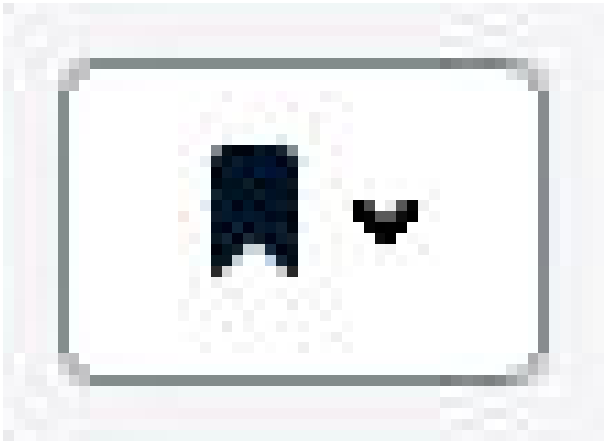

Figure 104: Label Icon

  - To assign an existing label (if there are existing labels), select the desired label and then click **Apply changes**.
  - To assign a new label, enter a meaningful label in the blank text field and then click **Create new label**.

  You can now perform any actions available on this labeled group from the **Actions** menu (see following section). You can edit or delete the label by clicking the **Label** icon and selecting **Manage Labels**.

The **Actions** menu appears when a cluster is selected. The pull-down list includes the following actions:

- Select **Launch Prism Element** to launch Prism element for that cluster in a separate tab or window (depending on your browser settings).

  Note: When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user drop-down list are **REST API Explorer**, **About Nutanix**, **Support Portal**, **Help**, **Nutanix Next Community**, and **Sign Out**.

-
- Select **Upgrade Software** to upgrade the AOS version on that cluster (see Installing or Upgrading on page 31).
- Select **Rack Configuration** to configure the rack awareness feature. This option appears only for clusters that satisfy the conditions for rack awareness. See the *Prism Web Console Guide* for instructions on configuring rack awareness.
- Select **Manage Categories** to configure cluster categories.
- Select **Disable Efficiency Measurement** to disable efficiency measurement for the cluster.

- Select **Disable Anomaly Detection** to disable anomaly detection for the cluster.
- Select **Run Playbook** to run or create a new playbook for the cluster entity type.
- Select **Download Encryption Keys** to download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just cluster-related alerts (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just cluster-related events (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the clusters. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information below the tab. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

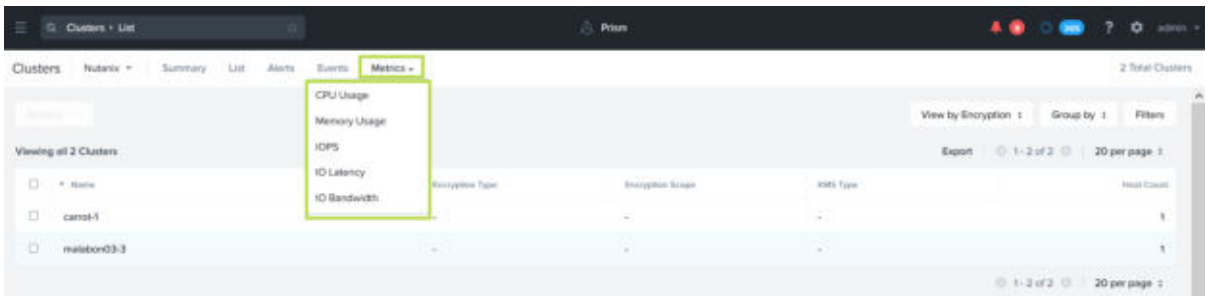

Figure 105: Clusters Metrics Tab

Table 53: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total clusters (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those clusters. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.

Metric	Description
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Cluster Details View

To access the details page for a cluster, go to the clusters dashboard **List** tab (see Clusters Summary View on page 146) and click the cluster name. You can also access the details page by clicking the cluster name wherever that name appears, such as in a dashboard widget or search result.

The cluster name and the following set of tabs appear on the top: **Summary**, **Alerts**, **Events**, **Metrics**, **Storage Usage**, **Hardware**, **Virtual Entities**, and **Capacity**. Click a tab to display that information below the tab. (Click the **Back to Clusters** link to return to the summary view.)

> Note:
>
> For Non-Nutanix clusters, the cluster name and the following set of tabs appear on the top: **Summary**, **Alerts**, **Events**, **Metrics**, and **Capacity**. Click a tab to display that information below the tab. (Click the **Back to Clusters** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the cluster (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.
- A **VM Efficiency** widget that displays the number of VMs, which are considered inefficient broken down by category (overprovisioned, inactive, constrained, and bully). See Behavioral Learning Tools on page 582 for more information about VM efficiency.
- Action buttons (only the applicable ones appear):
  - **Launch Prism Element**: Click this button to launch Prism Element for this cluster in a new tab or window.

    > Note: When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user drop-down list are

REST API Explorer, **About Nutanix**, **Support Portal**, **Help**, **Nutanix Next Community**, and **Sign Out**.

- **Upgrade Software**: Click this button to upgrade the AOS version on this cluster (see Installing or Upgrading on page 31).
- **Rack Configuration**: Click this button to configure the rack awareness feature (see the *Prism Web Console Guide*).
- Select **Manage Categories** to configure cluster categories (see Category Management on page 413).
- Select **Disable Efficiency Measurement** to disable efficiency measurement for the cluster.
- Select **Disable Anomaly Detection** to disable anomaly detection for the cluster.
- Select **Run Playbook** to run or create a new playbook for the cluster entity type (see Creating Playbooks using Triggers on page 596)..
- Select **Download Encryption Keys** to download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.

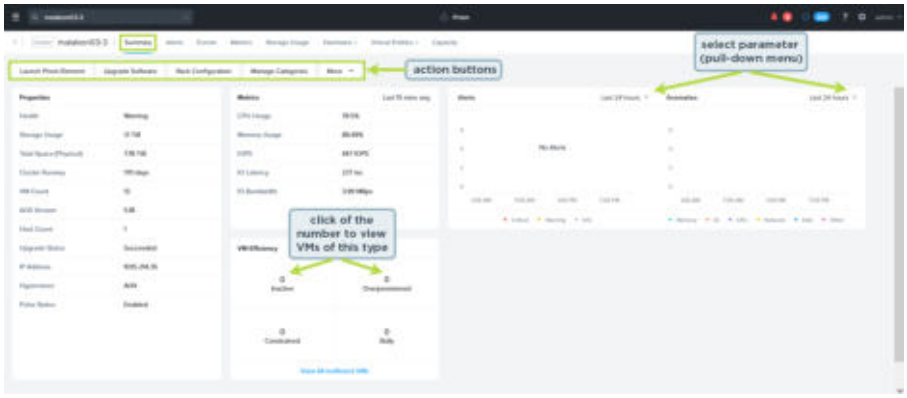

Figure 106: Cluster Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 54: Cluster Properties Fields

Parameter	Description	Values
Health	Displays the cluster health state (good, warning, or critical).	Critical, Warning, Good
Storage Usage	Displays the amount of storage used in the cluster	xxx [GiB|TiB]
Storage Capacity	Displays the total amount of storage capacity in this cluster.	xxx [GiB|TiB]

Parameter	Description	Values
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see the "Capacity Tab" section below).	(number of days)
VM Count	Displays the number of VMs in the cluster.	(number of VMs)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of hosts)
Upgrade Status	Displays the status of the last (or current) upgrade attempt.	(status condition)
IP Address	Displays the virtual IP address for the cluster (if defined).	(IP address)
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, or Hyper-V

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this cluster (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this cluster (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the cluster. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric below the tab. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 582). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).

- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 230).

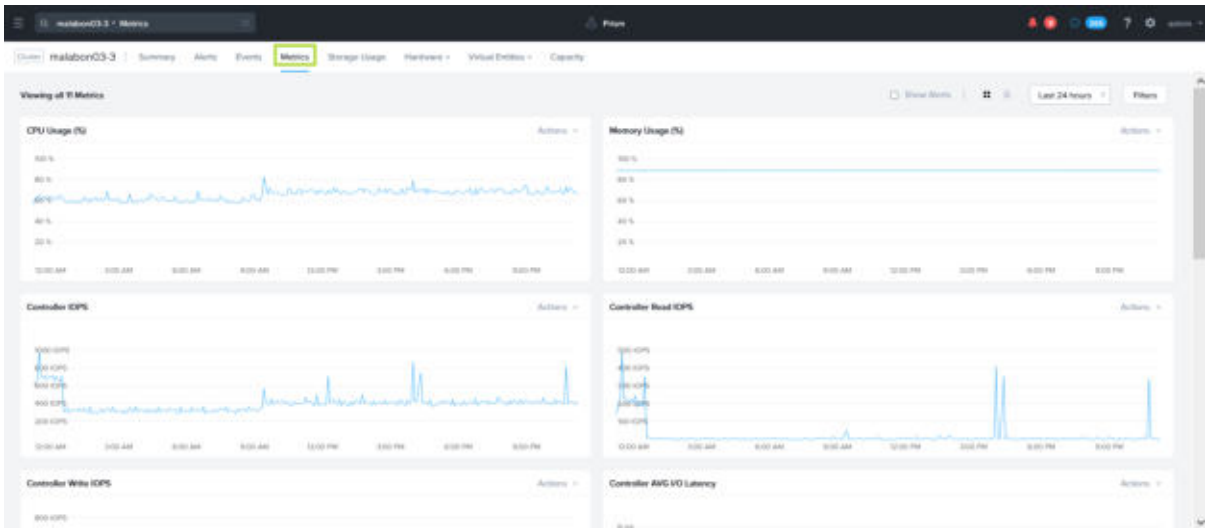

Figure 107: Cluster Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 55: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the cluster (0–100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the cluster (0–100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the cluster.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the cluster.
I/O Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the cluster.

Usage Tab

The **Usage** tab displays the following graphs:

- The **Cluster-wide Usage Summary** graph displays a rolling time interval monitor of total storage usage across the cluster that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of storage space used by each disk tier (SSD and DAS-SATA) across the cluster.

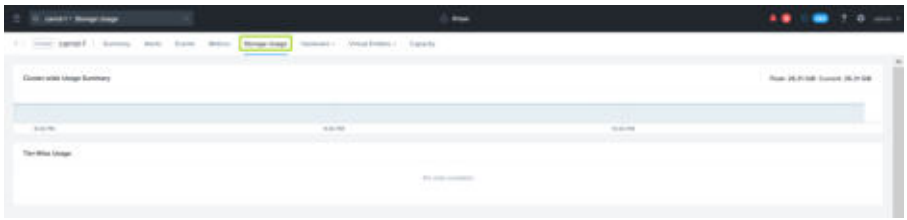

Figure 108: Cluster Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist in this cluster. Clicking a hardware entry such as **Hosts** or a virtual entity such as **Containers** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those in this cluster. See Hardware Entities on page 145 and Virtual Infrastructure Entities on page 82 for more information.

Capacity Tab

The **Capacity** tab displays current and historical usage information and provides resource planning tools. It includes the following sections:

Note:

- The capacity planning feature requires a Prism Pro license. If Prism Pro is disabled, the **Capacity** tab is grayed out and not available.
- Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
- Capacity planning checks the past 90 days of data for monthly seasonality, it also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
- For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 540 for more information.

- A **<cluster_name> Runway** area (upper left) that displays the current cluster runway overall and broken down by CPU, memory, and storage. Runway refers to how long the cluster can continue to run normally, based on the current consumption rate, before the existing resources are used to capacity. The overall runway value is the same as the lowest value for any single resource (storage, CPU, or memory). For example, if the storage runway is 14 days

while the CPU and memory runways are 45 days, the overall runway is 14 days (the storage value).

> Note: A plus sign in any runway value, for example the "89+" in the Cluster Capacity Tab figure, indicates that the predicted runway is more than that number of days, but the estimation stopped at that number.

- Clicking **Storage Runway** displays a storage usage graph and table (to the right).
- Clicking **CPU Runway** displays a CPU usage graph and table.
- Clicking **Memory Runway** displays a memory graph and table.

- An **Optimize Resource** button (middle left). Clicking this button opens the **Optimize Resources** window, which identifies areas to improve cluster resource allocation or capacity. Recommendations might include

  - Table of "inefficient" (over-provisioned, inactive, constrained, and bully) VMs based on the VM behavioral learning engine (see Behavioral Learning Tools on page 582). Click the appropriate link to see a list of the VMs in that state.
  - An add node recommendation. The recommendation is based on the current (historical) CPU, memory, and storage usage demand across the cluster.

  Click the **Get Report** link to run the Cluster Efficiency report (see Reports Management on page 617).

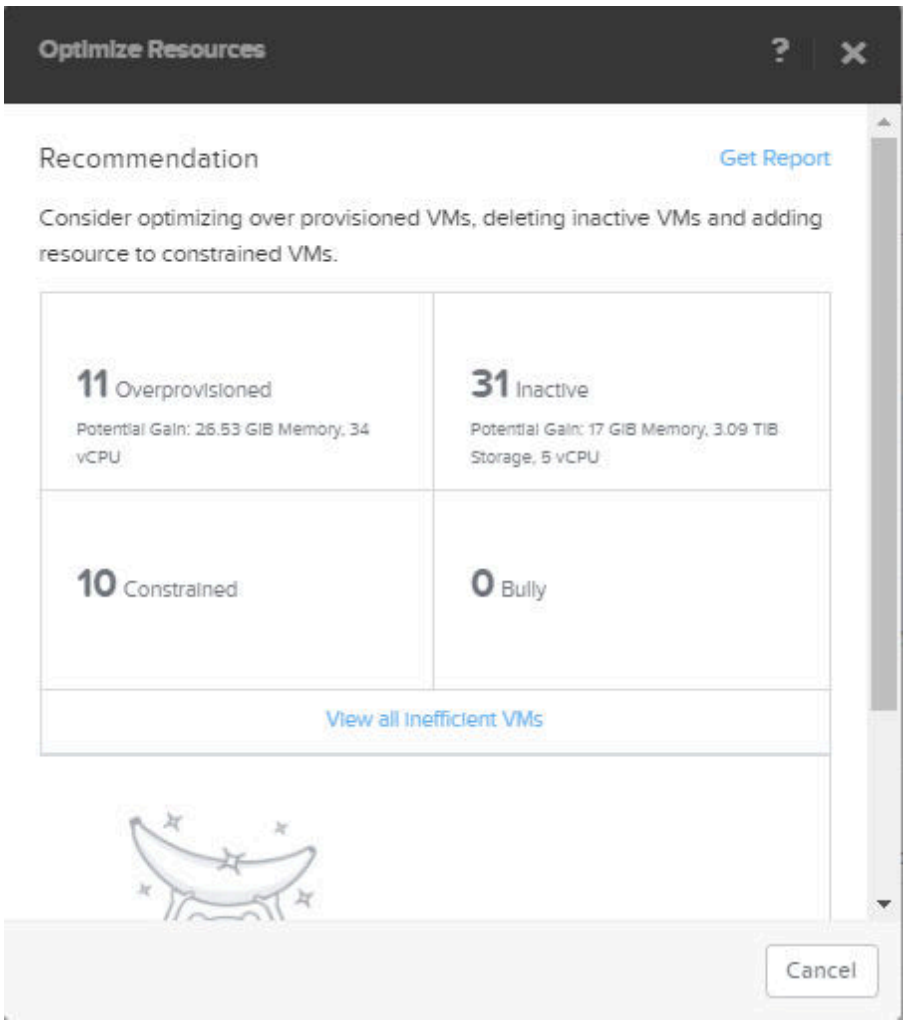

Figure 109: Optimize Resources Window

- A **Get Started** button (lower left). Clicking this button opens the resource planning scenario page. From this page you can create "what if" scenarios to test future workloads against current or hypothetical storage, CPU, and memory resource capacity (see Creating a Scenario on page 564).
- A usage history graph (upper right). The graph displays storage, CPU, or memory usage depending on which is selected in the left column. Placing the cursor anywhere on the horizontal axis displays the value(s) at that time. A (solid red) maximum capacity line

appears near the top of the graph, and an additional (dotted red) effective capacity line appears on the storage graph. A (vertical) "Today" line is displayed to indicate where the current date is located on the graph. When there is an alert or event, a bell icon appears either on the Today line or at the starting point of the alert. Clicking the bell icon (red for critical alert, orange for warning alert, or gray for event) displays more information about that alert or event.

- A table (lower right) that displays usage information. The information in the table varies depending on which runway (storage, CPU, or memory) is selected.

Figure 110: Cluster Capacity Tab (Nutanix Cluster)

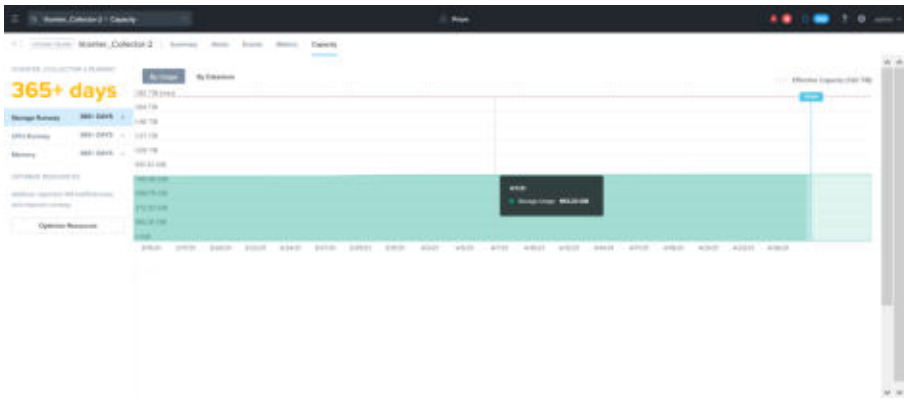

Figure 111: Cluster Capacity Tab (Non-Nutanix Cluster)

Capacity Tab: Storage View

When **Storage Runway** is selected, the graph and table display storage usage information. The maximum and effective storage capacity values (in TiBs) for the cluster appear above the graph. Effective capacity of a cluster is computed based on the configured reservations (from the capacity configurations settings) and the replication factor for all the containers in the cluster.

The storage view includes two tabs:

- Click the **By Usage** button to see cluster-wide storage information (see previous figure).
- Click the **By Storage Container** button to see storage information per storage container. You can select all storage containers from the pull-down list (upper left of graph) or a single

storage container from the pull-down list or by clicking the storage container name in the table.

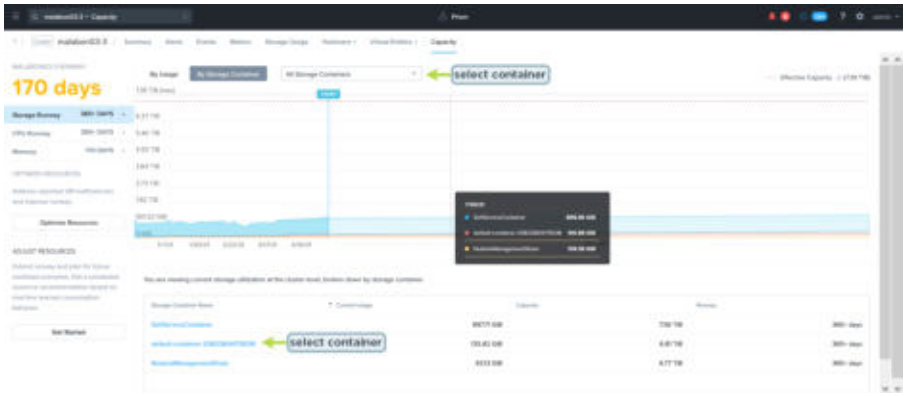

Figure 112: Storage Container View (Nutanix Cluster)

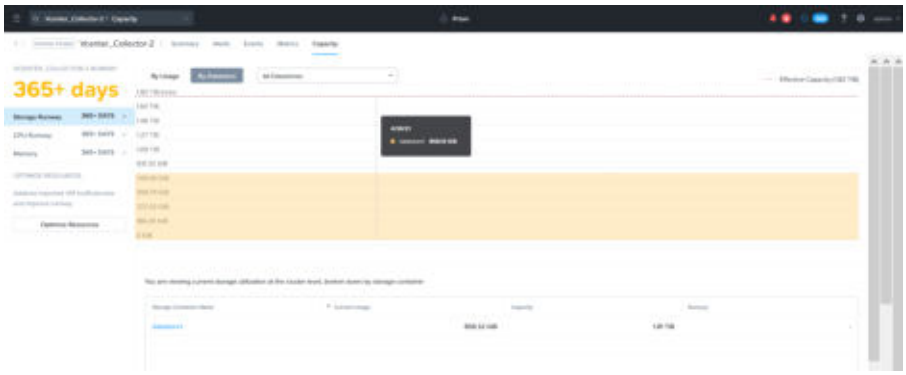

Figure 113: Storage Container View (Non-Nutanix Cluster)

The following table describes the fields in the storage table.

Table 56: Capacity Tab: Storage Fields

Parameter	Description	Values
"By Usage" tab (cluster-wide)		

Parameter	Description	Values
Name	Displays a descriptive name for a type of storage. There are four types: • Live Usage: The amount of live storage. • Snapshot Usage: The amount of storage used for snapshots. • System Usage: The amount of storage for everything else (total usage - (reserved + snapshot + live)). System usage includes garbage, parity, and other miscellaneous uses. • Reserved Usage: The amount of storage reserved for thick provisioned VMs (ESXi).	Live Usage, Reserved Usage, Snapshot Usage, System Usage
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB|TB]
"By Storage Container" tab: All Storage Containers		
Storage Container Name	Displays the name of the storage container.	(storage container name)
Current Usage	Displays the amount of storage space used by the storage container currently.	xxx [GB|TB]
Capacity	Displays the total amount of storage capacity available to the storage container.	xxx [TB]
"By Storage Container" tab: <storage container name>		
Storage Container Name	Displays one of the following: • If "All Storage Containers" is selected, the storage container names are displayed. Click a name to see the details for that storage container. • If a single storage container is selected, it displays the four storage types described for the Name field in the "By Usage" tab (Live Usage, Reserved Usage, Snapshot Usage, and System Usage).	(storage container name) OR Live Usage, Reserved Usage, Snapshot Usage, System Usage
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB|TB]

Capacity Tab: CPU View

When **CPU Runway** is selected, the graph and table display CPU usage information. The maximum and effective CPU capacity values (in GHz) for the cluster appear above the graph. The CPU view includes two tabs:

- Click the **Overall** button to see cluster-wide CPU information.

- Click the **By Host** button to see CPU information per host. You can select all hosts from the pull-down list (upper left of graph) or a single host from the pull-down list or by clicking the host name in the table.

The following table describes the fields in the CPU table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the CPU table.

Table 57: Capacity Tab: CPU Fields

Parameter	Description	Values
"Overall" tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM currently.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)
"By Host" tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average CPU Usage over last 1 hour	Displays the average host CPU usage during the last hour.	xxx [GHz]
Current CPU Capacity	Displays the CPU capacity of the host.	xxx [GHz]
Runway	Displays the current runway (estimated time interval that current CPU capacity is sufficient for the workload).	(number of days)
"By Host" tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)

Capacity Tab: Memory View

When **Memory Runway** is selected, the graph and table display memory usage information. The maximum and effective memory capacity values (in GiBs) for the cluster appear above the graph. The memory view includes two tabs:

- Click the **Overall** button to see cluster-wide memory information.
- Click the **By Host** button to see memory information per host. You can select all hosts from the pull-down list (upper left of graph) or a single host from the pull-down list or by clicking the host name in the table.

The following table describes the fields in the memory table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the memory table.

Table 58: Capacity Tab: Memory Fields

Parameter	Description	Values
"Overall" tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM.	(0–100%)
"By Host" tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average Memory Usage over last 1 hour	Displays the average host memory usage during the last hour.	xxx [GiB]
Runway	Displays the current runway (estimated time interval that current memory capacity is sufficient for the workload).	(number of days)
Current Memory Capacity	Displays the memory capacity of the host.	xxx [GiB]
"By Host" tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM currently.	(0–100%)

Metrics Tab

The **Metrics** tab allows you to view performance metrics for the cluster, both for actual past performance and for predicted future performance. The **Metrics** tab displays a section on the left with a list of metrics.

> Note: The predicted metrics and identified anomalies are based on sophisticated machine-learning capabilities. See Behavioral Learning Tools on page 582 for information about these capabilities and how they are used.

- Clicking a metric displays a graph on the right. (Some metrics have multiple graphs.) The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph. Placing the cursor anywhere on the horizontal axis displays the value at that time. To set the time interval (last 24 hours, last week, last 21 days), select the duration from the pull-down list on the right.

  > Note: The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph may not appear if less than 21 days of data is available.

- To display predicted performance, check the **Prediction** box. This displays projected performance for the next week.

- To create an alert for this cluster based on either behavioral anomalies or status thresholds, click the **Set Alerts** link above the graph (see Creating Custom Alert Policies on page 230).

The following table describes the available metrics.

> Note: Some of these metrics are not available on all hypervisors.

Table 59: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the VM (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the VM (0 - 100%).
I/O Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.

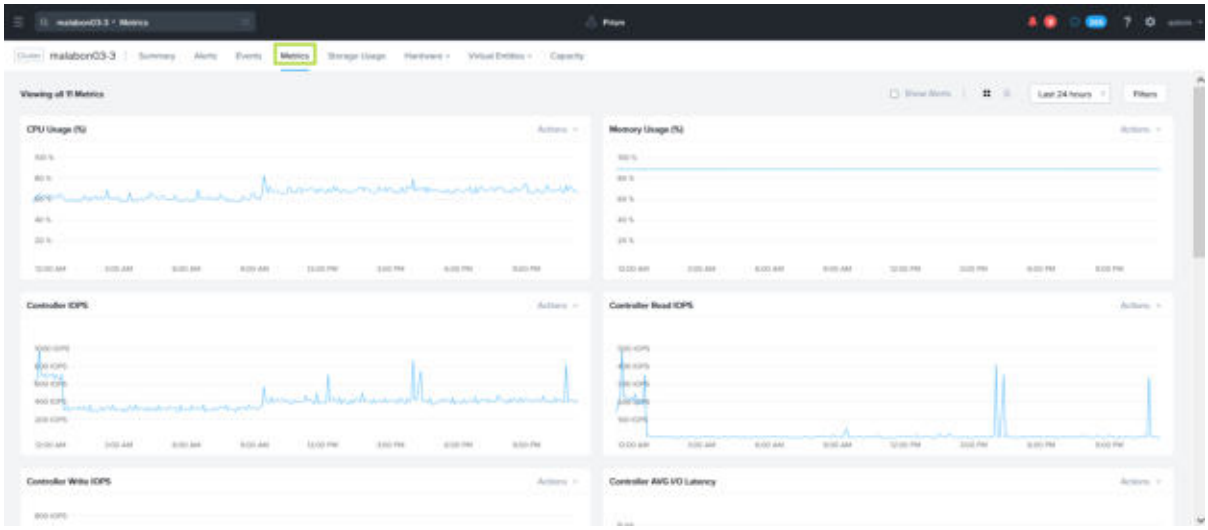

Figure 114: Metrics Tab: CPU Usage metric

Hosts Summary View

To access the hosts dashboard, select **Hardware** > **Hosts** from the entities menu (see Entities Menu on page 13). The hosts dashboard allows you to view summary information about hosts across the registered clusters and access detailed information about each host. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area to the right for the selected tab.

Note: This section describes the information and options that appear in the Hosts dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Suggested**: Displays a list of the hosts with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **Memory Usage**, **IO Latency**, and **IOPS**. Click the **View all XX Hosts** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of host-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

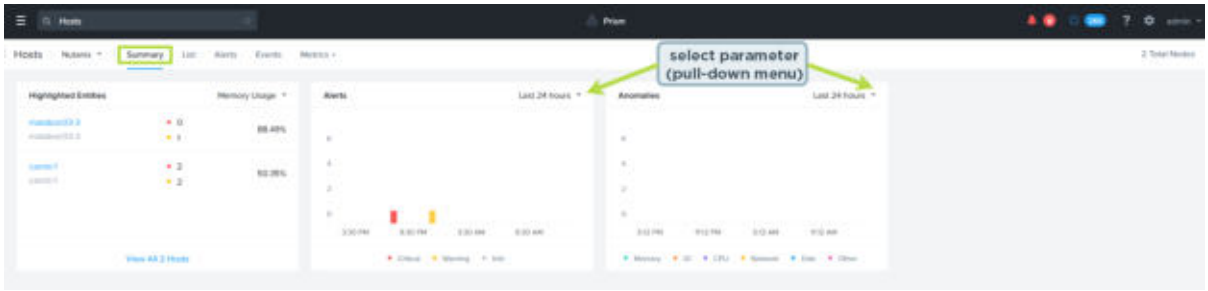

Figure 115: Hosts Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the hosts across the registered clusters. The following table describes the fields that appear in the hosts list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

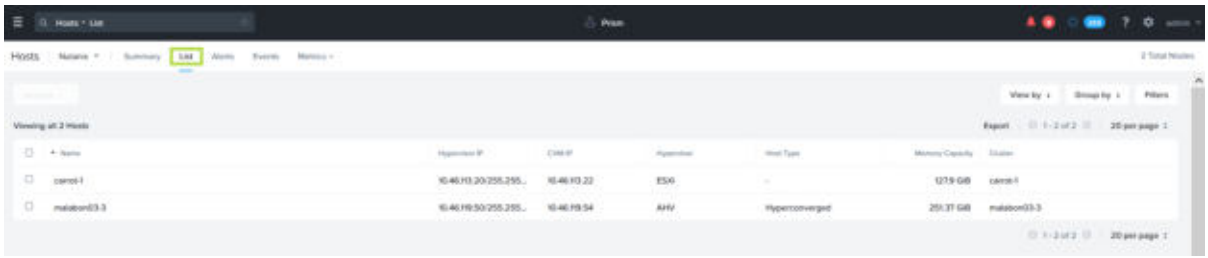

Figure 116: Hosts List Tab

Table 60: Hosts List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the name of the host. Clicking on the name displays the details page for that host (see Host Details View on page 170).	(host name)
Host IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Hypervisor	Displays the hypervisor type running on the host.	AHV, ESX, or Hyper-V
Memory Capacity	Displays the memory capacity of the host.	xxx [MB|GB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
"Performance" Focus Fields		
Name	Displays the name of the host.	(host name)
CPU Usage	Displays the percentage of CPU capacity currently being used by this host.	0 - 100%
Memory Usage	Displays the percentage of memory capacity currently being used by this host.	0 -100%
IOPS	Displays I/O operations per second (IOPS) for this host.	[0 - unlimited]
Disk IO Bandwidth	Displays I/O bandwidth used per second for this host.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this host.	xxx [ms]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)

You can filter the hosts list based on a variety of parameter values. The following table describes the filter options available when you open the Hosts view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 61: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the host name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of hosts that satisfy the name condition/ string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of hosts that reside in the clusters which satisfy the name string.	(cluster name string)
Health	Filters on the host health state (good, warning, or critical). Select one or more states to return a list of hosts in that state(s). The number of hosts currently in each state is displayed on the right of the line.	Critical, Warning, Good
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a pull-down list appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of hosts tagged to each selected category is displayed on the right of the line.	(category name)
Hypervisor	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, HyperV
Memory Capacity	Filters on the host memory capacity. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It will return a list of hosts with memory capacity in that range.	([xx] to [yy] GiB range)
CPU Usage	Filters on the amount of CPU being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of hosts utilizing CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)

Parameter	Description	Values
GPUs	Filters for GPU configuration information such as model name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. As you type a pull-down list appears to help you select the correct configuration information. It will return a list of hosts that satisfy the GPU condition/string.	(configuration info)
IOPS	Filters on the IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of hosts with IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of hosts with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of hosts with average I/O latency in that range.	([xx] to [yy] range)

You can group the hosts list in the following ways:

- The **Group** pull-down menu allows you to group the host entries by cluster, hypervisor type, or health state. (You can only choose one.)
- The **View by** pull-down menu allows you to group the host entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

The **Actions** menu appears when one or more hosts are selected. It includes a **Manage Categories** option (see Assigning a Category on page 415).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just host-related alerts across the registered clusters (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just host-related events across the registered clusters (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

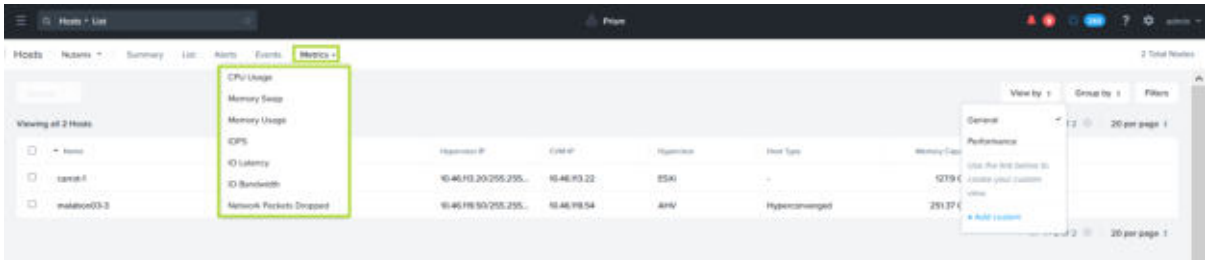

Figure 117: Hosts Metrics Tab

Table 62: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total hosts (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those hosts. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Swap	Displays memory swap-out and swap-in rate tables.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Host Details View

To access the details page for a host, go to the hosts dashboard **List** tab (see Hosts Summary View on page 165) and click the host name. You can also access the details page by clicking the host name wherever that name appear, such as in a dashboard widget or search result.

The host name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, **Usage**, **Hardware**, and **Virtual Entities**. Click a tab to display that information on the right. (Click the **Back to Hosts** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the host (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or

Last week from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

- A **Manage Categories** button to manage (set) categories for the host (see Category Management on page 413).

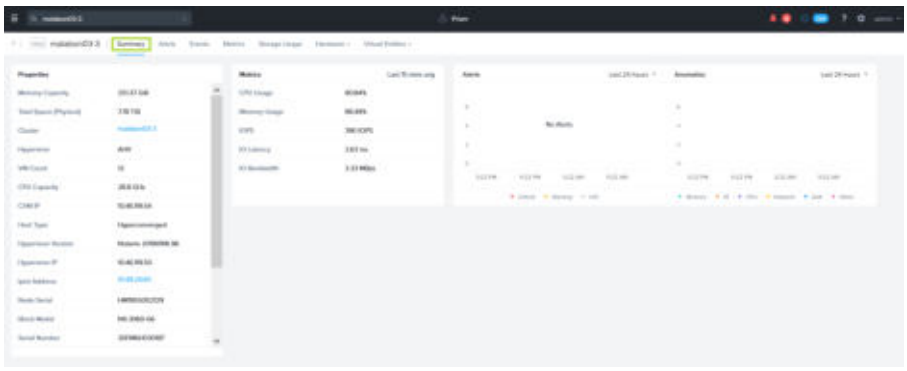

Figure 118: Host Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 63: Host Properties Fields

Parameter	Description	Values
Memory Capacity	Displays the total memory capacity for this host.	xxx [MB|GB]
Disk Capacity	Displays the total amount of disk capacity on this host.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the host resides. Clicking the name displays the details page for that cluster (see Cluster Details View on page 153).	(cluster name)
Host IP	Displays the host IP address.	(IP address)
Hypervisor	Displays the hypervisor name.	(hypervisor name)
VM Count	Displays the number of VMs running on this host.	(number)
Block Model	Displays the block model number.	(model series number)
Serial Number	Displays the block serial number.	(block serial number)
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Host Type		

Parameter	Description	Values
IPMI Address	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console.	(IP address)
Node Serial	Displays the node serial number. The node serial is a unique number passed through from the manufacturer. (The form can vary because it is determined by each manufacturer.)	(manufacturer serial number)
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on the metadata disk.	[0 - 100%]
Oplog Disk Size	Displays the current size of the operations log. (The Oplog maintains a record of write requests in the cluster.) A portion of the metadata disk is reserved for the Oplog, and you can change the size through the nCLI.	xxx [GB]
Monitor Enabled	Displays whether the host is high availability (HA) protected. A **Yes** value means HA is active for this host. A **No** value means VMs on this host are not protected (will not be restarted on another host) if the host fails. Normally, this value should always be **Yes**. A **No** value is likely a sign of a problem situation that should be investigated.	[Yes|No]
Disks	Displays the number of disks in each storage tier in the host. Tier types vary depending on the Nutanix model type.	DAS-SATA: (number), SSD-SATA: (number), SSD-PCIe: (number)
GPUs	Displays the number and type of GPUs in the host. For example, if the host contains four Tesla M10 GPUs, this field displays "Tesla M10 (4)".	(GPU type and number)
Datastore(s)	Displays the names of any datastores.	(names)

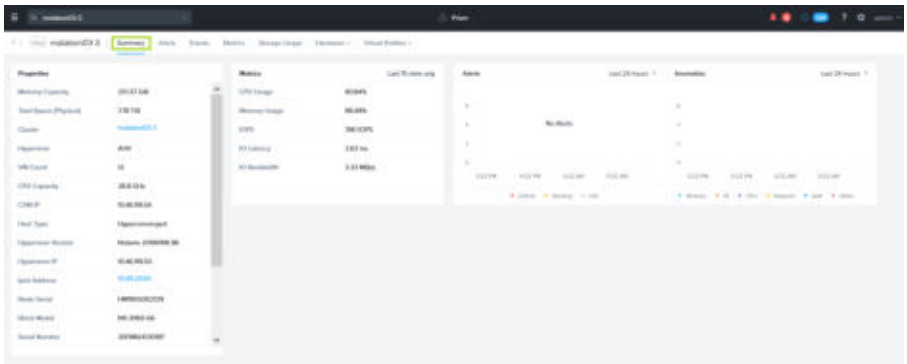

Figure 119: Host Summary Tab

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this host (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this host (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the host. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 582). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 230).

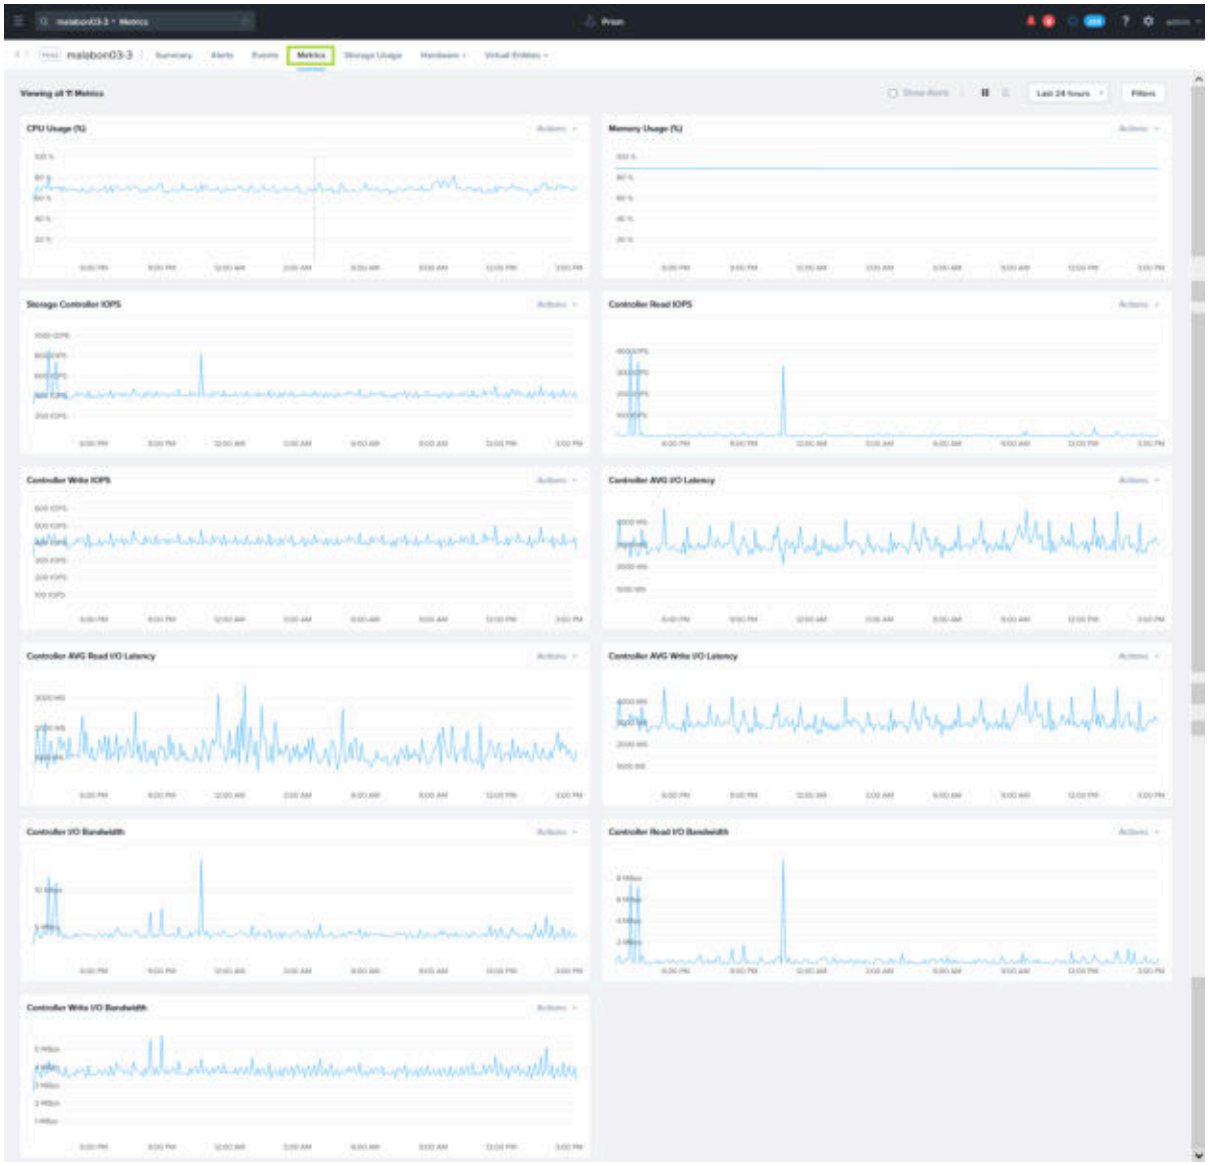

Figure 120: Host Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 64: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the host (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the host (0 - 100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the host.

Metric	Description
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the host.
IO Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the host.

Usage Tab

The **Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of host storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of host storage space used by each disk tier (SSD and DAS-SATA).

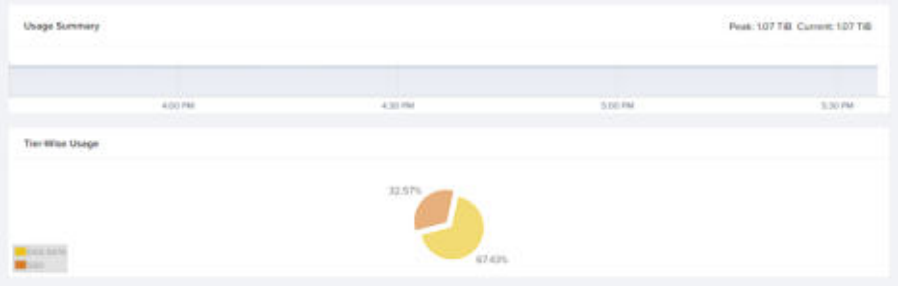

Figure 121: Host Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist on this host. Clicking a hardware entry such as **Disks** or a virtual entity such as **VMs** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those on this host. See Hardware Entities on page 145 and Virtual Infrastructure Entities on page 82 for more information.

Disks Summary View

To access the disks dashboard, select **Hardware** > **Disks** from the entities menu (see Entities Menu on page 13). The disks dashboard allows you to view summary information about disks across the registered clusters and access detailed information about each disk.

> Note: This section describes the information and options that appear in the Disks dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Suggested**: Displays a list of the disks with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Bandwidth**, **IOPS**, **IO Latency**, and **Disk Usage**. Click the **View all XX Disks** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of disk-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

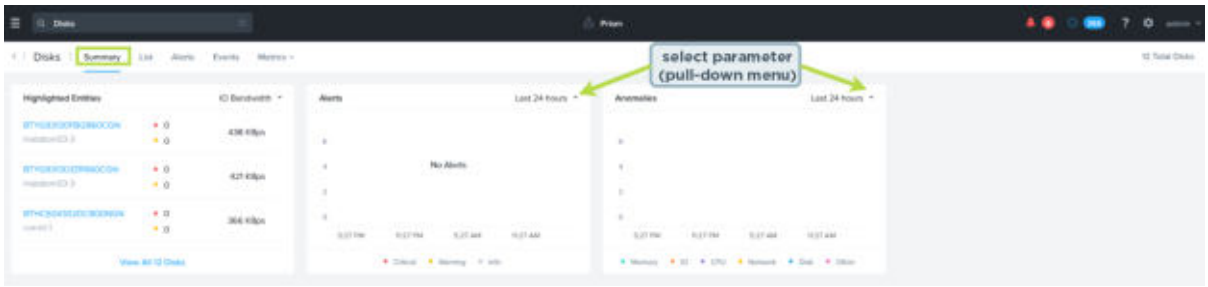

Figure 122: Disks Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the disks across the registered clusters. The following table describes the fields that appear in the disks list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

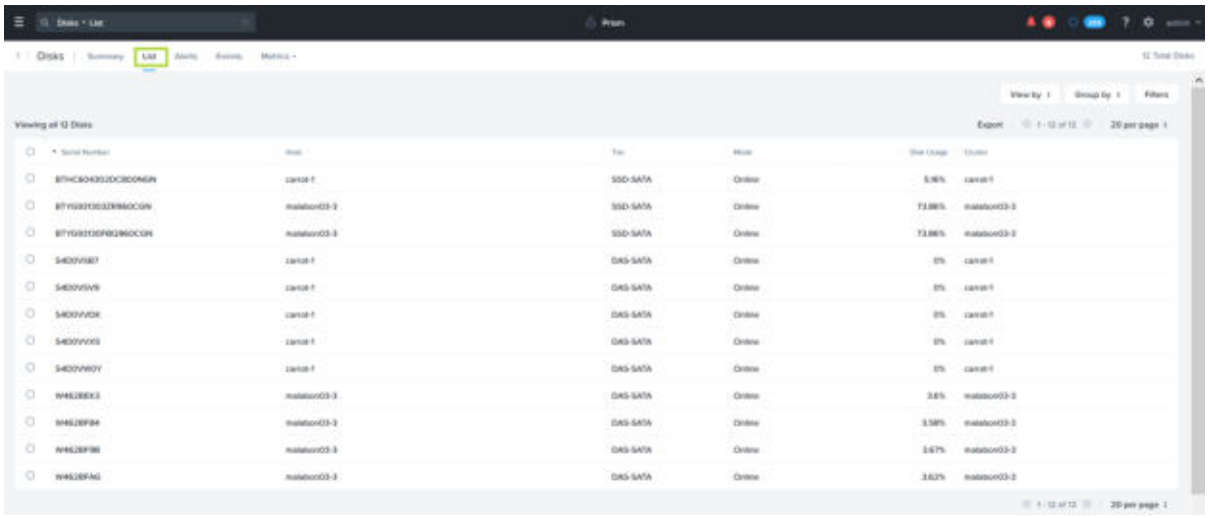

Figure 123: Disks List Tab

Table 65: Disks List Fields

Parameter	Description	Values
"General" Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Host	Displays the name of the host in which this disk resides.	(host name)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe | SSD-SATA | DAS-SATA]
Mode	Displays the operating state of the disk.	online, offline
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB|TB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
"Performance" Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB|TB]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB|TB]
IOPS	Displays I/O operations per second (IOPS) for this disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for this disk.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this disk.	xxx [ms]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)

You can filter the disks list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 66: Filter Pane Fields

Parameter	Description	Values
Serial Number	Filters on the disk serial number. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of disks that satisfy the serial number condition/string.	(serial number string)
Host	Filters on the host name. Enter a string in the field. It will return a list of disks in the hosts that satisfy the host name condition/string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of disks in the clusters that satisfy the cluster name condition/string.	(cluster name string)
Mode	Filters on whether the disk is online or offine. Check the box for one or both of these modes. The number of disks currently in each state is displayed on the right of the line.	Online, Offline
Tier	Filters on whether the disk is in the solid state (SSD-SATA) or hard disk (DAS-SATA) tier. Check the box for one or both of these modes. The number of disks currently in each tier is displayed on the right of the line.	DAS-SATA, SSD-SATA
Health	Filters on the disk health state (good, warning, or critical). Select one or more states to return a list of disks in that state(s). The number of disks currently in each state is displayed on the right of the line.	Critical, Warning, Good
Disk Usage	Filters on the used capacity. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of disks with used capacity in that range (0-100%).	([xx] to [yy]% range)
Disk Capacity	Filters on the total capacity. Enter an amount range in the "**from <low> to <high> GiB**" field. It will return a list of disks with total capacity in that range.	([xx] to [yy] GiB range)
IOPS	Filters on the IOPS. Enter a range in the "**from <low> to <high> iops**" field. It will return a list of disks with IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the I/O bandwidth used. Enter a range in the "**from <low> to <high> bps**" field. It will return a list of disks with I/O bandwidth usage in that range.	([xx] to [yy] range)

Parameter	Description	Values
IO Latency	Filters on the average I/O latency. Enter a range in the "**from <low> to <high> ms**" field. It will return a list of disks with average I/O latency in that range.	([xx] to [yy] range)

You can group the disks list in the following ways:

- The **Color** pull-down menu allows you to color code the disk entries by tier type, mode, or health state. (You can only choose one.) A legend appears at the bottom to indicate what each color means in that grouping.
- The **Group** pull-down menu allows you to group the disk entries by host, tier type, mode, cluster, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** pull-down menu allows you to group the disk entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just disk-related alerts across the registered clusters (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just disk-related events across the registered clusters (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

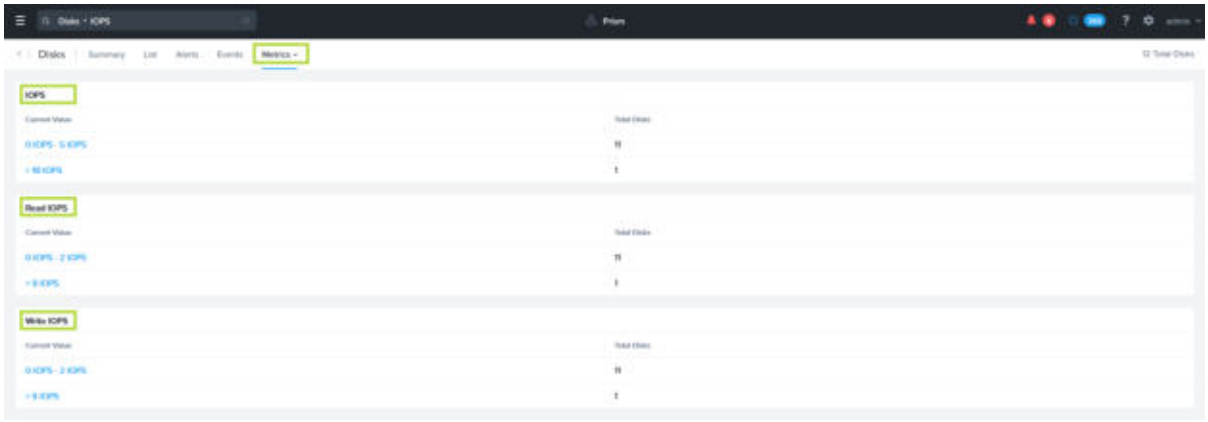

Figure 124: Disks Metrics Tab

Table 67: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS tables listing current values and total disks (number). The current values are split into intervals (for example, less than 20, 20-40, 40-60, more than 60). Clicking on an interval displays the Summary tab filtered to just those disk. Note: The same format also applies to the other metrics in this table.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Disk Details View

To access the details page for a disk, go to the disks dashboard **List** tab (see Disks Summary View on page 175) and click the disk name. You can also access the details page by clicking the disk name wherever that name appear, such as in a dashboard widget or search result.

The disk name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, and **Usage**. Click a tab to display that information on the right. (Click the **Back to Disks** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the disk (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

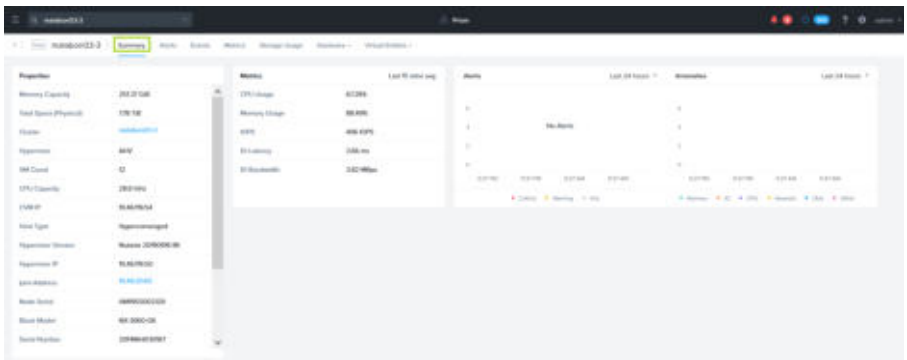

Figure 125: Disk Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 68: Disk Properties Fields

Parameter	Description	Values
Disk Usage	Displays the amount of used space on the drive.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)
Host	Displays the name of the host in which the disk resides.	(host name)
Host IP	Displays the IP address of the host.	(IP address)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe | SSD-SATA | DAS-SATA]
Mode	Displays whether the disk is currently online or offline.	[online|offline]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB|TB]
Status	Displays the operating status of the disk. Possible states include the following: • **Normal**. Disk is operating normally. • **Data migration initiated**. Data is being migrated to other disks. • **Marked for removal, data migration is in progress**. Data is being migrated in preparation to remove disk. • **Detachable**. Disk is not being used and can be removed.	Normal; Data migration initiated; Marked for removal, data migration is in progress; Detachable
Self Encryption Drive	Displays whether this is a self-encrypted drive.	Not Present, Present

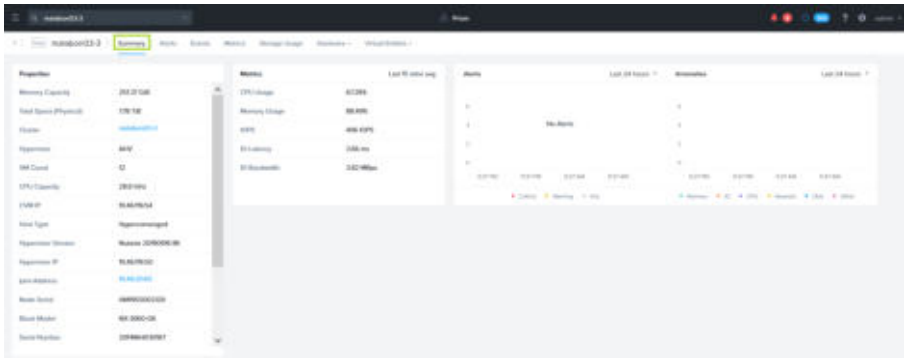

Figure 126: Disk Summary Tab

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this disk (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this disk (see Events Summary View on page 221).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the disk. Click the **Metrics** tab and then the desired metric name (**IOPS**, **IO latency**, and **IO Bandwidth**) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph.

> Note: The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 230).

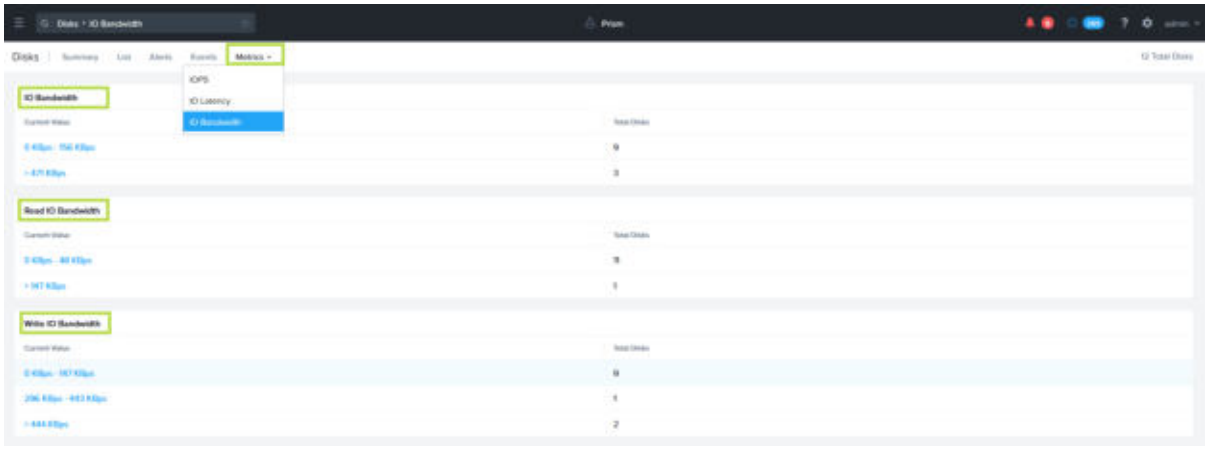

Figure 127: Disk Metrics Tab

Usage Tab

The **Usage** tab displays the following graph:

- The **Usage Summary** graph displays a rolling time interval monitor of disk storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

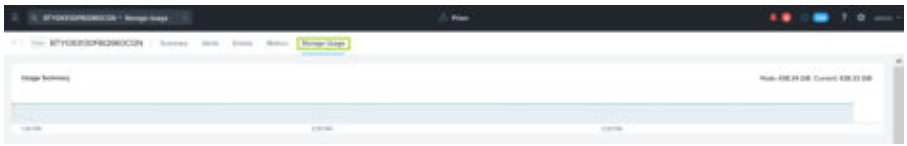

Figure 128: Disk Usage Tab

GPUs Summary View

To access the GPUs dashboard, select **Hardware** > **GPUs** from the entities menu (see Entities Menu on page 13). The GPUs dashboard allows you to view summary information about GPUs across the registered clusters and access detailed information about each GPU.

> Note: This section describes the information and options that appear in the GPUs dahsboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

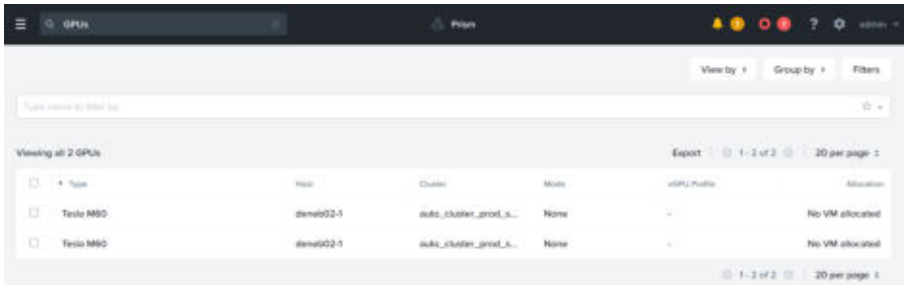

Figure 129: GPUs Dashboard

The following table describes the fields that appear in the GPUs list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 69: GPU List Fields

Parameter	Description	Values
Type	Displays the GPU model type.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster	Displays the name of the cluster in which the GPU resides.	(cluster name)
Mode	Displays the mode in which the GPU is operating.	None, vGPU, passthrough
Allocation	Indicates the number of VMs allocated to the GPU.	"No VM allocated", "x of y VMs allocated"

You can filter the GPUs list based on a several parameter values. The following table describes the filter options available when you open the GPUs view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 70: Filter Pane Fields

Parameter	Description	Values
Type	Filters on the GPU model type. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of GPUs that satisfy the type condition/string.	Tesla M10, Tesla M60, Tesla M60 compute
Host	Filters on the host name. Enter a string in the field. It will return a list of GPUs in the selected hosts.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of GPUs in the selected clusters.	(cluster name string)
Mode	Filters on the GPU operation mode. Check the box for one or more of these modes. The number of GPUs currently in each mode is displayed on the right of the line.	None, vGPU, Passthrough

You can group the GPUs list in the following ways:

- The **Group** pull-down menu allows you to group the GPU entries by cluster, host, or mode. (You can only choose one.)
- The **Sort** pull-down menu allows you to group the GPU entries by type, cluster, mode, or allocation. (You can only choose one parameter.)

GPU Details View

To access the details page for a GPU, click on the desired GPU type entry in the list (see GPUs Summary View on page 183). You can also access the details page by clicking the GPU name wherever that name appear, such as in a dashboard widget or search result. The GPU details page includes **Summary** and **VMs** tabs.

Summary Tab

Clicking the **Summary** tab, which appears by default, displays the following:

- A section on the left that displays summary information about the GPU (see following table).
- A section of the right that displays GPU performance metrics. The graphs are rolling time interval performance monitors that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. This section includes the following graphs:
  - **GPU Usage**: Displays the percentage of GPU capacity being used.
  - **GPU Framebuffer Usage**: Displays the percentage of GPU framebuffer (RAM) capacity being used.
- Action button on the upper right (which appears regardless of which tab is selected):
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

The following table describes the GPU summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 71: GPU Summary Fields

Parameter	Description	Values
GPU Type	Displays the GPU type for this entry.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster Name	Displays the name of the cluster in which the GPU resides.	(cluster name)
Host	Displays the name of the host in which the GPU resides.	(host name)
Mode	Displays the GPU operational mode. If it is vGPU, an additional field appears that displays the vGPU profile used.	None, vGPU, Passthrough
Allocation	Displays the number of VMs allocated to this GPU. This field does not appear if no VMs are allocated.	"x of y VMs allocated"
Framebuffer (RAM) Per GPU	Displays the framebuffer (RAM) size per GPU.	xx GiB
ID	Displays the GPU ID number.	(ID number)

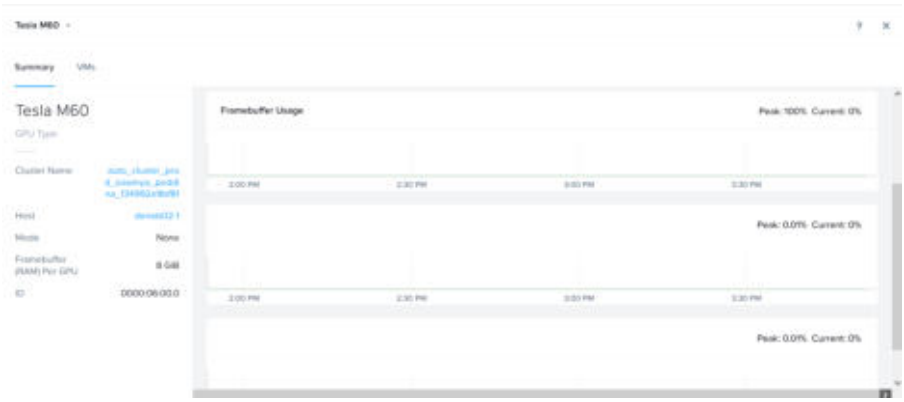

Figure 130: GPU Summary Tab

VMs Tab

Clicking the **VMs** tab displays a table of VMs allocated (attached) to the GPU. The table includes the following fields:

- **Name**: Displays the VM name. Click the name to display the details page for that VM (see VM Details View on page 90).
- **GPU Usage**: Displays the percentage of GPU capacity used by this VM.
- **GPU Framebuffer Usage**: Displays the percentage of GPU framebuffer (RAM) capacity used by this VM.

Activity Entities

You can access dashboards for the following activity monitors from the **Activity** category of the entities menu (see Entities Menu on page 13):

- Alerts (see Alerts Summary View on page 210)
- Events (see Events Summary View on page 221)
- Audits (see Audits Summary View on page 186)
- Tasks (see Tasks View on page 190)

Audits Summary View

To access the audits dashboard, select **Activity** > **Audits** from the entities menu (see Entities Menu on page 13). The audits dashboard allows you to view a list of actions performed across the registered clusters.

Note:

- This section describes the information and options that appear in the audits dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- Audit information appears only for those registered clusters running AOS 5.10 or later.
- Audit logs with default values are generated when updates to VMs are initiated, either by Prism Central **Self Service** users or by using Nutanix v3 API calls for the first time.

- The retention period for audit entries is four weeks by default.
- See Prism Central Logs on page 247 for all Prism Central related events. Additionally, you can refer to the Audit Log Events on page 250 for the complete list of captured audit events.

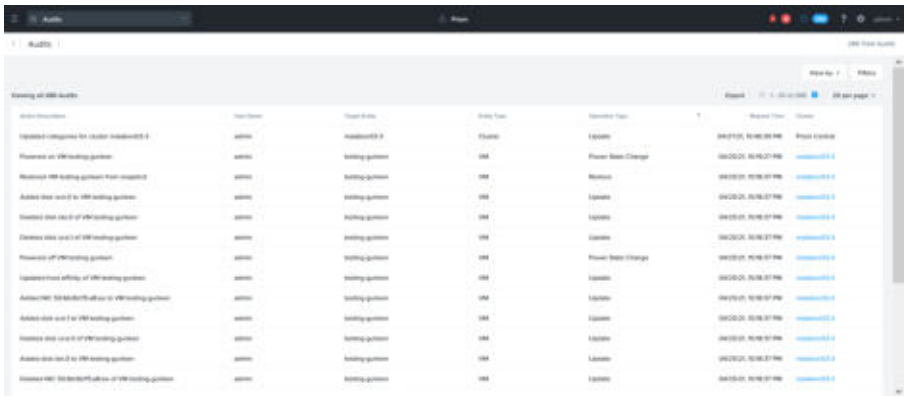

Figure 131: Audits Dashboard

The following table describes the fields that appear in the audits list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 72: Recovery Plans Fields

Parameter	Description	Values
Action Description	Describes the action taken such as "deleted VM *vm-name*" or "added disk *disk-name*"	(action description)
User Name	Displays the name of the user who requested the action.	(user name)
Target Entity	Displays the entity name. Click the name to go to the details page for that entity.	(entity name)
Entity Type	Displays the entity type such as VM or host.	(entity type)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)

To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. The following table describes the filter options available. You can apply multiple filters.

Table 73: Filter Pane Fields

Parameter	Description	Values
User Name	Enter a name string in the field to filter for users who requested an action.	(user name)
Entity Type	Check the boxes of one or more entities to filter for actions on those entity types.	VM, Storage Container, Catalog Item, Image, Cluster, Host, Disk, GPU, Security Policy, NGT Policy, Project, Role, User, Category, Availability Zone, Protection Policy, Recovery Plan, Recoverable Entity, Report
Operation Type	Check the boxes of one or more operations to filter on those operations.	Create, Update, Delete, Power State Change
Request Time	Check an interval box to filter for actions that were requested during that time period. For the custom interval option (**from** *xxx* **to** *xxx*), click in each field and select a date from the pop-up calendar.	Last 1 hour, Last 24 hours, Last week, From *xxx* to *xxx*
Cluster	Enter a cluster name in the field to filter for actions in the cluster.	(cluster name)
User IP	Enter a user IP address and then click the Add button to filter for actions requested by that user. You can add multiple user IP addresses.	(IP address)

Audit Details View

To access the details page for an action, go to the audits dashboard (see Audits Summary View on page 186) and click the action description. The details page includes the following:

- The action description (upper left). You can switch from one action details page to another by selecting from the pull-down list.
- A section on the left that displays summary information about the action (see following table).
- A section on the right that displays a table of information specific to that action. The table shows the attributes that were changed during the action and the current value that is set for the attribute. The attributes vary depending on the specifics of the action (see following examples).

The following table describes the action summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 74: Action Summary Fields

Parameter	Description	Values
Action Description	Displays the action description.	(description)
User Name	Displays the name of the user who requested the action.	(user name)
Target Entity	Displays the name of the entity that was the action target. Click the entity name to display the details for that entity.	(entity name)
Affected Entities	Displays the names of the entities that were affected by the action. Click an entity name to displays the details page for that entity.	(one or more entity names)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
User IP	Displays the IP address of the user.	(IP address)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)
Status	Displays the status of the action.	Succeeded, Failed

In the first example the action is adding more categories for the cluster, so the details simply include the names of the attached categories.

Figure 132: Audit Details Page (example 1)

In the second example the action is adding a disk, so the details include the VM name, disk address, size, and source disk UUID.

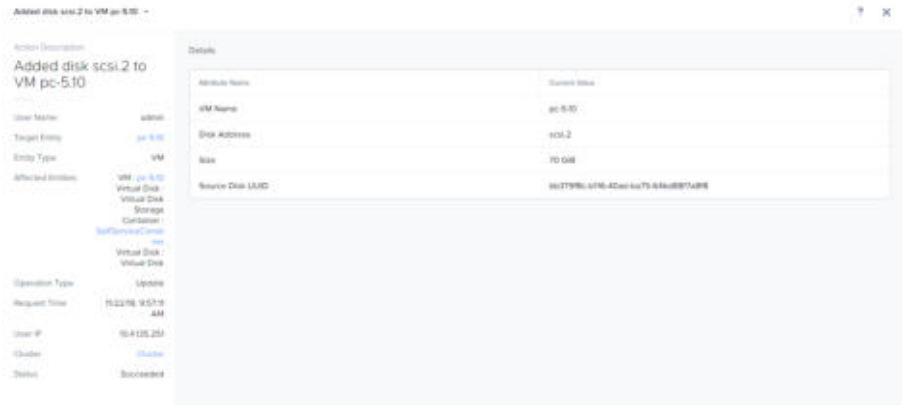

Figure 133: Audit Details Page (example 2)

In the third example the action is creating a VM, so the details include a variety of information about that VM.

Figure 134: Audit Details Page (example 3)

Tasks View

An icon appears in the main menu when there are current tasks (running or completed within the last 48 hours). The icon appears blue when a task runs normally, yellow when it generates a warning, or red when it fails. Clicking the icon displays a drop-down list of current tasks. You can see the Tasks dashboard, which displays information about all tasks across the registered clusters, by doing one of the following:

- Clicking the **View All Tasks** link at the bottom of the current tasks drop-down list.
- Clicking the **View All Tasks** link in the Tasks widget in the main dashboard (see Main Dashboard on page 60).
- Selecting **Activity** > **Tasks** in the entities menu (see Entities Menu on page 13).

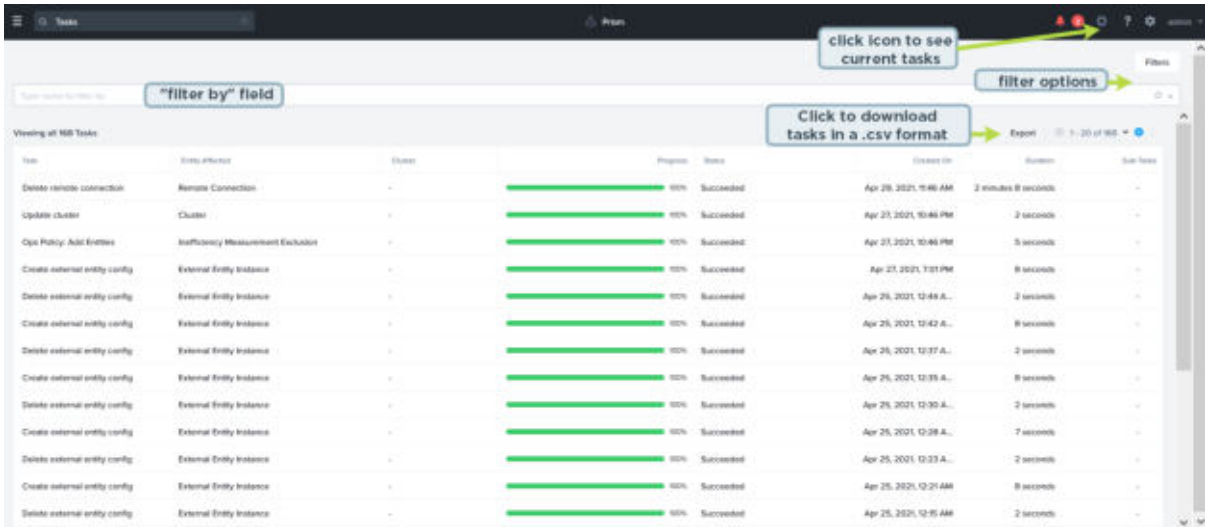

Figure 135: Tasks Screen

You can do the following in the Tasks dashboard:

- Filter the list by entering a name in the "filter by" field.
- Filter the list by clicking the **Filters** button. This displays a filter pane on the right of the screen. Click on one or more of the states to filter the list to just tasks in those states. The states are Aborted, Canceled, Failed, Queued, Running, Paused, Skipped, Succeeded, and

Scheduled. The number of tasks in each state currently appear to the right of the state name. Click the X icon (upper right) to close the filter pane.

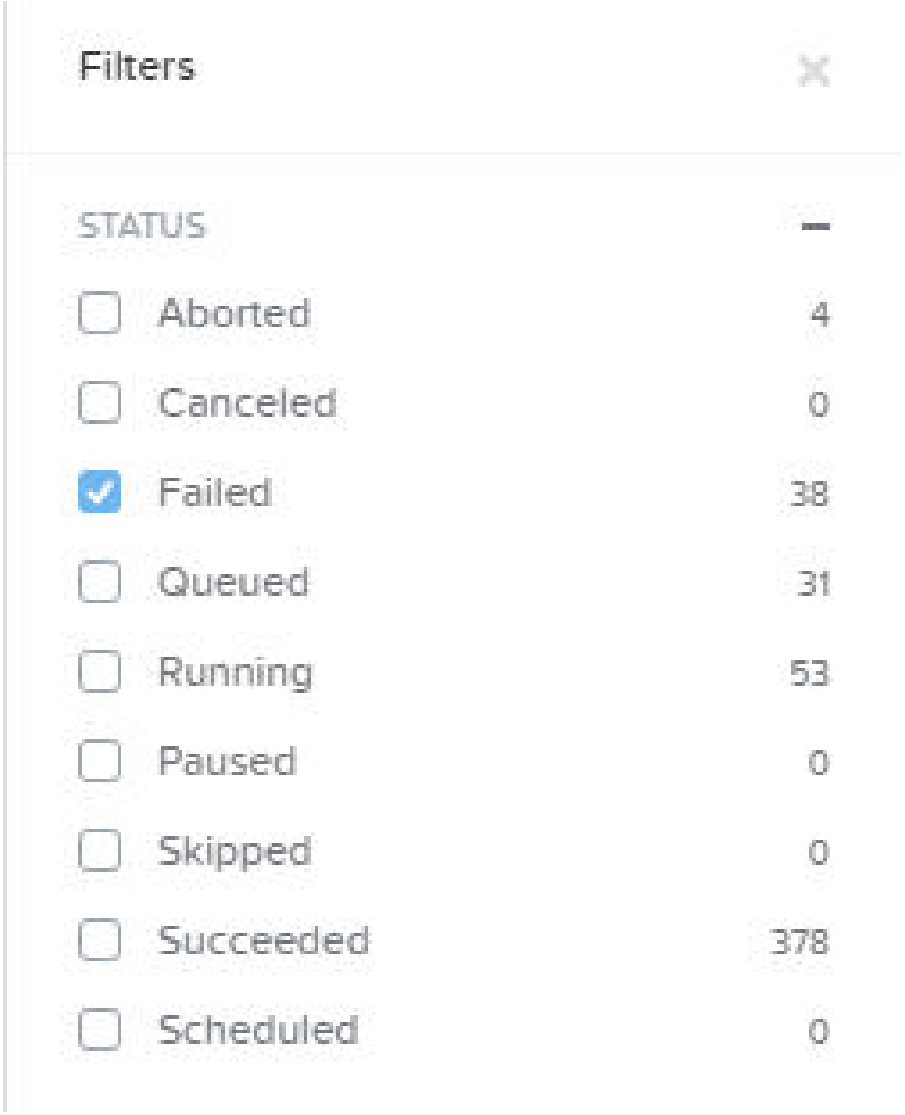

Figure 136: Filter Pane

- Download the task list in CSV format by clicking the download icon (on the right above the table).

The following table describes the fields in the tasks list. Each task appears in the list for a minimum of one hour after completion, but how long that task remains in the list depends on several factors. In general, the maximum duration is two weeks. However, tasks are rotated off the list as new tasks arrive, so a task might disappear from the list much sooner when activity is high. In some cases a task appears for longer than two weeks because the last task for each component is retained in the listing.

Table 75: Tasks List Fields

Parameter	Description	Values
Task	Specifies which type of operation the task is performing.	(name of operation such as "upgrade")
Entity Affected	Lists the entity name or describes the operation. If this is a live link, click it to either see more details or go to the entity details page.	(name or description)
Cluster	Specifies the cluster name(s) on which the task is running. A dash (-) appears when the task applies to Prism Central rather than one of the registered clusters.	(cluster names)
Progress	Indicates the current percentage complete for the task.	0%-100%
Status	Indicates the task status, which can be pending, running, completed, or failed.	Succeeded, Running,Failed, Queued
Created On	Displays when the task began.	xx [seconds| minutes| hours| days] ago
Duration	Displays how look the task has been running or took to complete.	xx [seconds| minutes| hours| days]
Sub-Tasks		

Operations Entities

You can access dashboards for the following tools from the **Operations** category of the entities menu (see Entities Menu on page 13):

- Analysis: Analyze cluster activity (see Analysis Dashboard on page 486)
- Cost Management: Analyze and manage costs (see Cost Management (Xi Beam) on page 539)
- App Discovery: Discover applications running in a cluster (see Application Discovery on page 518)
- Monitoring Configurations: Monitor select applications in a cluster (see Application Monitoring on page 540)
- Operations Policies: Manage operations policies (see Operations Policy Management on page 417)
- Planning: Plan for resource capacity and usage (see Resource Planning on page 560)
- Playbooks: Create playbooks for task automation (see Task Automation on page 585)
- Reports: Generate resource and activity reports (see Reports View on page 194)

Reports View

The Reports dashboard allows you to create and run reports about the entities and metrics of your choice. To access the Reports dashboard, select **Operations** > **Reports** from the entities menu (see Entities Menu on page 13).

Note: This section describes the information and options that appear in the Reports dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See Reports Management on page 617 for detailed information on how to create, configure, and run reports.

Figure 137: Reports Dashboard

The following table describes the fields that appear in the reports list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 76: Reports List Fields

Parameter	Description	Values
Report Name	Displays the name of the report. Note: The **SYSTEM** label appears against the report name for pre-defined reports.	(report name)
Owner	Displays the name of the user who created the report.	(user name)
Scheduled	Indicates whether this report is scheduled to run at specified times. • "Yes" means the report is scheduled to run automatically at specified times (see Frequency). • "No" means the report is not scheduled to run. To run the report you must either manually run it or update the report to specify a schedule.	Yes, No
Frequency	Displays when (how frequently) the report is scheduled to run.	Daily, Weekly, Monthly, Yearly

You can filter the reports by report name. To apply a filter, click the **Filter** button to open the **Filter** pane. Enter a string in the report name field. It will return a list of reports that satisfy the name string.

- Click the **New Report** button to create a new custom report (see Creating a New Report on page 619).
- Click the **Report Settings** button to configure the report appearance, e-mail, and retention settings (see Configuring Report Settings on page 645).

The **Actions** menu appears when one or more reports are selected. It includes the following actions (see Managing a Report on page 620):

- This action can be applied to multiple reports: **Delete**
- These actions can be applied to only one report at a time: **Run**, **Edit**, **Clone**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the selected report.

Administration Entities

You can access dashboards to manage the following objects from the **Administration** category of the entities menu (see Entities Menu on page 13):

- Users (see Users Summary View on page 195)
- Roles (see Roles Summary View on page 199)
- Projects (see Projects Summary View on page 203)
- Availability Zones (see Availability Zones View on page 208)
- LCM (see LCM View on page 209)

Users Summary View

To access the users dashboard, select **Administration** > **Users** from the entities menu (see Entities Menu on page 13). The users dashboard allows you to view a list of local and Active Directory-specified users.

> Note: This section describes the information and options that appear in the users dashboard.
>
> - See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - See “” () for information about creating and maintaining local user accounts.

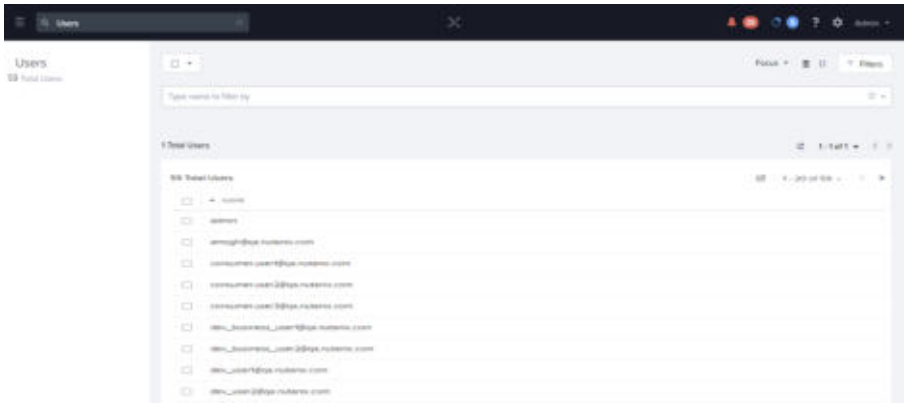

Figure 138: Users Dashboard

The Users view displays a list of local users and project-assigned Active Directory users. You can order the list alphabetically (A-Z or Z-A) by clicking (toggling) the **Name** column header. Clicking on a user name displays the details page for that user.

From the**Filter** pane you can filter the users list by name. Check the name box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field. It will return a list of users that satisfy the name condition/string.

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by name.

User Details View

To access the details page for a user, go to the users dashboard (see Users Summary View on page 195) and click the user name. The details page includes the following:

- The user name (upper left). You can switch from one user details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Trends**, **VM**, **Project**) to display that tab contents below.

Summary Tab

Clicking the **Summary** tab, which appears by default when you first open the page, displays the following:

- A section on the left that displays summary information about the user (see following table).
- A section of the right that displays vCPU, storage, and memory usage graphs for the top five VMs for this user. Place the cursor in a metric graph to see which VM is using that portion of the resource (vCPU, storage, or memory).

The following table describes the project summary information fields. A dash (-) is displayed in a field when a value is not available.

Table 77: User Summary Fields

Parameter	Description	Values
Display Name	Displays the name that appears in the Prism Central displays. This might be a truncated version of the name (next parameter) when that name is long.	(user display name)
Name	Displays the user name.	(user name)
Projects Count	Displays the number of projects to which this user is a member.	(number)
Total VMs	Displays the number of VMs assigned to this user.	(user name)
vCPU Usage	Displays the number of vCPUs used by this user.	(number)
Memory Usage	Displays the amount of memory used by this user.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this user.	xxx [GiB]

Figure 139: User Summary Tab

Trends Tab

The **Trends** tab displays the following graphs.

- The Total VMs graph displays a rolling time interval monitor of VMs owned by this user. Placing the cursor anywhere on the horizontal axis displays the number at that time.
- The **vCPU Usage** graph displays a rolling time interval monitor of user vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of user memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of user storage usage.

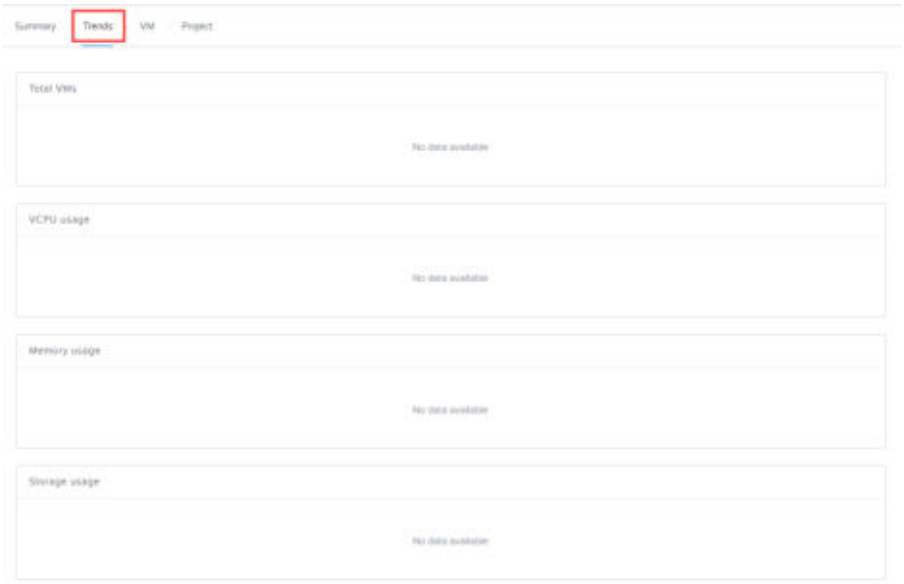

Figure 140: User Trends Tab

VM Tab

The **VM** tab displays a table of VMs owned by the user. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the VM information fields.

Table 78: User VM Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
Project Name	Displays the name of the project associated with this VM.	(project name)
Virtual CPU Count	Displays the number of vCPUs allocated to the VM.	(number)
Memory Capacity	Displays the amount of memory capacity allocated to the VM.	xxx [GiB]
Disk Capacity	Displays the amount of disk capacity allocated to the VM.	xxx [GiB]

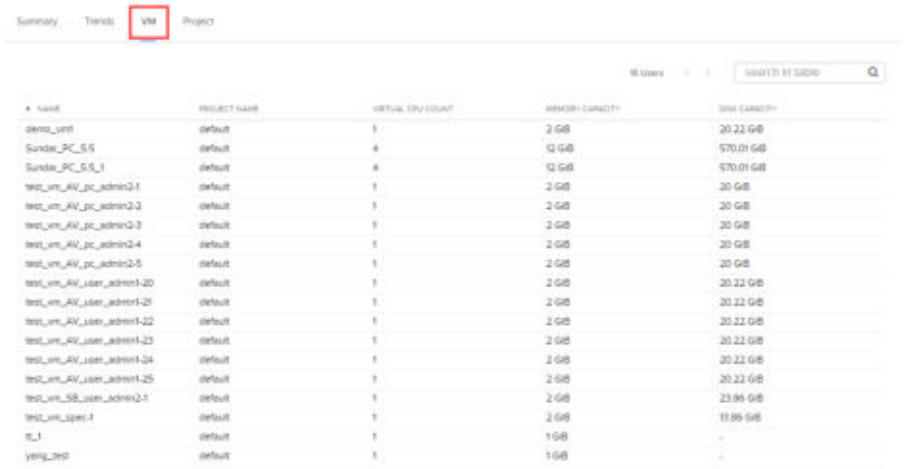

Figure 141: User VM Tab

Project Tab

The **Project** tab displays a table of projects to which the user is a member. You can filter the list by entering a string in the search field (upper right above the table). The following table describes the users information fields.

Table 79: User Project Fields

Parameter	Description	Values
Name	Displays the name of the project.	(project name)
vCPU Usage	Displays the number of vCPUs used by this project.	(number)
Memory Usage	Displays the amount of memory used by this project.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this project.	xxx [GiB]

Figure 142: User Project Tab

Roles Summary View

To access the roles dashboard, select **Administration** > **Roles** from the entities menu (see Entities Menu on page 13). The roles dashboard allows you to view summary information about built-in (default) and custom (created) roles and to access detailed information about each role.

Note: This section describes the information and options that appear in the roles dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize this information in a variety of ways.
- See “” () for information on how to create, manage, and apply roles.

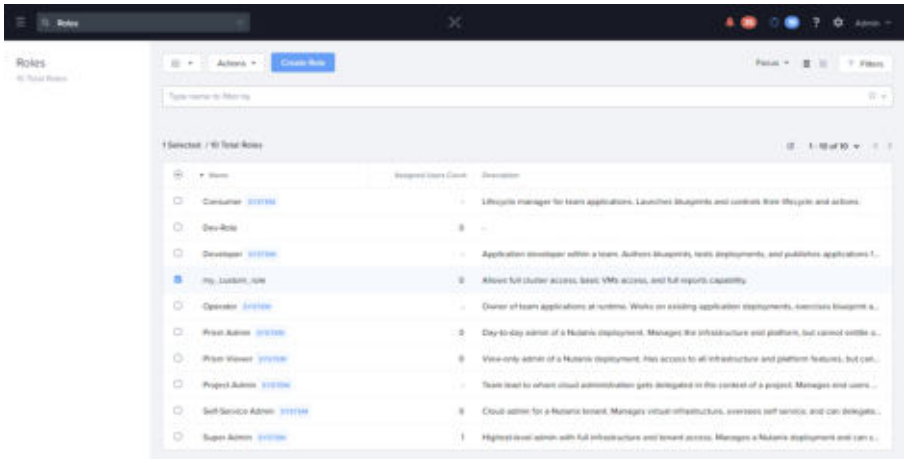

Figure 143: Roles Dashboard

The following table describes the fields that appear in the Roles list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 80: Roles List Fields

Parameter	Description	Values
Name	Displays the role name.	(role name)
Assigned Users Count	Displays the number of users assigned this role.	(number)
Description	Describes the role (if a description was provided when the role was created or updated).	(number)

You can filter the roles list based on several parameter values. The following table describes the filter options available when you open the roles dashboard **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 81: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the role name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of role that satisfy the name condition/ string.	(project name string)
Assigned Users Count	Filters on the number of assigned users. Enter a range in the "**from <low> to <high>**" field. It will return a list of roles that have an assigned user count within that range.	([xx] to [yy] range)

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by name or assigned user count. (You can only choose one parameter.)

The dashboard includes a **Create Role** button (see “” ()). The **Actions** menu appears when one or more roles are selected. It allows you to manage assignments for, duplicate, update, or delete a role (see Role Details View on page 201).

> Note: You cannot delete or modify the default roles.

Role Details View

To access the details page for a role, go to the roles dashboard (see Roles Summary View on page 199) and click the role name. The details page includes the following:

- The role name (upper left). You can switch from one role details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Users**, **User Groups**, **Role Assignment**) to display that tab contents below.
- Action buttons (upper right). The available actions appear in bold; other actions are grayed out.
  - **Manage Assignment**: Click this button to assign the role to users, groups, and entities (see “” ()). This button does not appear for certain default roles such as Prism Admin that you are not allowed to assign.
  - **Duplicate**: Click this button to duplicate the role. It opens the role configuration screen preconfigured with the same permissions as this role (see “” ()).
  - **Delete**: Click the button to delete the role. You cannot delete or modify the default roles.
  - **Update Role**: Click this button to modify the role permissions (see “” ()).
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

Summary Tab

Clicking the **Summary** tab, which appears by default when you first open the page, displays the following:

- A section on the left that displays summary information about the role (see following table).

- A section in the middle which displays a list of actions that the role is permitted to execute. Click the **Actions** header to display additional actions for those roles with more than a single page of permitted actions. You can filter the actions list by entering a string in the search field (upper right above the list).

The following table describes the role summary information fields. A dash (-) is displayed in a field when a value is not available.

Table 82: Role Summary Fields

Parameter	Description	Values
Name	Displays the role name.	(role name)
Assigned Users	Displays the number of users assigned this role.	(number)
Assigned User Groups	Displays the number of user groups assigned this role.	(number)

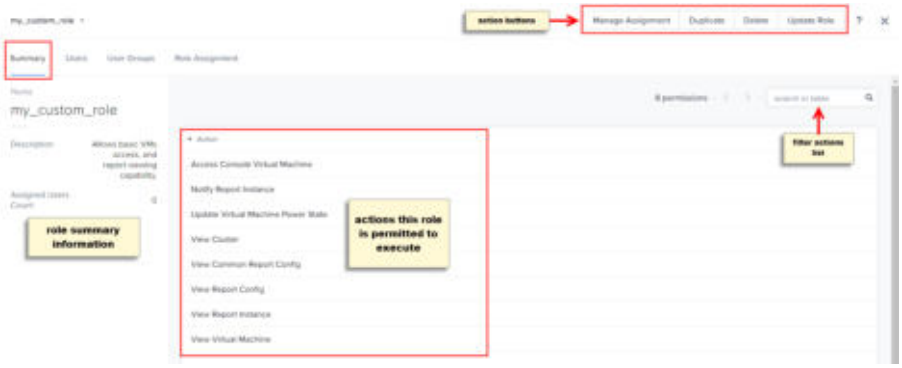

Figure 144: Role Summary Tab

Users Tab

The **Users** tab displays a list of users assigned to the role. You can filter the list by entering a string in the search field (upper right above the table).

User Groups Tab

The **User Groups** tab displays a list of user groups assigned to the role. You can filter the list by entering a string in the search field.

Roll Assignment Tab

The **Roll Assignment** tab displays the relationship between the users, groups, and entities assigned to this role. This tab does not appear for certain default roles such as Prism Admin that you are not allowed to assign.

Figure 145: Role Assignment Tab

Projects Summary View

To access the projects dashboard, select **Administration** > **Projects** from the entities menu (see Entities Menu on page 13). The projects dashboard allows you to view summary information about projects that have been created and access detailed information about each project.

Note: This section describes the information and options that appear in the projects view.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See Project Management on page 462 for information on how to create and manage projects.

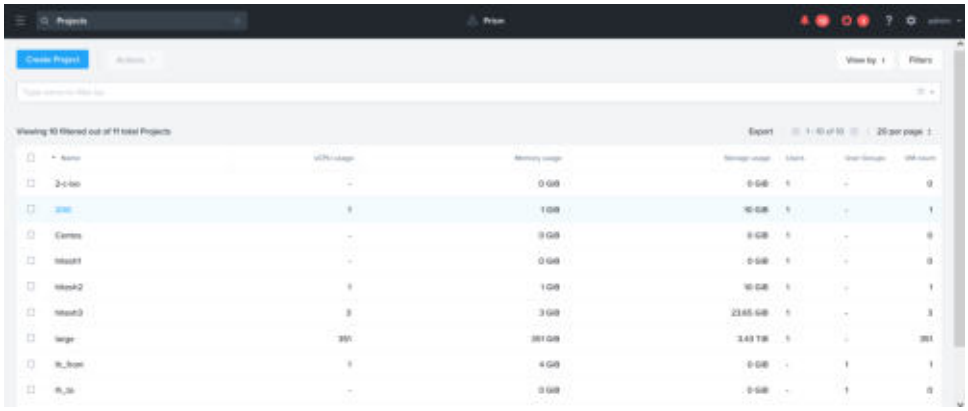

Figure 146: Projects Dashboard

The following table describes the fields that appear in the Projects list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 83: Projects List Fields

Parameter	Description	Values
Name	Displays the project name.	(project name)
vCPU Usage	Displays the number of vCPUs used by the project.	(number)
Memory Usage	Displays the amount of memory used by the project.	xxx [GiB]

Parameter	Description	Values
Storage Usage	Displays the amount of storage used by the project.	xxx [GiB]
Users	Displays the number of users in the project.	(number)
User Groups	Displays the number of user groups in the project.	(number)
VM Count	Displays the number of VMs in the project.	(number)

You can filter the projects list based on several parameter values. The following table describes the filter options available when you open the Projects view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 84: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the project name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of projects that satisfy the name condition/string.	(project name string)
vCPU Usage	Filters on the number of vCPUs. Enter a range in the "**from <low> to <high>**" field. It will return a list of projects within that range of vCPU usage.	([xx] to [yy] range)
Memory Usage	Filters on the amount of memory used. Enter a range in the "**from <low> to <high> GiB**" field. It will return a list of projects within that range of memory usage.	([xx] to [yy] GiB range)
Storage Usage	Filters on the amount of storage space used. Enter a range in the "**from <low> to <high> GiB**" field. It will return a list of projects within that range of storage usage.	([xx] to [yy] GiB range)
VM Count	Filters on the number of VMs. Enter a range in the "**from <low> to <high>**" field. It will return a list of projects within that range of total number of VMs.	([xx] to [yy] range)

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by name, vCPU usage, memory usage, storage usage, or VM count. (You can only choose one parameter.) This is the only grouping option available in the Projects view.

The Projects view includes a **Create Project** button (see Creating a Project on page 462). The **Actions** menu appears when one or more projects are selected. It allows you to update or delete the project (see Modifying a Project on page 465). You can update only one project at a time, but you can select multiple projects when deleting.

Note: You cannot delete the default project.

Project Details View

To access the details page for a project, go to the projects dashboard (see Projects Summary View on page 203) and click the project name. The details page includes the following:

- The project name (upper left). You can switch from one project details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Usage**, **VMs**, **Users**) to display that tab contents below.
- Action buttons (upper right). The available actions appear in bold; other actions are grayed out.
  - **Delete**: Click the button to delete the project. You cannot delete or modify the default project.
  - **Update Project**: Click this button to modify the role permissions (see *Modifying a Custom Role* in the Security Guide).
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

Summary Tab

Clicking the **Summary** tab, which appears by default when you first open the page, displays the following:

- A section on the left that displays summary information about the project (see following table).
- A section of the right that displays vCPU, storage, and memory usage graphs for the top five users and top five VMs in the project. Place the cursor in a metric graph to see which user or VM is using that portion of the resource (vCPU, storage, or memory).

The following table describes the project summary information fields. A dash (-) is displayed in any table field when a value is not available or applicable.

Table 85: Project Summary Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
VM Count	Displays the number of VMs in the project.	(number)
Users	Displays the number of users in the project.	(number)
User Groups	Displays the number of user groups in the project.	(number)
Networks	Displays the networks that the project is allowed to access.	(network names)
vCPU Usage	Displays the number of vCPUs used by the project.	(number)

Parameter	Description	Values
Memory Usage	Displays the amount of memory used by the project.	xxx [GiB]
Storage Usage	Displays the amount of storage used by the project.	xxx [GiB]
Description	Displays a project description, which is typically who created the project.	(description)

Figure 147: Project Summary Tab

Usage Tab

The **Usage** tab displays the following graphs. Set the graph time period by selecting the duration (1 day or 1 week) from the **Show** drop-down menu.

- The **vCPU Usage** graph displays a rolling time interval monitor of project vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of project memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of project storage usage.

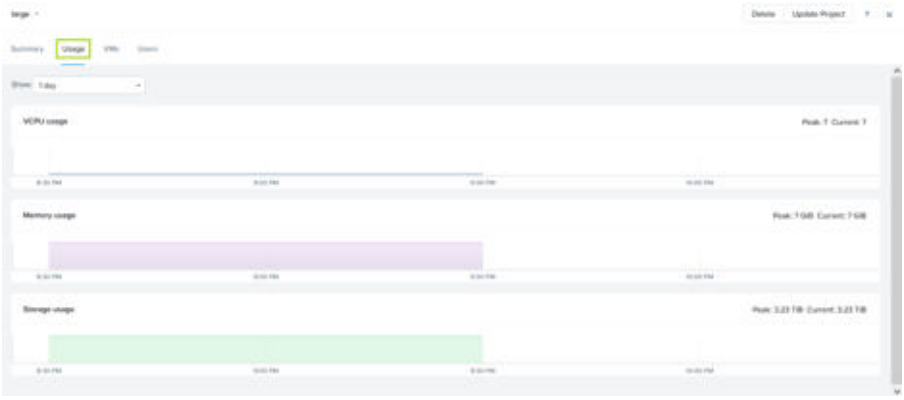

Figure 148: Project Usage Tab

VMs Tab

The **VMs** tab displays a table of VMs assigned to the project. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the VMs information fields.

Table 86: VMs Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
Power State	Displays whether the VM is currently on or off.	On (green circle), Off (red circle)
Owner	Displays the name of the user who owns this VM.	(user name)
Virtual CPU Count	Displays the number of vCPUs allocated to the VM.	(number)
Memory Capacity	Displays the amount of memory allocated to the VM.	xxx [GiB]
Disk Capacity	Displays the amount of disk space allocated to the VM.	xxx [GiB]

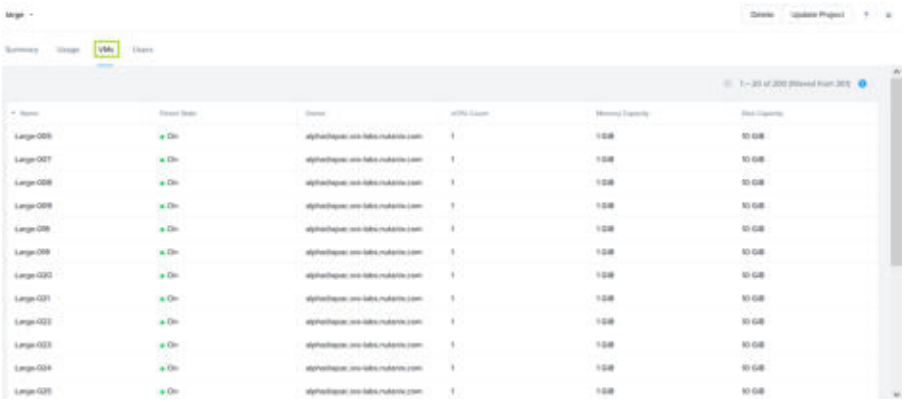

Figure 149: Project VMs Tab

Users Tab

The **Users** tab displays a table of users (individuals and group members) assigned to the project. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the users information fields.

Table 87: Users Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)

Parameter	Description	Values
Total VMs	Displays the number of VMs owned by this user.	(user name)
vCPU Usage	Displays the number of vCPUs used by this user.	(number)
Memory Usage	Displays the amount of memory used by this user.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this user.	xxx [GiB]

Figure 150: Project Users Tab

Availability Zones View

To access the availability zones dashboard, select **Administration** > **Availability Zones** from the entities menu (see Entities Menu on page 13). The availability zones dashboard allows you to view information about the availability zones connected to Prism Central currently and to connect to other zones.

Note: This section describes the information and options that appear in the availability zones dashboard.

- See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- See the *Xi Leap Administration Guide* for information about availability zones and how to connect to them.

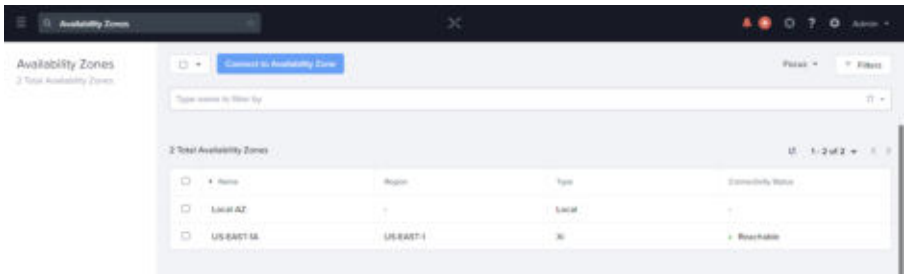

Figure 151: Availability Zones Dashboard

The following table describes the fields that appear in the availability zones list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 88: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the zone name.	(zone name)
Region	Displays the region in which the zone is located.	(region name)
Type	Displays the type or zone. The availability zone that you are logged in to is shown as a local availability zone. Availability zones that are backed by on-premises Prism Central instances are shown as type physical. Availability zones in Xi Cloud Services are shown as type Xi.	Local, Physical, Xi
Connectivity Status	Indicates whether Prism Central can access the zone currently.	Reachable, Unreachable

To filter the list based on the type of zone (local, physical, or Xi), click the **Filters** button and select the desired type(s) in the **Filters** pane.

Click the **Connect to Availability Zone** button to connect to a new availability zone. The **Actions** menu appears when one or more zones are selected. It allows you to disconnect from the selected zones.

LCM View

To access the LCM dashboard, select **Administration** > **LCM** from the entities menu (see Entities Menu on page 13). The life cycle manager (LCM) tracks software and firmware versions of the various components in a cluster. It allows you to view information about the current inventory and update the versions as needed. See the Life Cycle Manager Guide for more information.

Services Entities

You can enable the following services from the **Services** category of the entities menu (see Entities Menu on page 13):

- Calm (see Enabling Calm on page 651)
- Files (see Enabling Files on page 652)
- Foundation Central (see Enabling Foundation Central on page 652)
- Karbon (see Enabling Karbon on page 652)
- Objects (see Enabling Objects on page 653)

5

ALERT AND EVENT MONITORING

Prism Central provides several mechanisms to monitor and configure cluster alerts and events.

- The Alerts and Events dashboards, which you can access from the entities menu (see Entities Menu on page 13), allow you to view and monitor alert and event activity across the registered clusters.
  - The Alerts Summary View on page 210 displays a list of alerts that you filter in various ways. You can also drill down for detailed information and corrective actions for an alert (see Alert Details on page 218)
  - The Events Summary View on page 221 displays a list of events that you filter in various ways. You can also drill down for detailed information for an event (see Event Details on page 222).
- You can configure rules for who should receive email alerts and customize the alert messages (see Configuring Alert Emails on page 223).
- You can customize the list of events that generate an alert (see Configuring Alert Policies on page 226).
- You can review event log files when necessary (see Prism Central Logs on page 247).

> Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default. (Some DR events are retained for just 7 days by default.)

Alerts Summary View

The Alerts dashboard allows you to view summary information about alert messages across the registered clusters, access detailed information about each alert, and view alert policies from any source (user defined, system defined, or external defined).

To access the Alerts dashboard, select **Activities** > **Alerts** from the entities menu (see Entities Menu on page 13). The top pane includes **List** and **Alert Policies** tabs with **User Defined**, **System Defined**, and **External Defined** tabs under **Alert Policies**. Click the tab to display that content below the tab.

> Note: This section describes the information and options that appear in the Alerts dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways. See Alerts/Health checks on page 252 for a list of alert messages and corrective actions.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of active alerts across the registered clusters. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

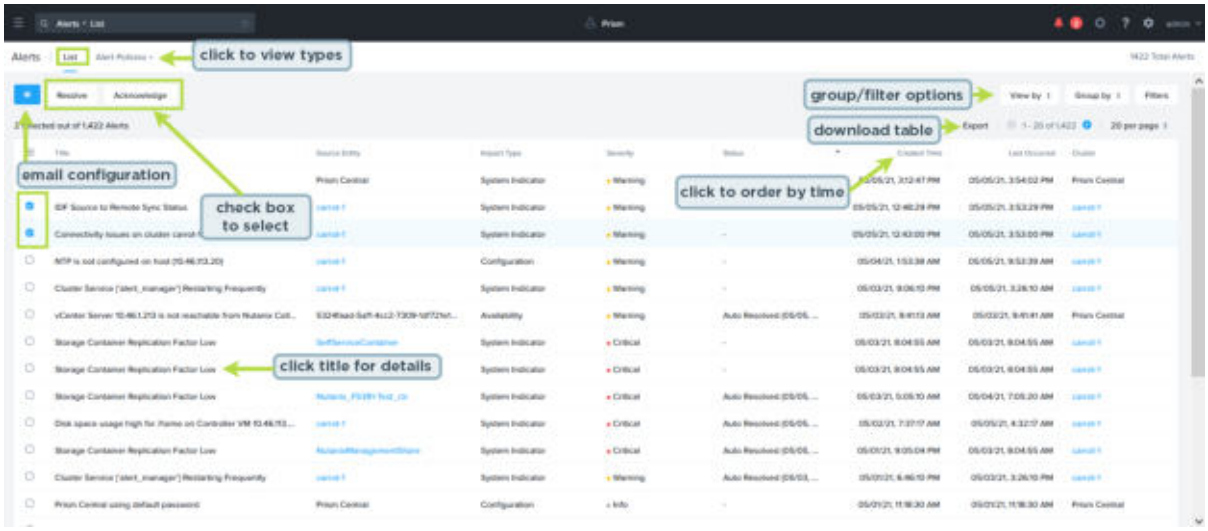

Figure 152: Alerts List Tab

- To order the alerts, click a column header, which toggles the alerts alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of alerts in CSV format, click **Export**. The maximum number of alerts you can export is 1000.
- To display the details page for an alert, click the alert title (see Alert Details on page 218).
- To enable alert emails and specify email addresses to which alerts should be sent, click the **Email Configuration** button (see Configuring Alert Emails on page 223).
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 77).
- To group the alerts, select an option from the **Group** pull-down menu. You can group the alerts by cluster, severity, or impact type. (You can only choose one.)
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Cluster: Enter name in search field.
  - Provider: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Severity: Critical, Warning, Info
  - Resolved: Yes (Manual), Yes (Auto), No
  - Acknowledged: Yes, No
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Create Time: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*
  - Last Occurred: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*
  - Title: Enter a string in the field.

- To select (or unselect) alerts in the list, do one of the following:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all** (or **Clear selections**) to select (unselect) all the alerts in the list.
- When one or more alerts are selected, **Resolve** and **Acknowledge** appear in the menu; click one or both to resolve and/or acknowledge the selected alerts. The user name and time appears in the appropriate field when an alert is acknowledged or resolved.

Table 89: Alerts List Fields

Parameter	Description	Values
(selection box)	Click this box to select the message for acknowledgement or resolution.	n/a
Title	Displays the name of the alert. Click the name to displays the details page for that alert.	(alert name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the alert applies. A comma separate list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Impact Type	Displays the category in which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Severity	Displays the severity level of this condition. There are three levels: **Critical** An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. **Warning** An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. **Info (Informational)** An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational

Parameter	Description	Values
Status	Indicates whether the alert has been resolved. Resolving an error means you set that error as fixed. (The alert may return if the condition is scanned again at a future point.) If you do not want to be notified about the condition again, turn off the alert for this condition (see Configuring Alert Policies on page 226). • A blank value means the alert is not resolved. • An "Acknowledged By" value means this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, check the box for that alert and then select the **Acknowledge** button. • An "Auto Resolved (*date_time*)" value means this alert was resolved automatically at the specified date and time. This requires that auto alert is enabled for this alert type. • A "Resolved By *user* (*date_time*)" value means this alert was resolved manually by the specified user at that date and time. To manually resolve an alert, check the box for that alert and then select the**Resolve** button.	(blank), Acknowledged By *user* (*date_time*), Auto Resolved (*date_time*), Resolved By *user* (*date_time*)
Create Time	Displays the date and time when the alert occurred.	(date and time)
Last Occurred	Displays the date and time when the alert last occurred (before this occurrence). If this is the first occurrence, both Create Time and Last Occurred display the same date and time.	(date and time)
Cluster	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 153).	(cluster name)

Alert Policies: User Defined Tab

Clicking the **User Defined** tab displays a list of the custom (user defined) alert policies. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

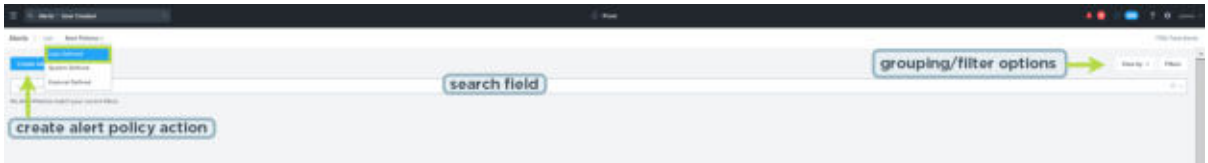

Figure 153: Alert Policies: User Defined Tab

- To create a custom policy, click the **Create Alert Policy** button (see Adding Custom Alert Policies on page 229).
- To view or modify a custom policy, click the policy name (see Modifying Custom Alert Policies on page 234).
- To delete a custom policy, select the policy and then select **Delete** from the **Actions** pull-down menu.
- To order the polices, click a column header, which toggles the list alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 77).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Auto Resolve: Enabled, Disabled
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 90: Alert User Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to displays the policy definition page (see Modifying Custom Alert Policies on page 234).	(policy name)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)

Parameter	Description	Values
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Auto Resolve	Displays whether auto resolved is enabled.	Enabled, Disabled
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Policies: System Defined Tab

Clicking the **System Defined** tab displays a list of the default (system defined) alert policies. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

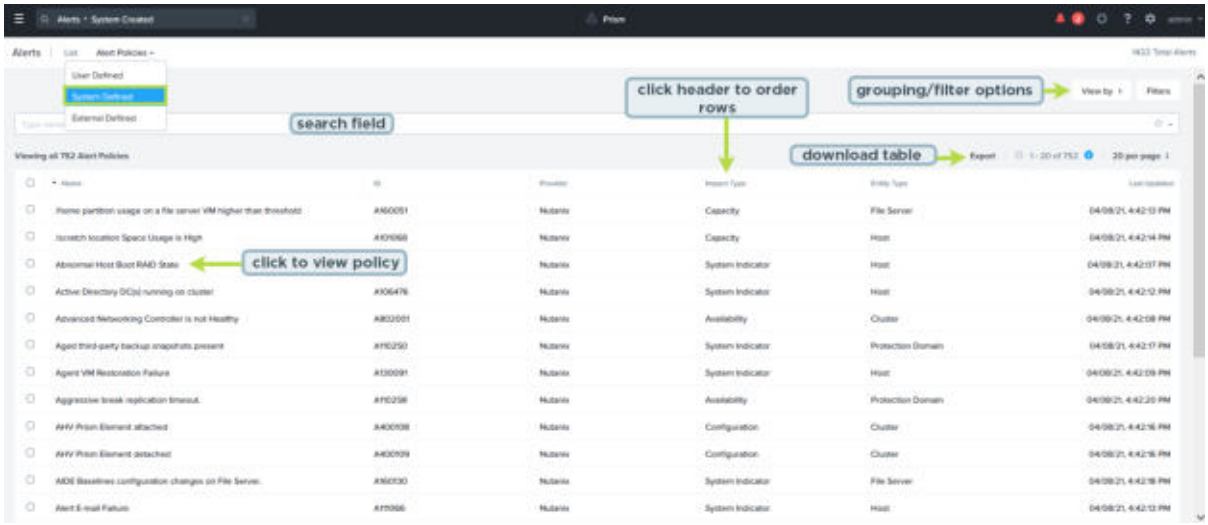

Figure 154: Alert Policies: System Defined Tab

- To view or modify a system policy, click the policy name (see Modifying System Alert Policies on page 227).
- To order the polices, click a column header, which toggles the list alphabetically, numerically, or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 77).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include

in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.

  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - ID: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator

- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 91: Alert System Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to display the Update Policy window (see Modifying System Alert Policies on page 227).	(policy name)
ID	Displays the identification number assigned to the policy.	(ID number)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Policies: External Defined Tab

Clicking the **External Defined** tab displays a list of the alert policies defined by external entities, typically an application through the API. The following figure is a sample view, and the table

describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

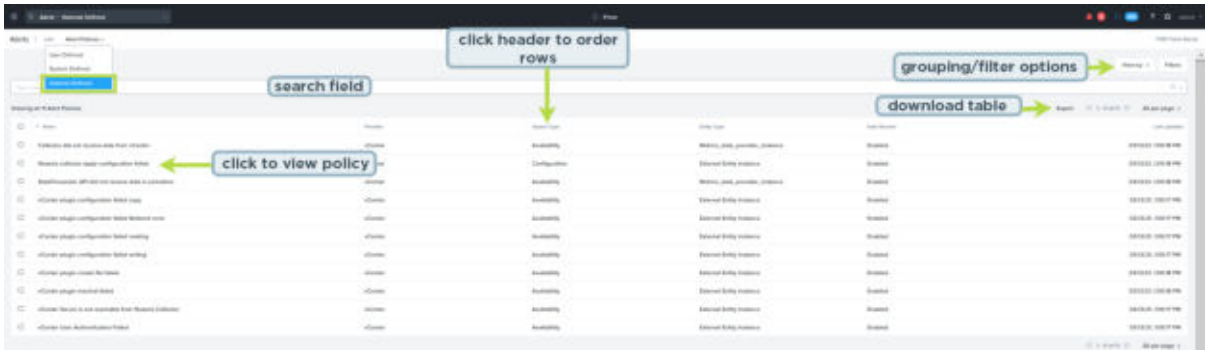

Figure 155: Alert Policies: External Defined Tab

- To view an external policy, click the policy name (see Viewing External Alert Policies on page 244).
- To order the polices, click a column header, which toggles the list alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 77).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Auto Resolve: Enabled, Disabled
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 92: Alert System Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to display the *policy_name* window (see Viewing External Alert Policies on page 244)..	(policy name)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Auto Resolve	Displays whether auto resolved is enabled.	Enabled, Disabled
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Details

Clicking on an alert message in the dashboard (see Alerts Summary View on page 210) or anywhere else the alert title appears, such as in a search list, displays detailed information about that alert. The alert details appear in the left column. Possible causes for the alert appear to the right. The most likely cause appears first with other possible causes (if any) appearing below in the order of likelihood. Each cause includes a Recommendation section that describes the recommended corrective action and in many cases a Details section that provides additional context and instructions.

- To switch among alert details pages, select the desired alert from the pull-down list in the upper left of the screen.
- To acknowledge or resolve the alert, click the **Resolve** or **Acknowledge** button in the upper right.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 93: Alert Details Fields

Parameter	Description	Values
Alert Title	Displays the title of the alert.	(alert title)

Parameter	Description	Values
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the alert. Clicking an entity link opens the detail page for that entity.	(entity names)
Severity	Displays the severity level of this condition. There are three levels: **Critical** An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. **Warning** An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. **Informational (Info)** An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational
Cluster Name	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 153).	(cluster name)
Cluster Version	Displays the AOS version running on the cluster.	(version number)
Created Time	Displays the date and time when the alert first occurred.	(time and date)
Last Occurred	Displays the date and time for the most recent occurrence of the alert.	(time and date)
Impact Type	Displays the impact type to which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Policy	Displays the name of the alert policy.	(alert policy name)
Status (for manually resolved alerts)	Displays the current alert status: acknowledged, resolved, or neither (blank).	(blank), Acknowledged, Resolved
Status (for auto resolved alerts)	Displays the current alert status: acknowledged, resolved, auto resolved, or neither (blank).	(blank), Acknowledged, Auto Resolved

Parameter	Description	Values
Acknowledged By	Indicates whether the alert has been acknowledged. • A blank (dash) value means the alert is not acknowledged. • An "N/A *date_time*)" entry indicates this alert was acknowledged automatically at the specified date and time. This requires that auto alert is enabled for this alert type (see Configuring Alert Policies on page 226). • A "*user date_time*" entry indicates this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, click the **Acknowledge** button (upper right). The field value changes from blank to the user name and date/time the alert was acknowledged.	(blank), N/A *date_time*), *user date_time*
Resolved By	Indicates whether the alert has been resolved. The options (blank, N/A, or user) are the same as for Acknowledged. To manually resolve an alert, click the **Resolve** button (upper right). The field value changes from blank to the user name and date/time the alert was resolved.	(dash), N/A *date_time*), *user date_time*
Description	Describes the reason for the alert.	(event description)

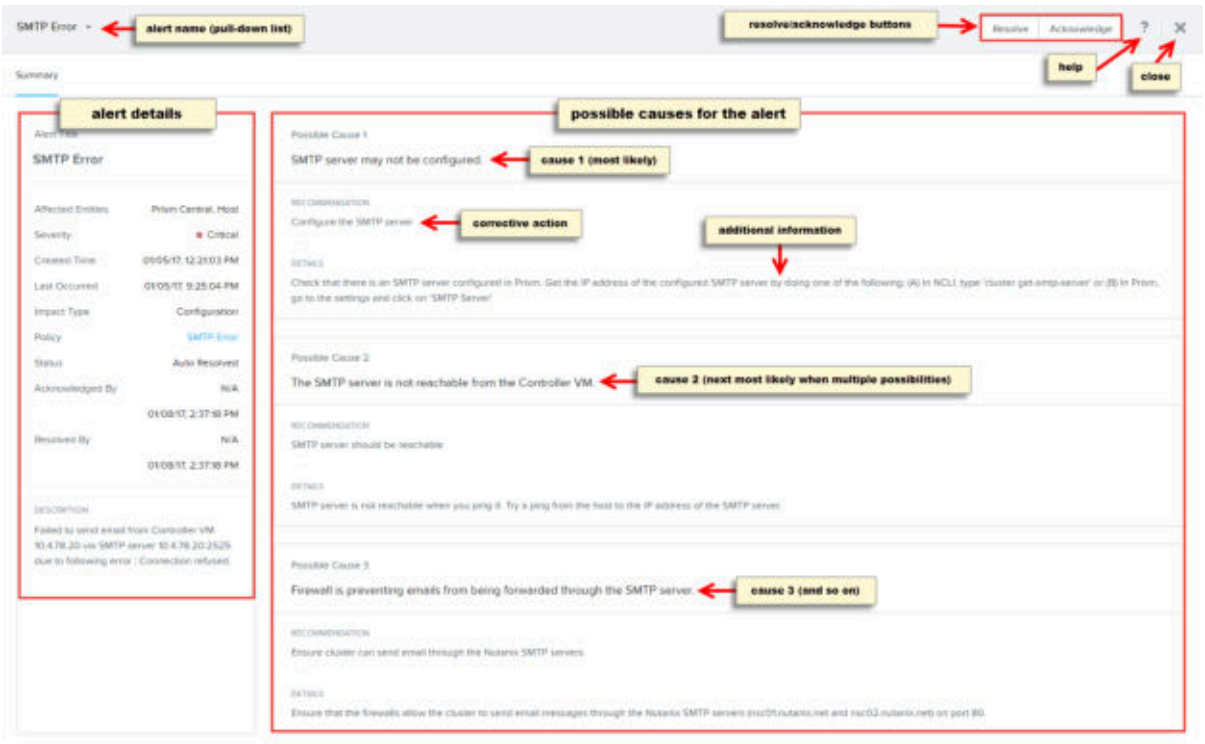

Figure 156: Alert Details Screen

Events Summary View

The Events dashboard summary view displays a list of event messages across the registered clusters. Event messages describe cluster actions such as adding a storage pool or taking a snapshot. Unlike alerts, event messages are simply informational without the need to acknowledge or resolve. The following figure is a sample view, and the table describes each field in this view. To access the Events dashboard, select **Activities** > **Events** from the entities menu (see Entities Menu on page 13).

> Note: This section describes the information and options that appear in the Events dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

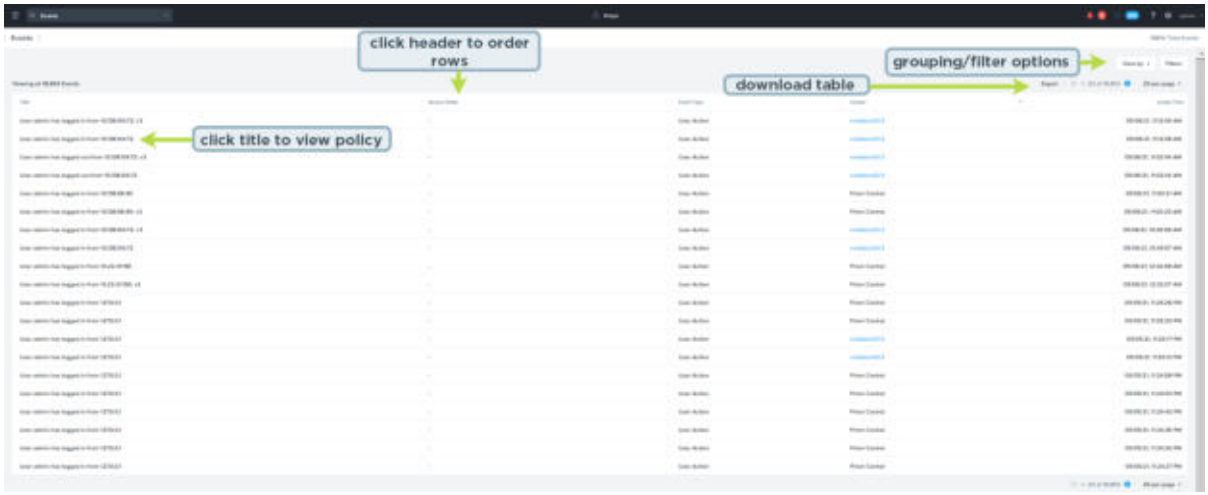

Figure 157: Events Dashboard

- To display the details page for an event, click the event title (see Event Details on page 222).
- To order the events chronologically, click the **Create Time** column header, which toggles the event list latest-to-earliest or earliest-to-latest.
- To download the table of alerts in CSV format, click the **Export** link.
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following event parameters and values.
  - Event Type: Behavioral Anomaly, System Action, User Action
  - Cluster: Enter name in search field.
  - Create Time: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*

Table 94: Events View Fields

Parameter	Description	Values
Title	Displays the name of the event. Click the name to displays the details page for that event (see Event Details on page 222).	(event name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the event applies. A comma separated list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Event Type	Displays the category in which the event is classified.	System Action, User Action, Behavioral Anomaly (see Behavioral Learning Tools on page 582), DR
Cluster	Displays the name of the cluster in which the event was issued.	(cluster name)
Create Time	Displays the date and time when the event occurred.	(date and time)

Event Details

Clicking on an event message in the dashboard (see Events Summary View on page 221) displays detailed information about that event. The event details appear in the left column. Additional information, such as a description or performance graph, appears to the right.

Figure 158: Event Details Screen

- To switch among event details pages, select the desired event from the pull-down list in the upper left of the screen.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 95: Event Details Fields

Parameter	Description	Values
Event Title	Displays the title of the event.	(alert title)

Parameter	Description	Values
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the event. Clicking an entity link opens the detail page for that entity.	(entity names)
Cluster Name	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 153).	(cluster name)
Cluster Version	Displays the AOS version running on the cluster.	(version number)
Created Time	Displays the date and time when the event occurred.	(time and date)
Event Type	Displays the type of event. Three types are defined, and each event is assigned to one of these types.	Behavioral Anomaly, System Action, User Action

Configuring Alert Emails

About this task

Prism Central allows you to configure the alert messages sent by Prism Central. To configure alert settings, reporting rules, and message templates, do the following:

Note:

- Alert emails sent by Prism Central are in addition to any alert emails you might have configured on individual clusters through the Prism Element web console. You will receive email from both entities in this case. Prism Central alert emailing is not enabled by default; you must explicitly enable it and specify the recipients (Nutanix customer support and/or supplied email addresses). If you enable alerts through Prism Central and do not want to receive double email notifications for the same alert, disable customer email notification for those alerts on the individual clusters through Prism Element (but keep email notification for Nutanix customer support enabled).
- Prism Central requires an SMTP server to send alert email messages (see Configuring an SMTP Server on page 297).

Procedure

1. Do one of the following:

   » Click the gear icon and select **Alert Email Configuration** from the **Settings** menu (see Settings Menu on page 17).

   » In the **Alerts** dashboard, click the **Configure** button and select **Email Configuration** from the pull-down list.

   The Alert Email Configuration window appears.

2. Click the **Settings** tab and do the following:

   a. Enable Prism Central to sent alert emails.

      - Check the **Every Single Alert** box to send an email whenever the event occurs.
      - Check the **Daily Digest** box to send a cumulative (24 hour) list of alerts once a day. (If you click the **Skip empty digest email** box, an email is not sent for the days when no alerts occur.)

      You can check one or both boxes. If neither box is checked, no alert emails are sent by Prism Central.

   b. To send alert notifications to others, enter their email addresses in a comma separate list in the**Email Recipients** field.

   c. Check the **Send alert emails to Nutanix Support** box to send alert notifications to Nutanix customer support.

   d. Click the **Save** button to apply the changes and close the window.

      If you want to configure rules or an email message template, continue to the following steps before clicking the **Save** button.

   The **Tunnel Connection** section displays mail transport status information.

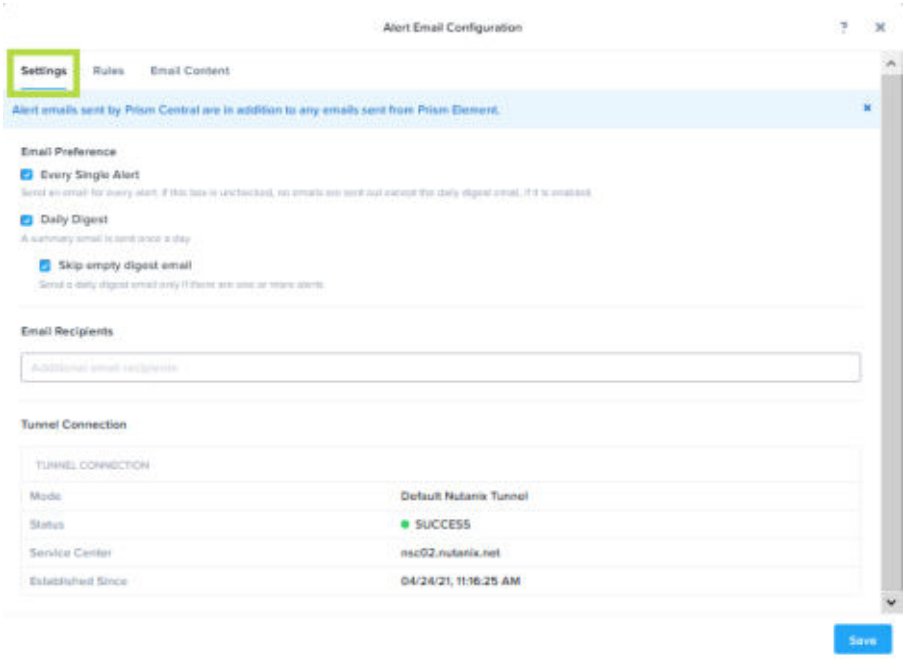

Figure 159: Alert Email Configuration Window (Settings tab)

3. To create a custom alert email rule, click the **Rules** tab and the **New Rule** button, and then configure the rule as follows:

   a. Specify the conditions for generating the alert:

      - **Alert Severity**: Select one or more of the severities from the pull-down list (Critical, Warning, Info, All).
      - **Impact Type**: Select one or more of the categories from the pull-down list (Availability, Capacity, Configuration, Performance, System Indicator, All).
      - **Cluster**: Select one or more of the clusters from the pull-down list (*cluster_name*, All).
      - **Alert Contains**: Enter a key phrase or word that should generate an email notification whenever the alert contains that phrase. For example, if you want to get an email notification when an alert contains the phrase "low memory", enter low memory in the field.

   b. Specify who should receive the alert email by entering recipient email addresses as a comma separated list in the **Send Email To** field.

   c. Click the **Save Rule** button to save the rule.

   d. Repeat these steps to create more custom rules.

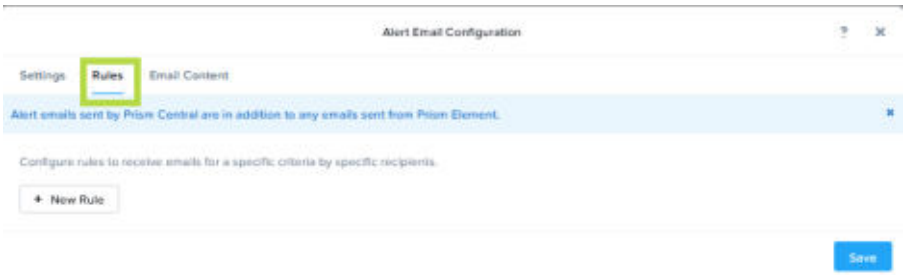

Figure 160: Alert Email Configuration Window (Rules tab)

4. To create a template for the email messages, click the **Email Content** tab and do the following:

   a. In the **Prepend Subject** field, enter the desired text.

      This text appears at the beginning of the subject field in each alert email. If the field is left blank, no prepended text appears in the subject.

   b. In the **Append Body** field, enter the desired text.

      This text appears at the end of the message body in each alert email. If the field is left blank, no appended text appears in the message body.

   c. Click the **Save** button to apply the changes and close the window.

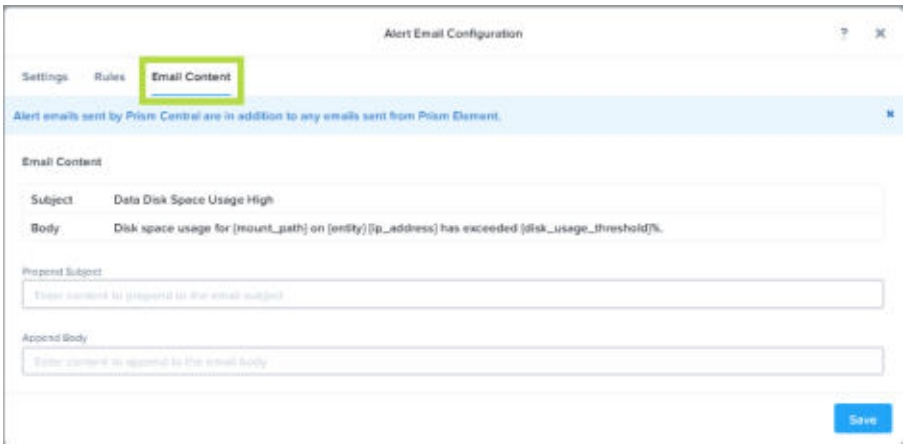

Figure 161: Alert Email Configuration Window (Template tab)

Configuring Alert Policies

About this task

The system monitors a variety of conditions and sends an alert whenever one of the alert conditions is detected (when alerting is enabled). There are default policies for these alerts, but you have the option to modify the default policies and add new policies.

> Note: To help ensure that Prism Central and each managed cluster are taking advantage of NCC features, ensure that:
>
> - Each node in your cluster is running the same NCC version.
> - Prism Central and each cluster managed by Prism Central are all running the same NCC version.

There are three categories of alert policies:

Procedure

- System Defined Alert Policies. You can view and customize policies for any of the system alerts (see Modifying System Alert Policies on page 227).
- User Defined Alert Policies. You have the option to create custom alert policies (see Adding Custom Alert Policies on page 229).

- External Defined Alert Policies. There is an API that external sources can use to define alert policies. These are policies you can view (but not modify) through the alerts dashboard (see Viewing External Alert Policies on page 244).

Modifying System Alert Policies

About this task

There are default policies for system alerts, but you can customize which system events should trigger an alert. To modify the system alert policies, do the following:

Procedure

1. Go to the **System Defined** tab of the alerts dashboard and click the name of the target system alert policy (see Alerts Summary View on page 210).
2. In the **Update Policy** window, do the following in the indicated fields:
   a. **Global Rule**: Uncheck (or check) the box next to the severity to disable (or re-enable) this as an alert.

      All the alerts are enabled (box checked) by default. In most cases this field includes just a single box with the word Critical, Warning, or Info indicating the severity level. Checking the box means this event will trigger an alert of that severity. Unchecking the box means an alert will not be issued when the event occurs. In some cases, such as in the example figure about disk space usage, the event can trigger two alerts, a warning alert when one threshold is reached (in this example 75%) and a critical alert when a second threshold is reached (90%). In these cases you can specify whether the alert should be triggered (check/uncheck the box) and at what threshold (enter a percentage in the box).

   b. **Auto Resolve These Alerts**: Uncheck (or check) the box to disable (or re-enable) automatic alert resolution.

      Automatic alert resolution is enabled for all alert types (where applicable) by default. When this is enabled, the system will automatically resolve alerts under certain conditions such as when the system recognizes that the error has been resolved or when the

initiating event has not reoccurred for 48 hours. (Automatic resolution is not allowed for some alert types, and this is noted in the policy window for those types.)

c. **Exceptions**: To exclude a cluster from this global rule, click the **Add Exceptions** link to open an exceptions line and then select the target cluster from the pull-down list.

   The same check boxes for severity and auto resolve appear. Create a custom rule for the cluster by checking (unchecking) the boxes that should apply to that cluster. Repeat for each cluster you want changed from the global rule.

d. Click the **Save** button to save the changes and return to the alerts dashboard.

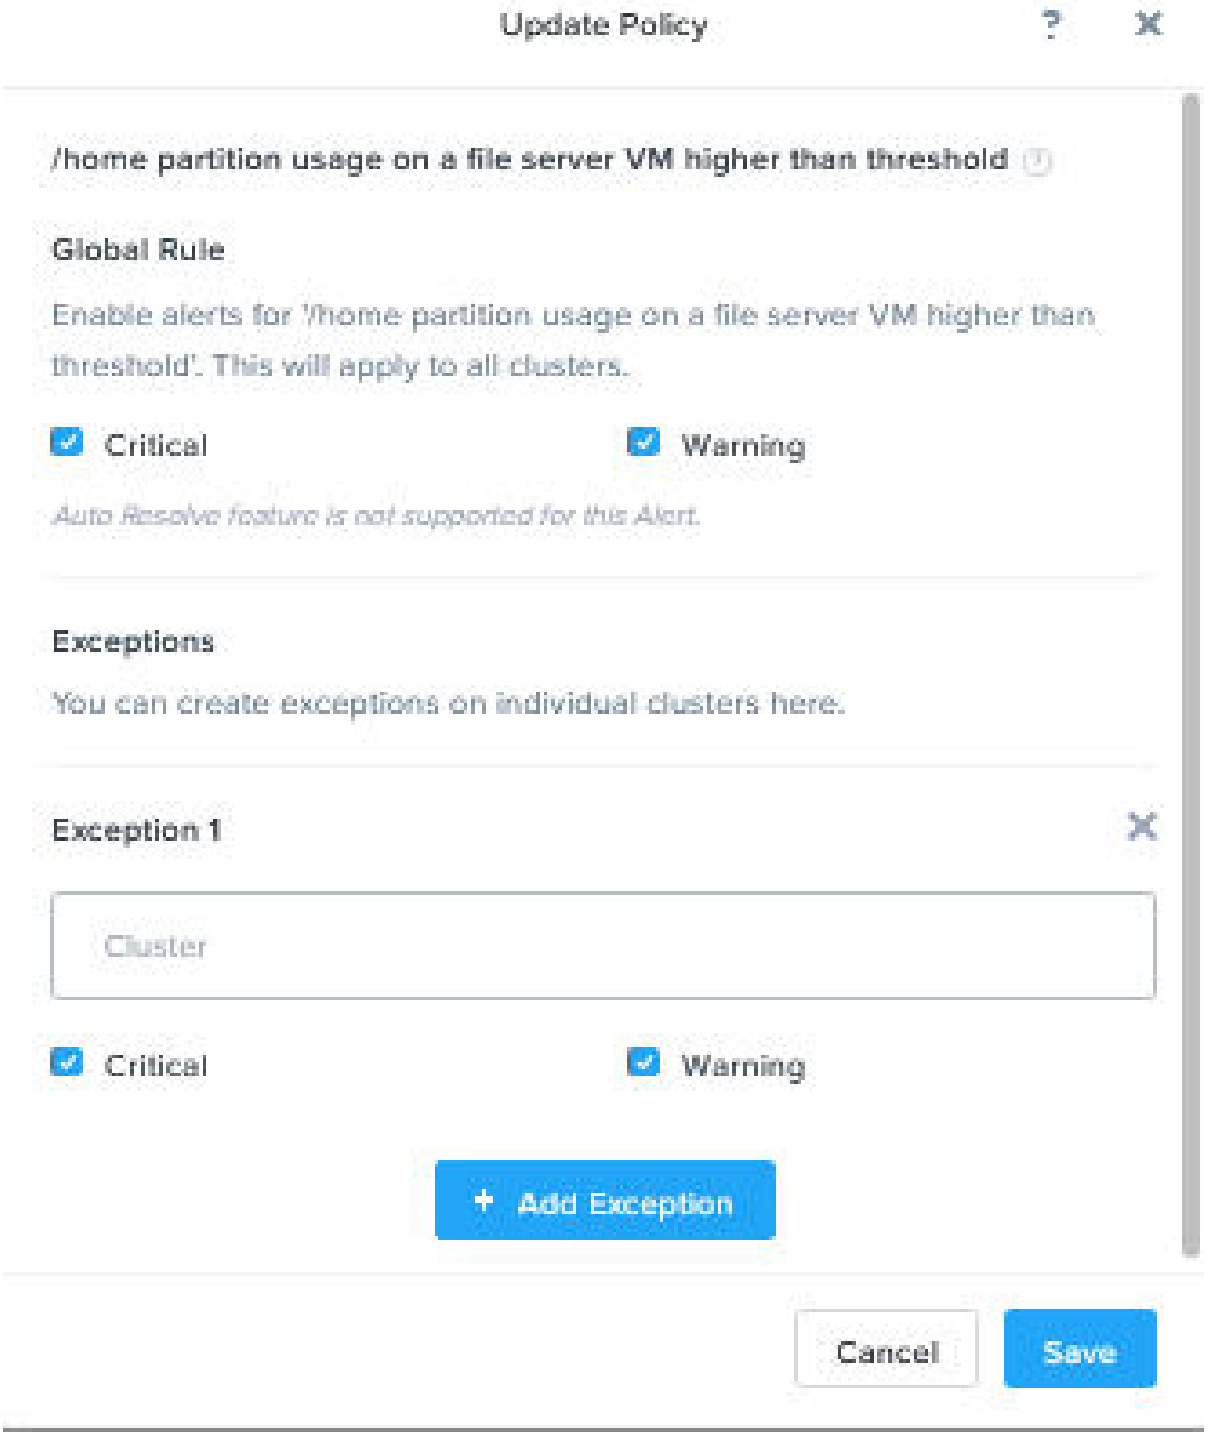

Figure 162: Update Policy Window (example)

Adding Custom Alert Policies

Prism Central provides a mechanism to define alert policies to monitor the entities with respect to the metrics that are critical to your virtual infrastructure. You can define alert policies to monitor entities such as VMs, hosts, and clusters. You can also define a single alert policy for either all VMs, all hosts, or all clusters that share some common criteria. You can create, update,

delete, enable, and disable the alert policies. Each enabled alert policy is executed every five minutes. You can also configure an alert policy so that Prism Central can automatically resolve the alert after the condition causing the alert is resolved.

For a complete description of the metrics available to create custom alert policies, see Alert Metrics on page 234.

Creating Custom Alert Policies

About this task

To create a custom alert policy, do the following.

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and click the **Create Alert Policy** button (see Alerts Summary View on page 210).
   The Create Alert Policy page appears.

Figure 163: Create Alert Policy page (example)

2. In the left pane, do the following in the indicated fields:
   a. **Entity Type**: Select the entity type (**Host**, **VM**, or **Cluster**) from the pull-down list.
   b. **Entity**: Select the target entity from the pull-down list. The choices allow you to specify a single entity or a group of entity. For example, if the entity type is **Host**, you can specify

All Hosts, **One Host**, or **All Hosts in a Cluster**. For all options except **All Hosts**, a search field appears to select the target host, category, or cluster.

> Note: The category option is available only if you select the **VM** entity type.

c. **Metric**: Select the desired metric from the pull-down list.

   The list of metrics varies depending on the entity type. When the metric is selected, more fields appear in the right pane to specify the conditions for the policy.

d. **Policy Name**: Enter a name for the policy.

   The system automatically generates a name for this field based on the policy criteria you have selected, but you can edit the name as desired.

e. **Description** (optional): Enter a description for the policy.

f. **Impact Type**: Select the type of impact the alert condition causes to the entity (**Performance**, **Capacity**, **Configuration**, or **Availability**) from the pull-down list. The default impact type is **Performance**.

g. **Auto resolve alerts**: Check this box to automatically resolve the alert generated by the policy after the condition causing the alert no longer persists (the condition must not occur again for 48 hours).

h. **Enable Policy**: Check this box to enable (or uncheck to disable) the policy.

3. In the right pane, do the following in the indicated fields:

   a. **<Metric Name>**: The top field displays an anomalies chart from the (machine learned) expected behavior for the specified metric. Hovering over the data displays additional information. Select the chart duration (last 12 hours, last 24 hours, or last week) from the pull-down list on the right.

      > Note: The following fields allow you to specify that an alert should be generated when the metric deviates from the expected behavior (**Behavioral Anomaly** fields), when the metric

crosses a set limit (**Static Threshold** fields), or both. See Behavioral Learning Tools on page 582 for more information about behavioral anomalies.

Note the following maximum limit of **Behavioral Anomaly** policies for VMs:

- For a small Prism Central VM (4 vCPUs, 16 GB memory) with more than 4000 VMs, you must not have more than 10 **Behavioral Anomaly** policies for VMs.
- For a large Prism Central VM (8 vCPUs, 32 GB memory) with more than 10,000 VMs, you must not have more than 5 **Behavioral Anomaly** policies for VMs.
- For a Prism Central VM in general, the total number of alert policies must not exceed 50.

If the number of policies exceeds the recommended maximum limit, the user-created alerts and behavioral anomaly alerts might be skipped.

b. **Every time there is an anomaly, alert**: Check this box to generate an alert when an anomaly occurs. Select **Critical** or **Warning** from the pull-down list to set this as a critical or warning level alert.

c. **Ignore all anomalies between**: Check this box to trigger the alert only when the anomaly occurs outside a certain range (percentages or units). Select the operator and value for the range.

d. **Alert critical if**: Check this box to generate a critical alert. Select the operator and value for the critical alert.

e. **Alert warning if**: Check this box to generate a warning alert. Select the operator and value for the warning alert. This option is only available if the behavioral anomaly threshold is not set.

f. **Trigger alert if condition persists for**: Select the duration from the pull-down list (0-30 minutes) for which the alert condition must persist before the system generates an alert.

4. Click **Save** to save the alert policy.

Overlapping Policies

If you are creating an alert policy specific to only an entity (a VM, host, or cluster), but you have earlier created a global policy with the same metric for all entities of that entity type (all VMs, all hosts, or all clusters), the policy that you are creating becomes an overlapping policy. For example, if you have created a policy (global-host-policy-1) that triggers a critical alert if the memory usage of any host exceeds 95%. However, you are now creating a policy (host1-policy) that triggers a critical alert if the memory usage (same metric) of host1 (a specific host) exceeds 90%, host1-policy is now an overlapping policy for global-host-policy-1.

When you are creating a new policy (see Creating Custom Alert Policies on page 230), if you select the same entity and metric as an existing policy but different values and you click **Save**, Prism Central displays a message that a similar policy exists as shown in the following screenshot:

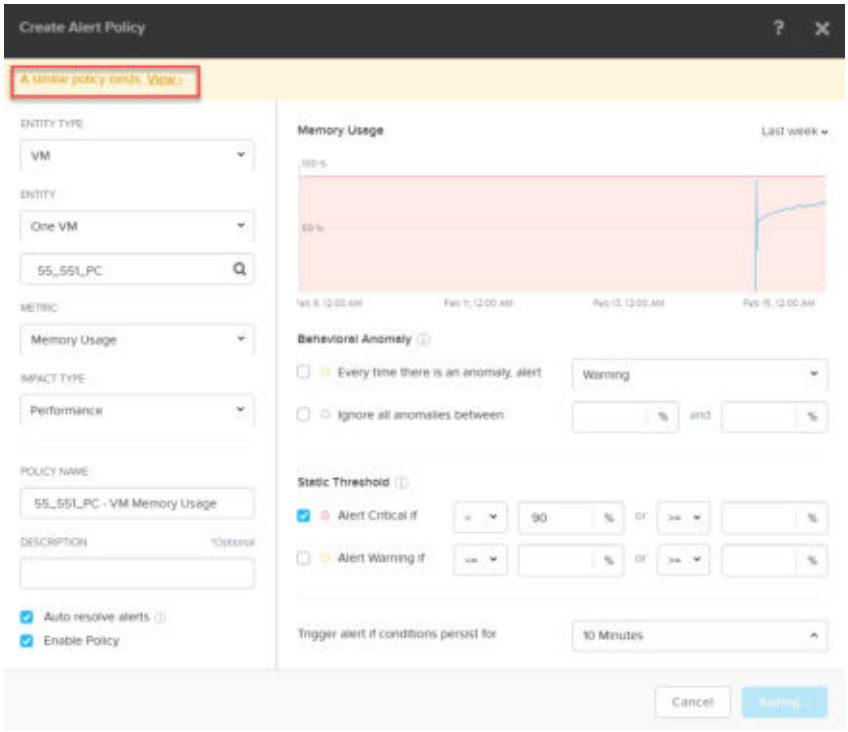

Figure 164: Overlapping Policy

Click **View** to display the overlapping policies and click **Save** to save the overlapping policy.

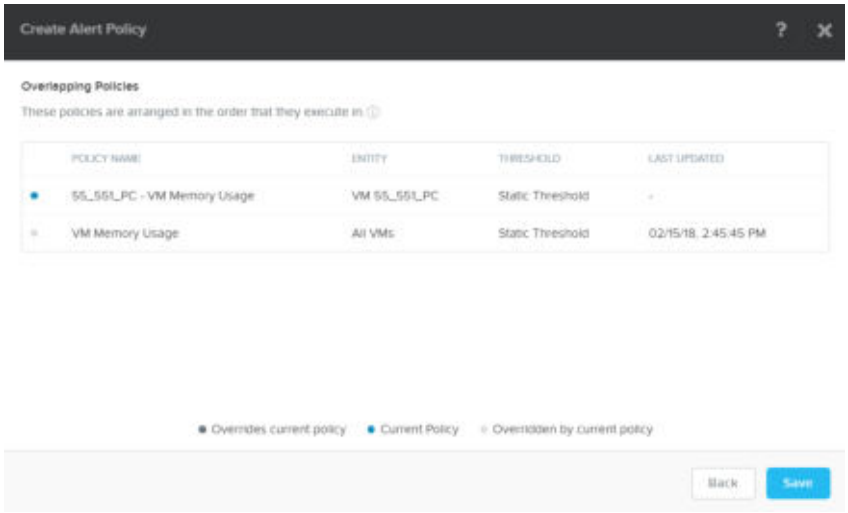

Figure 165: Overlapping Policies List

Precedence of Overlapping Polices

If an overlapping policy exists (policy that applies to a specific entity) for a global policy (policy that applies to all the entities of that entity type), the overlapping policy takes precedence over the global policy.

If multiple policies apply to an entity during policy execution (that is overlapping policies exist), the policy with the highest precedence for that entity is evaluated and all other policies that apply to that entity are ignored. Even if a policy is disabled, the policy might still overlap other

policies. If the policy that has the highest precedence is disabled, that policy is not applied and all other policies that apply to that entity are also ignored.

Policies are evaluated in the following order of precedence (highest precedence policy at the top):

1. Policy applied to a specific entity
2. Policy applied to an entity type in a category
3. Policy applied to an entity type in a cluster
4. Policy applied to all entities of an entity type

Precedence of Overlapping Polices in a Category

A category policy (policy that applies to all VMs in a category) takes precedence over a global policy (a policy that applies to all VMs in the cluster). However, a VM can belong to multiple categories. In this case, the category policy that is most recently updated takes precedence over any other category policy that applies to a VM.

For example, consider that you have created a global policy (global-policy-1) for all VMs in the cluster. You have VM-1 that belongs to both the categories: category-1 and category-2. You have created the category-1-pol policy for all VMs in category-1 and the category-2-pol policy for all VMs in category-2. In this case, all three policies apply to VM-1, but only one policy can be applied to a single VM. Because category-1-pol and category-2-pol are category policies, category-1-pol and category-2-pol take precedence over global-policy-1. However, you still have two policies (category-1-pol and category-2-pol) that apply to the same VM (VM-1). In this case, the policy that is updated most recently takes precedence over the other.

Modifying Custom Alert Policies

About this task

To modify a custom alert policy, do the following:

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and click the name of the target custom alert policy (see Alerts Summary View on page 210).
   The *policy_name* page appears. which includes the same fields as the Create Alert Policy page (see Creating Custom Alert Policies on page 230).
2. Make the desired changes in the appropriate fields and then click **Save**.

Deleting Custom Alert Policies

About this task

To delete a custom (user-created) alert policy, do the following:

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and select the target custom alert policy in the list (see Alerts Summary View on page 210).
2. Select **Delete** from the **Actions** pull-down menu.

Alert Metrics

The following table describes the VM metrics.

Table 96: VM Metrics

CPU Ready Time	
Metric	hypervisor.cpu_ready_time_ppm
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the time a virtual machine waits to use the physical CPU out of the total CPU time allotted to the VM. AHV: Ratio of the VM wait time to the total CPU time allotted to the VM in percentage. ESXi: Value of cpu.ready.summation for a VM in percentage.
Supported Hypervisors	AHV and ESXi
CPU Usage	
Metric	hypervisor_cpu_usage_ppm
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the total CPU cycles currently in use by the VM reported by the hypervisor. Note: This metric identical to the system created alert policy A3040. ESXi: Value of cpu.usage.average reported by ESXi for a VM. Hyper-V: Value of AvgCpuUtilizationPercent reported by Hyper-V for a VM.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Memory Usage	
Metric	hypervisor_memory_usage_ppm
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the memory used by the VM out of the total configured memory reported by the hypervisor. Note: This metric is identical to the system created alert policy A3021. AHV: Memory usage is always 100% because AHV does not support overcommit. ESXi: Value of mem.usage.average reported by ESXi for a VM.
Supported Hypervisors	ESXi and AHV
Memory Balloon*	
Metric	hypervisor_memory_balloon_reclaimed_bytes
Rollup	Average

Unit	Bytes
Description	Total guest memory in bytes reclaimed through ballooning. Note: This metric is similar to the system created alert policy A3024.
Supported Hypervisors	ESXi
Memory Balloon Target*	
Metric	hypervisor_memory_balloon_reclaim_target_bytes
Rollup	Average
Unit	Bytes
Description	Target value in bytes for a VM's memory balloon.
Supported Hypervisors	ESXi
Memory Swap In Rate*	
Metric	hypervisor_swap_in_rate_kBps
Rollup	Average
Unit	KBps
Description	Rate of the memory swapped in from the disk to active memory in kilobytes per second.
Supported Hypervisors	ESXi
Memory Swap Out Rate*	
Metric	hypervisor_swap_out_rate_kBps
Rollup	Average
Unit	KBps
Description	Rate of the memory swapped out from active memory to the disk in kilobytes per second.
Supported Hypervisors	ESXi
Controller AVG Read I/O Latency	
Metric	controller_avg_read_io_latency_usecs
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the read I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller AVG Write I/O Latency	
Metric	controller_avg_write_io_latency_usecs
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the write I/O latency of the VM in microseconds reported by the controller.

Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller AVG I/O Latency	
Metric	controller_avg_io_latency_usecs
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the read and write I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Read IOPS	
Metric	controller_num_read_iops
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM reads in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Write IOPS	
Metric	controller_num_write_iops
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM writes in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller IOPS	
Metric	controller_num_iops
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM read and writes in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Read IO Bandwidth	
Metric	controller_read_io_bandwidth_kBps
Rollup	Derived from summary rollup type.
Unit	KBps
Description	Read IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Write IO Bandwidth	
Metric	controller_write_io_bandwidth_kBps
Rollup	Derived from summary rollup type.
Unit	KBps

Description	Write IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller IO Bandwidth	
Metric	controller_io_bandwidth_kBps
Rollup	Derived from summary rollup type.
Unit	KBps
Description	Read and write IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Network Transmit Packets Dropped	
Metric	hypervisor_num_transmit_packets_dropped
Rollup	Summation
Unit	Packets
Description	Number of the network transmit packets dropped for the VM.
Supported Hypervisors	ESXi and AHV
Network Receive Packets Dropped	
Metric	hypervisor_num_receive_packets_dropped
Rollup	Summation
Unit	Packets
Description	Number of the network receive packets dropped for the VM.
Supported Hypervisors	ESXi and AHV
Controller Disk Usage	
Metric	controller_user_bytes
Rollup	Average
Unit	bytes
Description	Disk usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
I/O working set size	
Metric	controller.wss_3600s_union_MB
Rollup	Average
Unit	MB
Description	Size of the I/O working set size in megabytes.
Supported Hypervisors	All hypervisors
Write I/O working set size	
Metric	controller.wss_3600s_write_MB
Rollup	Average
Unit	MB

Description	Size of the I/O working set size writes in megabytes.
Supported Hypervisors	All hypervisors
Read I/O working set size	
Metric	controller.wss_3600s_read_MB
Rollup	Average
Unit	MB
Description	Size of the I/O working set size reads in megabytes.
Supported Hypervisors	All hypervisors
Snapshot usage	
Metric	controller.snapshot_usage_bytes
Rollup	Average
Unit	bytes
Description	Snapshot usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
Shared Data	
Metric	controller.shared_usage_bytes
Rollup	Average
Unit	bytes
Description	Snapshot usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
Network Bytes Received	
Metric	hypervisor_num_received_bytes
Rollup	Summation
Unit	bytes
Description	Number of network bytes received reported by the hypervisor.
Supported Hypervisors	All hypervisors
Network Bytes Transmitted	
Metric	hypervisor_num_transmitted_bytes
Rollup	Summation
Unit	bytes
Description	Number of network bytes transmitted reported by the hypervisor.
Supported Hypervisors	All hypervisors

* Experimental only. Might not be supported in the future releases.

Table 97: Host Metrics

The following table describes the host metrics.

CPU Usage	
Metric	hypervisor_cpu_usage_ppm
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the total CPU cycles currently in use by the host. Note: This metric is identical to the system created alert policy A6414. AHV: Ratio between the sum of kernal, user, and iowait and the sum of kernal, user, iowait, and idle over all the CPUs. For hosts with hyperthreads enabled, the host CPU usage is 10% more than the total value. ESXi: Value of cpu.usage.average for a host in percentage. Hyper-V: Value of LogicalProcessorTotalRuntimePercent from the output of the command Get-LocalHostPerformance | Convertto-Json -Depth 5 -Compress.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Memory Swap In Rate*	
Metric	hypervisor_swap_in_rate_kBps
Rollup	Average
Unit	KBps
Description	Rate of the swap-in memory in kilobytes per second. Note: This metric is similar to the system created alert policy A6413. ESXi: Value of mem.swapinRate.average for the host.
Supported Hypervisors	ESXi
Memory Swap Out Rate*	
Metric	hypervisor_swap_out_rate_kBps
Rollup	Average
Unit	KBps
Description	Rate of the swap-out memory in kilobytes per second. Note: This metric is similar to the system created alert policy A6413. ESXi: Value of mem.swapoutRate.average for the host.

Supported Hypervisors	ESXi
Memory Usage	
Metric	hypervisor_memory_usage_ppm
Rollup	Average
Unit	Percentage (%)
Description	Ratio of the total memory allocated to all VMs over the total memory size on the host. Note: This metric is identical to the system created alert policy A6414.
	AHV: Committed memory size from /proc/meminfo over the total memory size. Memory usage is always 100%.
	ESXi: Value of mem.usage.average for the host.
	Hyper-V: Amount of used memory over the total memory.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Network Transmit Packets Dropped	
Metric	hypervisor_num_receive_packets_dropped
Rollup	Summation
Unit	Packets
Description	Number of the network transmit packets dropped.
	ESXi: Value of net.droppedTx.summation for the host.
Supported Hypervisors	ESXi and AHV
Network Receive Packets Dropped	
Metric	hypervisor_num_receive_packets_dropped
Rollup	Summation
Unit	Packets
Description	Number of the network receive packets dropped.
	ESXi: Value of net.droppedRx.summation for the host.
Supported Hypervisors	ESXi and AHV
Controller AVG I/O Latency	
Metric	controller_avg_io_latency_usecs
Rollup	Average
Unit	Microseconds
Description	Average read and write I/O latency of the host in microseconds reported by the controller.

Controller AVG Read I/O Latency	
Metric	controller_avg_read_io_latency_usecs
Rollup	Average
Unit	Microseconds
Description	Average read I/O latency of the host in microseconds reported by the controller.
Controller AVG Write I/O Latency	
Metric	controller_avg_write_io_latency_usecs
Rollup	Average
Unit	Microseconds
Description	Average write I/O latency of the host in microseconds reported by the controller.
Controller Read IOPS	
Metric	controller_num_read_iops
Rollup	Average
Unit	IOPS
Description	Rate of the host reads in IOPS reported by the controller.
Controller Write IOPS	
Metric	controller_num_write_iops
Rollup	Average
Unit	IOPS
Description	Rate of the host writes in IOPS reported by the controller.
Controller IOPS	
Metric	controller_num_iops
Rollup	Average
Unit	IOPS
Description	Rate of the host reads and writes in IOPS reported by the controller.
* Experimental only. Might not be supported in the future releases.	

Table 98: Cluster Metrics

The following table describes the host metrics.

Controller Read IOPS	

Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller reads in IOPS.
Controller Write IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller writes in IOPS.
Controller IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller reads and writes in IOPS.
Controller AVG Read I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the read I/O latency of the controller in microseconds.
Controller AVG Write I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the write I/O latency of the controller in microseconds.
Controller AVG I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the read and write I/O latency of the controller in microseconds.
Controller I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Read and write IO bandwidth in kilobytes per second reported by the controller.
Controller Read I/O Bandwidth	
Rollup	Average
Unit	KBps

Description	Read IO bandwidth in kilobytes per second reported by the controller.
Controller Write I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Write IO bandwidth in kilobytes per second reported by the controller.
CPU Usage	
Rollup	Average
Unit	PPM
Description	CPU usage of the cluster in PPM reported by the hypervisor.
Memory Usage	
Rollup	Average
Unit	PPM
Description	Memory usage of the cluster in PPM reported by the hypervisor.

Viewing External Alert Policies

External alert policies are those created by applications.

Nutanix provides an API that applications can use to define alert policies. The **External Defined** tab displays a list of the application-defined alerts (see Alerts Summary View on page 210). You cannot modify an externally defined policy, but you can view the policy by clicking the policy name in the list.

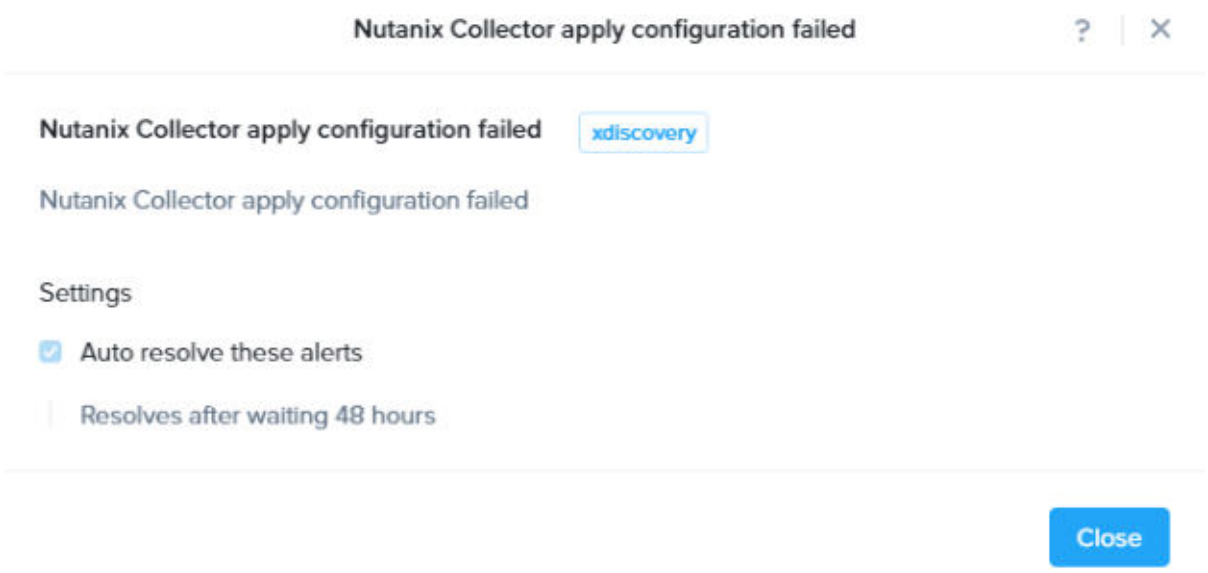

Figure 166: External Defined Policy Window (example 1)

Stopping Alerts During a Maintenance Window

You can stop alerts from being generated during a maintenance window by using Prism Central. You must log on to the Prism Central VM with SSH and create configurations to stop the alerts by using the nuclei command line interface.

About this task

You can stop the alerts of only the **cluster** entity type.

Perform the following procedure to create or modify a configuration to stop the alerts.

Procedure

1. Log on to the Prism Central VM with SSH.
2. Create a specifications file as shown in the following example to create or update a configuration.

```
{
 "spec": {
  "resources": {
   "scope_entity_list": [
    {
     "entity_type": "cluster",
     "uuid_list": ["123", "223"]
    }
   ],
   "schedule_list": [
    {
     "interval_multiple": 1,
     "duration_secs": 300,
     "end_time": "2017-08-20T23:47:45.141Z",
     "start_time": "2017-07-20T23:47:45.141Z",
     "interval_type": "DAILY"
    }
   ],
  }
},
"api_version": "3.0",
"metadata": {
  "kind": "blackout"
}
```

```
}
```

Specify the values for the following parameters:

- Under `scope_entity_list`, specify the scope of the entities for which you want to stop the alerts as follows:
  - entity_type. Specify the type of entities for which you want to stop the alerts. You can stop the alerts of only the **cluster** entity type in this release.
  - uuid_list. Specify the UUIDs of the entities for which you want to stop the alerts. For example, specify the UUIDs of the clusters for which you want to stop the alerts as shown in the example specification file.

    Note: If this field is empty, alerts for all entities of that entity type are stopped. For example, use this parameter if you want to stop the alerts for specific clusters. If you want to stop the alerts for all clusters, leave this field empty.

- Under `schedule_list`, define the schedule during which the alerts should be stopped as follows:
  - interval_type. Define the recurrence to stop the alerts (ONCE, DAILY, WEEKLY, or MONTHLY).
  - interval_multiple. Define the frequency to stop the alerts. This is related to interval_type. For example:
    - If interval_type=Daily and interval_multiple=1, the alert is stopped every day.
    - If interval_type=Daily and interval_multiple=2, the alert is stopped every alternate day.
  - start_time and end_time. Define start and end time to stop the alerts. The date time format is same as used in Nutanix REST API calls.

    ```
    date-fullyear(4digit) "-" date-month(2digit) "-" date-mday(2digit)
    "T" time-hour(2digit) ":" time-minute(2digit) ":" time-second(2digit)
    [.] "Z"  or ("+" / "-") time-hour(2digit) ":" time-minute(2digit)
    ```

  - duration_secs. Define the duration in seconds for which the alerts should be stopped.
  - day_of_week_list. Define the days of the week when the alerts should be stopped. You must include this parameter only if interval_type=WEEKLY.

3. Create a configuration to stop the alerts.

   ```
   $ nuclei --username admin --password password blackout.create spec_file=filepath
   ```

   Replace *password* with the password of the admin user account and *filepath* with the path of the specification file that you created.

4. List all the configurations that stop alerts.

   ```
   $ nuclei --username admin --password password blackout.list
   ```

   Replace *password* with the password of the admin user account.

   The output of this command lists the UUID of the each configuration that you have created.

5. Update a configuration.

   ```
   $ nuclei --username admin --password password blackout.put uuid spec_file=filepath
   ```

   Replace *password* with password of the admin user account, *filepath* with the path of the specification file that you created, and *uuid* with UUID of the configuration that you want to update (obtained in step 4).

6. Get a configuration.

   ```
   $ nuclei --username admin --password password blackout.get uuid
   ```

   Replace *password* with password of the admin user account and *uuid* with UUID of the configuration that you want to get (obtained in step 4).

Prism Central Logs

This section lists the logs that reside in a Prism Central VM. These logs record Prism Central-related events and actions. A Prism Central VM runs a subset of the services that run in a Controller VM plus a service about VMs (Metropolis) that is unique to Prism Central. Prism Central logs contain comparable information as the similarly named ones for a Controller VM. See the "logs" section in the *Acropolis Advanced Administration Guide* for more information.

> Note:
>
> - The timestamps for all Nutanix service logs are moved to UTC (in ISO 8601:2020-01-01 T00:00:00Z) from Prism version 5.18.
> - All operating system logs will not be moved to UTC, hence Nutanix recommends that you set the server local time to UTC.

Table 99: Location: **/home/nutanix/data/logs**

Log	Contents	Frequency
alert_manager.[out, ERROR, FATAL, INFO, WARNING]	Alert manager process output	
aplos.out	Aplos V3 gateway logs	
aplos_engine.out	Aplos V3 engine logs	
athena.[out, ERROR, INFO, FATAL]	Authentication and identity management activity	
cassandra_monitor.[out, ERROR, FATAL, INFO]	Cassandra database monitor process output	
catalina.out	Catalina/Tomcat for Prism process output	
connection_splicer.[out, ERROR, FATAL, INFO, WARNING]	Internal process connection status	
consolidated_audit.log	Audit logs for VM operations, see Audit Log Events on page 250 for details.	
cron_avahi_monitor.log	Avahi process status	

Log	Contents	Frequency
cron_time_check.log	Check time difference across Prism Central VMs when Prism Central is a multi-VM instance.	every 1 min
disk_usage.log	Disk and inode usage of all partitions on the Prism Central VM	every 1 min
dynamic_ring_changer.[out, ERROR, FATAL]	Metadata migration across nodes activity	
genesis.out	Nutanix software start process output	
hyperint_monitor.[out, ERROR, FATAL, INFO, WARNING]	Hypervisor integration activity	
karbon_core.out (formerly acs_controller.out)	Karbon controller output. This file applies only when Karbon is enabled (see Enabling Karbon on page 652).	
karbon_ui.out (formerly acs_ui.out)	Karbon user interface output. This file applies only when Karbon is enabled.	
metropolis.out	Internal Prism Central service about VMs (Metropolis) to validate cluster choice with scheduler, check out images from the catalog service, and forward fanout requests to Prism Element	
prism_gateway.[out, ERROR, FATAL, INFO]	Prism leader activity	
prism_monitor.[out, ERROR, FATAL, INFO]	Prism (web console, nCLI, REST API) monitor process output	
scavenger.out	Log and core file clean-up status	
search.[out, FATAL]		
search_monitor.[ERROR, INFO, WARNING]		
send-email.log	E-mail alerts sent from the Controller VM	every 1 min
ssh_tunnel.log	Connect status to nsc.nutanix.com for the remote support tunnel	
stats_aggregator.[out, ERROR, FATAL, INFO]	Statistics aggregator process output	
support-info.log	Daily automated support (ASUP) alerts	

Log	Contents	Frequency
using-gflags.log	gflags status	
zeus_config_printer.INFO	Contents of cluster configuration database	
zookeeper_monitor.[out, ERROR, INFO]	Cluster configuration and cluster state activity	

Table 100: Location: **/home/nutanix/data/logs/cassandra**

Log	Contents
system.log	Cassandra system activity

Table 101: Location:**/home/nutanix/data/logs/sysstats**

Log	Contents	Frequency	Command
df.info	Mounted filesystems	every 5 sec	df -h
disk_usage.INFO	Disk usage across disks	every 1 hour	du
interrupts.INFO	CPU interrupts	every 5 sec	
iostat.INFO	I/O activity for each physical disk	every 5 sec	sudo iostat
iotop.INFO	Current I/O in realtime	every 5 sec	sudo iotop
lsof.INFO	List of open files and processes that open them	every 1 min	sudo lsof
meminfo.INFO	Memory usage	every 5 sec	cat /proc/meminfo
metadata_disk_usage.INFO	Disk usage for metadata drives	every 5 sec	
mpstat.INFO	CPU activities per CPU	every 5 sec	mpstat
ntpq.INFO	NTP information	every 1 min	ntpq -pn
ping_gateway.INFO	Pings to the default gateway	every 5 sec	ping
ping_hosts.INFO	Pings to other Prism Central VMs (if Prism Central is a multi-VM instance)	every 1 min	ping
sar.INFO	Network bandwidth	every 5 sec	sar -n DEV, -n EDEV
top.INFO	Real-time CPU and memory activity	every 5 sec	top

Table 102: Location: **/home/nutanix/data/serviceability/alerts**

Log	Contents
num.processed	Alerts that have been processed

Table 103: Location: **/var/log**

Log	Contents
dmesg	OS start messages
kernel	OS kernel messages
messages	OS messages after starting

Audit Log Events

The following table provides the list of all the VM events that are audited and captured as part of the audit logs.

Table 104: Audit Event List

Event Type	Operation	Description
VmCreateAudit	Create	Created VM
VmReplicateAudit	Create	Replicated VM
VmCloneAudit	Create	Cloned VM
VmDeleteAudit	Delete	Deleted VM
VmRestoreAudit	Restore	Restored VM from snapshot
VmNicAddAudit	Update	Added NIC to VM
VmNicUpdateAudit	Update	Updated NIC of VM
VmNicDeleteAudit	Update	Deleted NIC of VM
VmDiskAddAudit	Update	Added disk to VM
VmDiskDeleteAudit	Update	Deleted disk of VM
VmPowerOnAudit	Power State Change	Powered on VM
VmPowerOffAudit	Power State Change	Powered off VM
VmResetAudit	Power State Change	Reset VM
VmPowerCycleAudit	Power State Change	Power cycled VM
VmPauseAudit	Power State Change	Paused VM
VmSuspendAudit	Power State Change	Suspended VM
VmResumeAudit	Power State Change	Resumed VM
VmGuestShutdownAudit	Power State Change	Guest shutdown on VM
VmGuestRebootAudit	Power State Change	Guest rebooted on VM

Event Type	Operation	Description
VmAcpiShutdownAudit	Power State Change	ACPI shutdown VM
VmAcpiRebootAudit	Power State Change	ACPI rebooted VM
VmPowerStateReconcileAudit	Power State Change	Reconciled VM power state
HAFailoverStartAudit	HA	Started HA failover on host
HAFailoverCompleteAudit	HA	Completed HA failover on host
EnteringMaintenanceModeAudit	Update	Entering maintenance mode on host
MaintenanceModeAudit	Update	Host in maintenance mode
ExitMaintenanceModeAudit	Update	Exited maintenance mode on host
VmHostAffinityUpdateAudit	Update	Updated host affinity of VM
VmDiskResizeAudit	Update	Resized disk of VM
VmDiskUpdateAudit	Update	Updated disk of VM
VmUpdateAudit	Update	Updated VM
VmExportAudit	Create	VM exported to OVA
OvaCreateAudit	Create	OVA created
OvaDeleteAudit	Delete	OVA deleted
OvaUpdateAudit	Update	OVA name updated
OvaConcatenateAudit	Update	OVA parts concatenated
VmMigrateAudit	Migrate	Migrated VM from one host to another host
VsCreateAudit	Create	Virtual switch created
VsUpdateAudit	Update	Virtual switch updated
VsDeleteAudit	Delete	Virtual switch deleted
MigrateBridgeToVirtualSwitchAudit	Migrate	Migrated bridge to virtual switch
IPAssignmentToVMAudit	Update	IP assigned to VM
IPUnassignmentToVMAudit	Update	IP unassigned from VM
SpanCreateAudit	Create	SPAN session created
SpanUpdateAudit	Update	SPAN session updated
SpanDeleteAudit	Delete	SPAN session deleted
VmHostAffinityPolicyCreateAudit	Create	VM Host Affinity Policy created
VmHostAffinityPolicyUpdateAudit	Update	VM Host Affinity Policy updated
VmHostAffinityPolicyDeleteAudit	Delete	VM Host Affinity Policy deleted

Event Type	Operation	Description
VmHostAffinityPolicyReEnforceAudit	Update	VM Host Affinity Policy Re-enforced
ImageCreateAudit	Create	Created Image
ImageDeleteAudit	Delete	Deleted Image
ImageUpdateAudit	Update	Updated Image

Alerts/Health checks

Cluster

Table 105: CVM Time Difference High [3025] [A1017]

Name	Time Drift Check
Description	Checks whether CVM times match.
Alert message	**Time differs between Controller VMs.**
Cause	Timedrift between CVMs is high. NTP connectivity may be lost.
Impact	The cluster may experience downtime.
Resolution	Check NTP/network configuration or connectivity.
Time Drift Threshold	Warning Threshold: 3 seconds
KB Article	1512
Severity	Warning

Table 106: Cluster services are down [3032] [A3032]

Name	Cluster Services Down Check
Description	Check that no services are down in the cluster.
Alert message	**Cluster Service: *service_name* is down on the Controller VM *ip_address*.**
Cause	One or more services in the cluster are down.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to serve I/O requests.
Resolution	[If the reason for the service(s) being down is unknown, run "ncc health_checks run_all", perform log collection covering the period of this alert timestamp, and provide the results to Nutanix Support via case attachment for further investigation. If due to planned maintenance or triage, please ensure the service is started again using "cluster start" as soon as possible. Report any reoccurring issues.]
KB Article	1518
Severity	Critical

Table 107: Alert Manager Service Check [3033]

Name	Alert Manager Service Check
Description	Check if Alert Manager service is available.
Cause	This check usually indicates that the Alert Manager service is not healthy, but there could be other causes.
Impact	Cluster issues may stay unnoticed.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.
KB Article	3377

Table 108: Cluster Service Restarting Frequently [3034] [A3034]

Name	Cluster Services Status
Description	Check if services have restarted recently across the cluster.
Alert message	There have been multiple service restarts of *services* within one day across all Controller VM(s). Latest crash of these services have occured at *timestamps* respectively.
Cause	This alert indicates that one or more services in the cluster were restarted.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.
KB Article	3378
Severity	Warning

Table 109: Cassandra service status check [21009]

Name	Cassandra service status check
Description	Check that all Cassandra nodes are online and and functioning normally
Cause	Cassandra service is not running in one or more hosts.
Impact	Cluster resilience might degrade.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.
KB Article	1547

Table 110: Cassandra service is running out of memory [21011] [A21011]

Name	Cassandra memory usage

Description	Check if the Cassandra service is running out of memory.
Alert message	**Cassandra service is running out of memory. Contact Nutanix support for assistance.**
Cause	Cassandra has crossed memory threshold more than 5 times in the last 3 minutes.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.
KB Article	1576
Severity	Critical

Table 111: Same timezone check [103085] [A103085]

Name	CVM same timezone check
Description	Check that all CVMs are in the same timezone
Alert message	*vm_type*s are not in the same timezone.
Cause	Some of the Controller VMs in the cluster are not in the same timezone.
Impact	It is recommended to have all the CVMs in the same timezone to avoid potential issues.
Resolution	Make sure all Controller VMs are in the same time zone. If they are not, then set the timezone using KB 1050
KB Article	2044
Severity	Warning

Table 112: CVM IP address is unreachable [103088] [A103088]

Name	Inter-CVM connectivity check
Description	Check that all CVMs are reachable via ping.
Alert message	**CVM *dest_ip* is unreachable**
Cause	The Controller VM is down or there is a network connectivity issue.
Impact	Storage may be unavailable and/or performance issues may be observed.
Resolution	If the Controller VM does not respond to ping, turn it on. Ensure that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	1635
Severity	Warning

Table 113: Host same timezone check [103092]

Name	Host same timezone check

Description	Check that all hosts are in the same timezone
Cause	Not all hosts in the cluster are in the same timezone.
Impact	Ensure all hosts in the cluster are in the same timezone to help ensure all guest VMs share the same timezone and log messages are timestamped consistently.
Resolution	Ensure that all hosts are in the same timezone by following the instructions in KB 1050
KB Article	2044

Table 114: Checks if AOS version of the target site configured for synchronous replication is on AOS version >= 5.17 [110021] [A110021]

Name	Remote AOS Version Check
Description	Checks if the target site configured is on AOS version >= 5.17 to support Synchronous Replication
Alert message	For protection rule: *protection_rule_name*, target cluster(s): *remote_site_list* have version older than 5.17
Cause	Target site is on AOS version < 5.17
Impact	Synchronous Replication will not be supported.
Resolution	Upgrade the cluster on the target site to AOS version >= 5.17. It is recommended to have both the source and target sites at the same version.
KB Article	7364
Severity	Warning

Table 115: Round-trip travel (RTT) to target Prism Element(s) is greater than 5ms, the maximum value allowed for synchronous replications. [110022] [A110022]

Name	Remote Site Latency Check
Description	Check if the Round-trip travel (RTT) is less than 5ms, the maximum supported for Synchronous Replication
Alert message	Round-trip travel (RTT) to target Prism Element(s) *err_msg* is greater than *max_permitted_latency*
Cause	Target cluster is unreachable or Round-trip travel (RTT) to target cluster is greater than 5ms.
Impact	Synchronous Replication will be affected.
Resolution	Ensure that the target cluster is reachable and the RTT is less than 5ms or choose another target cluster.
KB Article	7366
Severity	Warning

Table 116: Stale synchronous replication configuration found. [110456] [A110456]

Name	Stale Synchronous Replication configuration check
Description	Check if there are any stale synchronous replication configurations present on Prism Central.
Alert message	Stale synchronous replication configuration found for vms *err_msg*.
Cause	Stale synchronous replication configuration left behind.
Impact	Stale entries could cause unwanted tasks to be generated and they cannot be cleaned up automatically on the secondary as it could affect the VM stretch state hence require a manual cleanup.
Resolution	Cleanup the stale synchronous replication configuration left behind on Prism Central.
KB Article	9716
Severity	Warning

Table 117: Cassandra nodes up [111002]

Name	Cassandra nodes up
Description	Check that all Cassandra nodes are running
Cause	One or more Cassandra nodes are down which may be due to several issues including network issue, disk related issue, CPU or resource crunch.
Impact	Meta data load is inconsistent.
Resolution	Cassandra will repair itself. If other Cassandra checks are failing, contact Nutanix support.
KB Article	1369

Table 118: Cassandra tokens consistent [111003]

Name	Cassandra tokens consistent
Description	Check that all Cassandra nodes are consistent
Cause	The tokens may be inconsistent because of manual intervention or because of bad configuration.
Impact	Meta data load is inconsistent.
Resolution	If no other Cassandra checks are failing, repair nodetool by following KB 1369. If other Cassandra checks are failing, contact Nutanix support.
KB Article	1369

Table 119: Name server configuration [111022]

Name	Name server configuration
Description	Check for DNS Server configuration

Cause	Name server is not configured on CVMs and hypervisor hosts or is not able to resolve queries.
Impact	Name resolution will not work.
Resolution	Verify if Name server is configured on CVMs and on hypervisor hosts and if it can resolves queries. Review KB 3005 for more details.
KB Article	3005

Table 120: HTTP proxy check [111023]

Name	HTTP proxy check
Description	Check if HTTP proxy is working
Cause	No proxy configured.
Cause	Unable to connect to proxy on the port.
Impact	Cluster may not have access to Internet-based service such as 1-click auto and manual downloads, Pulse, LCM inventory, etc.
Resolution	Check the proxy server settings (network, port, user credentials, etc.)
KB Article	1708

Table 121: LDAP configuration [111024]

Name	LDAP configuration
Description	Check LDAP configuration.
Cause	LDAP not correctly configured in the cluster.
Impact	Directory users might not be able to log properly.
Resolution	Review KB 2997.
KB Article	2997

Table 122: Virtual IP Check [111027] [A111027]

Name	Virtual IP check
Description	Check if virtual IP is configured and reachable.
Alert message	*vm_type* **Virtual IP is configured but unreachable.**
Cause	Cluster virtual IP is not configured.
Impact	Nutanix features that use virtual IP address might be adversely affected.
Impact	
Resolution	Configure a valid virtual IP for the cluster.
KB Article	3708
Severity	Critical

Table 123: Advanced Vector Extensions (AVX) enabled [111035]

Name	Advanced Vector Extensions (AVX) enabled
Description	Checks AVX enabled on CPUs.
Cause	Intel's Advanced Vector Extensions (AVX) is disabled on the node.
Impact	If AVX is disabled, upgrade will fail.
Resolution	Enable AVX on the node. Review KB 8148.
KB Article	8148

Table 124: Prism Central using default password [111050] [A111050]

Name	PC using default password
Description	PC is using default password
Alert message	One or more Prism Central instances are still using the default credentials
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration.
Resolution	Ensure that Prism Central is using a custom password for improved security.
KB Article	6153
Severity	Info

Table 125: Notifications dropped check [111052]

Name	Notifications dropped check
Description	Some alert notifications were dropped
Cause	One or more cluster services are not healthy.
Impact	Some of the alerts will not be raised. Cluster issues may stay unnoticed.
Resolution	Contact Nutanix support.
KB Article	4905

Table 126: Failure to copy Image to Cluster [111072] [A111072]

Name	Image copy to Cluster check
Description	Detect if image copy to cluster is failing
Alert message	Image could not be copied to Cluster
Cause	Cluster may be unreachable from PC or Catalog service may be down on Cluster.
Impact	Image will not be available on the Cluster.
Resolution	Check Catalog service on cluster and connectivity to PC.

Severity	Warning

Table 127: SW Encryption Keys from {kms_name} are unavailable [111075] [A111075]

Name	SW Encryption Keys Unavailable
Description	Check to see if the KMS actually has the expected passwords for the container encryption keys.
Alert message	SW encryption keys for one or more containers are not available from external key manager *kms_name*.
Cause	If a 3rd party KMS gets restored from a backup or other incorrect management operation, it may have missing keys.
Impact	Cluster will have data unavailability if rebooted.
Resolution	Contact Nutanix and 3rd party key manager support.
KB Article	8223
Severity	Critical

Table 128: High number of tasks in the cluster [111082]

Name	High number of tasks in the cluster
Description	High number of tasks in the cluster
Cause	High number of tasks in the cluster
Impact	Operation on entities in the cluster might not progress.
Resolution	Contact Nutanix support for help.
KB Article	9866

Table 129: Pulse cannot connect to REST server endpoint [140001] [A140001]

Name	REST endpoint connection status
Description	Check if Pulse can connect to REST server endpoint.
Alert message	Pulse cannot connect to REST server endpoint. Connection Status: *connection_status*, Pulse Enabled: *enable*, Error Message: *message*
Cause	Pulse cannot connect to REST server endpoint.
Impact	Data driven serviceability and customer support cannot be performed.
Resolution	Ensure that the REST server endpoint is reachable from Pulse.
KB Article	5490
Severity	Info

Table 130: Pulse is disabled [140002] [A140002]

Name	Pulse Disabled
Description	Pulse is disabled on the paired Prism Central.
Alert message	Pulse is disabled on paired Prism Central *cluster_uuid*.
Cause	Pulse enablement might be missed or disabled.
Impact	Monitoring Nutanix Service Health cannot be performed.
Resolution	Enable Pulse on current cluster and all of its registered clusters.
KB Article	6377
Severity	Info

Table 131: Cluster Connectivity Status [200000] [A200000]

Name	Cluster Connectivity Status
Description	Tests whether the cluster connectivity is fine
Alert message	*component* data from cluster *cluster_name* is not up-to-date.
Cause	The connection between Prism Central and Prism Element is downor Services on Prism Central are crashing.
Impact	Cluster data shown in the Prism Central is not up to date.
Resolution	Ensure that cluster network connectivity is up and all CVM services are up.
KB Article	3379
Severity	Warning

Table 132: IDF Source to Remote Sync Status [200001] [A200001]

Name	IDF DB to DB sync heartbeat check
Description	Checks if IDF replica is receiving requests from IDF master.
Alert message	IDF data on this cluster is not in sync with the IDF data on cluster:*cluster_name*
Cause	Cluster network connectivity or cluster services such as insights server, insights uploader, insights receiver, Aplos, or Prism gateway could be down.
Impact	IDF data that is replicated from master cluster is not up to date.
Resolution	Ensure that cluster network connectivity is up and all cluster services are up.
KB Article	5582
Severity	Warning

Table 133: Prism Central vCPU Availability Check [200301] [A200301]

Name	Prism Central vCPU Availability Check

Description	Checks if the number of vCPUs is sufficient for the number of VM entities in Prism Central.
Alert message	The Prism Central does not have enough vCPUs for the number of VM entities it has.
Cause	Prism Central doesn't have enough vCPUs to support the current number of VMs
Impact	Prism Central performance may be degraded.
Resolution	1) Reduce the number of VM entities in Prism Central or 2) If using 1 node small Prism Central scale up to 1 node large Prism Central or scaleout to 3 node small Prism Central. 3) If using 1 node large Prism Central scaleout it to 3 node large Prism Central.
KB Article	4381
Severity	Warning

Table 134: Insufficient Space for UVMs deployed on PC [200302] [A200302]

Name	PC Sufficient Disk Space Check
Description	Checks if the amount of storage is sufficient for the number of VM entities in Prism Central.
Alert message	The PC does not have enough storage for the number of VM entities it has.
Cause	Too many VMs in Prism Central for the amount of storage in Prism Central.
Impact	Prism Central may run out of disk space to store data.
Resolution	Reduce the number of VM entities in Prism Central or provide more storage for Prism Central.
KB Article	4340
Severity	Warning

Table 135: Prism Central Memory Availability Check [200303] [A200303]

Name	Prism Central Memory Availability Check
Description	Checks if the amount of memory is sufficient for the number of VM entities and the services enabled in Prism Central.
Alert message	The Prism Central does not have enough memory for the current number of VMs.
Cause	Prism Central doesn't have enough memory to support thecurrent number of VMs
Impact	Services running in Prism Central may run out of memory and crash.
Resolution	1) Reduce the number of VM entities in Prism Central or 2) If using 1 node small Prism Central scale up to 1 node large Prism Central or scaleout to 3 node small Prism Central. 3) If using 1 node large Prism Central scaleout it to 3 node large Prism Central.
KB Article	4315
Severity	Warning

Table 136: Prism Central VM Limit Check [200304] [A200304]

Name	Prism Central VM Limit Check
Description	Checks if the number of VM entities is within the limit.
Alert message	The Prism Central cannot handle these many VM entities.
Cause	Too many VMs in Prism Central.
Impact	Prism Central performance may be degraded.
Impact	Prism Central may run out of disk space to store data.
Impact	Services running in Prism Central may run out of memory and crash.
Resolution	Reduce the number of VM entities in Prism Central.
KB Article	4296
Severity	Warning

Table 137: Prism Central cluster version check [200310]

Name	Prism Central cluster version check
Description	Checks if Prism Central is running GA version.
Cause	Non GA version detected on Prism Central.
Impact	Prism Central is running a Non GA version.
Resolution	Nutanix recommends running officially released GA versions of Prism Central software. For assistance refer to KB 6325
KB Article	6325

Table 138: Prism Central Version EOL [200311] [A200311]

Name	Prism Central version EOL check
Description	Checks if Prism Central is running an EOL version.
Alert message	*info_msg*. Upgrade Prism Central at the earliest.
Cause	EOL version running on Prism Central
Impact	Support may not be able to resolve issues since this version will no longer be maintained.
Resolution	Move to a recent released version immediately, to avoid disruption in support.You can refer to the EOL Schedule on the Nutanix support portal.
KB Article	6325
Severity	Info

Table 139: Inadequate Resource Configuration for the Prism Central VM [200327] [A200327]

Name	Prism Central Resource Resize Check

Description	Configured resource for the Prism Central VM is inadequate.
Alert message	*inadequate_resources* for the *pc_vm_type* Prism Central VM is inadequate on *non_resized_pc_vm_ips*.
Cause	The current feature capabilities of Prism Central require resource on the Prism Central VM to be increased for optimum performance.
Impact	Prism Central performance may be degraded.
Resolution	Increase the resource of the Prism Central VM. Please refer KB-8932 for guidance.
KB Article	8932
Severity	Warning

Table 140: Prism Central EXT4-fs Error Check [200329]

Name	Prism Central EXT4-fs Error Check
Description	Captures Prism Central EXT4-fs error messages
Cause	File system inconsistencies are present on the Prism Central VM node.
Impact	Inability for the Prism Central VM to boot or for the upgrade pre-checks to run.
Resolution	Look for any problems in the file system. Review KB 9622.
KB Article	9622

Table 141: Entity Sync Failure for the Protection Policy [500102] [A500102]

Name	Entity Sync failed for Protection Policy
Description	Entity Sync failed for the Protection Policy
Alert message	*message*
Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Severity	Warning

Table 142: Entity Sync Failure for the Recovery Plan [500103] [A500103]

Name	Entity Sync failed for Recovery Plan
Description	Entity Sync failed for the Recovery Plan
Alert message	*message*
Cause	Various
Impact	Data Protection and Recovery might be impacted.

Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Severity	Warning

Table 143: Entity Sync Failure for the Availability Zone [500104] [A500104]

Name	Entity Sync failed for Availability Zone
Description	Entity Sync failed for the Availability Zone
Alert message	*message*
Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Ensure Availability Zone is reachable and do a 'Forced Entity Sync' to ensure that all the entities are in sync. If the issue still persists, contact Nutanix support.
KB Article	10241
Severity	Critical

Table 144: Object Store deployment is not supported on this cluster. [805001] [A805001]

Name	Objects Store Deployment Check
Description	Checks if Object Store is deployed on an unsupported AOS cluster.
Alert message	The registered cluster(s) *affected_clusters_list* is/are on one of the versions between 5.18 and 5.18.0.4 which is unsupported for Object Store Deployment.
Cause	One or more Prism Element cluster(s) attached to this Prism Central is on version(s) between 5.18 and 5.18.0.4 which is unsupported for Object Store Deployment.
Impact	Object Store will not work as expected.
Resolution	Upgrade the cluster to 5.18.0.5 or above.
KB Article	10000
Severity	Critical

Controller VM

Table 145: CVM or Prism Central VM RAM Usage High [3023] [A1056]

Name	CVM Memory Usage
Description	Check that CVM or Prism Central VM memory usage is not high.
Alert message	Main memory usage in *vm_type ip_address* is high, *available_memory_kb* KB is free.
Cause	The RAM usage on the Controller VM or Prism Central VM has been high.
Impact	Cluster performance may be significantly degraded.

Resolution	Check the memory utilization of Prism Central VM or Controller VM. If abnormal behavior is seen, please collect logs and contact Nutanix Support.
CVM Available Memory Threshold	Critical Threshold: 589824 KB
KB Article	2473
Severity	Critical

Table 146: CVM/PCVM time not synchronized with external servers. [3026] [A3026]

Name	CVM/PCVM NTP Time Synchronized
Description	Checks that the {vm_type} is synchronizing time with an NTP server.
Alert message	*alert_msg*
Cause	External NTP servers are not configured or are not reachable
Impact	Workflows involving Kerberos may fail if the time difference between the Controller VM and the NTP server is greater than 5 minutes.
Resolution	Please verify that the external NTP servers are configured and are reachable.
KB Article	4519
Severity	Warning

Table 147: CVM Rebooted [3028] [A1024]

Name	CVM Rebooted Check
Description	Checks if the CVM is not rebooted recently.
Alert message	*vm_type* *ip_address* has been rebooted on *reboot_timestamp_str*.
Cause	CVM is rebooted.
Impact	During the time the Controller VM is down, cluster compute and storage capacity are reduced.
Resolution	Check CVM status.
KB Article	2474
Severity	Critical

Table 148: CVM Service(s) Restarting Frequently [3029] [A1032]

Name	CVM Services Status
Description	Check if services have restarted recently on the CVM.
Alert message	One or more cluster services have restarted within 15 minutes in the Controller VM *ip_address*, the total restarts being >= 10.The following are the service(s) that have restarted : *failed_components*.

Cause	Faulty behavior in one or more service(s), leading to frequent crashes,in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs more than once, please collect CVM logs and contact Nutanix support.
KB Article	2472
Severity	Warning

Table 149: CVM Python Service(s) Restarting Frequently [3037] [A3037]

Name	Python Services Fatal Check
Description	Check if python services have crashed/restarted recently in PCVM/CVM
Alert message	One or more cluster services have restarted within 15 minutes in the PCVM/CVM *ip_address*, the total restarts being >= 10. The following are the service(s) that have restarted : *failed_components*.
Cause	Faulty behavior in one or more service(s) leading to frequent crashes in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs more than once, please contact Nutanix support.
KB Article	2472
Severity	Critical

Table 150: Multiple cpu_unblock processes running [3042] [A3042]

Name	cpu_unblock Process Check
Description	Check that there are no stale cpu_unblock processes running.
Alert message	Multiple cpu_unblock processes are running on *svm_ip*.
Cause	Zookeeper restarting frequently on cluster node.
Impact	Cluster performance may be significantly degraded.
Resolution	Kill all cpu_unblock processes and restart cluster services on node.
KB Article	5581
Severity	Critical
Severity	Warning

Table 151: Prism Central VM CPU Load High [6517] [A6517]

Name	CVM CPU Load High

Description	Check that Prism Central VM CPU load is not high.
Alert message	CPU load in or Prism Central VM *ip_address* is high. Load average over 5 Minutes = #load_average_5min@, number of vCPUs in CVM or Prism Central VM = *num_procs*.
Cause	The CPU load on the Prism Central VM has been high.
Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix Support for diagnosis. CPU on the Prism Central VM may need to be increased.
KB Article	5227
Severity	Critical
Severity	Warning

Table 152: Prism Central VM disk usage high [101059] [A101059]

Name	Prism Central VM high disk space usage
Description	Check that disk space usage on the Prism Central VM is within limits.
Alert message	*alert_msg*
Cause	High disk usage in the Prism Central VM.
Impact	Prism Central VM may run out of storage space to store data.
Resolution	Refer to KB 5228 for further details.
Prism Central VM disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

Table 153: Core dumps are enabled on this CVM or Prism Central VM. [101069] [A101069]

Name	Saltstate coreoff check
Description	Checks if saltstate is correctly set to default coreoff state to ensure services do not produce core dumps unnecessarily.
Alert message	Core dumps are enabled on the *vm_type* with IP *service_vm_external_ip*.
Cause	Core dumps are enabled for services running on this CVM or Prism Central VM.
Impact	Core dumps are useful for isolated troubleshooting but leaving enabled may unnecessarily fill the node /home system partition, and this can affect node stability and availability.
Resolution	Follow KB7586 to validate and fix coreoff saltstate. Resolution may involve graceful node reboot, so ensure the cluster is otherwise healthy and resiliency status is OK prior to this action.

KB Article	7586
Severity	Critical

Table 154: Host time not synchronized with any external servers. [103090] [A103090]

Name	Host time synchronized
Description	Checks to ensure that the host is configured and synchronizing time with an NTP server.
Alert message	*alert_msg*
Cause	External NTP servers are not configured or are not reachable.
Impact	Logs may have different timestamps in the host and the CVMs/PCVM. Host may not work as expected.
Resolution	Please verify that the NTP servers are configured and are reachable from the host.
KB Article	4519
Severity	Warning

Table 155: CVM TCP checksum error rate check [103104]

Name	CVM TCP checksum error rate check
Description	Check if corrupted packets with compromised checksums are reaching the CVM and not dropped by the host NICs.
Cause	Hardware failure or NIC firmware/driver compatibility issue.
Impact	Cluster performance may be significantly degraded.
Resolution	Check for any hardware failures or issues with firmware/driver/cables.
Resolution	Please refer to KB 9623 for more details.
KB Article	9623

Table 156: Prism Central VM platform disk space usage high [200328] [A200328]

Name	Prism Central VM platform disk space usage check
Description	Check that platform disk space usage on the Prism Central VM is within limits.
Alert message	Platform disk space usage in Prism Central VM *svm_ip* exceeds *percentage_exceed*% for disk(s): *disk_paths*.
Cause	High disk usage in the Prism Central VM.
Impact	Prism Central VM may run out of storage space to store data.
Resolution	Refer to KB 5228 for further details.

Disk Usage Percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

Table 157: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed
Description	Programming a Flow rule failed.
Alert message	Flow policy operation failed for *rule_name*. *reason*
Cause	Microsegmentation service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.
Resolution	Check PC microsegmentation service, PE acropolis service, PC to PE connection, PE to AHV connection.
Severity	Warning

Table 158: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed
Description	Flow failure event.
Alert message	Flow operation failed on *remote_uuid*. *reason*
Cause	Microsegmentation service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microsegmentation service, Check PE acropolis service, AHV microsegmentation module.
Severity	Warning

Table 159: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	Flow mode change failed on *remote_uuid*. *reason*

Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect cvm_logs and contact Nutanix Support
Severity	Warning

DR

Table 160: Availability Zone configured in Protection Policy is not accessible [110401] [A110401]

Name	Availability Zone Accessibility Check
Description	Check if availability zones configured in protection policy are accessible.
Alert message	**Availability Zone(s)** *fault_azs* **configured in the Protection Policy** *protection_rule_name* **is/are not accessible.**
Cause	Availability zone configured in the protection rule is not accessible.
Impact	Availability zones of the backup site will not be reachable for taking backup.
Resolution	Check the health status of the configured availability zone.
KB Article	5685
Severity	Critical

Table 161: Protection Policy Max VMs per Category Check Failed. [110402] [A110402]

Name	Protection Policy Max VMs Per Category Exceeded
Description	Checks if the VM count for a category specified in Protection Policy exceeds the maximum allowed limit.
Alert message	**Maximum number of VMs for a category in a Protection Policy should not exceed** *max_vm_count*. **Following categories exceeds VMs limit :** *categories*.
Cause	Number of VMs for the specified categories in the Protection Policy exceeds the limit.
Impact	Specified category will not be considered for the recovery as the Recovery Plan supports categories with limited number of VMs.
Resolution	Reduce the protected VM count for the specified categories in the Protection Policy.
KB Article	6306
Severity	Warning

Table 162: Disconnected Availability Zones are affecting some entities. [110403] [A110403]

Name	Entities Affected By Disconnected Availability Zones
Description	Lists all entities which will be affected due to the disconnected Availability Zones.
Alert message	*entity_type* of name *name* has been affected by disconnected availability zones.
Cause	Availability zone configured in the protection rule has been disconnected.
Impact	Operations realted to the affected entities will be affected.
Resolution	Check the health status of the configured availability zone.
KB Article	6283
Severity	Info

Table 163: Subnet Creation Failure for the Recovery Plan [300405] [A300405]

Name	Subnet creation failure for Recovery Plan
Description	Subnet creation failed for the Recovery Plan.
Alert message	Subnet *network_cidr* creation in the virtual network *atlas_virtual_network_name* failed for the *operation_type* of the Recovery Plan *recovery_plan_name* executed through the recovery plan job *recovery_plan_job_name*.
Cause	Networks with conflicting CIDR might be present.
Impact	Some VMs will be recovered without vNICs.
Resolution	Resolve the issue. If you cannot resolve the issue, contact Nutanix support.
Severity	Critical

Table 164: Invalid Network Settings for the Recovery Plan [300407] [A300407]

Name	Recovery Plan Periodic Check for Invalid Network Settings
Description	Network settings have Availability Zones that are not in accordance with Availability Zone order list of Protection Policy applied to VMs in Recovery Plan.
Alert message	Invalid network settings found for the Recovery Plan *recovery_plan_name*. *reason*.
Cause	One or more Availability Zones listed in Protection Policy do not match the Availability Zones specified in the Recovery Plan network settings.
Impact	Recovery Plan will not be synced to Availability Zones defined in Protection Policy.
Resolution	Update Availability Zones in the Recovery Plan network settings to be in accordance with Availability Zone order in the Protection Policy.
Severity	Warning

Table 165: Recovery plans have conflicting network mappings [300410] [A300410]

Name	Recovery Plan network mapping conflict check
Description	Checks network mapping conflicts in multiple Recovery Plans.
Alert message	Recovery plans #rp_1@ and #rp_2@ have different source networks mapped to same destination network *target_network*
Cause	Multiple Recovery Plans have different source networks mapped to same destination network.
Impact	VMs isolated on source network might be able to reach each other on the destination after recovery.
Resolution	Avoid mapping multiple networks to same network in one or more Recovery Plans.
KB Article	7175
Severity	Info

Table 166: Test Failover on Recovery Plan has not been executed recently [300412] [A300412]

Name	Recovery Plan Last Test Failover run status check
Description	Checks when last Test Failover was run on the Recovery Plan.
Alert message	*time* days have elapsed since a Test Failover was run on Recovery Plan *recovery_plan_name*.
Cause	Test Failover on the Recovery Plan has not been executed.
Impact	There might be issues in the recovery of VMs, which may be difficult to identify if the Test Failover is not run periodically.
Resolution	Run a Test failover on the Recovery Plan periodically, to ensure Recovery Plan has no errors.
KB Article	5793
Severity	Info

Table 167: Multiple Recovery Plans associated with a Category [300413] [A300413]

Name	Same Category in multiple Recovery Plans check.
Description	Checks if same category belongs to multiple Recovery Plans.
Alert message	A specific category should be associated with only one Recovery Plan. Recovery Plans : *rp_list* are associated with the same category *category*.
Cause	Multiple Recovery Plans have Same Category
Impact	Execution of the following Recovery Plans will be affected and only one of those Recovery Plan will be able to recover VMs from specified category.
Resolution	Update the Recovery Plans such that the category is unique between them.
KB Article	5669
Severity	Warning

Table 168: Number of VMs in Recovery Plan exceeds the threshold [300414] [A300414]

Name	Recovery Plan VM Limit Check
Description	Checks if the VM count exceeds the threshold in Recovery Plan.
Alert message	Maximum number of VMs in a recovery plan should not exceed *max_vm_count*. Recovery Plan *recovery_plan_name* have *vm_count* VMs.
Cause	Number of VMs in the Recovery Plan exceeds the limit.
Impact	VM Recovery prone to failure.
Resolution	Reduce the number of VMs in the Recovery Plan.
KB Article	5717
Severity	Warning

Table 169: VMs are part of multiple stages in Recovery Plan [300415] [A300415]

Name	VM in Multiple Stages Check
Description	Checks whether a VM is part of more than one stage in Recovery Plan.
Alert message	A VM should be part of only one stage in a Recovery Plan. Recovery plan *recovery_plan_name* has entities in multiple stages. *duplicate_entities*
Cause	VMs are part of multiple stages in the Recovery Plan.
Impact	VM specified in the first stage will be considered for recovery, later stages will be skipped.
Resolution	Update the Recovery Plan to ensure that the VMs are not part of multiple stages.
KB Article	5715
Severity	Info

Table 170: The same floating IP is associated with multiple VMs belonging to different Recovery Plans [300416] [A300416]

Name	Conflicting Floating IP Check
Description	Checks if VMs which are part of different Recovery Plan have same Floating IPs.
Alert message	Same Floating IPs should not be assigned to multiple VMs and should not be part of multiple Recovery Plans. Floating IP *floating_ip* is assigned to VMs *alert_msg*.
Cause	VMs belonging to different Recovery Plans are assigned the same Floating IP.
Impact	Floating IP assignment post VM recovery may fail.
Resolution	Update the Recovery Plans to ensure that a Floating IP address is to be assigned to only one VM.
KB Article	5797
Severity	Warning

Table 171: Some of the VMs in the Recovery Plan are unprotected [300417] [A300417]

Name	Recovery Plan VMs Protection Policy Check
Description	Checks if VMs configured in the Recovery Plan are protected.
Alert message	VMs *vms* in the Recovery Plan *recovery_plan_name* are not protected.
Cause	VMs in the Recovery Plan are not protected.
Impact	VMs present in the Recovery Plan can't be recovered.
Resolution	Update associated VMs in Recovery Plan to have Protection Policy.
KB Article	5795
Severity	Warning

Table 172: VMs not cleaned up following the Test Failover for Recovery Plan [300418] [A300418]

Name	Recovery Plan Checks for Cleanup of VMs after Test Failover
Description	Checks that VMs recovered using Recovery Plan after Test Failover have been cleaned up.
Alert message	VMs recovered following the Test Failover of Recovery Plan *recovery_plan_name* have not been cleaned up. Manual clean up of the VMs is required.
Cause	VMs recovered following the Test Failover of the Recovery Plan have not been cleaned up.
Impact	Subsequent recovery of the VMs from Test Failover operation might get affected.
Resolution	Perform cleanup of the test VMs for the Recovery Plan.
KB Article	5794
Severity	Warning

Table 173: Subnets Deletion Failure for the Recovery Plan [300419] [A300419]

Name	Dynamic Subnets deletion failure for Recovery Plan
Description	Deletion of dynamic subnets failed.
Alert message	Deletion of the subnets *networks_list* failed for the *operation_type* of the Recovery Plan *recovery_plan_name* executed through the *recovery_plan_job_name* due to *reason*.
Cause	Some of the VMs might be using the subnets.
Impact	Unused subnets will be billable.
Resolution	Resolve the issue. If you cannot resolve the issue, contact Nutanix Support.
Severity	Critical

Table 174: Recovery Plan has multiple Availability Zone Orders [300424] [A300424]

Name	Recovery Plan Multiple Availability Zone Order Check
Description	Checks if the Recovery Plan has any entity protected to multiple Availability Zones.
Alert message	Recovery Plan *recovery_plan_name* has entities protected to multiple Availability Zones.
Cause	Recovery Plan contains entities protected to multiple Availability Zones.
Impact	The Recovery Plan update will not be allowed.
Resolution	Update associated entities in Recovery Plan to have single Availability Zone Order.
KB Article	6591
Severity	Warning

Table 175: Recovery Plan contains VMs with Unsupported CHDR VM configuration. [300425] [A300425]

Name	Recovery Plan CHDR VM Config Check
Description	Checks if Recovery Plan contains VMs with Unsupported CHDR VM configuration.
Alert message	Recovery Plan *recovery_plan_name* contains VMs *vm_list* with Unsupported CHDR VM configuration.
Cause	Recovery Plan contains VMs with unsupported configuration for Cross hypervisor disaster recovery.
Impact	The VMs will not be recovered.
Resolution	Install NGT on all the VMs in the Recovery Plan, ensure that you do not have UEFI boot, and have no SATA/PCI or delta disk attached.
KB Article	6835
Severity	Warning

Table 176: Incompatible Recovery Availability Zones for some VMs in the Recovery Plan [300426] [A300426]

Name	Recovery Plan Incompatible Recovery Availability Zone Check
Description	Checks if Recovery Plan contains VMware VMs and snapshots for these VMs are replicated to a recovery Availability Zone that doesn't support recovery of VMware VMs.
Alert message	Incompatible Recovery Availability Zones for Recovery Plan *recovery_plan_name*. Recovery of VMs *incompatibles_vms* on Recovery Availability Zones *incompatbile_target_availability_zone_names* will fail.
Cause	Target Availability Zone is running Prism Central version less than AOS 5.11 version and hence it does not support the recovery of VMware VMs.
Impact	The VM recovery will fail.

Resolution	Upgrade the Target Availability Zone to 5.11 or later version, or Remove the entity mentioned in the description of the alert from the Recovery plan.
KB Article	6872
Severity	Warning

Table 177: Recovery Plan contains VMs with special configuration and no empty CDROM [300428] [A300428]

Name	Recovery Plan NGT Config Empty CD ROM Check
Description	Checks if VMs part of the recovery plan have an empty CDROM available to mount nutanix guest tools, so NGT reconfiguration tasks can be completed on restore.
Alert message	Recovery Plan *recovery_plan_name* contains VMs *vm_list* with special configuration that require NGT reconfiguration on restore. Nutanix Guest tools cannot be mounted as there is no empty CD ROM slot on the VM.
Cause	VMs on Source Availability Zone do not have a free CDROM slot.
Impact	VM will be recovered however post recovery VM reconfiguration tasks, including NGT recofiguration, static IP reconfiguration will fail.
Resolution	Add an empty CDROM to the VM.
KB Article	7054
Severity	Warning

Table 178: Recovery Plan Execution Exceeded the Time Limit. [300429] [A300429]

Name	Recovery Plan Execution Exceeded the Maximum Expected Time Limit.
Description	Recovery Plan execution exceeded the time limit.
Alert message	Execution of Recovery Plan '*recovery_plan_name*' exceeded the estimated completion time by *excess_time_taken*.
Cause	Restore or Power On operation may have taken longer time.
Impact	Recovery Plan execution is prone to failure.
Resolution	Contact Nutanix customer service.
Severity	Warning

Table 179: VM(s) in category(s) present in multiple fault domains. [300432] [A300432]

Name	Category Protected VMs Multiple Fault Domain Check
Description	Checks if VM(s) protected by category(s) in Recovery Plan are present in multiple fault domains.
Alert message	*alert_message*
Cause	There may be multiple reasons for this issue.

Impact	Selection of VMs for recovery will be affected.
Resolution	Ensure categories present in Recovery Plan contain VMs from single fault domain only.
KB Article	9472
Severity	Warning

License

Table 180: Calm Trial License Expiry [14007] [A14007]

Name	NucalmTrialLicenseOver
Description	Nutanix Calm Trial License Expired
Alert message	Nutanix Calm 60 days Trial License *expiry_remaining_days*
Cause	Nutanix Calm Trial License of 60 days is set to expire or has expired.
Impact	The cluster is out of licensing compliance.
Resolution	Please upgrade your Calm license with additional license packs. If you cannot resolve the issue, please contact Nutanix support.

Guest VM

Table 181: VM {vm_name} Memory Overprovisioned [120240] [A120240]

Name	vm_memory_overprovisioned_alert
Description	Memory Over-provisioned VM
Alert message	*alert_message*
Cause	VM memory is over-provisioned.
Impact	Inefficient use of memory
Resolution	Reduce memory allocation for the VM.

Table 182: VM {vm_name} Memory Constrained [120241] [A120241]

Name	vm_memory_constrained_alert
Description	Memory constrained VM
Alert message	*alert_message*
Cause	VM is experiencing high memory usage.
Impact	Inefficient use of memory
Resolution	Add more memory resources.

Table 183: VM {vm_name} CPU Overprovisioned [120242] [A120242]

Name	vm_cpu_overprovisioned_alert
Description	CPU Overprovisioned VM
Alert message	*alert_message*
Cause	VM CPU is over-provisioned.
Impact	Inefficient use of CPU
Resolution	Reduce CPU allocation for the VM.

Table 184: VM {vm_name} CPU Constrained [120243] [A120243]

Name	vm_cpu_constrained_alert
Description	CPU constrained VM
Alert message	*alert_message*
Cause	VM is experiencing high CPU ready time.
Impact	Inefficient CPU usage
Resolution	Add more CPU resources.

Table 185: VM {vm_name} Inactive [120244] [A120244]

Name	inactive_vm_alert
Description	Inactive VM
Alert message	*alert_message*
Cause	VM is inactive.
Resolution	Activate the VM.

Table 186: VM Bully {vm_name} [120245] [A120245]

Name	bully_vm_alert
Description	Bully VM
Alert message	*alert_message*
Cause	VM is a bully.
Resolution	Troubleshot why this VM is a bully.

Node

Table 187: Latency between CVMs is high [6005] [A6005]

Name	Inter-CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	Latency between *vm_type*s is higher than *ping_threshold* ms.
Cause	Malfunctioning or misconfigured network infrastructure
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.
Ping Time Threshold	Critical Threshold: 15 seconds
KB Article	3690
Severity	Critical

Table 188: SSH keys missing [111089] [A111089]

Name	Node Missing SSH Keys check
Description	Check if /home/nutanix/ssh_keys directory contents are missing.
Alert message	*alert_msg*
Cause	Some of SSH keys are missing from /home/nutanix/ssh_keys directory.
Impact	Missing SSH keys can jeopardise ongoing infrastructure workflows in a silent manner such as node un-configuration.
Resolution	Please refer to KB 10736 for more information.
KB Article	10736
Severity	Critical

Table 189: Home partition is not expanded to 50GB. [200330] [A200330]

Name	Prism Central home partition expansion check
Description	Checks if home partition is expanded to 50GB after upgrade.
Alert message	Home partition is not expanded to 50GB on Prism Central *pc_ip* after upgrade.
Cause	During an upgrade, home partition extension has failed.
Impact	Prism Central performance may be degraded.
Resolution	Refer to KB 9525 for further details.
KB Article	9525
Severity	Warning

Network

Table 190: Latency between CVMs is high [6005] [A6005]

Name	Inter-CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	Latency between *vm_type*s is higher than *ping_threshold* ms.
Cause	Malfunctioning or misconfigured network infrastructure
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.
Ping Time Threshold	Critical Threshold: 15 seconds
KB Article	3690
Severity	Critical

Table 191: CVM TCP checksum error rate check [103104]

Name	CVM TCP checksum error rate check
Description	Check if corrupted packets with compromised checksums are reaching the CVM and not dropped by the host NICs.
Cause	Hardware failure or NIC firmware/driver compatibility issue.
Impact	Cluster performance may be significantly degraded.
Resolution	Check for any hardware failures or issues with firmware/driver/cables.
Resolution	Please refer to KB 9623 for more details.
KB Article	9623

Table 192: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed
Description	Programming a Flow rule failed.
Alert message	Flow policy operation failed for *rule_name*. *reason*
Cause	Microsegmentation service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.
Resolution	Check PC microsegmentation service, PE acropolis service, PC to PE connection, PE to AHV connection.
Severity	Warning

Table 193: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed

Description	Flow failure event.
Alert message	Flow operation failed on *remote_uuid*. *reason*
Cause	Microsegmentation service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microsegmentation service, Check PE acropolis service, AHV microsegmentation module.
Severity	Warning

Table 194: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	Flow mode change failed on *remote_uuid*. *reason*
Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect cvm_logs and contact Nutanix Support
Severity	Warning

Other

Table 195: Detected recovery points on prism central managing clusters with high density nodes [110452] [A110452]

Name	High density nodes recovery points check
Description	Checks if recovery points are detected on prism central managing clusters with high density nodes
Alert message	Recovery points are detected on prism central managing clusters with high density nodes
Cause	Detected recovery points on prism central managing clusters with high density nodes
Impact	Recovery points are not currently supported on clusters with high density nodes. They may impact cluster operations
Resolution	Delete recovery points on prism central hosted by clusters with high density nodes
KB Article	7003

Severity	Info

Table 196: Detected protection policies on Prism Central managing clusters with high density nodes [110453] [A110453]

Name	High density nodes protection rules check
Description	Checks if protection policies are detected on Prism Central managing clusters with high density nodes
Alert message	Protection policies are detected on Prism Central managing clusters with high density nodes
Cause	Protection policies are detected on a Prism Central managing clusters with high density nodes
Impact	Data protection is currently not supported for clusters with high density nodes, this may impact cluster operations
Resolution	Delete protection policies configured on Prism Central managing clusters with high density nodes
KB Article	7003
Severity	Info

Table 197: Memory configuration inconsistent. [200306] [A200306]

Name	Prism Central VM same memory level check
Description	Check all Prism Central VMs have the same memory level.
Alert message	The Prism Central VMs are not configured to have the same amount of memory.
Cause	Memory configuration among Prism Central VMs is inconsistent.
Impact	The Prism Central VM will not perform at the level necessary to manage the cluster.
Resolution	Fix Prism Central VM memory configuration by providing sameamount of memory for all the Prism Central VMs
KB Article	5377
Severity	Warning
Severity	Info

Table 198: Prism Central VM type or annotation not set. [200307] [A200307]

Name	PCVM type and annotation check
Description	Check the type and annotation of all PCVMs are set.
Alert message	The type or annotation of Prism Central VMs is not set
Cause	Prism Central VMs do not have annotation set.
Impact	Guardrails to avoid potentially disruptive user actions on a Prism Central will not be applied.

Resolution	Set the annotation for the Prism Central VM. Refer to KB 6027.
KB Article	6027
Severity	Warning

Table 199: Unequal metadata partition size(s) across Prism Central VMs. [200308] [A200308]

Name	Prism Central VM same disk size check
Description	Checks if metadata partition size(s) are the same among Prism Central VMs.
Alert message	The metadata partitions across Prism Central VMs are not consistent due to mismatch in sizes: *disks_info*
Cause	Metadata partition size(s) of one or more Prism Central VMs is inconsistent with other Prism Central VMs.
Impact	Manageability of the cluster may be lost. Cluster may be significantly degraded and ultimately fail.
Resolution	Contact Nutanix support to resize the Prism Central VM disks.
KB Article	6028
Severity	Warning

Table 200: High time difference between Prism Central and registered Prism Elements. [200309] [A200309]

Name	PC-PE time drift check
Description	Checks ntp sync between the PC and the registered PEs.
Alert message	Time differs between Prism Central and registered Prism Elements.
Cause	Time is not synchronized between Prism Central and registered Prism Elements.
Impact	Stale statistics or configuration may appear in Prism Central.
Impact	Statistics may not be available in Prism Central.
Impact	Components that make use of the statistics or configuration in Prism Central may not function as expected.
Resolution	Ensure that NTP is configured properly and that the cluster has connectivity with the NTP server. Refer to KB 6121.
KB Article	6121
Severity	Warning

Table 201: PE-PC incompatible AOS versions. [200312] [A200312]

Name	PC-PE incompatible versions
Description	Checks if the Prism Element and Prism Central AOS versions are compatible.

Alert message	The AOS version of the registered cluster is incompatible with the Prism Central version.
Cause	The AOS version of the registered cluster is incompatible with the Prism Central version.
Impact	Prism Central features will not work as expected for this cluster.
Resolution	Upgrade the AOS version of the registered cluster to a version compatible with this Prism Central.
KB Article	6323
Severity	Warning

Table 202: Prism Central VM active upgrade check [200313]

Name	Prism Central VM active upgrade check
Description	Checks if there are any active upgrades
Cause	An upgrade is in process for this entity
Impact	Maintenance activities will have to wait until upgrade is finished.
Resolution	Wait until the upgrade is completed before doing any other actions.Contact Nutanix support for assistance.
KB Article	6644

Table 203: Prism Central Upgrades are disabled on {cvm_ip} [200314] [A200314]

Name	Prism Central VM upgrades are disabled check
Description	Checks if upgrades are disabled on Prism Central VM
Alert message	Prism Central Upgrades are disabled on *cvm_ip*
Cause	Prism Central Upgrade has been paused manually.
Impact	Upgrades from Prism won't proceed.
Resolution	Upgrades have been disabled and need to be re-enabled.Contact Nutanix support for assistance.
KB Article	6644
Severity	Info

Table 204: Recovery Plan has conflicting IP mappings with other Recovery Plan(s) [300430] [A300430]

Name	Recovery Plan IP mappings check
Description	Checks for conflicting IP mappings across multiple Recovery Plans for the paired AZs.
Alert message	Recovery Plan *rp* has conflicting IP mappings with other Recovery Plans. *conflicting_ips* *conflicting_vms*

Cause	IPs mapped in Recovery Plans are conflicting.
Impact	IP mapping will fail if a VM has been recovered with the same IP from a different Recovery Plan.
Impact	VM will come up without the specified IP after recovery.
Resolution	Ensure that the same IP from a subnet is not specified for multiple VMs.
KB Article	7866
Severity	Warning

Table 205: Dataservice IP is unreachable [400105] [A400105]

Name	Data service IP connectivity check
Description	Data service IP is not reachable
Alert message	Data service ip '*data_service_ip*' is unreachable from prism central *node_name*
Cause	Invalid or empty data service ip
Impact	Nutanix Calm containers will be inaccessible
Resolution	Please provide correct data service ip. Review KB 5199
KB Article	5199
Severity	Critical

Table 206: Karan service(s) are unreachable [400106] [A400106]

Name	Karan service connectivity check
Description	Karan service(s) are not reachable
Alert message	Karan service on '*karan_service_ip*:*karan_service_port*' is unreachable from prism central *node_name*
Cause	Invalid or empty karan service ip and port
Impact	Nutanix Calm windows automation will fail
Resolution	Please provide correct karan service ip and port. Review KB 5209.
KB Article	5209
Severity	Critical

Table 207: Calm Containers are unhealthy [400112] [A400112]

Name	Calm containers health check
Description	Check for Calm Container's state
Alert message	Nutanix Calm or Epsilon container is in unhealthy state
Cause	Internal services of each calm containers are down, Docker plugin is not working properly

Impact	Nutanix Calm services may be inaccessible or performs incorrectly.
Resolution	Please check internal services of each calm containers.
KB Article	7343

Prism Central VM

Table 208: XPilot Retry Approval Pending [111081] [A111081]

Name	XPilot Retry Approval Pending
Description	XPilot Playbook is waiting for Approval to retry
Alert message	XPilot Playbook *action_rule_name* is awaiting approval to begin retry no. *retry_count*
Cause	XPilot Playbook is waiting for approval to begin retry
Impact	Playbook will fail after wait period of 7 days
Resolution	Resume playbook to give approval. Refer to KB 9961 for further details
KB Article	9961
Severity	Info

Storage

Table 209: Prism Central VM System Root Partition Space Usage High [200316] [A200316]

Name	Prism Central VM System Root Partition Usage
Description	Checks if Prism Central system root partition usage is within threshold to ensure uninterrupted operations.
Alert message	Disk space usage for root partition *mount_path* on *entity-ip_address* has exceeded *threshold%*.
Cause	Increased Prism Central VM system root partition usage due to excessive logging or incomplete maintenance operation.
Impact	If Prism Central VM system root partition is highly utilised then certain maintenance operations, such as upgrades, may be impacted. If Prism Central VM system root partition is 100% utilized then services may stop and impact Prism Central cluster management functions.
Resolution	Reduce Prism Central VM system root partition usage by removing any known temporary or unneeded files. Refer to KB 5228 for further details.
Prism Central VM System Root Partition Usage Threshold	Critical Threshold: 95 %
KB Article	5228
Severity	Critical

Table 210: Prism Central VM home partition disk usage high [200317] [A200317]

Name	Prism Central VM home partition disk usage
Description	Checks if Prism Central home partition usage is within threshold to ensure uninterrupted operations.
Alert message	*alert_msg*
Cause	Increased Prism Central VM home partition usage due to excessive logging or incomplete maintenance operation.
Impact	If Prism Central VM home partition is highly utilised then certain maintenance operations, such as upgrades, may be be impacted. If Prism Central VM home partition is 100% utilised then services may stop and impact cluster storage availability.
Resolution	Reduce Prism Central VM home partition usage by removing any known temporary or unneeded files. Refer to KB 5228 for further details.
Prism Central VM home partition disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

System Indicator

Table 211: Tomcat is restarting frequently. [700101] [A700101]

Name	Tomcat Frequent Restart
Description	Tomcat is restarting frequently.
Alert message	Tomcat is restarting frequently on *cvm_ip*. Reason: *message*.
Cause	Tomcat is restarting frequently, please check prism, prism_monitor and catalina logs for exact cause of failure.
Impact	Performance and availability impact for Prism
Resolution	Check the logs for error information or contact Nutanix support.
KB Article	8524
Severity	Critical

6

SYSTEM MANAGEMENT

You can configure various system settings for Prism Central.

- You can configure an HTTP proxy, which can be used to communicate with a Nutanix service center (see Configuring an HTTP Proxy on page 290
- You can specify one or more name servers (see Configuring Name Servers on page 296).
- You can specify one or more NTP servers for setting the system clock (see Configuring NTP Servers on page 296).
- You can configure SNMP (see Configuring SNMP on page 298).
- You can configure a login banner page (see Configuring a Banner Page on page 311).
- You can change the language setting for displayed text (see Changing the Language Settings (Prism Central) on page 312).

Configuring Prism Central When a Cluster Uses Proxy Servers

Note:

- Prism Central and its managed clusters are not supported in environments deploying Network Address Translation (NAT).
- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see Prism Central Proxy for Pulse Data on page 659.

To communicate with the Nutanix service center directly, you need to configure the HTTP proxy. See Configuring an HTTP Proxy on page 290 for information on how to configure the HTTP proxy through Prism Central web console. After the HTTP proxy is configured, all the HTTP or HTTPS communication initiated by Prism Element or the Prism Central is routed through the proxy server.

The communication from Prism is routed through the proxy server until a whitelist target entry indicates otherwise. See Who Needs to Use the Whitelist Method on page 289 for information on who needs to whitelist the IP addresses. To overcome a communication failure between a Prism Element and the registered Prism Central, you need to manually whitelist the target entries from the Prism web console or nCLI. These options enable you to add Prism Central and its managed/registered clusters to a whitelist, where any HTTP proxy settings are ignored. This configuration allows network traffic between them, bypassing any proxy servers configured in the cluster. The whitelist also enables you to register new clusters with Prism Central successfully where clusters are using an HTTP proxy.

See *Step 3* of Configuring an HTTP Proxy on page 290 for information on how to configure the whitelist IP address through the Prism Central web console.

Alternatively, you can use the **ncli http-proxy add-to-whitelist** and **ncli http-proxy delete-from-whitelist** nCLI command options. See Whitelisting Prism Central and Its Managed Clusters (nCLI) on page 295 for information on how to configure the whitelist IP address through the nCLI.

- You can add or delete one whitelist entry at a time.
- Each whitelist entry cannot exceed 253 characters.
- A maximum of 1000 whitelist entries are supported.
- When deleting an entry from a whitelist, delete the target, not the target type.
- The commands do not support the IPv4 network mask network prefix or * (asterisk) prefix notation. You can specify a subnet and netmask, such as 10.0.0.0/255.0.0.0 or 192.168.1.0/255.255.255.0. This results in adding the specified subnet to the whitelist.
- When applying a whitelist, domain names like **contoso.com** are not processed the same as **www.contoso.com** and are treated as separate, distinct entities.
- Use fully qualified domain names to whitelist.

Who Needs to Use the Whitelist Method

SSL port 9440 needs to be open in both directions between the Prism Central VM and any registered clusters or clusters to be registered. For the complete list of required ports, see Port Reference.

If you are implementing a proxy server in your cluster environment with this port open as shown in this simple graphic, you do not need to allow Prism Central and its managed/registered clusters.

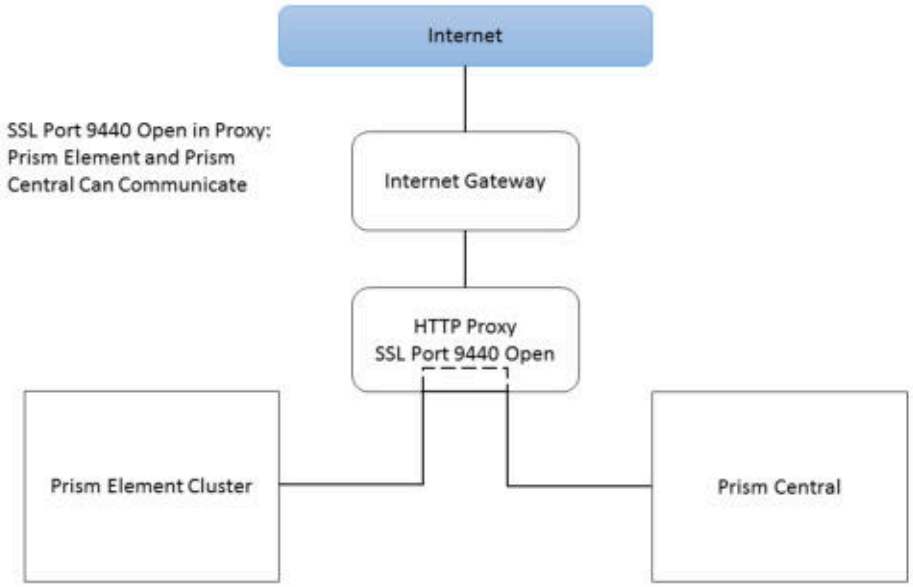

Figure 167: Prism Central, Proxy Server with SSL Port 9440 Open

If you are implementing a proxy server in your cluster environment with this port closed as shown in this simple graphic, you must allow direct communication between Prism Central and its managed/registered clusters as described in *Step 3* of Configuring an HTTP Proxy on

page 290 or alternatively in Whitelisting Prism Central and Its Managed Clusters (nCLI) on page 295.

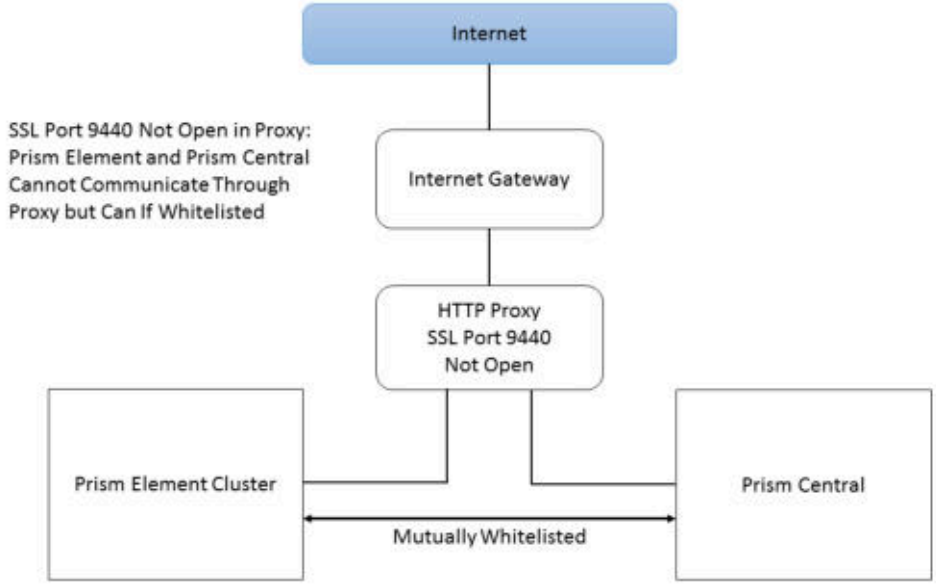

Figure 168: Prism Central, Proxy Server with SSL Port 9440 Closed and with Whitelisting

Configuring an HTTP Proxy

About this task

Note:

- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see Prism Central Proxy for Pulse Data on page 659.

About this task

If Prism Central cannot send traffic to a Nutanix Service center directly, an HTTP proxy is required. To configure an HTTP proxy, do the following:

Procedure

1. Click the gear icon in the main menu and then select **HTTP Proxy** in the **Settings** page. The HTTP Proxies dialog box appears.

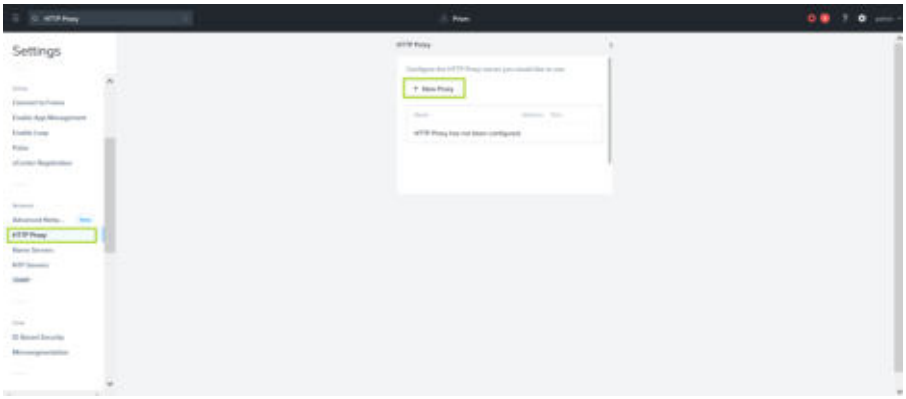

Figure 169: HTTP Proxy Window

2. To add an HTTP proxy, click the **New Proxy** button and do the following in the displayed fields:

   > Note: You can configure only one HTTP proxy at a time. If one exists currently, you must first delete it before creating a new one.

   a. **Name**: Enter a proxy server name.

   b. **Address**: Enter an IP address or host name for the proxy server.

   c. **Port**: Enter the port number to use.

   d. **Username**: Enter a user name.

   e. **Password**: Enter a password.

   f. **Protocols**: Select (check) the protocol to use, **HTTP**, **HTTPS**, or both.

      Select **HTTP** to route all HTTP requests through the proxy.

      Select **HTTPS** to route all HTTPs requests through the proxy.

      Select **HTTP** and **HTTPS** to route both HTTP and HTTPS requests through the proxy.

   g. When all the fields are correct, click the **Save** button (lower right).

      This saves the configuration and redisplays the dialog box with the new HTTP proxy entry appearing in the list.

      > Note: To return to the HTTP Proxy window without saving, click the **Cancel** button.

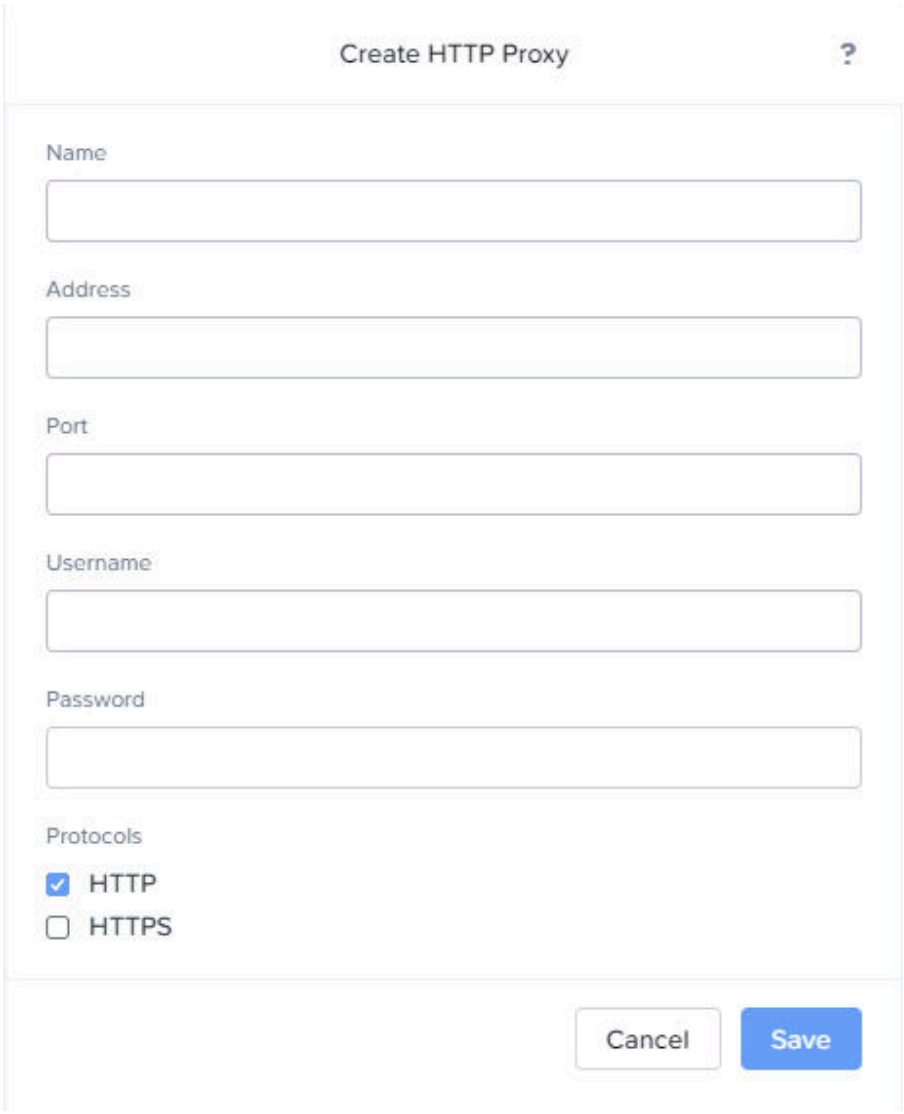

Figure 170: Create HTTP Proxy Window

3. To edit an HTTP proxy entry, click the pencil icon on the line for that entry, update one or more of displayed field entries as desired, and then click the **Save** button.

   The Update HTTP Proxy dialog box appears with the same fields as the Create HTTP Proxy dialog box plus the option (below the protocol check boxes) to add whitelist entries.

   To configure HTTP proxy whitelist entries, do the following:

   - To add a whitelist target, click the **+ Create** link. This step opens a line to enter a target address. Enter the target IP address or the network address and the subnet mask and then click the **Save** link in that field.

     Prism Central sends traffic to the whitelist IP addresses directly rather than through the HTTP or HTTPS proxy.
   - To edit a whitelist target, click the pencil icon for that target and update as needed.
   - To delete a whitelist target, click the X icon for that target.

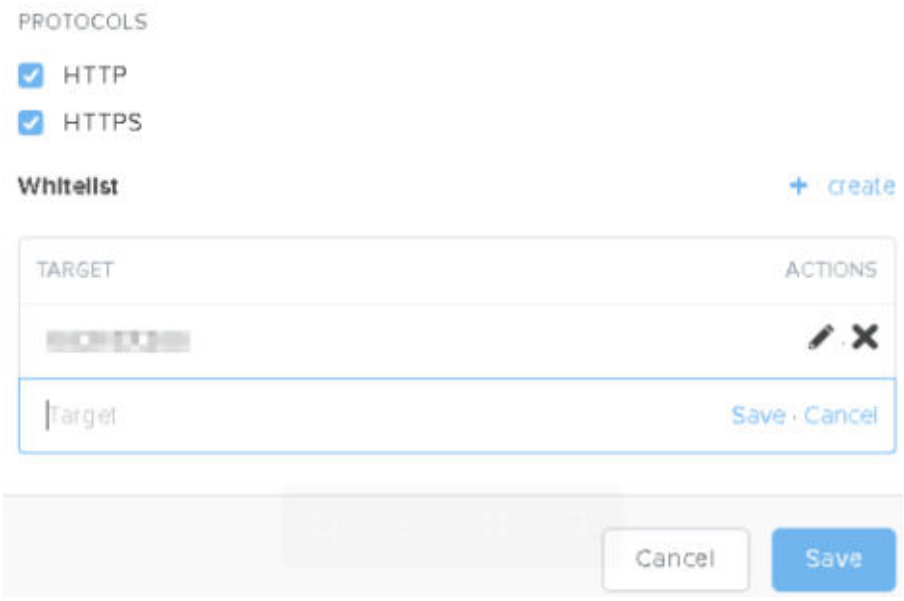

Figure 171: Whitelist Targets

> Note: Where proxy whitelist configuration is needed, it is recommended that you add the cluster virtual IP and all external CVM IPs of each registered cluster to the proxy whitelist for Prism Central. The virtual IP of Prism Central and the IP of each Prism Central VM (PCVM) should be added to the proxy whitelist on the registered Prism Element clusters. For easier management, you can add the subnets instead of adding the individual IPs.

4. To delete an HTTP proxy entry, click the X icon for that entry.

   A window prompt appears to verify the action; click the **OK** button. The entry is removed from the HTTP proxy list.

Whitelisting Prism Central and Its Managed Clusters (nCLI)

About this task

In this example, bypass a proxy server used by a managed Prism Element cluster and allow network traffic between Prism Central and the cluster. Previously, if you attempted to register a cluster that implemented a proxy server, the registration failed.

Procedure

1. Open a SSH session to any Controller VM in the cluster to be managed by Prism Central.
2. In this example, add the Prism Central VM IP address to the whitelist, then ensure the Prism Central VM IP address was added to the whitelist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address target=10.4.52.40
```

```
nutanix@cvm$ ncli http-proxy get-whitelist
Target Type : IPV4_ADDRESS
Target : 10.4.52.40
```

   Note: Repeat this step for additional Prism Central VM IP addresses and the virtual IP address in the case of scaleout Prism Central.

3. Open a SSH session to the Prism Central VM managing the cluster where you just modified the HTTP whitelist.
4. Add the cluster virtual IP address to the whitelist, then ensure the IP address was added to the whitelist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address target=10.4.52.10
```

```
nutanix@cvm$ ncli http-proxy get-whitelist
Target Type : IPV4_ADDRESS
Target : 10.4.52.10
```

   Note:

   - Repeat this step for individual CVM IP addresses.
   - For large clusters or for multiple clusters in a network it can be more efficient to add a network to the http-proxy whitelist. For example,

     The following is a sample command to add a network to the http-proxy whitelist.

     ```
     nutanix@PCVM:~$ ncli http-proxy add-to-whitelist target-type=ipv4_network_mask
      target=10.4.52.0/255.255.255.0
     ```

     The following is sample output of the **get-whitelist** command, which displays the whitelisted network added above.

     ```
     nutanix@CVM:~$ ncli http-proxy get-whitelist

     Target Type : IPV4_NETWORK_MASK Target : 10.4.52.0/255.255.255.0
     ```

   In this case, Prism Central and its managed cluster can communicate, with network traffic bypassing any proxy servers configured in the cluster.

Configuring Name Servers

About this task

Name servers are computers that host a network service for providing responses to queries against a directory service, such as a DNS server. To add (or delete) a name server, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Name Servers** in the **Settings** page. The Name Servers dialog box appears.

Figure 172: Name Servers Window

2. To add a name server, enter the server IP address in the **Server** field and then click the **Add** button to the right of that field.

   The server is added to the **IP Address** list (below the **Server** field).

   > Note: Changes in name server configuration may take up to 5 minutes to take effect. Functions that rely on DNS may not work properly during this time. You can configure a maximum of three name servers.

3. To delete a name server entry, click the X icon for that server in the **Host Name or IP Address** list.

   A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

4. Click the **Close** button to close the Name Servers window.

Configuring NTP Servers

About this task

Network Time Protocol (NTP) is a protocol for clock synchronization between computers, and Prism Central must connect to an NTP server to synchronize the system clock. To add (or delete) an NTP server entry, do the following:

> Note: If Prism Central is running on Hyper-V, you must specific the IP address of the Active Directory Domain Controller server, not the hostname. Do not use DNS hostnames or external NTP servers.

Procedure

1. Click the gear icon in the main menu and then select **NTP Servers** in the **Settings** page.

   The NTP Servers dialog box appears.

Figure 173: NTP Servers Window

2. To add an NTP server entry, enter the server IP address or fully qualified host name in the **NTP Server** field and then click the **Add** button to the right of that field.

   The name or address is added to the **Server** list (below the **NTP Server** field).

3. To delete an NTP server entry, click the delete icon for that server in the **Servers** list.

   A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

4. Click the **Close** button to close the NTP Servers window.

Configuring an SMTP Server

About this task

Simple Mail Transport Protocol (SMTP) is an Internet standard protocol for electronic mail transmission across Internet Protocol (IP) networks, and Prism Central uses SMTP to send alert emails and to exchange emails with Nutanix customer support. To configure an SMTP server entry, do the following:

Procedure

1. Click the gear icon in the main menu and then select **SMTP Server** in the **Settings** page.

   The SMTP Server Settings dialog box appears.

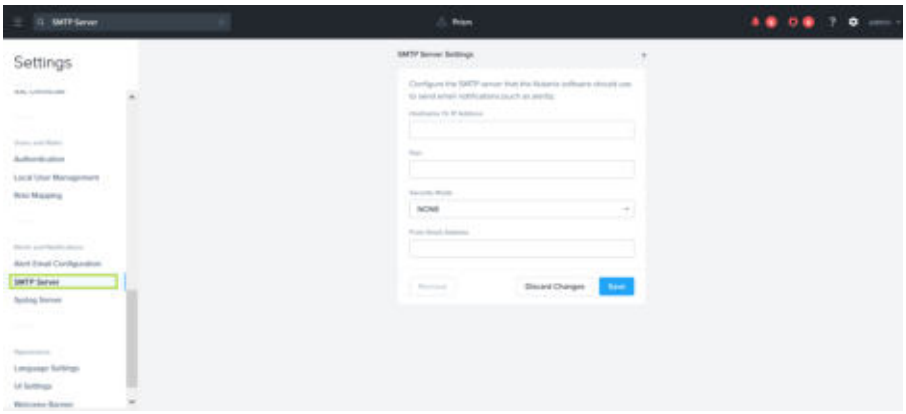

Figure 174: SMTP Server Settings Window

2. Do the following in the indicated fields:

   a. **Host Name or IP Address**: Enter the IP address or fully qualified domain name for the SMTP server.

   b. **Port**: Enter the port number to use.

      The standard SMTP ports are 25 (unencrypted), 587 (TLS), and 465 (SSL). For the complete list of required ports, see Port Reference.

   c. **Security Mode**: Enter the desired security mode from the pull-down list.

      The options are **NONE** (unencrypted), **STARTTLS** (use TLS encryption), and **SSL** (use SSL encryption).

   d. **User**: Enter a user name.

      The **User** and **Password** fields apply only when a secure option (**STARTTLS** or **SSL**) is selected. The user name might need to include the domain (*user@domain*) depending on the authentication process.

   e. **Password**: Enter the user password.

   a. **From Email Address** (optional): Enter an e-mail address that appears as the sender address.

      By default, alert and status information e-mails display "cluster@nutanix.com" as the sender address. You have the option to replace that address with a custom address by entering a sender address in this field.

3. When all the fields are correct, click the **Save** button.

Configuring SNMP

About this task

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. To configure SNMP on Prism Central, do the following:

Note: Prism Element (individual cluster) supports both the SNMP service (agent) and SNMP traps, but Prism Central only supports SNMP traps. Prism Central does not support the SNMP service (unlike Prism Element), so its SNMP capability is limited to just sending traps. See the *Prism Web Console Guide* for instructions on configuring SNMP for an individual cluster and for details about the Nutanix MIB.

Procedure

1. Click the gear icon in the main menu and then select **SNMP** in the **Settings** page.

   The SNMP Configuration dialog box appears.

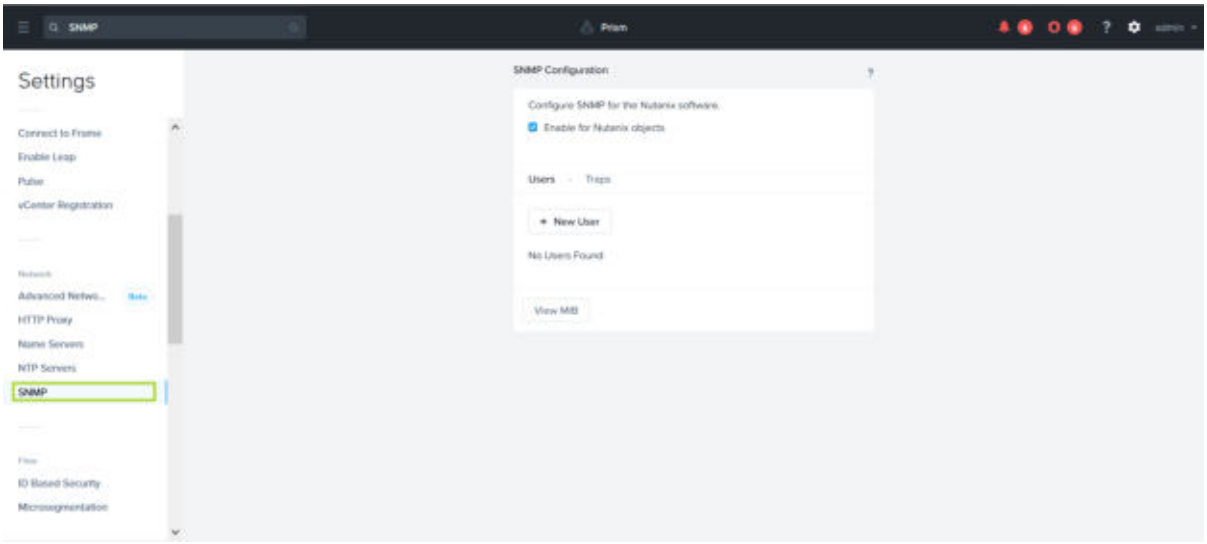

Figure 175: SNMP Configuration Window

2. To enable SNMP for Prism Central, check the **Enable SNMP** box. To disable SNMP, uncheck the box.

3. To view the Nutanix MIB (NUTANIX-MIB.txt), click the **View MIB** link. To download NUTANIX-MIB.txt, right-click and select the appropriate download action for your browser.

4. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:

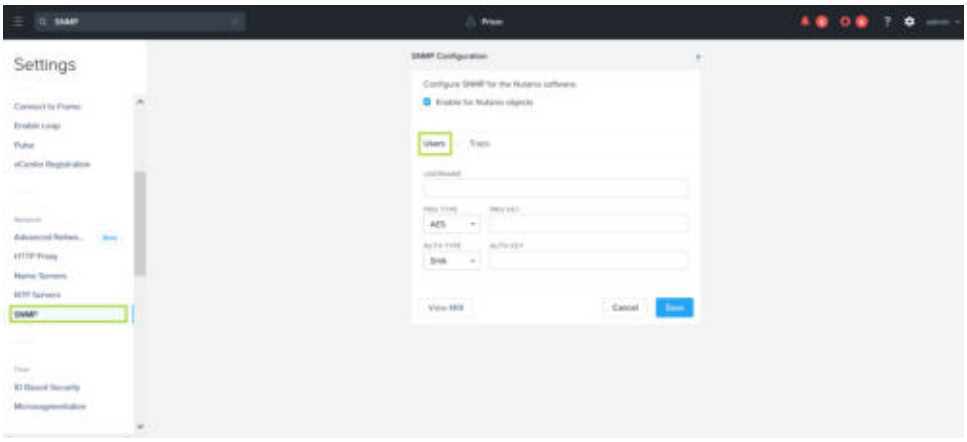

Figure 176: SNMP Configuration: Users Tab

a. **Username**: Enter a user name.

b. **Priv Type**: Select the privacy encryption type from the pull-down list.

   The only option is **AES** (Advanced Encryption Standard).

c. **Priv Key**: Enter a privacy key phrase (password) into this field.

   The key phrase is AES encrypted when the user is created.

d. **Auth Type**: Select the authentication hash function type from the pull-down list.

   The only option is **SHA** (Secure Hash Algorithm).

e. **Auth Key**: Enter an authentication key phrase (password) into this field.

   The key phrase is SHA-1 encrypted when the user is created.

f. When all the fields are correct, click the **Save** button (lower right).

   This saves the configuration and redisplays the dialog box with the new user entry appearing in the list.

5. To add an SNMP trap receiver, click the **Traps** tab and the **New Trap Receiver** button, and then do the following in the indicated fields:

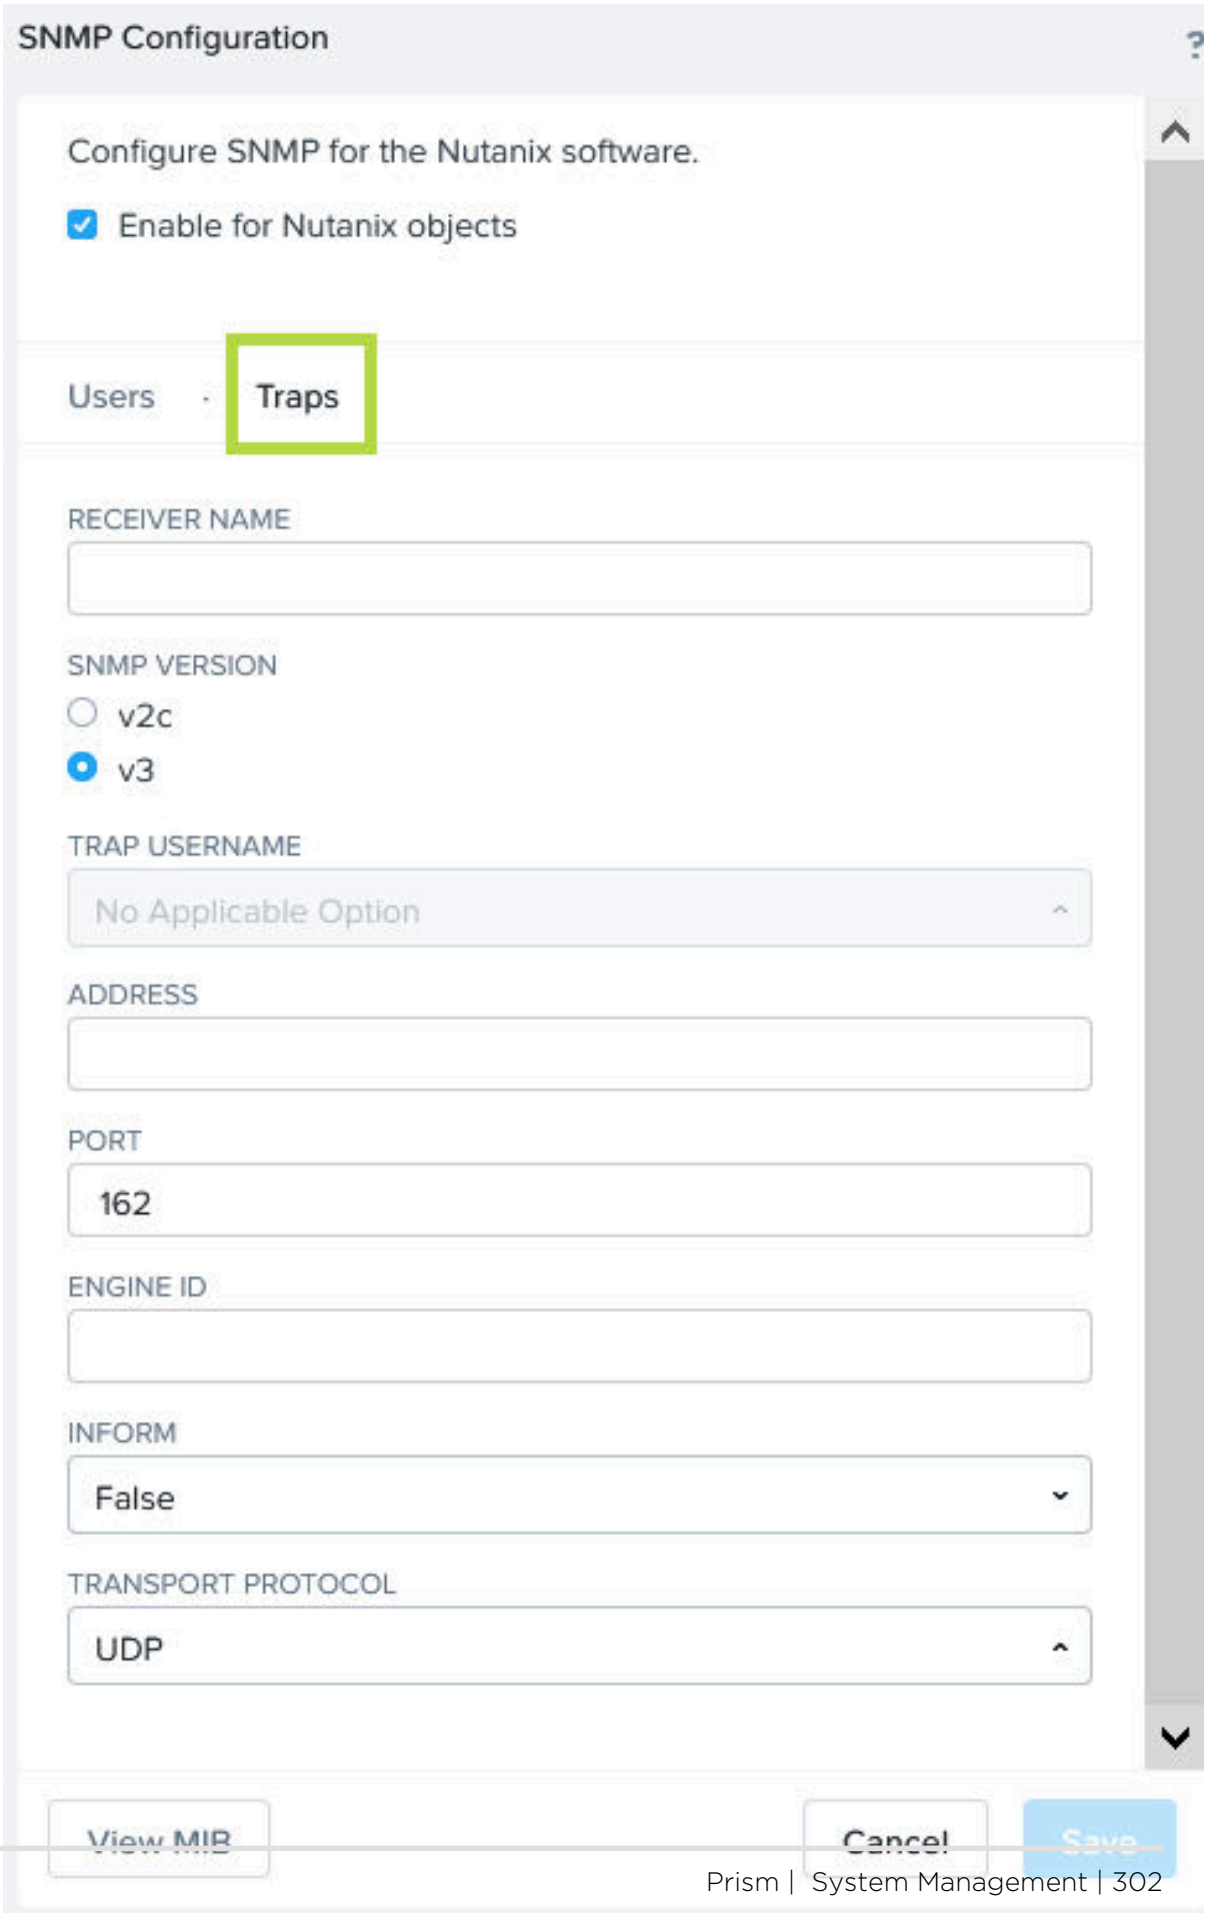

Figure 177: SNMP Configuration: Traps Tab

a. **Receiver Name**: Enter the receiver name.

b. **SNMP Version**: Select the button for the version to use, **v2c** or **v3**.

c. **Trap Username**: Select a user from the pull-down list.

   All users added previously (see step 4) appear in the pull-down list. You cannot add a trap receiver entry until at least one user has been added.

d. **Address**: Enter the target address.

   An SNMP target address specifies the destination and user that receives outgoing notifications, such as trap messages. SNMP target address names must be unique within the managed device.

e. **Port**: Enter the port number to use.

   The SNMP trap receiver uses UDP port number 162. For the complete list of required ports, see Port Reference.

f. **Engine ID** (optional): Enter an engine identifier value, which must be a hexadecimal string between 5 and 32 characters long.

   If you do not specify an engine ID, an engine ID is generated for you for use with the receiver. Every SNMP v3 agent has an engine ID that serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

g. **Inform**: Select **True** from the pull-down list to use inform requests as the SNMP notification method; select **False** to use traps as the SNMP notification method.

   SNMP notifications can be sent as traps or inform requests. Traps are one-way transmissions; they do not require an acknowledgment from the receiver. Informs expect a response. If the sender never receives a response, the inform request can be sent again. Therefore, informs are more reliable than traps. However, informs consume more resources. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform may be retried several times. The retries increase traffic and add overhead on the network. Thus, traps and inform requests provide a trade-off between reliability and resources.

h. **Transport Protocol**: Select the protocol to use from the pull-down list.

   The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.

i. When all the fields are correct, click the **Save** button (lower right).

   This saves the configuration and redisplays the dialog box with the new trap entry appearing in the list.

j. To test all configured SNMP traps, click the **Traps** tab, and then click **Test All**.

   The Nutanix cluster sends test alerts to all the SNMP trap receivers configured on the cluster.

6. To edit a user or trap receiver entry, click the appropriate tab (**Users** or **Traps**) and then click the pencil icon for that entry in the list.

   An edit window appears for that user or trap receiver entry with the same fields as the add window. (Transport entries cannot be edited.) Enter the new information in the appropriate fields and then click the **Save** button.

7. To delete an SNMP entry, click the appropriate tab (**Users** or **Traps**) and then click the X icon for that entry in the list.

   A window prompt appears to verify the delete action; click the **OK** button. The entry is removed from the list.

8. Click the **Close** button to close the SNMP Configuration window.

Configuring Syslog Monitoring

Using Prism Central, you can configure syslog monitoring to forward system logs (API Audit, Audit, Security Policy Hitlogs, and Flow Service Logs) of the registered clusters to an external syslog server. You can configure multiple remote syslog servers for Prism Central. Additionally, you can configure separate log modules to be sent to each of the rsyslog servers.

About this task

To configure syslog monitoring, do the following.

Before you begin

- You must have the IP address of the syslog server that is deployed in your environment.
- For forwarding Flow logs, the Flow feature must be enabled.

Procedure

1. Click the gear icon in the main menu and then select **Syslog Server** in the **Settings** page.
2. Click **Configure Syslog Server**.

   Note: You can configure only one syslog server per cluster.

3. In **Server Name**, enter a descriptive name for the server.
4. Enter the **IP Address** and **Port**.
5. Select the **Transport Protocol** (TCP or UDP).
6. Optionally, click the check box to enable RELP (Reliable Logging Protocol).
7. Click **Edit** against **Data Sources**.
8. Select one or more log modules from the following log types. See Syslog Modules on page 305 for details.

   - **API Audit**
   - **Audit**
   - **Security Policy Hitlogs** (policy hitlog files logs)
   - **Flow Service Logs** (Flow processes logs)

   CAUTION: Users recommend that you only configure rsyslog modules listed in Syslog Modules on page 305. Configuration of other modules using ncli commands will be over-written with any modifications done on the Prism Central web console.

9. Select the **Severity Level** from the following available options. See Syslog Modules on page 305 for details.

   Click **Save** to complete.

Syslog Modules

Refer to the following table to understand the information that is sent to the syslog server based on the selected log modules and severity levels.

Table 212: Log Modules and Severity Levels

Module Name	Severity Level	Description
API Audit	0-7	API Audit logs send information about REST API endpoints that are called and who called them from both Prism Central and Prism Element. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent.
Audit	0-7	Audit logs send information about VM, Category, and Security Policy creation, update, and delete operations. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent. See Table 214: Syslog Module: AUDIT Fields on page 306 for information on Syslog Audit fields.
Security Policy Hit Logs	All Levels	Security Policy Hit Logs send information about network flows that are acted on by a security policy. This information includes the source and destination IP address, protocol, and port. It also includes the action taken by the policy such as allow, monitor, or drop. Network statistics are also included in each message for bytes sent and received in the flow. Messages are always sent to the syslog server as level 6 - informational.
Flow Service Logs		Flow service logs contain detailed debug information used for troubleshooting the backend processes used by Flow. Note: Do not enable these unless instructed to do so by Nutanix support for troubleshooting purposes. Unlike the other data sources, the syslog level selected for Flow Service Logs has an impact on the types of data sent to the syslog server.
	0 - 5	Critical service failure messages as well as log rotation indications. Sent at the selected log level to syslog.
	6 - 7	Detailed log messages from the various components that make up the Flow product on both Prism Central and Prism Element.

The following table provides information on the syslog severity levels.

Table 213: Syslog Severity Levels

Value	Severity	Keyword	Description
0	Emergency	emerg	System is unusable.
1	Alert	alert	Should be corrected immediately.
2	Critical	crit	Critical conditions.
3	Error	err	Error Conditions.
4	Warning	warning	May indicate that an error will an action is not taken.
5	Notice	notice	Events that are unusual, but not error conditions.
6	Informational	info	Normal operational messages that require no action.
7	Debug	debug	Information useful to developers for debugging the application.

The following table provides information on the AUDIT fields of the Syslog Module.

Table 214: Syslog Module: AUDIT Fields

AUDIT Field	Description
entityType	Indicates the resource type of the entity that is associated with the audit record. Example. The entity type is *vm* for operations on virtual machines.
name	Indicates the name of the entity that is associated with the audit record. Example. Name of the virtual machine for audits or operations on virtual machines.
uuid	Indicates the unique identifier of the entity that is associated with the audit. uuid is also used to fetch the details of the entity instance.
alertUid	Indicates the name of the audit record. Example. *VmMigrateAudit* and *VmCreateAudit*.
creationTimestampUsecs	Indicates the time (in microseconds) from Epoch when the audit record was created.
defaultMsg	Provides a brief description of the current state of the operation that is captured by the audit.
opStartTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit begins.

AUDIT Field	Description
opEndTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit is completed.
operationType	Indicates the category of the audit record. Example. *Create*, *Delete*, and *Migrate*.
originatingClusterUuid	Indicates the cluster (if applicable) where the operation of is performed.
params	Additional key value pairs that capture information pertaining to the operation that is captured by the audit record.
userName	Indicates the name of the user who triggered the operation.
userUuid	Indicates the unique identifier of the user who triggered the operation.

The following table provides information on the different fields of the Security Policy Hit Logs.

Table 215: Security Policy Hit Logs Fields

Field	Description
<Timestamp>	Indicates the time of the hit log capture.
<Policy UUID>	Indicates the unique identifier of the policy that is associated with the hit log.
<Policy Name>	Indicates the name of the policy as configured by the user.
<Session Information> Create, Update, or Destroy	Indicates the session information. • Create - New session started. TCP handshake or UDP bidirectional packets. • Update - Sent every 30 seconds for active sessions. • Destroy - Session ended by TCP handshake or UDP timeout.
SRC	Indicates the source IP address.
DST	Indicates the destination IP address.
PROTO	Indicates the transport protocol.

Field	Description
ACTION	Indicates the action taken on traffic by the applied policy. • Allow - traffic allowed by policy • Drop - traffic dropped by policy • Monitor - traffic allowed by policy in monitor mode
<Packet Stats>	Indicates the packet statistics.

Modifying UI Settings

About this task

The Prism Central login page includes background animation by default, and users are logged out automatically after being idle for 15 minutes. You can change one or both of these settings.

Procedure

1. Click the gear icon and select **UI Settings** from the **Settings** menu (see Settings Menu on page 17).
   The UI Settings window appears.

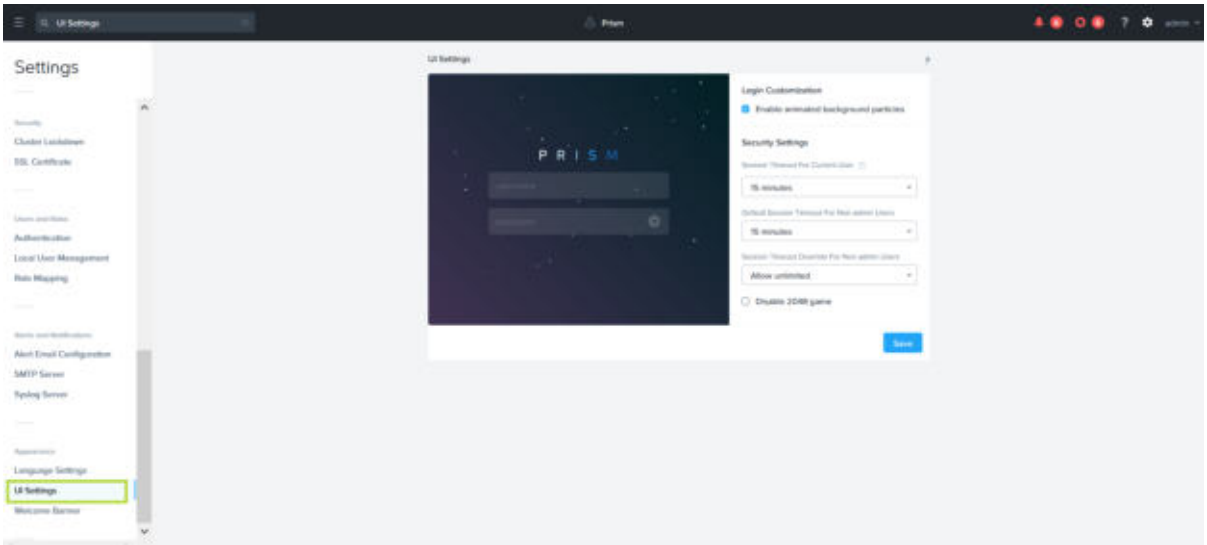

Figure 178: UI Settings Window

2. To disable the logon page background animation, under **Login Customization**, uncheck the **Enable animated background particles** box (or check it to enable).

   Unchecking the **Enable animated background particles** box in the Prism UI Settings dialog box disables the creation or drawing of particles entirely. This action stops the drawing of the particles on the Prism Element logon page.

   > Note: This setting is not persistent. In other words, if the Prism service restarts, this setting is lost and must be disabled again.

   Disabling the particles allows you to conserve critical CPU resources that are used in creating and maintaining the particles.

   > Note: Disabling or enabling this setting in Prism Web Console does not propagate to Prism Central or vice versa. The setting must be disabled in Prism Web Console and Prism Central UI separately.

   You can disable the particle animation from the logon page by clicking **Freeze space-time continuum!** at the right bottom of the logon page. This action stores a setting in the local browser to stop the animation. However, this action does not stop creation or drawing of the particles itself.

   > Note: You can enable the particle animation by clicking **Engage the warp drive!** .

3. To customize the theme, background color, title text, or blurb text on the logon page, do the following:

   - Click on the top bar (displayed in the following image) in the **UI Settings** dialog box and simultaneously press the option key on the MAC system or Alt key on the Windows system. Options for customizing the theme, title text, and blurb text are displayed.
   - Select the theme from the options displayed for **Theme**. You can change the HEX codes to create your own custom gradient background color for the logon page.
   - In the **Title Text** field, enter the text to create your custom title.
   - In the **Blurb Text** field, enter the text to create your custom blurb text. This text is displayed below the password field.

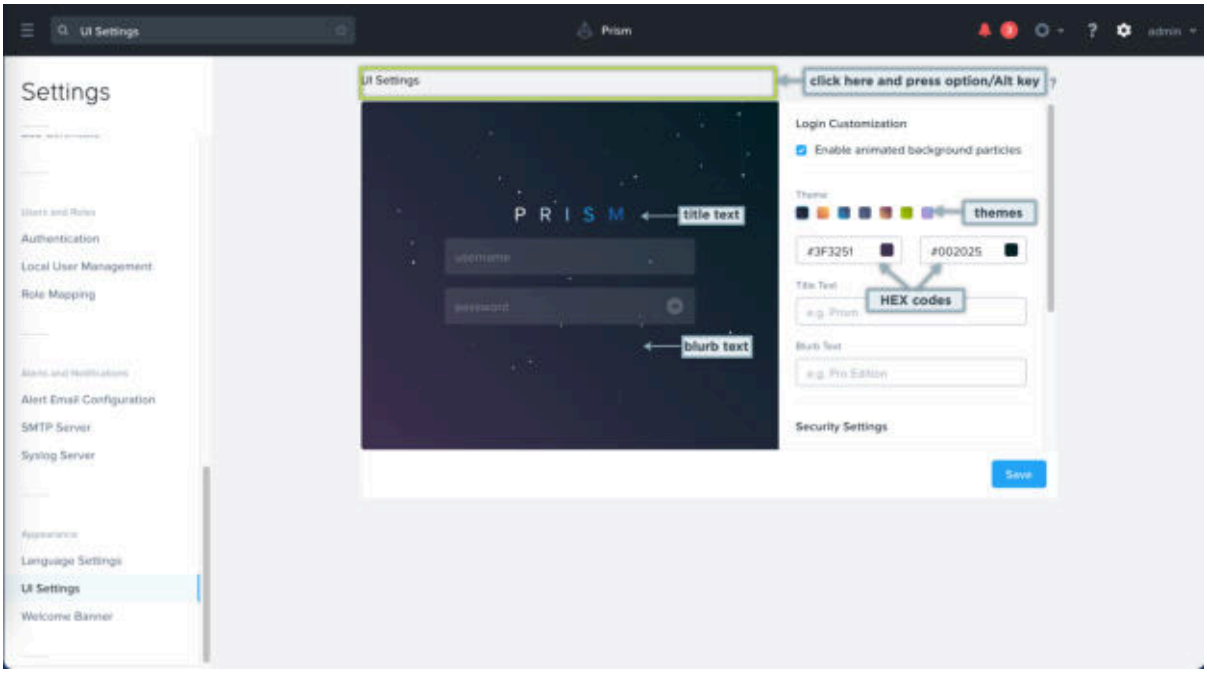

Figure 179: UI Settings Window for customizing the theme, title text, and blurb text

4. To configure session timeout, do the following under **Security Settings**:

   - Select the session timeout for the current user from the **SESSION TIMEOUT FOR CURRENT USER** drop-down list.
   - Select the default session timeout for all non-administrative users from the **DEFAULT SESSION TIMEOUT FOR NON-ADMIN USERS** drop-down list.
   - Select the appropriate option from the **SESSION TIMEOUT OVERRIDE FOR NON-ADMIN USERS** drop-down list to override the session timeout for non-administrative users.

   Note: The timeout interval for an administrator cannot be set for longer than 1 hour.

5. Clear the **Disable 2048 game** option to disable the 2048 game.
6. Click **Save** to save your changes and close the window.

Configuring a Banner Page

You have the option to create a welcome banner, which will be the first screen that appears when a user attempts to log into Prism Central. The content of the banner page is configurable, and it can include a custom message and graphics.

About this task

To configure a banner page, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Welcome Banner** in the **Settings** page.

   The Edit Welcome Banner dialog box appears.

2. Enter (paste) the desired content in HTML format in the pane on the left.

   Only "safe" HTML tags are supported. Inline event handlers, scripts, and externally-sourced graphics are not allowed.

Figure 180: Welcome Banner Window

3. Click the **Preview** button to display the banner in the pane on the right.
4. If the banner is not correct, update the HTML code as needed until the preview pane displays the desired message.
5. When the preview is correct, check the **Enable Banner** box (lower left) and then the **Save** button.

   A live banner page includes an "Accept terms and conditions" bar at the bottom. Clicking on this bar sends the user to the login page.

   You can disable (or enable) the banner at any time by unchecking (checking) the **Enable Banner** box.

Internationalization (i18n) (Prism Central)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 216: Internationalization Support for Prism Central

Supported Entities	Unsupported Entities
User-defined dashboard name	Password fields
Custom widget name	Static dashboard name
First and last name under **Update Profile**	Static widget name
User name, first name, and last name under **User Management**	
Chart name	

Localization (L10n) (Prism Central)

Nutanix localizes the user interface in Simplified Chinese and Japanese language. All the static screens are translated to the selected locale language.

You have an option to change the language settings of the cluster from English (default) to Simplified Chinese or Japanese. For more information, see Changing the Language Settings (Prism Central) on page 312.

If the Prism Element instance is launched from Prism Central, language settings of Prism Central takes precedence over Prism Element.

The dashboards (including tool tips) and menus of Prism Central are localized.

Guidelines and Limitations

- Logical entities that do not have a contextual translation available in the localized language are not translated.
- The AOS generated alerts and events are not localized to the selected locale language.
- Following strings are not localized: VM, CPU, vCPU, Language Settings, licensing details page, hardware names, storage denominations (GB, TB), About Nutanix page, EULA, service names (SNMP, SMTP), hypervisor types.

Changing the Language Settings (Prism Central)

Perform the following procedure to change the language settings in Prism Central. You can change the language setting to Simplified Chinese or Japanese.

Procedure

1. Click the gear icon in the main menu and then select **Language Settings** in the **Settings** page.

   The Language Setting window appears. The **English** language is selected by default.

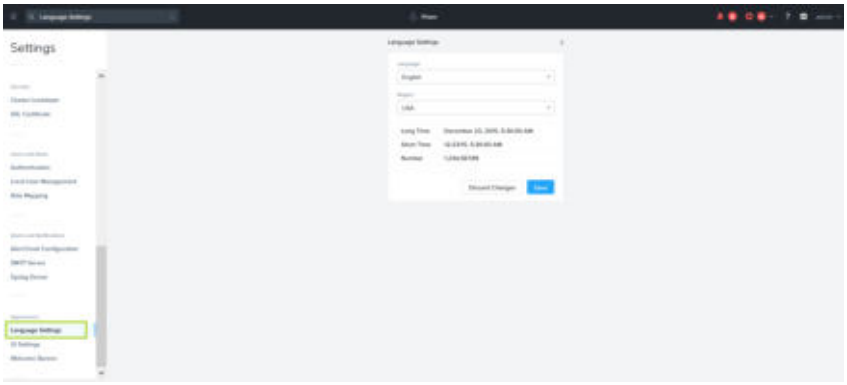

Figure 181: Language Settings Window

2. To change the setting, do the following:

   » To change the language setting of the cluster to Simplified Chinese, select **Simplified Chinese** from the drop-down menu.

   » To change the language setting of the cluster to Japanese, select **Japanese** from the drop-down menu.

   » To change the locale settings (date, time, calendar), select the appropriate region from the **Region** drop-down menu.

   By default, the locale is set to the language setting that you have set in the **Language** drop-down menu. However, you can change the **Region** to display the date, time, or calendar in some other format. This format for date, time, and calendar is applied for the entire cluster.

3. Click **Save**.

   The language and locale settings (date, time, calendar) is changed according to the selection. For example, in the below screen shot, once you click **Save** the language setting for the cluster is changed to Chinese and locale setting is changed to Russian. For more information on the entities that are supported in Simplified Chinese, see Internationalization (i18n) (Prism Central) on page 311. Also, the user interface is localized according to the

selection. For more information about localization, see Localization (L10n) (Prism Central) on page 312.

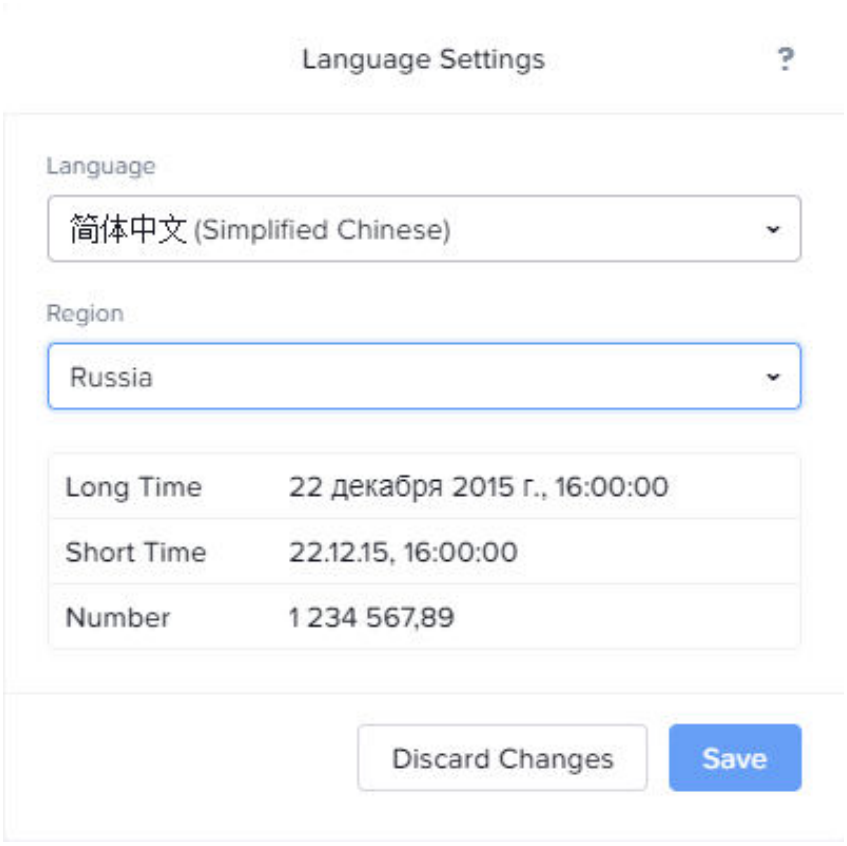

Figure 182: Localized Settings (Chinese/Russian)

7

SECURITY AND USER MANAGEMENT

Prism Central provides several mechanisms to manage security and control user access. See the following topics in the "Security Management using Prism Central (PC)" chapter of the *Nutanix Security Guide*.

- Set the user authentication method to local, directory service, or both (see "Configuring Authentication").
- Add, edit, or delete local user accounts (see "Managing Local User Accounts").
  - Update your account (see "Updating My Account").
- Assign roles to users (see "Controlling User Access (RBAC)").
- Install SSL certificates (see "Installing an SSL Certificate").
- Control SSH access to Prism Central (see Controlling Remote (SSH) Access "Controlling Remote (SSH) Access").
- Create and apply security-related policies (see "Security Policies" in the *Flow Microsegmentation Guide*).

8

VIRTUAL INFRASTRUCTURE (CLUSTER) ADMINISTRATION

You can view hardware and infrastructure information about registered (individual) clusters through the Entities menu (see Hardware Entities on page 145 and Virtual Infrastructure Entities on page 82). In addition, you can configure certain components on a registered cluster directly through Prism Central.

- To create or manage a VM on an AHV or ESXi cluster, see
  - AHV: Creating a VM (AHV) on page 317, Creating a VM (Self Service) on page 327, Managing a VM (AHV) on page 332, and Managing a VM (Self Service) on page 342
  - ESXi: Creating a VM (ESXi) on page 357 and Managing a VM (ESXi) on page 362
- To upload disk or ISO images that can by applied when creating a VM, see Image Management on page 394.
- To add images or VMs to a catalog that can be applied by a self-service portal administrator to project users, see Catalog Management on page 390.
- To configure affinity polices for an AHV cluster, see Affinity Policies for AHV on page 416
- To configure network connections for a cluster, see Configuring Network Connections on page 423.
- To perform other administrative tasks for a cluster, see Performing Other Administrative Tasks on page 451.
- To register (or unregister) an ESXi cluster with vCenter, see vCenter Server Integration on page 452.

VM Management

You can create and manage VMs directly from Prism Central when the hypervisor is either ESXi or AHV.

- ESXi
  - To create a VM, see Creating a VM (ESXi) on page 357.
  - To manage or modify a VM, see Managing a VM (ESXi) on page 362.
- AHV
  - To create a VM, see Creating a VM (AHV) on page 317.
  - To manage or modify a VM, see Managing a VM (AHV) on page 332.

- Prism Self Service (AHV)
  - To create a VM, see Creating a VM (Self Service) on page 327.
  - To manage or modify a VM, see Managing a VM (AHV) on page 332.

Creating a VM (AHV)

You can create virtual machines (VMs) in Acropolis managed clusters through Prism Central.

About this task

Figure 183: Creating a VM (AHV)

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275780958001 on Brightcove.

Note: For instructions about how to install Nutanix VirtIO on Windows VMs, see *Windows VM Provisioning* in the *AHV Administration Guide*.

To create a VM, do the following:

Note: If you are logged in as the self-service administrator or a project member, see Creating a VM (Self Service) on page 327.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Create VM** button.
   The **Create VM** wizard appears.

2. In the Configuration step, do the following in the indicated fields:
   a. **Name**: Enter a name for the VM.
   b. **Description** (optional): Enter a description for the VM.
   c. **Cluster**: Select the target cluster from the drop-down list on which you intend to create the VM.
   d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.
   e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
   f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
   g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

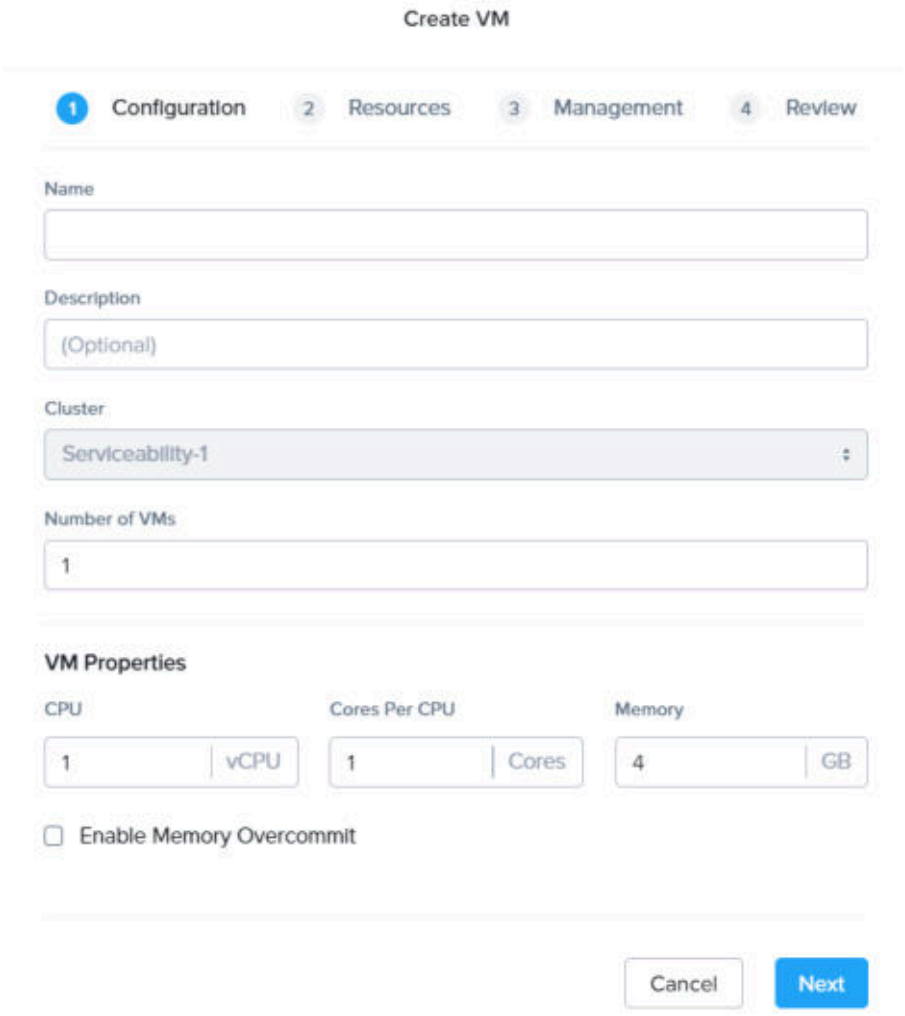

Figure 184: Create VM Window (Configuration)

3. In the **Resources** step, do the following.

   Disks: To attach a disk to the VM, click the **Attach Disk** button. The Add Disks dialog box appears. Do the following in the indicated fields:

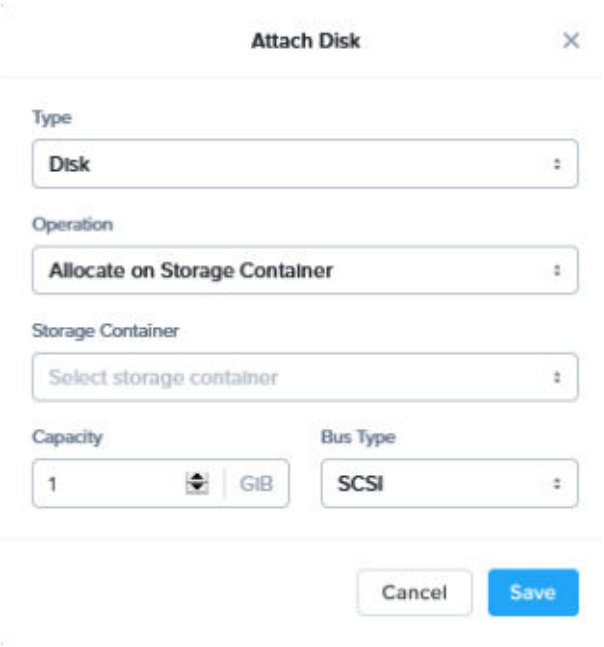

Figure 185: Add Disk Dialog Box

   a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the drop-down list.

   b. **Operation**: Specify the device contents from the drop-down list.

      - [CD-ROM only] Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
      - [Disk only] Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
      - Select **Clone from Image** to copy an image that you have imported by using image service feature onto the disk.

   c. **Bus Type**: Select the bus type from the drop-down list. The choices are **IDE**, **SCSI**, **PCI**, or **SATA**.

   d. **Clone frm Image**: Select the image that you have created by using the image service feature.

      This field appears only when **Clone from Image Service** is selected. It specifies the image to copy.

      > Note: If the image you created does not appear in the list, see KB-4892.

   e. **Storage Container**: Select the storage container to use from the drop-down list.

      This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.

f. **Capacity**: Enter the disk size in GiB.

g. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the Create VM dialog box.

h. Repeat this step to attach additional devices to the VM.

4. In the Resources step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The Attach to Subnet dialog box appears.

   Do the following in the indicated fields:

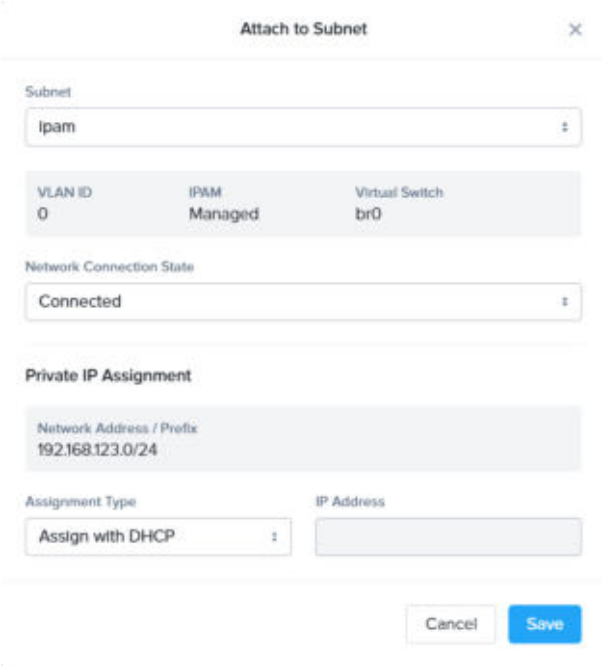

Figure 186: Attach to Subnet Dialog Box

   a. **Subnet Name**: Select the target subnet from the drop-down list.

      The list includes all defined networks (see *Network Configuration for VM Interfaces* in *Prism Web Console Guide*).

   b. **VLAN ID**: This is a read-only field that displays the VLAN ID.

   c. **IPAM**: This is a read-only field that informs you if the subnet is IPAM managed or not.

   d. **Virtual Switch**: This is a read-only field that displays the name of the virtual switch associated with the subnet.

   e. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are *Connected* or *Disconnected*.

   f. **Network Address/Prefix**: This is a read-only field that displays the network IP address and prefix.

   g. **Assignment Type**: Select the IP assignment type from the drop-down list. The options are *Assign with DHCP*, *Assign Static IP*, and *No Private IP*.

   h. **IP Address**: Enter an IP address for the VLAN.

      This field appears only if the NIC is placed in a managed network. Entering an IP address in this field is optional when the network configuration provides an IP pool. If the field is left blank, the NIC is assigned an IP address from the pool.

   i. When all the field entries are correct, click the **Save** button to create a network interface for the VM and return to the Create VM dialog box.

j. Repeat this step to create additional network interfaces for the VM.

5. In the Resources step, do the following.

   Boot Configuration: Select one of the following firmware to boot the VM.

   » **Legacy BIOS**: Select legacy BIOS to boot the VM with legacy BIOS firmware.

   » **UEFI**: Select UEFI to boot the VM with UEFI firmware. UEFI firmware supports larger hard drives, faster boot time, and provides more security features. For more information about UEFI firmware, see *UEFI Support for VM* section in the *AHV Administration Guide*.

6. In the Resources step, do the following.

   GPUs: (For GPU-enabled AHV clusters only) To add a GPU, click the **Add GPU** button. The Add GPU dialog box appears. Do the following in the indicated fields:

   a. **GPU Mode**: Click the radio button for the desired mode, either **vGPU** or **Passthrough**.

   b. If you selected **vGPU**, do the following:

   - **NVIDIA Virtual GPU License**: Select a license type from the drop-down list. This sets (filters the list of) available profiles. Click the help link (circled question mark) for information about the license types.

     Note:

     You can add multiple vGPUs to the VM only if you select the license for NVIDIA Virtual GPU software version 10.1 (440.53) or later.

     Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.

   - **vGPU Profile**: Click the radio button for the desired profile. Click the help links (Virtual Slice column and end of line) for more information about the profiles.

     Note:

     Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.

   - After you add the first vGPU, to add multiple vGPUs, see Adding Multiple vGPUs to the Same VM on page 350.

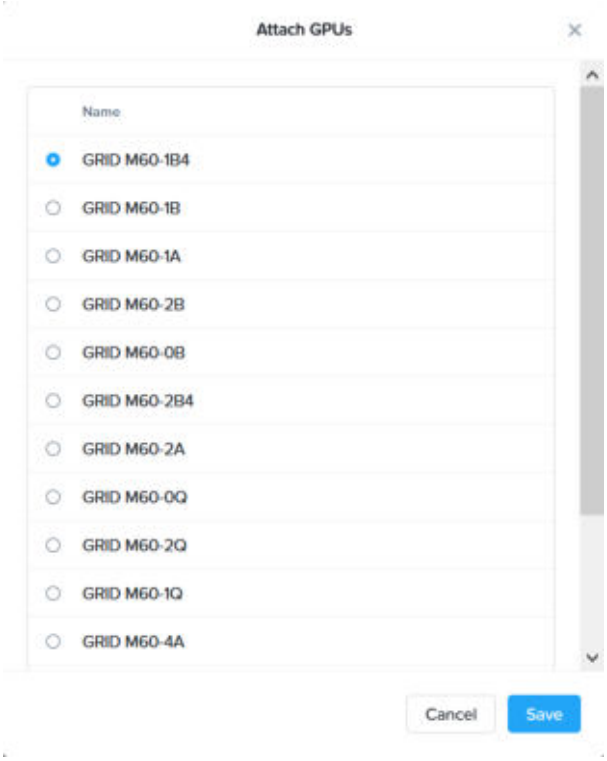

Figure 187: Add GPU Dialog Box: vGPU

c. If you selected **Passthrough**, click the radio button for the desired type.

Figure 188: Add GPU Dialog Box: Passthrough

d. Click the **Add** button.

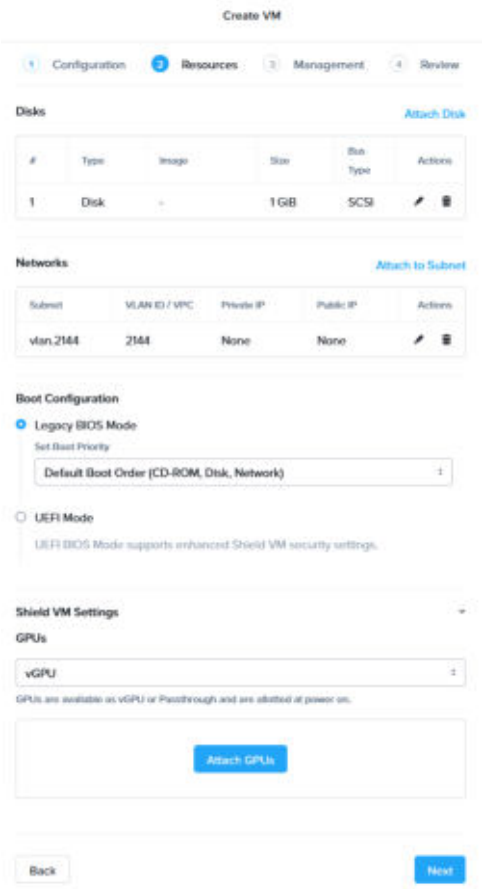

Figure 189: Create VM Window (Resources)

7. In the **Management** step, do the following.

   a. **Categories**: Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

   b. **Timezone**: Select the local timezone to use from the drop-down list.

   c. **Use this VM as an agent VM**: Select this option to make this VM as an agent VM.

      You can use this option for the VMs that must be powered on before the rest of the VMs (for example, to provide network functions before rest of VMs are powered on the

host) and must be powered off and migrated after rest of the VMs (for example, during maintenance mode operations).

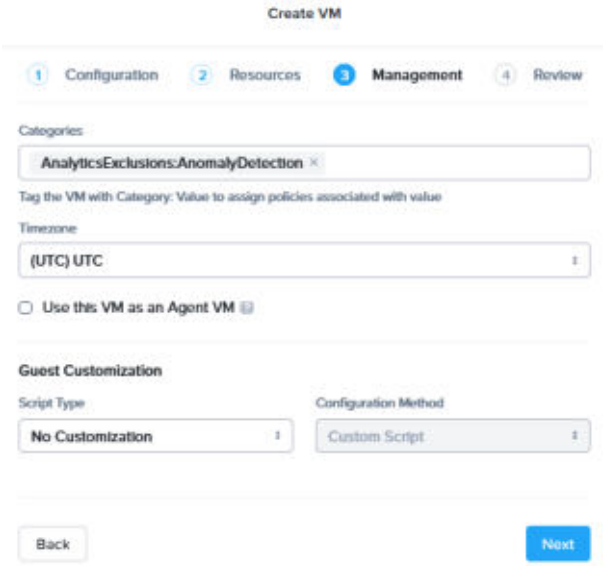

Figure 190: Create VM Window (Management)

8. In the **Management** step, do the following.

   Guest Customization: To customize the VM, select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs).

   Script Type: Select the VM script customization type from the drop-down list. The options are *No Customization*, *Sysprep(Windows)*, or *Cloud-init (Linux)*.

   Configuration Method: Select the VM script customization method from the drop-down list. This field is activated when the **Script Type** is either *Sysprep(Windows)* or *Cloud-init (Linux)*. The options are *Custom Script* or *Guided Script*.

   Field required for configuring Cloud-init and Sysprep, such as options for specifying a configuration script and option to upload script appears.

9. To specify a user data file (Linux VMs) or answer file (Windows VMs) the script for unattended provisioning, do one of the following.

   » If the file is available on your local computer, click the **Upload Script** button, choose and upload the file.

   » Create or paste the contents of the file in the text box below the **Upload Script** button.

   Note: The script type supports the following file formats.

   - Sysprep: XML
   - Cloud-init (Linux): YAML, JSON, or Shell.

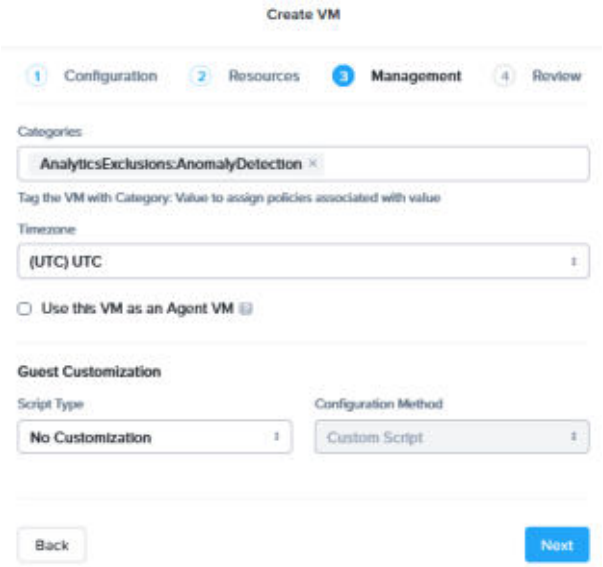

Figure 191: Create VM Window (Management)

10. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the Create VM dialog box.

    The new VM appears in the VMs entity page list.

Creating a VM (Self Service)

About this task

When you are logged in as a self-service administrator or a project member with permission to create a VM (see Prism Self Service Overview on page 457), the procedure to create a VM is different because the VM will be based on a source file stored in the Prism Central catalog (see Catalog Management on page 390). To create a VM in this case, do the following:

Note: If you are not logged in as the self-service administrator or a project member, see Creating a VM (AHV) on page 317.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Create VM** button.

   The Create VM page appears.

2. In the **Select Source** tab, select the source image for the VM.

   » Click **VM Template** if you want to create a VM from a template in the catalog, and then click **Next**.

   A VM template includes all the configuration information required to create a VM. You can use a template to quickly deploy a VM without having to specify all configuration information.

   » Click **From Disk Images** if you want to create a VM from a mounted disk image, and then click **Next**.

   Disk images can be CD-ROM images such as installer ISO images or images of hard drives that contain pre-installed applications and data. Disk images enable you to share data with other VMs, but you need to specify configuration information to create the VM that will use the data.

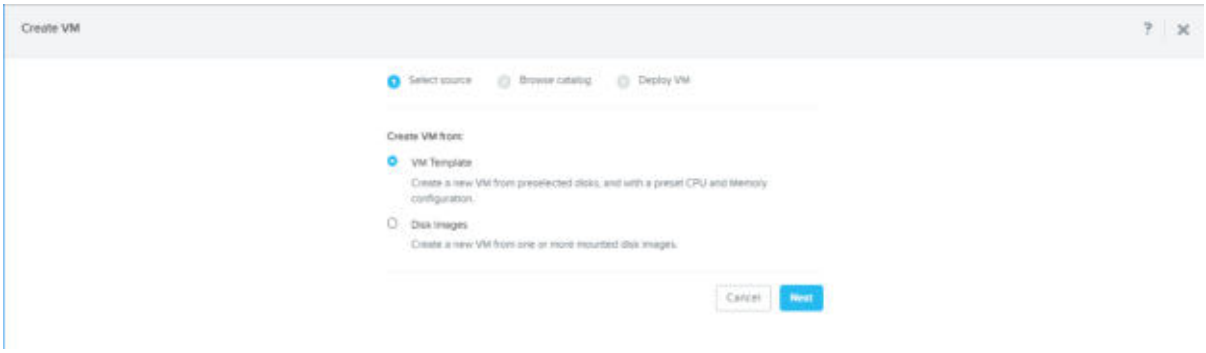

Figure 192: Create VM page: Select Source tab

3. In the **Browse Catalog** tab, do one of the following:

   - For **VM Template** select (click on) the target VM template, and then click **Next**. The available VM templates from the catalog are listed. If the target VM image does not appear in the list, you can search for it by name in the search field.

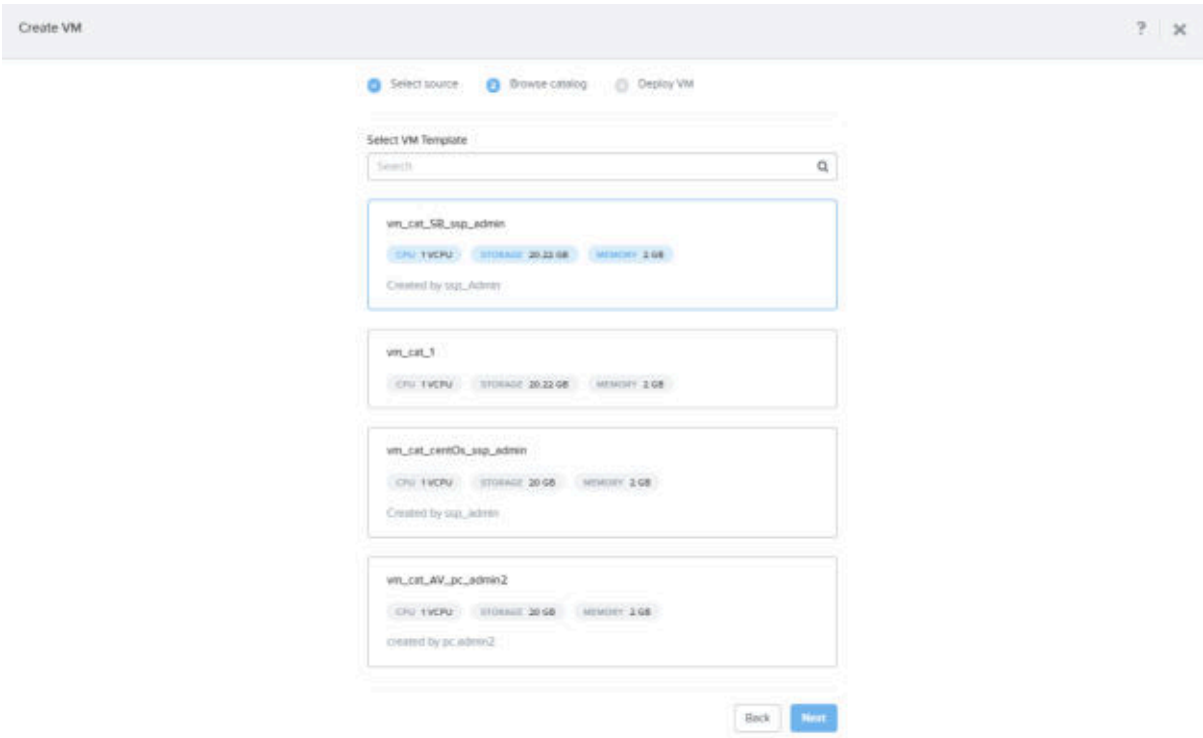

Figure 193: Create VM page: Browse Catalog tab (VM templates)

   - For **Disk Images** select (click the radio button for) one or more of the disk images, and then click **Next**. The available disk images from the catalog are listed. If the target disk image does not appear in the list, you can search for it by name in the search field.

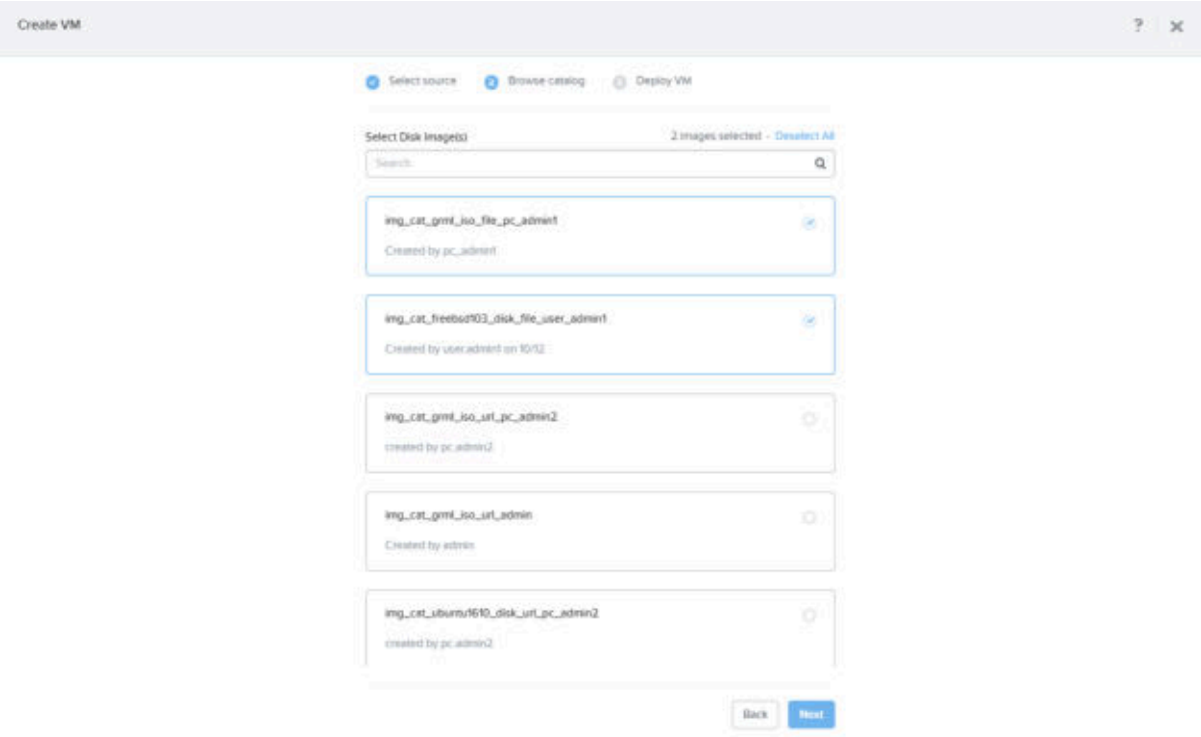

Figure 194: Create VM page: Browse Catalog tab (disk images)

4. In the **Deploy VM** tab, do the following in the indicated fields:

   a. **Name**: Enter a name for the VM.

   b. **Target Project**: Select the project associated with this VM from the pull-down list.

   c. **Disks**: Displays the disks and CD-ROMs attached to this VM.

      - For **VM Template** the list is fixed. Click the radio button for the device to boot from.
      - For **Disk Images** there are **+ New Disk** and **+ New CDROM** links above the list.
        - Click **+ New Disk** to add a disk. This displays a new line at the bottom of the list. Specify the disk size and then click **Save** at the end of line to add the disk. (The name and type field values are entered automatically.)
        - Click **+ New CDROM** to add a CD-ROM. This displays a new line at the bottom of the list. Click **Save** at the end of line to add the CD-ROM. (All field values are entered automatically.)

      Repeat this step until you have added all the desired disks and CD-ROMs. When the list is complete, click the radio button for the device to boot from.

   d. **Selected VM Template** (VM template only): Displays information about the selected VM template image. (There is nothing to do in this field.)

   e. **Guest Customization** (VM template only): If customizations were implemented when the VM template was added to the catalog, they appear here (see Adding a Catalog Item on page 391).

   f. **Network**: Select a network for the VM.

      If there is a default network specified for your project (see Creating a Project on page 462), that network is selected by default. However, you can add networks or change the default selection.

   g. **Categories** (optional): To assign the VM a category value, enter the category name and then click the plus sign (+) to the right of the field.

      This field acts like a search field; it provides a list of matching categories as you enter a string. Select the desired category value when you see it in the list. You can repeat this step for as many categories as desired.

   h. **Advanced Settings**: Check this box to specify the number of vCPUs and amount of memory allocated to this VM. Otherwise, the default values apply (1 vCPU and 2 GB of memory).

   i. When all values are correct, click the **Save** button.

      This closes the page and deploys the VM. The new VM appears in the VMs entity page list.

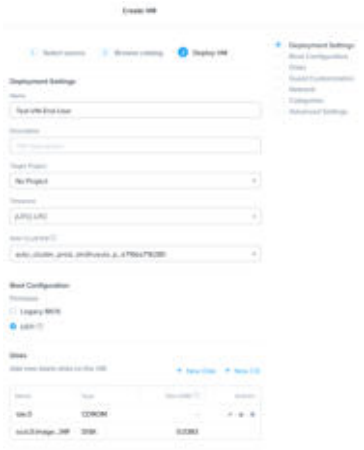

Figure 195: Create VM page: Deploy VM tab (disk images)

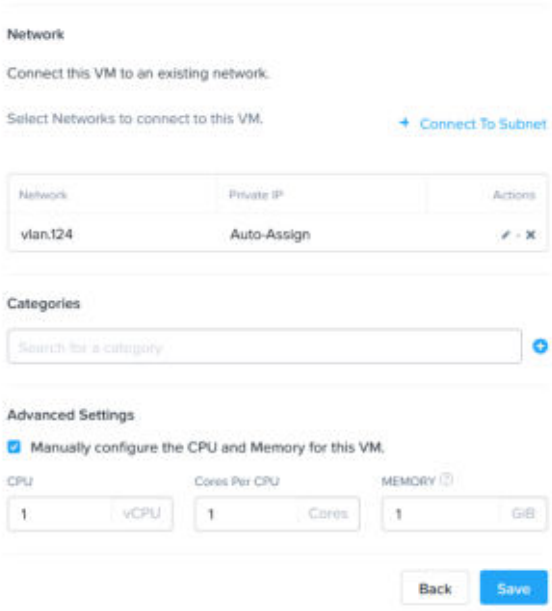

Figure 196: Create VM page: Deploy VM tab (network)

Managing a VM (AHV)

About this task

After creating a VM (see Creating a VM (AHV) on page 317) you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, update the network configuration, disable efficiency measurement, disable anomaly detection, power on (or off) the VM, enable flash mode for a VM and disks, assign the VM to a protection policy, migrate the VM, create VM recovery point, add the VM to a recovery plan, run a playbook, manage categories, install, manage, and upgrade Nutanix Guest Tools (NGT), configure VM host affinity, add the VM image to the catalog, manage the VM ownership, configure QoS settings and export VM as OVA.

> Note: You can perform only those operations for which you have permissions from the admin.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 82) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
- Go to the details page of a selected VM (see VM Details View on page 90) and select the desired action.

> Note: The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Make the desired changes and then click the **Save** button in the **Review** step.

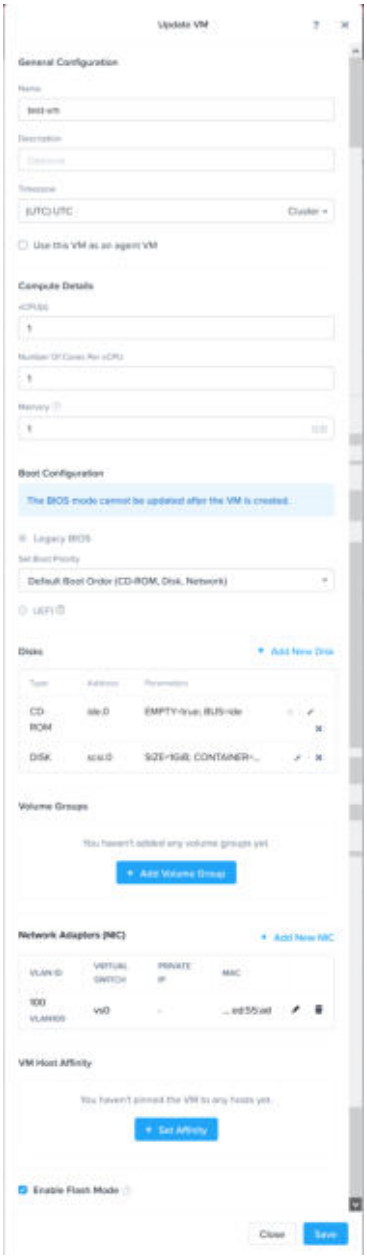

Figure 197: Update VM Window (Configuration)

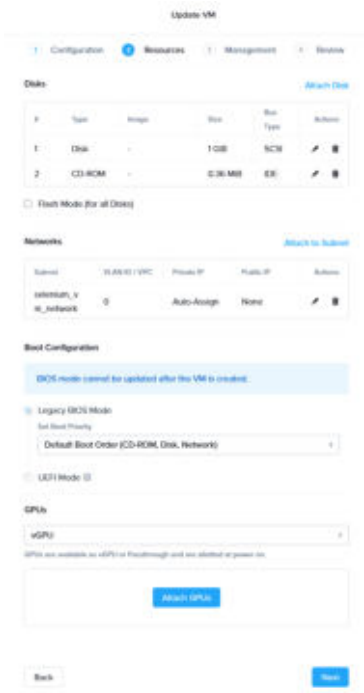

Figure 198: Update VM Window (Resources)

- **Disks:** You can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See

Creating a VM (AHV) on page 317 before you modify the disk or create a new disk for a VM. You can enable or disable the flash mode settings for the VM and VM disks.

a. To enable flash mode on the VM, click the **Enable Flash Mode** check box.

   » After you enable this feature on the VM, the status is updated in the VM table view. To view the status of individual virtual disks (disks that are flashed to the SSD), click the update disk icon in the **Disks** pane in the Update VM window.

   » You can disable the flash mode feature for individual virtual disks. To update the flash mode for individual virtual disks, click the update disk icon in the **Disks** pane and deselect the **Enable Flash Mode** check box.

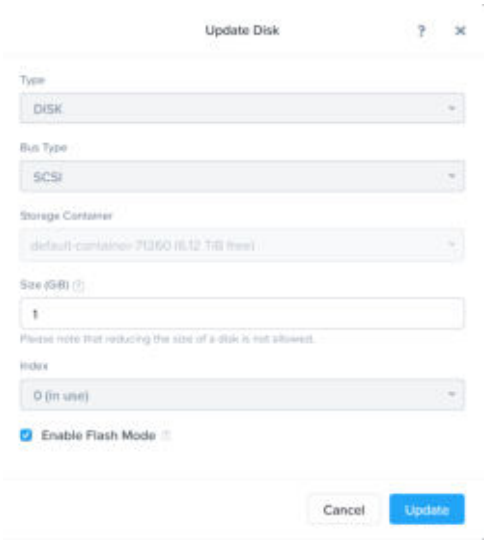

Figure 199: Update Disk Window (VM Disk Flash Mode)

Note: The flash mode setting is not supported for **CD-ROM** storage device.

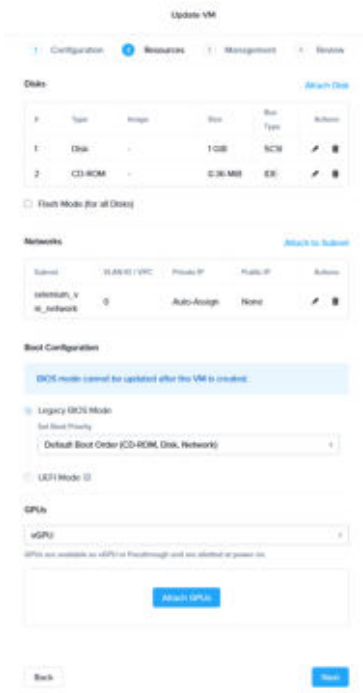

Figure 200: Update VM Window (Resources)

- **Networks:** You can attach new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (AHV) on page 317 before you modify NIC network or create a new NIC for a VM.
- **Boot Configuration:** You cannot modify the BIOS mode after the VM is created. However, if you have enabled the Secure Boot option under UEFI Mode while creating the VM, you can modify this setting during the VM update.
- **GPUs:** (For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the specified vGPU license. However, you need to power off the VM before you perform these operations.

  - Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.
  - Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
  - After you add the first vGPU, to add multiple vGPUs, see *Adding Multiple vGPUs to the Same VM* in the *AHV Admin Guide*.

- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

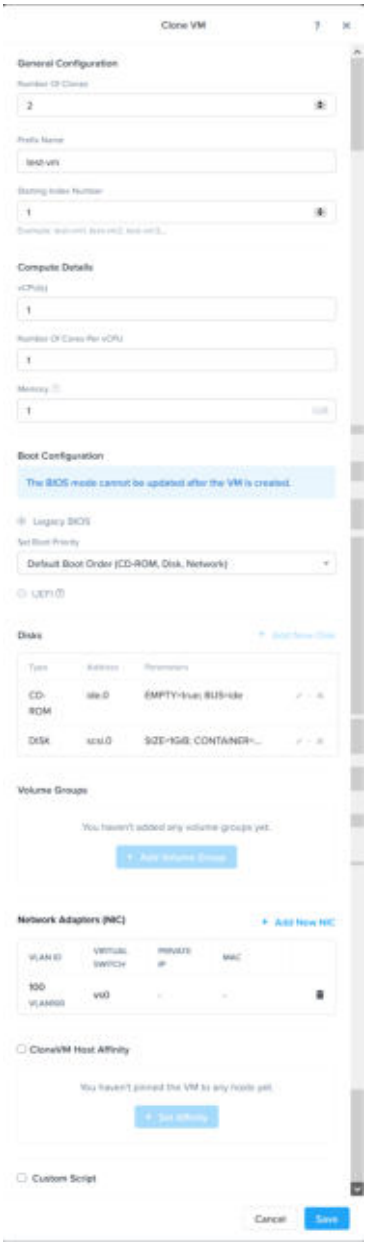

Figure 201: Clone VM Window

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM

table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

> Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

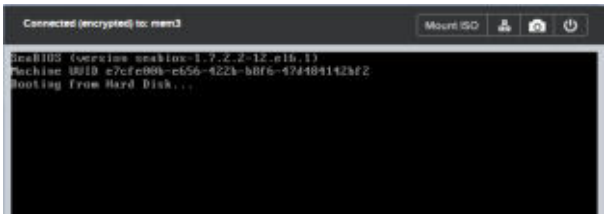

Figure 202: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To assign the VM to a protection policy, select **Protect**. This opens a page to specify the protection policy to which this VM should be assigned (see Policies Management on page 478). To remove the VM from a protection policy, select **Unprotect**.

  > Note: You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Leap and configuring the Availability Zone.

- To create VM recovery point, select **Create Recovery Point**.

  This displays the Create VM Recovery Point dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

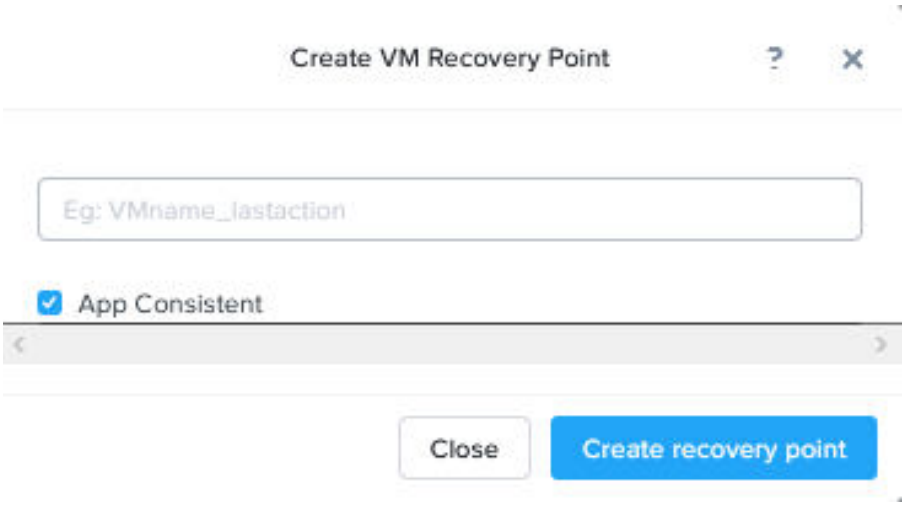

Figure 203: Create VM Recovery Point Window

- To migrate the VM to another host, select **Migrate**.

  This displays the Migrate VM dialog box. Select the target host from the pull-down list (or select the **System will automatically select a host** option to let the system choose the host) and then click the **Migrate** button to start the migration.

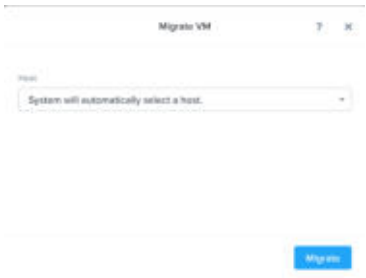

Figure 204: Migrate VM Window

> Note: Nutanix recommends to live migrate VMs when they are under light load. If they are migrated while heavily utilized, migration may fail because of limited bandwidth.

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the "Explicitly Adding a VM to a Recovery Plan" section in the Leap Administration Guide).
- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 606).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 415).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 372).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 374).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > Note: If you clone a VM, by default NGT is not enabled on the cloned VM. You need to again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the *Prism Web Console Guide*.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

```
ncli> ngt mount vm-id=virtual_machine_id
```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

```
ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-
```

c1601e759987

Note:

- Self-service restore feature is not enabled by default on a VM. You need to manually enable the self-service restore feature.
- If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.

- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 375).
- To establish VM host affinity, select **Configure VM Host Affinity**.

  A window appears with the available hosts. Select (click the icon for) one or more of the hosts and then click the **Save** button. This creates an affinity between the VM and the selected hosts. If possible, it is recommended that you create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure. See Affinity Policies for AHV on page 416 for more information about the VM affinity policies.

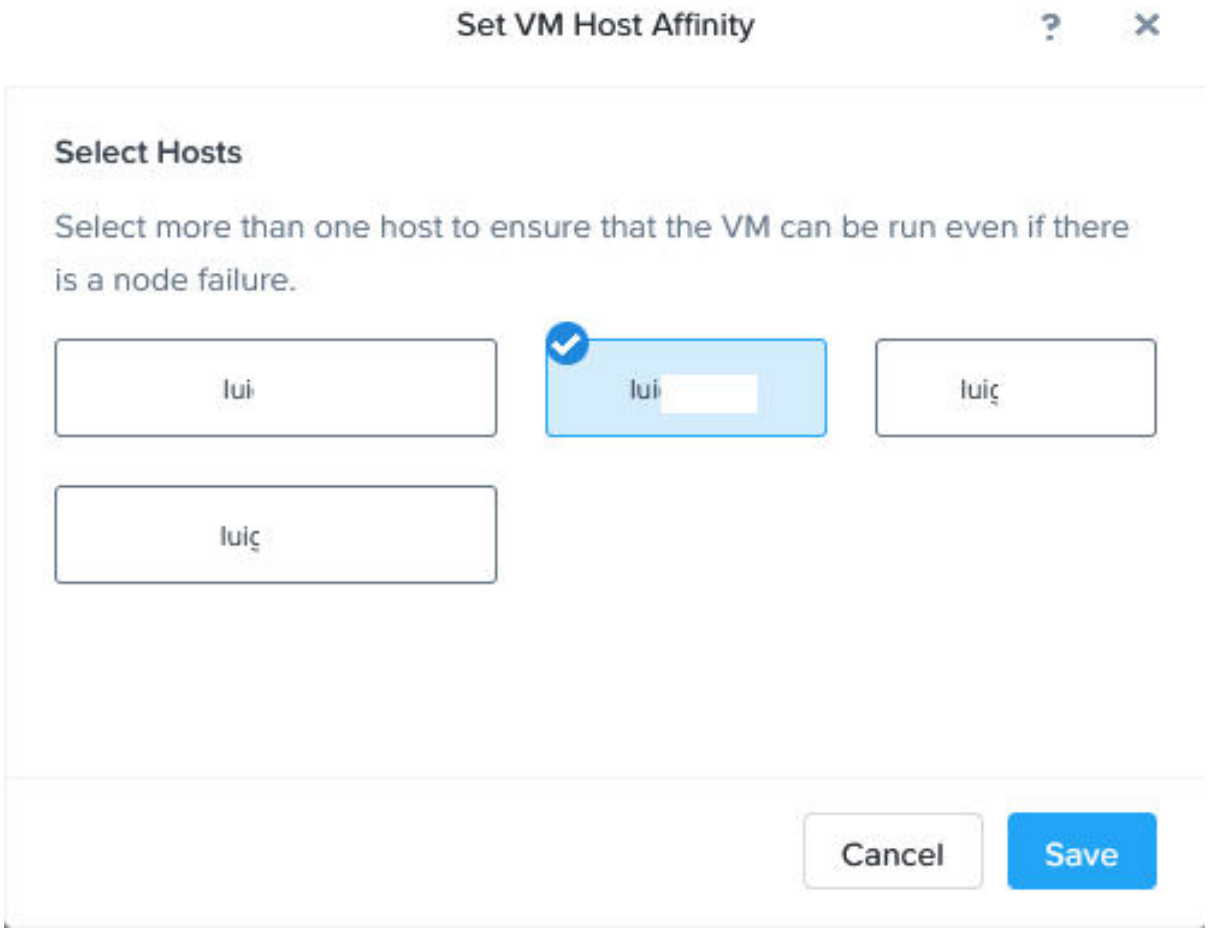

Figure 205: Set VM Host Affinity Window

- To add a VM to the catalog, select **Add to Catalog**. This displays the **Add VM to Catalog** page (see Adding a Catalog Item on page 391).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

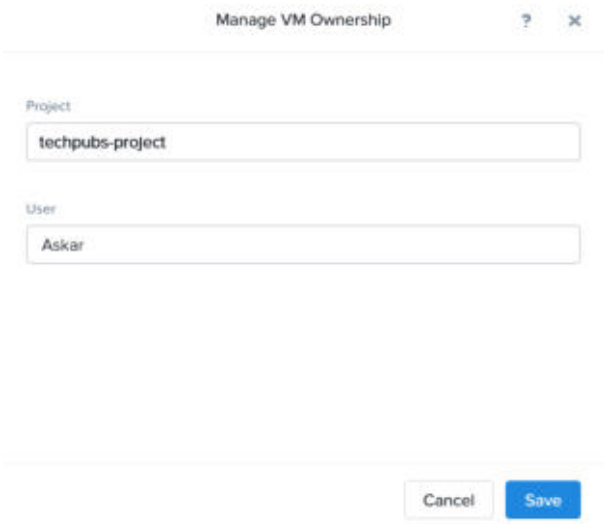

Figure 206: VM Ownership Window

- To configure quality of service (QoS) settings, select **Set QoS Attributes** (see Setting QoS for an Individual VM on page 378).
- To export as OVA, select **Export as OVA** (see OVA Management on page 433 and Uploading an OVA on page 434).

Managing a VM (Self Service)

About this task

After creating a VM (Creating a VM (Self Service) on page 327), you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, power on (or off) the VM, disable efficiency measurement, disable anomaly detection, set flash mode settings for a VM and VM disks, assign the VM to a protection policy, migrate the VM, add the VM to a recovery plan, run a playbook, manage categories, install and manage Nutanix Guest Tools (NGT), configure VM host affinity, add the VM image to the catalog, manage the VM ownership, configure QoS settings, or export VM as OVA.

> Note: You can perform only those operations for which you have permission.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 82) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.

- Go to the details page of a selected VM (see VM Details View on page 90) and select the desired action.

> Note: The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The Edit VM dialog box appears, which includes the same fields as the Create VM dialog box. Make the desired changes and then click the **Save** button.

  (For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the specified vGPU license. However, you need to power off the VM before you perform these operations.

  - Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.
  - Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
  - After you add the first vGPU, to add multiple vGPUs, see *Adding Multiple vGPUs to the Same VM* in the *AHV Admin Guide*.

Figure 207: Edit VM Window

- **Disks:** You can add new disks to the VM using the **+ New Disk** option. You can also add new CDs to the VM using the **+ New CD** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM (Self Service) on page 327 before you modify the disk or create a new disk for a VM.

- **Network:** You can attach new network to the VM using the **Connect to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (Self Service) on page 327 before you modify NIC network or create a new NIC for a VM.
- **Configuration:**You can update the VM category value, enter the category name and then click the plus sign (+) to the right of the field. See Creating a VM (Self Service) on page 327 before you modify or assign new categories for a VM.
- **Categories:**You can update the .
- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot enabled.

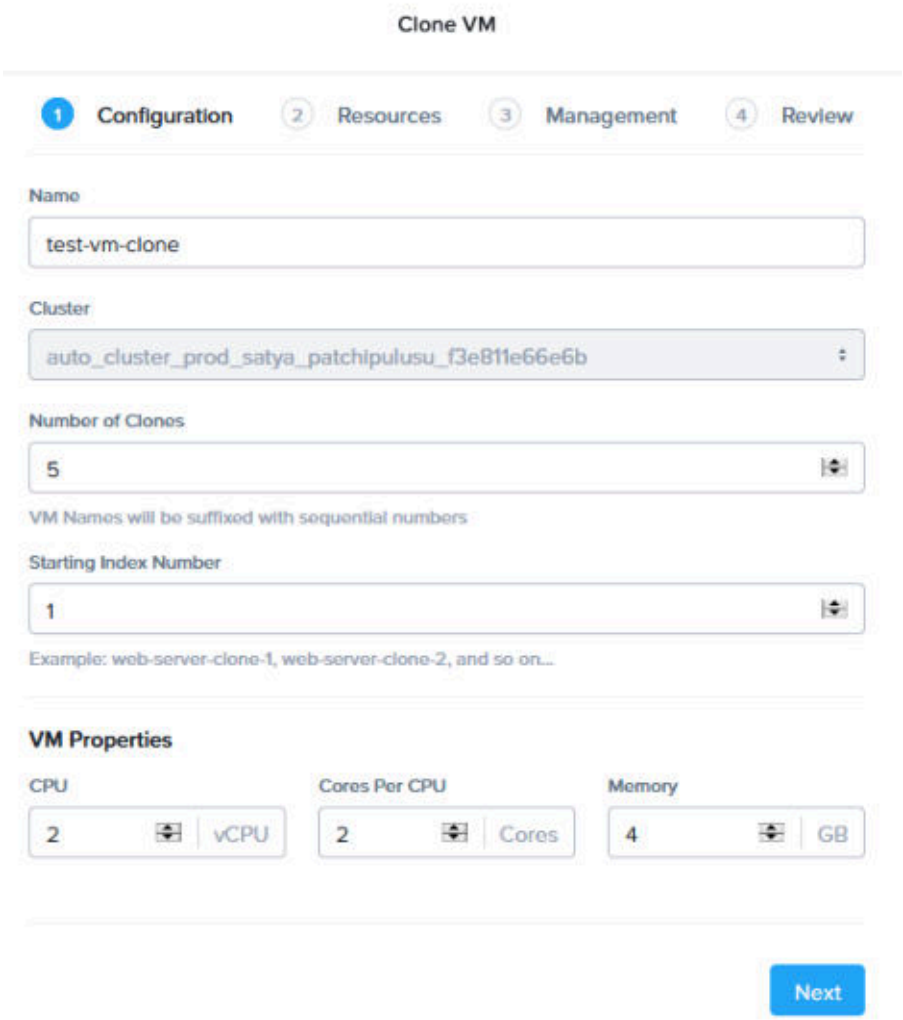

Figure 208: Clone VM Window (Configuration)

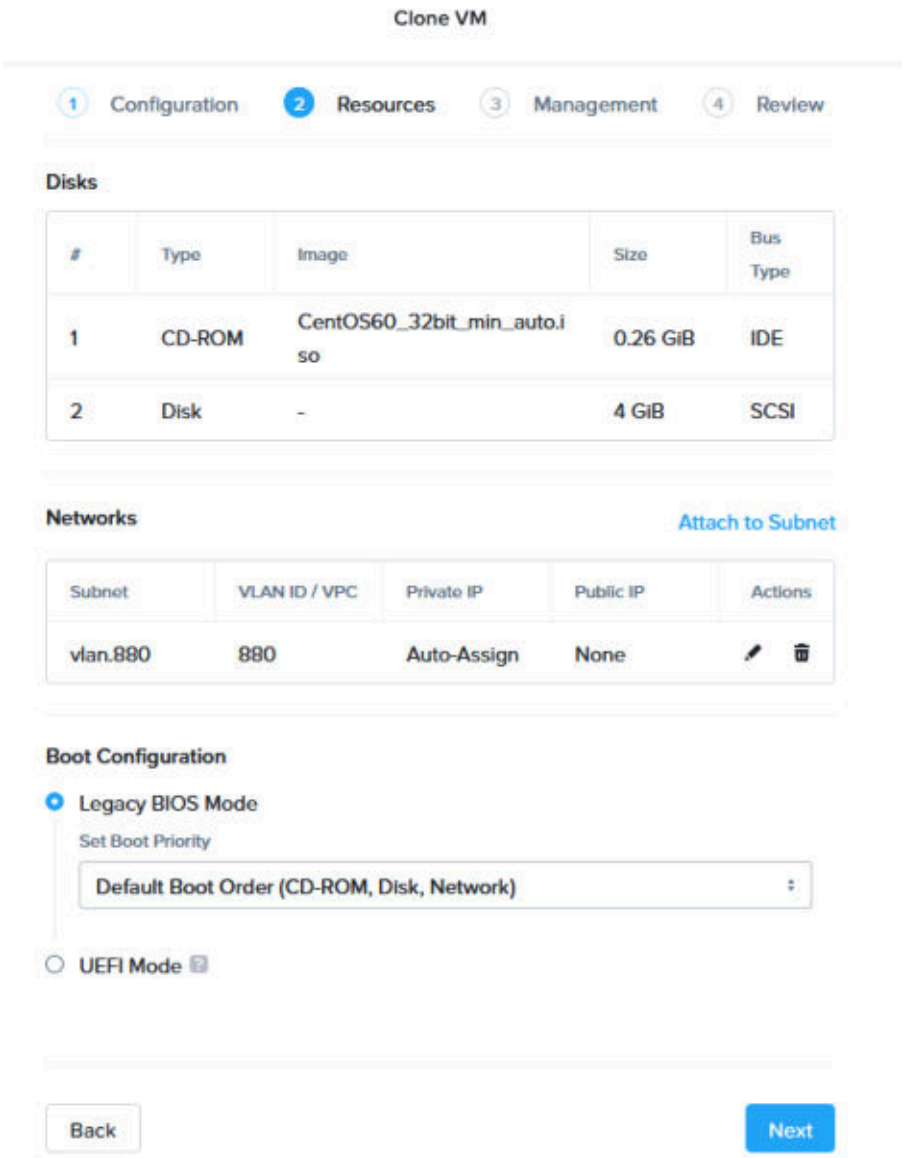

Figure 209: Clone VM Window (Resources)

Figure 210: Clone VM Window (Management)

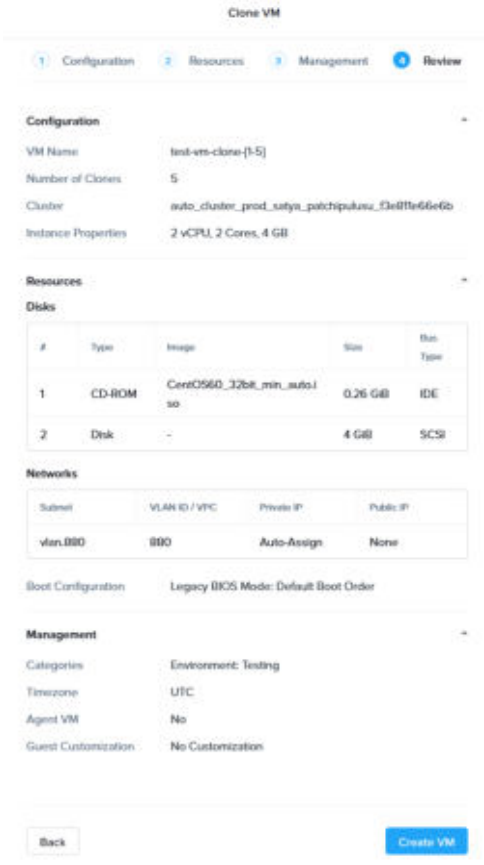

Figure 211: Clone VM Window (Review)

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

  > Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

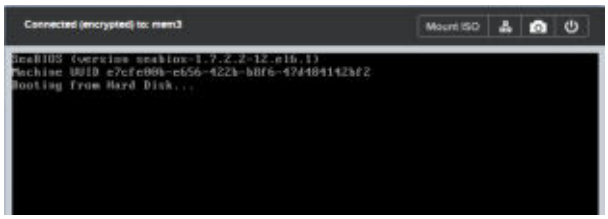

Figure 212: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To create VM recovery point, select **Create Recovery Point**.

  This displays the Create VM Recovery Point dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

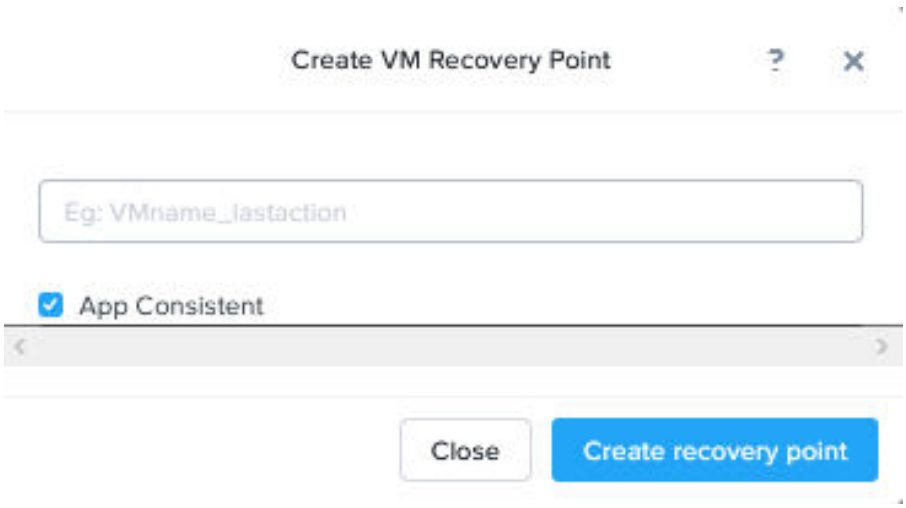

Figure 213: Create VM Recovery Point Window

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the "Explicitly Adding a VM to a Recovery Plan" section in the Leap Administration Guide).

- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 606).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 415).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 372).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 374).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > Note: If you clone a VM, by default NGT is not enabled on the cloned VM. You need to again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the *Prism Web Console Guide*.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

  ```
  ncli> ngt mount vm-id=virtual_machine_id
  ```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

  ```
  ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-
  c1601e759987
  ```

  > Note:
  >
  > - Self-service restore feature is not enabled by default on a VM. You need to manually enable the self-service restore feature.
  > - If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.
- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 375).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

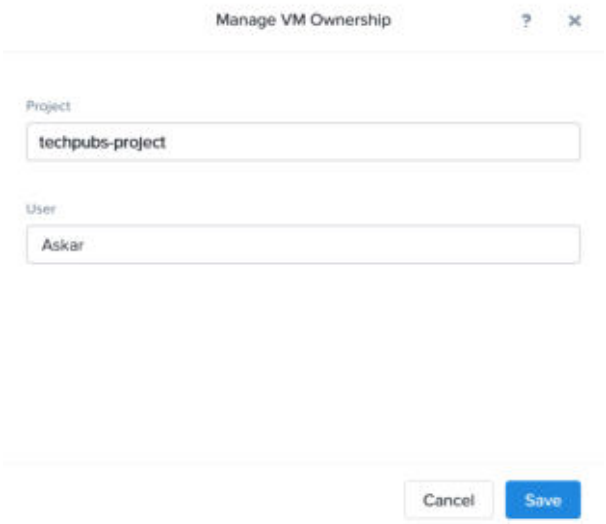

Figure 214: VM Ownership Window

- To export as OVA, select **Export as OVA** (see OVA Management on page 433 and Uploading an OVA on page 434).

Adding Multiple vGPUs to the Same VM

About this task

You can add multiple vGPUs of the same vGPU type to:

- A new VM when you create it.
- An existing VM when you update it.

> Important:
>
> Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.

After you add the first vGPU, do the following on the Create VM or Update VM dialog box (the main dialog box) to add more vGPUs:

Procedure

1. Click **Add GPU**.

2. In the Add GPU dialog box, click **Add**.

   The License field is grayed out because you cannot select a different license when you add a vGPU for the same VM.

   The **VGPU Profile** is also auto-selected because you can only select the additional vGPU of the same vGPU type as indicated by the message at the top of the dialog box.

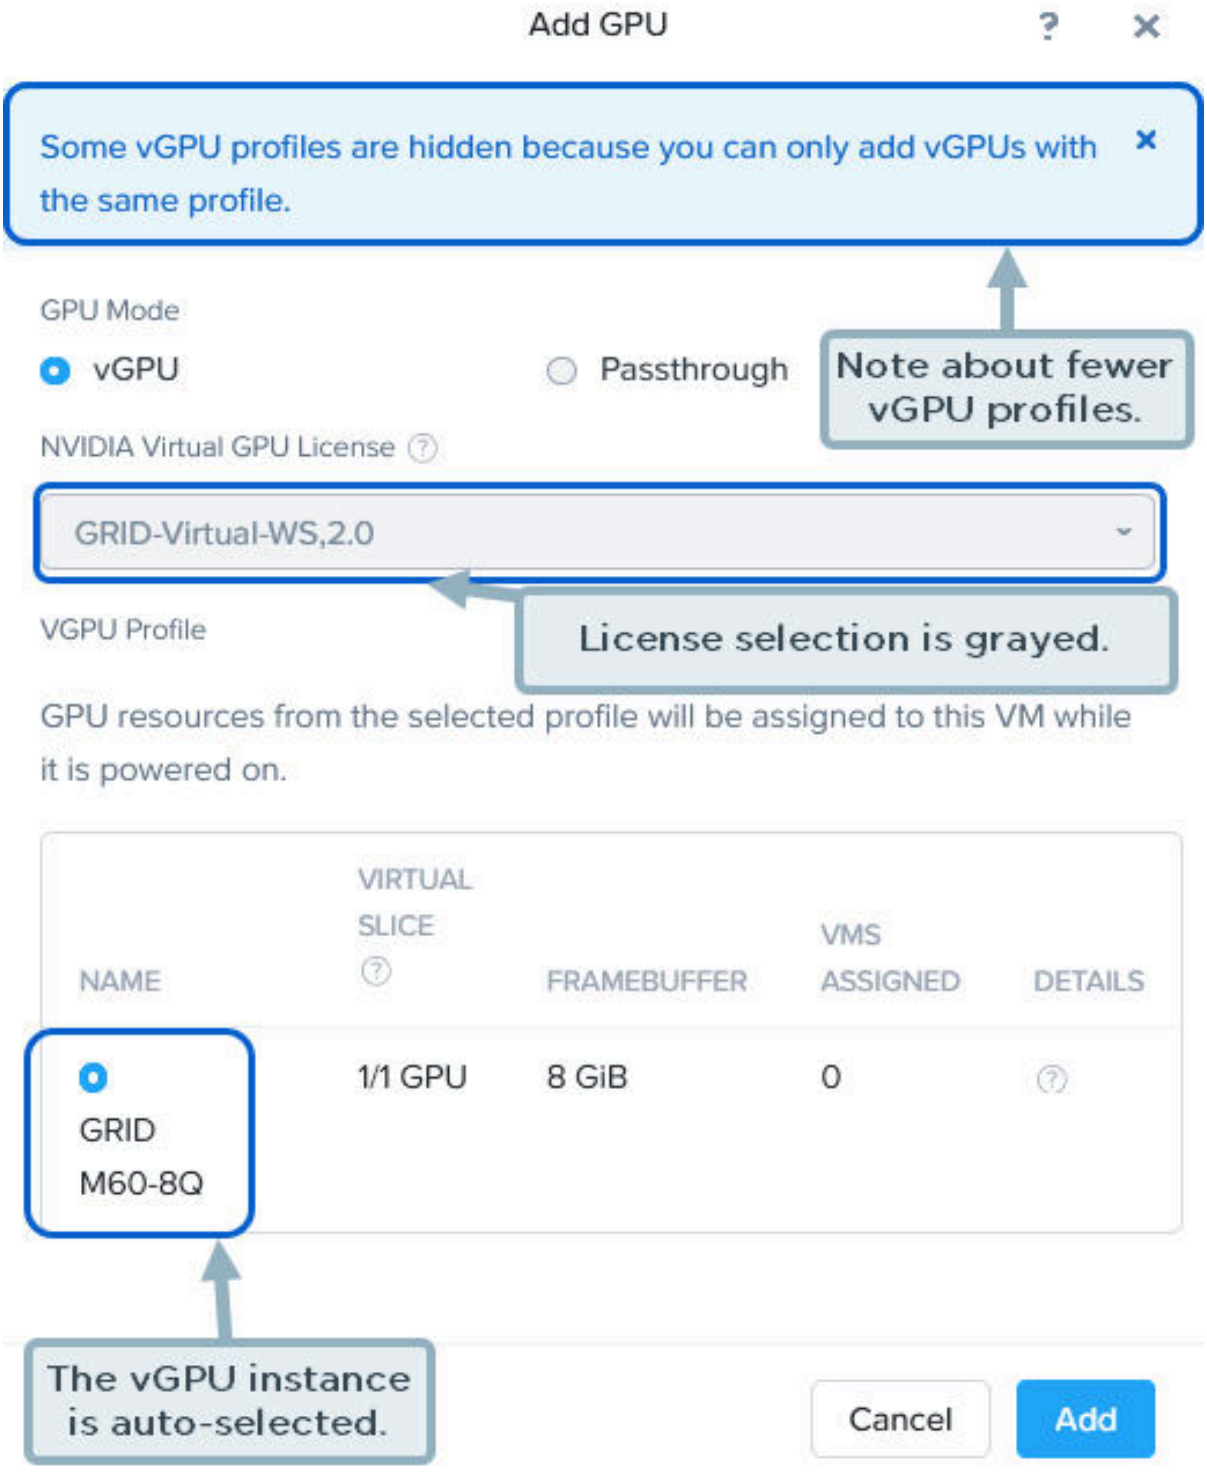

Figure 215: Add GPU for multiple vGPUs

3. In the main dialog box, you see the newly added vGPU.

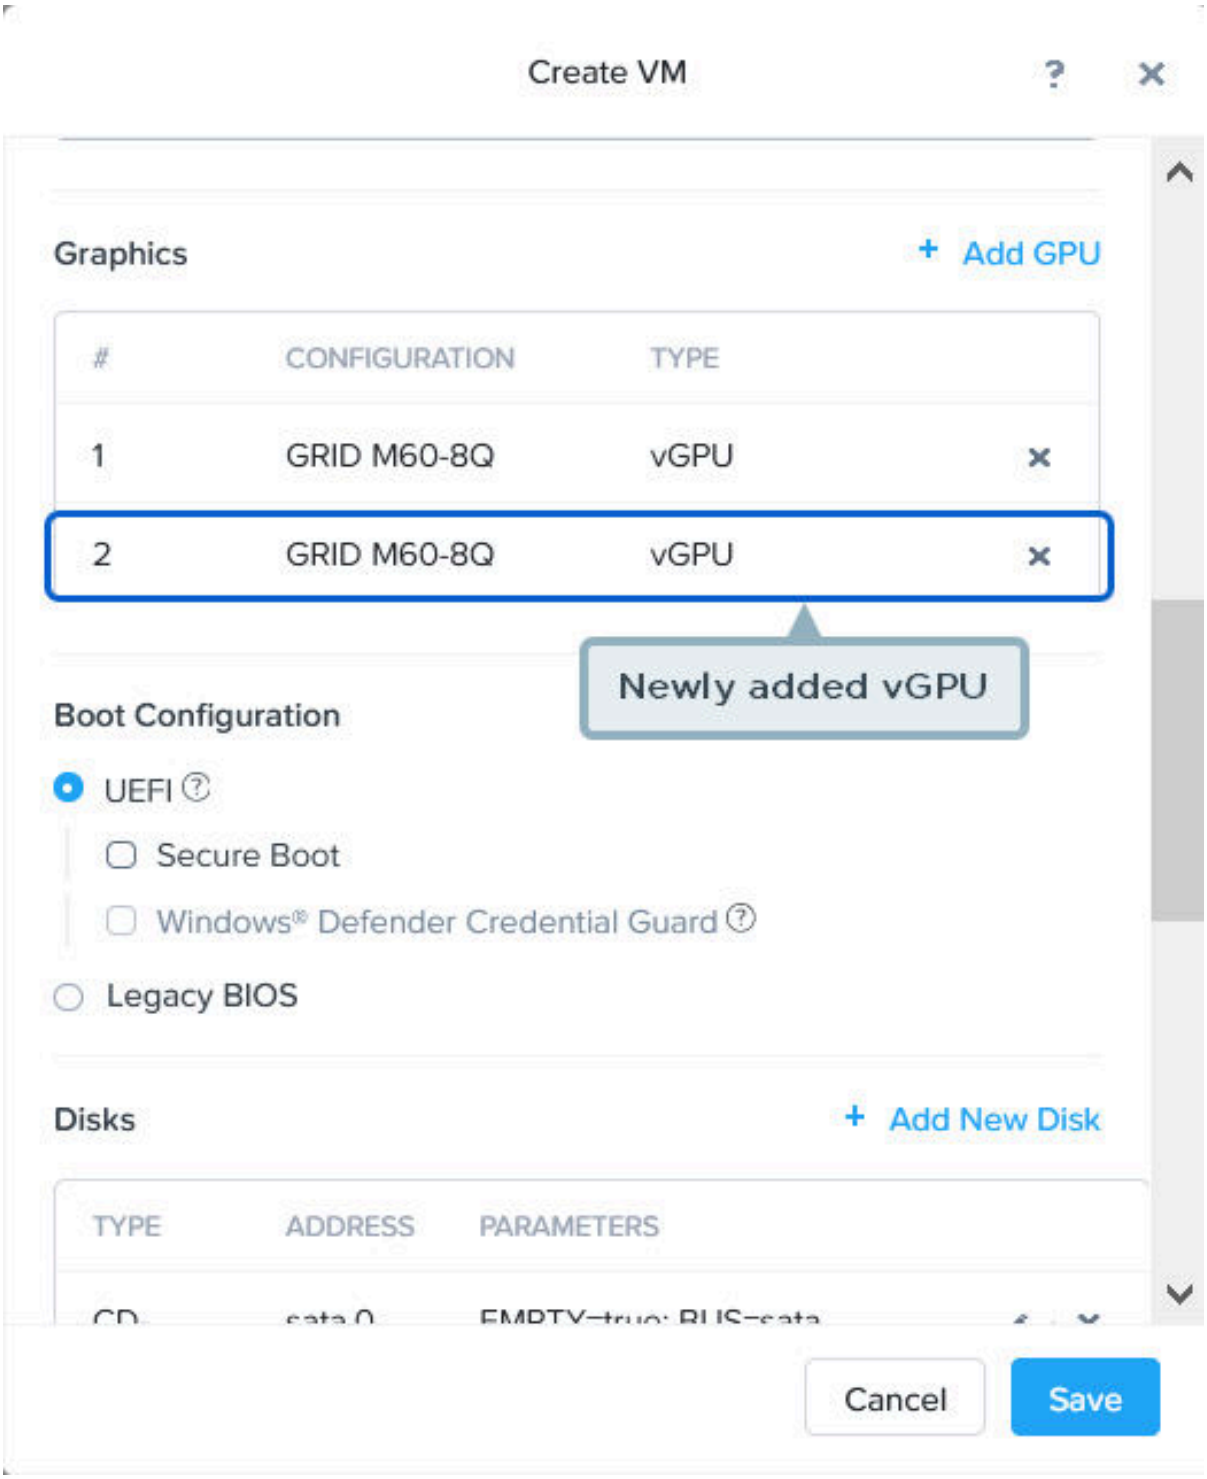

Figure 216: New vGPUs Added

4. Repeat the steps for each vGPU addition you want to make.

Migrating Live a VM with Virtual GPUs

About this task

A vGPU-enabled VM can be migrated on both the Prism Element and the Prism Central. You can migrate the vGPU-enabled VM to another host in the same cluster.

> Note:
>
> You cannot migrate the VM to a host in another cluster. The **Migrate Outside Cluster** option is not available in the Migrate VM dialog box, if the VM you want to migrate has vGPUs configured.
>
> Live migration of VMs with vGPUs is supported:
>
> - For vGPUs created with NVIDIA Virtual GPU software version 10.1 (440.53).
> - On AOS 5.18.1 or later.

Before you begin

Ensure the following:

1. The VM is not powered off.
2. The host affinity is not set for the VM.

   If the host affinity of the VM is set to only one host, you cannot migrate the VM. However, based on the GPU resources required, if the host affinity is set to multiple hosts with the same or similar GPU resources, you can migrate the VM among the hosts with which the affinity is set.
3. You have another host in the same cluster to migrate the VM.

For limitations applicable to live migration of vGPU-enabled VMs, see *Limitations for Live Migration* in the *NVIDIA GRID Virtual GPU Support on AHV* section of the *AHV Administration Guide*.

Procedure

1. To migrate the VM, go to the **Virtual Infrastructure** > **VMs** > **List**.

2. Select the VM you want to migrate live. Click **Actions** > **Migrate** > **.**

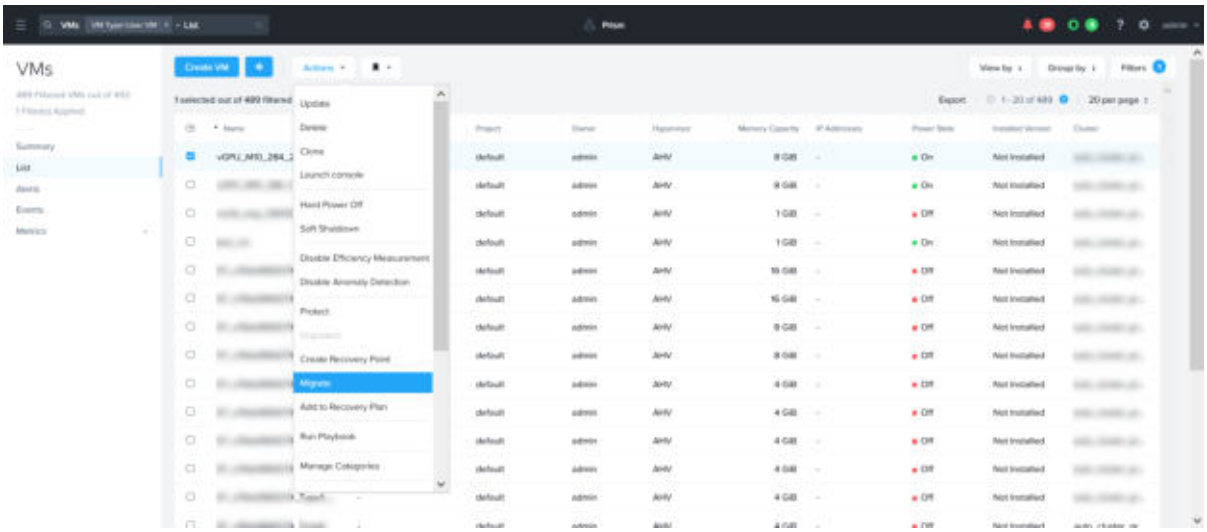

Figure 217: VM Selection

3. In the Migrate VM dialog box, do the following:

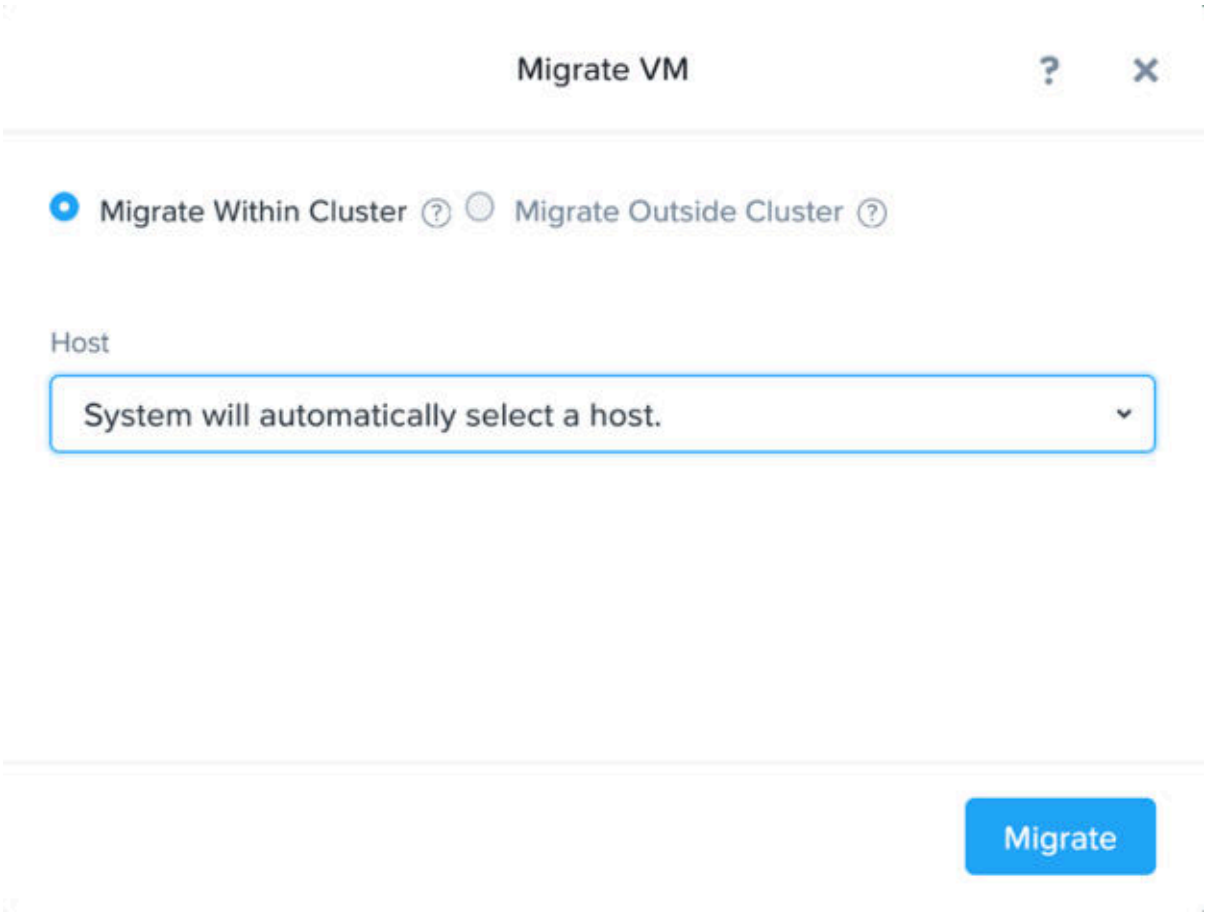

Figure 218: Migrate VM dialog box

a. The **Migrate Within Cluster** option is auto-selected.

b. In the **Host** drop-down list, do one of the following:

   » Retain the *System will automatically select a host* option if you want to migrate the VM to a host selected by the system.

   The system selects a host based on the GPU resources available with the host as appropriate for the VM to be migrated live.

   » Select the host to which you want to migrate the VM

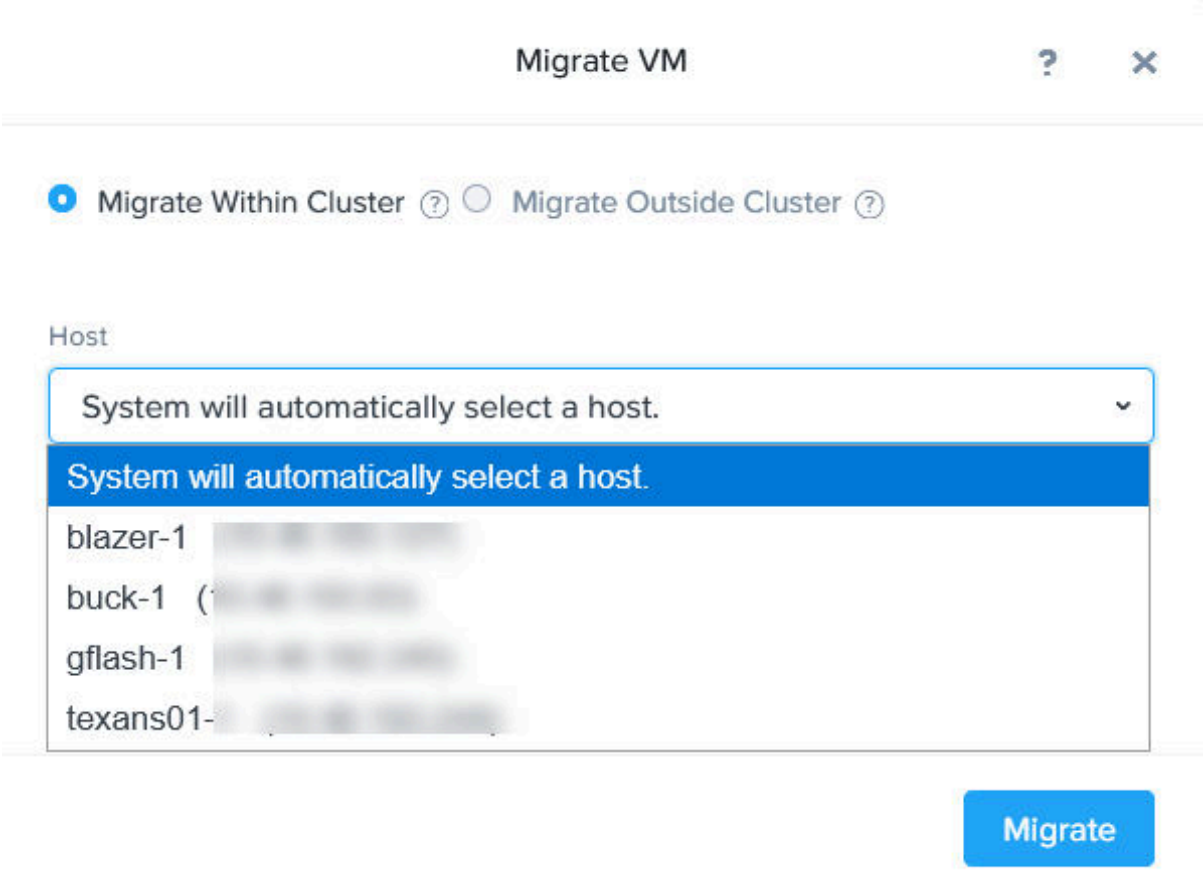

Figure 219: Host Selection

4. Click **Migrate**.

   Prism submits the task and displays the following message:

   ```
   Successfully submitted migrate operation.
   Task details
   ```

   Task details is a link to the Tasks page. Click the **Task details** link to monitor the migration task in the Tasks page.

   The host name of the VM in the List view changes to the host name to which you migrated the VM.

Creating a VM (ESXi)

In ESXi clusters, you can create a new virtual machine (VM) through the Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration on page 452 before proceeding.
- Register the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 453.

About this task

To create a VM, do the following:

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Create VM** button.
   The **Create VM** wizard appears.
2. In the Configuration step, do the following in the indicated fields:
   a. **Name**: Enter a name for the VM.
   b. **Description** (optional): Enter a description for the VM.
   c. **Cluster**: Select the target cluster from the pull-down list on which you intend to create the VM.
   d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.
   e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
   f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
   g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

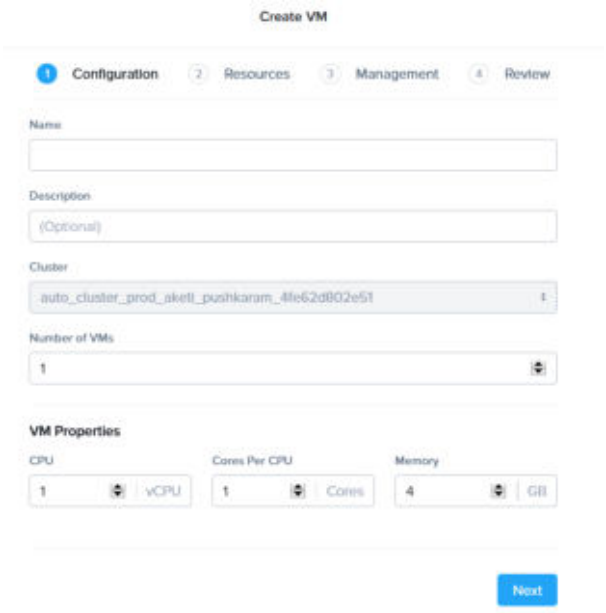

Figure 220: Create VM Window (Configuration)

3. In the **Resources** step, do the following.

   Disks: To attach a disk to the VM, click the **Attach Disk** button. The Add Disks dialog box appears. Do the following in the indicated fields:

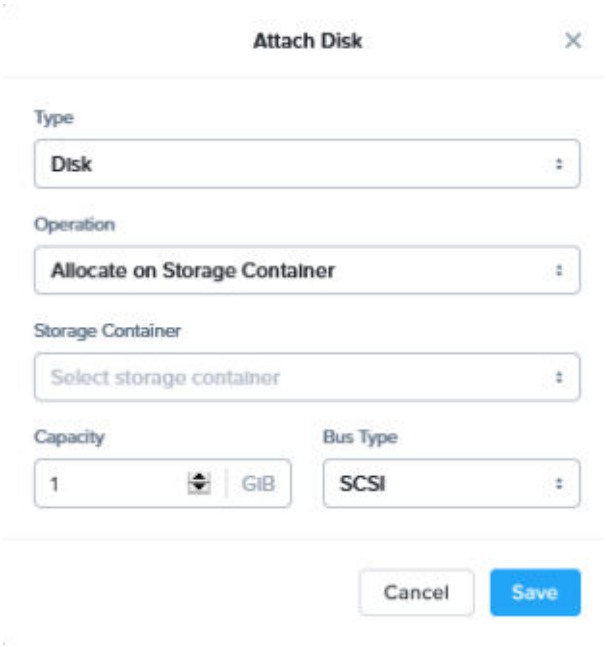

   Figure 221: Add Disk Dialog Box

   a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the pull-down list.

   b. **Operation**: Specify the device contents from the pull-down list.

      - Select **Clone from NDSF file** to copy any file from the cluster that can be used as an image onto the disk.
      - [CD-ROM only] Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
      - [Disk only] Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
      - Select **Clone from Image** to copy an image that you have imported by using image service feature onto the disk.

   c. **Bus Type**: Select the bus type from the pull-down list. The choices are **IDE**, **SCSI**, **PCI**, or **SATA**.

   d. **Path**: Enter the path to the desired system image.

   e. **Clone from Image**: Select the image that you have created by using the image service feature.

      This field appears only when **Clone from Image Service** is selected. It specifies the image to copy.

      Note: If the image you created does not appear in the list, see KB-4892.

f. **Storage Container**: Select the storage container to use from the pull-down list.

   This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.

g. **Capacity**: Enter the disk size in GiB.

h. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the Create VM dialog box.

i. Repeat this step to attach additional devices to the VM.

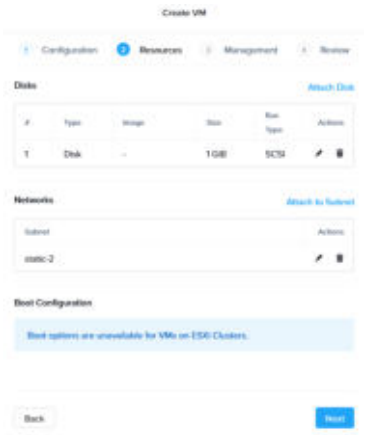

Figure 222: Create VM Window (Resources)

4. In the Resources step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The Attach to Subnet dialog box appears.

   Do the following in the indicated fields:

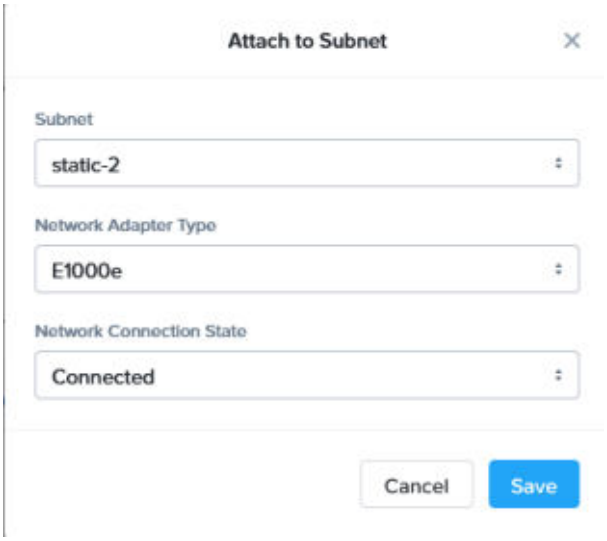

Figure 223: Attach to Subnet Dialog Box

   a. **Subnet**: Select the target virtual LAN from the pull-down list.

      The list includes all defined networks (see Configuring Network Connections on page 423).

   b. **Network Adapter Type**: Select the network adapter type from the pull-down list.

      See vCenter Server Integration on page 452 for the list of supported adapter types.

   c. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are *Connected* or *Disconnected*.

   d. When all the field entries are correct, click the **Add** button to create a network interface for the VM and return to the Create VM dialog box.

   e. Repeat this step to create more network interfaces for the VM.

5. In the **Management** step, do the following.

   a. **Categories**: Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

   b. **Guest OS**: Type and select the guest operating system.

      The guest operating system that you select affects the supported devices and number of virtual CPUs available for the virtual machine. The Create VM wizard does not install the

guest operating system. See the list of supported operating systems in vCenter Server Integration on page 452 topic.

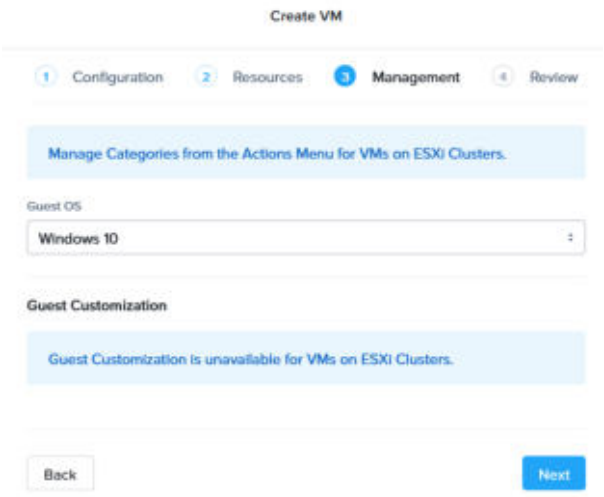

Figure 224: Create VM Window (Management)

6. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the Create VM dialog box.

   The new VM appears in the VMs entity page list.

Managing a VM (ESXi)

You can manage virtual machines (VMs) in an ESXi cluster through Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration on page 452 before proceeding.
- Ensure that you have registered the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 453.

About this task

After creating a VM (see Creating a VM (ESXi) on page 357), you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, power on (or off) the VM, enable flash mode for a VM, assign the VM to a protection policy, create VM recovery point, add the VM to a recovery plan, run a playbook, manage categories, install and manage Nutanix Guest Tools (NGT), manage the VM ownership, or configure QoS settings.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 82) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
- Go to the details page of a selected VM (see VM Details View on page 90) and select the desired action.

Note: The available actions appear in bold; other actions are unavailable. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Make the desired changes and then click the **Save** button in the **Review** step.

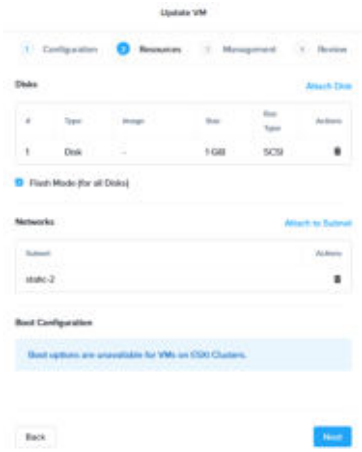

Figure 225: Update VM Window (Resources)

- **Disks:** You can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM (ESXi) on page 357 before you create a new disk for a VM. You can enable or disable the flash mode settings for the VM. To enable flash mode on the VM, click the **Enable Flash Mode** check box. After you enable this feature on the VM, the status is updated in the VM table view.
- **Networks:** You can attach new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (ESXi) on page 357 before you modify NIC network or create a new NIC for a VM.
- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

> Note:
>
> - You can clone up to 250 VMs at a time.
> - You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

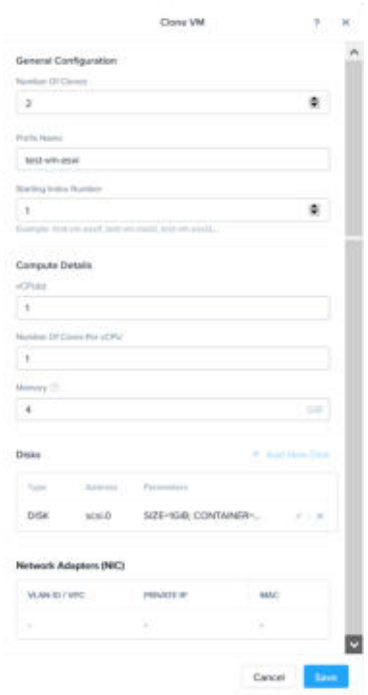

Figure 226: Clone VM Window

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

> Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

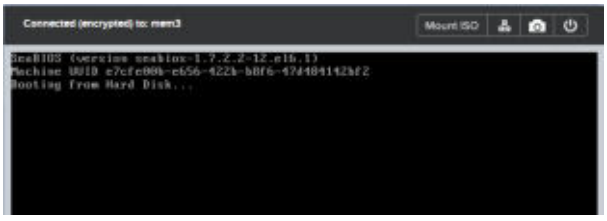

Figure 227: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).

- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To assign the VM to a protection policy, select **Protect**. This opens a page to specify the protection policy to which this VM should be assigned (see Policies Management on page 478). To remove the VM from a protection policy, select **Unprotect**.

  Note: You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Leap and configuring the Availability Zone.

- To migrate the VM to another host, select **Migrate**.

  This displays the Migrate VM dialog box. Select the target host from the pull-down list (or select the **System will automatically select a host** option to let the system choose the host) and then click the **Migrate** button to start the migration.

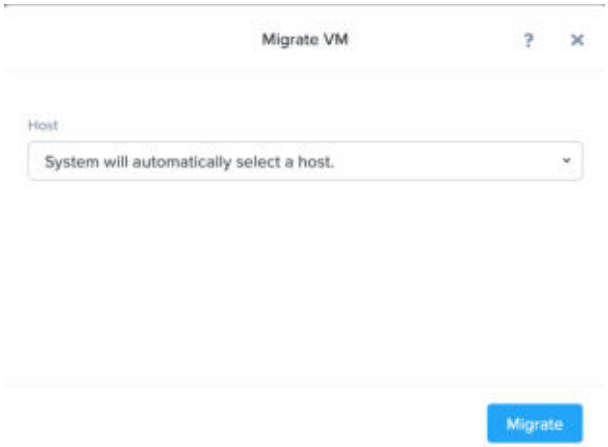

Figure 228: Migrate VM Window

Note: Nutanix recommends to live migrate VMs when they are under light load. If they are migrated while heavily utilized, migration may fail because of limited bandwidth.

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the "Explicitly Adding a VM to a Recovery Plan" section in the Leap Administration Guide).
- To create VM recovery point, select **Create Recovery Point**.

  This displays the Create VM Recovery Point dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling

the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

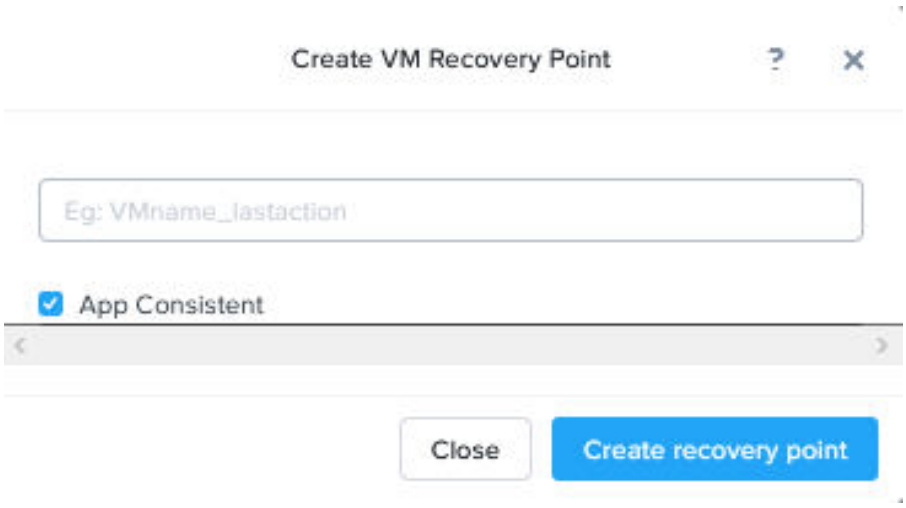

Figure 229: Create VM Recovery Point Window

- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 606).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 415).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 372).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 374).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > Note: If you clone a VM, by default NGT is not enabled on the cloned VM. You need to again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the *Prism Web Console Guide*.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

  ```
  ncli> ngt mount vm-id=virtual_machine_id
  ```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

  ```
  ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-
  c1601e759987
  ```

  > Note:
  >
  > - Self-service restore feature is not enabled by default on a VM. You need to manually enable the self-service restore feature.
  > - If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST

APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.

- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 375).
- To establish VM host affinity, select **Configure VM Host Affinity**.

  A window appears with the available hosts. Select (click the icon for) one or more of the hosts and then click the **Save** button. This creates an affinity between the VM and the selected hosts. If possible, it is recommended that you create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure. See Affinity Policies for AHV on page 416 for more information about the VM affinity policies.

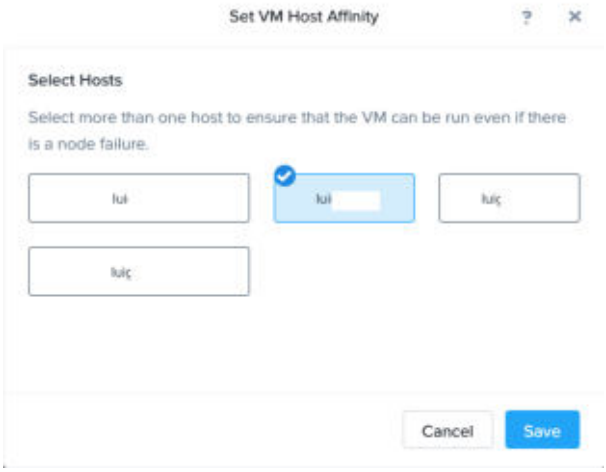

Figure 230: Set VM Host Affinity Window

- To add a VM to the catalog, select **Add to Catalog**. This displays the **Add VM to Catalog** page (see Adding a Catalog Item on page 391).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

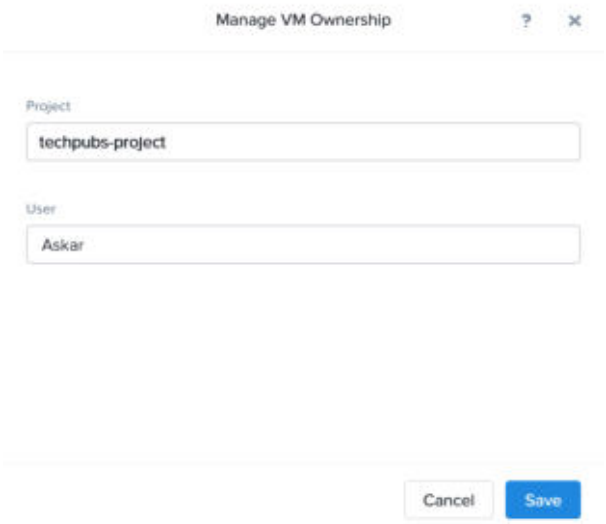

Figure 231: VM Ownership Window

- To configure quality of service (QoS) settings, select **Set QoS Attributes** (see Setting QoS for an Individual VM on page 378).

Exporting a VM as an OVA

Before you begin

If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the destination hypervisor.

About this task

As a Super Admin, you can export a VM as an OVA by selecting the **Export** as OVA action from the **Actions** dropdown list. You can export a VM as an OVA with a specified disk format for the disks configured on the VM.

> Note:
>
> - The Super Admin role was formerly called as User Admin.
> - If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the target destination hypervisor.

To export a VM as an OVA, do the following in the VMs dashboard:

Procedure

1. Select the VM in List tab.

   You can click the VM and open the VM details page.

2. In the **Actions** list, select **Export as OVA**.

   Alternatively, on the details page of the VM, select **More** and click **Export as OVA**.

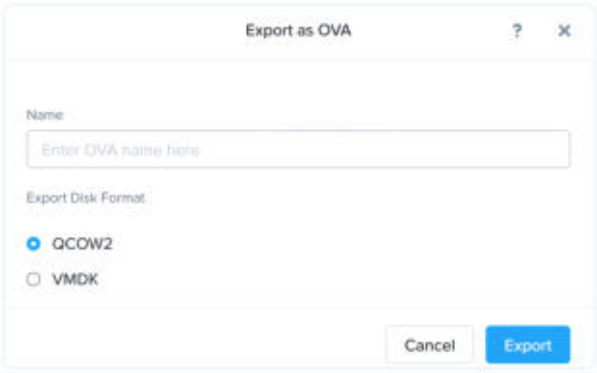

Figure 232: Export as OVA dialog box

3. In the Export as OVA window, do the following in the indicated fields:

   a. Enter a name for the OVA, in the **Name** field.

   b. Select the disk format that you want to format the disks in the VM to, in the **Export Disk Format** section.

      The default format is QCOW2. If you want to use the OVA to deploy a VM with disks formatted as VMDK, then select VMDK as the disk format.

4. Click **Export**.

   You can check the progress of the *Export as OVA* task in Tasks.

   The exported OVA is available on the OVA dashboard.

NGT Management in Prism Central

The NGT management in Prism Central feature lets you to select multiple VMs in the Prism Central from the VM Entity browser and then install, upgrade, or manage NGT applications on these VMs. You can use this feature for large-scale deployments. If you use this feature, you do not have to manually login to the VM and do the NGT installation as the installation of NGT is fully automated. Following are other important characteristics of this feature.

- Ability to enable VSS, SSR applications as part of the install workflow.
- Ability to select multiple VMs in the Prism Central from the VM Entity browser and upgrade NGT on these VMs.
- Ability to defer restarting of the VMs to a specified later time and date after NGT install or upgrade.
- Ability to define policies that allow you to defer restarting of the VMs to a specified later time and date.

For more information on the Nutanix Guest Tools, see the *Prism Web Console Guide*.

NGT Management in Prism Central Requirements

Following are the general requirements, specific requirements for Linux and Windows.

General Requirements

- Supported on the AHV and ESXi VMs.
- For AHV, both PC and PE should at least be on 5.9 version.
- For ESXi, both PC and PE should at least be on 5.11 version.
- VMs and virtual IP address of the cluster should be able to communicate with each other.
- The antivirus scanners must exclude the C:\Program Files\Nutanix directory for successful NGT installation.
- If you clone the VM from the VM where NGT was installed by using NGT Management in Prism Central feature, see Cloning from VMs With NGT on page 374 topic for how to enable NGT on the cloned VM.

The following TCP ports must be open between the Controller VMs and user VMs with NGT installed:

Port	Description	Inbound/ Outbound	Source	Destination
CVM Ports				
9440	To run Prism Element, Prism Central, and APIs.	Inbound and Outbound	Prism Central and Prism Element	Prism Central and Prism Element
2100	For Prism Central cluster creation.	Inbound	Prism Element	Prism Central
2073	For communication between NGT and other services in the cluster.	Inbound and Outbound	NGT leader node IP address and any other service in the cluster	NGT leader node IP address and any other service in the cluster
2074	Requests from user VMs come to this port in the CVM.	Inbound and Outbound	User VM IP address and CVM virtual IP address	User VM IP address and CVM virtual IP address
User VM Ports				
5986	Applies to Windows VMs only. The firewall rule must allow inbound connection over SSL on this port.	Inbound	CVM IP address	Target Windows user VM IP address

Port	Description	Inbound/ Outbound	Source	Destination
22 or any SSH port	Applies to Linux VMs only.	Inbound	NGT leader node IP address	Target Linux user VM IP address
23578	This is an internal port and required for inter-process communication (IPC) inside the user VM.	NA	NA	NA

Note: All ports are TCP.

For the complete list of required ports, see Port Reference.

Requirements for the Linux VMs

- The lvm2 package version should be 2.02.109 or later version.
- CVM should have SSH access to the VM.
- Ensure that you have at least 5 free SSH session configured in MaxSessions by navigating to /etc/ssh/sshd_config.
- The user should have no password sudo access on the VM.
- Ensure that you have disabled requiretty for a particular user.

```
sudo sed -i 's/Defaults    requiretty/ #Defaults    requiretty/g' /etc/sudoers
```

- For Linux operating systems, NGT is installed at /usr/local and NGT ISO is mounted on /mnt/nutanix/ngt. Hence, ensure that the user has write permissions to these locations.

Requirements for the Windows VMs

- PowerShell version on the Windows VMs should at least be 3.0.
- Local VM accounts with administrative privileges are required.
- PowerShell path must be available in the environment for NGT installation. Otherwise, NGT install fails.
- Windows Remote Manager Service (winrm) should be running.
- The winrm service with ssl configuration should be configured. Run the following commands from the Windows PowerShell as an administrator to configure winrm with SSL.

  For example, you can use the following commands to configure the winrm service.

```
> $certificate = New-SelfSignedCertificate -DnsName $env:computername -CertStoreLocation cert:\LocalMachine\My
```

```
> winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$env:computername`";CertificateThumbprint=`"$($certificate.ThumbPrint)`"}"
```

```
> cmd /c 'winrm set winrm/config/service/auth @{Basic="true"}'
```

```
> netsh advfirewall firewall add rule name=\"WinRM-HTTPS\" dir=in localport=5986 protocol=TCP action=allow
```

Installing NGT on Multiple VMs

Perform the following procedure to install NGT on multiple VMs at the same time by using the NGT management in Prism Central feature.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 82).
   The NGT status column displays whether the NGT is installed on the VM, if NGT is installed whether it is the latest version, or whether an upgrade is available or not.
2. Select the VMs on which you want to install NGT and from the **Actions** drop-down menu, and then click **Install NGT**.
   The Install NGT windows displays the number of VMs on which you can install NGT. For example, if you have selected a VM on which NGT is already installed, this VM will be skipped during this operation.

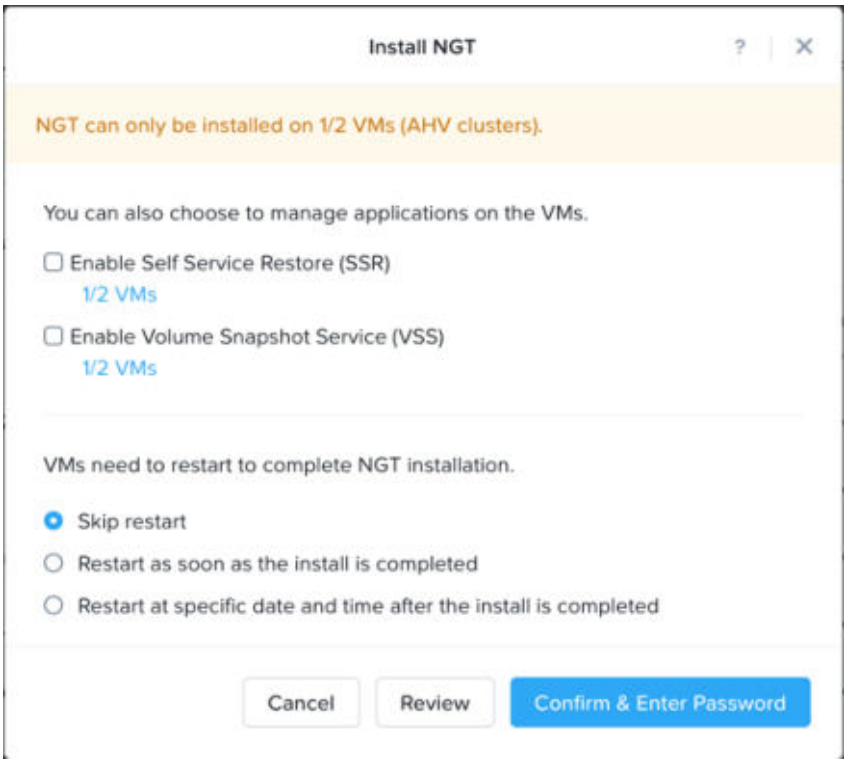

Figure 233: Installing NGT

3. To enable SSR and VSS applications, select the **Enable Self Service Restore (SSR)** and **Enable Volume Snapshot Service (VSS)** checkboxes.
   If you do not want to enable these features, you can do it later by following Managing NGT Applications on page 374.

   The VMs may have to be restarted to complete NGT installation because virtIO drivers are packaged with NGT ISO. When you install NGT on the VM, the virtIO drivers are also installed and hence VM restart is required. NGT operations does not require a restart. If you choose

to restart the VM at later point of time, all the NGT functionality will be available except the virtIO driver. VirtIO driver functionality comes into effect only after restart.

Select the restart schedule of the VMs.

- **Skip restart**: If you do not want the VMs to be restarted post NGT installation or if you want the category policy to be applied, select this option. For more information on the policies, see Creating a New Restart Policy on page 478.
- **Restart as soon as the install is completed**: The VMs will be restarted after the installation is completed.
- **Restart at specific date and time after the install is completed**: The VMs will be restarted at the specific date and time that you have configured in the **Date** and **Time** fields.

4. (Optional) Click **Review** to view your configurations.
5. Click **Confirm and Enter Password** to proceed with the installation.
   If all the VMs have same user names and password, you can enter it in the administrator **Username** and **Password** field. If the user names and passwords are different, you have to upload a JSON in the specified format.
6. Click **Done** to finish the installation process.
   Once the installation is completed, the **NGT Status** column displays NGT status as latest. This is an Async operation and takes some time for the status to be reflected in the Prism Central.

   - If you have not conform to the NGT prerequisites or if there is any issue with the installation process, the NGT installation fails and an alert message is displayed. Perform the installation process again to install NGT on the VMs.
   - You can also filter the list by clicking the **Filters** button. This displays a filter pane on the right of your screen. Click on the appropriate filter to list out the tasks that corresponds to that filter. The filters and their description are as follows.
     - Installed: Filters the VMs on which only NGT is installed.
     - SSR: Filters the VMs on which only SSR application is installed.
     - VSS: Filters the VMs on which only VSS application is installed.
     - VSS and SSR: Filters the VMs on which both VSS and SSR application is installed.
     - Not Installed: Filters the VMs on which NGT is not installed.

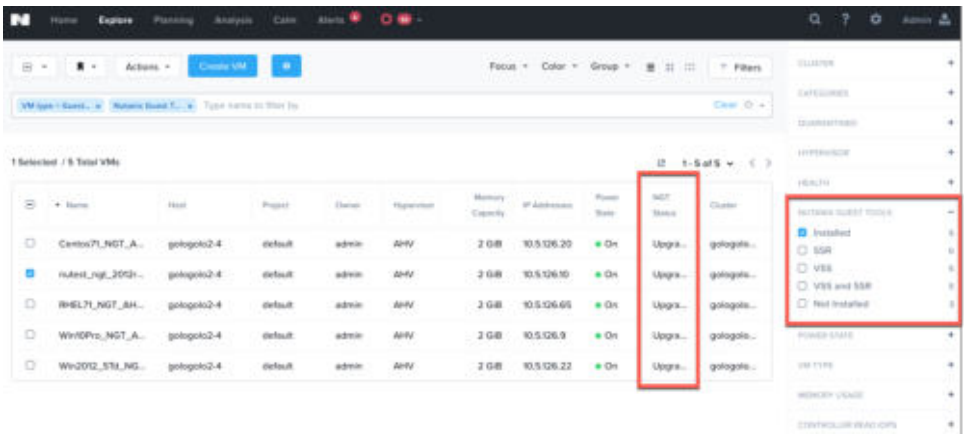

Figure 234: NGT Status and Filter Option

7. (Optional) If you do not want to install NGT and just want to mount NGT on the VM, click **Skip and Mount**.

   The VMs are registered with the NGT service. NGT is enabled and mounted on the selected virtual machines. A CD with volume label NUTANIX_TOOLS gets attached to the VMs. You can manually install the NGT by logging into the VM.

Cloning from VMs With NGT

For a VM that is cloned from an image or a VM that has NGT installed, perform the following procedure to enable NGT on these VMs. Since NGT is already present on the VM, you do not have to install NGT and only refresh of certifications is required.

Procedure

1. Select the cloned VMs and follow the NGT install process again to install NGT on the cloned VMs.
2. During the installation of the NGT on the cloned VMs, you can only mount NGT on the cloned VM by using **Skip and Mount** operation.

   For more information on selecting VMs and mounting NGT, see Installing NGT on Multiple VMs on page 372. NGT will be enabled and mounted on the cloned VMs.

   The NGT ISO will be mounted on the VM.
3. After the ISO is mounted, restart the VM.

   The NGT status is refreshed and updated after the VM starts.

Managing NGT Applications

If you have not enabled the VSS and SSR applications during the install workflow or if you want to change the application state, you can do it by using the following procedure. You can also disable these applications (if already enabled) by using the same workflow.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 82).

2. Select the VM or VMs on which you want to enable SSR and VSS applications by clicking **Manage NGT Applications**.

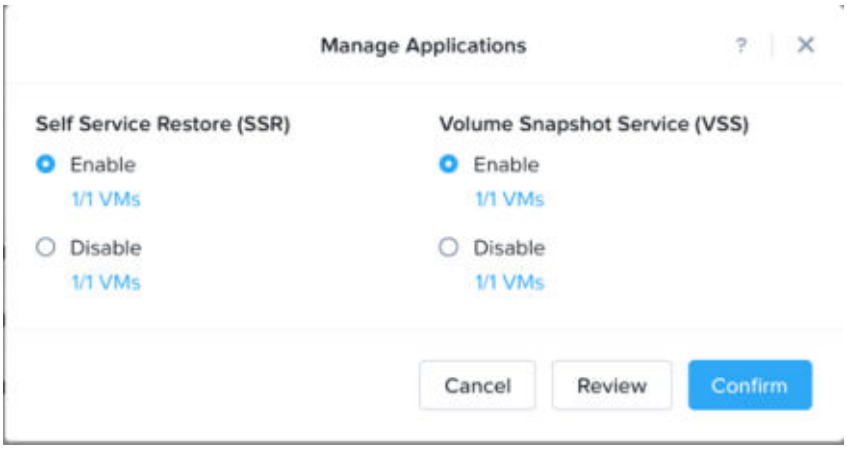

Figure 235: Manage NGT Applications

   The Manage Applications page displays whether you can enable these applications on the selected VMs or not.

3. Select **Enable or Disable** for the application, and click **Confirm**.

Upgrading NGT

Perform the following procedure to upgrade NGT on the VMs.

Before you begin

You can only upgrade from VMs that have NGT version is 1.2.0.3 and later (excluding 1.2.1, 1.2.2). If you have older versions of NGT, first upgrade NGT by using the procedure mentioned in the *Prism Web Console* guide.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 82).
   The NGT status column displays whether an upgrade is available for the selected VMs or not.

2. Select the VMs for which you want to upgrade NGT and from the **Actions** drop-down menu click **Upgrade NGT**.

   The Upgrade VM window displays the number of VMs on which you can upgrade NGT. For example, if you have selected a VM on which NGT is already installed and is latest, this VM is skipped during this operation.

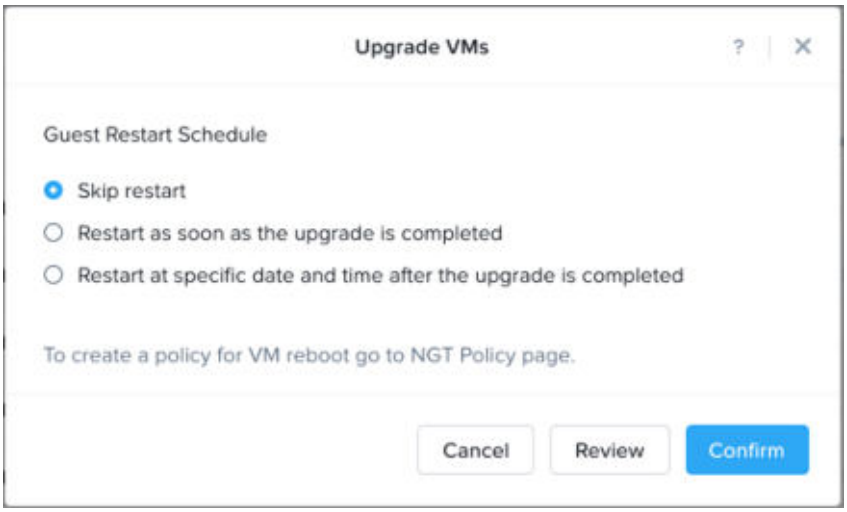

Figure 236: Upgrade NGT

3. Select the restart schedule of the VMs.

   - **Skip restart**: If you do not want the VMs to be restarted post NGT installation or if you want the category policy to be applied, select this option. For more information on the policies, see Creating a New Restart Policy on page 478.
   - **Restart as soon as the install is completed**: The VMs will be restarted after the upgrade is completed.
   - **Restart at specific date and time after the install is completed**: The VMs will be restarted at the specific date and time that you have configured in the **Date** and **Time** fields.

     Note: If you do not restart the VMs post upgrade and if the virtIO version has been upgraded, the upgraded functionality of the virtIO drivers will not be available.

4. (optional) Click **Review** to view your configurations.
5. Click **Confirm** to start the upgrade process.

   Note: If you have not conform to the prerequisites or if there is any issue with the upgrade process, the NGT upgrade fails and an alert message is displayed. Perform the upgrade process again to upgrade NGT on the VMs.

NGT Uninstallation

To uninstall NGT, you have to login into the VM and then perform the uninstallation process.

For more information on NGT uninstallation, see the *Disabling and Removing Nutanix Guest Tools* topic for the *Prism Web Console Guide*.

Storage Quality of Service (QoS)

Storage QoS provides administrators granular control to manage the performance of virtual machines and ensure that the system delivers consistent performance for all workloads. You

can use a controllable knob to limit the IOPS that the storage layer would serve for individual virtual machines. IOPS is the number of requests the storage layer can serve in a second. You can set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources.

Note:

- Storage QoS is applicable only if you have an AOS Pro or above license. See the *Nutanix Licensing Guide* for more information on AOS licenses.
- Setting storage QoS is not allowed while the AOS upgrade is in progress.

Terminology

Throttled IOPS

The maximum IOPS that the storage layer tries to admit over a sustained period. Applications that are running on the VM are not permitted to exceed this level. This prevents VMs from affecting a system beyond the set limits. However, if the system is in a state of flux, your application might not be throttled. To prevent application crashing due to misconfiguration, the throttled IOPS setting must not be less than one hundred. Otherwise, an error is raised.

Setting maximum IOPS on a VM may affect the cluster latency for that VM and may result in higher latency being displayed in the cluster latency charts. Performance of any other VM is not affected by this setting. You can view individual VM latency graph to ensure that the other VMs are not affected. You can also create a custom focus view to display the throttled IOPS for all the VMs.

Throughput

Throughput is the amount of data that an application sees from the storage layer underneath per second. The maximum throughput is calculated by multiplying the number of IOPS with the block size.

Calculation of the IOPS and Throughput

The default block size used for IOPS calculation is 32 KiB. If the block size is less than or equal to 32 KiB, then the system is rate limited by the number of IOPS and the bandwidth is calculated by multiplying the number of IOPS with the block size of the application. Otherwise, if the block size is greater than 32 KiB, the system is rate limited by the throughput and the number of IOPS is calculated by dividing the throughput by block size of the application.

For example, if you have set a throttle limit of 800 IOPS on a VM and the application is admitting IOPS with 8 KiB block size, then the application would see 800 IOPS and throughput of 6.25 MBps. However, if application is admitting IOPS with 64 KiB block size, then the application would see 400 IOPS and a throughput of 25 MBps. For more information about the block size, IOPS, and throughput relationship when throttle limit it set to 800 IOPS, see *Relationship between Block Size, IOPS, and Throughput* table.

Table 217: Relationship between Block Size, IOPS, and Throughput

Block size	IOPS	Throughput
8 KiB	800	6.25 MBps
16 KiB	800	12.50 MBps
32 KiB	800	25 MBps

Block size	IOPS	Throughput
64 KiB	400	25 MBps
128 KiB	200	25 MBps

Limitations

Storage QoS has the following limitations:

- Storage QoS is not supported for an in-place restore, out-of-place restore, and snapshot operation on VM.
- While creating a clone of a VM, any QoS attributes throttle limit set on the original VM might not be applied to the new cloned VM. QoS throttle limit has to be set on the new cloned VM separately.
- For linked clones, when you are sharing the vDisks across multiple VMs, shared disks are not a part of QoS policy and shared disks would not be throttled. Each VM might exceed the maximum limit depending on the shared vDisk usage. For example, during bootstorm in a VDI setup.
- If a VM has volume groups attached, then QoS is not applicable.
- Storage QoS is not supported for Metro Availability, synchronous replication, and Nutanix Files virtual machines.

Setting QoS for an Individual VM

About this task

Storage QoS allows you to set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources. Perform the following procedure to set the QoS for an individual VM.

Figure 237: Setting QoS for an Individual VM

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275779955001 on Brightcove.

Procedure

1. Go to the **List** tab of the VM dashboard (see VMs Summary View on page 82) and select the target VM.
2. Do one of the following:
   - » From the **Actions** drop-down list, select **Set QoS Attributes**.
   - » Click the VM and then select **Set QoS Attributes** from the **More** drop-down list.
3. Under **Input Type**, select one of the following:
   - » IOPS (For more information about IOPS, see *Throttled IOPS* in *Storage Quality of Service (QoS)*.)
   - » Throughput (For more information about throughput, see *Throughput* in *Storage Quality of Service (QoS)*.)

4. If you selected **IOPS**, enter a value for the throttled IOPS in the **Throttled IOPS**.
5. If you selected **Throughput**, enter a value for the throttled throughput in the **Throttled Throughput**.

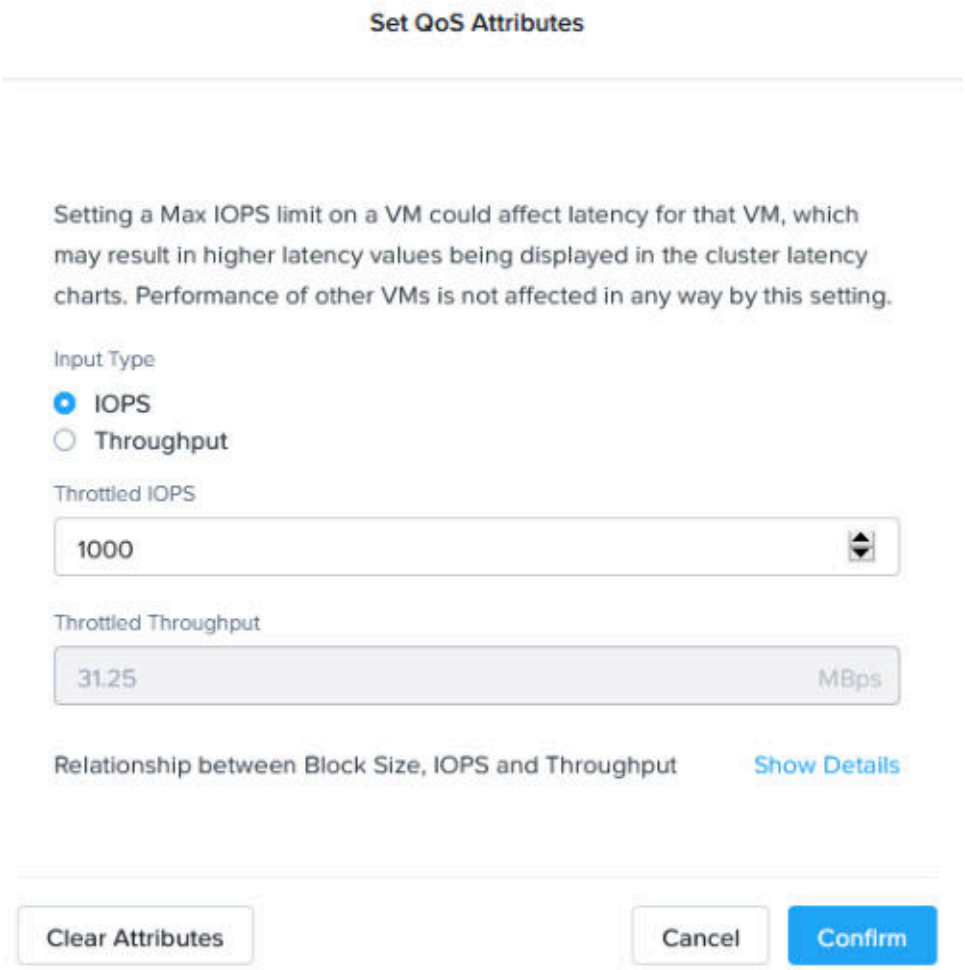

Figure 238: Setting QoS Attributes

6. Click **Show Details** to populate a table that describes the relationship between block size, IOPS, and throughput.
7. Click **Clear Attributes** to delete all the values.
8. To save details, click **Confirm**.
   This step stores the number of IOPS or throughput that would be served by a storage layer. To edit the storage QoS setting, follow the previous process and enter a new value for IOPS or throughput.

Storage Management

Storage in a Nutanix cluster is organized into several components that allow you to manage capacity and performance. Nutanix clusters include storage pool, storage container, volume group, and virtual disk components. Prism Element allows you to manage all aspects of storage for a cluster, but you can also manage some storage components directly from Prism Central.

- The storage containers dashboard displays information about storage containers across the registered clusters (see Storage Containers Summary View on page 97).
- The volume groups dashboard displays information about volume groups across the registered clusters (see Volume Groups Summary View on page 108). Through Prism Central you can do the following:

  Note: Some actions are unavailable for clusters running a pre-6.0 AOS version.

  - Create a volume group (see Creating a Volume Group on page 380)
  - Update or delete a volume group (see Modifying a Volume Group on page 384)
  - Add, update, or remove connections to external clients for a volume group (see Managing Volume Group Connections on page 385)
  - Add, update, or delete virtual disks for a volume group (see Managing Volume Group Virtual Disks on page 389)
- For information about storage management tasks that are performed through Prism Element, see the "Storage Management" chapter in the *Prism Web Console Guide*.

Creating a Volume Group

About this task

To create a volume group through Prism Central, do the following:

Figure 239: Creating a Volume Group

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275775660001 on Brightcove.

Procedure

1. Go to the **List** tab in the volume groups dashboard (see Volume Groups Summary View on page 108) and click the **Create Volume Group** button.
   The **Create Volume Group** page appears.

2. In the Configuration screen, do the following in the indicated fields:

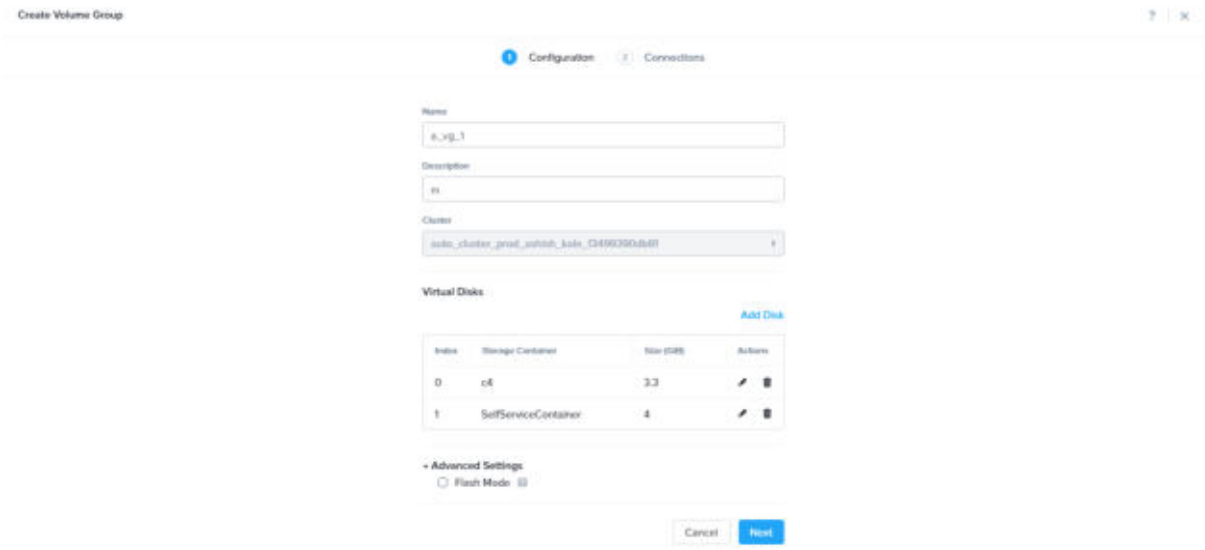

Figure 240: Create Volume Group: Configuration

a. **Name**: Enter a name for the volume group.

b. **Description**: Enter an optional description for the volume group.

c. **Cluster**: Select the cluster in which to create the volume group from the pull-down list.

   The list includes all clusters registered with this Prism Central.

d. **Virtual Disks**: Click the **Add Disk** button. In the **Add Virtual Disk** window, select the target volume group from the pull-down list, enter the disk size (in GiBs), and then click the **Add** button.

   The Configuration screen reappears with a table of virtual disks. The table includes index (assigned number), storage container (name), size, and action fields with a line for each disk in the volume group.

   - To add another disk to the volume group, click the **Add Disk** link above the table. This redisplays the **Add Virtual Disk** window. Repeat this step for the new disk.
   - To change the size of a disk, click the pencil mark icon for that disk. An **Edit Virtual Disk** window appears in which you can change the size (but not the storage container).
   - To delete a disk, click the garbage can icon for that disk.

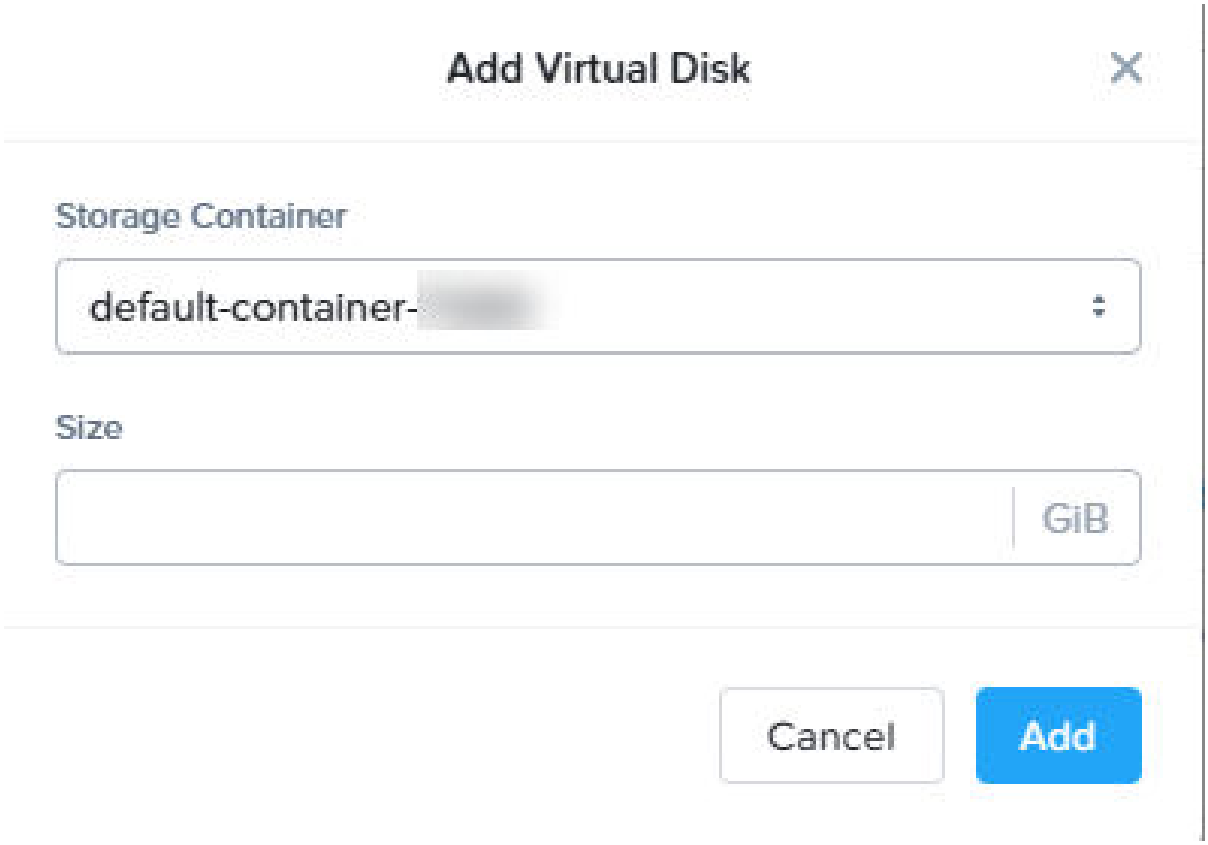

Figure 241: Add Virtual Disk Window

e. **Advanced Settings** (optional): Check the **Flash Mode** box to ensure no down migration of data occurs from the flash tier.

   Clicking **Advanced Settings** toggles (display/hide) the **Flash Mode** box.

f. Click the **Next** button.

3. In the Connections screen, do the following in the indicated fields:

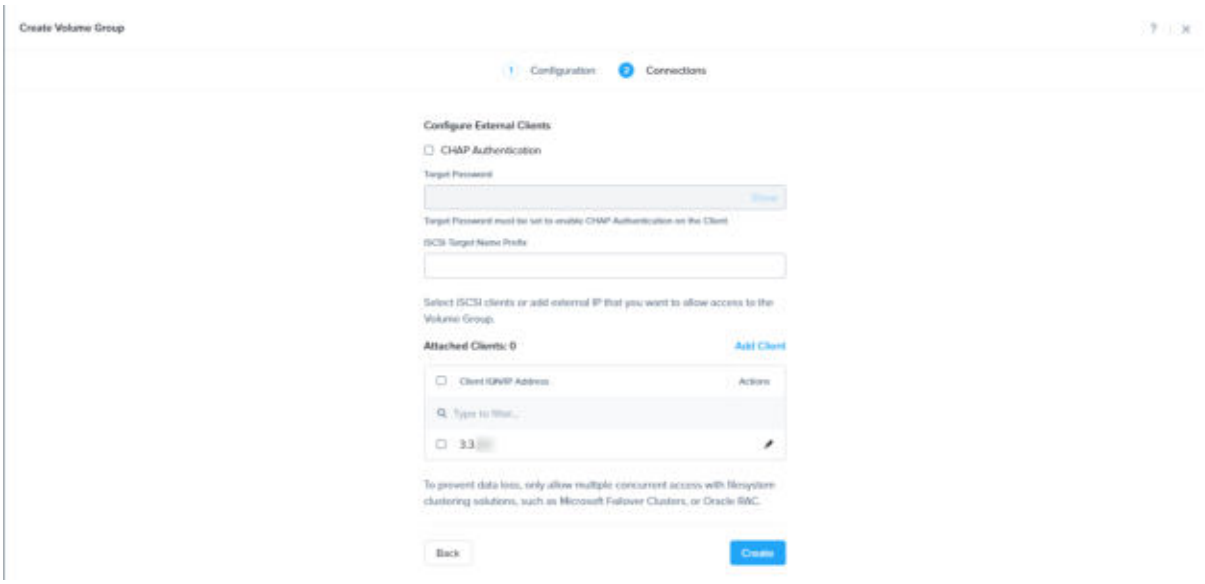

Figure 242: Create Volume Group: Connections

a. **CHAP Authentication**: Check the box to enable the Challenge Handshake Authentication Protocol (CHAP) and then enter the password in the indicated field.

b. **iSCSI Target Name** (read only): The iSCSI target name appears in the field. (You cannot change the name.)

c. **Attached Client**: Configure the clients as desired. The configured clients appear in the table.

   - To attach (unattach) a client, check (uncheck) the box for that client in the table.
   - To add a client, click the **Add Client** link. In the **Add External Client** window, enter the client IP address or IQN designator name in the indicated field, check the **CHAP**

Authentication box and password to enable CHAP authentication, and then click the **Add** button. The added client appears in the table.

- To update a client, click the pencil icon for that client in the table. The **Edit External Client** window appears. Check (uncheck) the box to enable (disable) CHAP authentication and then click the **Save** button. (You cannot change the client address.)

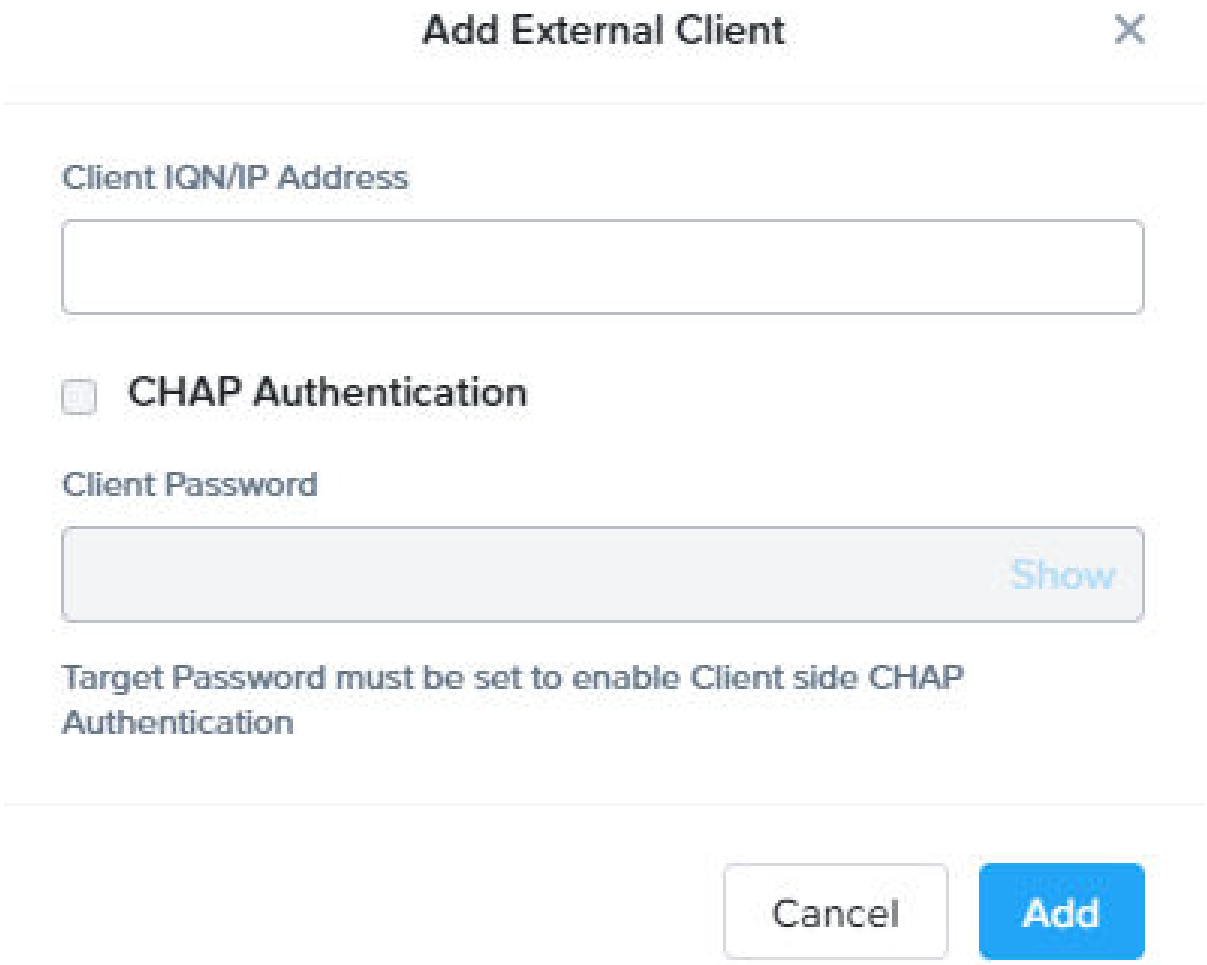

Figure 243: Add External Client Window

d. Click the **Create** button.

The volume group is created, and the new volume group appears in the list of volume groups.

Modifying a Volume Group

About this task

You can modify a volume group (update or delete) at any time after creating it.

Procedure

- To update a volume group, do the following:

  a. Do one of the following:

     » Go to the **Summary** tab in the details page of the target volume group (see Volume Group Details View on page 112) and click the **Update** button.

     » Select the target volume group in the **List** tab of the volume groups dashboard and then click the **Update** button (see Volume Groups Summary View on page 108).

  b. In the **Update Volume Group** page, update the target parameters and then click the **Save** button.

     The fields are the same as when creating the volume group (see Creating a Volume Group on page 380).

- To delete a volume group, either select the volume group in the **List** tab of the dashboard or go to the **Summary** tab in the details page and then click the **Delete** button. (In the dashboard **Delete** is an option in the **Action** drop-down menu.)

  If there are active connections, you must first remove the connections before deleting the volume group.

Managing Volume Group Connections

About this task

Connections to external clients are configured when the volume group is created (see Creating a Volume Group on page 380), but you configure authentication and add or remove connections to a volume group at any time.

Procedure

- To manage connections as a whole, do the following:

  a. Do one of the following:

     » Go to the **Summary** tab in the details page of the target volume group (see Volume Group Details View on page 112) and click the **Manage Connections** button.

     » Select the target volume group in the **List** tab of the volume groups dashboard and then select **Manage Connections** from the **Actions** pull-down menu (see Volume Groups Summary View on page 108).

     The **Manage Connections** page appears.

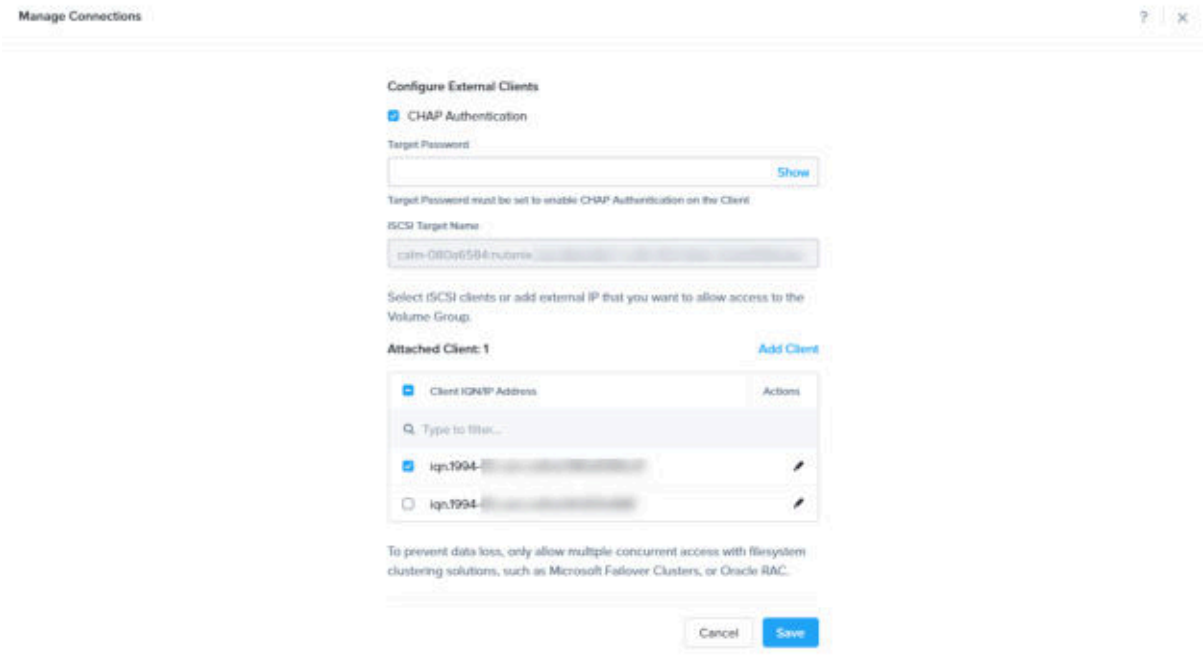

Figure 244: Manage Connections Page

  b. Check the **CHAP Authentication** box to enable the Challenge Handshake Authentication Protocol (CHAP) and then enter the password in the indicated field.

  c. The iSCSI target name appears in the **iSCSI Target Name** field. This is a read-only field because you cannot change the name.

  d. Configure the clients as desired in the **Attached Clients** section. The configured clients appear in the table.

     - To attach (unattach) a client, check (uncheck) the box for that client in the table.
     - To add a client, click the **Add Client** link. In the **Add External Client** box, enter the client IP address or IQN designator name in the indicated field, check the **CHAP Authentication** box and enter the password to enable CHAP authentication, and then click the **Add** button. The added client appears in the list.
     - To change the CHAP authentication for a client, click the pencil icon for that client in the table. The **Edit External Client** box appears. Check (uncheck) the box to enable

(disable) CHAP authentication and then click the **Save** button. (You cannot change the client address.)

e. Click the **Save** button.

- To add an external client connection to a volume group, do the following:

  a. Go to the **Connections** tab in the details page of the target volume group (see Volume Group Details View on page 112) and then click the **Add Connection** button.

  b. In the **Add External Client** window, do the following in the indicated fields:

  - **Add External Client**: Enter the client IP address or IQN designator name
  - **CHAP Authentication**: Check the box to enable CHAP authentication and enter the client password in the indicated field.
  - Click the **Add** button.

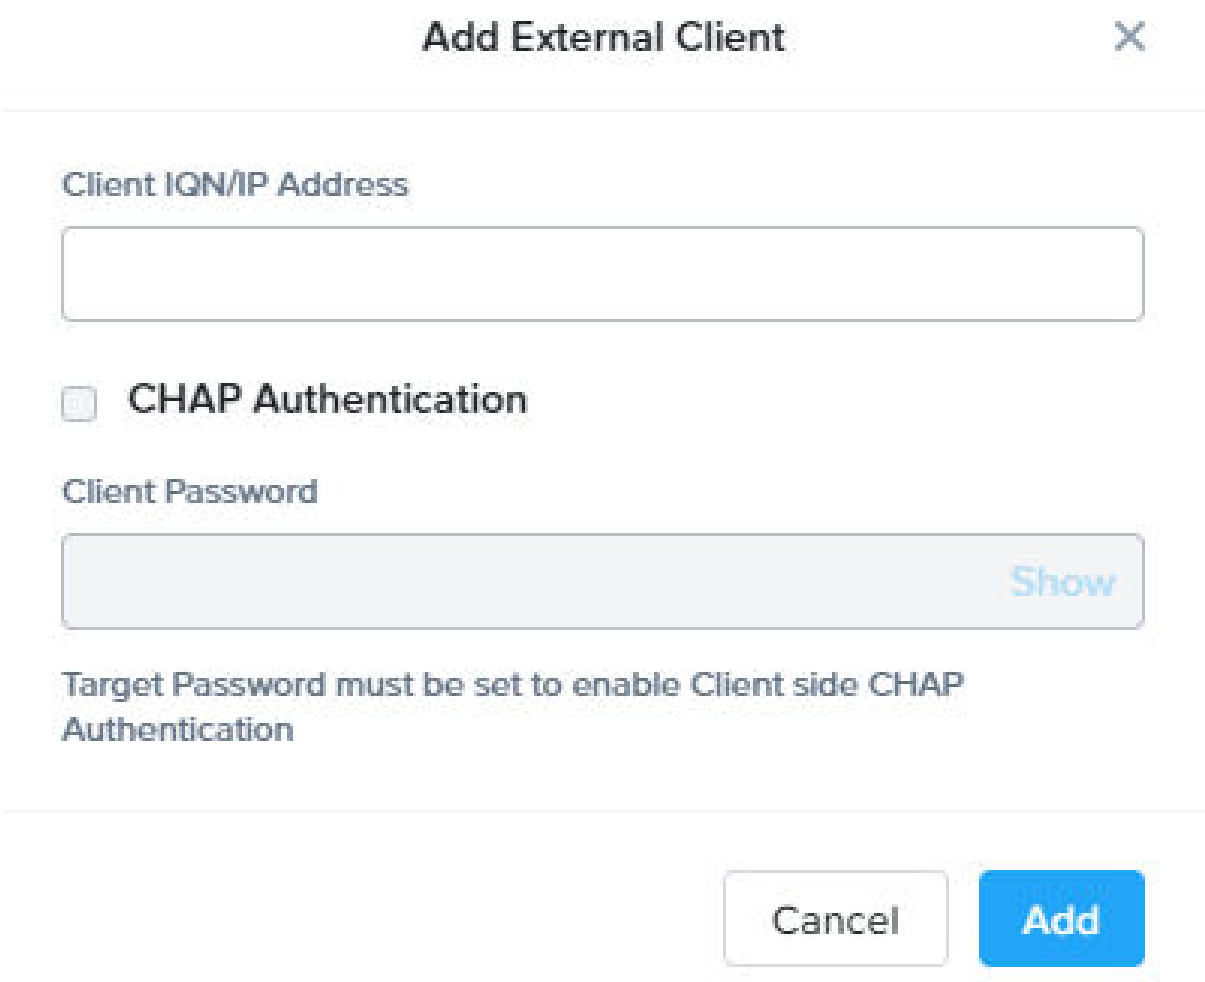

Figure 245: Add External Client Window

The connection is added, and the new client appears in the list of connections for the volume group.

- To update a client connection, do the following:

  Note: A separate CHAP authentication setting for an individual client is not allowed when CHAP authentication is not enabled for the volume group, and you cannot change the client address.

  a. Go to the **Connections** tab in the details page of the target volume group, select the target connection, and then click the **Update** button.

  b. In the **Edit External Client** window, uncheck the **CHAP Authentication** box to disable or check to enable. If you check the box, enter the client password in the indicated field.

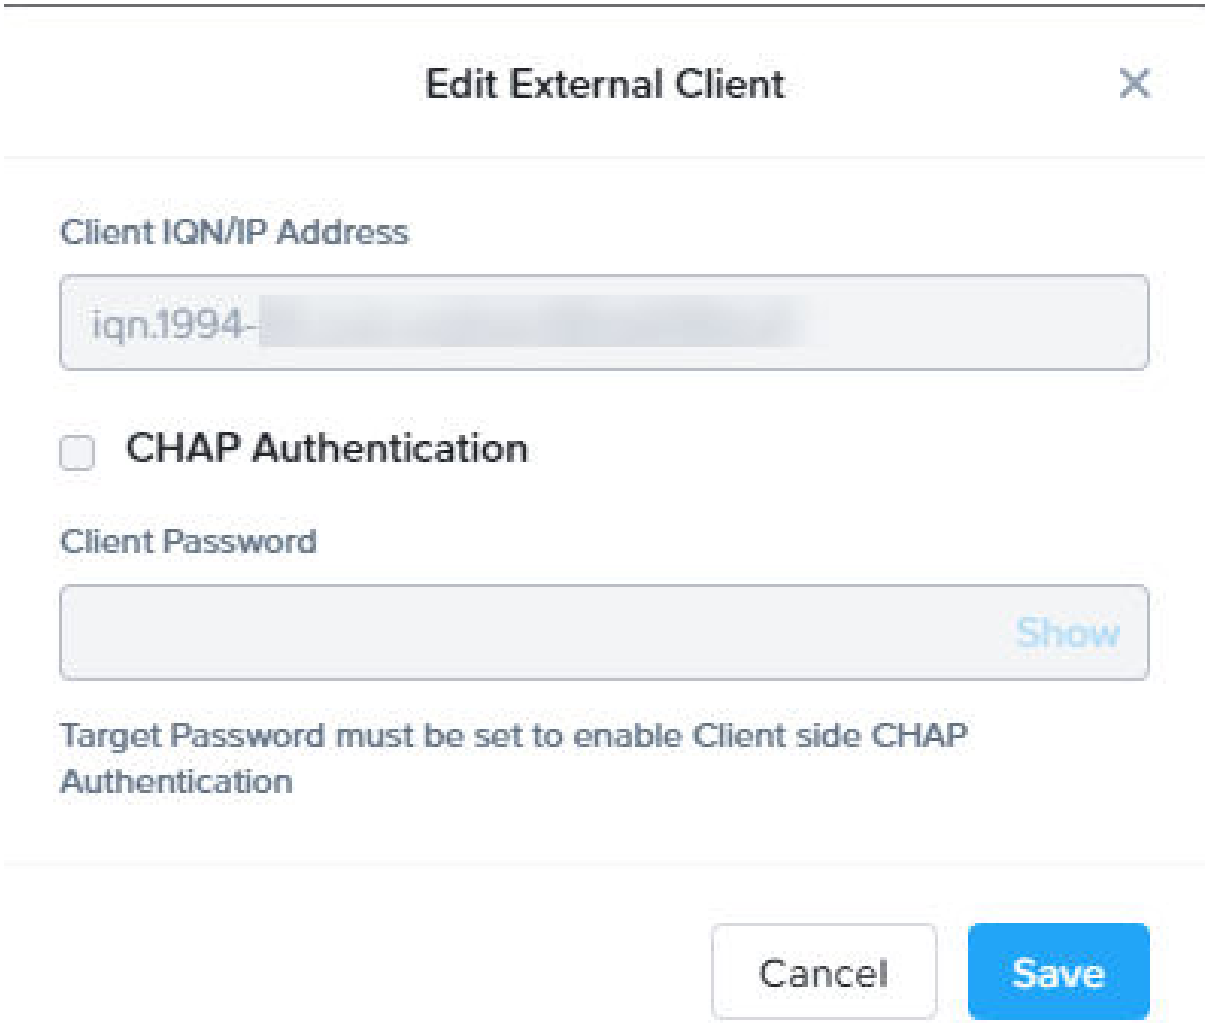

Figure 246: Edit External Client Window

- To remove a client connection, select the **Connections** tab in the details page of the target volume group, select the target connection, and then click the **Remove** button.
  The connection is removed, and the entry is deleted from the list of connections.

Managing Volume Group Virtual Disks

About this task

Virtual disks are added when the volume group is created (see Creating a Volume Group on page 380), but you add, update, or delete disks for a volume group at any time.

Procedure

- To add a virtual disk to a volume group, do the following:
  a. Go to the **Virtual Disks** tab in the details page of the target volume group (see Volume Group Details View on page 112) and then click the **Add Virtual Disk** button.
  b. In the **Add Virtual Disk** window, do the following in the indicated fields:
     - **Storage Container**: Select the target storage container from the pull-down list.
     - **Size**: Enter the desired size (in GiBs) for the virtual disk.
     - Click the **Add** button.

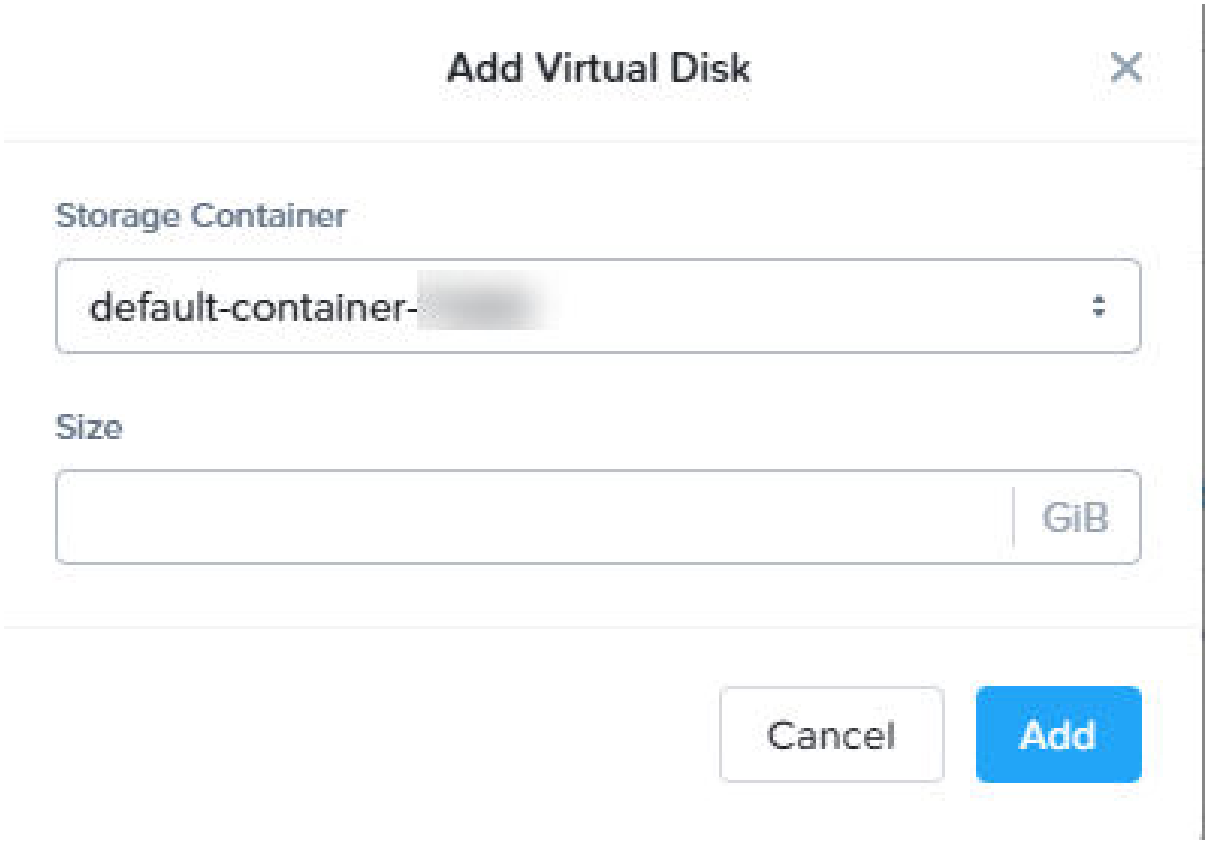

Figure 247: Add Virtual Disk Window

The disk is created, and the new disk appears in the list of virtual disks for the volume group.

- To update a virtual disk, do the following:

  Note: You can only change the disk size, not the storage container.

  a. Go to the **Virtual Disks** tab in the details page of the target volume group, select the target disk, and then click the **Update** button.

  b. In the **Edit Virtual Disk** window, update the disk size (in GiBs) as desired and then click the **Save** button.

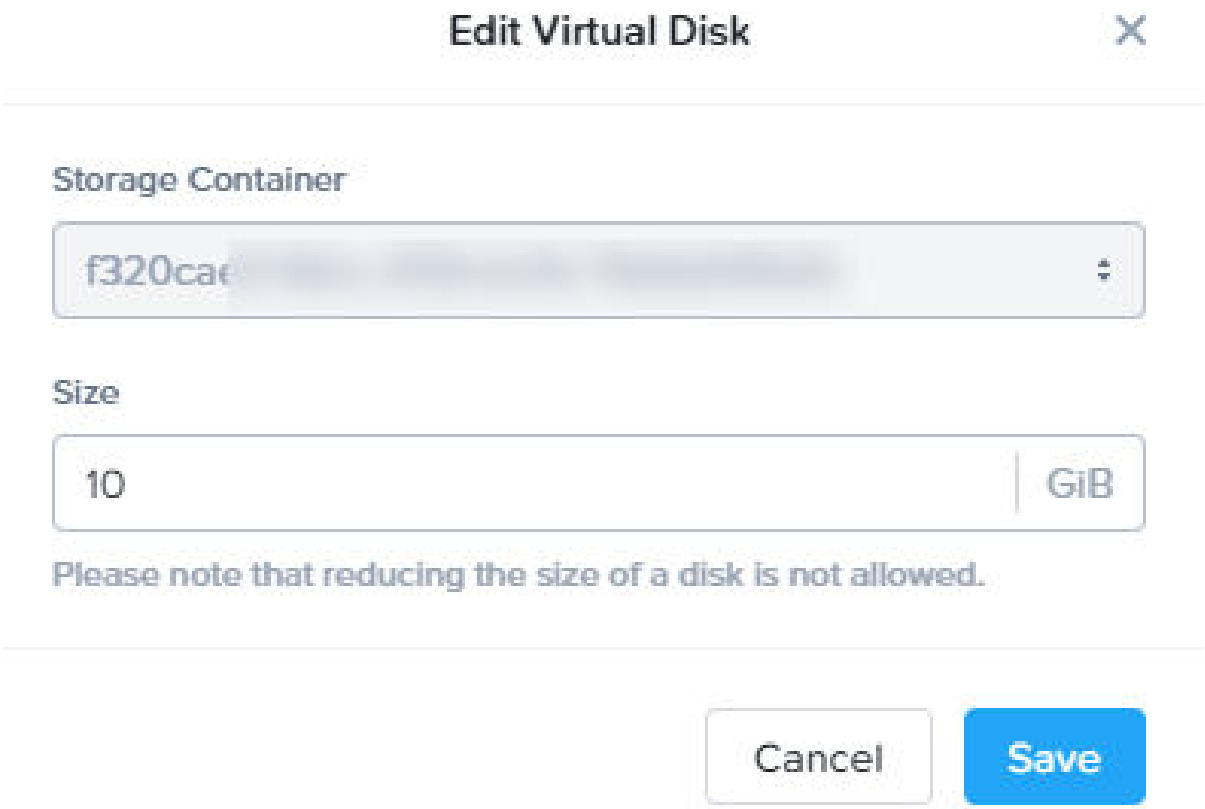

Figure 248: Edit Virtual Disk Window

- To delete a virtual disk, select the **Virtual Disks** tab in the details page of the target volume group, select the target disk, and then click the **Delete** button.

Catalog Management

Prism Central includes a catalog service for storing VM snapshots and images. A Prism Central or self-service administrator creates this catalog of objects so that self-service users who have permissions to create a VM can use them.

Note:

- The catalog service is a self-service feature that appears in Prism Central only when Prism Self Service is enabled (see Prism Self Service Overview on page 457).
- You can manage the Catalog Item entity permissions by enabling or disabling the Marketplace Item permissions while creating a custom role. For example, if you

select **No Access** permission for the Marketplace Item entity while creating the custom role, the custom role will not have access to the Catalog Item entity as well. See the *Creating A Custom Role* section of the *AOS Security Guide* on how to create a custom role with access to catalog items entity.

Adding a Catalog Item

About this task

When you add a VM to the catalog, a snapshot of the VM is created. The snapshot is available to users across all the self-service projects. Users who have the requisite permissions can create VMs from the snapshot. You can continue to use the VM or delete the VM after you add it to the catalog. These actions do not affect the snapshot.

Adding an image to the catalog makes the image available to self-service users who have permissions to create VMs. A copy of the image is added to the catalog, so you can delete the image without affecting the copy in the catalog.

To add a VM snapshot or image to the catalog:

Note: Nutanix recommends that you power off the VM before adding a snapshot of it to the catalog.

Procedure

1. Do one of the following:
   - » Go to the **List** tab of the VMs dashboard and select the VM that you want to add to the catalog (see VMs Summary View on page 82).
   - » Go to the images dashboard and select the image that you want to add to the catalog (see Images Summary View on page 125).

2. To add an image, select **Add Image to Catalog** from the **Actions** menu and do the following in the Add Image to Catalog window in the indicated fields:

   a. **Name**: Enter a name for the image.

   b. **Description** (optional): Enter a description of the image.

   c. Click the **Save** button.

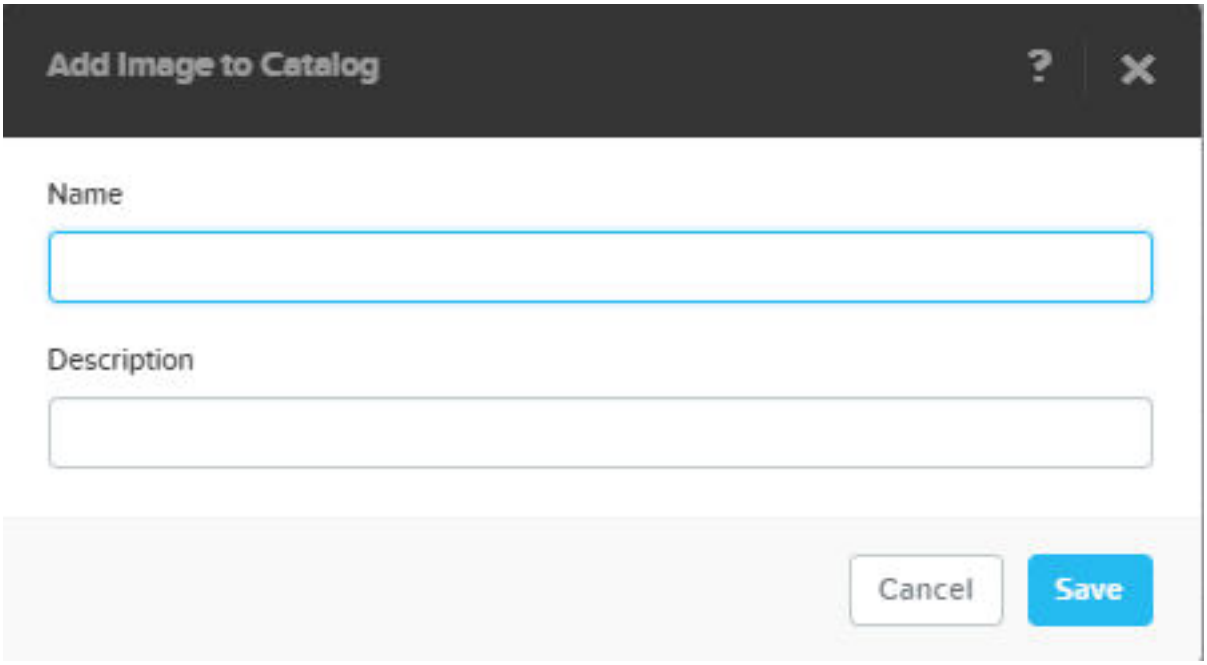

Figure 249: Add Image to Catalog Window

3. To add a VM, select **Add to Catalog** from the **Actions** menu and do the following in the Add VM to Catalog page in the indicated fields:

   a. **Name**: Enter a name for the VM.

   b. **Description** (optional): Enter a description of the VM.

   c. **Guest Customization**: Select **No Customization** in the left pull-down menu if you do not want to allow any customizations. If you do want to allow customizations, select either

SysPrep (Windows) or **Cloudinit (Linux)** in the left pull-down menu and then either **Guided Setup** or **Existing Script** in the right pull-down menu.

If you select **Existing Script**, a **Custom Script** field appears. Click the **Upload File** button and select the target script file.

If you select **Guided Setup**, a set of fields appears. Do the following in the indicated fields:

- **Authentication**: Check the box to allow the user deploying the VM to set a username and password.
- **SSH Key** [Cloudinit (Linix) only]: Check the box to allow the user deploying the VM to provide an SSH key.
- **Locale**: Check the box to allow the end user to specify the locale (language).
- **Hostname**: Check the appropriate radio button to specify the host name source: the deployed VM name, a name provided by the person deploying the VM, or restricted hostname access.
- **License Key**: Specify the license key source. Check the **Enter License Key** radio button and enter the key in the field to set the license key, check the **Allow end user to input License Key** radio button to let the user do it, or check the **No License Key** radio button to not require a license key.

d. Click the **Save** button.

Figure 250: Add VM to Catalog Page

Deleting a Catalog Item

About this task

VM snapshots and images deleted from a catalog makes those items unavailable to project members. To delete a VM snapshot or image from the catalog, do the following:

Procedure

1. Do one of the following:

   » Go to the **List** tab of the VMs dashboard and select the VM that you want to delete from the catalog (see VMs Summary View on page 82).

   » Go to the images dashboard and select the image that you want to delete from the catalog (see Images Summary View on page 125).

2. From the **Actions** menu, select **Delete Catalog Item**.

Image Management

Prism Central provides a centralized location to manage the images you require on registered AHV clusters. Prism Central enables you to upload images to the clusters and maintains an inventory of the images on them.

> Note: You can configure policies that govern which clusters receive the images that you upload (see Image Placement Policies on page 480).

Figure 251: Image Management Demo

A video on this topic is available https://youtu.be/UlqM8jjvr4s on YouTube.

Uploading Images from Prism Central to Registered Clusters

From Prism Central, you can upload images by using the following methods:

- Uploading images from a workstation
- Specifying URLs to images on a remote server
- Uploading images from a VM disk

For information about these methods, see Creating an Image on page 395.

Selecting Target Clusters When Uploading Images from Prism Central

When uploading images by either of the methods described earlier, you can select target clusters. Your image placement options are as follows:

- If you have specific target clusters in mind, you can select those clusters when uploading the images.
- If you have image placement policies in place, you can associate one or more categories with the images. Prism Central uploads the images to the clusters identified by the image placement policies. For information about image placement policies, see Image Placement Policies on page 480.

Upload Behavior When Target Clusters Are Not Selected

If you do not specify target clusters, the choice of target clusters depends on whether you choose to upload the images from a workstation or from a remote server. The respective procedures describe the choice of target clusters. See Uploading Images from a Workstation on page 396 and Uploading Images from a Remote Server on page 400, and then choose the option that works for you. Regardless of the method you choose, Prism Central retains the ability to manage the images that you upload from Prism Central.

Importing Images from Registered Clusters to Prism Central

In addition to uploading images, you can also import images from registered AHV clusters. You can import images only from registered clusters that are running AOS 5.8 or later. For more information, see Importing Images to Prism Central on page 412.

Requirements

The image management feature in Prism Central has the following requirements:

- The version of both Prism Central and the version of AOS on registered clusters must be 5.8 or later.
- Clusters must be registered with Prism Central and must have AHV installed.
- The image service uses port 2007, so that port must be open.

Limitations

Image management from Prism Central has the following limitations:

- Prism Central cannot update images that another Prism Central instance managed earlier. However, the images are displayed in Prism Central, and you can use the images to create VMs or delete the images you no longer need.
- You cannot choose a container when uploading images from Prism Central. Prism Central uploads images to the container named `SelfServiceContainer`.
- When uploading images from a workstation, If multiple clusters are selected during the initial image upload then the uploaded image becomes active (in Prism Element) on only one of the selected clusters and is inactive on the other clusters registered to Prism Central.
- If an image is not active on a cluster, as an admin user it is possible that you are unable to use Prism Central to create that VM on that cluster. The cluster on which you want to create the VM cannot check out the image from a remote cluster. In this scenario, you must upload the image manually to the cluster on which you want to create the VM.

  When you create a VM, Prism Central uses API v2 workflows. API v2 does not have the check-out workflow. The image checkout workflow is used to copy an image on demand from cluster B to cluster A if the image does not exist or is inactive on cluster A.

  > Note:
  >
  > This limitation is removed by applying an image placement policy and a category to the image that is available on multiple remote clusters. When you apply an image placement policy to an image, Prism Central propagates the image to all the clusters that are included in the image placement policy and category. Prism Central can then check-out the image from the remote cluster where the image is active and create the VM.
  >
  > Prism Central can take up to 15 minutes to switch the image propagation from an alternative cluster. For example, if an image is located on clusters A and B. As a result of an image placement policy you applied to that image, cluster A starts copying the image to cluster C. During the copy operation, cluster A becomes unavailable. In such a case, Prism Central takes up to 15 minutes to switch to cluster B and start copying the image from cluster B to cluster C.

Creating an Image

You can add an image in one of two ways.

About this task

To add an image, do one of the following:

Procedure

- Upload Images from a workstation. See Uploading Images from a Workstation on page 396.
- Uploading Images from a Remote Server. See Uploading Images from a Remote Server on page 400.
- Uploading Images from a VM disk. See Uploading Images from a VM Disk on page 405.

Uploading Images from a Workstation

You can browse the workstation from which you are accessing Prism Central and upload images from it.

About this task

You can select multiple images and upload them as part of a single operation. With this method of upload, Prism Central identifies a single registered cluster for each specified image and uploads the image to it. The choice of cluster is random. The image becomes active (in Prism Element) on the selected cluster and inactive on other clusters registered to Prism Central. Image placement policies take effect after a short duration.

> Note: Most modern browsers impose file size limitations that affect this upload method. If you must upload images larger than 2 GB, upload the images from a remote server instead. Also, the browser type and CPU and RAM utilization on the workstation limit the number of concurrent uploads. Concurrent uploads exceeding the default limit of the browser are queued or throttled by the browser and can take more time. Large file uploads and high CPU and memory utilization can slow down the browser.

To upload an image from a workstation, do the following:

Procedure

1. In the **Entities** menu, go to **Virtual Infrastructure** > **Images**, and then click the **Add Image** button (see Images Summary View on page 125).
   The Add Images page appears.

2. Under **Image Source**, click **Image File**, and then click **Add File**.

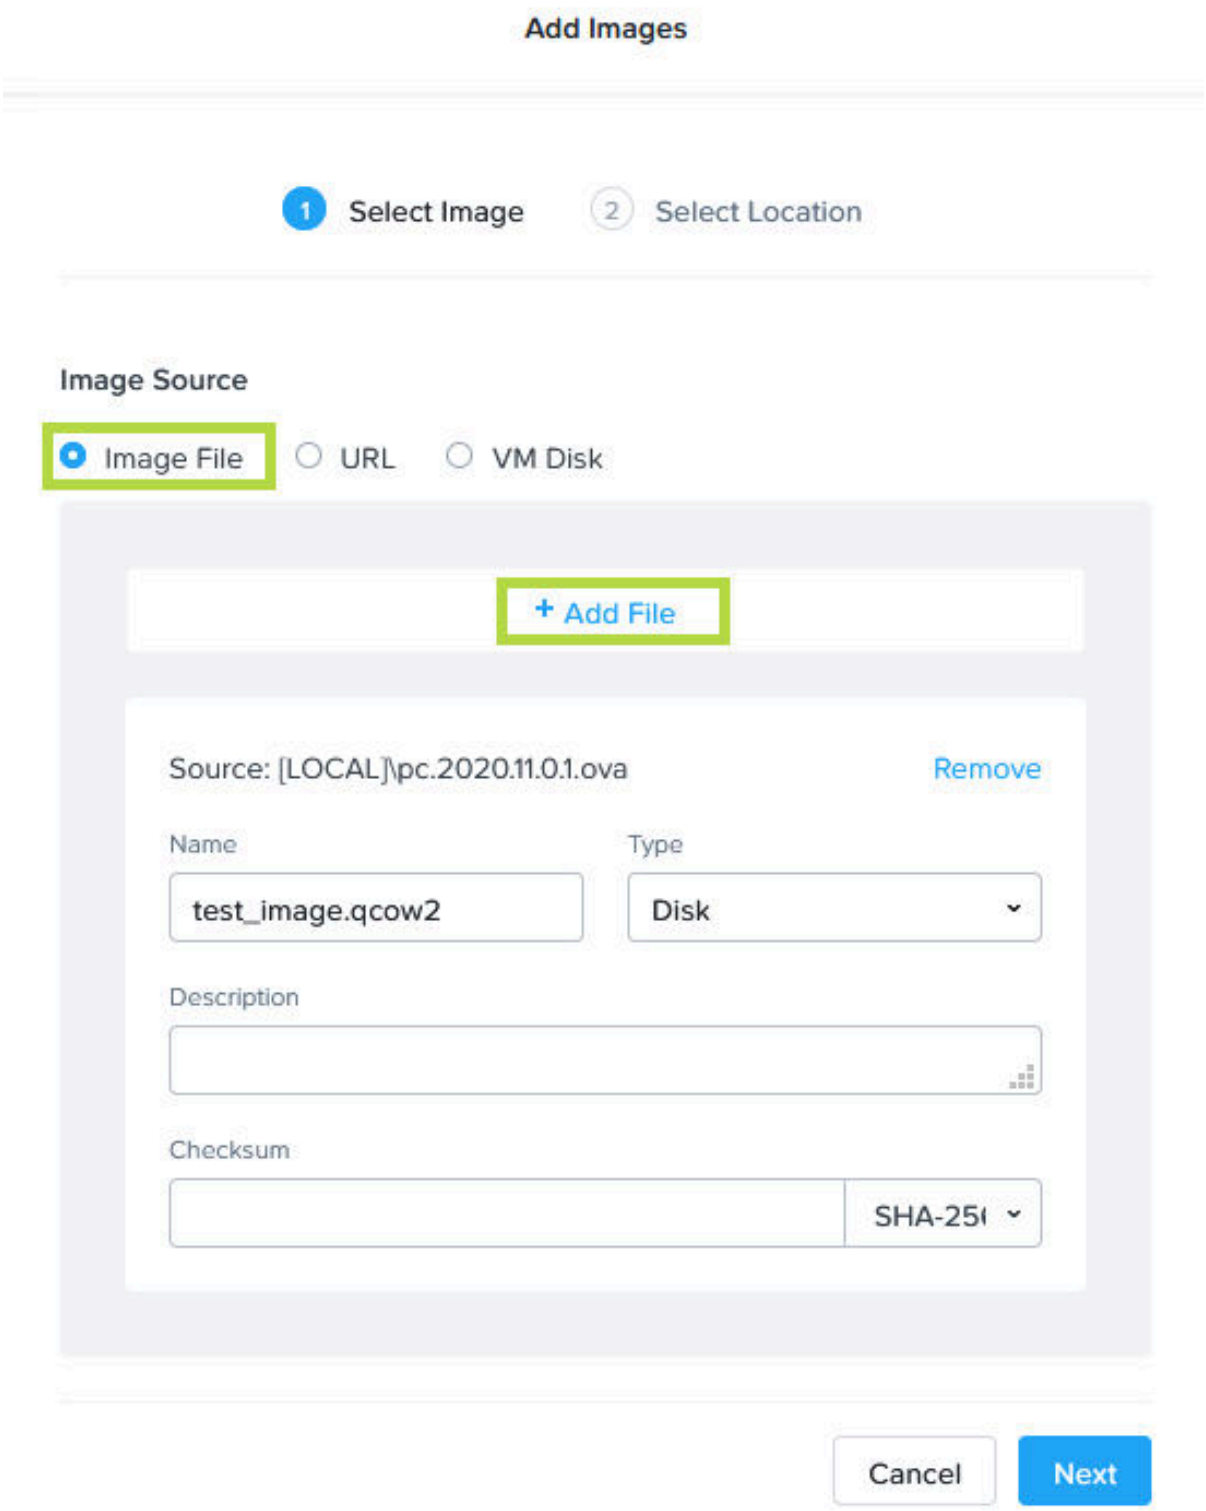

Figure 252: Add Images: File

3. To specify one or more files for upload, do the following:

   a. Browse to the location of the image file, and then click **Open**.
      Prism Central adds the image file to a list on the page.

   b. In **Image Name**, make sure that the name of the image is unique across all the images in Prism Central.
      The name defaults to the name of the file you selected, but you can change that name here.

   c. In **Image Type**, select the type of image.

   d. In **Image Description**, describe the image file.

   e. In **Checksum**, select the hashing algorithm.

   Repeat this step to add as many image files as you want.

4. To remove an image file entry, locate the entry and click **Remove**.

5. After you have added all the image files you want, click **Next**.
   The **Select Location** page appears.

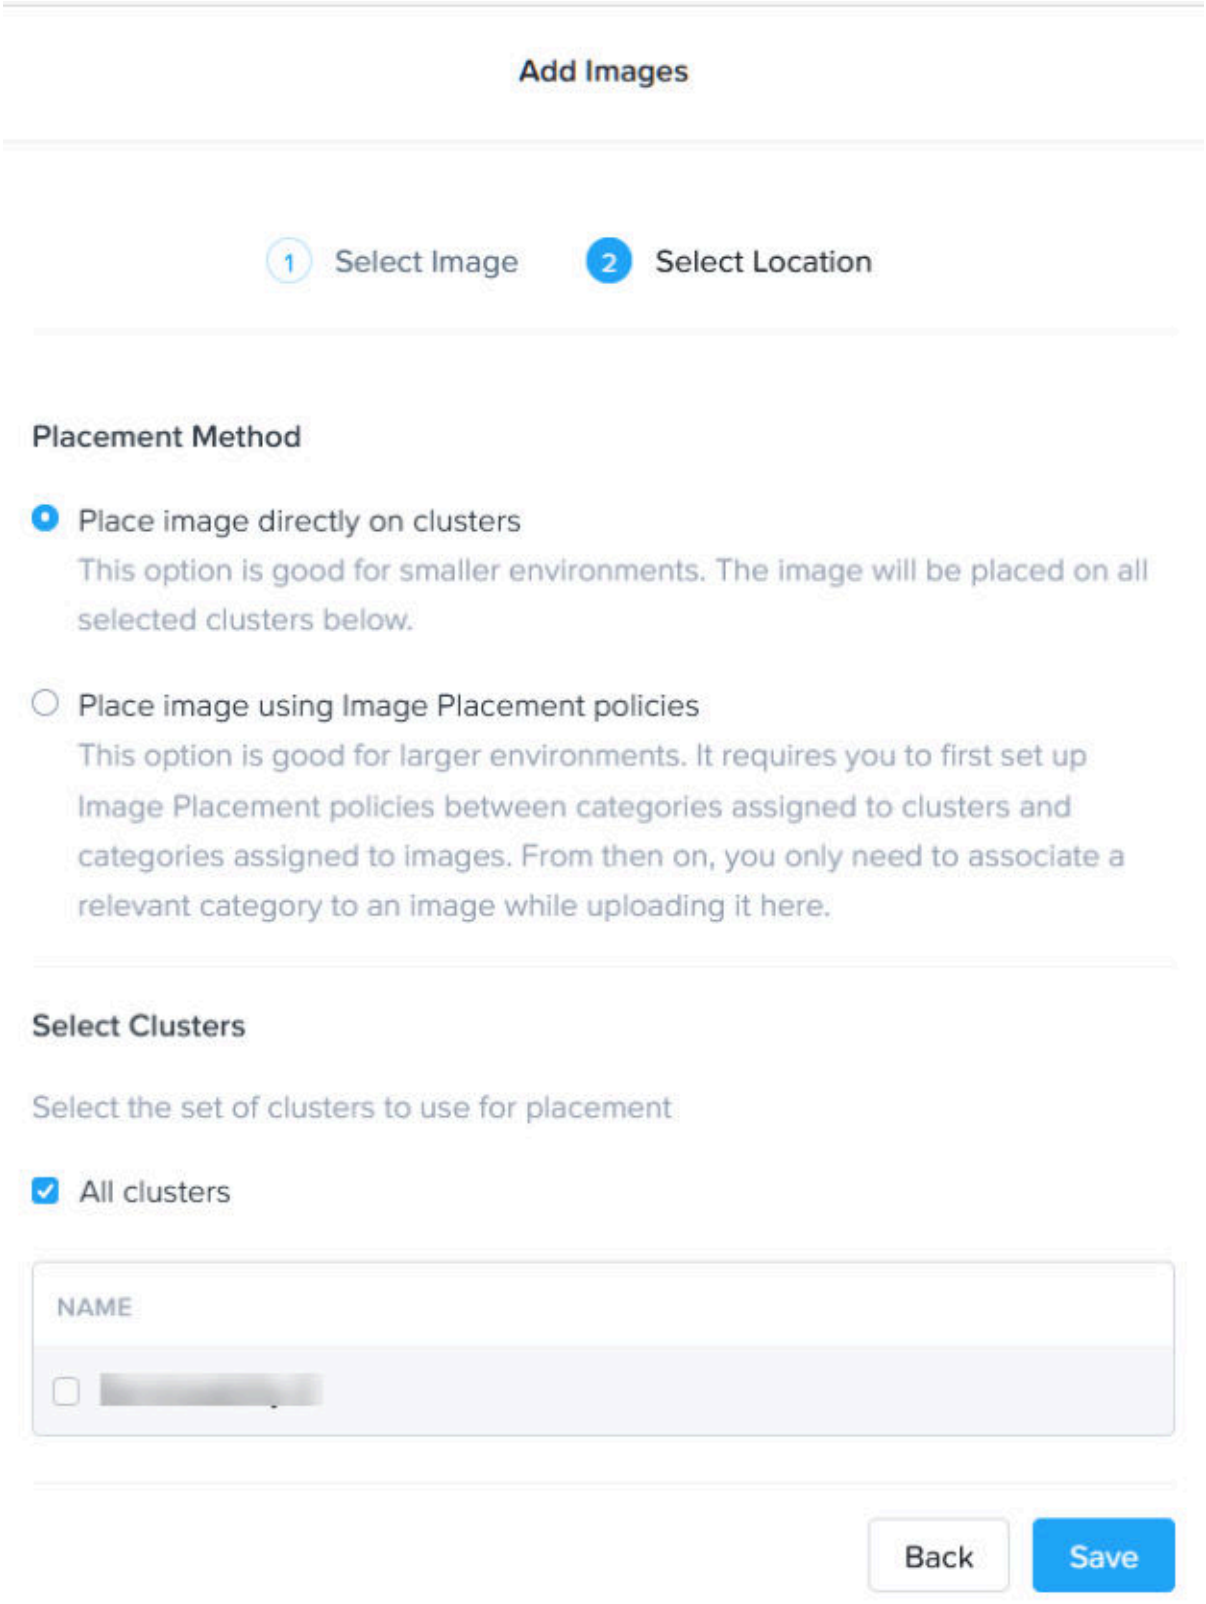

Figure 253: Add Images: Select Location

6. In **Placement Method**, do one of the following:

   - To manually select the target clusters, click **Place image directly on clusters**, and then do one of the following in the **Cluster Details** section:

     - If you want to upload the images to all registered clusters, make sure to select **All clusters**, and then click **Save**.
     - If you want to upload to only a subset of the registered clusters, clear **All clusters**, select the clusters you want from the list, and then click **Save**.

   - To delegate image placement decisions to configured policies and assign categories to the images, click **Place image using Image Placement policies**, and then do the following in the **Categories** section:

     1. Click inside the **Categories** search box and select the category you want from the list. You can also start typing the name of the category to reduce the list to matching names.
     2. To specify another category, click the add icon beside the search box. Repeat this step to add as many categories as you need.

   Prism Central uploads the image files in batches and takes some time to enforce the image placement policies.

Uploading Images from a Remote Server

You can upload an image from a remote server to registered Nutanix clusters. To upload an image from a remote server, you need the URL of the image. You can also specify URLs to multiple images as part of a single operation. When you specify image URLs, Prism Central uploads the images to all registered clusters. Uploading from a remote server also allows you to overcome file size limitations imposed by modern browsers. The file size limitation is usually 2 GB.

About this task

To upload an image from a remote server, do the following:

Procedure

1. In the **Entities** menu, go to **Virtual Infrastructure** > **Images**, and then click the **Add Image** button (see Images Summary View on page 125).
   The Add Images page appears.

2. Under **Image Source**, click **URL**.

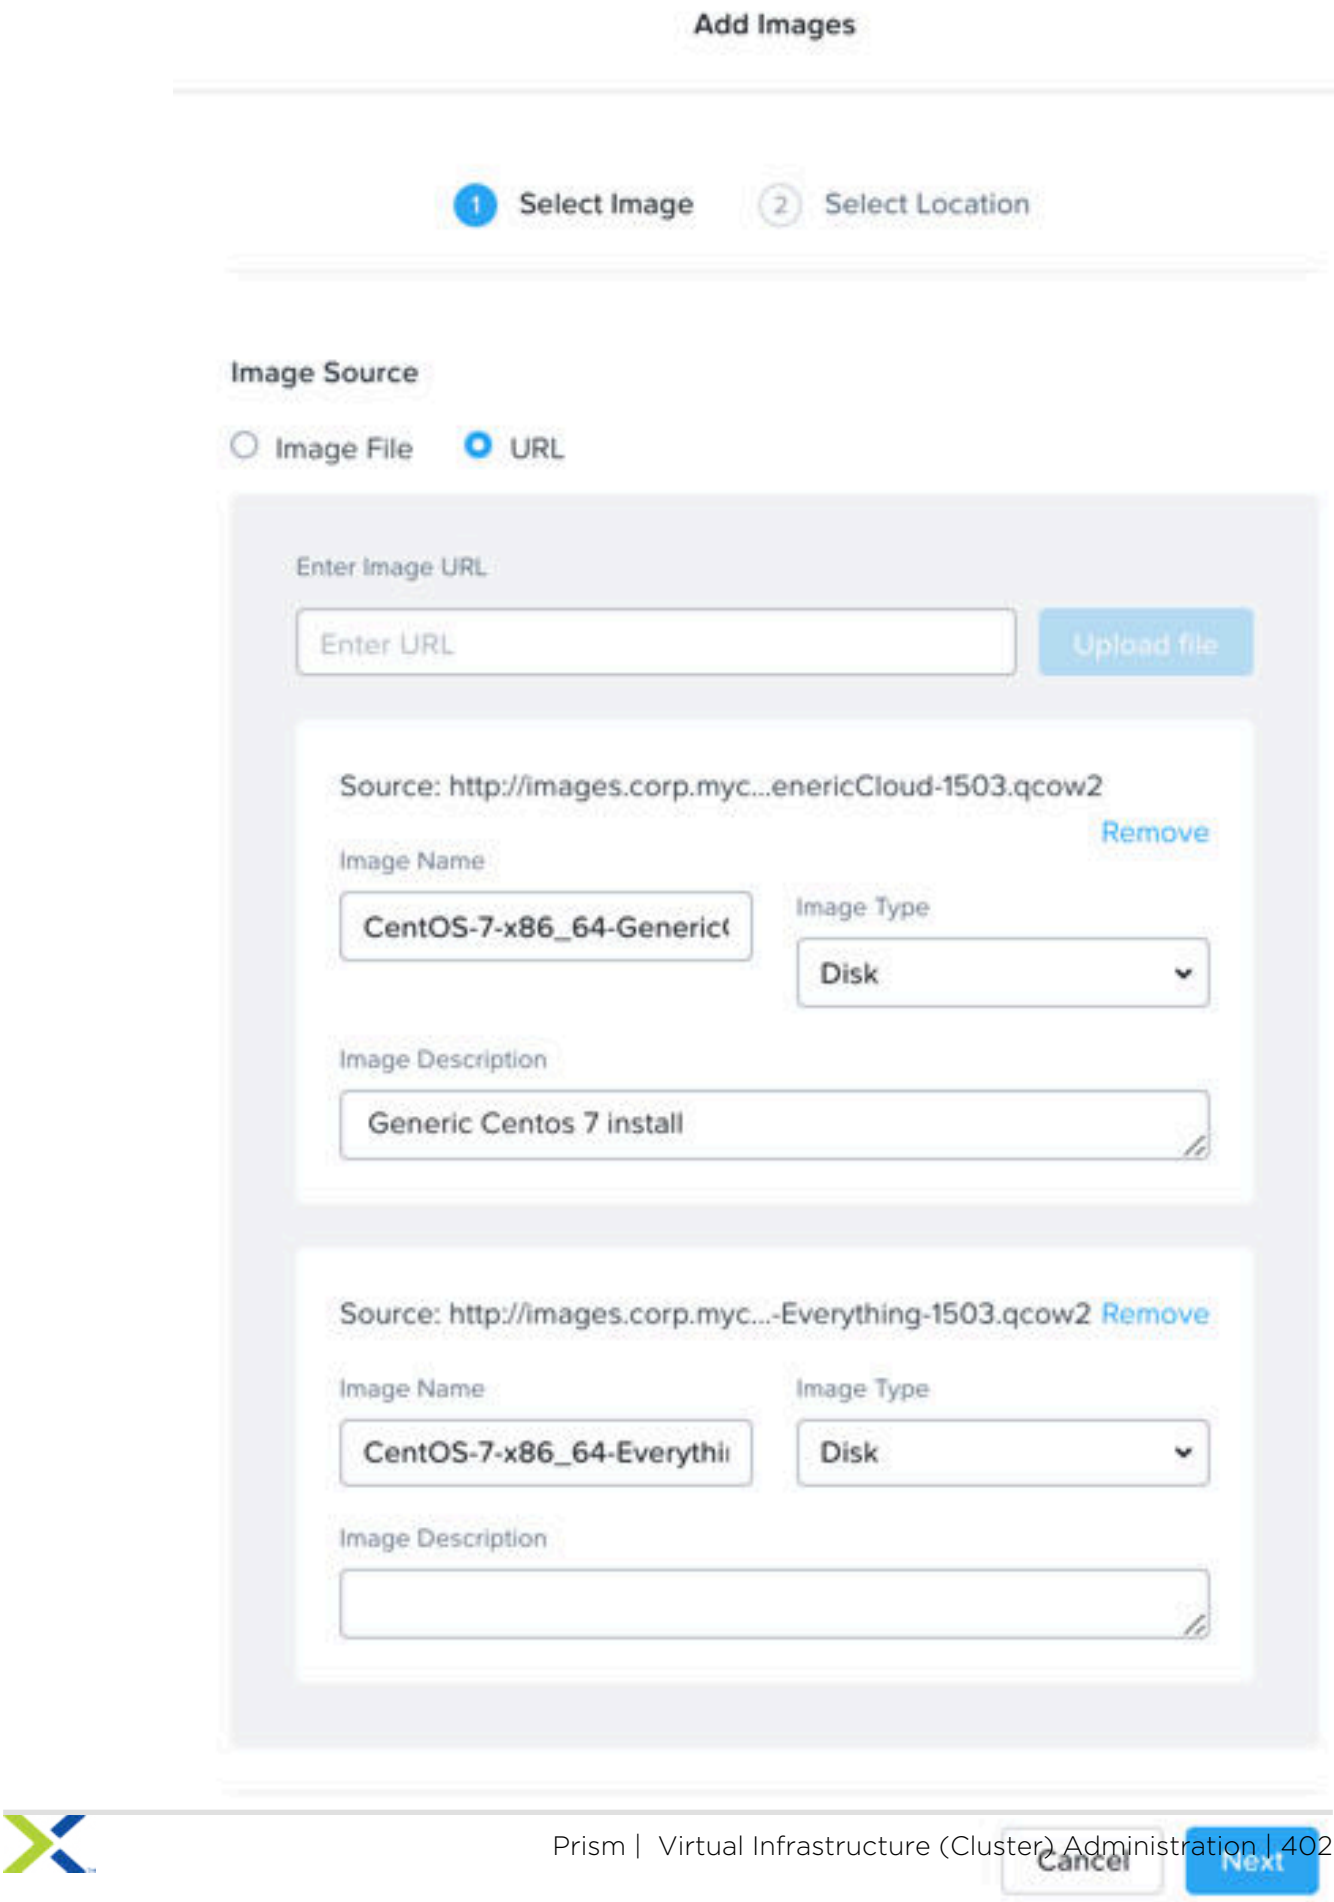
Add Images
1 Select Image
2 Select Location
Image Source
Image File
URL
Enter Image URL
Enter URL
Upload file
Source: http://images.corp.myc...enericCloud-1503.qcow2
Remove
Image Name
CentOS-7-x86_64-Generic(
Image Type
Disk
Image Description
Generic Centos 7 install
Source: http://images.corp.myc...-Everything-1503.qcow2 Remove
Image Name
CentOS-7-x86_64-Everythi
Image Type
Disk
Image Description
Cancel
Next

3. To specify one or more image URLs, do the following:

   a. In **Enter Image URL**, enter the appropriate URL address in the field using the following syntax for either NFS or HTTP, and then click **Upload File**. (NFS and HTTP are the only supported protocols.)

      ```
      nfs://[hostname|IP_addr]/path
      http://[hostname|IP_addr]/path
      ```

      Enter either the name of the host (*hostname*) or the host IP address (*IP_addr)* and the path to the file. If you use a *hostname*, the cluster must be configured to point at a DNS server that can resolve that name (see Configuring Name Servers on page 296). A file uploaded through NFS must have 644 permissions.

      Prism Central adds the URL to a list on the page and clears the **Enter Image URL** text box for another URL.

   b. In **Image Name** in the list entry, make sure that the name of the image is unique across all the images in Prism Central.

      The name defaults to the name of the file you selected, but you can change that name here.

   c. In **Image Type**, select the type of image.

   d. In **Image Description**, describe the image file.

   Repeat this step to specify as many URLs as you want.

4. To remove an image file entry, locate the entry and click **Remove**.

5. After you have added all the URLs you want, click **Next**.
   The **Select Location** page appears.

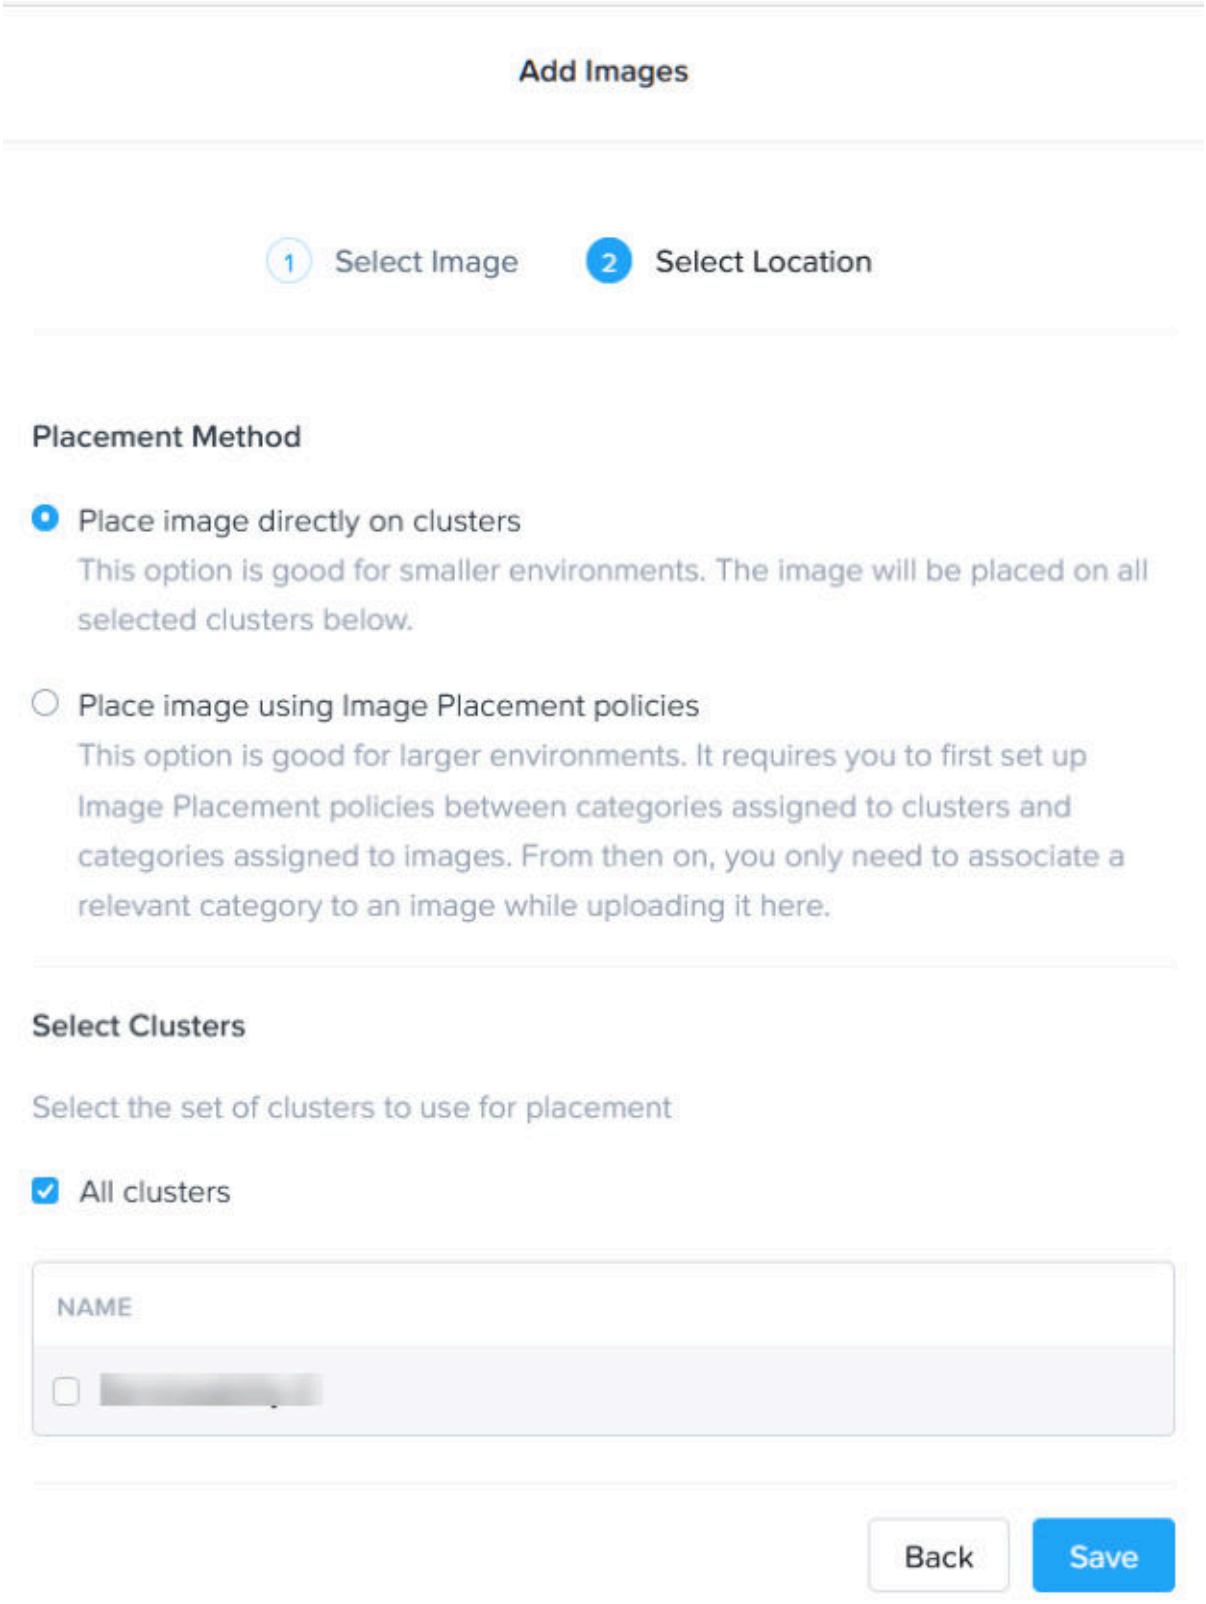

Figure 255: Add Images: Select Location

6. In **Placement Method**, do one of the following:

   - To manually select the target clusters, click **Place image directly on clusters**, and then do one of the following in the **Cluster Details** section:

     - If you want to upload the images to all registered clusters, make sure to select **All clusters**, and then click **Save**.
     - If you want to upload to only a subset of the registered clusters, clear **All clusters**, select the clusters you want from the list, and then click **Save**.

   - To delegate image placement decisions to configured policies and assign categories to the images, click **Place image using Image Placement policies**, and then do the following in the **Categories** section:

     1. Click inside the **Categories** search box and select the category you want from the list. You can also start typing the name of the category to reduce the list to matching names.
     2. To specify another category, click the add icon beside the search box. Repeat this step to add as many categories as you need.

   Prism Central uploads the image files in batches.

Uploading Images from a VM Disk

You can clone an existing VM disk to add it to the image list.

About this task

Figure 256: Uploading Images from a VM Disk

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6271311948001 on Brightcove.

- You select a powered-off source VM on an AHV cluster.
- You then select a disk attached to this source VM to add it as an image.

You can also upload the added image to the registered clusters based on the image placement method. You can either:

- Place the image on the cluster corresponding to where the source VM resides.
- Associate relevant categories to the image so that the corresponding clusters are internally selected based on categories and associated policies.

Before you begin

Ensure that the following requirements are met before you add and upload an image from a VM disk.

- The source VM belongs to an AHV cluster.
- The source VM on which the disk resides is in the powered-off state.

- To place the image on clusters associated with relevant categories, you must first set up the image placement policies between categories assigned to clusters and categories assigned to images.

  See the Creating a Category on page 414 and Assigning a Category on page 415 sections for information on how to create a category and assign a category value to an image respectively.

Procedure

1. In the **Entities** menu, go to **Virtual Infrastructure** > **Images**, and then click the **Add Image** button (see Images Summary View on page 125).
   The Add Images page appears.

2. Under **Image Source**, click **VM Disk**.

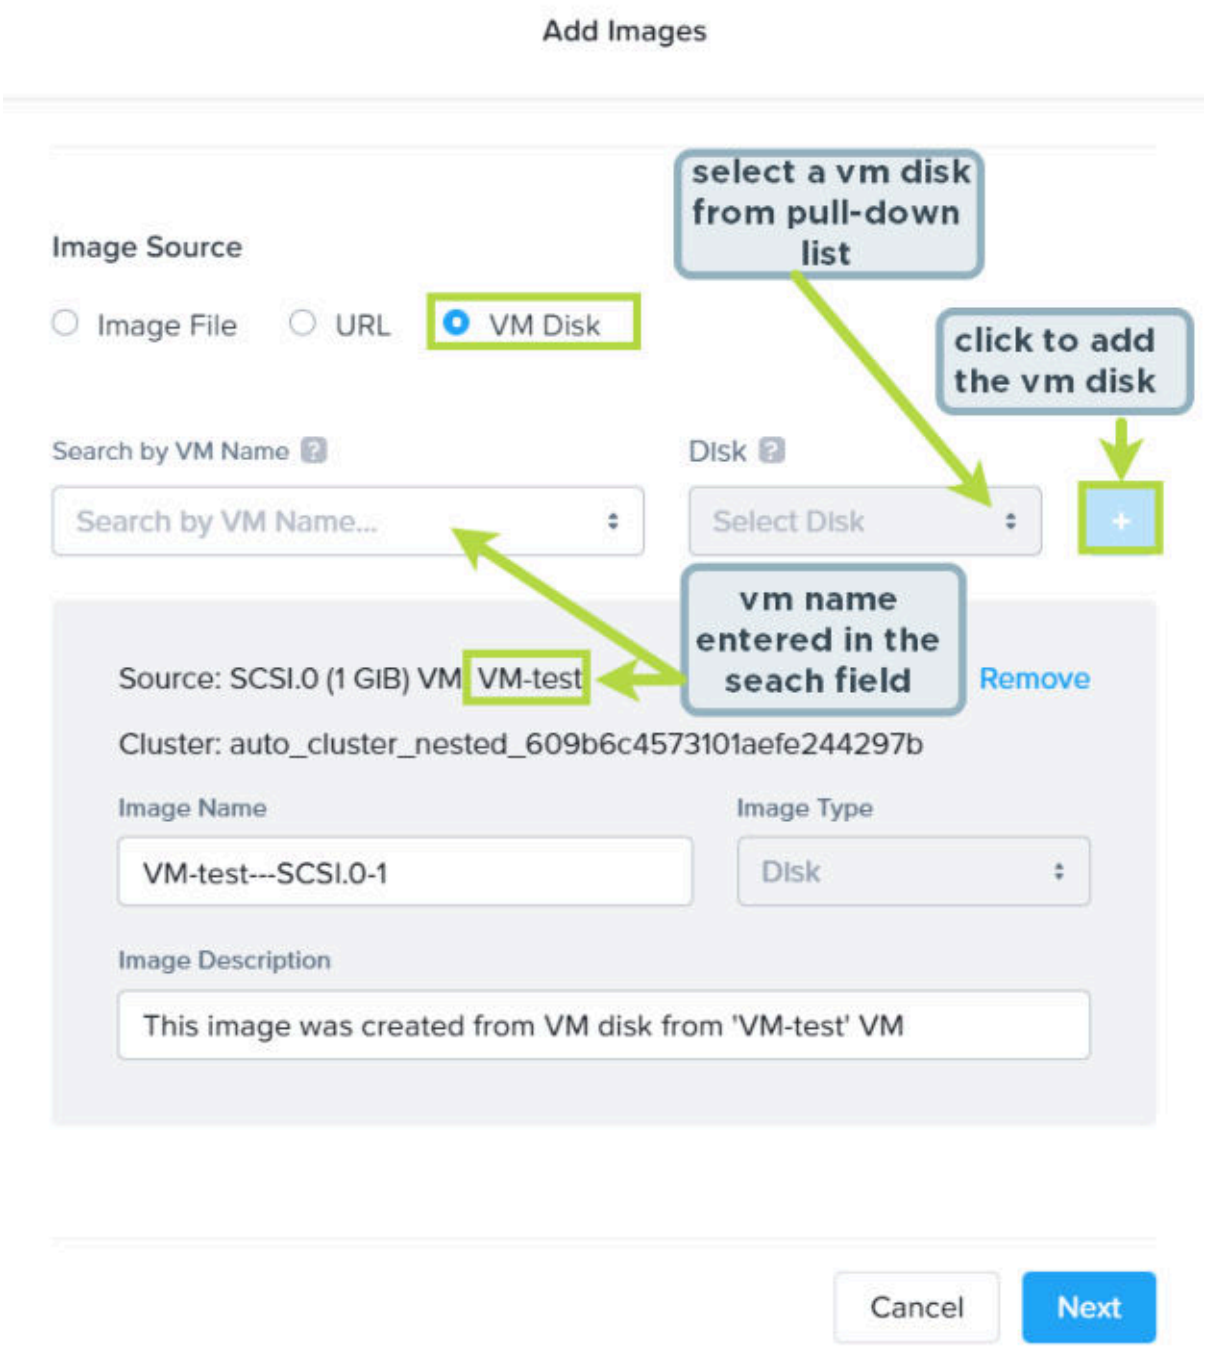

Figure 257: Add Images: VM Disk

3. To specify the VM disk, do the following:

   a. In **Search by VM Name**, click this field to populate the list of VMs from which you intend to clone the disk. The list displays all the powered-off VMs on the cluster. Alternatively, you can enter the VM name from which you intend to clone the disk.
   Prism Central adds the VM to a list on the page and populates the attached disks to this VM in the **Disks** pull-down list.

   b. In **Disk**, select the disk from the pull-down list and click the **+** button.

      The name defaults to the name of the file you selected, but you can change that name here.

      Prism Central displays the disk, and cluster information of the VM source.

4. To specify the image information, do the following:

   a. In **Image Name**, you can optionally change the auto-generated image name that defaults to the combination of the selected VM and disk name.

   b. In **Image Type**, Prism Central displays the type as *Disk* by default. You cannot modify this field.

   c. In **Image Description**, enter a description for the image.

5. To remove an image file entry, locate the entry and click **Remove**.

6. After you have added all the VM disks, click **Next**.
   The **Select Location** page appears.

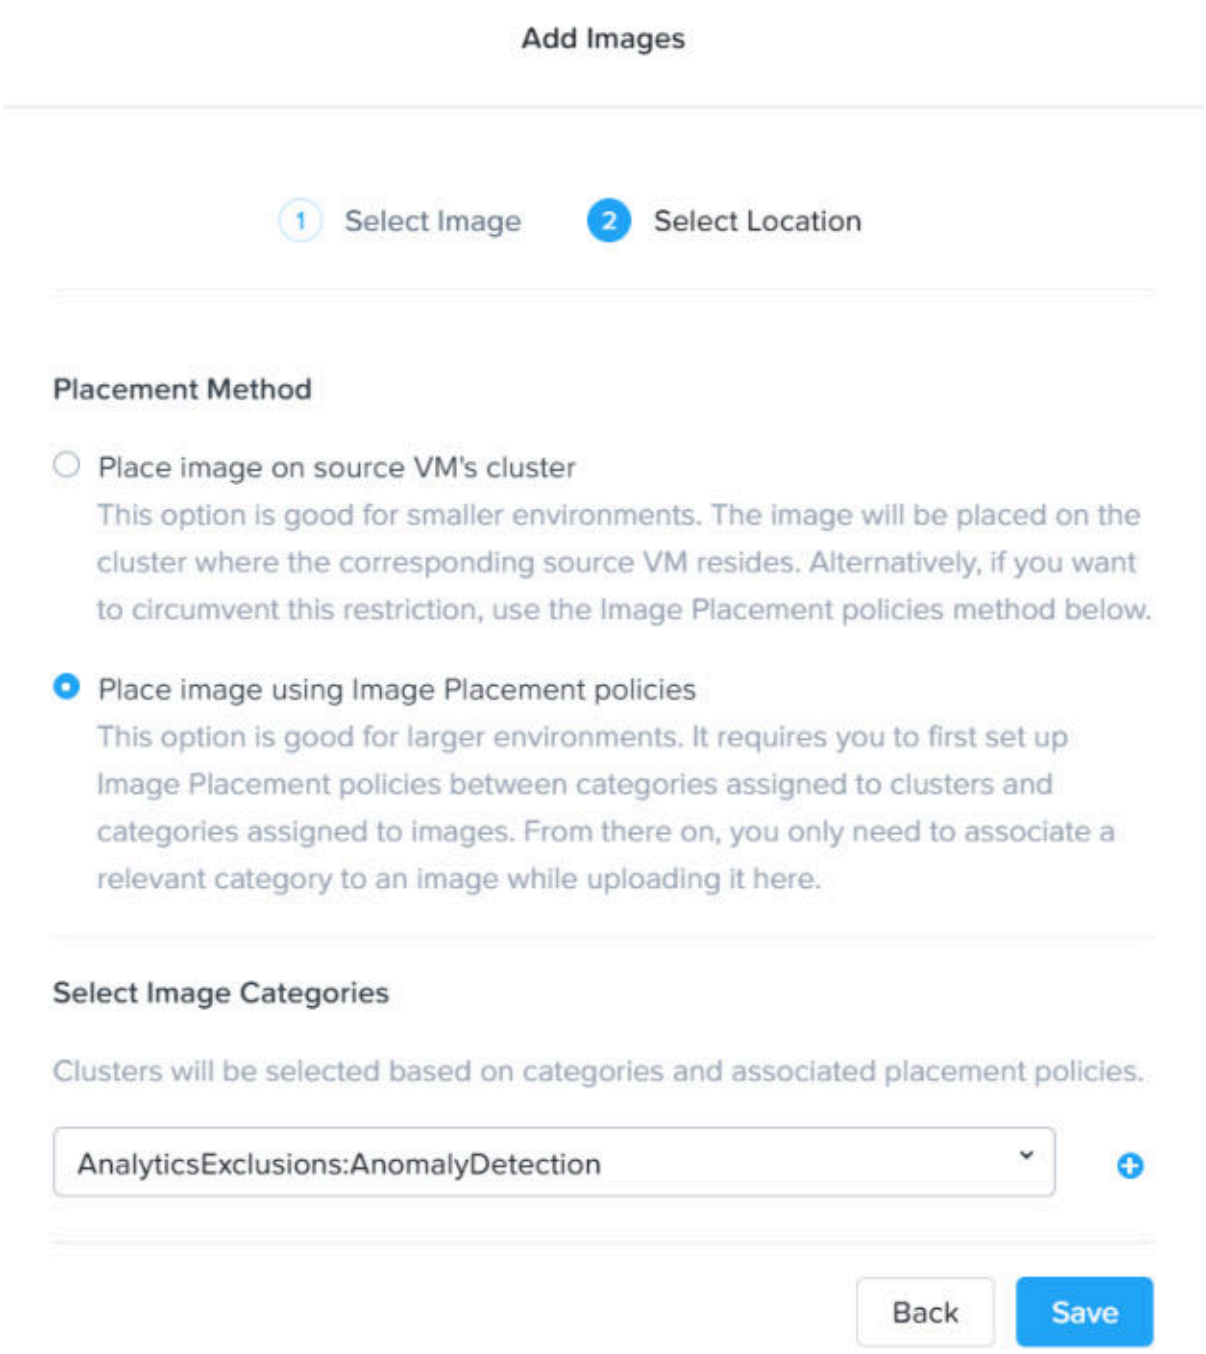

Figure 258: Add Images: Select Location

7. In **Placement Method**, do one of the following:

   » To select the source VM cluster, click **Place image on source VMs cluster** .

   » To delegate image placement decisions to configured policies and assign categories to the images, click **Place image using Image Placement policies**, and then do the following in the **Categories** section.

8. If you selected **Place image using Image Placement policies**, do the following:
   a. To assign a category, click the **Select Image Categories** search box.
   b. Select the desired category from the list.
      You can also start typing the name of the category to reduce the list to matching names.
   c. Click the add icon beside the search box.
   d. To specify another category, repeat these steps.
9. Click **Save**.
   Prism Central uploads the image files in batches.

Image Upload Verification

Verify whether an image upload to multiple clusters succeeded.

When you add an image and select multiple clusters for initial placement of the image (using the **Place image directly on clusters**), the status message box shows that the task succeeded. The status message displays the successful completion of the task even if the image is successfully uploaded to only one cluster. Therefore, go to **Tasks** and check whether the Image upload and Image update subtasks succeeded for all the clusters that you selected. It is possible that the Image upload or update subtasks for some of the clusters failed.

Nutanix recommends that you use Categories and Image Placement Policy based image placement option when you add images to multiple clusters. Image Placement policy based placement ensures that images propagate successfully to multiple clusters in a single operation. See Creating an Image on page 395 and Associating an Image with Categories on page 481 for more details about adding an image.

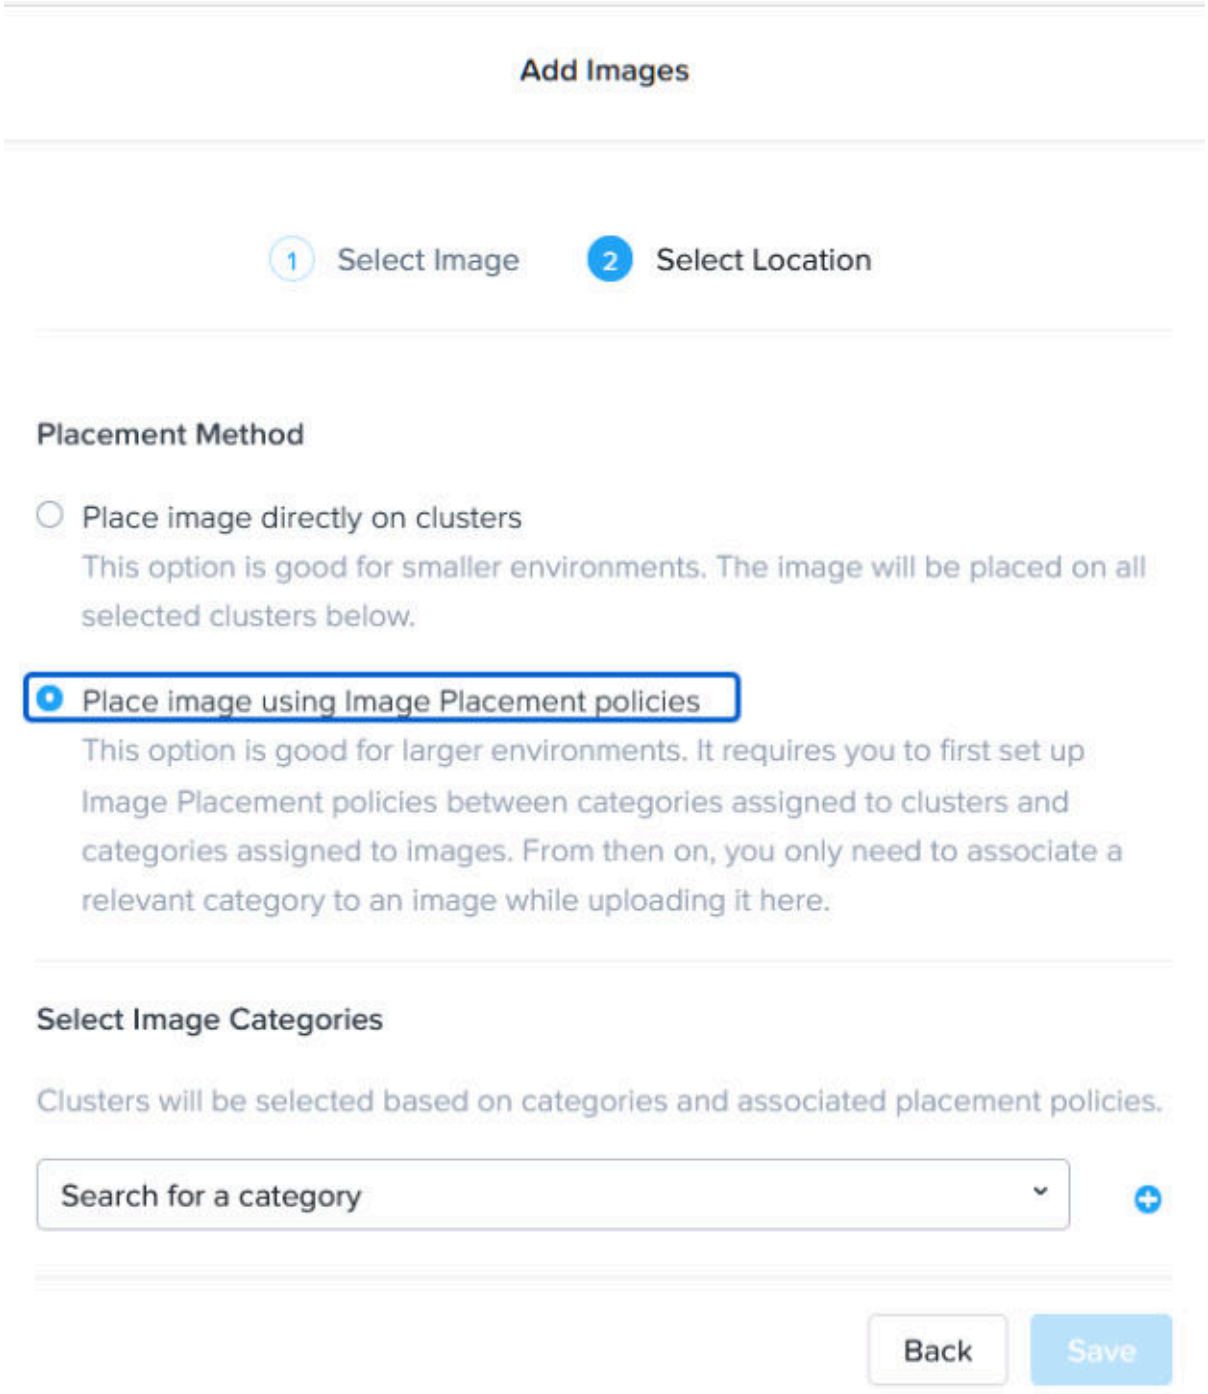

Figure 259: Image Placement Method selection for multiple clusters

Modifying an Image

About this task

To modify an image (delete, update, or add to catalog), in the **Entities** menu, go to **Virtual Infrastructure** > **Images** (see Images Summary View on page 125), and then do the following:

Procedure

- To delete an image, select (check the box for) the target image and then select **Delete** from the **Action** pull-down menu.
- To update an image, select the target image and then select **Update Image** from the **Action** pull-down menu. In the Edit Image page, update the name, description, or type as desired, and then click the **Save** button.

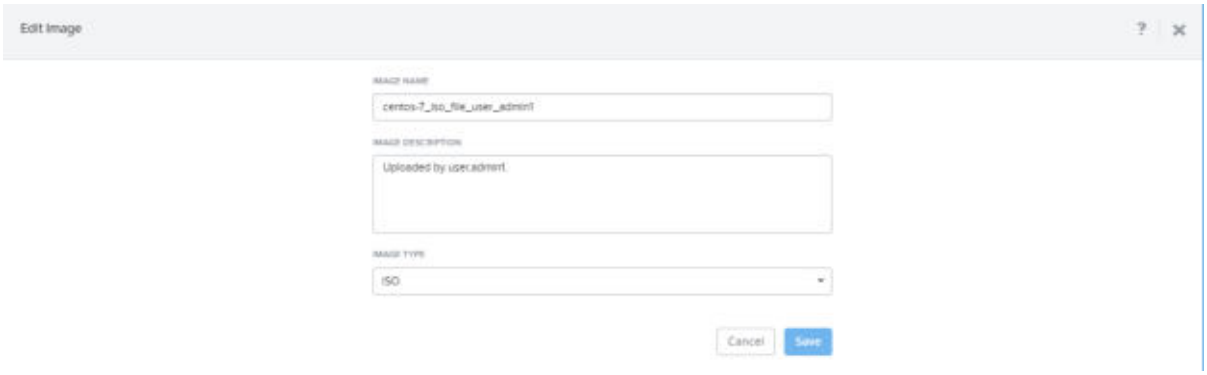

Figure 260: Update Image

- To add an image to the catalog (see Catalog Management on page 390), select the target image and then select **Add Image to Catalog** from the **Action** pull-down menu.

Importing Images to Prism Central

You can import images from registered clusters and manage the images centrally from Prism Central. An image imported to Prism Central continues to reside on the cluster that owns it. Prism Central only creates and stores image metadata locally, and it uses that metadata when you perform an action on the image. After you import an image, the image remains visible on the cluster from which you imported it but you cannot update the image from the cluster. You can update the image only from Prism Central.

About this task

The image import feature does not affect the following images:

- Any images that you choose not to import.
- Any images that you upload subsequently to the cluster from its web console. These images remain editable on the cluster until you import them to Prism Central.

In a single operation, you can import all images from all registered clusters, all images from a selection of clusters, or a selection of images from some of the clusters.

To import images from registered clusters, do the following:

Procedure

1. In the **Entities** menu, go to **Virtual Infrastructure** > **Images**, and then click the **Import Images** button (see Images Summary View on page 125).

2. In the **Import Images** dialog box, under **Select Import Type**, do one of the following:

   » To import all images from all registered clusters, click **All Images**.

   » To import all images from a selection of registered clusters, click **Images On a Cluster**, and then, in the list of clusters, select the clusters.

   » To import specific images from a given cluster, click the **Select Images** link provided for the cluster. In the dialog box that is displayed, select the images that you want to import, and then click **Done**.

   Repeat this step for all the clusters from which you want to import images.

3. To begin the import, click **Save**.
   Prism Central imports the metadata of the selected images and marks the images as read-only entities on the clusters.

Category Management

A category is a grouping of entities into a key value pair. Typically, new entities are assigned to a category based on some criteria. Policies can then be tied to those entities that are assigned (grouped by) a specific category value.

For example, you might have a Department category that includes values such as engineering, finance, and HR. In this case you could create one backup policy that applies to engineering and HR and a separate (more stringent) backup policy that applies to just finance. Categories allow you to implement a variety of policies across entity groups, and Prism Central allows you to quickly view any established relationships.

The following hypothetical example illustrates the relationship of four policies (Backupbasic, Engenvironment, Hourly alerts, and Daily backup) tied to three departments (Eng, Fin, and HR) that apply to 30+ VMs in each department.

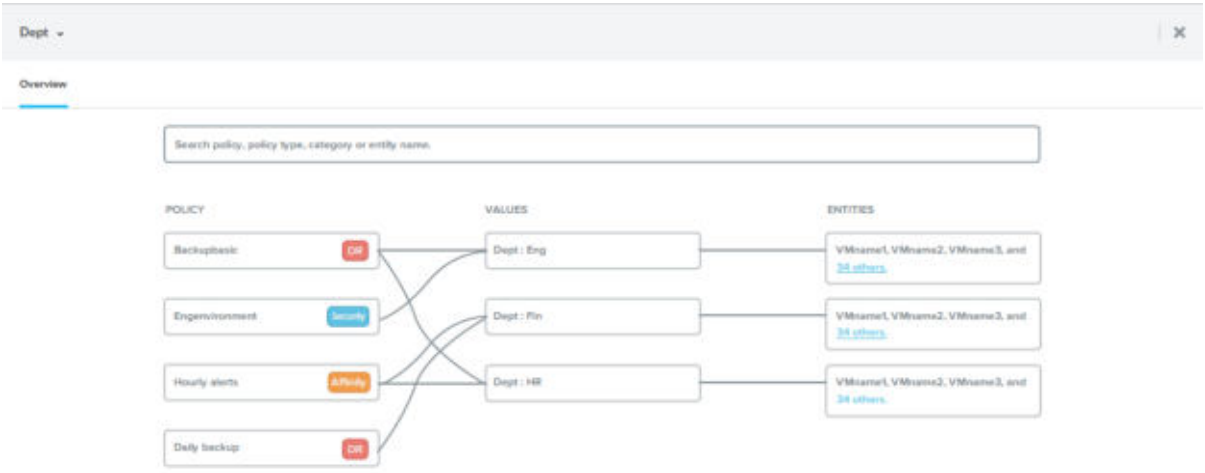

Figure 261: Dept Details View (example)

Creating a Category

About this task

To create a category, do the following:

Figure 262: Creating a Category

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275774441001 on Brightcove.

Procedure

1. Go to the categories dashboard (see Categories Summary View on page 128) and click the **Create Category** button.
   The Create Category page appears.

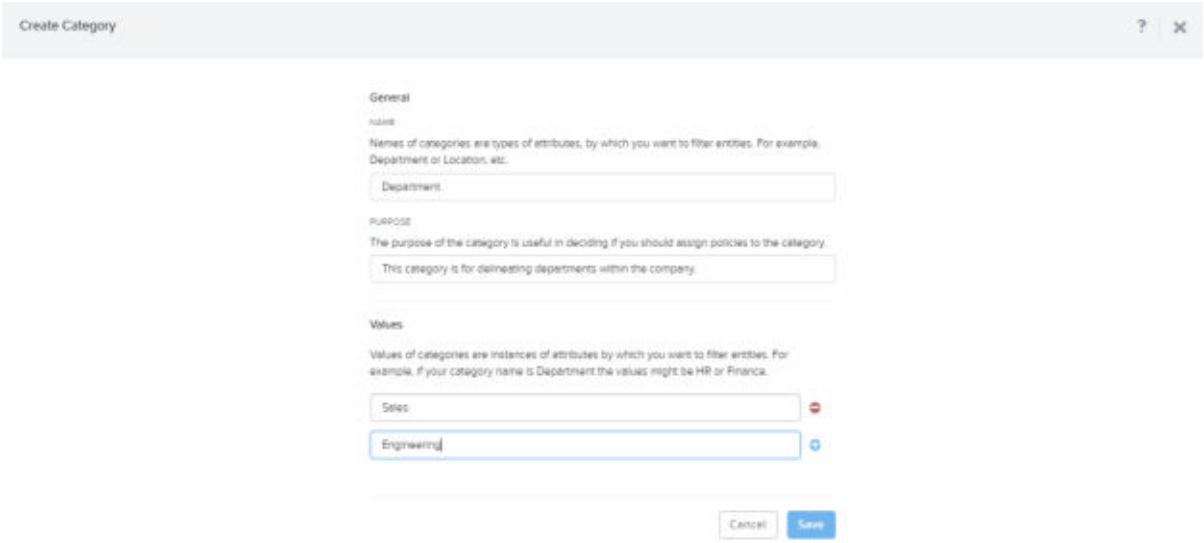

Figure 263: Create Category Page

2. Do the following in the indicated fields:
   a. **Name**: Enter a name for the new category.

      Note: The **Name** and **Values** fields are case sensitive.

   b. **Purpose** (optional): Enter a description of the category purpose.
   c. **Values**: Enter a category value. To add a second (and subsequent) value, click the plus sign (+) to the right. This opens another line; enter the next value in the new field.

      Repeat this step for all the values you want to include in the category. For example, if the category name is Departments, values might include Engineering, HR, Sales, Marketing, and so on.

      Duplicate values are discarded. Character case is considered when determining duplicates. For example, multiple instances of the value Sales are considered duplicates and

all instances but one are dropped; the values Sales and sales are considered unique and are accepted.

d. Click the **Save** button.

This creates the category and closes the page. The new category now appears in the category list.

Modifying a Category

About this task

To update or delete an existing category, do the following:

> Note: System (built-in) categories cannot be modified or deleted.

Procedure

1. Go to the categories dashboard (see Categories Summary View on page 128) and select (check the box for) the desired category from the list.
2. Do one of the following:

» To modify that category, select **Update** from the **Actions** pull-down menu. The Update Category page appears, which includes the same fields as the Create Category page (see Creating a Category on page 414). Update the field values as desired and then click the **Save** button.

» To delete that category, select **Delete** from the **Actions** pull-down menu. You are prompted to verify the delete (click the **OK** button). The category is then deleted and removed from the list.

> Note: You cannot delete a category if it is used in an existing policy. All associations with existing policies must be removed before a category can be deleted.

Assigning a Category

About this task

You can assign a category value to an entity of the following types: cluster, VM, image, or subnet. To assign a category value to one or more entities, do the following:

Procedure

1. Go to the entity type dashboard and select all the entities of that type you want to tag with the same category value(s).

- Cluster: Go to the **List** tab in the Clusters dashboard (see Clusters Summary View on page 146) and select (check the boxes for) the target clusters.
- VM: Go to the **List** tab in the VMs dashboard (see VMs Summary View on page 82) and select the target VMs.
- Image: Go to the Images dashboard (see Images Summary View on page 125) and select the target images.
- Subnet: Go to the Subnets dashboard (see Subnets View on page 135) and select the target subnets.

2. Select **Manage Categories** from the **Actions** pull-down menu.
   This displays the Manage [Cluster|VM|Image|Subnet] Categories page.

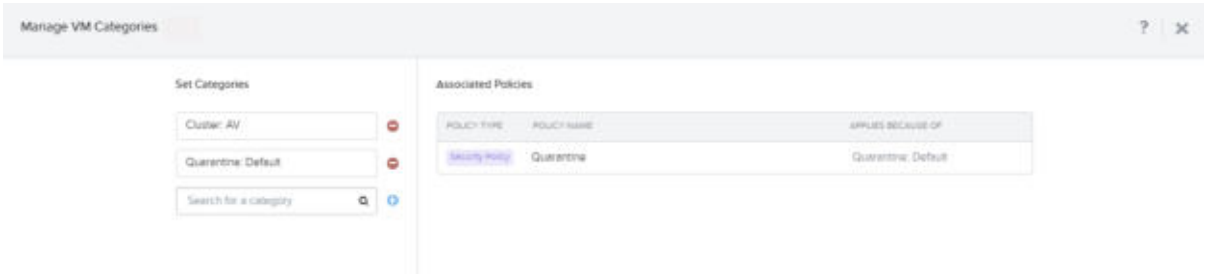

Figure 264: Manage VM Categories Page

3. In the Manage [Cluster|VM|Image|Subnet] Categories page, do the following:

   a. Enter a category name in the **Set Categories** field, select the target value from the list, and then click the plus sign (+) to the right of the field to assign that category value to the selected entities.

      The **Set Categories** field acts like a search field; it provides a list of matching categories as you enter a string. Select the desired category value when you see it in the list. Any policies associated with the selected category value appear in the **Associated Polices** section to the right.

   b. Repeat the first step to assign a value for a second category.

      You can repeat this step for as many categories as desired. To illustrate, in the figure above the VMs are assigned two values, "AV" from the "Cluster" category and "Default" from the "Quarantine" category. In this example "Cluster:AV" has no policies associated with it currently, but "Quarantine:Default" is associated with the Quarantine security policy.

      > Note: Categories support multi-cardinality, which means you can assign multiple category values to the same entity. In this case you can assign multiple values to the same VM.

   c. Click the **Save** button (bottom right).

      This saves the category assignments and closes the page.

Affinity Policies for AHV

As an administrator in an Acropolis managed clusters, you can specify scheduling policies for virtual machines on an AHV cluster. By defining these policies, you can control placement of the virtual machines on the hosts within a cluster.

You can define two types of affinity policies.

VM-Host Affinity Policy

The VM-host affinity policy controls the placement of the VMs. You can use this policy to specify that a selected VM can only run on the members of the affinity host list. This policy checks and enforces where a VM can be hosted when you power on or migrate the VM.

> Note:
>
> - If you choose to apply the VM-host affinity policy, it limits Acropolis HA and Acropolis Dynamic Scheduling (ADS) in such a way that a virtual machine cannot

be powered on or migrated to a host that does not conform to requirements of the affinity policy as this policy is enforced mandatorily.

- The VM-host anti-affinity policy is not supported.

You can define the VM-host affinity policies by using Prism Central during the VM create or update operation. For more information, see Creating a VM (AHV) on page 317.

Note: If a VM is cloned that has the affinity policies configured, then the policies are not automatically applied to the cloned VM. However, if a VM is restored from a DR snapshot, the policies are automatically applied to the VM.

Limitations of Affinity Rules

- Even if a host is removed from a cluster, the host UUID is not removed from the host-affinity list for a VM.
- The VM-host affinity cannot be configured on a cluster that has HA configured by using reserved host method.

Operations Policy Management

Prism Central implements a variety of system services to analyze cluster integrity and performance. You have the option to apply policies to customize some services in your environment. You can apply the following policies:

- **Anomaly Detection Exclusion**. Prism Central includes machine-learning capabilities that analyze resource usage and identify potential problems (see Behavioral Learning Tools on page 582). One is anomaly detection that records when performance or resource usage is outside an expected range. The anomaly detection exclusion policy allows you to specify VMs to ignore when looking for anomalies. For example, you might have some VMs with expected peak load on certain days, and you do not want that usage surge counted as an anomaly. In this case you can either define a category for those VMs and enable anomaly detection exclusion for that category or enable the exclusion for each VM individually.
- **Inefficiency Measurement Exclusion**. Another machine-learning capability is VM "right sizing" that identifies VMs which exhibit inefficient profiles. Again, you have the option to exclude selected VMs from the inefficiency analysis.

Operations Policies View

The operations policies dashboard allows you to view and modify operations policies that can be applied to VMs or other entities.

To access the operations policies dashboard, select **Operations** > **Ops Policies** from the entities menu (see Entities Menu on page 13). The dashboard displays a list of the operations policies.

Note: This section describes the information and options that appear in the operations policies dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

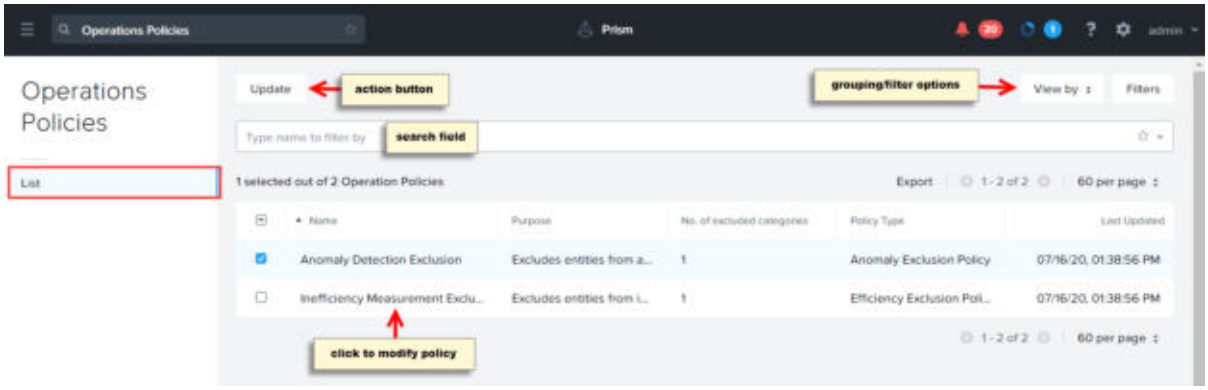

Figure 265: Operations Policies Dashboard

The following table describes the fields that appear in the operations policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 218: Operations Policy Fields

Parameter	Description	Values
Name	Displays the application name for the instance. Click the name to display the Application Info window. The window displays the application name, VM name, TCP ports, and UDP ports.	(app name)
Purpose	Describes the purpose of the policy.	(VM name)
No. of Excluded Categories	Displays the number of categories the policy applies to (see Category Management on page 413).	(integer)
Policy Type	Describes what type of policy this is.	(type name)
Last Updated	Displays the date and time when the policy was last updated.	(date and time)

You can filter the list either by entering a string in the search field or by clicking the **Filters** button, checking the **Name** box, selecting a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and entering a string in the field.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

To modify a policy, select that policy and then select **Update** from the **Actions** pull-down menu (see Modifying an Operations Policy on page 418).

Modifying an Operations Policy

You can customize the list of categories to which an operations policy applies.

About this task

No categories are assigned to the Anomaly Detection Exclusion and Inefficiency Measurement Exclusion polices by default. To add (or delete) one or more categories to an operations policy, do the following:

Procedure

1. Go to the operations policy dashboard, select the desired policy, and then select **Update** from the **Actions** pull-down menu (see Operations Policies View on page 417).
2. In the policy page, click the **Add** link.

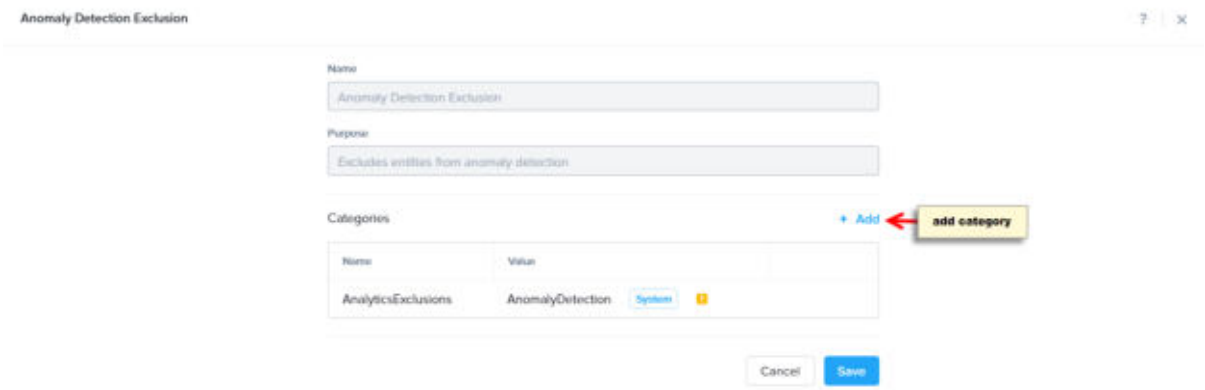

Figure 266: Anomaly Detection Exclusion Page (add category)

3. In the Add window, enter a category name and then click the plus icon. Repeat this step for all categories you want to add.

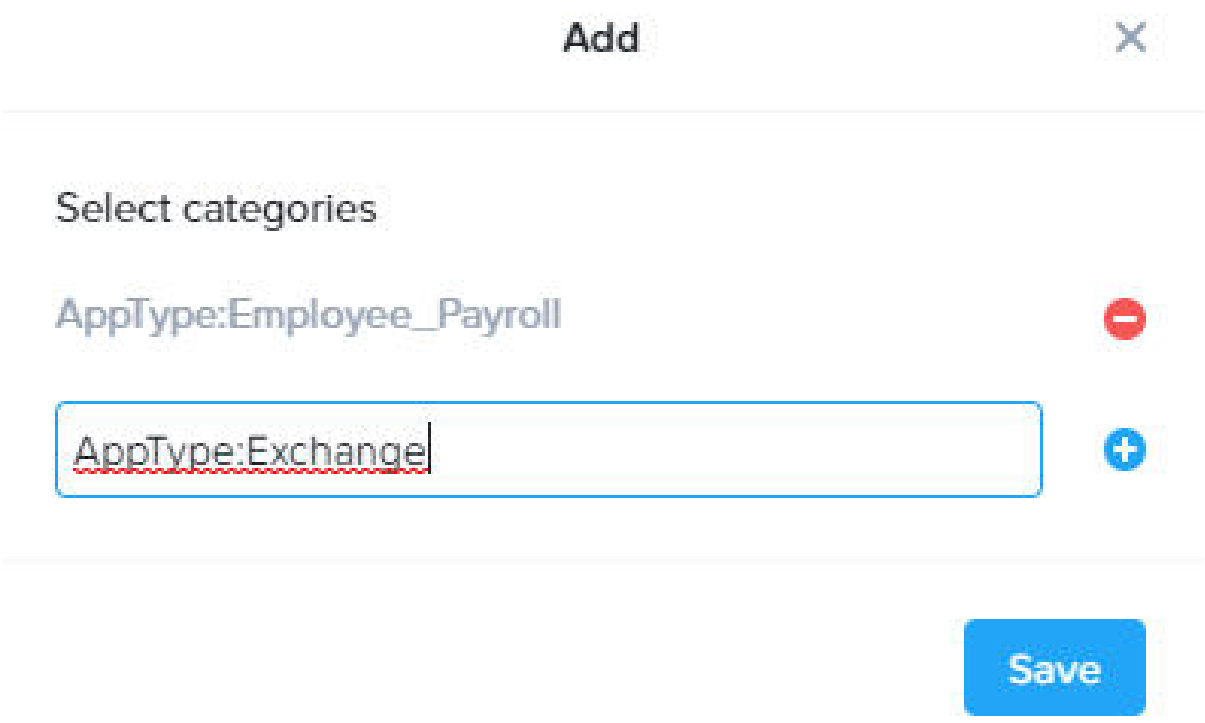

Figure 267: Add Window

4. Click the **Save** button.

   The policy is now enabled for entities in the selected categories.

5. To delete a category from the list, click the garbage pail icon for that category and then click the **Save** button.

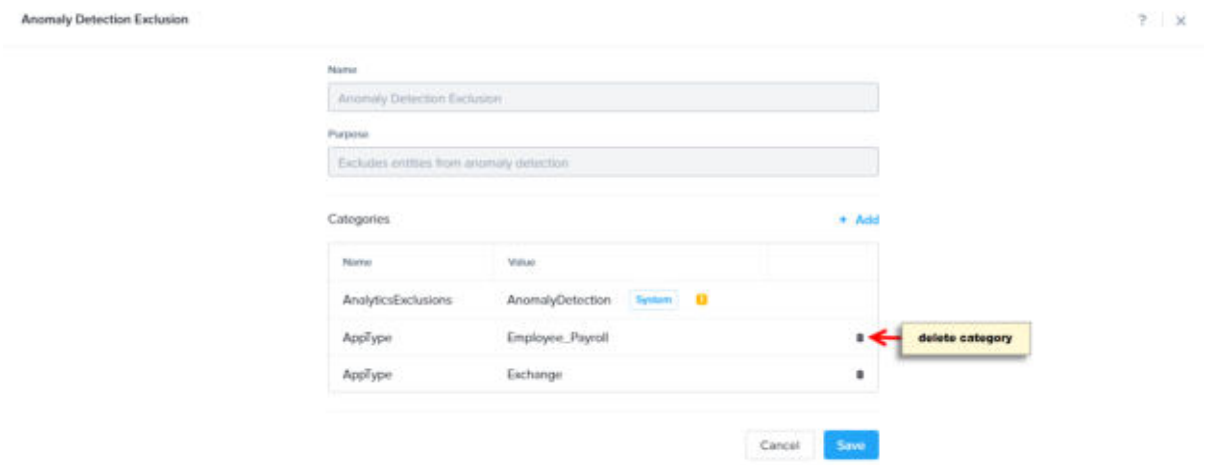

Figure 268: Anomaly Detection Exclusion Page (delete category)

Enabling an Operations Policy

You can enable (or disable) an operations policy for selected VMs.

About this task

To disable anomaly detection or inefficiency measurement for selected VMs or for all the VMs in selected clusters, do the following:

Procedure

1. Do one of the following:

   » Go to the List tab of the VMs dashboard, select the target VMs, and then select **Disable Efficiency Measurement** or **Disable Anomaly Detection** from the **Actions** pull-down menu (see VMs Summary View on page 82).

   » Go to the List tab of the clusters dashboard, select the target clusters, and then select **Disable Efficiency Measurement** or **Disable Anomaly Detection** from the **Actions** pull-down menu (see Clusters Summary View on page 146).

   This immediately disables the specified operation (enables the exclusion policy) for the target VMs.

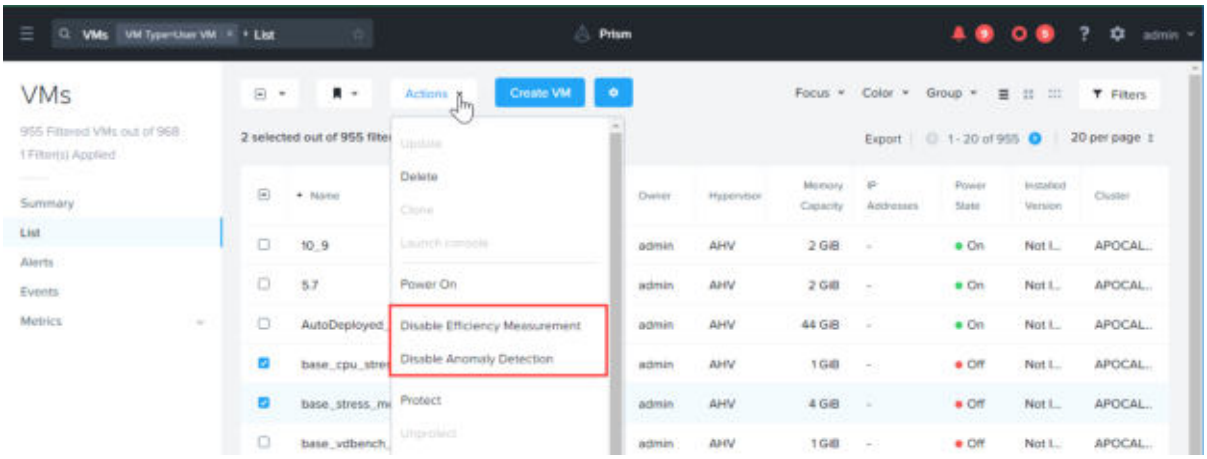

Figure 269: VM Actions Menu

2. To re-enable the operation, select the target VMs or clusters, and then select **Enable Efficiency Measurement** or **Enable Anomaly Detection** from the **Actions** pull-down menu.

   This immediately enables the specified operation (disables the exclusion policy) for the target VMs.

Network and Security Management

You can view and configure network connections for the cluster in the **Network Configuration**. To display the **Network Configuration** page, click the **Network Config** button in either the VMs or Subnets dashboard.

Network Connections

Each VM network interface is bound to a virtual network, and each virtual network is bound to a single VLAN. Information about the virtual networks configured currently appears in the Network Configuration page.

The Network Configuration page includes three tabs.

- The **Subnets** tab displays a list of the configured networks.

Figure 270: Networks Tab

- The **Internal Interfaces** tab displays a list of LAN interfaces.

Figure 271: Internal Interfaces Tab

- The **Virtual Switch** tab displays a list of virtual switches configured, including the default system-generated virtual switch vs0.

Figure 272: Virtual Switch Tab

The following table describes the fields in each tab.

Table 219: Network Configuration Fields

Parameter	Description	Values
Networks Tab		
Name	Displays the name of the virtual network.	(name)
Virtual Switch	Displays the name of the virtual switch in the form **vs#**, for example **vs0** for virtual switch 0 which is the default virtual switch.	(vs<number>)
VLAN ID	Displays the VLAN identification number for the network in the form **vlan.#**, for example **vlan.27** for virtual LAN number 27.	(ID number)
Internal Interfaces Tab		

Parameter	Description	Values
Descriptive Name	Displays a name for the LAN.	(LAN name)
Subnet (Gateway IP / Prefix Length)	Displays the subnet that the internal interface belongs to in the form *<IP Address>*/<number (prefix)>	(IP Address/prefix number)
Features	Displays the features available on the internal interfaces.	
Interface	Displays the interface designation such as eth0 or eth1.	(interface name)
Virtual Switch		
Name	Displays the name of the switch in the form vs#.	(vs<number>)
Bridge	Displays the name of the bridge associated with the virtual switch in the form br#, for example br0 for the default bridge.	(br<number>)
MTU (bytes)	Displays the MTU set for the virtual switch in bytes. The default MTU is 1500.	(number)
Bond Type	Displays the uplink bond type associated with the virtual switch. See the *Bond Type* table in . For example, Active-Backup	(<bond_type>)

LAG and LACP on the ToR Switch

For more information, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 431.

Configuring Network Connections

Prism Central allows you to configure network connections for a selected cluster.

About this task

To create one or more network configurations, do the following:

Procedure

1. Do one of the following:
   - » Go to the List tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Network Config** button.
   - » Go to the Network and Security dashboard (see Subnets View on page 135 and click the **Network Config** button.

2. In the Subnets tab of the Network Configuration page, click the **Create Network** button.
   This displays the Create Network page.

Figure 273: Create Network Page (expanded)

3. Do the following in the indicated fields:

   a. **Subnet Name**: Enter a name for the network.

   b. **Virtual Switch**: Select the virtual switch for the network from the drop-down list.

      You can associate the default virtual switch vs0 or create a virtual switch. To create a virtual switch, see Creating or Updating Virtual Switch on page 426.

      For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

   c. **VLAN ID**: Enter the number of the VLAN.

      Enter just the number in this field, for example 1 or 27. Enter 0 for the native VLAN. The value appears as **vlan.1** or **vlan.27** in displays.

   d. **Enable IP Address Management**: Check the box to have the cluster control IP addressing in the network.

      Checking this box display additional fields. If this box is not checked, no network management is attempted. In this case it is assumed management for this virtual LAN is handled outside the cluster.

   e. **Network IP Address/Prefix Length**: Enter the IP address of the gateway for the network and prefix with the network prefix (CIDR notation, for example, 10.1.1.0/24).

   f. **Gateway IP Address**: Enter the VLAN default gateway IP address.

   g. **Configure Domain Settings**: Check this box to display fields for defining a domain.

      Checking this box displays fields to specify DNS servers and domains. Unchecking this box hides those fields.

   h. **Domain Name Servers (comma separated)**: Enter a comma-delimited list of DNS servers.

   i. **Domain Search (comma separated)**: Enter a comma-delimited list of domains.

   j. **Domain Name**: Enter the VLAN domain name.

   k. **TFTP Server Name**: Enter the host name or IP address of the TFTP server from which virtual machines can download a boot file. Required in a Pre-boot eXecution Environment (PXE).

   l. **Boot File Name**: Name of the boot file to download from the TFTP server.

4. To define a range of addresses for automatic assignment to virtual NICs, click the **Create Pool** button (under **IP Address Pools**) and enter the following in the Add IP Pool page:

   Note: If no pool is provided, the user must assign IP addresses to VMs manually.

Figure 274: Add IP Pool Page

   a. Enter the starting IP address of the range in the **Start Address** field.

   b. Enter the ending IP address of the range in the **End Address** field.

   c. Click the **Submit** button to close the page and return to the Create Network page.

5. To configure a DHCP server, check the **Override DHCP server** box and enter an IP address in the **DHCP Server IP Address** field.

   This address (reserved IP address for the Acropolis DHCP server) is visible only to VMs on this network and responds only to DHCP requests. If this box is not checked, the **DHCP Server IP Address** field is not displayed and the DHCP server IP address is generated automatically. The automatically generated address is *network_IP_address_subnet*.254, or if the default gateway is using that address, *network_IP_address_subnet*.253.

6. Click the **Save** button to configure the network connection, close the Create Network page, and return to the Network Configuration page.

7. Click the **Close** button to close the Network Configuration page.

Creating or Updating Virtual Switch

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can create or update a VS using the Create Virtual Switch or Update Virtual Switch dialog box. The Create Virtual Switch and Update Virtual Switch dialog boxes have identical parameters and settings.

To open the Create Virtual Switch dialog box in Prism Central, do the following.

Procedure

1. Click **Subnets** > **Network Config**.

2. On the Network Configuration page, click the **Virtual Switch** > **Create VS** link.

About this task

To create a virtual switch using the Create Virtual Switch dialog box, do the following.

Procedure

1. In the Create Virtual Switch dialog box, provide the necessary information on the **General** tab.

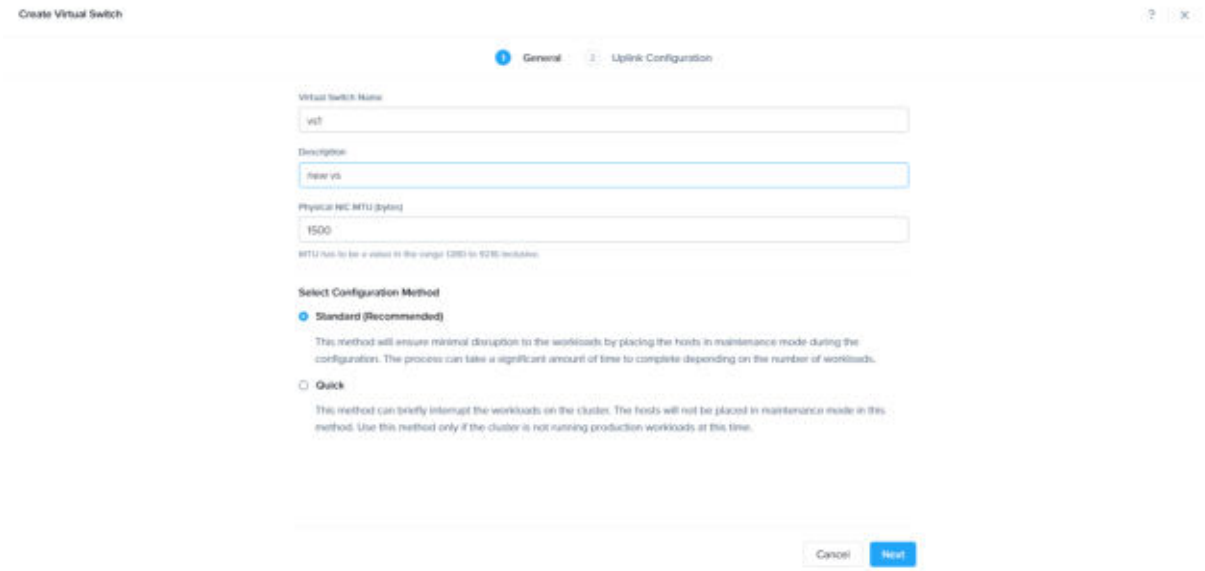

Figure 275: Create Virtual Switch - General tab

Field	Description
Virtual Switch Name	Enter a name for the virtual switch.
Description	Provide a description for the virtual switch that helps identify the virtual switch.
Physical NIC MTU (bytes)	MTU must be a value in the range 1280 to 9216 inclusive.

Field	Description
Select Configuration Method	Select one of the two methods that you can use to implement the VS configuration: • **Standard (Recommended)**: This method ensures no disruptions occur to the workloads by putting the hosts in maintenance mode and migrating the VMs out of the host before applying the configuration. This process requires a longer duration of time to complete. The time required depends on the number and configuration of VMs. In this method, the VS configuration is deployed in the rolling update process. • **Quick**: This method interrupts the workloads running on the hosts. Use this method only if you are not running production workloads because it may result in network interruptions. In this method, the VS configuration is deployed in a rolling update process but the nodes are not put in maintenance mode before modifying the VS configuration on the node.

2. Click **Next** to go to the Uplink Configuration tab.

   On the Uplink Configuration tab, provide the following details:

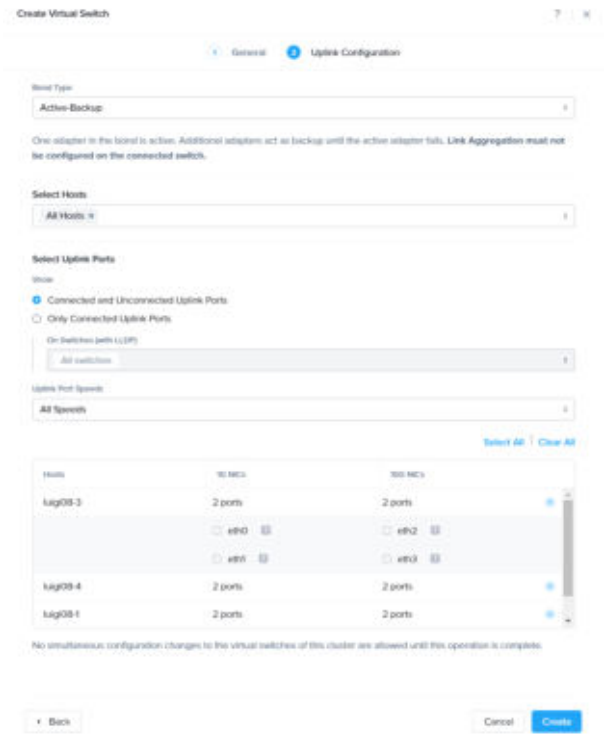

Figure 276: Create Virtual Switch - Uplink Configuration tab

Field	Description and Value
Bond Type	Select an appropriate bond type. See the *Bond Types* table for details about the bond types.
Select Hosts	Select the hosts that can host VMs.
Select Uplink Ports	Select the criteria that need to be satisfied for the uplink ports. The available uplink ports that satisfy the criteria are displayed in the (Host port) table at the bottom of this tab.
Show	(Port Type) Select one of the following: **Connected and Unconnected Uplink Ports**: Select this if you want to use ports that are not currently connected but may be connected later. **Only Connected Uplink Ports**: Select this option if you want to use only the connected ports. You must also select in the **On Switches (with LLDP)** drop down box the switches with LLDP.
Uplink Port Speeds	Select a speed to display the ports that have the selected speed. You can select speeds such as **1G**, **10G** or both (**All Speeds**). The speeds displayed depend on the NIC type that is installed on the host. Based on your selection the columns in the (Host Port) table change dynamically to display the ports with the speeds you selected.
(Host Port) table	Based on the Selections you made in this Select Uplink Ports section, a table displays the hosts that have the uplink ports that satisfy the selected criteria. Select the ports you need for this configuration from the list. Click the down arrow on the right side of the table to display the ports listed for each host. Note: A port listing is greyed out if it is unavailable because it is already associated with another virtual switch. Click **Select All** to select all the ports available and listed. Click **Clear All** to unselect all the ports available and listed.

Table 220: Bond Types

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Backup	Recommended. Default configuration, which transmits all traffic over a single active adapter.	10 Gb	10 Gb

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Active with MAC pinning Also known as balance-slb	Works with caveats for multicast traffic. Increases host bandwidth utilization beyond a single 10 Gb adapter. Places each VM NIC on a single adapter at a time. Do not use this bond type with link aggregation protocols such as LACP.	10 Gb	20 Gb
Active-Active Also known as LACP with balance-tcp	LACP and link aggregation required. Increases host and VM bandwidth utilization beyond a single 10 Gb adapter by balancing VM NIC TCP and UDP sessions among adapters. Also used when network switches require LACP negotiation. The default LACP settings are: • Speed—Fast (1s) • Mode—Active fallback-active-backup • Priority—Default. This is not configurable.	20 Gb	20 Gb
No Uplink Bond	No uplink or a single uplink on each host. Virtual switch configured with the No uplink bond uplink bond type has 0 or 1 uplinks. When you configure a virtual switch with any other bond type, you must select at least two uplink ports on every node.	-	-

Note: The Maximum VM NIC Throughput and Maximum Host Throughput values are not restricted to the value provided in this table. The values in the table are indicated for an assumption of 2 x 10 Gb adapters.

For more information about uplink configuration, see the *Uplink Configuration (AHV Only)* topic in the *AHV Administration Guide*.

3. Click **Create** to create the virtual switch.

   Click **Cancel** to exit without creating the virtual switch.

   Click **Back** to go back to the General tab.

What to do next

> Important:
>
> You can migrate or convert the bridges other than br0 in the cluster to virtual switches after you upgraded the minimum or compatible version of AOS and AHV. You can convert the other bridges only on Prism Web Console or using aCLI. You can convert only one bridge at a time. You need to repeat the workflow for every bridge that you want to convert to a virtual switch.
>
> For more information about converting the bridges other than br0, see *Migrating Bridges after Upgrade*in *Prism Web Console Guide*.

- If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

  For information about enabling LAG and LACP, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 431

Deleting a Virtual Switch

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can delete a VS on the Virtual Switch tab in the Network Configuration dialog box. When you delete a virtual switch, the bridge created by the virtual switch is also removed completely.

You can also disable all the virtual switches. However, you can do so only using CLI. Prism does not support disabling virtual switches. Use the net.disable_virtual_switch to disable the virtual switch functionality and remove all the virtual switches including the default virtual switch vs0. However, the bridge and uplink bond or interface configurations are not affected by this command. This action cannot be undone either.

To delete a virtual switch, do the following:

Procedure

1. Click the **Delete** icon (trash icon) next to the virtual switch listed in the list on the Virtual Switch tab.

   The Delete Virtual Switch dialog box opens. Deletion cannot be undone or reversed. Therefore, you must confirm the deletion in this dialog box.

2. Click **Delete** to confirm deletion.

   Click **Cancel** to close the Delete Virtual Switch dialog box without deleting the virtual switch.

What to do next

Check if the virtual switch is deleted from the list on the Virtual Switch tab.

Enabling LAG and LACP on the ToR Switch (AHV Only)

Active-Active with MAC pinning does not require link aggregation since each source MAC address is balanced to only one adapter at a time. If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

About this task

In an **Active-Active** NIC-teaming configuration, network traffic is balanced among the members of the team based on source and destination IP addresses and TCP and UDP ports. With link aggregation negotiated by LACP, the uplinks appear as a single layer-2 link so a VM MAC address can appear on multiple links and use the bandwidth of all uplinks. If you do not enable LAG and LACP in this **Active-Active** configuration, network traffic might be dropped by the switch.

Perform the following procedure to enable LAG and LACP on a TOR switch after enabling the Active-Active policy by using the Prism Element web console.

Procedure

1. Put the node in maintenance mode. This is in addition to the previous maintenance mode that enabled Active-Active on the node.
2. Enable LAG and LACP on the ToR switch connected to that node.
3. Exit maintenance mode after LAG and LACP is successfully enabled.
4. Repeat steps 1 to 3 for every node in the cluster.

Modifying Network Connections

About this task

To modify or delete a network configuration (defined on an Acropolis managed cluster), do the following:

Procedure

1. Do one of the following:
   - » Go to the List tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Network Config** button.
   - » Go to the Networks dashboard (see Subnets View on page 135 and click the **Network Config** button.
2. To modify a network configuration, in the Subnets tab, select the target network, and click the pencil icon.

   The Update Subnet page appears, which contains the same fields as the Create Subnet page (see Configuring Network Connections on page 423). Do the following:

   a. Update the field values as desired.

   b. When the configuration is correct, click the **Save** button to close the page, update the network configuration, and return to the Network Configuration page.
3. To delete a network configuration, in the Subnets tab, select the target network, and click the X icon (on the right).

   A window prompt appears to verify the action; click the **OK** button. The network configuration is removed from the list.

   > Note:
   >
   > You can also select the target network on the Networks dashboard and select **Delete** in the **Actions** dropdown list.

4. Click the **Close** button to close the Network Configuration page.

OVA Management

You can export a VM as an OVA, upload and manage OVAs only from Prism Central.

Of the following OVA operations, you can export a VM only from the VMs dashboard (see VMs Summary View on page 82). You can perform the other actions from the OVAs dashboard (see OVAs View on page 131).

- To export a VM in the VM dashboard as an OVA , see Exporting a VM as an OVA on page 368.
- To upload an OVA, see Uploading an OVA on page 434.
- To deploy an OVA file as a VM, see Deploying an OVA as VM on page 442.
- To download an OVA file to your local machine, see Downloading an OVA on page 448.
- To rename an OVA file, see Renaming an OVA on page 448.
- To delete an OVA file, see Deleting an OVAs on page 449 .

To perform any action an OVA in the OVAs dashboard, select the target OVA, then select the desired action in the **Actions** menu.

The task progress for these actions are displayed in the Tasks page.

> Note:
>
> The **Actions** menu appears only if you select an OVA in the list. If the OVAs list is empty, then the **Actions** menu is available only after you create an OVA by uploading it or exporting an existing VM as an OVA.

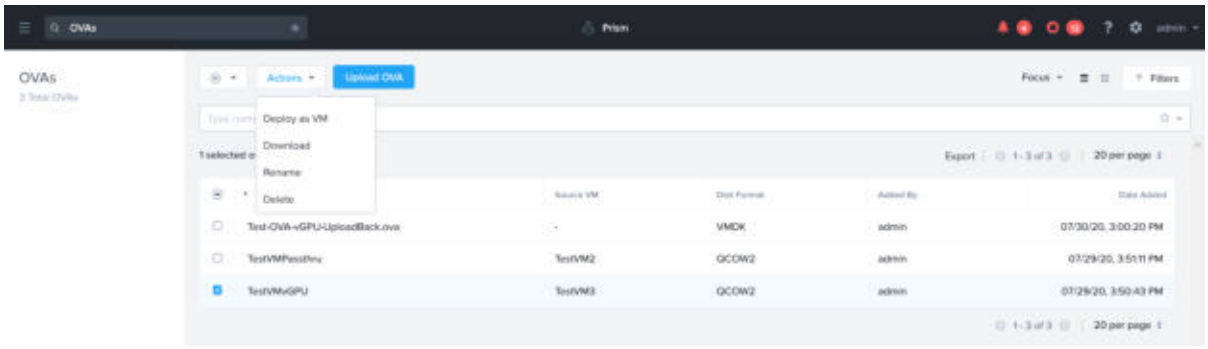

Figure 277: Actions menu

Exporting a VM as an OVA

Before you begin

If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the destination hypervisor.

About this task

As a Super Admin, you can export a VM as an OVA by selecting the **Export** as OVA action from the **Actions** dropdown list. You can export a VM as an OVA with a specified disk format for the disks configured on the VM.

> Note:
>
> - The Super Admin role was formerly called as User Admin.
> - If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the target destination hypervisor.

To export a VM as an OVA, do the following in the VMs dashboard:

Procedure

1. Select the VM in List tab.

   You can click the VM and open the VM details page.

2. In the **Actions** list, select **Export as OVA**.

   Alternatively, on the details page of the VM, select **More** and click **Export as OVA**.

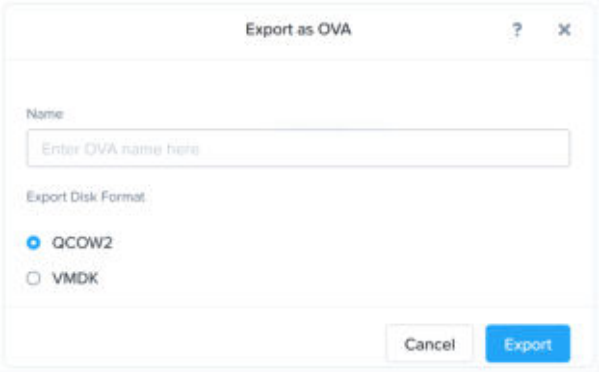

Figure 278: Export as OVA dialog box

3. In the Export as OVA window, do the following in the indicated fields:

   a. Enter a name for the OVA, in the **Name** field.

   b. Select the disk format that you want to format the disks in the VM to, in the **Export Disk Format** section.

      The default format is QCOW2. If you want to use the OVA to deploy a VM with disks formatted as VMDK, then select VMDK as the disk format.

4. Click **Export**.

   You can check the progress of the *Export as OVA* task in Tasks.

   The exported OVA is available on the OVA dashboard.

Uploading an OVA

About this task

You can upload an OVA from the following sources:

- From a local folder: When you upload an OVA from a local source, you can upload it to only one target cluster.
- From a URL: When you upload an OVA from a URL, you can upload it to multiple clusters. The upload operation runs concurrently on all the selected clusters.

Select the appropriate upload source considering the multiple-cluster upload option.

> Note:
>
> Prism Central processes up to three OVA file uploads concurrently for OVA file uploads from local folders. Prism Central queues other OVA files that you upload until one or more of the concurrent uploads are completed.
>
> The is no restriction on the number of concurrent uploads for OVA uploads from URL.

For OVA upload from local folders:

- The upload process can be interrupted due to the following reasons:
  - Network issues when the upload is running.
  - You close the browser window or tab in which the upload is running.
- The concatenate and validate operations of the OVA upload operation fail when an error occurs in the upload services because the checksum provided for the OVA file when it was created is incorrect.

  > Note:
  >
  > If the upload process is interrupted due to any reason, you can resume the process. When the validation operation is interrupted due to a concatenation error then you must run the concatenate API to resume concatenation and validation. See Concatenating Upload using APIs on page 440.

- You can also check the status of the upload in Tasks. Prism Central displays the following three sub-tasks in Tasks:
  - OVA Create
  - OVA Upload
  - OVA Validate
- Prism Central uses indicators to display the status of OVAs that you may have tried to upload. The status message is displayed when you hover the mouse or pointing device cursor on the indicator. See *OVA Alerts* table in OVAs View on page 131.

To upload an OVA, do the following:

Procedure

1. To upload an OVA, click **Upload OVA**.

2. In the Upload OVA dialog box, do the following:

   a. Select the **OVA Source**. As an example, select *OVA File* to upload from a local folder.

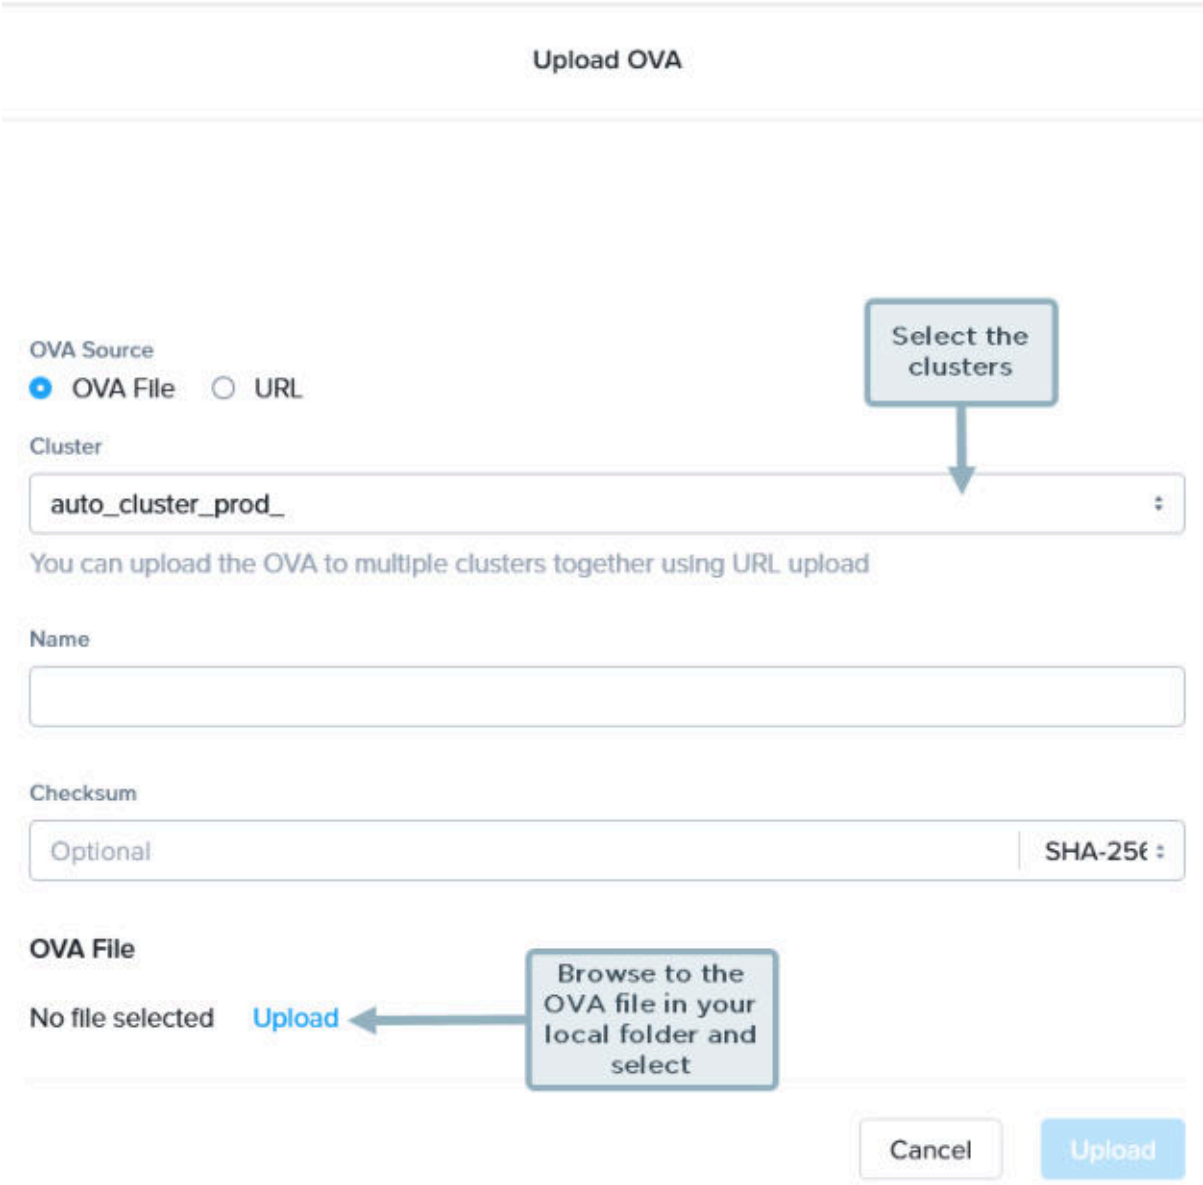

Figure 279: Upload OVA dialog box

   b. One cluster is already selected in the **Cluster** field.

      If you choose *URL* as **OVA Source** then you can select multiple clusters from the **Cluster** dropdown list that is displayed.

   c. Enter a **Name** for the OVA you are uploading.

      By default, Prism uses the file name of the OVA file that you upload, in the **Name** field if you do not enter a name and directly select the file for upload.

   d. Optionally, enter the **Checksum** value of the OVA file after selecting the encryption as SHA-1 or SHA-256 hash algorithm.

e. For **OVA File**, click **Select File** to navigate to the location of the OVA file in your local folder and open it.

   In case of upload from URL, in the **OVA URL** field, enter the source URL from where you want to upload the OVA file.

Figure 280: OVA URL field

> Note:
>
> Ensure that the URL you provide is not redirected to another URL. The upload fails for redirected URLs.

f. Click **Upload**.

3. The OVA upload starts. Wait until the **Continue in Background** button is displayed.

Note: The **Continue in Background** button is not displayed for uploads from URL.

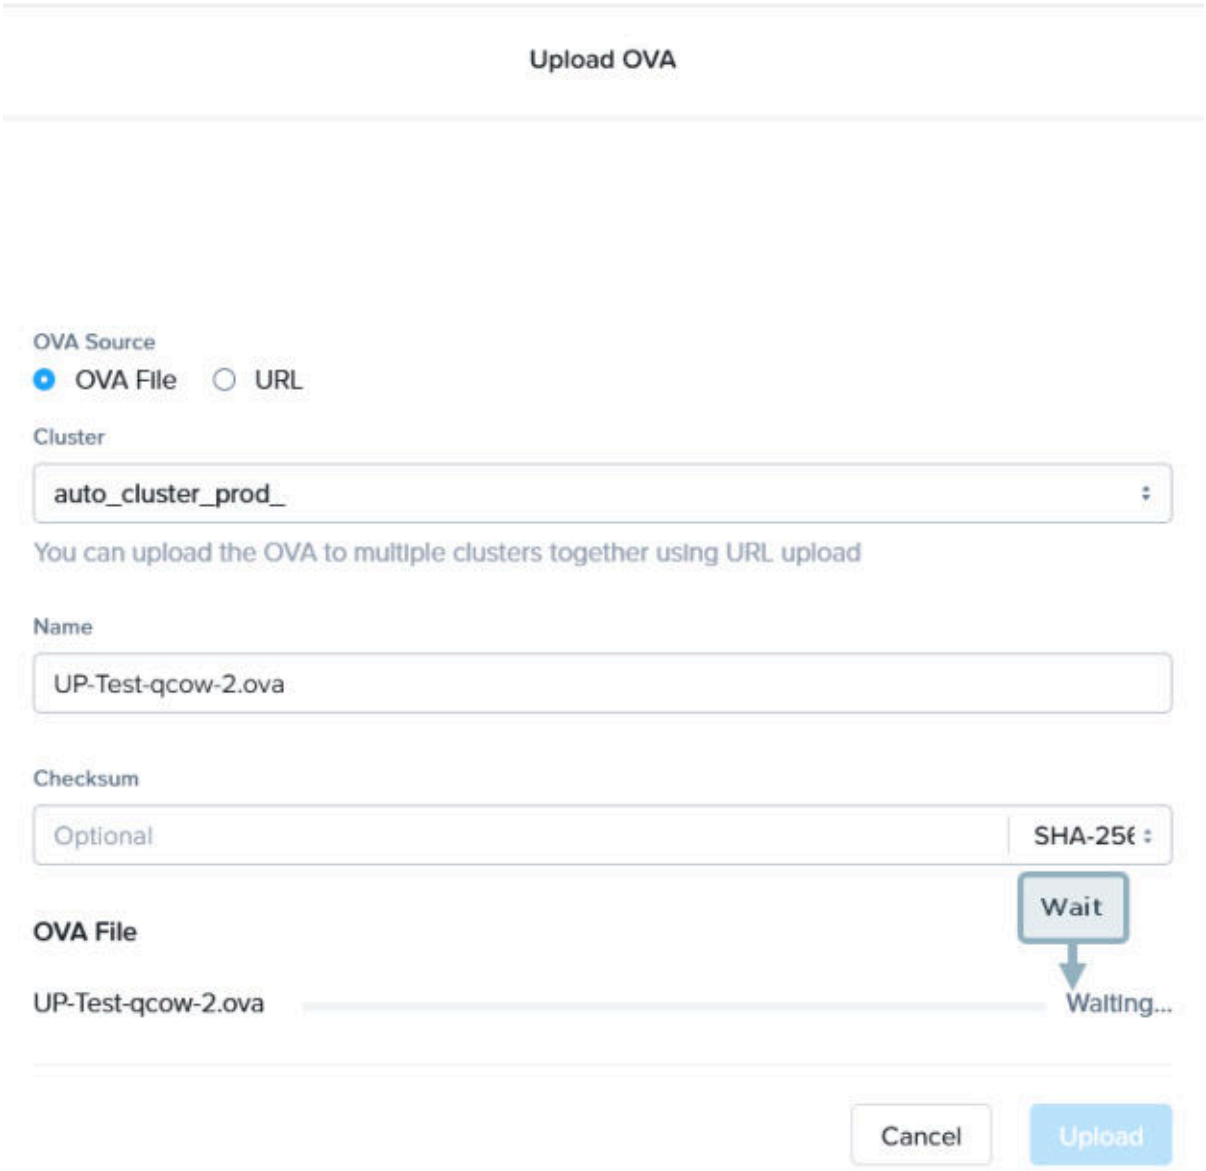

Figure 281: Upload OVA - Waiting

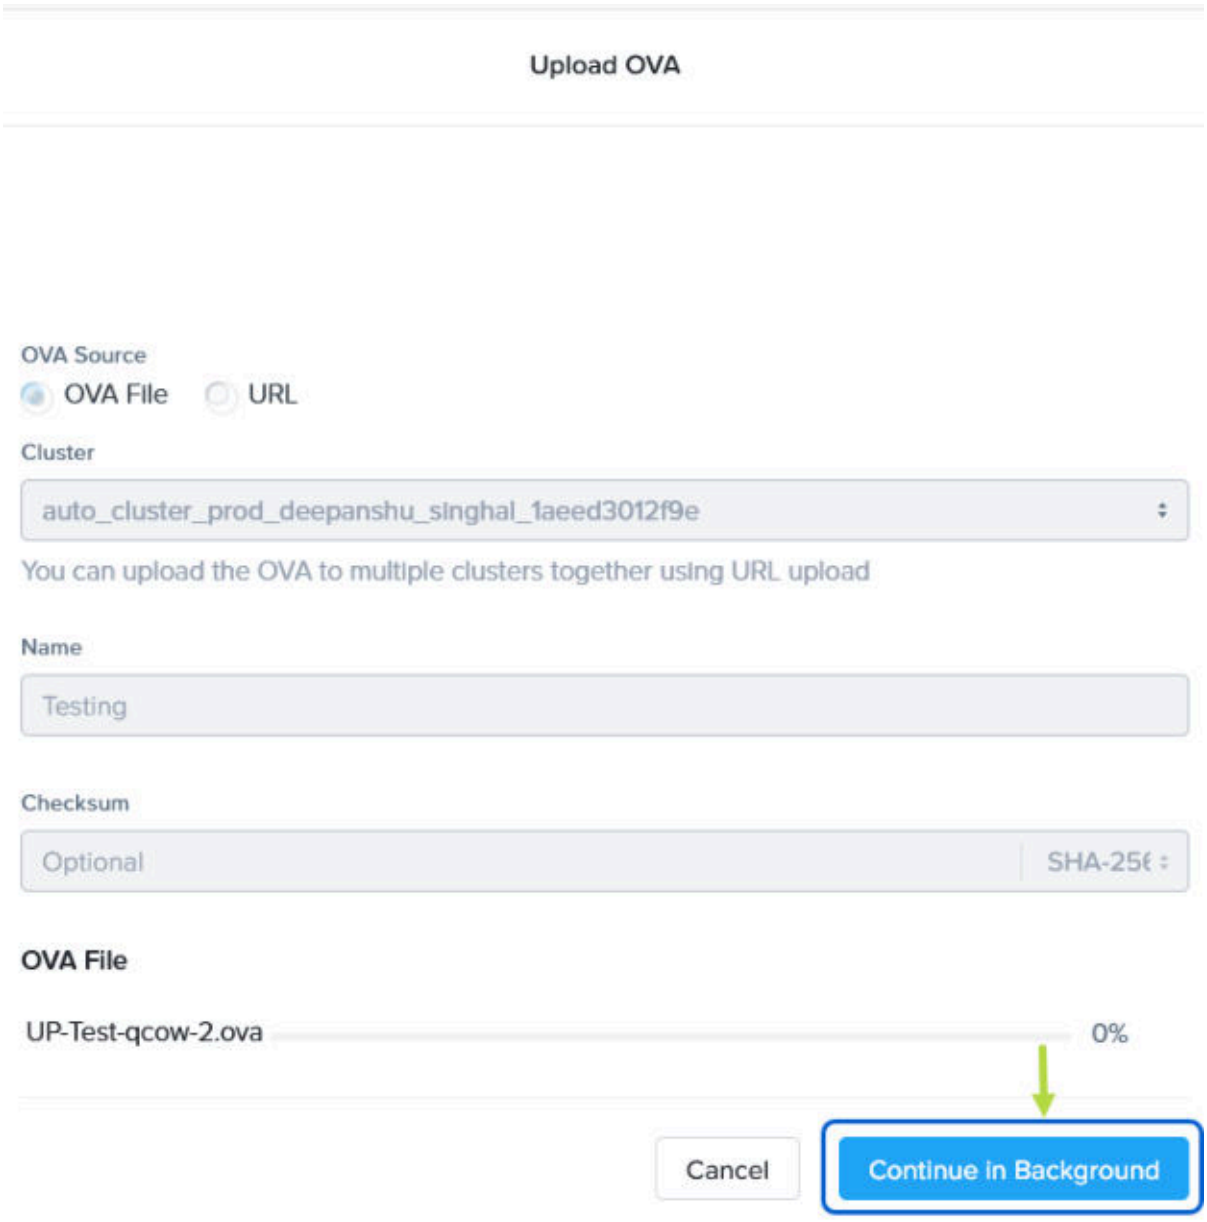

Figure 282: Continue in Background button

4. When the upload task is successfully completed, the following message is displayed:

   Upload of the file is successfully completed. Verification of file is progressing in background. Please track it from Tasks page.

5. To resume an upload process that was interrupted, select the OVA in the OVAs dashboard. In the **Actions** menu and click the **Resume Upload** option.

   The status message that appears when you hover the mouse over the status icons for each OVA indicates whether you can use the **Resume Upload** option to resume the upload.

   Note:

   Check the reason for failure of **OVA Validate** task in the Tasks page. When the upload process is interrupted because of a concatenation and validation failure due

an incorrect checksum, the **Resume Upload** option is not available. Upload the OVA again in case the upload process was interrupted due to a checksum error.

Concatenating Upload using APIs

About this task

During an OVA upload, the OVA file is chunked and the chunks are uploaded. After upload, during the validation phase, the chunks are concatenated and validated.

When the OVA upload process is interrupted because of a concatenation error, you need to run the concatenation API to resume the upload process.

To run the API successfully, you need the UUID of the OVA for which the upload is interrupted. The UUID is derived using the name of the OVA.

Procedure

1. Click the user icon in the upper-right corner of the web console, and click **REST API Explorer**.

   The REST API Explorer displays a list of the objects that can be managed by the API.

2. In the REST API Explorer browser tab, select **ovas** in the list and click **List Operations**.

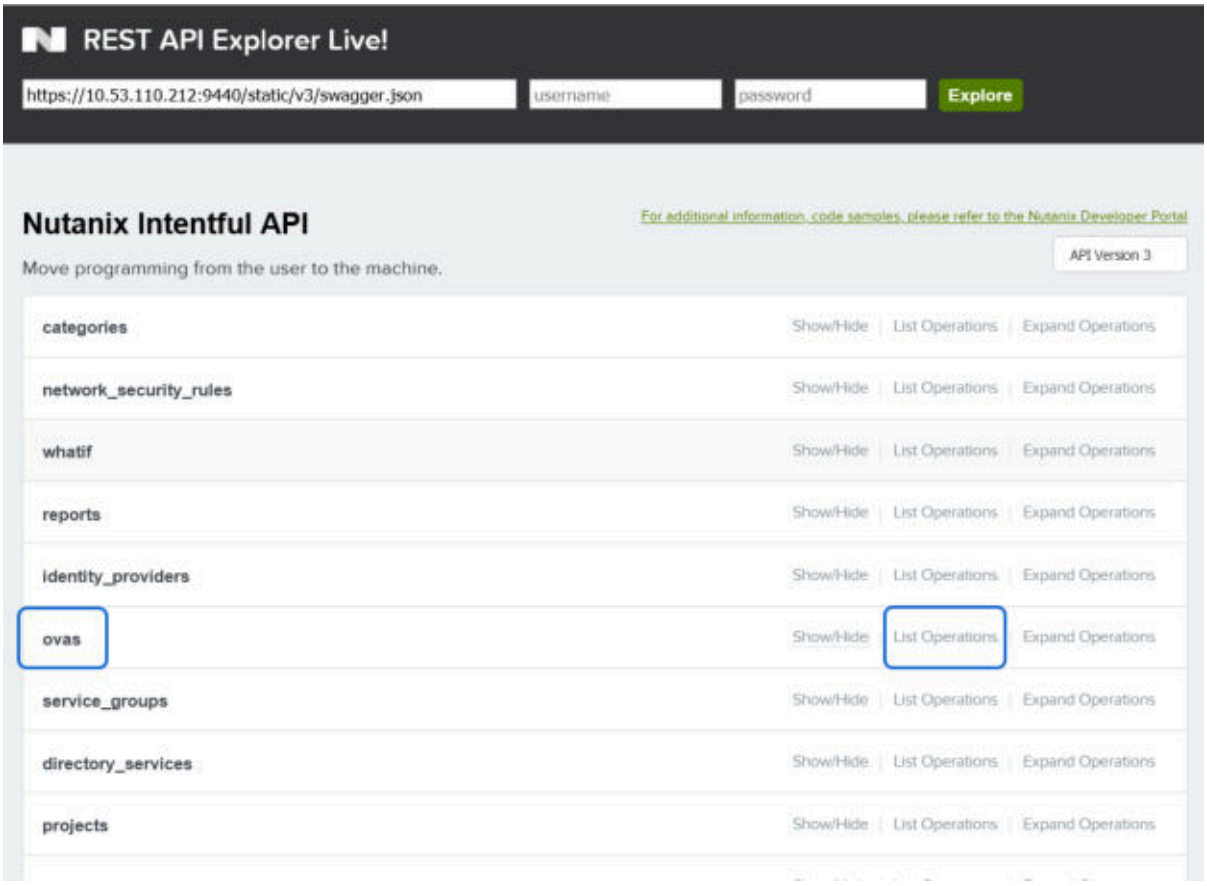

Figure 283: ovas List Operations

3. In the expanded **ovas** list of operations, click **Post /ova/list**.

   To use the concatenate API, you need the UUID of the OVA. **Post /ova/list** helps you filter the OVAs and get the details of the details of the interrupted OVA.

4. Enter the name of the OVA in the **get_entities_request** parameter field in the following format and click **Try it out!**:

```
{
 “filter”:”name==<name of ova>”
}
```

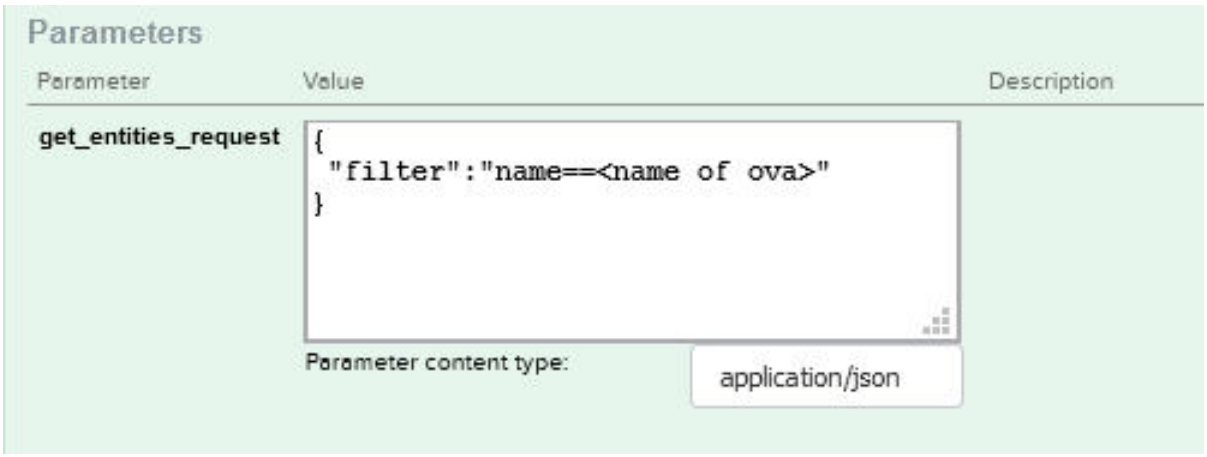

Figure 284: Parameter Filter

5. In the **Response Body** box, the UUID of the queried OVA is listed. Note or copy the UUID.

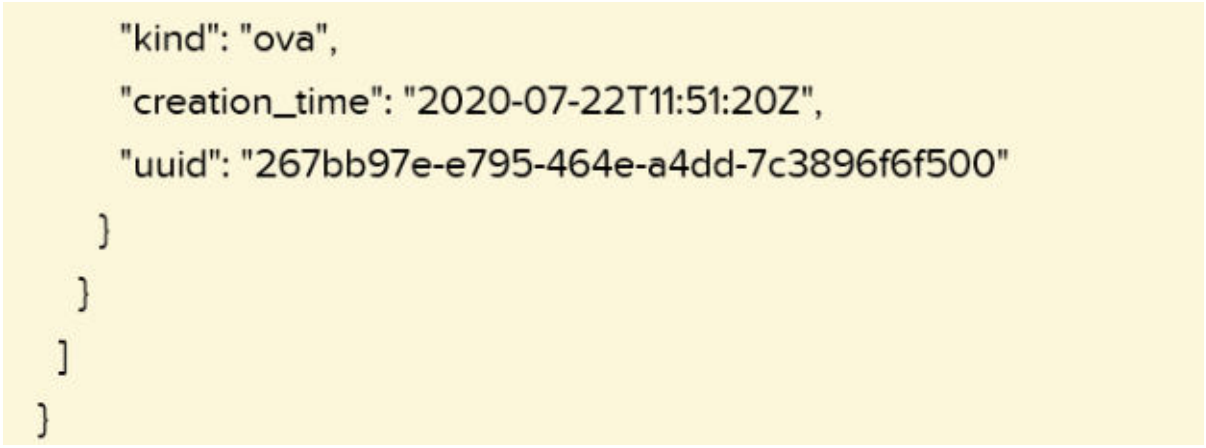

Figure 285: OVA Upload UUID

6. In the expanded **ovas** list of operations, click **Post /ova/{uuid}/chunks/concatenate** and enter the UUID of the OVA in the UUID parameter value field.

7. Click **Try it out!**.

   The **Response Code** field displays 202 indicating that the concatenation and validation request has been accepted.

Deploying an OVA as VM

About this task

To perform this task in the OVAs dashboard, select the target OVA. Some of the steps that you need to complete in this procedure are the same as the steps provided in Creating a VM (AHV) on page 317. The Deploy as VM dialog box has four tabs. Ensure that information in the first three tabs are appropriate for the VM you want to deploy. The Review tab helps you review the information in the first three tabs in one tab. You can navigate back and forth between the tabs using the **Back** and **Next** buttons.

> Note:
>
> In case of OVAs that are not created from AHV, ensure that the OVA has pre-installed VirtIO drivers.

To deploy an OVA as a VM, do the following:

Procedure

1. In the **Actions** menu, select **Deploy as VM**.

   You need to use the tabs in the Deploy as VM dialog box to configure the deployment. Use the **Back** and **Next** buttons to move between the tabs in the dialog box.

   > Note:
   >
   > The Actions menu is displayed only if you select an OVA in the list.

2. In the Deploy as VM dialog box, complete the indicated details in the Configuration tab:

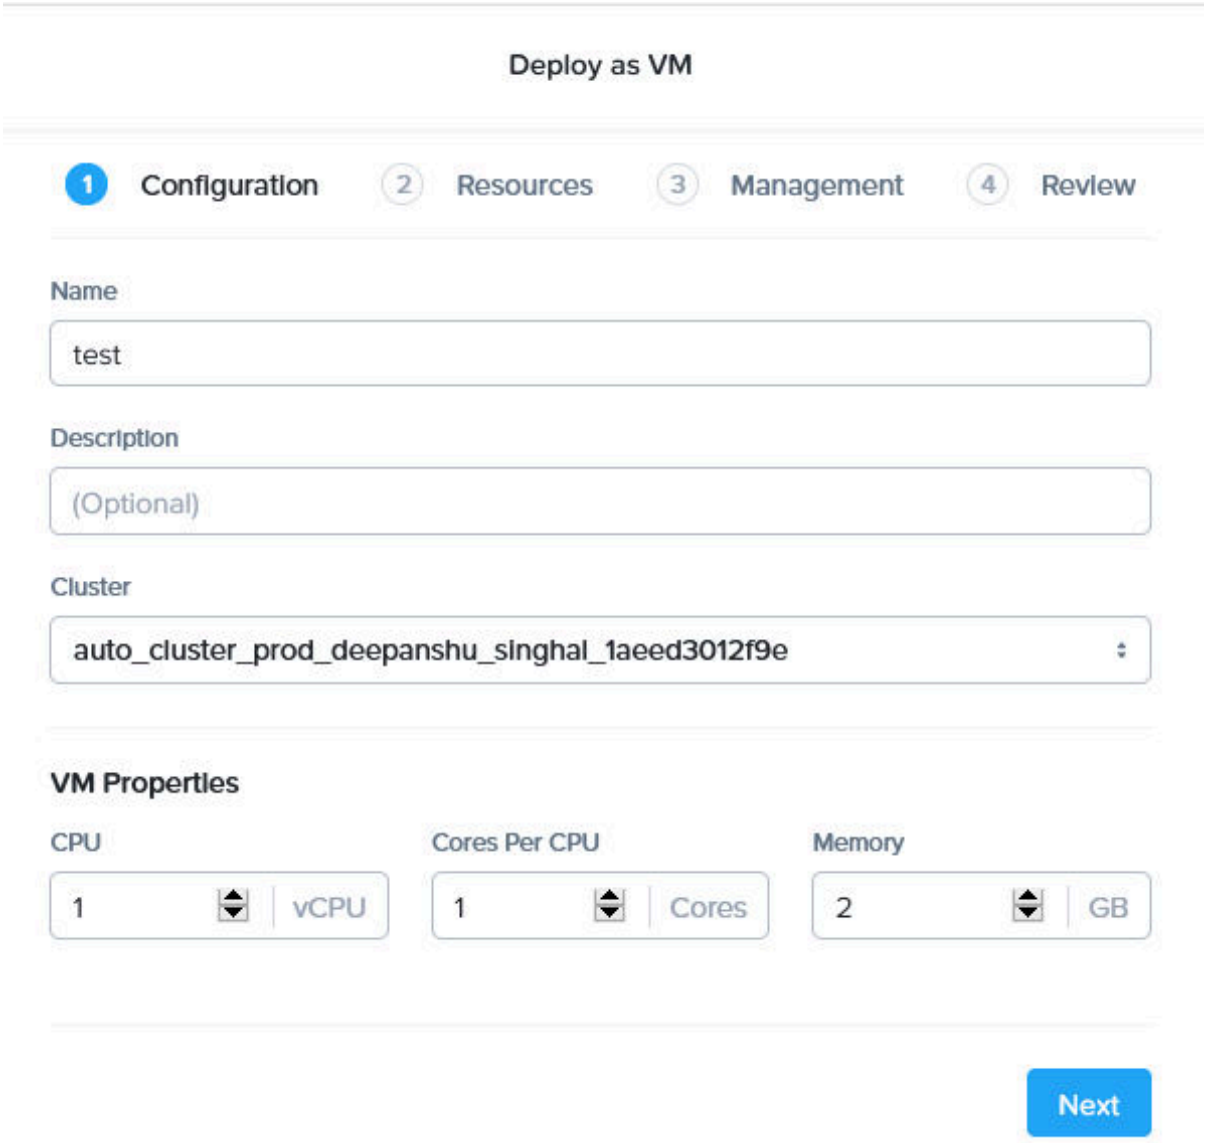

Figure 286: Deploy as VM dialog box

a. **Name** for the VM that you are deploying.

b. **Description** that are informative like Backup VM for Prism.

c. **Cluster** that you want to deploy the VM on. Select the cluster from the drop down list.

   Only clusters, where OVAs are available and you have access to, are displayed in the dropdown list. Prism Central also recommends a cluster in the list for the deployment.

d. The number of **CPU**, **Cores Per CPU** and **Memory** in **VM Properties**.

e. Select **Next**.

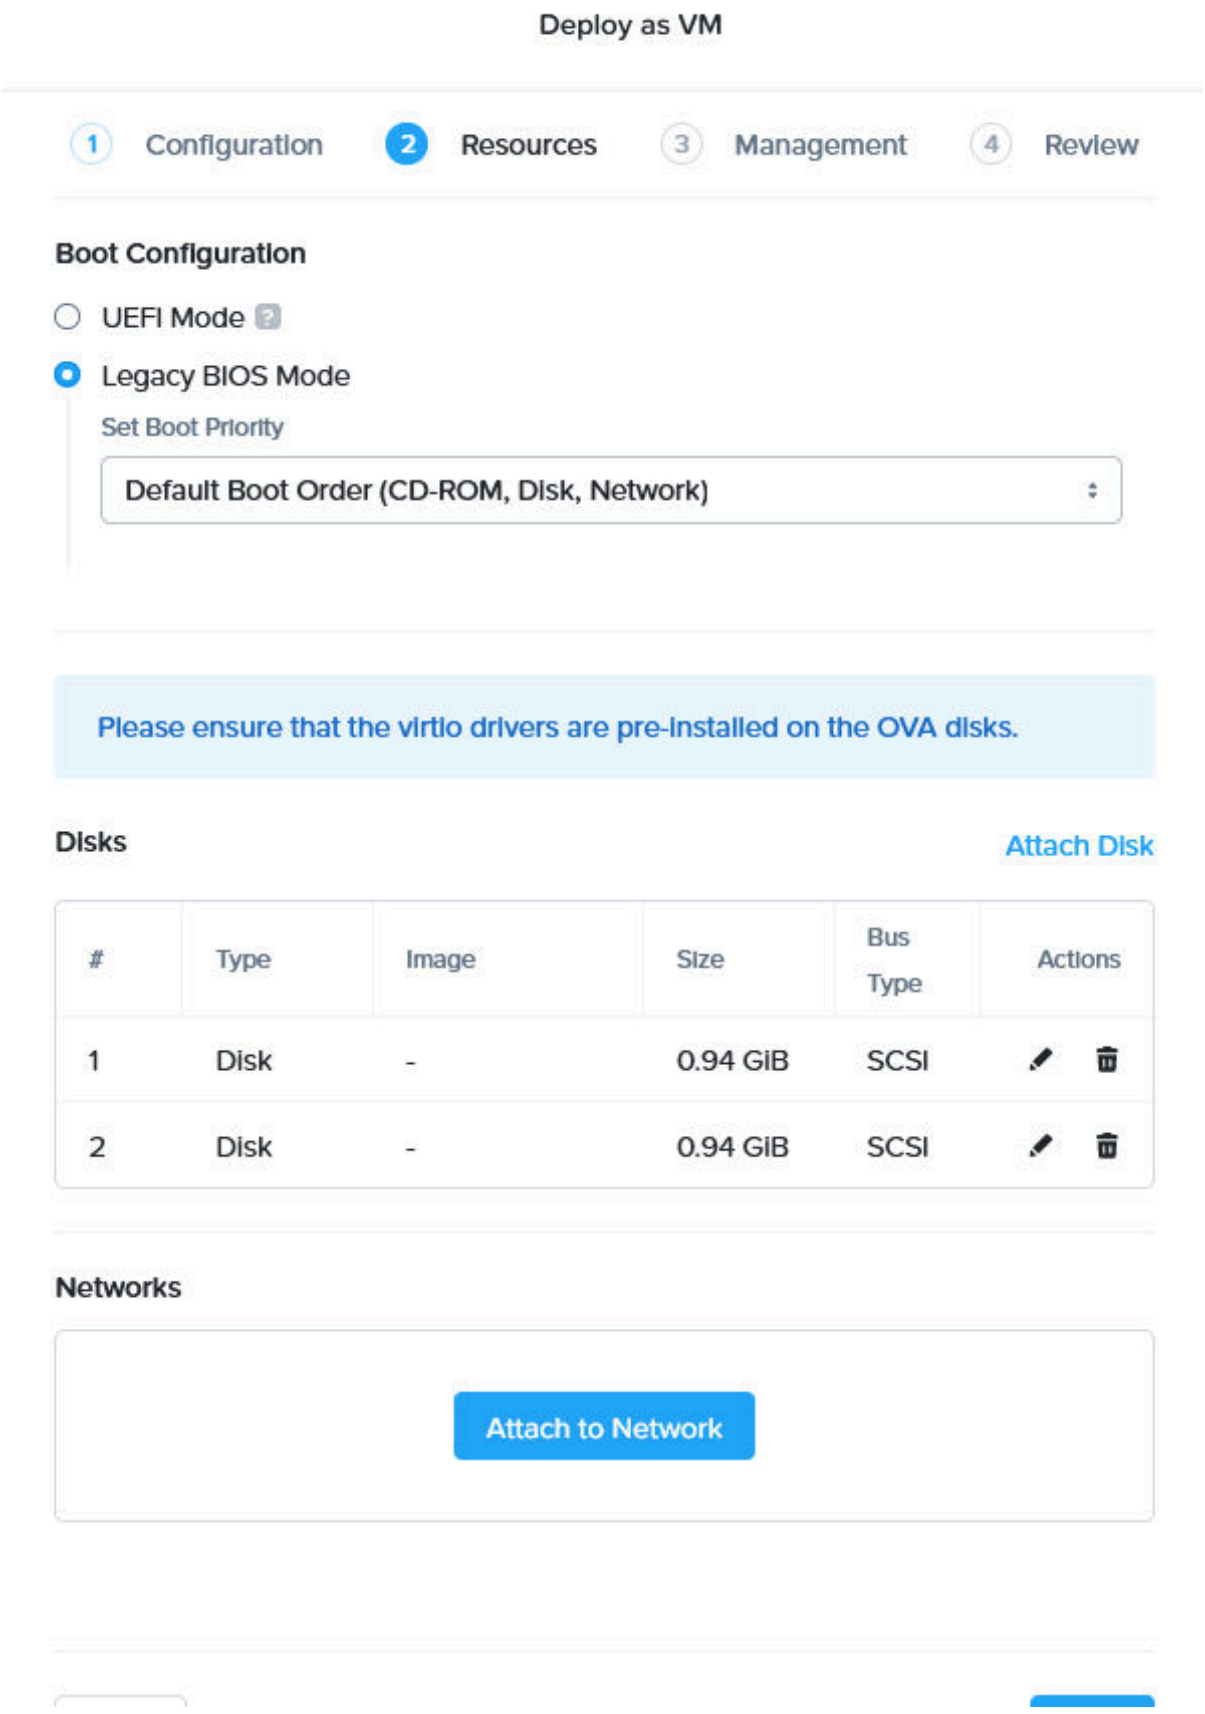

Figure 287: Deploy A VM - Resources tab

3. Complete the indicated details in the **Resources** tab.

   Note:

   You can add new disks, NICs and GPUs.

   a. **Boot Configuration** mode.
      Select **UEFI Mode** only if your hardware supports UEFI boot mode.

   b. **Set Boot Priority**.
      The default priority is already set.

   c. **Attach Disk** in addition to the disks already available in the OVA, if so required.

      In the Attach Disk dialog box, complete the indicated details and select **Save**.

      See Creating a VM (AHV) on page 317 for more information about attaching disks.

   d. Configure the network by selecting **Attach to Network**.

      In the Attach to Network dialog box, select the **Subnet** and **Network Connection State** from the drop down lists.

      See Creating a VM (AHV) on page 317 for more information about configuring networks.

   e. If the selected cluster has GPU resources, then you can add multiple vGPUs. Select **Add GPU**.

      See Creating a VM (AHV) on page 317 and Adding Multiple vGPUs to the Same VM on page 350 for more information about configuring networks..

      Note:

      You can add multiple vGPUs to the same VM. Addition of multiple vGPUs is based on the installed GPU resources. See *Multiple Virtual GPU Support*in *AHV Admin Guide*.

   f. Select **Next**.

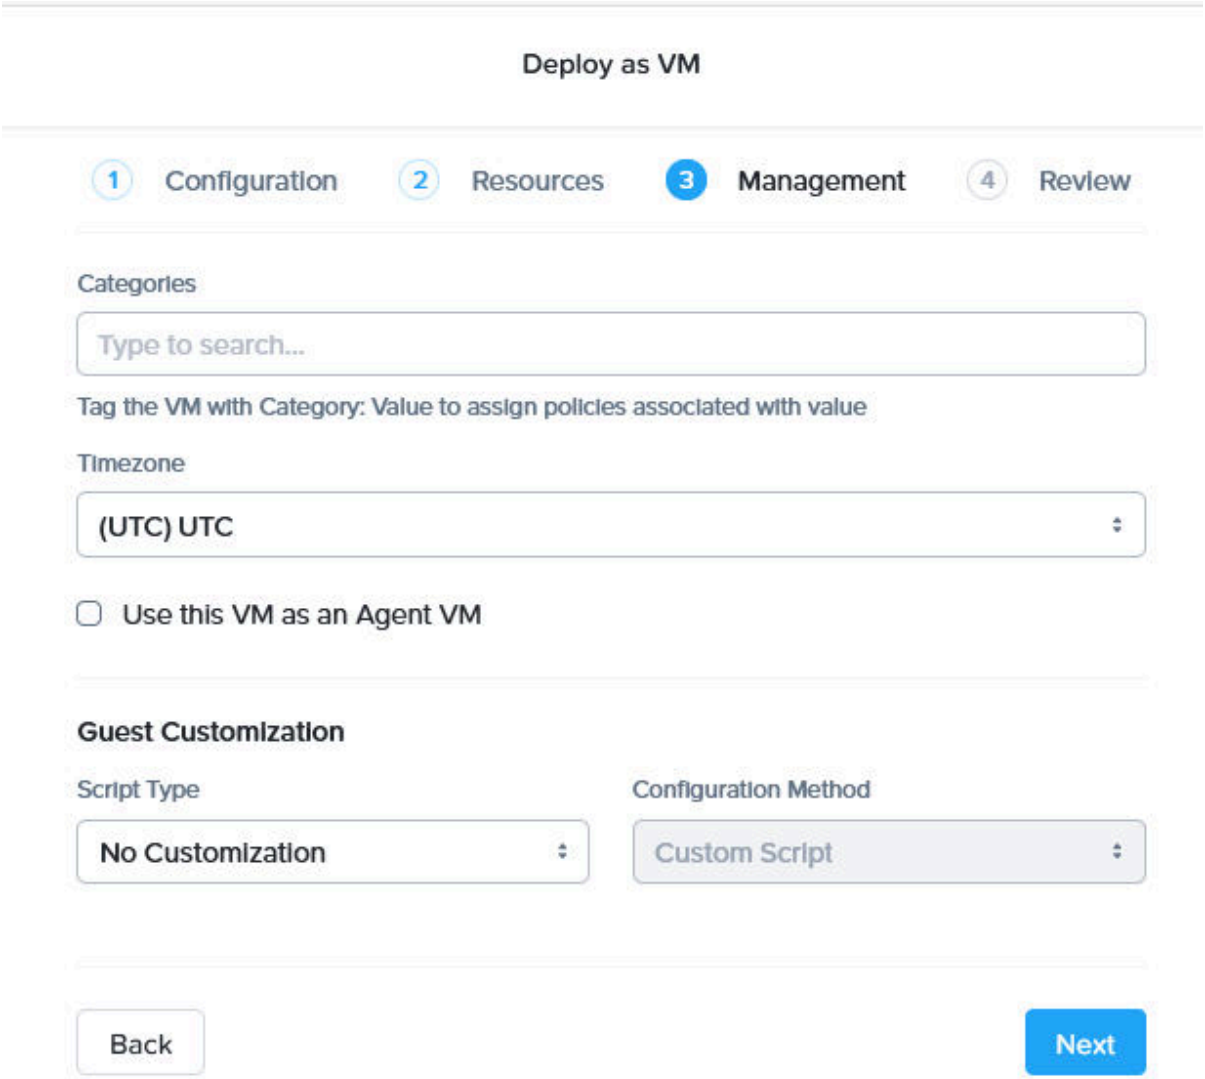

Figure 288: Deploy A VM - Management tab

g. Select the **Category**. Type in this field to display a list of available categories. For more information about categories, see Category Management on page 413.

h. Set the **Timezone**.

i. If you want to use the deployed VM as an agent VM, check the **Use this VM as an Agent VM** checkbox.

j. Configure **Guest Customizations** if you want to run scripts for configurations on Microsoft Windows or Linux VMs.

   You can choose from *No Customization*, *Sysprep (Windows)* and *Cloud-init (Linux)* as **Script Type**. You can choose from *Custom Script* or *Guided Script* as **Configuration Method**.

   If you choose *No Customization* as **Script Type**, then the **Configuration Method** field is grayed out.

k. Select **Next**.

Deploy as VM

1 Configuration 2 Resources 3 Management 4 Review

Configuration

Name	test
CPU	1 vCPU
Cores Per CPU	1 Cores
Memory	2 GB

Resources

Disks

#	Type	Image	Size	Bus Type
1	Disk	-	0.94 GiB	SCSI
2	Disk	-	0.94 GiB	SCSI

Boot Configuration: Default Boot Order (CD-ROM, Disk, Network)

Networks

Subnet	Private IP
pc_network	Auto-Assign

Management

Categories	None
Timezone	UTC
Guest Customization	No Customization

4. Review the deployment configuration in the Review tab.
5. Click **Launch** to deploy the VM.

   You can check the deployment tasks in Tasks. The deployed VM is displayed in the VMs dashboard.

Downloading an OVA

About this task

To perform this task in the OVAs dashboard, select the target OVA.

> Note:
> You can download upto five OVAs in parallel.

To download an OVA, do the following:

Procedure

1. In the **Actions** menu, select **Download**.

   > Note:
   > The Actions menu is displayed only if you select an OVA in the list.

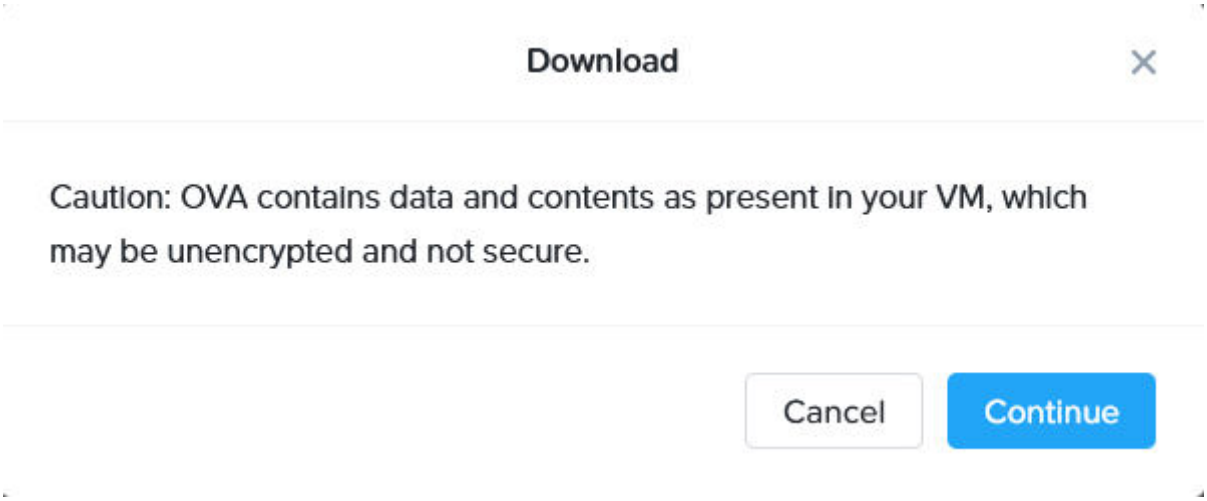

Figure 290: Download message box

2. To confirm the download action, select **Continue**.

   To cancel the download action, select **Cancel**.

   The downloaded file is saved in a local folder based on the download settings of your browser.

Renaming an OVA

About this task

To perform this task in the OVAs dashboard, select the target OVA.

To rename an OVA, do the following:

Procedure

1. In the **Actions** menu, select **Rename**.

   Note:
   The Actions menu is displayed only if you select an OVA in the list.

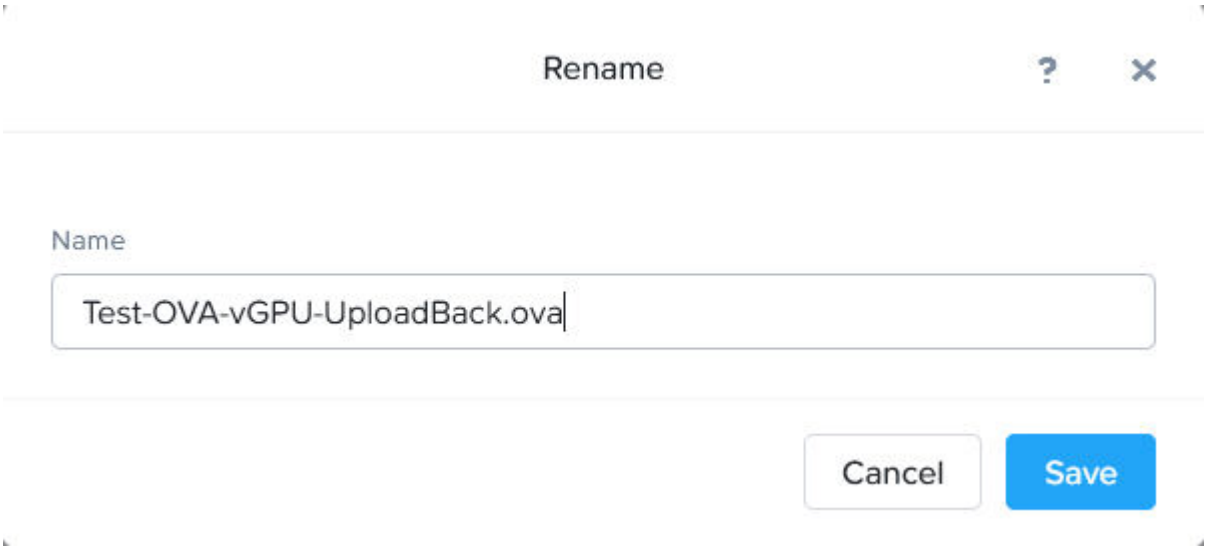

   Figure 291: Rename dialog box

2. In the **Name** field of the Rename dialog box, delete the current name and enter the new name for the OVA.
3. To confirm the rename action, select **Save**.
   To cancel the delete action, select **Cancel**.

Deleting an OVAs

About this task

To perform this task in the OVAs dashboard, select the target OVA.

To delete an OVA, do the following:

Procedure

1. In the **Actions** menu, select **Delete**.

   Note:
   The Actions menu is displayed only if you select an OVA in the list.

   The **Delete** action deletes the OVA permanently. You cannot recover a deleted OVA.

   Therefore, in the Delete dialog box, confirm if you want to delete the OVA <name>.

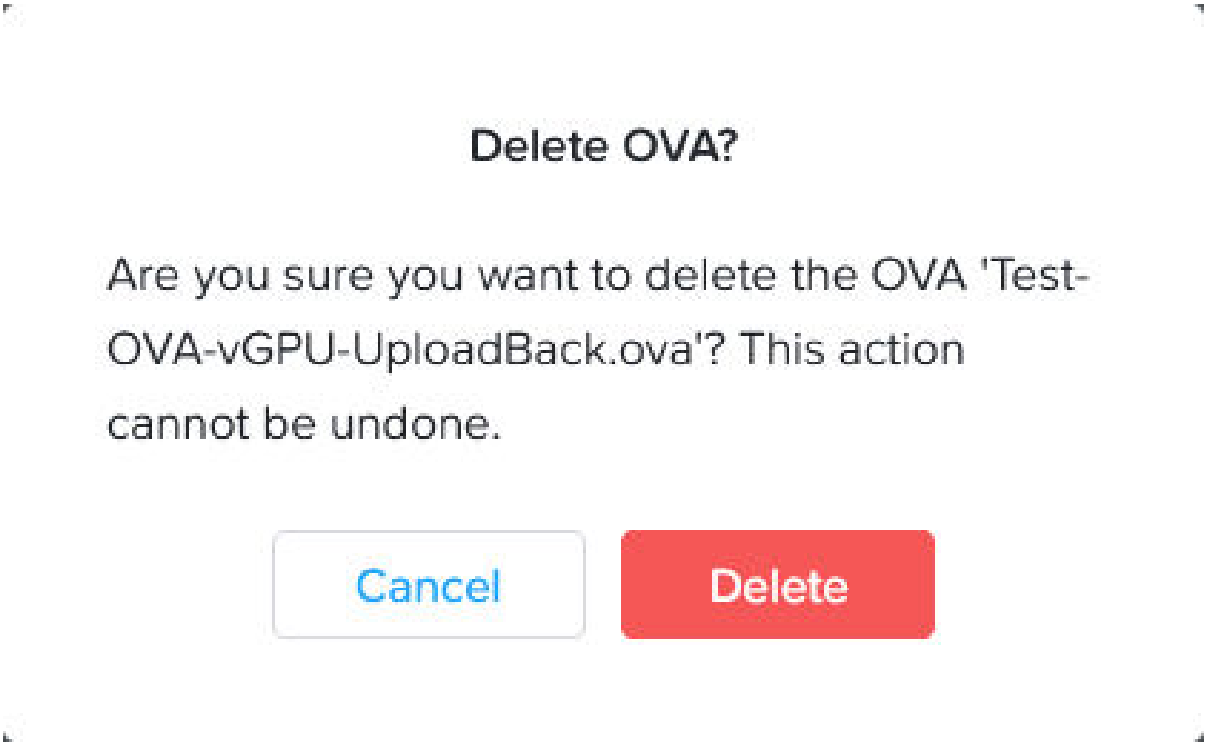

   Figure 292: Delete OVA dialog box

2. To confirm the delete action, select **Delete**.
   To cancel the delete action, select **Cancel**.

Affinity Policies for AHV

As an administrator in an Acropolis managed clusters, you can specify scheduling policies for virtual machines on an AHV cluster. By defining these policies, you can control placement of the virtual machines on the hosts within a cluster.

You can define two types of affinity policies.

VM-Host Affinity Policy

The VM-host affinity policy controls the placement of the VMs. You can use this policy to specify that a selected VM can only run on the members of the affinity host list. This policy checks and enforces where a VM can be hosted when you power on or migrate the VM.

Note:

- If you choose to apply the VM-host affinity policy, it limits Acropolis HA and Acropolis Dynamic Scheduling (ADS) in such a way that a virtual machine cannot be powered on or migrated to a host that does not conform to requirements of the affinity policy as this policy is enforced mandatorily.
- The VM-host anti-affinity policy is not supported.

You can define the VM-host affinity policies by using Prism Central during the VM create or update operation. For more information, see Creating a VM (AHV) on page 317.

Note: If a VM is cloned that has the affinity policies configured, then the policies are not automatically applied to the cloned VM. However, if a VM is restored from a DR snapshot, the policies are automatically applied to the VM.

Limitations of Affinity Rules

- Even if a host is removed from a cluster, the host UUID is not removed from the host-affinity list for a VM.
- The VM-host affinity cannot be configured on a cluster that has HA configured by using reserved host method.

Performing Other Administrative Tasks

About this task

You can perform some administrative tasks for a selected cluster directly from Prism Central (see Virtual Infrastructure (Cluster) Administration on page 316). Other tasks can only be performed through Prism Element, which is the web console interface for an individual cluster. To access Prism Element from Prism Central, do the following:

Procedure

1. Go to the clusters dashboard (see Clusters Summary View on page 146).
2. Do one of the following:
   - » Check the box for the target cluster and then select **Launch Prism Element** from the **Action** pull-down menu (located above the list of clusters). The **Action** menu appears only after selecting the target cluster.
   - » Click the target cluster name to display the details page, and then click the **Launch Prism Element** button (top right of screen).

   The Prism Element web console for that cluster appears in a new tab or window. You can now perform any administrative task for that cluster. See the Prism Web Console Guide for more information.

What to do next

When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The following table describes the user icon options when accessing a cluster from Prism Central.

Table 221: User Menu List

Name	Description
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page.
About Nutanix	Opens a window that displays the AOS version running on the cluster.
Support Portal	Opens a new browser tab (or window) at the Nutanix support portal login page.
Help	Opens the online help.
Nutanix Next Community	Opens a new browser tab (or window) at the Nutanix Next Community entry page.
Sign Out	Logs out the user from the web console.

vCenter Server Integration

The VM management through Prism Central for ESXi provides a unified management interface for all of the vCenter Server instances. All the vCenter Server instances that are registered or not registered with any of the clusters are auto-discovered and displayed. If you have not registered the vCenter Server in Prism Element, you can register the clusters to the vCenter Server instances directly from Prism Central.

During the registration process, you have an option to select the ESXi clusters that you want to manage using Prism Central. For more information about registering vCenter Server, see Registering vCenter Server (Prism Central) on page 453. After you successfully register vCenter Server instances, you can perform the following operations directly from Prism Central.

- Create, clone, update, and delete VMs.
- Create and delete NICs.
- Attach and delete disks.
- Power operations: Power on or off, reset, suspend, resume.
- Open and launch VM console.
- Enable and disable NGT.

  Note: Managing VMware guest tools is not supported through Prism Central.

Note:

- You can perform the power operations and launching of VM console even when vCenter Server is not registered.
- If you are creating VM through Prism, configuration changes to the VM when it is powered on is enabled by default and it depends on the guest operating system that is deployed on the VM.

Rules and Guidelines

- Ensure that all the hosts in the cluster are managed by a single vCenter Server.

- Ensure that DRS is enabled on all the vCenter Server instances.
- Ensure that you are running ESXi and vCenter Server 5.5 or later releases.
- Ensure that you have homogeneous network configuration for all the vCenter Server instances. For example, network should have either 1G or 10G NICs.
- Ensure that you unregister the vCenter Server from the cluster before changing the IP address of the vCenter Server. After you change the IP address of the vCenter Sever, you must register the vCenter Server again with the new IP address.
- The vCenter Server Registration page displays the registered vCenter Server. If for some reason the Host Connection field changes to **Not Connected**, it implies that the hosts are being managed by a different vCenter Server. In this case, there will be new vCenter entry with host connection status as Connected and you need to register to this vCenter Server. For more information about registering vCenter Server again, see Managing vCenter Server Registration Changes (Prism Central) on page 456.

  CAUTION: If multiple vCenter Servers are managing the hosts of a single Nutanix cluster, you will not be able to perform the VM management operations. Move all the hosts into one vCenter Server.

Requirements and Limitations

- The E1000, E1000e, PCnet32, VMXNET, VMXNET 2, VMXNET 3 network adapter types (NICs) are supported.
- Only SCSI and IDE disks are supported. SATA and PCI disks are not supported.
- Creating a VM by using a template is not supported.
- Creating a VM by using image service is not supported.
- If a VM is deleted, all the disks that are attached to the VM get deleted.
- Network configuration (creation of port groups or VLANs) is not supported.

Registering vCenter Server (Prism Central)

All the vCenter Server instances that are registered to Prism Element are listed in Prism Central. If you do not want to manage your cluster through Prism Central, you have an option to de-select the clusters from Prism Central.

Before you begin

Ensure that you have vCenter Server Extension privileges as these privileges provide permissions to perform vCenter registration for the Nutanix cluster.

About this task

Note the following points about registering vCenter Server.

- Nutanix does not store vCenter Server credentials.
- Whenever a new node is added to a cluster, vCenter Sever registration for the new node is automatically performed.

Procedure

1. Click the gear icon in the main menu and select **vCenter Registration** from the Settings menu (see Settings Menu on page 17).

   - The **vCenter Registration** window lists all the vCenter Server instances that are registered or not registered with the clusters.
   - If you have not registered the vCenter Server in Prism Element, you can register the clusters to the vCenter Server instances directly from Prism Central.
   - You also have an option to de-select any ESXi clusters that you do not want to manage using Prism Central.

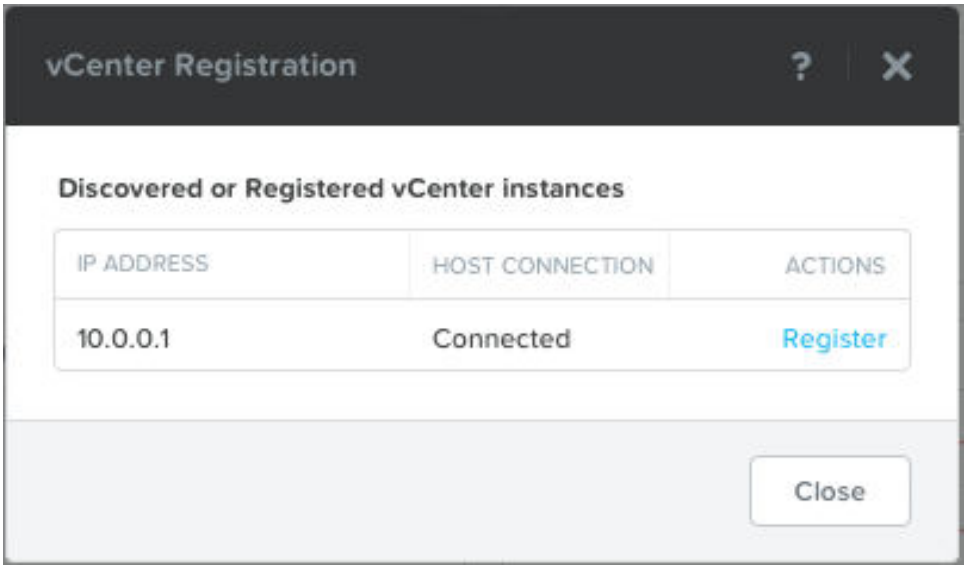

Figure 293: vCenter Registration

2. To register the vCenter Server, click **Register**.

   - The vCenter Server that is managing the hosts in the cluster is auto-discovered and its IP address is auto-populated in the **Address** field.
   - The port number field is also auto-populated with 443. Do not change the port number. For the complete list of required ports, see Port Reference.

3. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.

4. If you do not wish to manage all your ESXi clusters from Prism Central, select the check box of just the clusters that you want to manage and click **Save**.

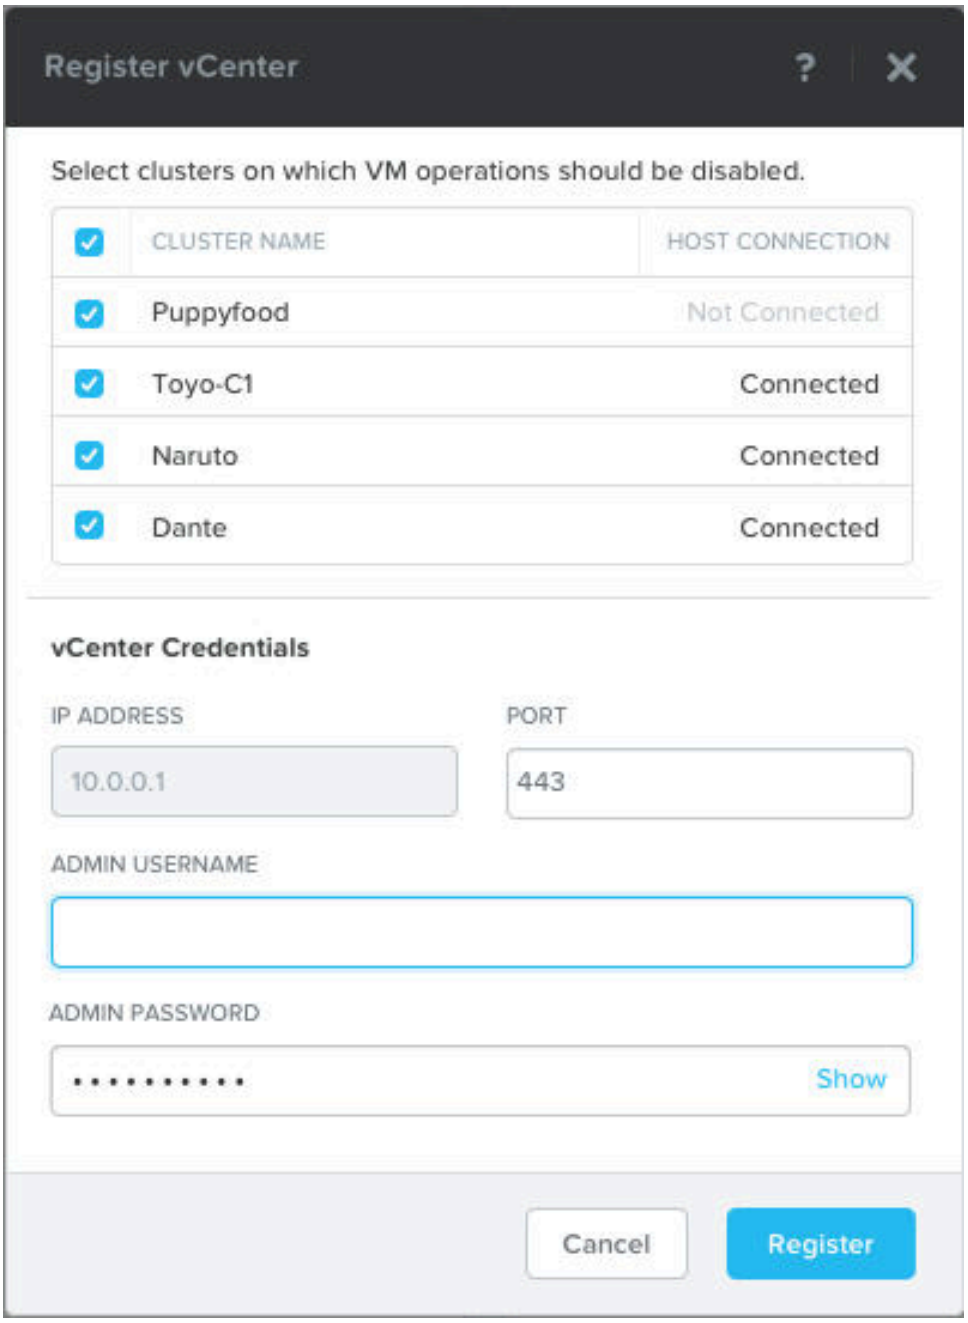

Figure 294: Cluster Selection

During the registration process a certificate is generated to communicate with the vCenter Server. If the registration is successful, a relevant message is displayed in the **Tasks** dashboard. The Host Connection field displays as Connected, which implies that all the hosts are being managed by the vCenter Server that is registered.

5. Click **Register**.

Unregistering a Cluster from the vCenter Server (Prism Central)

To unregister your clusters from the vCenter Server, perform the following procedure.

About this task

- Ensure that you unregister the vCenter Server from the cluster before changing the IP address of the vCenter Server. After you change the IP address of the vCenter Sever, you should register the vCenter Server again with the new IP address with the cluster.
- The vCenter Server Registration page displays the registered vCenter Server. If for some reason the Host Connection field changes to **Not Connected**, it implies that the hosts are being managed by a different vCenter Server. In this case, there will be new vCenter entry with host connection status as Connected and you need to register to this vCenter Server. For more information about registering vCenter Server again, see Managing vCenter Server Registration Changes (Prism Central) on page 456.

Procedure

1. Click the gear icon in the main menu and select **vCenter Registration** from the Settings menu (see Settings Menu on page 17).

   The list of vCenter Server instances already registered to one or more vCenter Server instances are displayed.

2. To unregister a cluster from a particular vCenter Server, click the edit icon under **Action** field.
3. Clear the check box of the clusters that you want to unregister from the vCenter Server.
4. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.
5. Click **Unregister**.

   If the credentials are correct, the vCenter Server is unregistered from the cluster and a relevant message is displayed in the **Tasks** dashboard.

Managing vCenter Server Registration Changes (Prism Central)

After you change the IP address of the vCenter Sever or if the Host Connection field changes to **Not Connected** (may occur if current vCenter Server is not managing the hosts), you must register your cluster with the new vCenter Server. Perform the following procedure to re-register the vCenter Server.

Procedure

1. Click the gear icon in the main menu and select **vCenter Registration** from the Settings menu (see Settings Menu on page 17).

   The vCenter Server that is managing the hosts in the cluster is displayed along with the vCenter Server that you must unregister.

2. Unregister the existing vCenter Server instance by clicking **Unregister**.

   For more information about unregistering vCenter Server, click Unregistering a Cluster from the vCenter Server (Prism Central) on page 456.

3. Register the new vCenter Server instance by clicking **Register**.

   For more information about registering vCenter Server, click Registering vCenter Server (Prism Central) on page 453.

9

PRISM SELF SERVICE ADMINISTRATION

The Prism Self Service feature allows you to create projects where consumers of IT infrastructure within an enterprise—individual users or teams such as development, test, and DevOps—can provision and manage VMs in a self-service manner, without having to engage IT in day-to-day operations.

> Note: This chapter describes how to configure and administer Prism Self Service. See the *Prism Self Service Administration Guide* for instructions on the role and functions of a self-service administrator. See the *Prism Self Service User Guide* for information about the self-service capabilities available to a project user.

Prism Self Service Overview

Prism Self Service represents a special view within Prism Central. While Prism Central enables infrastructure management across clusters, Prism Self Service allows end users to consume that infrastructure in a self-service manner. Prism Self Service uses the resources provided by a single AHV cluster. (Other hypervisors are not supported platforms for Prism Self Service.)

> Note: A SelfServiceContainer storage container (see Storage Containers Summary View on page 97) is created on the target cluster and used by Prism Self Service for storage and other feature operations. To ensure proper operation of these features, do not delete this storage container.

There are three roles to consider when configuring Prism Self Service:

1. Prism Central administrator. The Prism Central administrator adds an Active Directory that includes the pool of self-service users and (optionally) creates one or more self-service administrators. Prism Central administrators also create VMs, images, and network configurations that may be consumed by self-service users.
2. Self-service administrator. The self-service administrator performs the following tasks:

   - Creates a project for each team that needs self service and adds users and groups to the projects.
   - Configures roles for project members. A project member can access only the entities or perform only the actions defined in the role assigned to that project member.
   - Publishes VM templates and images to the catalog.
   - Monitors resource usage by various projects and its VMs and members, and then adjusts resource quotas as necessary.

   It is optional to configure a separate self-service administrator because a Prism Central administrator can perform any of these tasks. However, if you would like to authorize users to administer end-user VM or application provisioning, you can use this role to give them access to virtual infrastructure without giving them access to physical infrastructure.

   > CAUTION: Self-service administrators have full access to all VMs running on the Nutanix cluster, including infrastructure VMs not tied to a project. Self-service administrators can assign infrastructure VMs to project members, add them to the catalog, and delete them

even if they do not have administrative access to Prism Central. Consider these privileges when appointing self-service administrators, and make sure to communicate to self-service administrators the need to exercise caution when working with infrastructure VMs.

After a Prism Central administrator has designated a user as a self-service administrator, the Prism Central administrator cannot limit the user's privileges. Therefore, if you plan to delegate self-service administration responsibilities to an Active Directory group, be sure that you want to delegate the responsibility to all the users in the group. If the user group is large or includes users that must not have self-service administrator privileges, Nutanix recommends that you create a separate Active Directory group for the users to whom you want to delegate self-service administration responsibilities.

3. Project user. These are the users assigned to a project by a self-service administrator. They can perform any action that the self-service administrator grants them. The permissions are determined by the roles assigned to the users and groups in the project. When project users log in, they see a custom self-service GUI interface that shows only what the role permissions allow. Project users create and manage only what they need.

To configure Prism Self Service, perform the following tasks:

- Specify an Active Directory for self service and (optionally) add one or more self-service administrators (see Configuring Prism Self Service on page 458).
- Create a project for each team that needs self service and add users and groups to the projects (see Project Management on page 462).
- Configure roles for project members (see “” ()).

  Note: RBAC is an independent feature, so you can configure RBAC without configuring self service or creating projects. However, projects allow you to enforce RBAC in a more granular way.

- Publish VM templates and images to the catalog (see Catalog Management on page 390).
- Monitor resource usage for the projects and adjust resource quotas as needed. You can view resource usage across and within projects through the **Explore** tab (see Projects Summary View on page 203 and Project Details View on page 205).
- Create VMs as needed and assign them to project members as appropriate (see VM Management on page 316). You can also allow project members to create their own VMs.

Configuring Prism Self Service

About this task

To configure Prism Self Service, do the following:

Note: Prism Self Service was hosted through Prism Element (the single-cluster Prism interface) prior to AOS 5.5. However, starting with AOS 5.5 Prism Self Service is hosted through Prism Central only. If your current Prism Self Service is on Prism Element, see Migrating Prism Self Service on page 461.

Procedure

1. Configure an active directory that includes the pool of self-service users if one is not already configured (see “” ()).

2. Click the gear icon in the main menu and select **Self-Service Admin Management** from the **Settings** menu.

   The Self-Service Admin Management window appears.

   > Note: The **Self-Service Admin Management** option does not appear in the **Settings** menu until you configure a directory service (step 1).

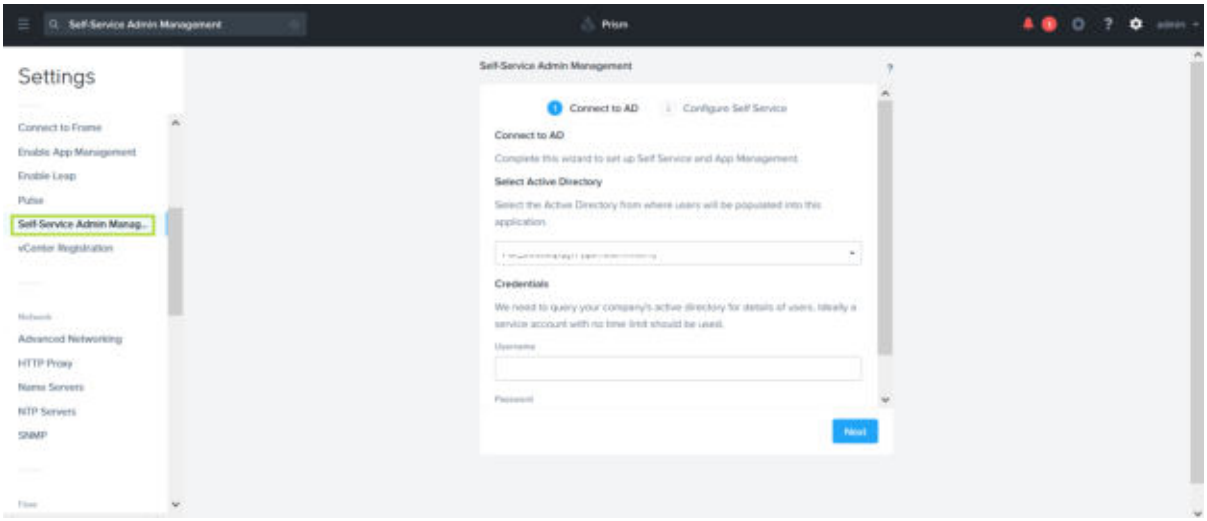

Figure 295: Self-Service Admin Management Window

3. Do the following in the **Connect to AD** tab:

   a. Select the desired Active Directory from the pull-down list.

   b. Enter the user name and password of an Active Directory user with administrator permissions.

      > Note: To prevent future access problems, it is recommended that you use an administrator account with no time limit.

   c. Click the **Next** button.

4. Do one or more of the following in the **Configure Self Service** tab and then click the **Save** button.

   - To add an administrator, click the **Add Admins** link. This displays a line to specify a self-service administrator. In the **Name** field, enter an individual or group name (in the directory service) that you want to give self-service administrator privileges and then

click **Save** in the **Actions** field. To add additional self-service administrators, click the **Add Admins** link and repeat this process.

- To change an existing self-service administrator, click the pencil icon in the **Actions** field for that user, update the user name in the **Name** field, and then click **Save** in the **Actions** field.
- To remove an existing self-service administrator, click the X icon in the **Actions** fields for that user.

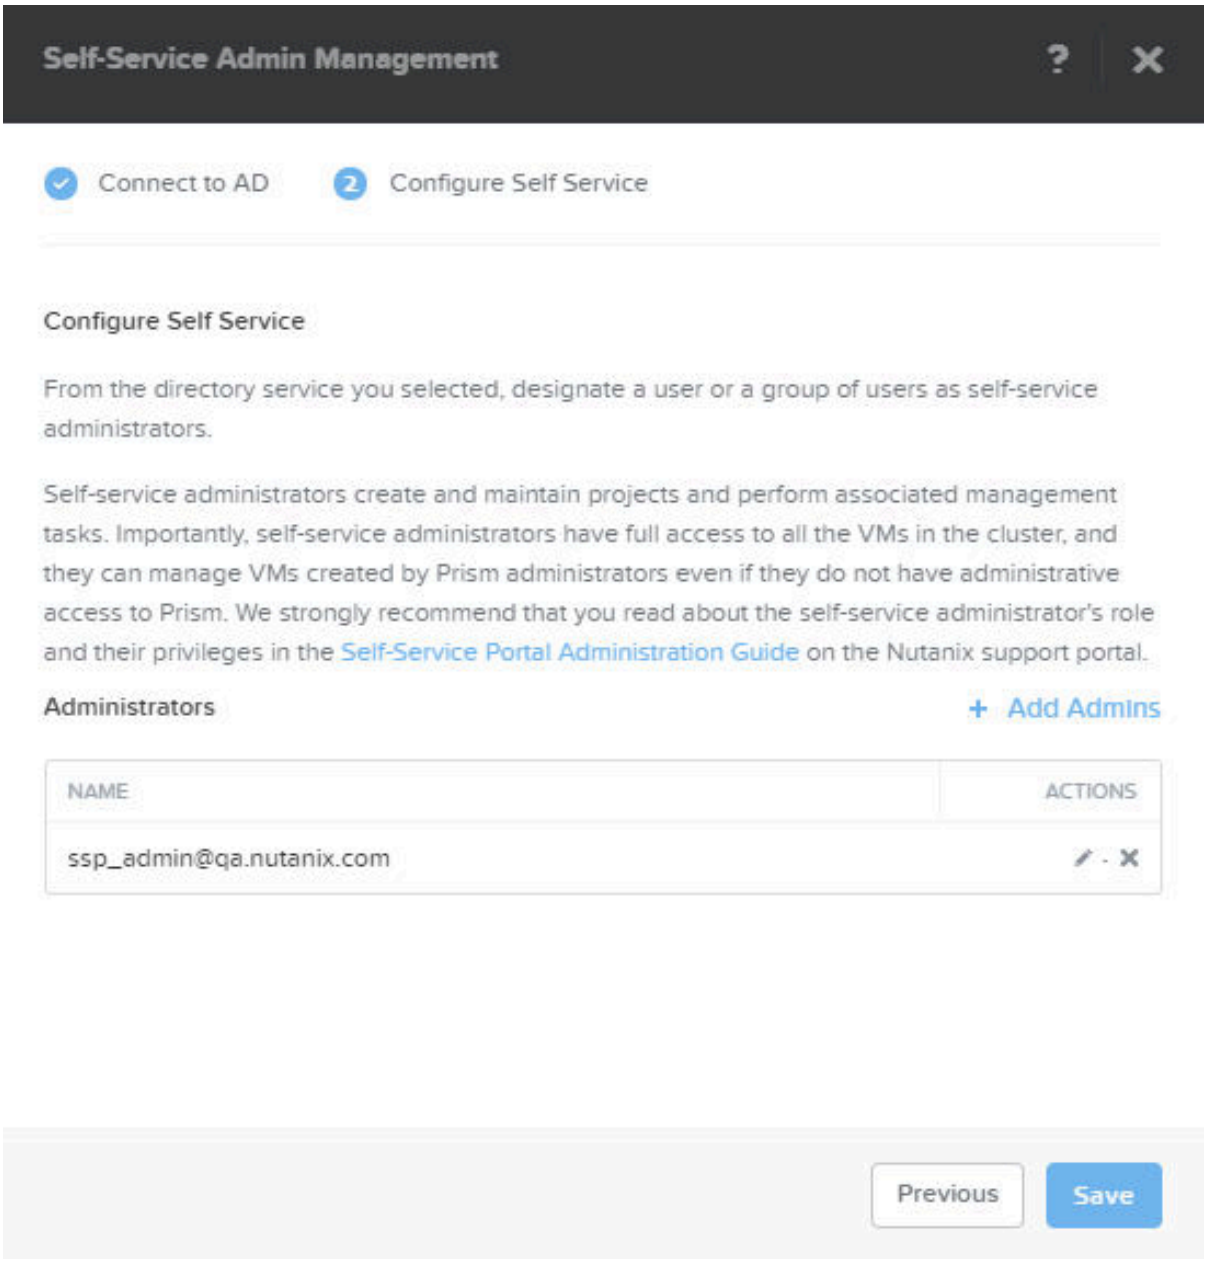

Figure 296: Configure Self Service tab

Migrating Prism Self Service

About this task

If you configured the Prism Self Service feature (previously referred to as the self-service portal or SSP) on a cluster that runs a pre-5.5 version of AOS, that Prism Self Service instance must first be migrated to Prism Central before you can use it. To migrate a Prism Self Service instance, do the following:

Procedure

1. Click the gear icon in the main menu and select **Migrate SSP** from the Settings menu.

   The Migrate Prism Self Service to Prism Central window appears.

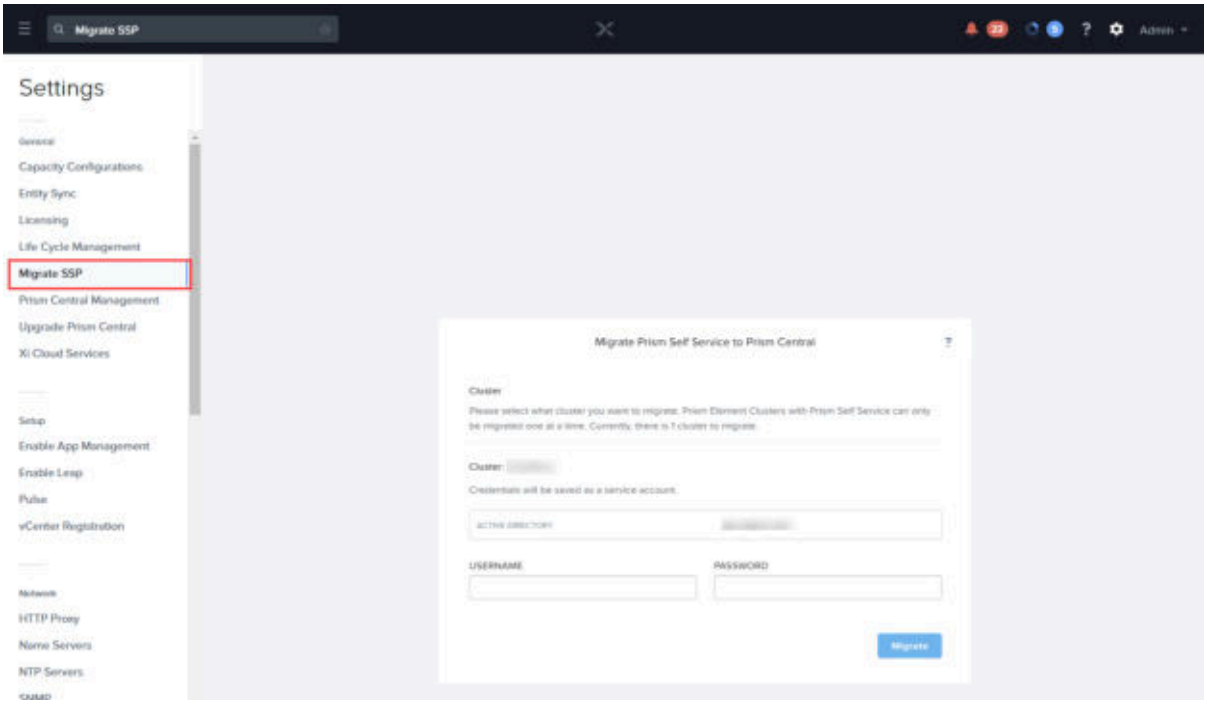

Figure 297: Migrate Prism Self Service to Prism Central Window

2. Select the cluster that hosts the Prism Self Service instance to migrate from the pull-down list.

   If there is only one cluster with a Prism Self Service instance to migrate, that cluster is selected automatically (no pull-down list). If no cluster is found that has a Prism Self Service instance to migrate, a message to that effect appears.

3. Enter the credentials (username and password) for an Active Directory user with administrator permissions.

4. Click the **Migrate** button.

   This starts the migration process. Click the **Close** button to close the window. You can view progress of the migration through the Tasks dashboard (see Tasks View on page 190).

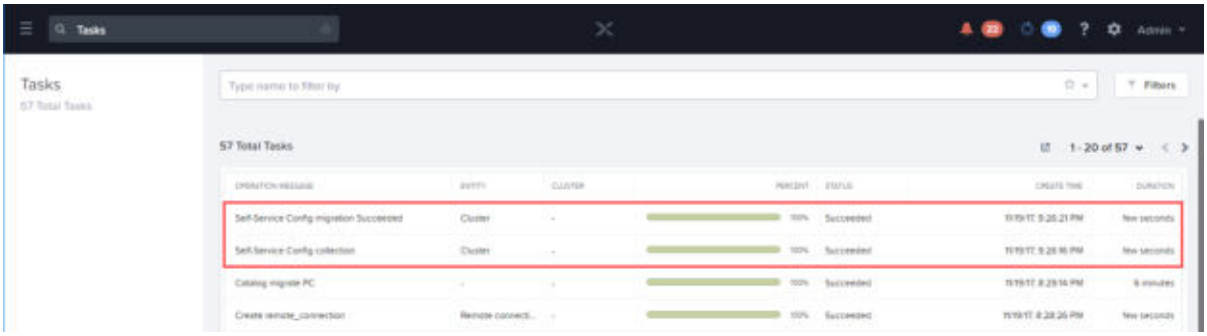

Figure 298: SSP Migration Tasks

Project Management

A project defines Active Directory users or groups to manage common set of requirements or functions. For example, a project can define a team collaborating on an engineering project. A project specifies roles to associate its members, select existing networks that the deployed VMs can use, and (optionally) set usage limits on infrastructure resources.

Creating a Project

Before you begin

Authentication must be configured before you can create a project (see “” ()).

About this task

To create a project, do the following:

Procedure

1. Go to the projects dashboard (see Projects Summary View on page 203) and click the **Create Project** button.
   The Create Project page appears.

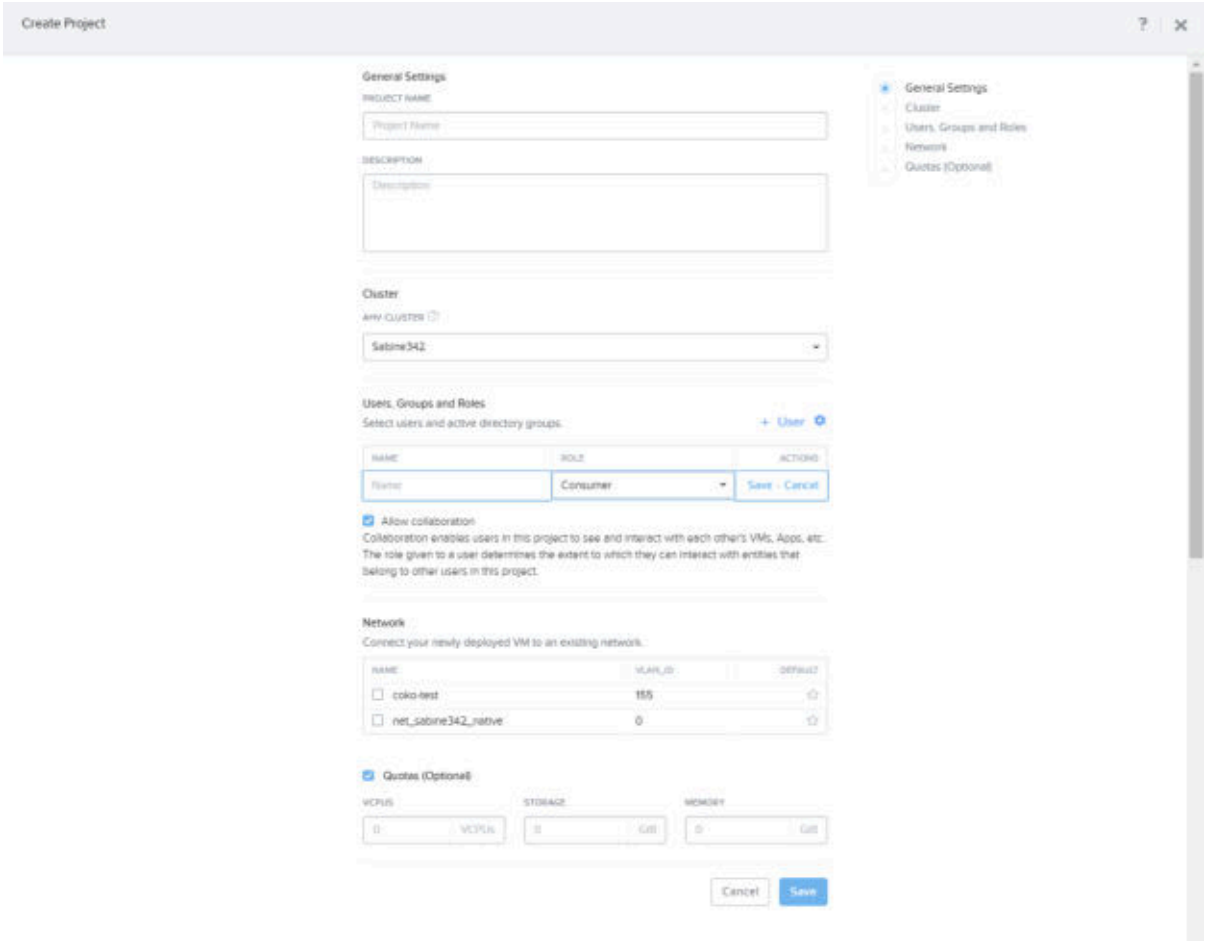

Figure 299: Create Project Page

2. Do the following in the indicated fields:

   a. **Project Name**: Enter a name for the new project.

   b. **Description** (optional): Enter a description of the project.

   c. **Cluster**: Select the target cluster from the pull-down list.

      Note: Only clusters running AHV appear in the list.

   d. **Users, Groups and Roles**: Click the **+ Users** link to open a blank row in the table. Enter the Active Directory name of a user or group (typically in the form *name@domain*) in the **Name** field, select the role for that user from the pull-down list in the **Role** field, and then click **Save** in the **Actions** field.

      If multiple Active Directory domains are configured, you can specify the Active Directory to use by clicking the gear icon next to the **+ Users** link. This displays a **Search Directories**

window. Select the radio button for the Active Directory to use and then click the **Save** button.

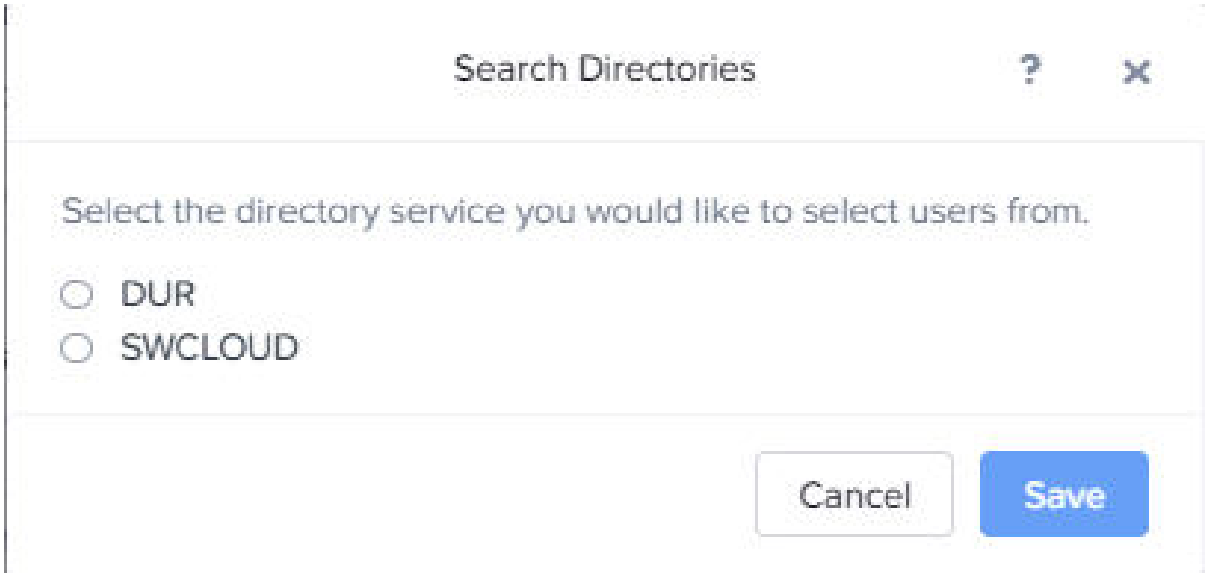

Figure 300: Search Directories Window

Repeat this step for each user or group that you want to add. You can assign one of the predefined roles or a custom role that you previously created (see “” ()). Nested groups (groups within a group) are not supported, that is if a selected group includes a group name (instead an individual name) within that group, the group name is not expanded so those group members are not included in the project.

Note:

- Local users are not supported in a project. You can only add users from your configured directory service.
- If you add a group to a project, users in the group might not appear in the project members list until they log in.

e. **Allow collaboration**: Check the box to allow any group member to see the VMs, applications, and other objects created by other members of the group. If this box is not checked, group members can see only the objects they create. The role assigned a group member determines the permissions that user has on objects created by other group members.

f. **Network**: Select (check the **Name** field box for) the usable network(s) and select (check the **Default** field star icon for) the default network for the project.

g. **Quotas** (optional): Check the box to specify usage limits for compute, storage, and memory in the displayed **vCPUs**, **Storage**, and **Memory** fields.

A quota specifies a usage limit on an infrastructure resource (compute, memory, or storage) for the project. Project members cannot use more than the specified limit. A quota does not guarantee the project a certain amount of infrastructure resources. Instead, it ensures that a single project or a small number of projects do not overrun the infrastructure. If the Nutanix cluster runs out of a resource, project members might not be

able to use the resource even if the project has not reached its specified limit. However, if a project requires more resources, you can increase its quota.

If you do not specify a quota, no usage limit is applied. However, usage statistics are collected even if you do not specify a quota.

h. Click the **Save** button (lower right) to add the project. The page closes and the new project appears in the **Projects** view list.

Modifying a Project

About this task

To update or delete an existing project, do the following:

Procedure

1. Go to the projects dashboard (see Projects Summary View on page 203) and select (check the box for) the target project.
2. Do one of the following:

   » To modify that project, select **Update Project** from the **Actions** pull-down menu. The Edit Project page appears, which includes the same fields as the Create Project page (see Creating a Project on page 462). Update the field values as desired and then click the **Save** button.

   » To delete that project, select **Delete** from the **Actions** pull-down menu. You are prompted to verify the delete (click the **OK** button). The project is then deleted and removed from the list.

   Note: Before you can delete a project, you must first remove any VMs and networks, in that order, from the project. You cannot delete the **default** project.

10

NETWORK AND SECURITY MANAGEMENT

You can view and configure network connections for the cluster in the **Network Configuration**. To display the **Network Configuration** page, click the **Network Config** button in either the VMs or Subnets dashboard.

Network Connections

Each VM network interface is bound to a virtual network, and each virtual network is bound to a single VLAN. Information about the virtual networks configured currently appears in the Network Configuration page.

The Network Configuration page includes three tabs.

- The **Subnets** tab displays a list of the configured networks.

Figure 301: Networks Tab

- The **Internal Interfaces** tab displays a list of LAN interfaces.

Figure 302: Internal Interfaces Tab

- The **Virtual Switch** tab displays a list of virtual switches configured, including the default system-generated virtual switch vs0.

Figure 303: Virtual Switch Tab

The following table describes the fields in each tab.

Table 222: Network Configuration Fields

Parameter	Description	Values
Networks Tab		
Name	Displays the name of the virtual network.	(name)
Virtual Switch	Displays the name of the virtual switch in the form **vs#**, for example **vs0** for virtual switch 0 which is the default virtual switch.	(vs<number>)
VLAN ID	Displays the VLAN identification number for the network in the form **vlan.#**, for example **vlan.27** for virtual LAN number 27.	(ID number)
Internal Interfaces Tab		
Descriptive Name	Displays a name for the LAN.	(LAN name)
Subnet (Gateway IP / Prefix Length)	Displays the subnet that the internal interface belongs to in the form *<IP Address>*/<number (prefix)>	(IP Address/prefix number)
Features	Displays the features available on the internal interfaces.	
Interface	Displays the interface designation such as eth0 or eth1.	(interface name)
Virtual Switch		
Name	Displays the name of the switch in the form **vs#**	(vs<number>)
Bridge	Displays the name of the bridge associated with the virtual switch in the form **br#**, for example **br0** for the default bridge.	(br<number>)
MTU (bytes)	Displays the MTU set for the virtual switch in bytes. The default MTU is 1500.	(number)
Bond Type	Displays the uplink bond type associated with the virtual switch. See the *Bond Type* table in . For example, **Active-Backup**	(<bond_type>)

LAG and LACP on the ToR Switch

For more information, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 431.

Configuring Network Connections

Prism Central allows you to configure network connections for a selected cluster.

About this task

To create one or more network configurations, do the following:

Procedure

1. Do one of the following:
   - » Go to the List tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Network Config** button.
   - » Go to the Network and Security dashboard (see Subnets View on page 135 and click the **Network Config** button.
2. In the Subnets tab of the Network Configuration page, click the **Create Network** button.

   This displays the Create Network page.

Figure 304: Create Network Page (expanded)

3. Do the following in the indicated fields:

   a. **Subnet Name**: Enter a name for the network.

   b. **Virtual Switch**: Select the virtual switch for the network from the drop-down list.

      You can associate the default virtual switch vs0 or create a virtual switch. To create a virtual switch, see Creating or Updating Virtual Switch on page 426.

      For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

   c. **VLAN ID**: Enter the number of the VLAN.

      Enter just the number in this field, for example 1 or 27. Enter 0 for the native VLAN. The value appears as **vlan.1** or **vlan.27** in displays.

   d. **Enable IP Address Management**: Check the box to have the cluster control IP addressing in the network.

      Checking this box display additional fields. If this box is not checked, no network management is attempted. In this case it is assumed management for this virtual LAN is handled outside the cluster.

   e. **Network IP Address/Prefix Length**: Enter the IP address of the gateway for the network and prefix with the network prefix (CIDR notation, for example, 10.1.1.0/24).

   f. **Gateway IP Address**: Enter the VLAN default gateway IP address.

   g. **Configure Domain Settings**: Check this box to display fields for defining a domain.

      Checking this box displays fields to specify DNS servers and domains. Unchecking this box hides those fields.

   h. **Domain Name Servers (comma separated)**: Enter a comma-delimited list of DNS servers.

   i. **Domain Search (comma separated)**: Enter a comma-delimited list of domains.

   j. **Domain Name**: Enter the VLAN domain name.

   k. **TFTP Server Name**: Enter the host name or IP address of the TFTP server from which virtual machines can download a boot file. Required in a Pre-boot eXecution Environment (PXE).

   l. **Boot File Name**: Name of the boot file to download from the TFTP server.

4. To define a range of addresses for automatic assignment to virtual NICs, click the **Create Pool** button (under **IP Address Pools**) and enter the following in the Add IP Pool page:

   Note: If no pool is provided, the user must assign IP addresses to VMs manually.

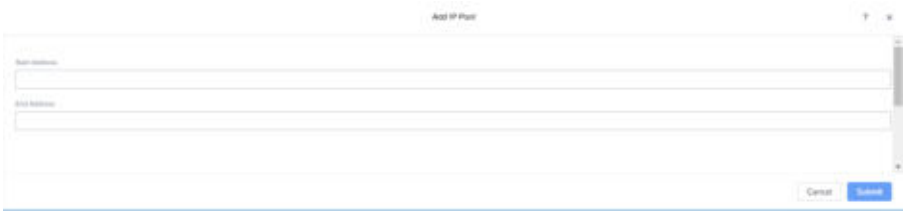

   Figure 305: Add IP Pool Page

   a. Enter the starting IP address of the range in the **Start Address** field.

   b. Enter the ending IP address of the range in the **End Address** field.

   c. Click the **Submit** button to close the page and return to the Create Network page.

5. To configure a DHCP server, check the **Override DHCP server** box and enter an IP address in the **DHCP Server IP Address** field.

   This address (reserved IP address for the Acropolis DHCP server) is visible only to VMs on this network and responds only to DHCP requests. If this box is not checked, the **DHCP Server IP Address** field is not displayed and the DHCP server IP address is generated automatically. The automatically generated address is *network_IP_address_subnet*.254, or if the default gateway is using that address, *network_IP_address_subnet*.253.

6. Click the **Save** button to configure the network connection, close the Create Network page, and return to the Network Configuration page.

7. Click the **Close** button to close the Network Configuration page.

Creating or Updating Virtual Switch

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can create or update a VS using the Create Virtual Switch or Update Virtual Switch dialog box. The Create Virtual Switch and Update Virtual Switch dialog boxes have identical parameters and settings.

To open the Create Virtual Switch dialog box in Prism Central, do the following.

Procedure

1. Click **Subnets** > **Network Config**.

2. On the Network Configuration page, click the **Virtual Switch** > **Create VS** link.

About this task

To create a virtual switch using the Create Virtual Switch dialog box, do the following.

Procedure

1. In the Create Virtual Switch dialog box, provide the necessary information on the **General** tab.

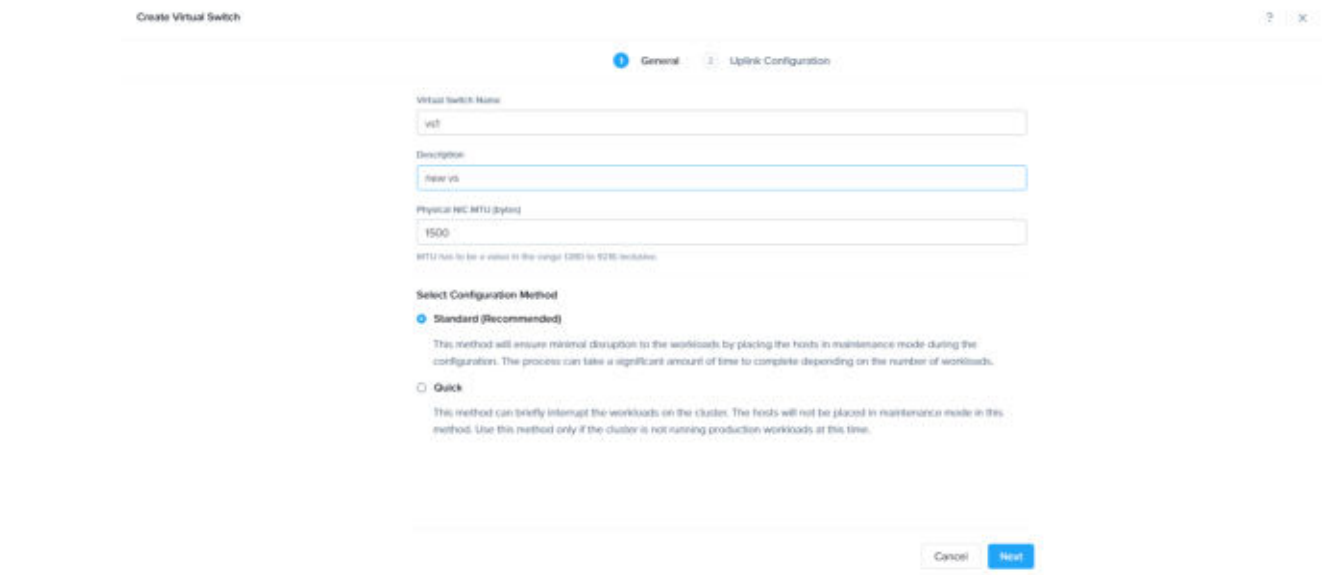

Figure 306: Create Virtual Switch - General tab

Field	Description
Virtual Switch Name	Enter a name for the virtual switch.
Description	Provide a description for the virtual switch that helps identify the virtual switch.
Physical NIC MTU (bytes)	MTU must be a value in the range 1280 to 9216 inclusive.

Field	Description
Select Configuration Method	Select one of the two methods that you can use to implement the VS configuration: • **Standard (Recommended)**: This method ensures no disruptions occur to the workloads by putting the hosts in maintenance mode and migrating the VMs out of the host before applying the configuration. This process requires a longer duration of time to complete. The time required depends on the number and configuration of VMs. In this method, the VS configuration is deployed in the rolling update process. • **Quick**: This method interrupts the workloads running on the hosts. Use this method only if you are not running production workloads because it may result in network interruptions. In this method, the VS configuration is deployed in a rolling update process but the nodes are not put in maintenance mode before modifying the VS configuration on the node.

2. Click **Next** to go to the Uplink Configuration tab.

   On the Uplink Configuration tab, provide the following details:

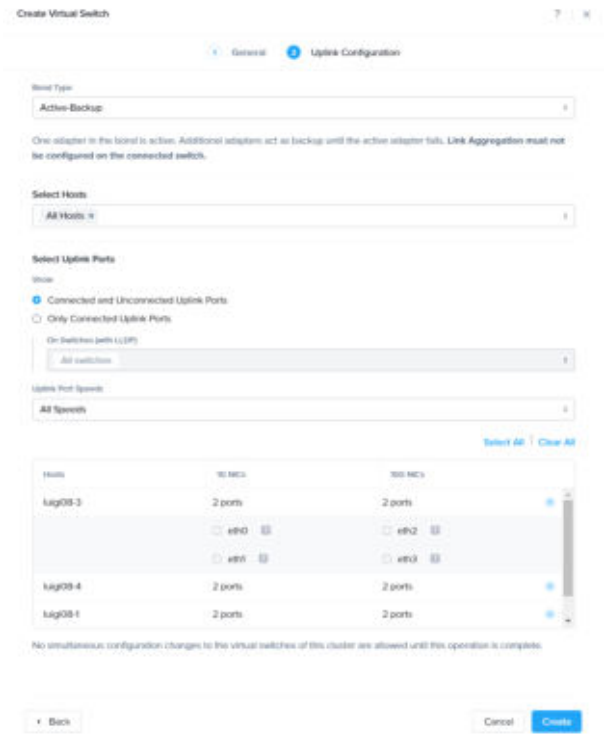

Figure 307: Create Virtual Switch - Uplink Configuration tab

Field	Description and Value
Bond Type	Select an appropriate bond type. See the *Bond Types* table for details about the bond types.
Select Hosts	Select the hosts that can host VMs.
Select Uplink Ports	Select the criteria that need to be satisfied for the uplink ports. The available uplink ports that satisfy the criteria are displayed in the (Host port) table at the bottom of this tab.
Show	(Port Type) Select one of the following: **Connected and Unconnected Uplink Ports**: Select this if you want to use ports that are not currently connected but may be connected later. **Only Connected Uplink Ports**: Select this option if you want to use only the connected ports. You must also select in the **On Switches (with LLDP)** drop down box the switches with LLDP.
Uplink Port Speeds	Select a speed to display the ports that have the selected speed. You can select speeds such as **1G**, **10G** or both (**All Speeds**). The speeds displayed depend on the NIC type that is installed on the host. Based on your selection the columns in the (Host Port) table change dynamically to display the ports with the speeds you selected.
(Host Port) table	Based on the Selections you made in this Select Uplink Ports section, a table displays the hosts that have the uplink ports that satisfy the selected criteria. Select the ports you need for this configuration from the list. Click the down arrow on the right side of the table to display the ports listed for each host. Note: A port listing is greyed out if it is unavailable because it is already associated with another virtual switch. Click **Select All** to select all the ports available and listed. Click **Clear All** to unselect all the ports available and listed.

Table 223: Bond Types

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Backup	Recommended. Default configuration, which transmits all traffic over a single active adapter.	10 Gb	10 Gb

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Active with MAC pinning Also known as balance-slb	Works with caveats for multicast traffic. Increases host bandwidth utilization beyond a single 10 Gb adapter. Places each VM NIC on a single adapter at a time. Do not use this bond type with link aggregation protocols such as LACP.	10 Gb	20 Gb
Active-Active Also known as LACP with balance-tcp	LACP and link aggregation required. Increases host and VM bandwidth utilization beyond a single 10 Gb adapter by balancing VM NIC TCP and UDP sessions among adapters. Also used when network switches require LACP negotiation. The default LACP settings are: • Speed—Fast (1s) • Mode—Active fallback-active-backup • Priority—Default. This is not configurable.	20 Gb	20 Gb
No Uplink Bond	No uplink or a single uplink on each host. Virtual switch configured with the No uplink bond uplink bond type has 0 or 1 uplinks. When you configure a virtual switch with any other bond type, you must select at least two uplink ports on every node.	-	-

Note: The Maximum VM NIC Throughput and Maximum Host Throughput values are not restricted to the value provided in this table. The values in the table are indicated for an assumption of 2 x 10 Gb adapters.

For more information about uplink configuration, see the *Uplink Configuration (AHV Only)* topic in the *AHV Administration Guide*.

3. Click **Create** to create the virtual switch.

   Click **Cancel** to exit without creating the virtual switch.

   Click **Back** to go back to the General tab.

What to do next

> Important:
>
> You can migrate or convert the bridges other than br0 in the cluster to virtual switches after you upgraded the minimum or compatible version of AOS and AHV. You can convert the other bridges only on Prism Web Console or using aCLI. You can convert only one bridge at a time. You need to repeat the workflow for every bridge that you want to convert to a virtual switch.
>
> For more information about converting the bridges other than br0, see *Migrating Bridges after Upgrade*in *Prism Web Console Guide*.

- If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

  For information about enabling LAG and LACP, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 431

Deleting a Virtual Switch

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can delete a VS on the Virtual Switch tab in the Network Configuration dialog box. When you delete a virtual switch, the bridge created by the virtual switch is also removed completely.

You can also disable all the virtual switches. However, you can do so only using CLI. Prism does not support disabling virtual switches. Use the `net.disable_virtual_switch` to disable the virtual switch functionality and remove all the virtual switches including the default virtual switch vs0. However, the bridge and uplink bond or interface configurations are not affected by this command. This action cannot be undone either.

To delete a virtual switch, do the following:

Procedure

1. Click the **Delete** icon (trash icon) next to the virtual switch listed in the list on the Virtual Switch tab.

   The Delete Virtual Switch dialog box opens. Deletion cannot be undone or reversed. Therefore, you must confirm the deletion in this dialog box.

2. Click **Delete** to confirm deletion.

   Click **Cancel** to close the Delete Virtual Switch dialog box without deleting the virtual switch.

What to do next

Check if the virtual switch is deleted from the list on the Virtual Switch tab.

Enabling LAG and LACP on the ToR Switch (AHV Only)

Active-Active with MAC pinning does not require link aggregation since each source MAC address is balanced to only one adapter at a time. If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

About this task

In an **Active-Active** NIC-teaming configuration, network traffic is balanced among the members of the team based on source and destination IP addresses and TCP and UDP ports. With link aggregation negotiated by LACP, the uplinks appear as a single layer-2 link so a VM MAC address can appear on multiple links and use the bandwidth of all uplinks. If you do not enable LAG and LACP in this **Active-Active** configuration, network traffic might be dropped by the switch.

Perform the following procedure to enable LAG and LACP on a TOR switch after enabling the Active-Active policy by using the Prism Element web console.

Procedure

1. Put the node in maintenance mode. This is in addition to the previous maintenance mode that enabled Active-Active on the node.
2. Enable LAG and LACP on the ToR switch connected to that node.
3. Exit maintenance mode after LAG and LACP is successfully enabled.
4. Repeat steps 1 to 3 for every node in the cluster.

Modifying Network Connections

About this task

To modify or delete a network configuration (defined on an Acropolis managed cluster), do the following:

Procedure

1. Do one of the following:
   - » Go to the List tab of the VMs dashboard (see VMs Summary View on page 82) and click the **Network Config** button.
   - » Go to the Networks dashboard (see Subnets View on page 135 and click the **Network Config** button.
2. To modify a network configuration, in the Subnets tab, select the target network, and click the pencil icon.

   The Update Subnet page appears, which contains the same fields as the Create Subnet page (see Configuring Network Connections on page 423). Do the following:

   a. Update the field values as desired.

   b. When the configuration is correct, click the **Save** button to close the page, update the network configuration, and return to the Network Configuration page.

3. To delete a network configuration, in the Subnets tab, select the target network, and click the X icon (on the right).

   A window prompt appears to verify the action; click the **OK** button. The network configuration is removed from the list.

   > Note:
   >
   > You can also select the target network on the Networks dashboard and select **Delete** in the **Actions** dropdown list.

4. Click the **Close** button to close the Network Configuration page.

11

POLICIES MANAGEMENT

Prism Central allows you to create and apply policies to manage and control access and operations in a cluster. This chapter describes how to implement certain policy types.

- “” () describes how to create and apply security-related policies.
- NGT Policies on page 478 describes how to create and apply NGT-related policies.
- Image Placement Policies on page 480 describes how to create and apply image placement policies.

> Note: See the *Leap Administration Guide* for information about creating and applying protection policies and recovery plans.

NGT Policies

If a restart is required post installation or upgrade, you can define a policy on when the restart should happen. You can use these policies when you have different set of VMs which you would like to be restarted at different times, for example if they are in different time zones or if their down times are expected to be different. Policy workflow are defined on the categories. You have to attach the category to the VM and then you can create a policy on the VM.

Creating a New Restart Policy

Perform the following procedure to create a new restart policy.

Procedure

1. Go to the NGT policies dashboard (see NGT Policies View on page 139) and click the **New VM Reboot Policy** button.

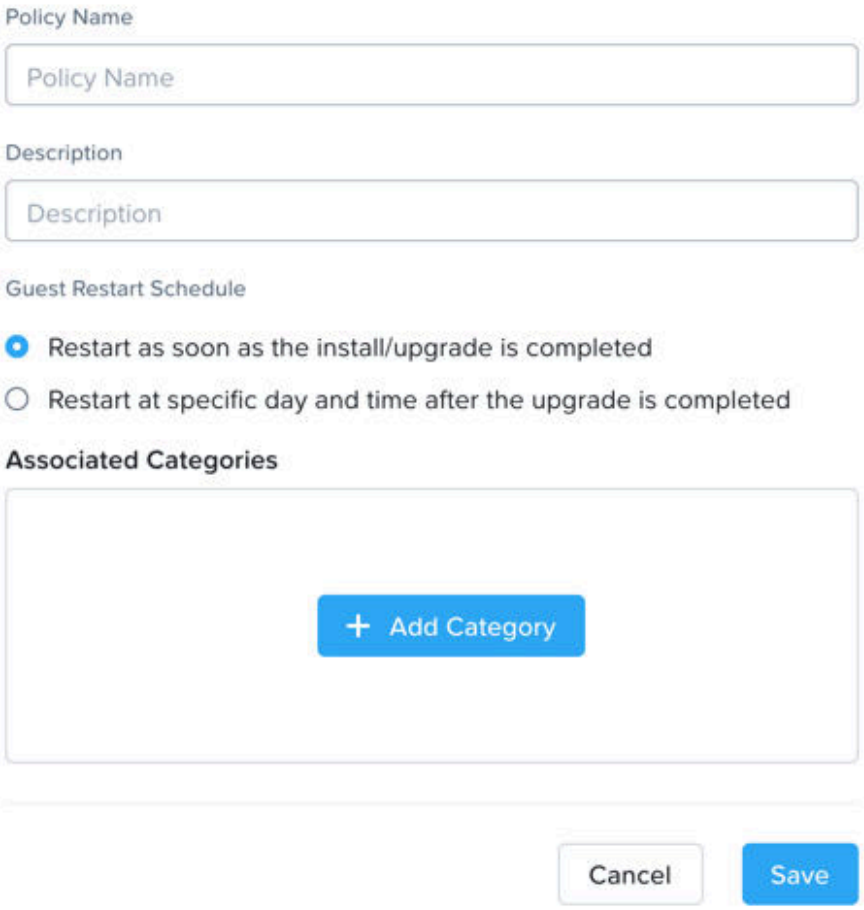

Figure 308: NGT Policy

2. In the **Policy Name** text box, type a policy name.
3. In the **Policy Description** text box, type a description of the policy.
4. Configure the guest restart schedule.

   - **Restart as soon as the install/upgrade is completed**: Select this option if you want to restart the VMs of the selected category right after the install process is completed.
   - **Restart at specific day and time after the upgrade is completed**: Select this option and choose the date and time on which you want to the restart to happen.

     Note: The VMs will be restarted depending on the Prism Element timezone settings.

5. Click **Add Category** and select the category that you want to apply this policy.

   Note: You can select multiple categories and the application of policy is dependent on the following rules.

   - 1. If you choose to restart the VM by using the **Restart as soon as the install/upgrade is completed** option either by using NGT install/upgrade workflow or by using this procedure, it will always take precedence over other workflows. Also, if you select either of these options, the policy will be created as a System policy and you will not be able to modify this policy. Once these policies are applied, it is removed from the user interface.
     2. If you schedule a restart during the NGT install or upgrade workflow, it will take precedence over other policies, for example, the policy that you have created by using this procedure.
     3. If there are multiple VM restart policies that are created at the same time by using this procedure, an alert message will be displayed.

6. Click **Save**.
   The policy is created and saved in the NGT Policies pane.

Modifying a Restart Policy

You can update and delete the restart policy that you have created by following this procedure.

Procedure

1. Go to the NGT policies dashboard (see NGT Policies View on page 139).
2. Check the box for the policy to update and then select **Update** from the **Actions** drop-down menu.
   a. Make the necessary changes to the policy and click **Save**.
3. To delete a policy, select the policy and from the **Actions** drop-down menu click **Delete**.
   a. Click Delete in the confirmation dialog box to delete the policy.

Image Placement Policies

Prism Central enables you to upload VM images from your workstation or from a remote server. It also enables you to manually specify a subset of registered clusters as targets for those images. However, image placement decisions are often driven by compliance requirements, policies, and regulations. For example, you might have to confine a VM image to a Nutanix cluster in a specific location. Or, a cluster in one region might require a Windows VM image with a specific set of applications while a cluster in another region might require a Linux VM with its own set of applications.

Prism Central enables you to configure policies that govern which clusters receive the images that you upload. These policies, called image placement policies, map images to target clusters using categories associated with both those entities. For example, an image placement policy states that the Linux images associated with the `OS:Linux` and `distribution:Centos` categories must be uploaded to the clusters associated with the `Location:SFO` category.

Image placement policies also specify how strictly the policy must be enforced. Soft enforcement allows clusters not identified by the policy to use the images, when required, by checking them out. Hard enforcement disallows such usage. By using both soft and hard enforcement on the same category of clusters, you can address more specific use cases. For more information, see Sample Scenarios.

Image placement policies maintain a state of compliance across the clusters registered with Prism Central. They detect and correct violations introduced by changes in category assignments and updates to their configuration (including a switch from soft enforcement to hard enforcement). For more information, see Sample Scenarios.

To view the location of, the category and Image Placement policy associated with an image, see Images Summary View on page 125 and Image Details View on page 127.

Configuration Workflow

About this task

To set up an image placement policy, do the following:

Procedure

1. Create categories for the following entities:
   a. Clusters (for example, create categories based on cluster location or region)
   b. Images (for example, create categories based on operating system and size)
   For information about creating a category, see Creating a Category.
2. Apply the cluster categories to clusters and image categories to images.
   For information about associating categories with a Nutanix cluster, see Associating a Registered Nutanix Cluster with Categories. For information about associating categories with an image, see Associating an Image with Categories.
3. Configure the image placement policy. See Configuring an Image Placement Policy.

Associating a Registered Nutanix Cluster with Categories

About this task

To associate categories with a registered Nutanix cluster, do the following:

Procedure

1. In Prism Central, in the **Entities** menu, go to **Hardware** > **Clusters**.
2. Select the cluster that you want to associate with categories.
3. In the **Actions** menu that is displayed, click **Manage Categories**.
4. In the **Search for a category** field, type the name of the category or value that you want to add, and then select the category and value from the list that is displayed.
5. Click the add button to add as many search fields and repeat this step for each category that you want to add. To remove a category, click the remove button beside the field.
6. Click **Save**.

Associating an Image with Categories

About this task

To associate categories with an Image, do the following:

Procedure

1. In Prism Central, in the Entities menu, go to **Virtual Infrastructure** > **Images**.
2. Select the image that you want to associate with categories.
3. In the **Actions** menu that is displayed, click **Manage Categories**.
4. In the **Search for a category** field, type the name of the category or value that you want to add, and then select the category and value from the list that is displayed.
5. Click the add button to add as many search fields and repeat this step for each category that you want to add. To remove a category, click the remove button beside the field.
6. Click **Save**.

Configuring an Image Placement Policy

About this task

To configure an image placement policy, do the following:

Before you begin

Complete the following tasks:

- Configure the categories that you need to associate with the images.
- Configure categories for the clusters to which you want Prism Central to upload images, and then associate the categories with the clusters.
- Review existing image placement policies to avoid causing conflicts. For information about issues caused by conflicting policies, see How Prism Central Handles Conflicting Policies.

Procedure

1. In the entity menu, go to **Policies** > **Image Placement Policies**, and then click **Create Image Placement Policy**.
2. On the **Create Image Placement Policy**, enter a name and description for the policy.
3. Click inside the **Assign Images From The Following Categories** search field and select the category you want to associate with the images.

   To associate additional categories with the images, click the add icon beside the search field, and select the categories you need.
4. Click inside the **To The Clusters From The Following Categories** search field and select the category you associated with the target clusters.

   To specify additional cluster categories, click the add icon beside the search field, and select the categories you need.
5. In the **Policy Enforcement** section, from the **Enforcement** list, select one of the following:

   » **Soft**. The images are uploaded to the set of clusters identified by the image placement policy. In addition, no restriction is placed on using those images on clusters outside the identified set. For example, if clusters A, B, and C are registered with Prism Central, and an image placement policy with a soft enforcement policy uploads an image to clusters A

and B, you can use the image to create a VM on cluster C. The checkout to cluster C is not blocked.

» **Hard**. The images are uploaded to the set of clusters identified by the image placement policy. You cannot use the VM on any cluster outside the identified set. For example, if clusters A, B, and C are registered with Prism Central, and an image placement policy with a hard enforcement policy uploads an image to clusters A and B, the image is available only on clusters A and B. Cluster C cannot check out the image.

For various scenarios involving these enforcement policies, see Sample Scenarios.

6. Click **Save**.

Updating an Image Placement Policy

You can update an image placement policy. Updates to an image placement policy can result in a policy violation. Prism Central attempts to correct the violation by executing a series of actions.

About this task

For example, if you change the cluster category such that the policy now identifies a different set of clusters as targets, Prism Central first copies the images to the clusters in the new category, and then removes the images from the clusters in the previously mentioned category. If you change the policy enforcement setting from Soft to Hard, and if any clusters that are not identified as targets had previously checked out the images, Prism Central corrects the violation by deleting the images from those clusters. An update to a policy might also result in a conflict with one or more of the other policies that apply to the same category of images. For more information about how Prism Central handles conflicting policies, see How Prism Central Handles Conflicting Policies.

To update an image placement policy, do the following:

Procedure

1. In the entity menu, go to **Policies** > **Image Placement Policies**, and then click the image placement policy that you want to update.
2. From the **Actions** menu that is displayed, click **Update**.
3. The **Edit Image Placement Policy** page that is displayed includes the same options and settings as the **Create Image Placement Policy** page.
4. Update the settings that you want, and then click **Save**.

How Prism Central Handles Conflicting Policies

Policies are said to be in conflict if they identify the same category of images but their resulting cluster placements are different.

For example, the following policies are conflicting policies if they apply to the same category of images:

- Policy P1, which identifies clusters C1 and C2 as target clusters, and the policy enforcement setting is Soft.
- Policy P2, which identifies clusters C3 and C4 as target clusters, and the policy enforcement setting is Hard.

Prism Central handles such conflicts in the following manner:

- Prism Central ignores conflicting policies and does not perform any action on the images. If policy P1 was configured first, Prism Central stops enforcing P1 as soon as P2 is configured, and it does not enforce P2 either. When there is a policy conflict, the conflict is noted in the **Policies** tab of the details view for that image.
- Prism Central continues to enforce other policies that are not in conflict. For example, if policies P3 and P4 apply to the same categories of images as P1 and P2, Prism Central continues to enforce P3 and P4 as long as they are not in conflict with P1 and P2.

Sample Scenarios and Configurations

The following table describes how to configure image management policies to achieve a desired result. The examples use a Prism Central instance to which three clusters A, B, and C and registered (possibly among other clusters).

Table 224: Sample Scenarios and Configurations

Desired Result	Image Placement Policy Configuration
The images must be available on any one of the clusters (A, B, or C) and can be checked out to the other two clusters and to any other clusters that might be registered with Prism Central now or in the future.	Create an image placement policy that identifies any one of the clusters (A, B, or C) as the target cluster (by the use of categories) and set the policy enforcement to Soft.
The images must be available on all of the clusters A, B, and C and can be checked out to any other clusters that are registered with Prism Central now or in the future.	Create an image placement policy that identifies clusters A, B, and C as the target clusters (by the use of categories) and set the policy enforcement to Soft.
The images must be available on all of the clusters A, B, and C and cannot be checked out to any other clusters that are registered with Prism Central now or in the future.	Do the following: • Create an image placement policy that identifies clusters A, B, and C as the target clusters (by the use of categories) and set the policy enforcement to Soft. (This policy ensures that the images are available on all of the three clusters) • Create a second image placement policy that identifies clusters A, B, and C as the target clusters (by the use of categories) and set the policy enforcement to Hard. (This policy makes sure that the images are available only on these three clusters)
The images must be available only on cluster A and cannot be checked out to clusters B and C now, and the images also cannot be checked out to any other clusters that are registered with Prism Central in the future.	Create an image placement policy that identifies cluster A as the target cluster (by the use of categories) and set the policy enforcement to Hard.

Desired Result	Image Placement Policy Configuration
The images are available on cluster A and must now be available on clusters B and C. Cluster A can use the images if required. The policy enforcement is currently Soft.	Do the following: • Dissociate the categories specified in the image placement policy from cluster A and associate them with clusters B and C. • Do not update the image placement policy (no changes are required). Prism Central takes the following corrective actions: 1. Prism Central copies the images to clusters B and C. If clusters B and C had checked out the images before the reassignment of categories to the clusters, the images are not copied. 2. Additionally, because the Soft policy enforcement setting is retained, Prism Central does not remove the image from cluster A.
The images are available on cluster A and must now be available only on clusters B and C. Cluster A must not be permitted to use the images. The policy enforcement is currently Soft.	Do the following: • Dissociate the categories specified in the image placement policies from cluster A and associate them with clusters B and C. • Change the policy enforcement setting in the image placement policy to Hard. Prism Central takes the following corrective actions: 1. Prism Central copies the images to clusters B and C. If clusters B and C had checked out the images before the reassignment of categories to the clusters, the images are not copied. 2. To honor the Hard policy enforcement setting, Prism Central removes the image from cluster A.

12

OPERATIONS MANAGEMENT

- Analysis: Analyze cluster activity (see Analysis Dashboard on page 486)
- App Discovery: Discover applications running in a cluster (see Application Discovery on page 518)
- Cost Management: Analyze and manage costs (see Cost Management (Xi Beam) on page 539)
- Integrations: Monitor select applications in a cluster (see Application Monitoring on page 540)
- Ops Policies: Manage operations policies (see Operations Policy Management on page 417)
- Planning: Plan for resource capacity and usage (see Resource Planning on page 560)
- Playbooks: Create playbooks for task automation (see Task Automation on page 585)
- Reports: Generate resource and activity reports (see Reports View on page 194)

Performance Monitoring

Prism Central allows you to monitor performance across the registered clusters.

- You can view performance graphs for selected elements from the Analysis dashboard (see Analysis Dashboard on page 486).
- You can create custom entity and metric performance charts (see Chart Management on page 496).

> Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

Analysis Dashboard

The Analysis dashboard allows you to create sessions with charts that can monitor dynamically a variety of performance measures. To view the Analysis dashboard at any time, select **Operations** > **Analysis** from the entities menu (see Entities Menu on page 13).

When you open the Analysis dashboard, it displays the session that you were last working in as the default session. A session helps you correlate the metrics with the alerts and events for troubleshooting. You can create new sessions to troubleshoot new situations and preserve the old sessions created for old issues.

When you open the Analysis dashboard for the first time after installing or upgrading Prism Central, the Analysis dashboard displays a session named **Analysis Session**. Prism displays the label **System** with the name of this session in the Sessions dashboard to indicate that this session is a system-generated session. Prism displays the label next to the **Analysis Session** session name to indicate that this session is system-generated. You can set the default session

for display in the Analysis dashboard in the Sessions dashboard. For more information about the Sessions dashboard, see Sessions Dashboard on page 512.

> Note:
>
> The updated Analysis dashboard and Sessions dashboard are available only with the Prism Pro and Prism Ultimate licenses. With the Prism Starter license, only the system-generated session is available.

Prism Central saves all session changes and new sessions automatically including **Session Name** changes, except changes made using specific options that open respective dialog boxes to facilitate the changes. For example, if you change the name of the session in the Analysis dashboard, then the change is automatically saved. However, when you edit a chart using the **Edit** option and the **Edit Chart** dialog box on the Analysis dashboard, you need to click the **Update** button to save the changes made in the dialog box. When you modify the session name or description in the Sessions dashboard using the **Edit Session Details** action, you need to click the **Update** button to save the changes made in the dialog box.

The following figures are sample views of the Analysis dashboard. Some fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated.

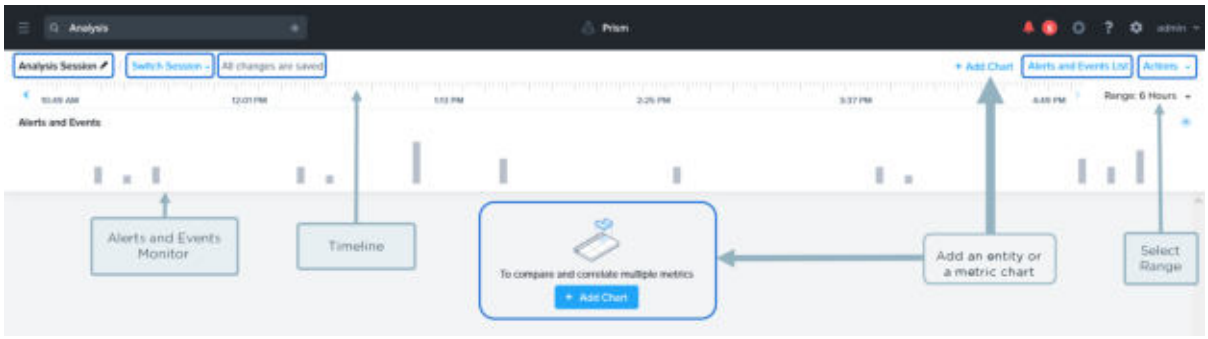

Figure 309: Analysis Dashboard - Prism-generated Session

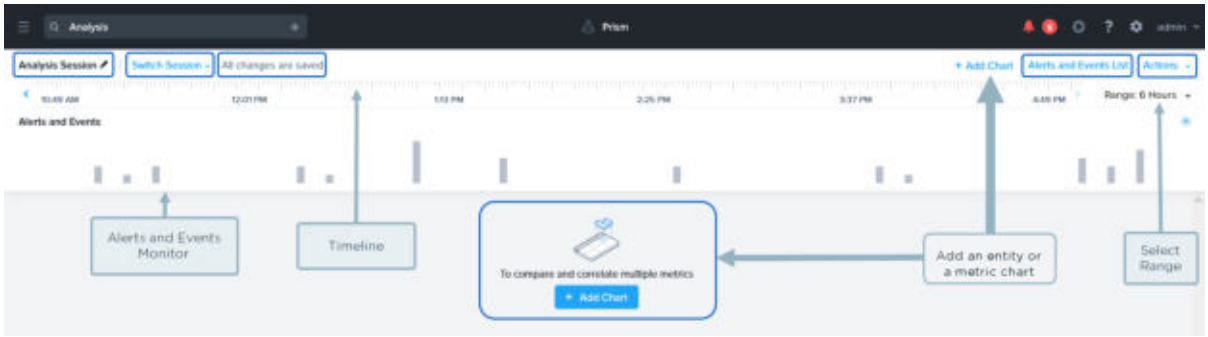

Figure 310: Analysis Dashboard - User-defined Session

The Analysis dashboard elements are described in the table starting with the elements in the top left corner to the right side of the dashboard.

Table 225: Analysis Dashboard Elements

Field	Description
Session	The name of the session that is displayed in the Analysis dashboard. The Analysis dashboard displays the session you viewed last by default. You can change the default session in **Settings** in the Sessions dashboard.
Switch Session	A dropdown menu of all available sessions created and saved. Select the session you want to view and modify the details for. Select **Switch Sessions** > **View All Sessions** to view the Sessions dashboard that lists all the sessions created and saved.
Changes saved status	The status whether the changes made to the session are saved. • Displays **Saving** when Prism is saving the changes. • Displays **All changes are saved** when Prism has saved all the changes.
Add Chart	You can add a chart by clicking: • The **Add Chart**option at the top left of the dashboard. • The **Add Chart** button at the bottom center of the dashboard. You can define a chart to monitor many entities with the same single metric. You can define a chart to monitor several metrics for a single entity. Note: You cannot define a single chart with many entities and many metrics.

Field	Description
Alerts and Events monitor	Displays the alerts and events that occurred. The alerts and events occurring at any specific point of time are displayed as a stacked bar with three colored segments. The hover image of the stacked bar displays the following: • Time range for the stacked bar. • Red segment depicting **Critical Alert** with the number of alerts. • Yellow segment depicting **Warning Alert** with the number of alerts. • Grey segment depicting **Events** with the number of events. Click the stacked bar to open the Alerts and Events List panel on the right. The Alerts and Events List panel displays the details of the alerts and events represented by the stacked bar.
Alerts and Events List	Provides more information about the alerts and events. When you click: • The Alerts and Events List option at the top right of the Analysis dashboard, the Alerts and Events List panel provides information about all the alerts and events that have occurred during the selected time range. • A stacked bar in the **Alerts and Events** monitor, the Alerts and Events List panel provides information about all the alerts and events that occurred in the time range indicated by the stacked bar that you clicked on the timeline. See the *Alerts and Events List Panel Elements* table for descriptions of the tabs and fields in the **Alerts and Events List** panel.

Field	Description
Actions	Provides the following three options for the current session: 1. **Edit Session Details** 2. **Delete** (session) This action is not available and appears dimmed if the current session is the system-defined session, because the system-defined session cannot be deleted. 3. **Close** (session) When you close the current session, the Sessions dashboard is displayed.
(Defined chart monitors)	Displays the monitors for the charts that you defined for the session.

Table 226: Alerts and Events List Panel Elements

Field	Description
Alerts tab	Displays the list of alerts. The tab name includes colored icons with inscribed numbers. The color and number provide the following information: • Color depicts the severity type of the alert. Red indicates critical and yellow indicates warning. • Number indicates the number of respective alerts. In the example depicted in the figure, the red icon has 1 inscribed indicating that there is one critical alert in the stacked bar instance.
Group by field in Alerts tab	Provides the grouping criteria that you can use to group the alerts. You can group the alerts by: • **Entity Type** • **Impact** • **Severity**

Field	Description
Events tab	Displays the list of events. The tab name includes a grey-colored icon with an inscribed number. The inscribed number indicates the number of events. You can see the event details page by clicking the event in the list. The events in the list are grouped into two default groups: • **System Action** • **User Action**
(Time range for the stacked bar)	The time range considered for the stacked bar plotted on the **Alerts and Events** monitor. The time range is specifically displayed only when you open the Alerts and Events List panel by clicking a stacked bar in the **Alerts and Events** monitor.
Hide/Show toggle	Displayed for each group of alerts or events in the list. This toggle allows you to hide or show the alerts or events in a group.

Modifying the Time Range

You can select a time range for the session and the charts you choose in the Analysis dashboard.

About this task

The time range you select applies to the session as a whole including the charts you add to the session. The default time range is **1 Day**.

> Note: The time range selection affects the whole session and all the charts added to the session.

You can select a predefined time range or customize a range. You can customize a range by moving the start and end time selectors on the time-line or you can set a range in the Select a Time Range dialog box.

To select a predefined time range for the session, do the following:

Procedure

1. Click the Range drop-down list.
2. Select one of the following:
   - » A predefined range from the list.
   - » **Custom Range**.

3. If you selected **Custom Range**, then in the Select a Time Range dialog box:
   a. Select the date and time in the **From** section to define the start of the range.
   b. Select the date and time in the **To** section to define the end of the range.
   c. Click **Select**.

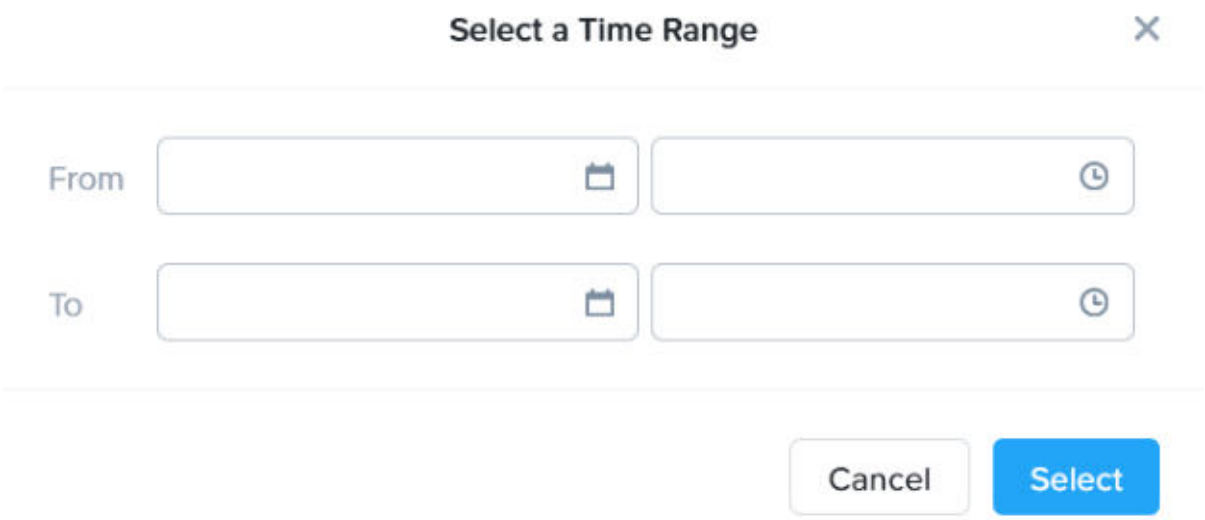

Figure 311:

If you selected **Custom Range**, you can also select the range dynamically on the time-line by sliding the range start and end sliders.

There is no minimum time range for **Custom Range**, and the maximum range is one year.

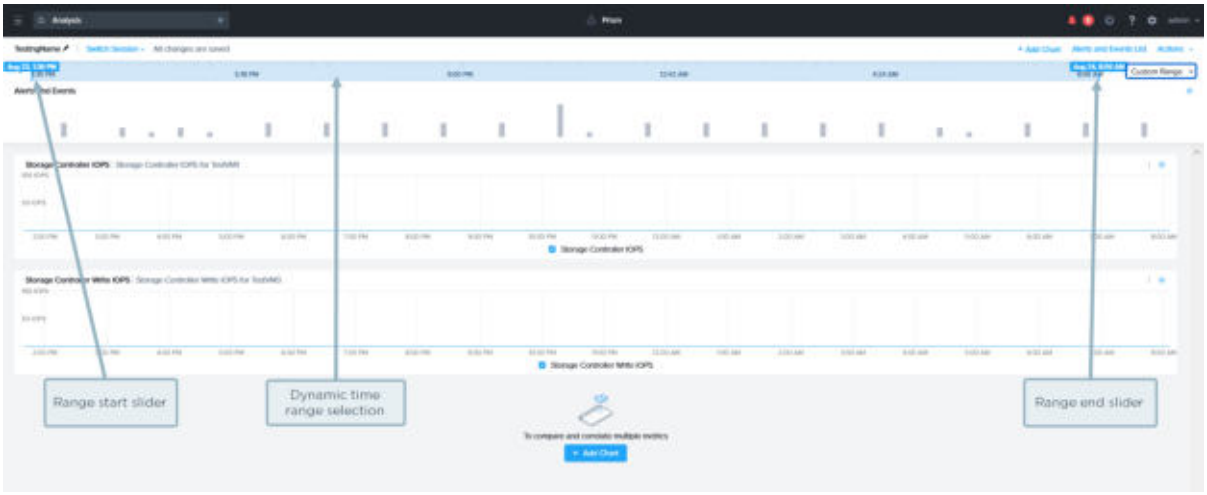

Figure 312: Time Range

Viewing Alerts and Events

About this task

You can view the alerts and events occurring during the selected time range in the following ways:

- Displayed on the **Alerts and Events** monitor available below the time-line.
- By clicking the **Alerts and Events List** option at the top right of the Analysis dashboard

Procedure

1. Click the stacked bar to see the details of the alerts and events in the Alerts and Events List panel.

   Point to a stacked bar on the monitor to see the hover image. The hover image provides more information about the alerts and events that occurred in the time range that the stacked bar depicts.

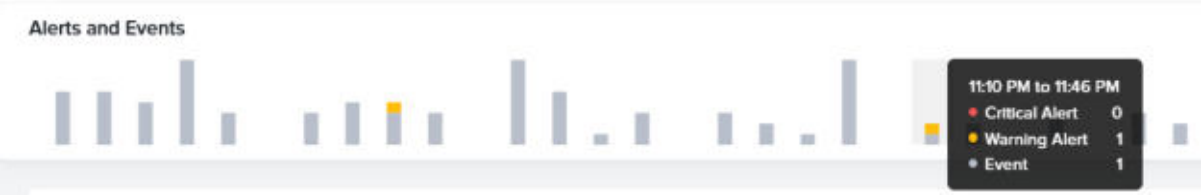

Figure 313: Stacked bar - Alerts and Events Information

2. Click the **Alerts and Events List** option at the top right of the Analysis dashboard.

   The Alerts and Events List panel displays the details of the alerts and events that occurred during the time range selected.

   See the Table 226: Alerts and Events List Panel Elements on page 490 table for more information.

   - The Alerts tab provides the details of the alerts. The alerts can be filtered.

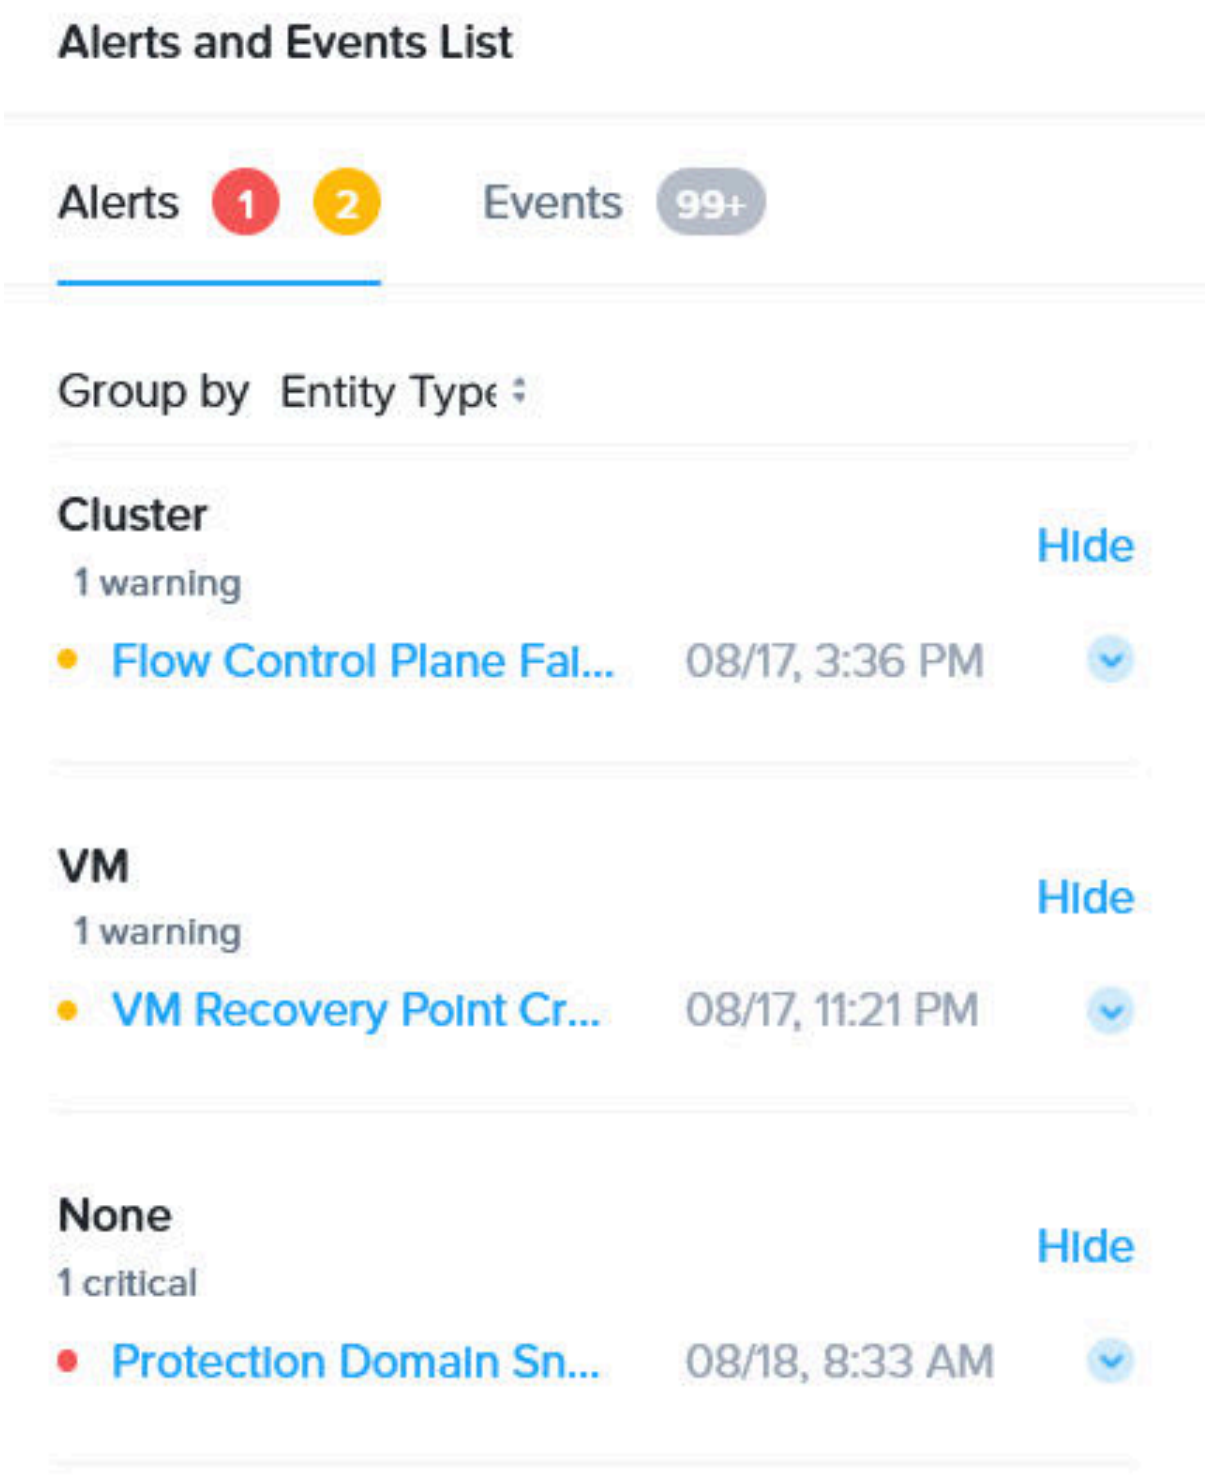

Figure 314: Alerts tab

Use the filters provided in the **Group by** drop down list to filter the category of alerts.

- The Events tab provides the details of the events. The events cannot be filtered.

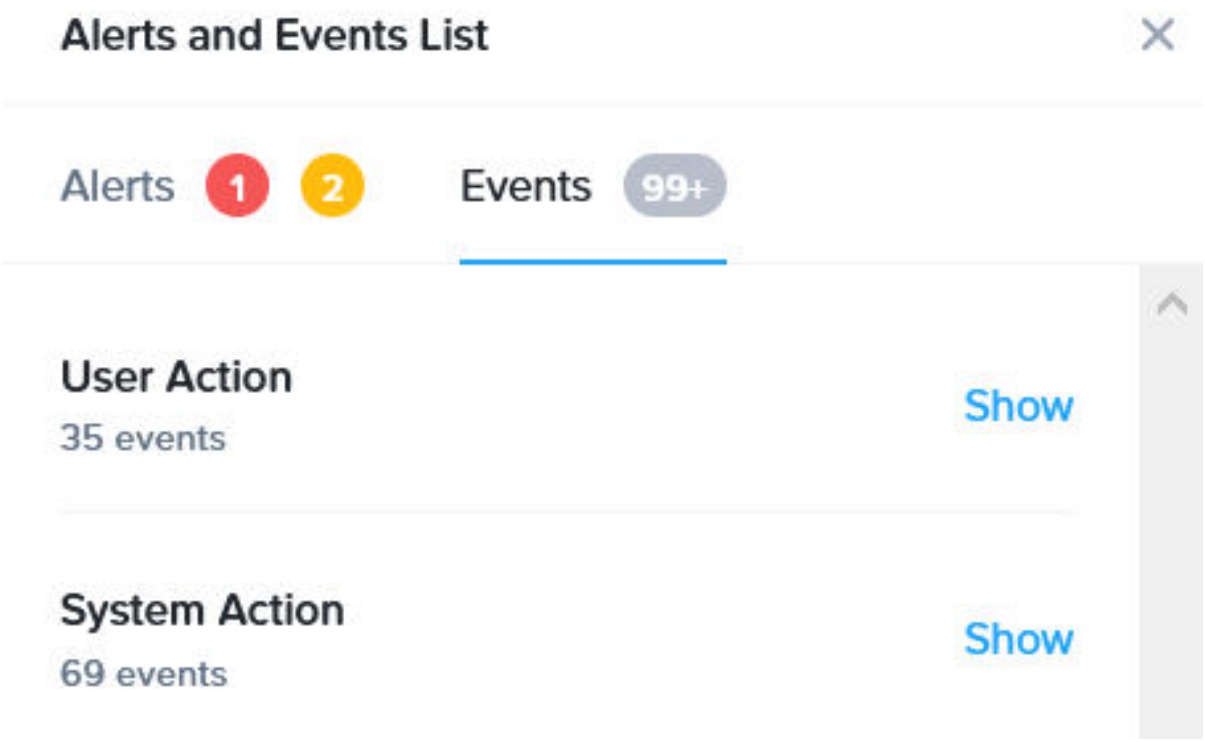

Figure 315: Events tab

Click **Show** to expand the category of alerts or events.

3. Click an alert or event to open the alert or event details page.

Chart Management

On the Analysis dashboard, you can:

Add charts	Adding Charts on page 497
Edit a chart	Editing a Chart on page 499
Delete a chart	Deleting a Chart on page 500
Clone a chart	Cloning a Chart on page 501
Export a chart as a JSON or CSV file	Exporting a Chart on page 501

Click the options button (three vertical dots) on the top left corner of a chart to see the aforesaid options, except **Add Chart**. The following sample figure shows the options.

Figure 316: Sample Chart - Chart Options

Adding Charts

This is a Nutanix task.

About this task

You can add charts in a session by clicking:

- The **Add Chart** option at the top left of the dashboard.
- The **Add Chart** button at the bottom center of the dashboard.

The **Add Chart** option and button opens the Add Chart dialog box. You can add multiple charts by opening the Add Chart dialog box once.

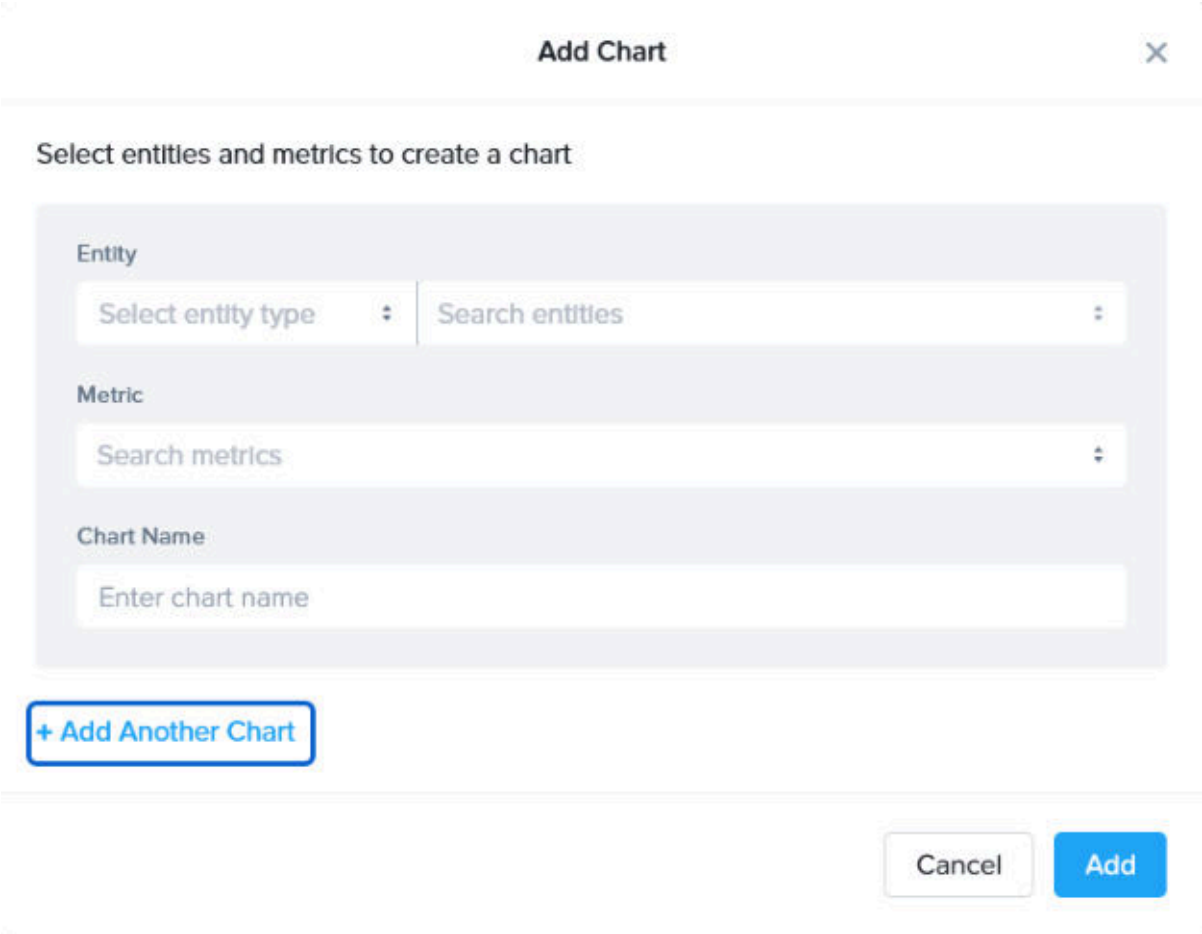

Figure 317: Add Chart dialog box

You can define a single chart to monitor a single metric for multiple entities of the same entity type. You can define a chart to monitor multiple metrics for a single entity.

> Note:
>
> You cannot define a single chart to monitor multiple metrics for multiple entities.

Procedure

1. Click one of the following:
   - » The **Add Chart** option at the top left of the dashboard.
   - » The **Add Chart** button at the bottom center of the dashboard.
2. In the Add Chart dialog box, do the following:
   a. Under **Entities**, select an entity type.
   b. Select the entities by selecting the check boxes for the entities in the list of entities.
   c. Under **Metric**, select the metrics you want by selecting the check boxes for the metrics in the list of metrics.

      See Chart Metrics on page 502 for the list of metrics you can select.

      > Note:
      >
      > In one chart, select metrics that use the same unit of measurement. Add a metric with a different unit of measurement to a separate chart.
      >
      > See the step to add another chart below.
      >
      > When adding a metric for multiple entities hosted on different hypervisors, select a metric only if it is supported on the hypervisors that are hosting the selected entities. Prism validates the list of available metrics provided for the **Metric** field. Prism removes from the list any metric that is not supported by all the hosting hypervisors for the selected entities. For example, for two entities hosted on ESXi and AHV hypervisors respectively, the list includes only the metrics that are supported on both ESXi and AHV.

   d. Provide an appropriate **Chart Name** with appropriate keywords.

      Prism provides a default name in the **Chart Name** field based on the entities and metrics you select in the preceding steps.

   e. To add another chart, click **Add Another Chart**.

      Follow the preceding steps a to d, to define the chart.

      Click the **Add Another Chart** option multiple times in the Add Chart dialog box. Each time you click the **Add Another Chart** option, provide the necessary details (steps a to d).

   f. Click **Add** after adding all the charts.

3. You can also add any chart or metric from anywhere in Prism Central. For example, you can add a chart from the Metrics tab of a Disk details page. On the chart or metric widget, do one of the following:

   » Click **Actions** > **Add to Analysis** to add the metric to the current session in the Analysis dashboard.

   » Click **Actions** > **Select Analysis Session** to add the metric to any available session or a new session.

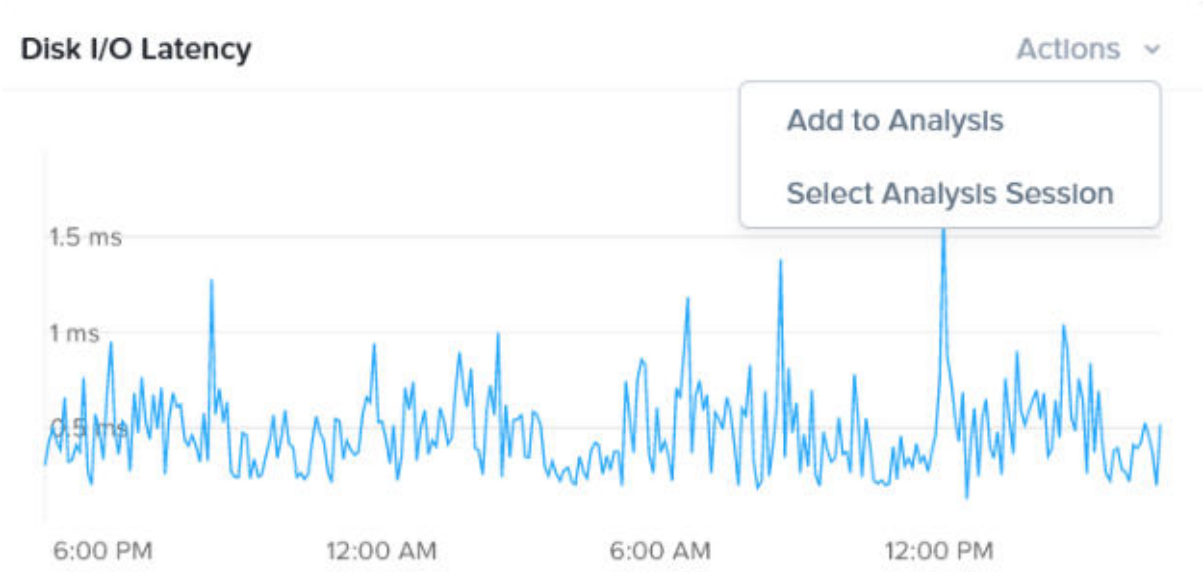

Figure 318: Example of chart or metric widget

Editing a Chart

About this task

You can edit a chart in the Analysis dashboard to modify the entity or the metric you are monitoring with the chart. You can also modify the name of the chart.

To edit a chart, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.

2. Click **Edit**.

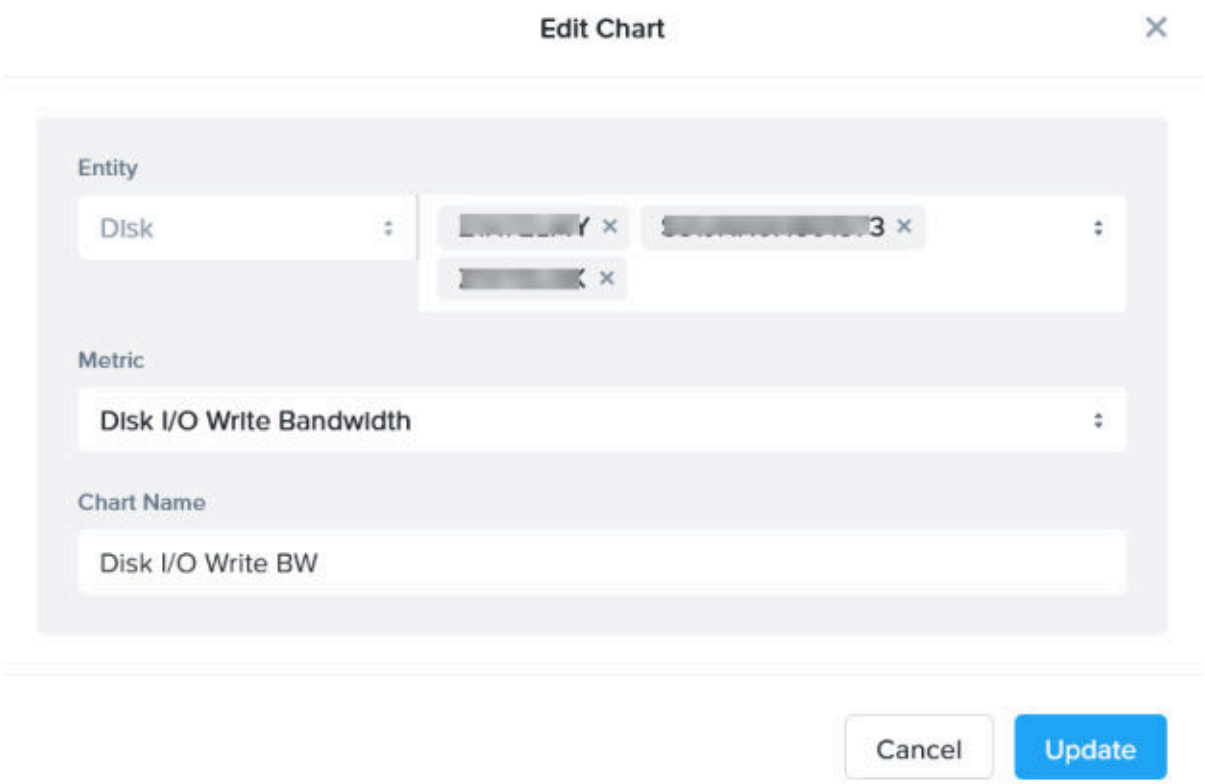

Figure 319: Edit Chart dialog box

3. In the **Edit Chart** dialog box, make the necessary changes in the fields.
4. Click **Update**.
   The changes are displayed in the chart on the Analysis dashboard.

Deleting a Chart

About this task

To delete a chart, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.

2. Click **Delete**.

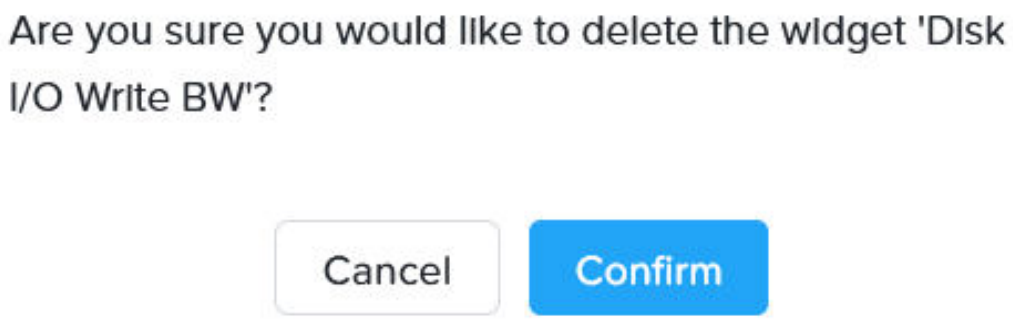

Figure 320: Delete confirmation dialog box

3. Do one of the following:
   - » Click **Confirm** to delete the chart widget.
   - » Click **Cancel** to return to the Analysis dashboard without deleting the chart.

Cloning a Chart

About this task

To clone a chart, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.
2. Click **Clone**.
   The cloned chart is the last chart in the row of charts in the Analysis dashboard.

Exporting a Chart

About this task

You can export a chart as a CSV or a JSON file.

To export a chart as CSV or JSON file, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.

2. Click one of the following:
   - » **Export CSV** to save as a CSV file.
   - » **Export JSON** to save as a JSON file.
3. Select **Save File**.
4. Browse to the location you want to save the file in the local folder and click **Save**.

Chart Metrics

These metrics can be added to charts.

> Note: The mapping between a metric and an entity type is hypervisor dependent.

Metric	Entity Type(s)	Description
Content Cache Hit Rate (%)	• Host • Cluster	Content cache hits over all lookups. ID: `CONTENT_CACHE_HIT_PPM`
Content Cache Hits	• Host • Cluster	Number of hits on the content cache. ID: `CONTENT_CACHE_NUM_HITS`
Content Cache Logical Memory Usage	• Host • Cluster	Logical memory (in bytes) used to cache data without deduplication. ID: `CONTENT_CACHE_LOGICAL_MEMORY_USAGE_BYTES`
Content Cache Logical SSD Usage	• Host • Cluster	Logical SSD memory (in bytes) used to cache data without deduplication. ID: `CONTENT_CACHE_LOGICAL_SSD_USAGE_BYTES`
Content Cache Lookups	• Host • Cluster	Number of lookups on the content cache. ID: `CONTENT_CACHE_NUM_LOOKUPS`
Content Cache Memory Saved	• Host • Cluster	Memory (in bytes) saved due to content cache deduplication. ID: `CONTENT_CACHE_SAVED_MEMORY_USAGE_BYTES`
Content Cache Physical Memory Usage	• Host • Cluster	Real memory (in bytes) used to cache data by the content cache. ID: `CONTENT_CACHE_PHYSICAL_MEMORY_USAGE_BYTES`
Content Cache Reference Count	• Host • Cluster	Average number of content cache references. ID: `CONTENT_CACHE_NUM_DEDUP_REF_COUNT_PPH`

Metric	Entity Type(s)	Description
Content Cache SSD Usage	• Host • Cluster	Real SSD usage (in bytes) used to cache data by the content cache. ID: CONTENT_CACHE_PHYSICAL_SSD_USAGE_BYTES
Content Cache SSD Usage Saved	• Host • Cluster	SSD usage (in bytes) saved due to content cache deduplication. ID: CONTENT_CACHE_SAVED_SSD_USAGE_BYTES
Deduplication Fingerprints Cleared	• Host • Cluster	Number of written bytes for which fingerprints have been cleared. ID: DEDUP_FINGERPRINT_CLEARED_BYTES
Deduplication Fingerprints Written	• Host • Cluster	Number of written bytes for which fingerprints have been added. ID: DEDUP_FINGERPRINT_ADDED_BYTES
Disk I/O Bandwidth	• Host • Cluster • Disk • Storage Pool	Data transferred per second in KB/second from disk. ID: STATS_BANDWIDTH
Disk I/O Bandwidth - Read	• Host • Cluster • Disk • Storage Pool	Read data transferred per second in KB/second from disk. ID: STATS_READ_BANDWIDTH
Disk I/O Bandwidth - Write	• Host • Cluster • Disk • Storage Pool	Write data transferred per second in KB/second from disk. ID: STATS_WRITE_BANDWIDTH
Disk I/O Latency	• Host • Cluster • Disk • Storage Pool	I/O latency in milliseconds from disk. ID: STATS_AVG_IO_LATENCY

Metric	Entity Type(s)	Description
Disk IOPS	• Host • Cluster • Disk • Storage Pool	Input/Output operations per second from disk. ID: STATS_NUM_IOPS
Disk IOPS - Read	• Host • Cluster • Disk • Storage Pool	Input/Output read operations per second from disk. ID: STATS_NUM_READ_IOPS
Disk IOPS - Write	• Host • Cluster • Disk • Storage Pool	Input/Output write operations per second from disk. ID: STATS_NUM_WRITE_IOPS
GPU Framebuffer Usage	Virtual Machine	Framebuffer usage in percentage. ID: FRAMEBUFFER_USAGE_PPM Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU Usage	Virtual Machine	GPU compute usage in percentage. ID: GPU_USAGE_PPM Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU video decoder Usage	Virtual Machine	GPU video decoder usage in percentage. ID: DECODER_USAGE_PPM Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU video encoder usage	Virtual Machine	GPU video encoder usage in percentage ID: ENCODER_USAGE_PPM Note: The Virtual Machine entity is applicable to AHV hypervisor.
Hypervisor CPU Ready Time (%)	Virtual Machine	Percentage of time that the virtual machine was ready, but could not get scheduled to run. ID: STATS_HYP_CPU_READY_TIME

Metric	Entity Type(s)	Description
Hypervisor CPU Usage (%)	• Host • Cluster • Virtual Machine	Percent of CPU used by the hypervisor. ID: STATS_HYP_CPU_USAGE
Hypervisor I/O Bandwidth	• Host • Cluster • Virtual Machine	Data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_BANDWIDTH Note: Virtual Machine entity is applicable for ESXi. hypervisor
Hypervisor I/O Bandwidth - Read	• Host • Cluster • Virtual Machine	Read data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_READ_BANDWIDTH Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Bandwidth - Write	• Host • Cluster • Virtual Machine	Write data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_WRITE_BANDWIDTH Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency	• Host • Cluster • Virtual Machine	I/O latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_IO_LATENCY Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency - Read	• Host • Cluster	I/O read latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_READ_IO_LATENCY Note: Cluster entity is applicable only for AHV hypervisor.
Hypervisor I/O Latency - Write	• Host • Cluster	I/O write latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_WRITE_IO_LATENCY Note: Cluster entity is applicable only for AHV hypervisor.

Metric	Entity Type(s)	Description
Hypervisor IOPS	• Host • Cluster • Virtual Machine	Input/Output operations per second from Hypervisor. ID: STATS_HYP_NUM_IOPS Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor IOPS - Read	• Host • Cluster • Virtual Machine	Input/Output read operations per second from Hypervisor. ID: STATS_HYP_NUM_READ_IOPS Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor IOPS - Write	• Host • Cluster • Virtual Machine	Input/Output write operations per second from Hypervisor. ID: STATS_HYP_NUM_WRITE_IOPS Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor Memory Usage (%)	• Host • Cluster • Virtual Machine	Percent of memory used by the hypervisor. ID: STATS_HYP_MEMORY_USAGE
Logical Usage	Storage Container	Logical usage of storage (physical usage divided by replication factor). ID: STATS_UNTRANSFORMED_USAGE
Memory Usage(%)	• Virtual Machine	Percent of memory used by the hypervisor. Consolidated guest memory usage in percentage. ID: MEMORY_USAGE_PPM

Metric	Entity Type(s)	Description
Network Rx Bytes	Virtual Machine	Write data transferred in KB/second from the Storage Controller. ID: HYPERVISOR_NUM_RECEIVED_BYTES Note: Virtual Machine entity is applicable for AHV hypervisor.
Network Tx Bytes	Virtual Machine	Write data transferred per second in KB/second. ID: HYPERVISOR_NUM_TRANSMITTED_BYTES Note: Virtual Machine entity is applicable for AHV hypervisor.
Physical Usage	• Host • Cluster • Disk • Storage Pool • Storage Container	Actual usage of storage. ID: STATS_TRANSFORMED_USAGE
Read IOPS (%)	• Host • Cluster • Disk • Storage Pool	Percent of IOPS that are reads. ID: STATS_READ_IO_PPM
Replication Bandwidth - Received	• Cluster • Remote Site • Protection Domain • Replication Link	Replication data received per second in KB/second ID: STATS_REP_BW_RECEIVED
Replication Bandwidth - Transmitted	• Cluster • Remote Site • Protection Domain • Replication Link	Replication data transferred per second in KB/second ID: STATS_REP_BW_TRANSFERRED

Metric	Entity Type(s)	Description
Replication Bytes - Received	• Cluster • Remote Site • Protection Domain • Replication Link	Number of bytes received. ID: STATS_REP_NUM_RECEIVED_BYTES Note: Cluster entity is applicable for AHV hypervisor.
Replication Bytes - Total Received	Replication Link	Total number of bytes received. ID: STATS_REP_TOT_RECEIVED_BYTES
Replication Bytes - Total Transmitted	Replication Link	Total number of bytes transmitted. ID: STATS_REP_TOT_TRANSMITTED_BYTES
Replication Bytes - Transmitted	• Cluster • Remote Site • Protection Domain • Replication Link	Number of bytes transmitted. ID: STATS_REP_NUM_TRANSMITTED_BYTES Note: Cluster entity is applicable for AHV hypervisor.
Storage Controller Bandwidth	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_BANDWIDTH
Storage Controller Bandwidth - Read	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Read data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_READ_BANDWIDTH

Metric	Entity Type(s)	Description
Storage Controller Bandwidth - Write	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Write data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_WRITE_BANDWIDTH
Storage Controller IOPS	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_IOPS Note: The Host entity is applicable for AHV hypervisor only.
Storage Controller IOPS - Read	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output read operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_READ_IOPS Note: The Host entity is applicable for AHV hypervisor.
Storage Controller IOPS - Read (%)	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Percent of Storage Controller IOPS that are reads. ID: STATS_CONTROLLER_READ_IO_PPM Note: The Host entity is applicable for AHV hypervisor.

Metric	Entity Type(s)	Description
Storage Controller IOPS - Write	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output write operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_WRITE_IOPS
Storage Controller IOPS - Write (%)	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Percent of Storage Controller IOPS that are writes. ID: STATS_CONTROLLER_WRITE_IO_PPM
Storage Controller Latency	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	I/O latency in milliseconds from the Storage Controller. ID: STATS_CONTROLLER_AVG_IO_LATENCY Note: • Custer entity is applicable only for AHV hypervisor. • Virtual Machine entity is applicable for AHV and Hyper-V hypervisors.
Storage Controller Latency - Read	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Storage Controller read latency in milliseconds. ID: STATS_CONTROLLER_AVG_READ_IO_LATENCY Note: Custer entity is applicable for AHV hypervisor.

Metric	Entity Type(s)	Description
Storage Controller Latency - Write	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Storage Controller write latency in milliseconds. ID: STATS_CONTROLLER_AVG_WRITE_IO_LATENCY Note: Custer entity is applicable for AHV hypervisor.
Storage container own usage	Storage Container	Storage container's own usage + Reserved (not used). ID: NEW_CONTAINER_OWN_USAGE_LOGICAL
Swap In Rate	Virtual Machine	Rate of data being swapped in. ID: STATS_HYP_SWAP_IN_RATE Note: Virtual Machine entity is applicable for ESXi and Hyper-V hypervisors.
Swap Out Rate	Virtual Machine	Rate of data being swapped out. ID: STATS_HYP_SWAP_OUT_RATE Note: Virtual Machine entity is applicable for ESXi and Hyper-V hypervisors.
Virtual NIC bytes received packets with error.	Virtual Machine	Virtual NIC bytes received packets with error. STATS_NETWORK_ERROR_RECEIVED_PACKETS Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC bytes received rate.	Virtual Machine	Virtual NIC bytes received rate in kbps. STATS_NETWORK_RECEIVED_RATE Note: Virtual Machine entity is applicable only for ESXi hypervisor.
Virtual NIC bytes transmitted rate.	Virtual Machine	Virtual NIC bytes transmitted rate in kbps. STATS_NETWORK_TRANSMITTED_RATE Note: Virtual Machine entity is applicable for ESXi hypervisor.

Metric	Entity Type(s)	Description
Virtual NIC dropped transmitted packets.	Virtual Machine	Number of dropped transmitted packets by the Virtual NIC. STATS_NETWORK_DROPPED_TRANSMITTED_PACKETS Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC receive packets dropped.	Virtual Machine	Number of receive packets dropped by the Virtual NIC. STATS_NETWORK_DROPPED_RECEIVED_PACKETS Note: Virtual Machine entity is applicable for ESXi hypervisor.
Write IOPS (%)	• Host • Cluster • Disk • Storage Pool	Percent of IOPS that are writes. ID: STATS_WRITE_IO_PPM

Sessions Dashboard

You can view the Sessions dashboard by selecting **Switch Sessions** > **View All Sessions** in the Analysis dashboard. You can also view the Sessions dashboard by closing the current session. The Sessionsdashboard provides a list of sessions including the system-generated session. After an upgrade, the old sessions are also available in the Sessions dashboard.

> Note:
>
> Prism saves all session changes and new sessions automatically including **Session Name** changes, except changes made using specific options that open respective dialog boxes to facilitate the changes.

For example, when you edit a chart using the **Edit** option and the **Edit Chart** dialog box on the Analysisdashboard, you need to click the **Update** button to save the changes made in the dialog box. When you modify the session name or description in the Sessionsdashboard using the **Edit Session Details** action, you need to click the **Update** button to save the changes made in the dialog box. However, if you change the name of the session in the Analysis dashboard, then the change is automatically saved.

> Note:
>
> The updated Analysis dashboard and Sessions dashboard are available only with the Prism Pro and Prism Ultimate licenses. With the Prism Starter license, only the system-generated session is available.

The following figure is a sample view of the Sessions dashboard.

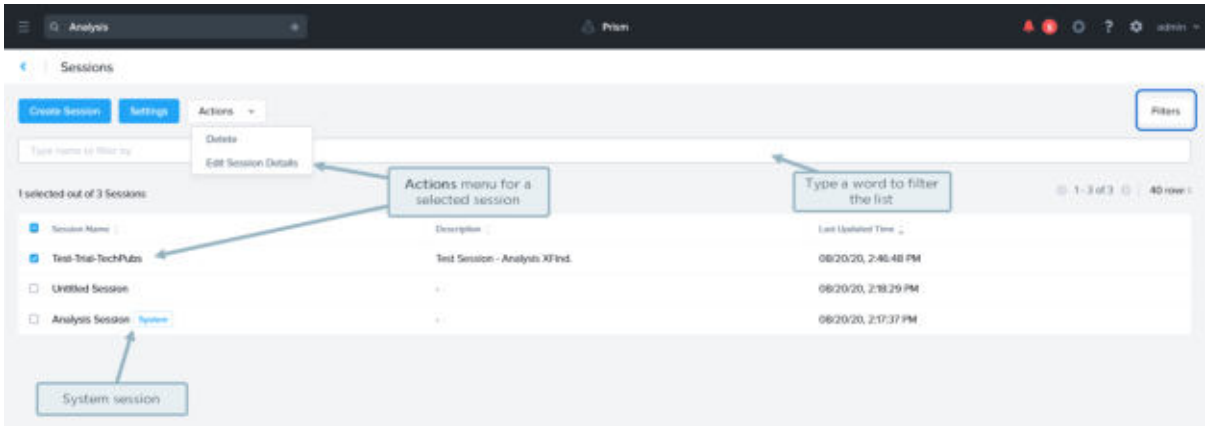

Figure 321: Sessions dashboard

Table 227: Sessions Dashboard Elements

Field	Allows you to...
Create Session	Create a new session.
Settings	Set the default session for the Analysis dashboard.
Actions	• **Delete** a session you select in the list. • **Edit Session Details** of a session you select in the list. The Actions menu is displayed only when you select a session in the list.

Field	Allows you to...
Filters	Filter the list of sessions. You can filter the list by: 1. **Session Name**: Filters on the session name. Select a condition from the pull-down list and enter a string in the field. It returns a list of sessions that satisfy the session name condition/string. Note: Adding a name for a session is mandatory. The default name is Untitled Session. Change the default name to a suitable name with appropriate keywords to facilitate filtering the sessions by session name. 2. **Description**: Filters on the description of the session. Select a condition from the pull-down list and enter a string in the field. It returns a list of sessions that satisfy the description condition/string. Note: Adding a description for a session is not mandatory. However, to facilitate filtering the sessions by description, enter a description with appropriate keywords for the sessions you create.
Type name to filter by	Allows you to filter the sessions by name. Type a word used in the session names that you want to filter to in this search box. Prism Central filters the list of all the sessions to provide a list of the sessions which have the word in their names. The filter is not case sensitive.
Session Name	Enter a name for the session you create or update.
Description	Provide a description for the session.
Last Updated Time	View the date and time when the session was last updated.

Viewing Sessions Dashboard

About this task

You can view the Sessions dashboard and manage the sessions.

To view the Sessions dashboard, do one of the following:

Procedure

1. Click **Switch Session** in the Analysis dashboard and select **View All Sessions**.

   This switches the view to the Sessions dashboard.

2. Click **Actions** in the Analysis dashboard and select **Close**.

   This closes the active session and switches to the Sessions dashboard.

Creating a New Session

About this task

You can create a new session from the Analysis dashboard or the Sessions dashboard.

To create a new session, do the following:

Procedure

1. In the Analysis dashboard, click **Switch Session** > **Create New Session**.

   > Note: To create a new session from the Sessions dashboard, click **Create Session** in the Sessions dashboard.

2. In the new session displayed in the Analysis dashboard, change the **Session Name** and add a **Description** help quickly identify the session.

   To change the **Session Name** and add a **Description**, see Editing Session Details on page 517.

What to do next

You can add charts to the new session.

Setting the Default Session View

About this task

You can set the default session view for the Analysis dashboard in the Sessions dashboard.

To set the default session view, do the following:

Procedure

1. Click **Settings** in the Sessions dashboard.

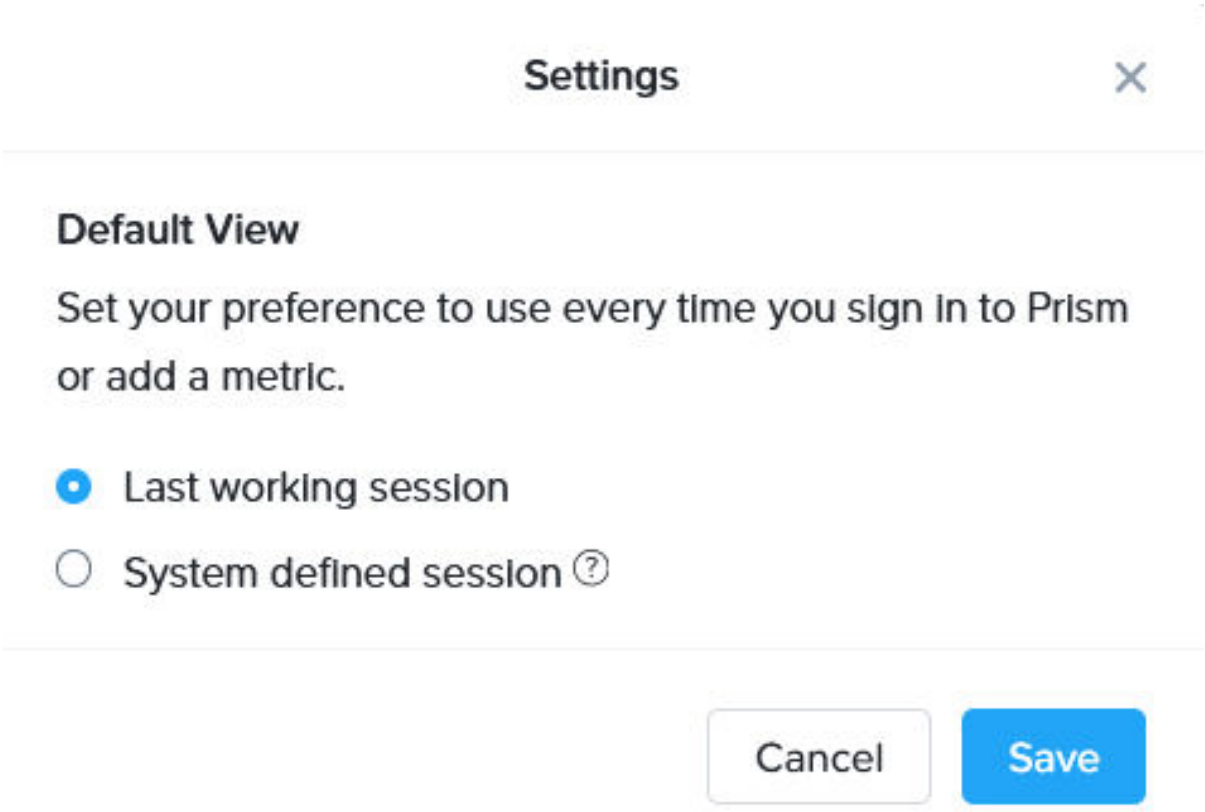

Figure 322: Settings dialog box

2. Select one of the following options:

   » **Last working session**

   This session is displayed by default in the Analysis dashboard. This is the last session that you worked on before signing out of Prism. This option is pre-selected.

   » **System defined session**

   This session is named **Analysis Session**. Prism creates and opens this session after installation of or upgrade to Prism Central 5.18 or later. Prism migrates and persists all the charts that you created before upgrading, to **Analysis Session** in Prism Central 5.18.

   Prism also opens this session in the Analysis dashboard after Prism services are restarted.

3. Click **Save**.

Deleting a Session

About this task

To delete a session in the Sessions dashboard, do the following:

Procedure

1. Select the session you want to delete.
2. Click **Actions** > **Delete**.
3. In the confirmation dialog box, do one of the following:
   - » Click **Confirm** to delete the session.
   - » Click **Cancel** to return to the Sessions dashboard without deleting the session.

Editing Session Details

About this task

To edit the session details like **Session Name** and **Description** of a session, do the following:

Procedure

1. Click **Actions** > **Edit Session Details**.

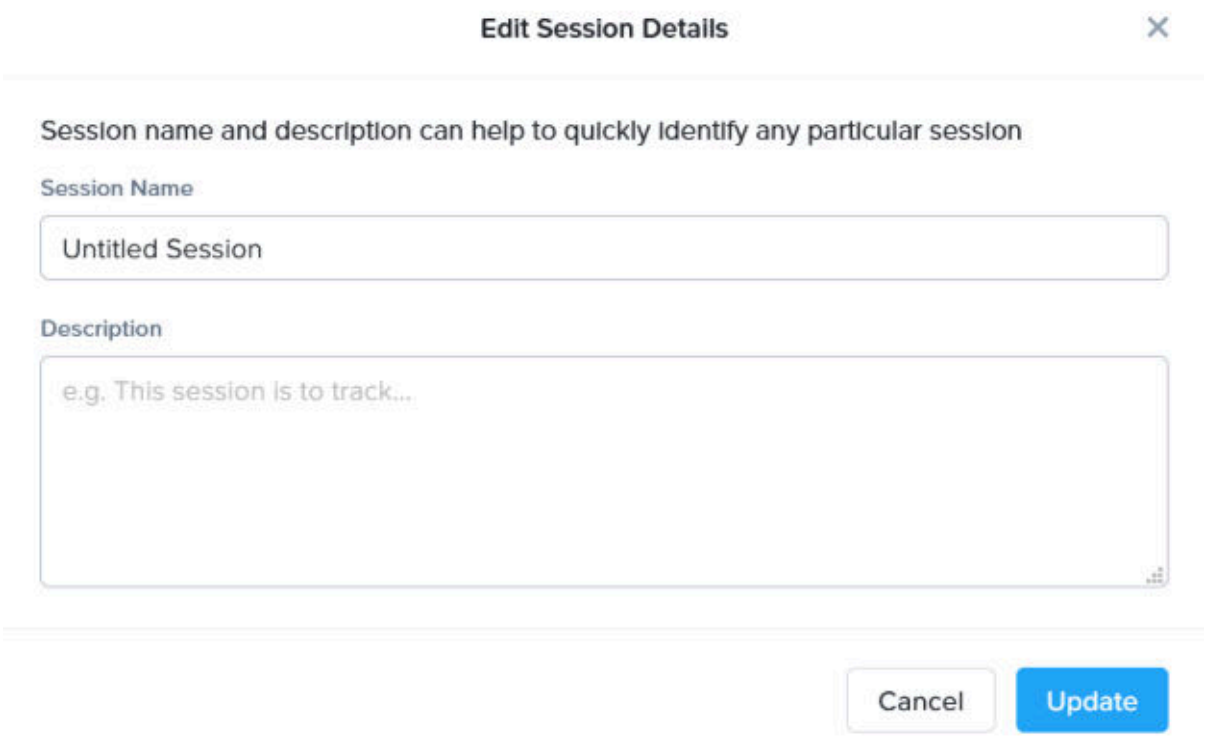

Figure 323: Edit Session Details dialog box

2. Provide appropriate **Session Name** and **Description** to help quickly identify the session.
3. Click **Update**.

Closing the Session

About this task

To close the current session in the Analysis dashboard, do the following:

Procedure

Click **Actions** > **Close**.

Switching Between Sessions

About this task

To switch between sessions, do the following:

Procedure

1. Click **Switch Session** in the Analysis dashboard.
2. From the drop-down list, select the session you want to view.
3. If you do not find the session you want to switch to in the list, do the following:
   a. Click **Switch Session** > **View All Sessions** to go to the Sessions dashboard.
   b. Select the session you want to open in the Analysis dashboard.

Application Discovery

Application discovery provides the capability to discover applications in a specified set of clusters, identify in which VM each application runs and what ports it uses, and publish that data to an internal database which is accessible through an external API call. You can monitor applications in both Nutanix and non-Nutanix (vCenter managed) clusters.

Application discovery sends data about the VM network traffic stats to a SaaS-like service in the cloud for analysis. The information sent to the cloud includes network flow data (source IP, destination IP, destination port, protocol, packet count, byte count) and VM data (cluster UUID, source VM name, UUID, VLAN ID, and IP addresses).

- To configure application discovery:
  - enable application discovery (see Enabling Application Discovery on page 526)
  - connect to the cloud-based discovery service (see Connecting to the Cloud on page 532)
  - add clusters to monitor (see Adding (Removing) Clusters to Monitor on page 533)
  - authenticate with vCenter for ESXi cluster access (see Configuring vCenter Authentication on page 534)
- To generate current findings at any time:
  - start an application discovery (see Discovering Applications on page 535)
  - publish data for identified applications (see Publishing Applications on page 536)
  - create policies for unidentified applications (see Creating Discovery Policies on page 536)

- To view discovery findings, see the application discovery dashboard (see Application Discovery View on page 519).

Application Discovery View

The application discovery dashboard allows you to view information about discovered applications and associated policies.

To access the application discovery dashboard, select **Operations** > **App Discovery** from the entities menu (see Entities Menu on page 13). The dashboard includes six tabs at the top (**Summary**, **Discovered Apps**, **Published Apps**, **Policies**, **Monitored Clusters**, and **Alerts**) with a display area below for the selected tab.

> Note:
>
> - This section describes the information and options that appear in the application discovery dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
> - If application discovery is not enabled, an enablement page appears instead of the dashboard (see Enabling Application Discovery on page 526).

Summary Tab

Clicking the **Summary** tab displays the following four widgets:

- **Discovered Highlights**: Displays the total number of application instances, the number of identified (known application to port mapping) and unidentified instances, the number of application types, and the number of ports.
- **Status & Settings**: Displays the number of monitored clusters, the number of discovery policies, and the status (active or inactive) of the discovery process, vCenter authentication, vCenter monitoring, and cloud connection. Click either **View Details** to see the vCenter authentication and monitoring details (see Configuring vCenter Authentication on page 534) or click**Manage** to see the cloud connection details (see Connecting to the Cloud on page 532). If vCenter monitoring is not configured, an **Add vCenter** link appears (instead of the **View Details** link); click the link to add clusters from any configured vCenter instances (see Adding (Removing) Clusters to Monitor on page 533).
- **Top Discovered Apps**: Lists the applications with the most instances.
- **Top Discovered Ports**: Lists the ports with the most instances.

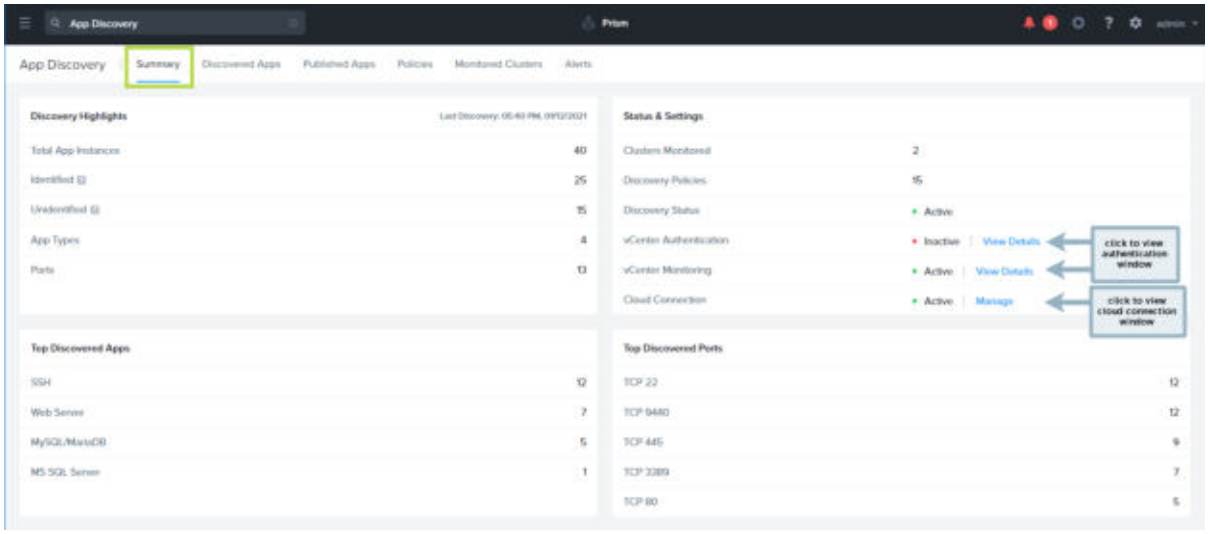

Figure 324: App Discovery Summary Tab

Discovered Apps Tab

Clicking the **Discovered Apps** tab, which appears by default when you first open the page, displays a list of the discovered applications.

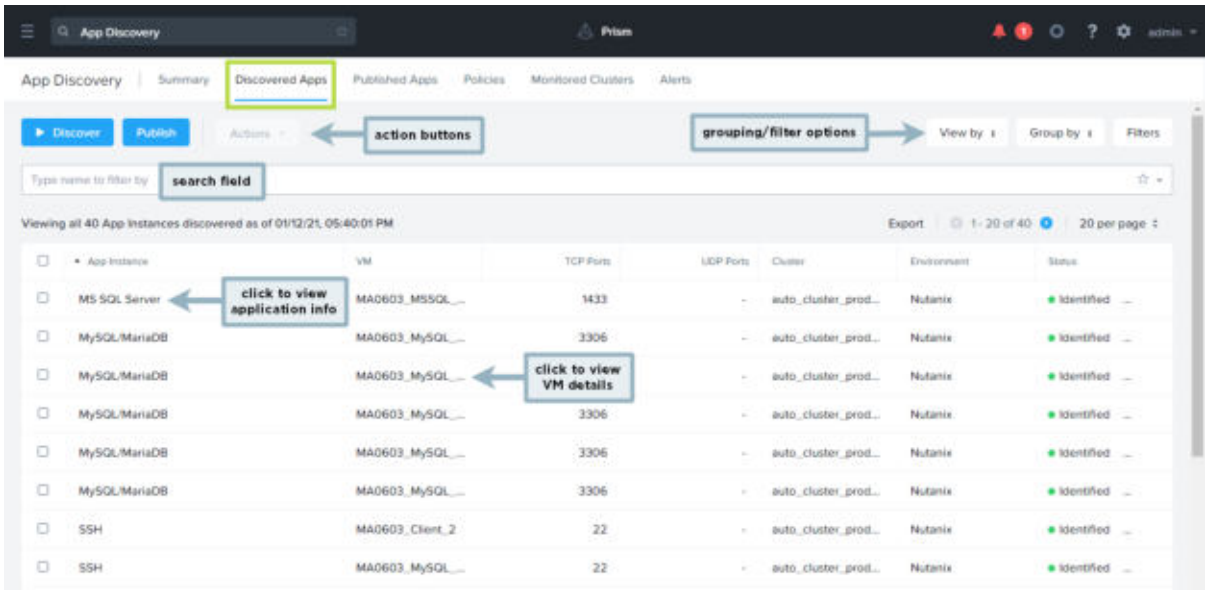

Figure 325: Discovered Apps Tab

The following table describes the fields that appear in the discovered applications list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 228: Discovered Apps Fields

Parameter	Description	Values
App Instance	Displays the application name for the instance. Click the name to display the Application Info window. The window displays the application name, VM name, TCP ports, and UDP ports.	(app name)
VM	Displays the VM name. Click the name to displays the details page for the VM (see VM Details View on page 90).	(VM name)
TCP Ports	Displays the TCP port(s) used by the application.	(port numbers)
UPD Ports	Displays the UDP port(s) used by the application.	(port numbers)
Cluster	Displays the name of the cluster in which the applications runs. Click the name to view the details for that cluster (see Cluster Details View on page 153).	(cluster name)
Environment	Displays whether the cluster resides in a Nutanix or non-Nutanix (vCenter managed) environment.	Nutanix, vCenter
Status	Indicates whether the instance is an identified application type and whether it is newly discovered ("new app" meaning discovered during the most recent pass) or discovered previously ("existing app" meaning discovered in an earlier pass). An "identified" application is one that is a known type. The discovery service can classify many applications, but for those it cannot identify, you can assign them to a type through a custom discovery policy (see Creating Discovery Policies on page 536).	Identified, Unidentified

You can filter the list either by entering a string in the search field or by opening the **Filter** pane (open by default) to select a filter option. The following table describes the filter options. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 229: Filter Pane Fields

Parameter	Description	Values
Status	Filters by whether the instance has been identified as a known application type.	Identified, Unidentified
App State	Filters by whether the instance is new.	New, Existing

Parameter	Description	Values
Ignored	Filters for ignored applications.	Show Ignored Apps (check box)
App Instance	Filters on the application instance name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.	(app instance name string)
Cluster	Filters on the cluster name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.	(cluster name string)
Environment	Filters based on whether the applications are in a Nutanix or non-Nutanix (vCenter) managed environment.	Nutanix, vCenter
VM	Filters on the VM name (same options as App Instance).	(VM name string)
TCP Ports	Filters on the TCP port numbers. Enter a range in the "**From <low> to <high>**" field.	([xx] to [yy] range)
UDP Ports	Filters on the UDP port numbers. Enter a range in the "**From <low> to <high>**" field.	([xx] to [yy] range)

To download the table contents in CSV format, click the **Export** link.

You can group the list in the following ways:

- The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.
- To group the list by application name, select **App Instance** from the **Group By** pull-down menu.

The discovered apps tab includes the following action buttons:

- Click the **Discover** button to start a discovery search (see Discovering Applications on page 535).
- Click the **Publish** button to publish application information (see Publishing Applications on page 536).
- An **Actions** menu appears when one or more instances are selected. The menu allows you to set the selected instances to [**Identify**|**Unidentify**] or [**Ignore**|**Unignore**] status (as appropriate). The available actions appear in bold; other actions are grayed out.

Published Apps Tab

Clicking the **Published Apps** tab displays a list of the published applications. A "published" application is one that is made available to the API. Publishing involves copying the data for the identified (and not ignored) applications into an internal database that you can access through an external API call. An "identified" application is one where the VM in which it runs is known and the used ports are mapped.

- The listed information for published applications is the same as for discovered applications, except there is no status field.

- The **Focus** , **Group**, and **Export** options are the same, but the **Filter** pane include just **App Instance** and **VM** options.

The published apps tab includes the following action buttons:

- Click the **Get API Link** button to copy the API link to your clipboard. The link is the access point to consume the published applications via the API.
- An **Actions** menu appears when one or more instances are selected; select **Delete** from the pull-down list to deleted the selected instances from the published list.

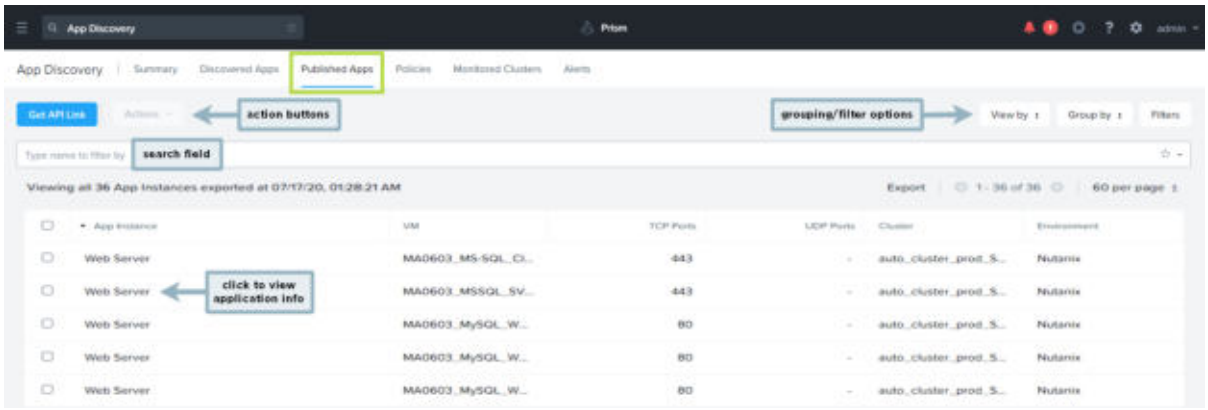

Figure 326: Published Apps Tab

Policies Tab

Clicking the **Policies** tab displays a list of application discovery policies. The policies define the application types.

> Note: Existing discovered applications are remapped whenever there is a policy action change, based on the existing VM traffic statistics.

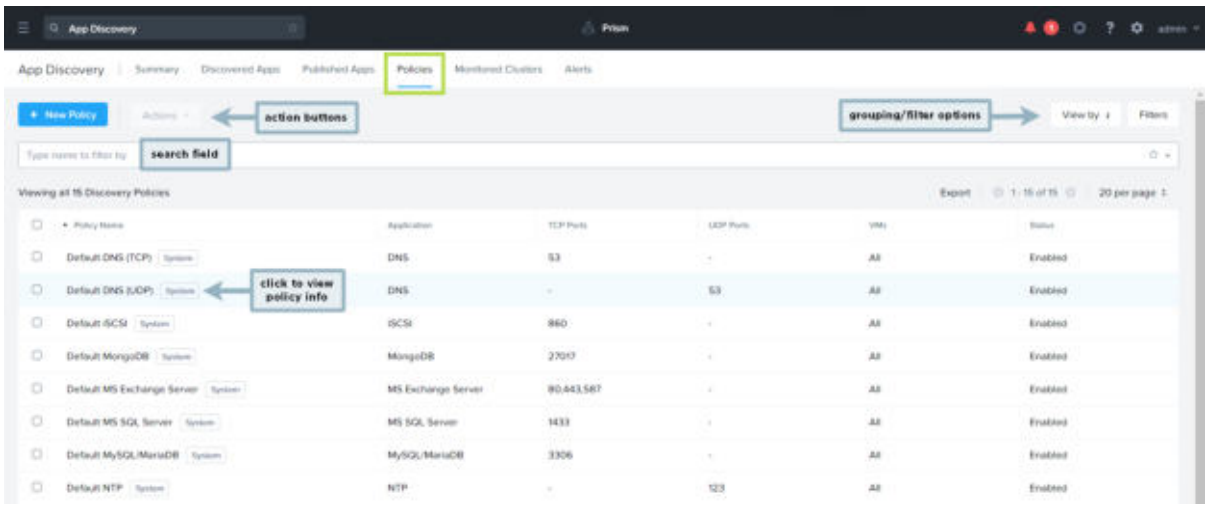

Figure 327: Policies Tab

The following table describes the fields that appear in the policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 230: Policies Fields

Parameter	Description	Values
Policy Name	Displays the name of the policy. Click the name to display the Edit Discovery Policy window for user-defined policies or the View Discovery Policy window for system-defined policies (see Modifying (Viewing) Discovery Policies on page 538).	(policy name)
Application	Displays the application type. Application discovery can identify a variety of common applications by default (referred to as system-defined applications) such as DNS, NTP, SSH, Web Server, MS Exchange Server, MS SQL Server, Oracle DB, SAP HANA, and so on. For user-defined applications, the user provides the name (see Creating Discovery Policies on page 536).	(application type)
TCP Ports	Displays the TCP port(s) used by the application.	(port numbers)
UPD Ports	Displays the UDP port(s) used by the application.	(port numbers)
VMs	Displays the scope of the policy, that is to which VMs the policy applies. A value of "All" indicates it applies to all VMs. One or more names indicates the policy applies to the named VMs only.	All, (VM names)
Status	Displays the current application operational status.	Enabled, Disabled

To download the table contents in CSV format, click the **Export** link.

You can filter the list either by entering a string in the search field or by opening the **Filter** pane to select one or both of the filter options.

- **Policy Name**: Check the box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field to filter the list for policies with that string in the name.
- **Status**: Check the **Enabled** or **Disabled** box to filter the list for enabled or disabled policies.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

The policies tab includes the following action buttons:

- Click the **New Policy** button to create a custom discovery policy (see Creating Discovery Policies on page 536).
- An **Actions** menu appears when one or more policies are selected. The menu allows you to **Enable**, **Disable**, or **Delete** the selected policies (as appropriate). You can only delete user-defined policies (not system policies).

Monitored Clusters Tab

Clicking the **Monitored Clusters** tab displays the list of clusters that application discovery monitors when it does a discovery search.

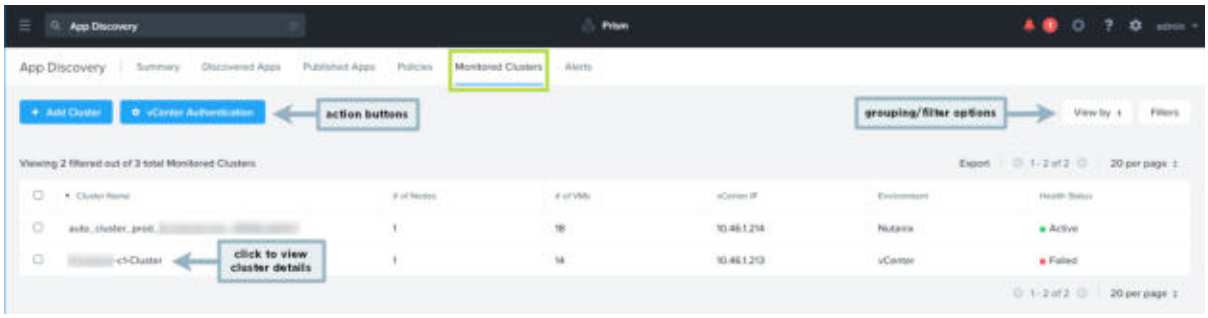

Figure 328: Monitored Clusters Tab

The following table describes the fields that appear in the monitored clusters list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 231: Monitored Clusters Fields

Parameter	Description	Values
Cluster Name	Displays the cluster name. Click the name to displays the details page for the cluster (see Cluster Details View on page 153).	(cluster name)
# of Nodes	Displays the number of nodes in the cluster.	(integer)
# of VMs	Displays the number of VMs in the cluster.	(integer)
vCenter IP	Displays the IP address of the vCenter for this cluster.	(IP address)
Environment	Displays whether the cluster resides in a Nutanix or non-Nutanix (vCenter managed) environment.	Nutanix, vCenter
Health Status	Displays the current VM network traffic monitoring status.	Active, Inactive

To download the table contents in CSV format, click the **Export** link.

You can filter the list either by entering a string in the search field or by opening the **Filter** pane to select a filter option. The following table describes the filter options. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 232: Filter Pane Fields

Parameter	Description	Values
vCenter IP	Filters on the vCenter IP address (same options as Cluster Name).	(IP address string)
Health Status	Filters on the selected health status. Check the box for each status to include.	Deploying, Deployed, Running, Stopped, Failed

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

The monitored clusters tab includes the following action buttons:

- Click the **Add Cluster** button to add a cluster for monitoring (see Adding (Removing) Clusters to Monitor on page 533).
- Click the **vCenter Authentication** button to authenticate application discovery on the monitored clusters with vCenter (see Configuring vCenter Authentication on page 534).
- A **Remove** button appears when one or more clusters are selected; click the button to remove the selected clusters from application discovery monitoring.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard (List tab), except it is filtered to display just application discovery-related alerts (see Alerts Summary View on page 210).

Enabling Application Discovery

Enabling application discovery (disabled by default) involves connecting to the discovery service in the cloud, specifying clusters to monitor, and authenticating to the vCenter instance(s) for the specified clusters.

About this task

When you select application discovery for the first time, a wizard appears to guide you through the setup process. To configure application discovery initially, do the following:

Before you begin

Complete the Prerequisites on page 531.

Procedure

1. Select **Operations** > **App Discovery** from the entities menu (see Entities Menu on page 13).

   If application discovery is not enabled, a "Welcome to App Discovery" screen appears.

2. In the "Welcome to App Discovery" screen, check the "I have increased the PC VM's memory and vCPUs" box (as specified in the prerequisites) and then click the **Enable App Discovery** button.

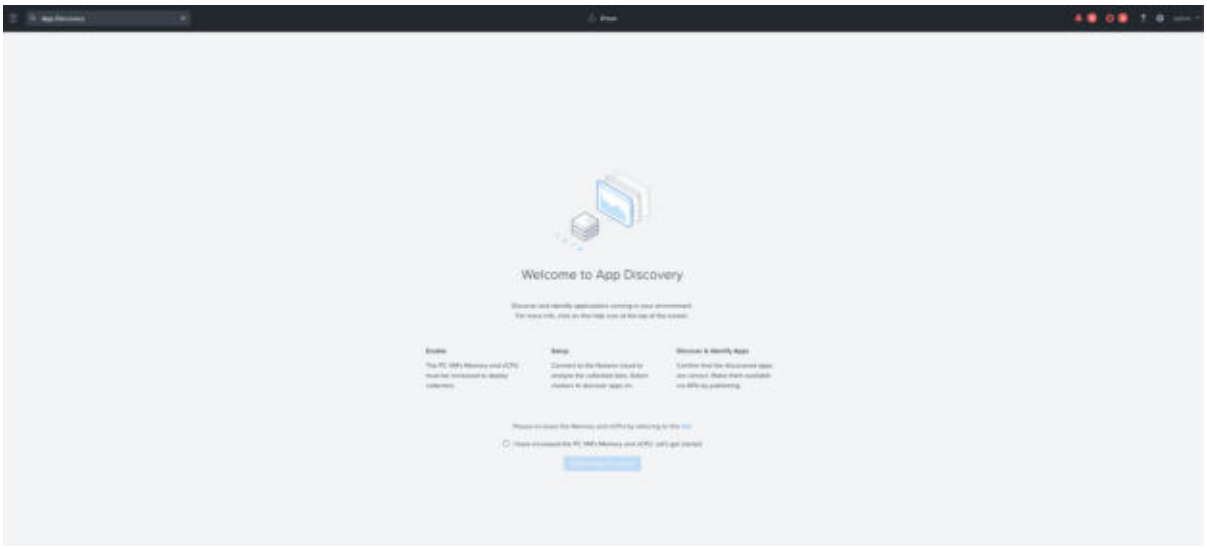

Figure 329: Welcome to App Discovery screen

3. In the "App Discovery Setup: Cloud Connection" screen, enter your API key and key authentication values in the indicated fields and then click the **Next** button.

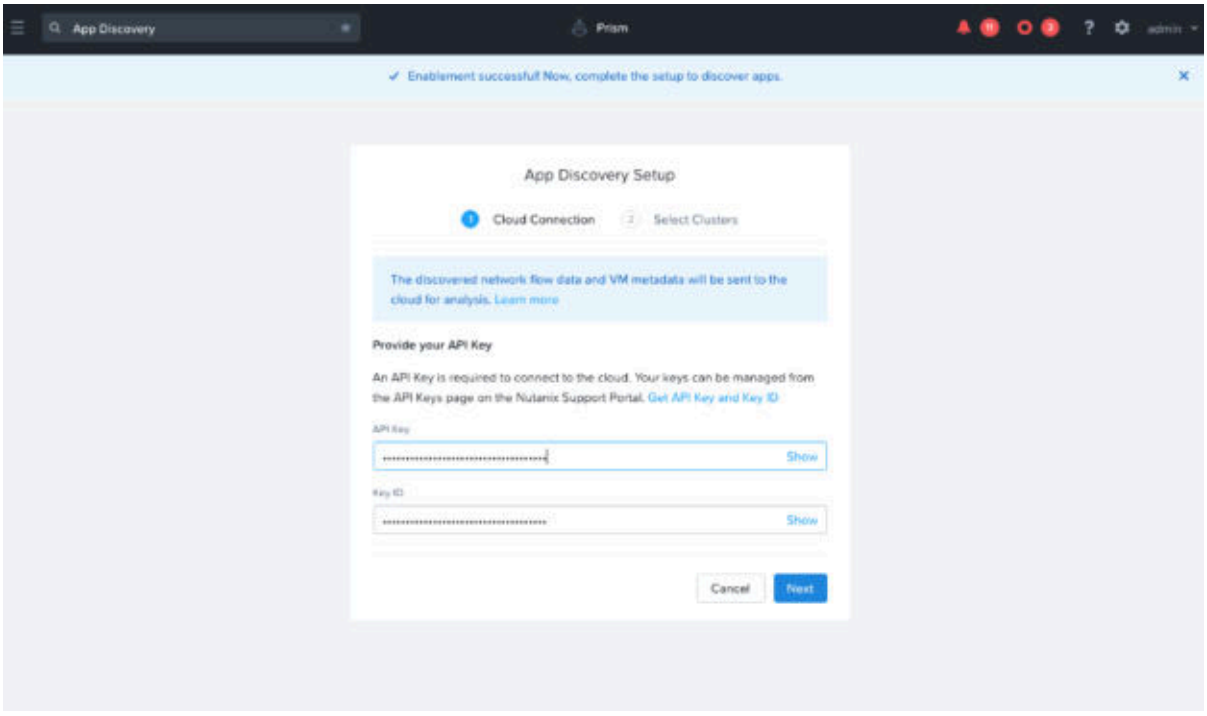

Figure 330: App Discovery Setup: Cloud Connection screen

4. In the "App Discovery Setup: Select Clusters" screen, do the following:

   » If the clusters to select are in a Nutanix environment, continue with this step.

   » If the clusters to select are in a non-Nutanix (vCenter managed) environment, click the **Configure vCenter Monitoring** link and then skip to step 5.

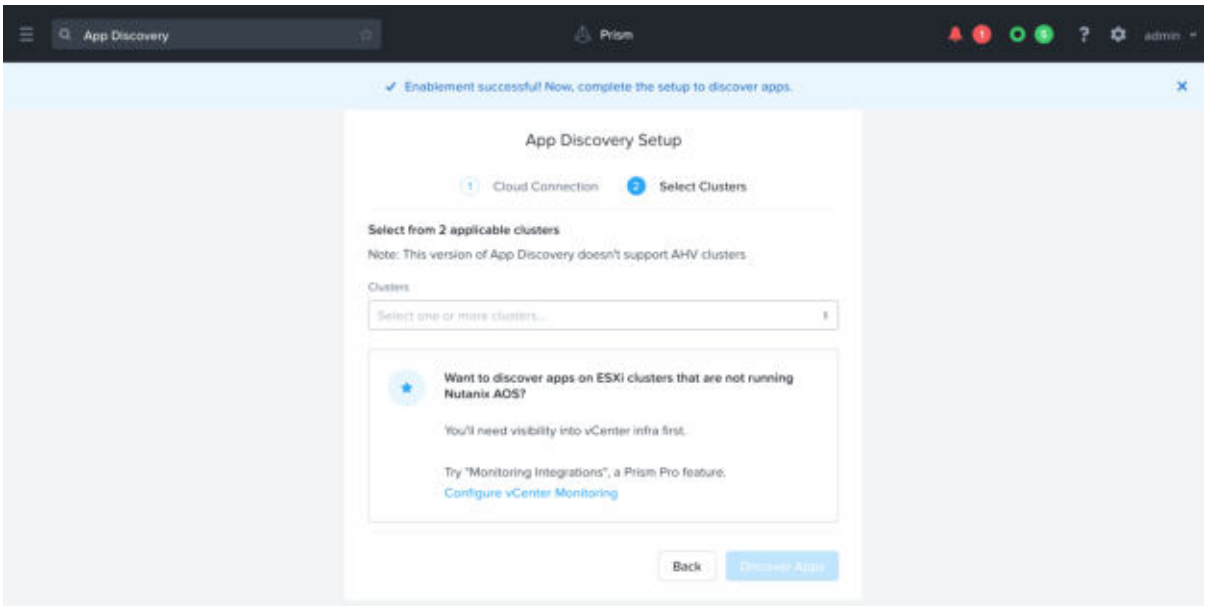

Figure 331: App Discovery Setup: Select Clusters (initial)

   a. Check the boxes for the target clusters from the pull-down list.

   b. Enter the vCenter administrator credentials (user name and password) in the indicated fields. The administrator must have write permission.

      An entry for each registered vCenter appears here (see Registering vCenter Server (Prism Central) on page 453).

   c. Click the **Discover Apps** button.

      This starts a discovery search. When it completes the App Discovery window and dashboard appear (see Application Discovery View on page 519). This completes the setup for a Nutanix environment.

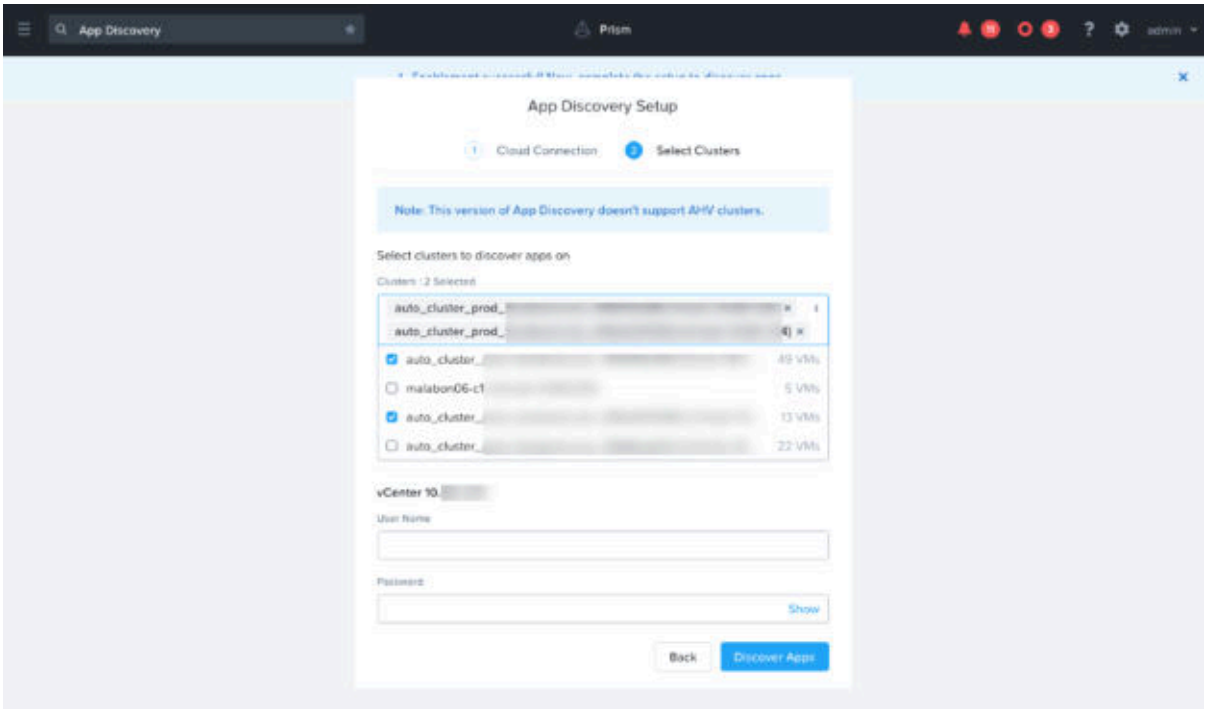

Figure 332: App Discovery Setup: Select Clusters

5. [vCenter only] Select **vCenter** in the **External Entity Type** field and then click the **Enable** button.

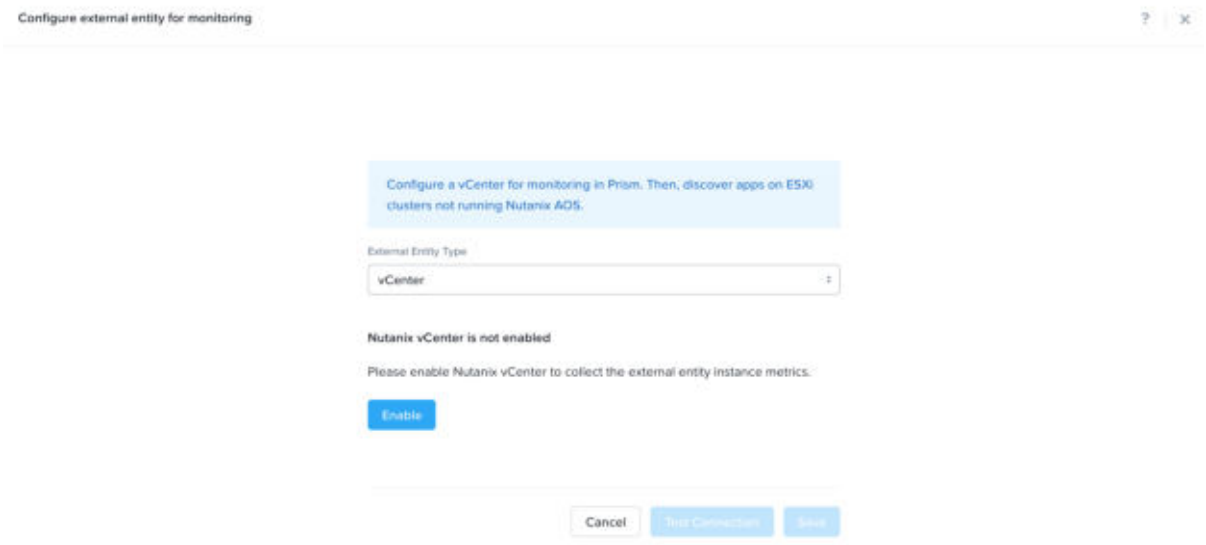

Figure 333: Enable vCenter Discovery

6. [vCenter only] Do the following in the indicated fields to configure a vCenter instance:

   a. **External Entity Type**: Select **vCenter** from the pull-down menu.

   b. **vCenter Host**: Enter the IP address of the vCenter host VM.

      If the IP address does not resolve automatically to the proper port number, enter the address in the form *ip_address:port#*. The standard vCenter server port is 443.

   c. **Username**: Enter the vCenter administrator user name.

   d. **Password**: Enter the vCenter administrator password.

   e. **SSL Verify**: Select the SSL mode to use (**Verify** or **Don't Verify**) when connecting to the target.

      This setting determines whether to verify (or not verify) the target's SSL certificate as part of the connection process.

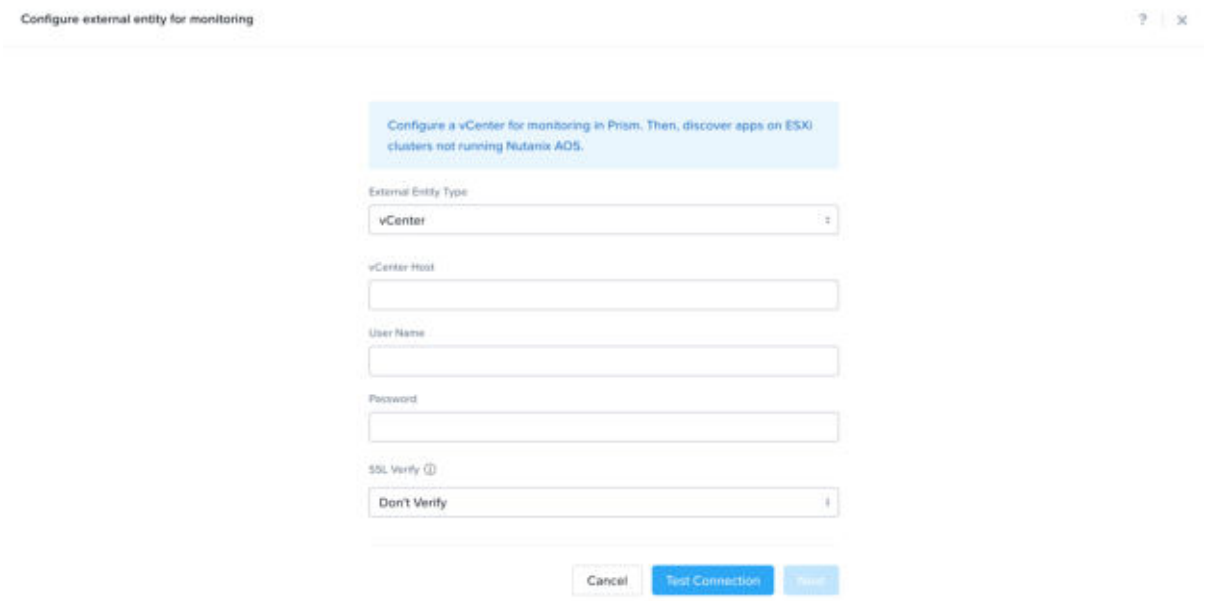

Figure 334: Configure vCenter Instance

7. [vCenter only] Click the **Test Connection** button to verify (test) the connection.

   A "Test connection success" message appears after a successful test. After a successful test, click the **Next** button.

8. [vCenter only] Select the clusters to monitor.

   A list of discovered clusters for the vCenter instance appears. Select (check the boxes for) the clusters you want to monitor from the list and then click the **Save** button. The instance

configuration and data collection begins. When it completes (which can take a few minutes), a "vCenter Configuration successful" message appears.

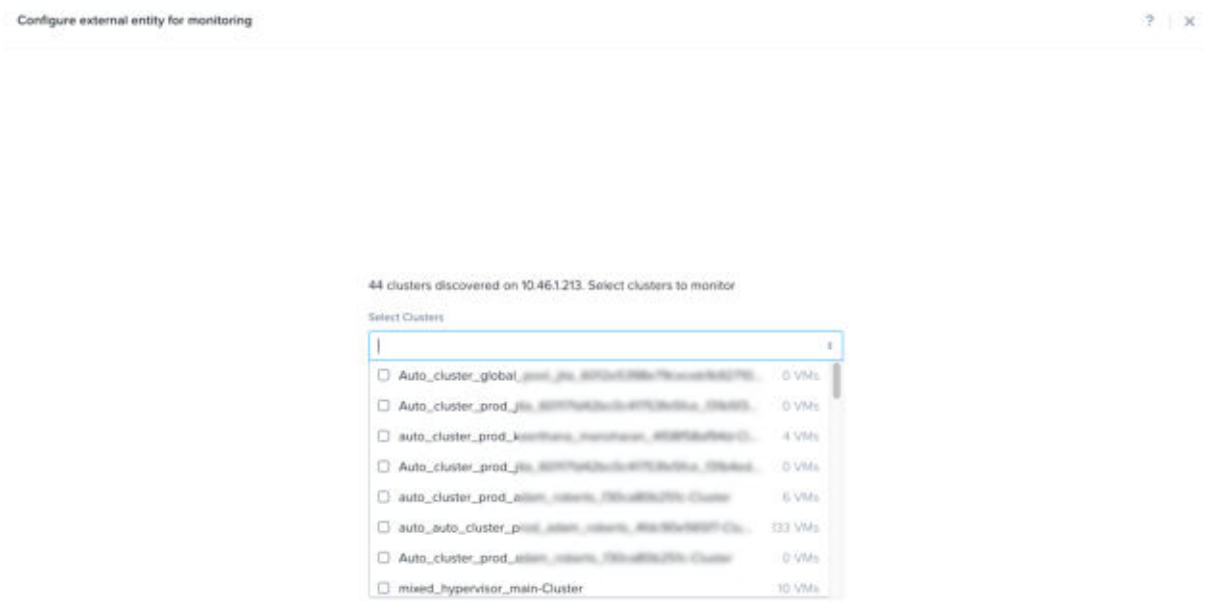

Figure 335: Select Clusters (vCenter)

9. [vCenter only] After the vCenter configuration completes, click the **Back to App Discovery** button.

   This takes you to the monitored clusters page and opens a cluster selection window. The vCenter clusters specified in step 8 will appear in the list for you to select.

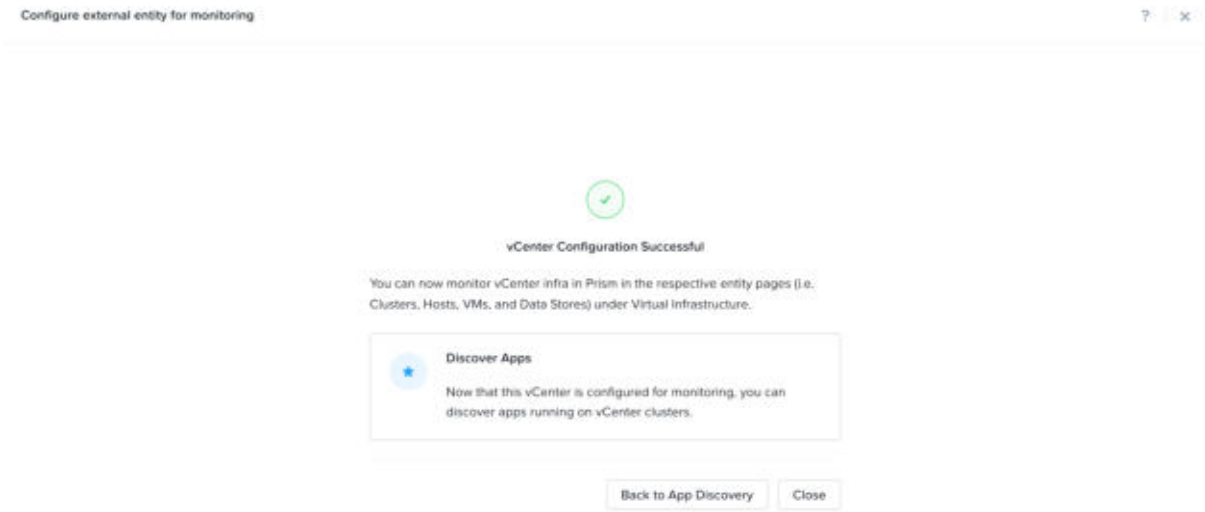

Figure 336: vCenter Configuration Complete

Prerequisites

Address the following before enabling application discovery:

Note: Application discovery is supported currently for ESXi clusters only. (An ESXi cluster that includes AHV storage only nodes is not supported.)

- Sufficient Prism Central VM resources. Application discovery is one of the Prism Ultimate license features that requires additional Prism Central VM resources (more memory and vCPUs). The requirements are listed in KB 8932. If your Prism Central VM resources are below the recommended amount, manually increase the memory and vCPUs to the recommended amount.
- Internet connection. An Internet connection is required to connect to the Nutanix cloud.
- Up and running firewall in Prism Central.
- VDS enabled on vCenter. This requires a vSphere version (Enterprise Plus license) that supports vSphere Distributed Switch (VDS) and netflow version 10.
- My Nutanix account that is either admin role or account admin.
- API key and key ID. Log on to the Nutanix support portal and get your API key information. See the *Nutanix Licensing Guide* for instructions on how to generate a new key or view information about an existing key.

  Each Prism Central instance requires a unique API key/ID pair. If you have multiple Prism Central instances (or a key was revoked for any reason), create additional API key/ID pairs as needed so you have a unique and valid pair for each Prism Central instance.

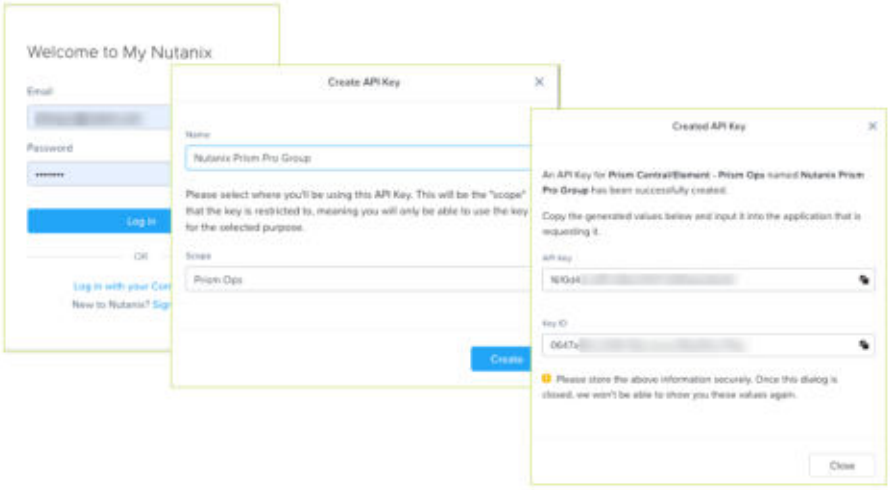

Figure 337: Create API Key Flow (example)

Connecting to the Cloud

Using application discovery requires connecting to a cloud-based discovery service.

About this task

Nutanix provides a SaaS-like service in the cloud that analyzes application data from the monitored clusters. The connection is established during initial setup (see Enabling Application Discovery on page 526), but to update the connection after setup, do the following:

Before you begin

Log on to the Nutanix support portal and get your API key information. See the *Nutanix Licensing Guide* or *License Manager Guide* for instructions on how to generate a new key or view information about an existing key.

Procedure

1. Go to the **Summary** tab of the application discovery dashboard and click **Manage** in the Cloud Connection line of the Status & Settings widget (see Application Discovery View on page 519).
2. In the Cloud Connection window, do the following in the indicated fields:
   a. **API Key**: Enter the API key value.
   b. **Key ID**: Enter the key identification number.
   c. Click the **Save** button.

   Only one API key/ID pair is allowed per Prism Central instance. If you have multiple Prism Central instances (or your initial key was revoked for any reason), create additional API key/ID pairs as needed (through the Nutanix support portal) so you have a unique and valid pair for each Prism Central instance.

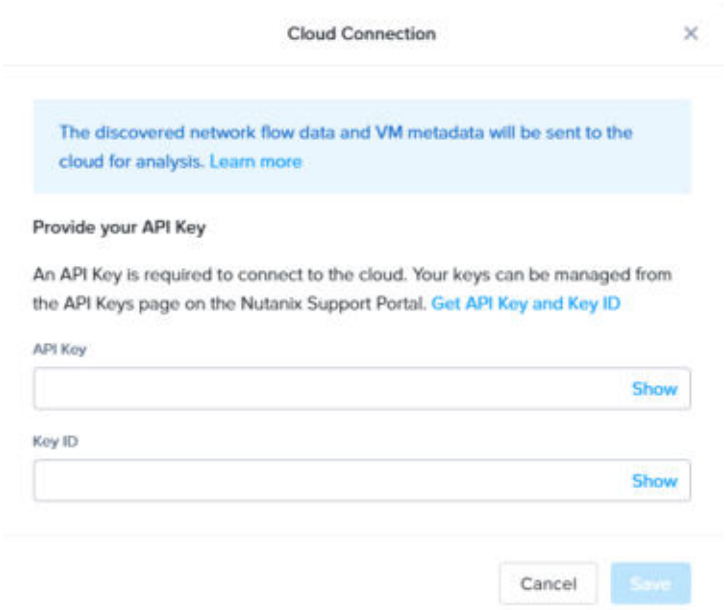

Figure 338: Cloud Connection Window

Adding (Removing) Clusters to Monitor

Using application discovery requires first specifying the clusters to monitor.

About this task

Specifying clusters to monitor is done during initial setup (see Enabling Application Discovery on page 526), but you can add or remove clusters after setup at any time.

Procedure

- To add a cluster to the monitored list, do the following:

  a. Go to the **Monitored Clusters** tab of the application discovery dashboard and click the **Add Cluster** button (see Application Discovery View on page 519).

  b. In the Add Cluster to App Discovery window, check the boxes for the target clusters from the pull-down list and then click the **Add** button.

  Network data collection begins immediate and may take some time to complete when multiple clusters are added.

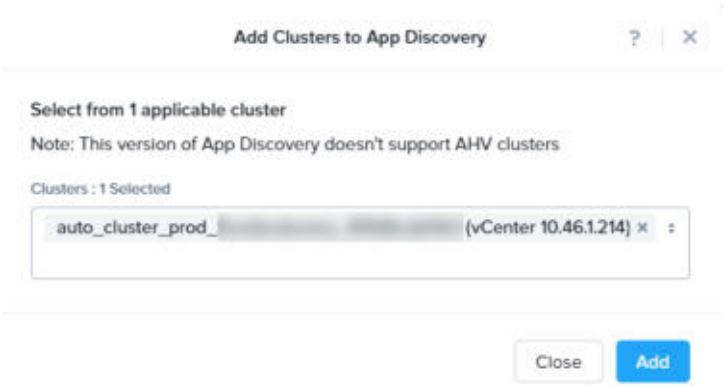

Figure 339: Add Clusters to App Discovery Window

  c. Run discovery to identify applications on the new clusters (see Discovering Applications on page 535).

- To remove a cluster from the monitored list, select the target cluster(s) in the list and then click the **Remove** button.

> Note: Removing a cluster stops application monitoring for that cluster. You can effectively disable application monitoring by removing all clusters, which means no data is sent between Prism Central and the Nutanix cloud.

Configuring vCenter Authentication

Using application discovery with ESXi clusters requires first authenticating with the associated vCenter.

About this task

Authenticating with vCenter is done during initial setup (see Enabling Application Discovery on page 526), but to update authentication after setup, do the following:

Procedure

1. Go to the **Monitored Clusters** tab of the application discovery dashboard and click the **vCenter Authentication** button (see Application Discovery View on page 519).

2. In the vCenter Authentication window, do the following in the indicated fields:

   a. **User Name**: Enter the user name for a vCenter administrator or any user with write permission to the clusters.

   b. **Password**: Enter the user password.

   c. Click the **Save** button.

   An entry for each registered vCenter appears here (see Registering vCenter Server (Prism Central) on page 453). The authentication and monitoring status for each vCenter is displayed. If a vCenter is not monitored currently, you can enable monitoring by clicking the **Configure** button (see Configuring Application Monitoring on page 553).

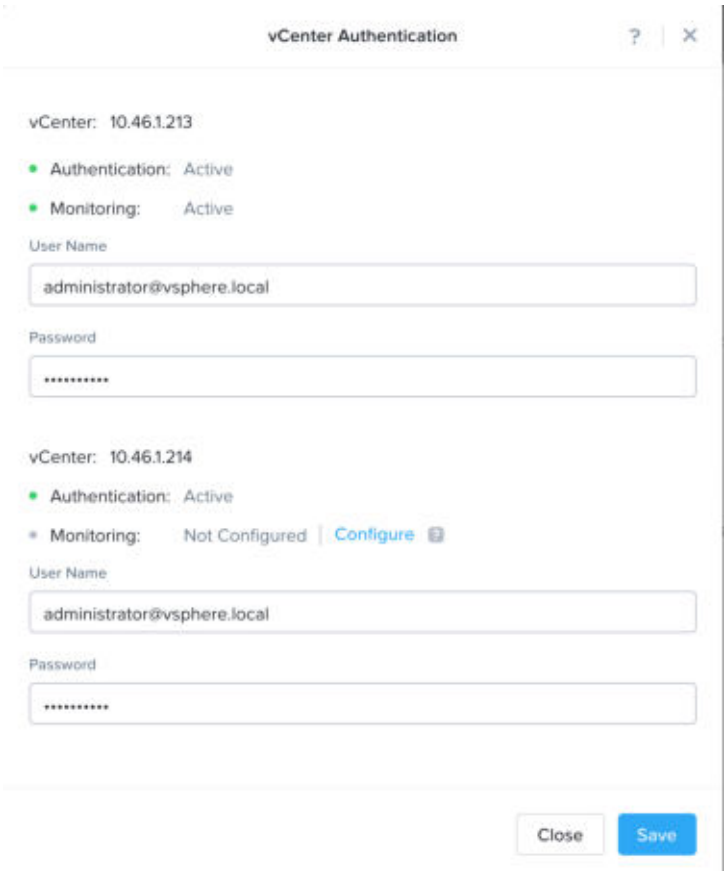

Figure 340: vCenter Authentication Window

Discovering Applications

You can check for the current applications on the monitored clusters at any time.

About this task

To start an application discovery search, do the following:

Procedure

1. Go to the **Discovered Apps** tab of the application discovery dashboard and click the **Discover** button (see Application Discovery View on page 519).
   The search begins immediate and may take a few minutes to complete depending on the number and size of the monitored clusters. When the search completes, the **App Discovery**

window appears. The window displays the search status (successful, unsuccessful); number of discovered, identified, and unidentified applications; number of application types; and number of new application instances.

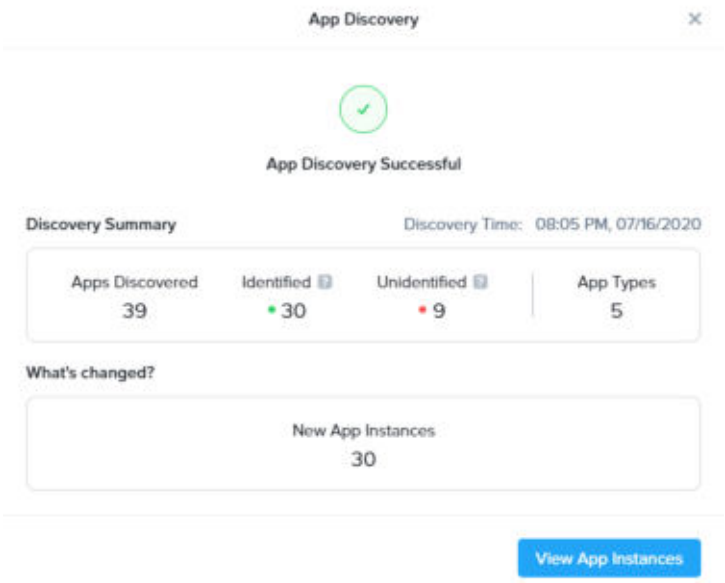

Figure 341: App Discovery Window

2. Click the **View App Instances** button to redisplay the tab.

   If some applications are misclassified, you can modify the results by manually identifying, unidentifying, or ignoring (hiding) selected applications (see the **Actions** menu options in Application Discovery View on page 519).

Publishing Applications

You can publish the application discovery findings to a database.

About this task

Publishing means copying the VM and port mapping data for the identified applications to an internal database, which is accessible through an external API call. To publish data for the current identified applications, do the following:

Procedure

1. Go to the **Discovered Apps** tab of the application discovery dashboard and click the **Publish** button (see Application Discovery View on page 519).
   Information for the identified applications is published immediately.

2. Go to the **Published Apps** tab to view the list of published applications.

   You can delete applications from the published list that you do not want in the database (see Application Discovery View on page 519).

Creating Discovery Policies

You can define custom discovery policies for unidentified applications.

About this task

Application discovery can identify a variety of common (system defined) application types such as DNS, NTP, SSH, Web Server, MS Exchange Server, MS SQL Server, Oracle DB, and SAP HANA. However, it cannot determine the type for some applications. To create a custom (user defined) discovery policy for an unidentified application, do the following:

Procedure

1. Go to the **Policies** tab of the application discovery dashboard and click the **New Policy** button (see Application Discovery View on page 519).

2. In the New Discovery Policy window, do the following in the indicated fields:

   a. **Application**: Enter the application name.

   b. **Policy Name**: Enter a name for the policy.

   c. **Description**: Enter a description for the policy (optional).

   d. **VMs**: Do one of the following:

      » If you want the policy to apply to all VMs, click the **All VMs** radio button.

      » If you want the policy to apply to a limited set of VMs, click the **Specific VMs** radio button and then select the target VMs from the pull-down list.

   e. **TCP Ports**: Enter the TCP port number(s) used by the application. Enter a comma separated list for multiple port numbers.

   f. **UDP Ports**: Enter the UDP port number(s) used by the application. Enter a comma separated list for multiple port numbers.

   g. Click the **Save and Apply** button.

   Information about existing discovered applications is refreshed, reflecting any changes from the new policy.

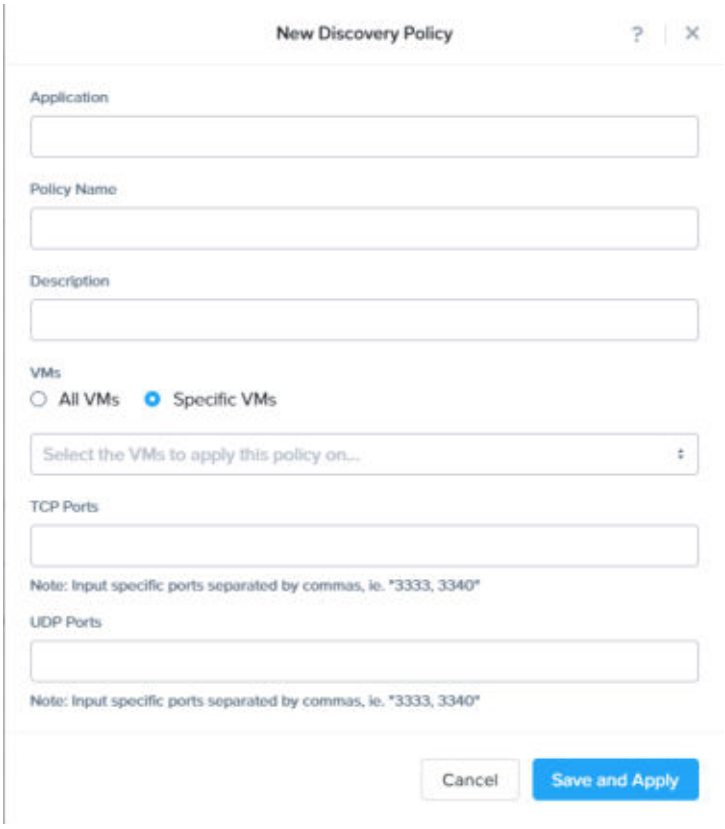

Figure 342: New Discovery Policy Window

Modifying (Viewing) Discovery Policies

You can view any discovery policy, enable or disable any policy, or modify a user-defined policy at any time.

About this task

To view or modify a policy, do the following:

Procedure

1. Go to the **Policies** tab of the application discovery dashboard, find the target policy in the list, and click the name (see Application Discovery View on page 519).

   If it is a system policy, the View Discovery Policy window appears. If it is a user-defined policy, the **Edit Discovery Policy** window appears. The displayed information is the same, but the view window is read-only.

2. To modify one or more parameters in a user-defined policy (see Creating Discovery Policies on page 536), change the entry in the appropriate field(s) and then click the **Save and Apply** button.

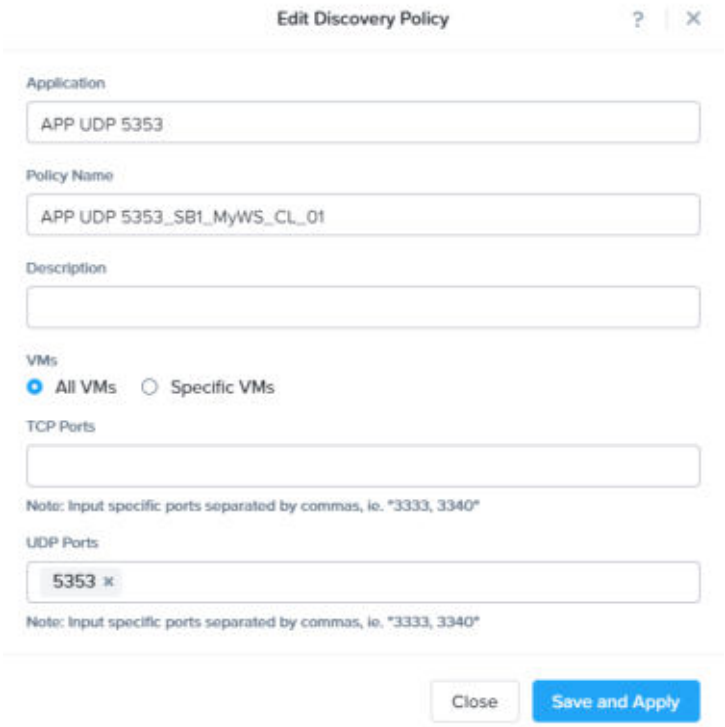

Figure 343: Edit Discovery Policy Window

3. To enable, disable, or delete a policy, select the target policy in the list and then select the desired option (**Enable**, **Disable**, or **Delete**) from the **Actions** pull-down menu. You can only delete user-defined policies (not system policies).

Cost Management (Xi Beam)

About this task

Xi Beam is a cost governance and security compliance SaaS product offering by Nutanix. Xi Beam helps cloud-focused organizations to gain visibility into cloud spend and security compliance status across multiple cloud environments.

To enable Xi Beam, do the following:

Procedure

1. Enable Pulse if it is not already enabled (see Configuring Pulse on page 656).
2. Set up a Xi Beam (SaaS) account.
3. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Operations** > **Cost Management** from the entities menu (see Entities Menu on page 13).
4. In the Welcome to Cost Management page, click the **Manage Cost** button.

   The Xi Beam site appears. See the *Xi Beam User Guide* for information about setting up a Xi Beam account and using Xi Beam to configure and monitor cost governance and security compliance in your environment.

Application Monitoring

Application monitoring provides visibility into integrated applications by collecting application metrics using Nutanix and third-party collectors, providing a single pane of glass for both application and infrastructure data, correlating application instances with virtual infrastructure, and providing deep insights into applications performance metrics.

Application monitoring provides visibility into the following application:

- Microsoft SQL Server
- VMware vCenter Server

Application Monitoring Summary View

The monitoring integrations dashboard allows you to view information about select applications, such as SQL Server or vCenter instances, running in the cluster.

To access the monitoring integrations dashboard, select **Operations** > **Integrations** from the entities menu (see Entities Menu on page 13). The monitoring integrations dashboard allows you to view summary information about application instances and access detailed information about each instance.

Note:

- This section describes the information and options that appear in the application discovery dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.
- If monitoring integrations is not enabled, a "Welcome to Monitoring Integrations" page appears instead of the dashboard. Click the **Get Started** button. This enables monitoring integrations, after which the dashboard appears. However, the dashboard is blank until you configure monitoring (see Configuring Application Monitoring on page 553).

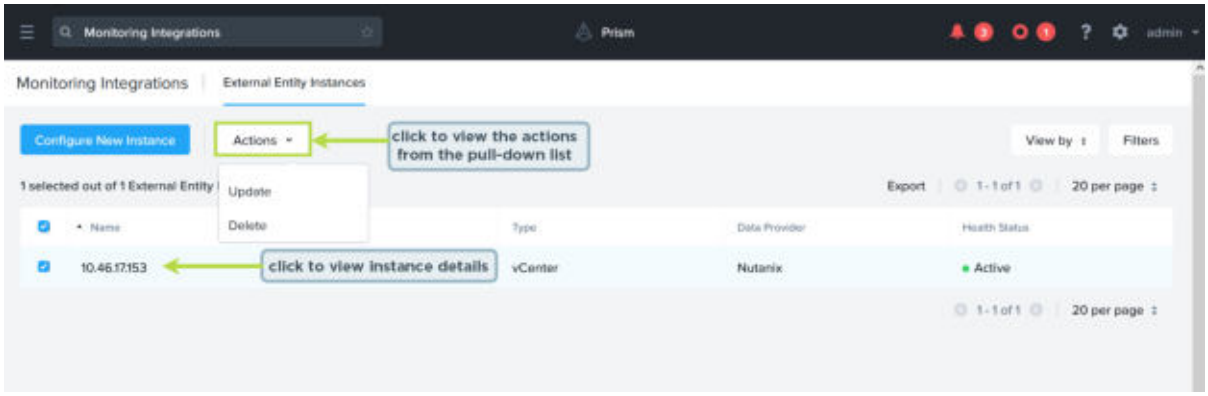

Figure 344: Monitoring Integrations Dashboard

The following table describes the fields that appear in the list of application instances. A dash (-) is displayed in a field when a value is not available or applicable.

Table 233: Application Instance Fields

Parameter	Description	Values
Name	Displays the name of the application instance. Click the name to display the details for that instance (see Application Instance Details View (SQL Server) on page 542 or Application Instance Details View (vCenter) on page 551).	(instance name)
Type	Displays what type of application the instance is.	Microsoft SQL Server, vCenter
Data Provider	Application metrics are supplied through an agentless collector, and this field displays who provided that collector.	Nutanix, (third-party name)
Health Status	Displays the current health status of the instance.	Active, Failed

To download the table contents in CSV format, click the **Export** link.

You can filter the list by opening the **Filter** pane to select one or more of the filter options.

- **Name**: Check the box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field to filter the list for instances with that string in the name.
- **Type**: Check the **Microsoft SQL Server** or **vCenter** box to filter the list for that application type.
- **Health Status**: Check the **Active** or **Failed** box to filter the list for active or failed application instances.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

The dashboard includes the following action buttons:

- Click the **Configure Instance** button to add an application instance for monitoring (see Configuring Application Monitoring on page 553).
- An **Actions** menu appears when an instance is selected:
  - Select **Update** from the pull-down list to update the instance configuration (Modifying Application Monitoring on page 556).
  - Select **Delete** from the pull-down list to stop monitoring the instance. The instance entry disappears from the list.

Application Instance Details View (SQL Server)

The application instance details page provides a range of data and performance metrics for an SQL Server instance.

To access the details page for an SQL Server instance, go to the application integrations dashboard and click the application instance name (see Application Monitoring Summary View on page 540). The instance name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Databases**, **Queries**, and **Metrics**. Click a tab to display that information on the right. (Click the **Back to External Entity Instances** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the cluster (see following table).
- A **Metrics** widget that displays capacity and performance metrics (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, disk or other anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
- Click the **Back to External Entity Instances** link to return to the monitoring applications dashboard.

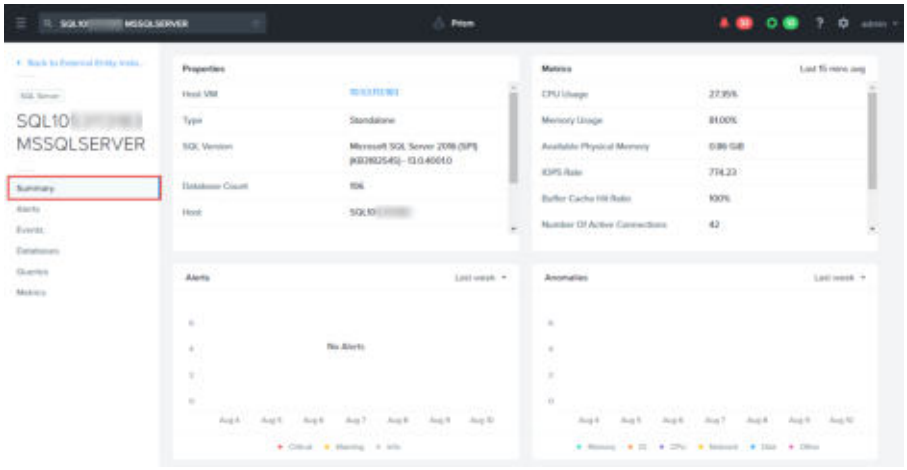

Figure 345: Application Instance Summary Tab

The following table describes the fields in the **Properties** and **Metrics** widgets. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 234: Instance Properties Fields

Parameter	Description	Values
Properties		
Host VM	Displays the host VM IP address. Click the IP address to displays the details page for that VM (see VM Details View on page 90).	(IP address)
Type	Displays the type of SQL server.	Standalone
SQL Version	Displays the SQL server version number for the instance.	(version number)
Database Count	Displays the number of databases managed by the SQL server instance.	(integer)
Host	Displays the SQL server host name.	(host name)
Instance	Displays the instance type.	
VM Name	Displays the host VM name.	(name)
VM Alert Count	Displays the number of alerts assciated with the instance.	(integer)
Metrics		
CPU Usage	Displays the percentage of CPU capacity currently being used by the instance.	0 - 100%
Memory Usage	Displays the percentage of memory capacity currently being used by the instance.	0 - 100%
Available Physical Memory	Displays the amount of available memory.	xx [GiB]

Parameter	Description	Values
IOPS Rate	Displays the average number of combined read and write operations per second for the past 15 minutes.	(number)
Buffer Cache Hit Ratio	Displays the buffer cache hit rate for the past 15 minutes.	0-100%
Number of Active Connections	Displays the number of active connections.	(integer)
Read Ops Rate	Displays the average number of read operations per second for the past 15 minutes.	(number)
Write Ops Rate	Displays the average number of write operations per second for the past 15 minutes.	(number)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard (List tab), except it is filtered to display just SQL Server-related alerts (see Alerts Summary View on page 210).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just SQL Server-related events (see Events Summary View on page 221).

Databases Tab

The **Databases** tab displays a list of databases supported by this SQL Server instance. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

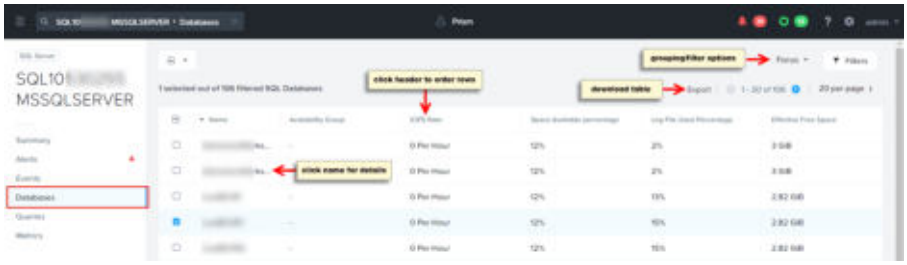

Figure 346: Application Instance Databases Tab

- To order the rows, click a column header, which toggles the list alphabetically or numerically (as appropriate) based on that parameter.
- To download the table in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 77).
- To filter the list by database name, open the **Filter** pane, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field.

- To select (or unselect) databases in the list:
  - Click the check box on the far left for each database to select. (Click the check box again to unselect that database.)
  - Click the check box in the header row to select (unselect) all the databases in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of databases.

Table 235: Databases Fields (SQL Server)

Parameter	Description	Values
Name	Displays the database name. Click on the name to display the details for that database (described below).	(database name)
Availability Group	Displays the name of the availability group to which the database belongs.	(availability group name)
IOPS Rate	Displays the average number of combined disk read and write read operations per hour.	<number> Per Hour
Space Available Percentage	Displays the percentage of space available on disk used by the database.	0-100%
Log File Used Percentage	Displays the percentage of the log file used for the database.	0-100%
Effective Free Space	Displays the amount of space on disk that is usable by the database. This includes the free space of the entire disk and any space being used by the database.	xxx [MiB|GiB]

Clicking a database name displays the details page for that database. The following figure is a sample view. The page displays four graphs. Each graph is a rolling time interval performance or usage monitor for a metric. The metrics are IOPS rate, space available, log file utilization, and effective free space.

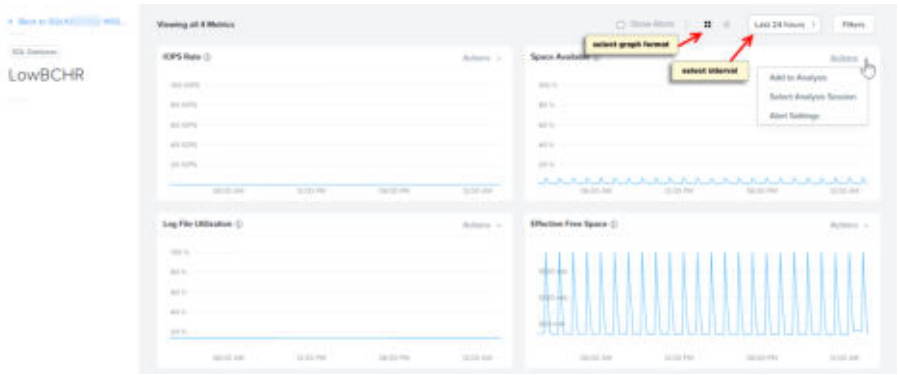

Figure 347: Database Instance Page

- Each graph includes a pull-down actions menu with the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard on page 486).
  - Select **Select Analysis Session**, which opens the **Select A Session** window, to add the metric to an existing or new analysis session. Choose an option from the pull-down list and then click the Add to Session button.

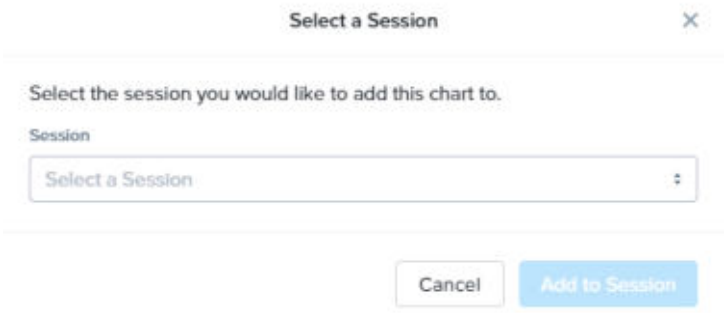

Figure 348: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 230).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired graphs.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to <instance_name>** link to return to the application instance details page.

Queries Tab

The **Queries** tab displays a list of SQL query statements. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

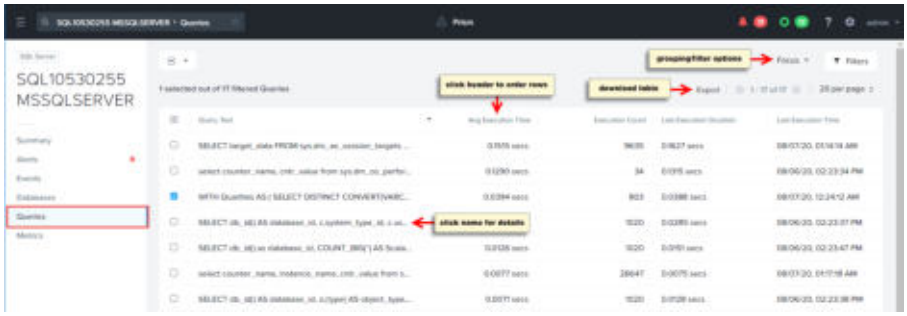

Figure 349: Application Instance Queries Tab

- To order the rows, click a column header, which toggles the list alphabetically, numerically, or chronologically (as appropriate) based on that parameter.

- To download the table in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 77).
- To filter the list for a text string in the query, open the **Filter** pane, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field.
- To select (or unselect) queries in the list:
  - Click the check box on the far left for each query to select. (Click the check box again to unselect that query.)
  - Click the check box in the header row to select (unselect) all the queries in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of queries.

Table 236: Query Fields

Parameter	Description	Values
Query Text	Displays the query statement text. Click on the text to display details about that query statement (described below).	(query text)
Avg Execution Time	Displays the average amount of time it took to run the query.	xxx secs
Execution Count	Displays the number of times the query statement has been executed.	(integer)
Last Execution Duration	Displays the amount of time it took to run the last query.	xxx secs
Last Execution Time	Displays the date and time when the query was last executed.	(date and time)

Clicking a query text displays details about that query, including a Properties widget that includes the Execution Count, Last Execution Duration, and Query text fields.

- To see the full query statement, click the **View** link.

Figure 350: Query Statement Window

- To copy the query statement, click the **Copy** link.
- Click the **Back to <instance_name>** link to return to the application instance details page.

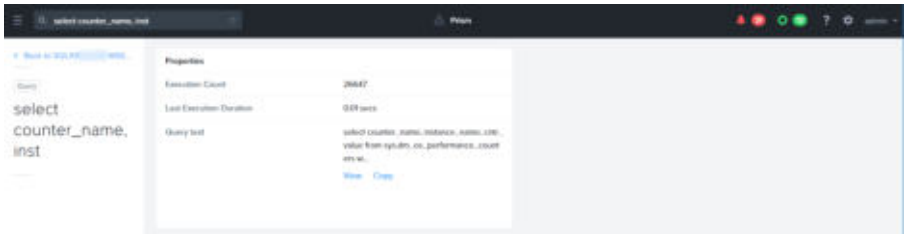

Figure 351: Query Details Page

Metrics Tab

The **Metrics** tab displays a set of graphs for various metrics. Each graph is a rolling time interval performance or usage monitor for that metric.

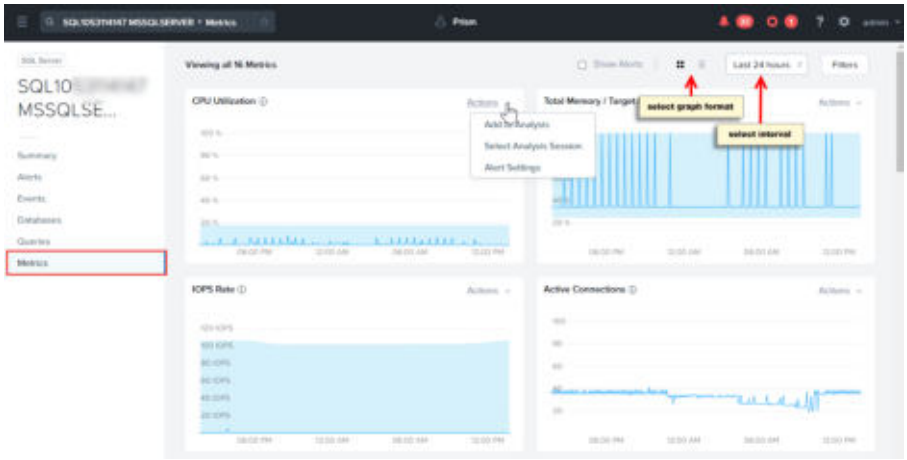

Figure 352: Application Instance Metrics Tab

- Each graph includes a pull-down actions menu with the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard on page 486).
  - Select **Select Analysis Session**, which opens the **Select A Session** window, to add the metric to an existing or new analysis session. Choose an option from the pull-down list and then click the Add to Session button.

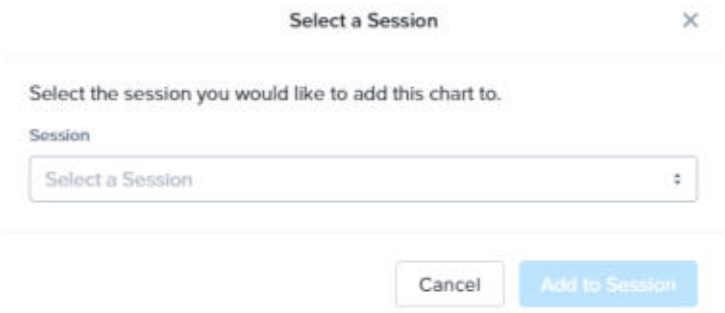

Figure 353: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 230).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired graphs.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to External Entity Instances** link to return to the monitoring applications dashboard.

The following table describes the metric graphs.

Table 237: Metric Performance Graphs

Metric	Description
CPU Utilization	Displays the CPU bandwidth usage by all requests in all workload groups belonging to all related resource pools.
Total Memory / Target Memory	Displays the percentage of total server memory divided by target server memory. When the total server memory and target server memory values are similar (close to 100%), there is no memory pressure on the server. Total server memory significantly lower than target server memory during normal SQL Server operations indicates there is memory pressure on the server, so SQL Server cannot get as much memory as needed.

Metric	Description
IOPS Rate	Displays the total amount of disk write and read operations. This is a calculated metric displaying the IOPS that occurred per second since the previous collection.
Active Connections	Displays the number of active connections to this instance.
Buffer Cache Hit Ratio	Displays the percentage of pages found in the buffer cache without having to read from disk. The ratio is the total number of cache hits divided by the total number of cache lookups over the last few thousand page accesses. Because reading from the cache is much less expensive than reading from disk, you want this ratio to be high. Generally, you can increase the buffer cache hit ratio by increasing the amount of memory available to SQL Server or by using the buffer pool extension feature.
Batch Request Rate	Displays the rate of transact-SQL command batches received. This statistic is affected by all constraints (such as I/O, number of users, cache size, complexity of requests, and so on). High batch requests mean good throughput. This is a calculated metric displaying the batch requests that occurred per second since the previous collection.
SQL Compilations Rate	Displays the rate of SQL compilations of the instance. This is a calculated metric displaying the SQL compilations that occurred per second since the previous collection.
SQL Recompilations Rate	Displays the rate of statement recompiles. Counts the number of times statement recompiles are triggered. Generally, you want the recompiles to be low. This is a calculated metric displaying the SQL Recompilations that occurred per second since the previous collection.
Buffer Page Life Expectancy	Displays the number of seconds a page will stay in the buffer pool without references.
Checkpoint Pages Rate	Displays the number of pages flushed to disk by a checkpoint or other operation that requires all dirty pages to be flushed. This is a calculated metric displaying the checkpoint pages that occurred per second since the previous collection.
Available Physical Memory	Displays the amount of physical memory available.
Buffer Pool Size	Displays the total size of the buffer pool.
Lock Wait Time	Displays the total wait time for locks in the last second.
User Connections	Displays the number of connections that are currently being used for user sessions.
Read Operations Rate	Displays the total amount of disk read operations. This is a calculated metric displaying the write operations that occurred per second since the previous collection.
Write Operations Rate	Displays the total amount of disk write operations. This is a calculated metric displaying the write operations that occurred per second since the previous collection.

Application Instance Details View (vCenter)

The application instance details page provides a range of data and performance metrics for a vCenter instance.

To access the details page for a vCenter instance, go to the application integrations dashboard and click the application instance name (see Application Monitoring Summary View on page 540). The instance name and tabs for **Summary** and **Alerts** appear on the left. Click a tab to display that information on the right. (Click the **Back to External Entity Instances** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays a **Properties** widget that shows summary information about the instance. Click the **Back to External Entity Instances** link to return to the monitoring applications dashboard.

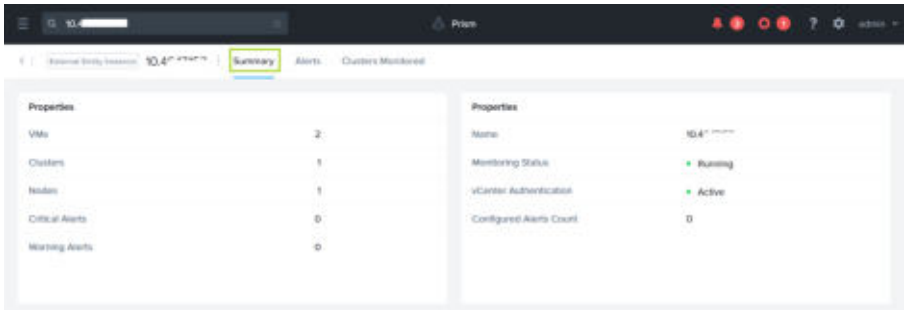

Figure 354: Application Instance Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 238: Instance Properties Fields

Parameter	Description	Values
VMs	Displays the number of VMs managed by the vCenter instance.	(integer)
Clusters	Displays the number of clusters managed by the vCenter instance.	(integer)
Nodes	Displays the number of clusters managed by the vCenter instance.	(integer)
Critical Alerts	Displays the current number of critical alerts associated with the vCenter instance.	(integer)
Warning Alerts	Displays the current number of warning alerts associated with the vCenter instance.	(integer)
Name	Displays the name of the vCenter instance.	(name)
Configured Alerts Count	Displays the number of alerts configured to monitor the vCenter instance.	(integer)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard (List tab), except it is filtered to display just vCenter-related alerts (see Alerts Summary View on page 210).

Infrastructure Tabs (VMs, Clusters, and Hosts)

In addition to the application details view, tabs for non-Nutanix entities appear in the VMs, clusters, and hosts dashboards after enabling vCenter application monitoring.

- To view information about VMs managed by vCenter, select **Virtual Infrastructure** > **VMs** from the entities menu and then click the **non-Nutanix** tab (see VMs Summary View on page 82). The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed VMs. The **Metrics** tab displays a subset of the full list (10 of the 12 metrics). The **List** tab displays fields for name, node name, hypervisor, memory, IP addresses, power state, and cluster name.

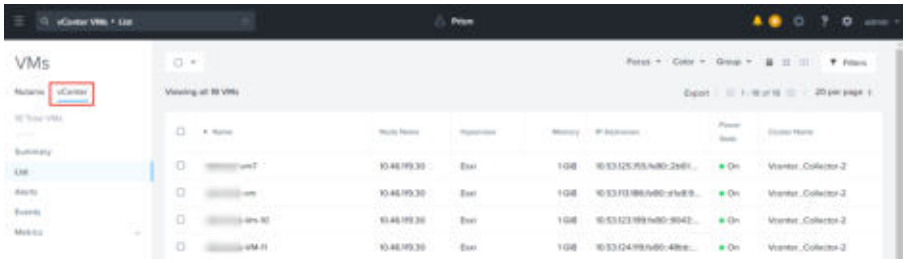

Figure 355: VMs dashboard (vCenter tab)

- To view information about clusters managed by vCenter, select **Hardware** > **Clusters** from the entities menu and then click the **non-Nutanix** tab (see Clusters Summary View on page 146). The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed clusters. The **Metrics** tab displays CPU usage, memory usage, and VM operations. The **List** tab displays fields for name, host count, VM count, storage usage, and storage capacity.

Figure 356: Clusters dashboard (vCenter tab)

- To view information about hosts managed by vCenter, select **Hardware** > **Hosts** from the entities menu and then click the **non-Nutanix** tab (see Hosts Summary View on page 165). The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed

hosts. The **Metrics** tab displays a subset of the full list (5 of the 7 metrics). The **List** tab displays fields for name, host IP, VM count, memory capacity, and cluster name.

Figure 357: Hosts dashboard (vCenter tab)

Configuring Application Monitoring

Monitoring an application instance requires configuring an instance profile.

Before you begin

Complete the Prerequisites on page 555.

About this task

To configure an instance profile, do the following:

Procedure

1. Go to the monitoring integrations dashboard and click the **Configure Instance** button (see Application Monitoring Summary View on page 540).

2. In the Configure external entity for monitoring page, do the following in the indicated fields:

   a. **External Entity Type**: Select the target type from the pull-down menu.

      The options are **Microsoft SQL Server** or **vCenter**. If this is the first time configuring an instance, an **Enable** button appears. Click the **Enable** button to provision instance monitoring.

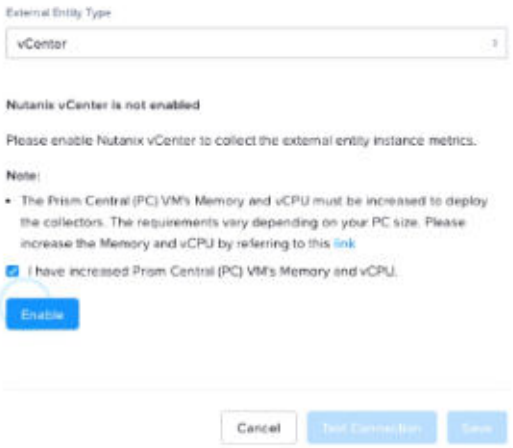

Figure 358: Configuration Page (enable)

When the provisioning is complete (which can take a few minutes), the following fields appear.

Figure 359: Configuration Page (SQL server example)

   b. [**Microsoft SQL Server** | **vCenter**] **Host**: Enter the IP address of the SQL Server or vCenter host VM.

   c. [**Microsoft SQL Server** | **vCenter**] **Port**: Enter the target port number.

      The port number defaults to the standard port number for SQL Server (1433) or vCenter (443).

   d. **Username**: Enter the SQL Server or vCenter administrator user name.

   e. **Password**: Enter the SQL Server or vCenter administrator password.

   f. [vCenter only] **SSL Verify**: Select the SSL mode to use (**Verify** or **Don't Verify**) when connecting to the target.

      This setting determines whether to verify (or not verify) the target's SSL certificate as part of the connection process.

g. Click the **Test Connection** button to verify (test) the connection.

   A "Test connection success" message appears after a successful test.

h. After a successful test, do one of the following:

   » SQL Server: Click the **Save** button. The instance configuration and data collection begins. When it completes (which can take a few minutes), click the **Close** button to close the page. The monitoring integrations dashboard reappears with the new instance in the list.

   » vCenter: Click the **Next** button.

i. [vCenter only] Select the clusters to monitor.

   A list of discovered clusters for the vCenter instance appears. Select (check the boxes for) the clusters you want to monitor from the list and then click the **Save** button.

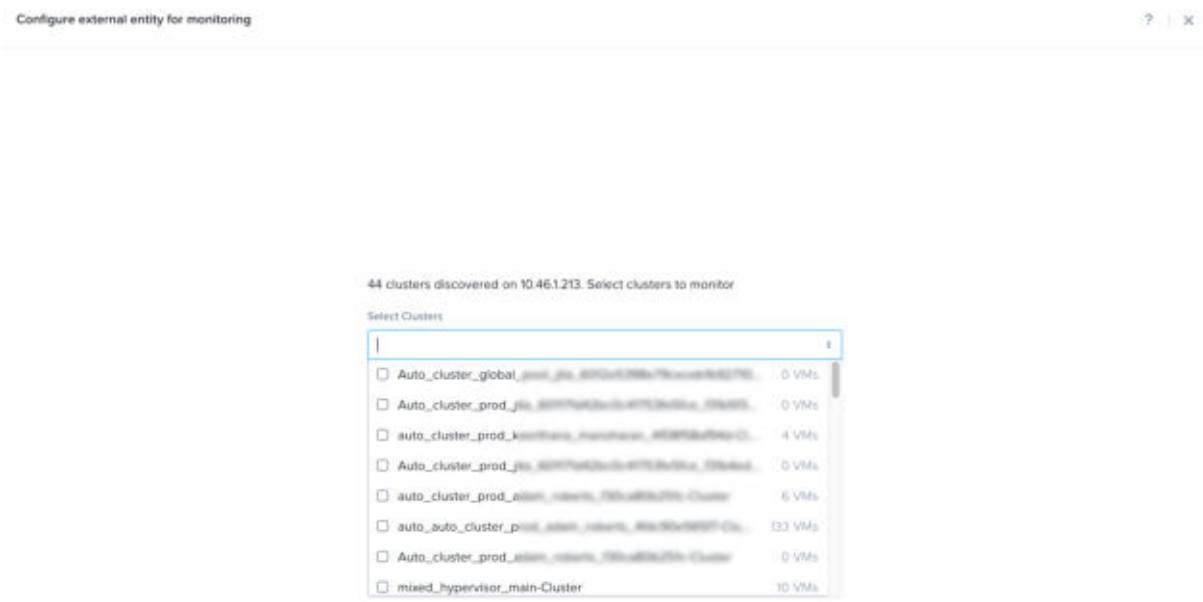

Figure 360: Select Clusters

> Note: Nutanix clusters monitored by vCenter that are also registered with this Prism Central instance do not appear in the discovered clusters list. Conversely, non-registered Nutanix clusters selected here but later registered with Prism Central appear in both the Nutanix and non-Nutanix tabs of the clusters dashboard (see Application Instance Details View (vCenter) on page 551).

The instance configuration and data collection begins. When it completes (which can take a few minutes), a "vCenter Configuration Successful" screen appears. Either click the **Close** button to close the page or click the **App Discovery** link (if present) to enable application discovery for the selected clusters (see Enabling Application Discovery on page 526).

Prerequisites

Address the following before enabling application monitoring:

- Sufficient Prism Central VM resources. Application monitoring is a feature that requires additional Prism Central VM resources (more memory and vCPUs). The requirements are listed in KB 8932 (see instructions for Prism Ultimate features). If your Prism Central VM resources are below the recommended amount, manually increase the memory and vCPUs to the recommended amount.
- SQL Server-specific requirements:
  - Each SQL Server instance port must be accessible through your network, but no other special network configuration is required.
  - Create a read-only login to connect to your server:

    ```
    CREATE LOGIN <username> WITH PASSWORD = '<PASSWORD>';
    CREATE USER <username> FOR LOGIN <username>;
    GRANT SELECT on sys.dm_os_performance_counters to <username>;
    GRANT VIEW SERVER STATE to <username>;
    ```

    After creating the login make sure that your SQL Server instance supports SQL Server authentication by enabling "SQL Server and Windows Authentication mode" in the server properties (**Server Properties** > **Security** > **SQL Server and Windows Authentication mode**).
  - Active Directory users and Windows authentication are not supported for SQL Server monitoring.

Modifying Application Monitoring

You can view, update, or delete the application configuration for an instance at any time.

About this task

To modify application monitoring for an instance, do the following:

Procedure

1. Go to the monitoring integrations dashboard, select the target instance in the list, and then select **Update** from the **Actions** menu pull-down list (see Application Monitoring Summary View on page 540).

   The Configure external entity for monitoring page for that instance appears (see Configuring Application Monitoring on page 553).
2. To modify one or more parameters, change the entry in the appropriate field(s), click the **Test Connection** button to verify the changes, and then click the **Save** button.

   Note: [vCenter] Modifying the cluster selection occurs after testing the connection.
3. To remove application monitoring for an instance, select the target instance in the list, and then select **Delete** from the **Actions** menu pull-down list.

   The entry for that instance disappears from the list.

Operations Policy Management

Prism Central implements a variety of system services to analyze cluster integrity and performance. You have the option to apply policies to customize some services in your environment. You can apply the following policies:

- **Anomaly Detection Exclusion**. Prism Central includes machine-learning capabilities that analyze resource usage and identify potential problems (see Behavioral Learning Tools on

page 582). One is anomaly detection that records when performance or resource usage is outside an expected range. The anomaly detection exclusion policy allows you to specify VMs to ignore when looking for anomalies. For example, you might have some VMs with expected peak load on certain days, and you do not want that usage surge counted as an anomaly. In this case you can either define a category for those VMs and enable anomaly detection exclusion for that category or enable the exclusion for each VM individually.

- **Inefficiency Measurement Exclusion**. Another machine-learning capability is VM "right sizing" that identifies VMs which exhibit inefficient profiles. Again, you have the option to exclude selected VMs from the inefficiency analysis.

Operations Policies View

The operations policies dashboard allows you to view and modify operations policies that can be applied to VMs or other entities.

To access the operations policies dashboard, select **Operations** > **Ops Policies** from the entities menu (see Entities Menu on page 13). The dashboard displays a list of the operations policies.

> Note: This section describes the information and options that appear in the operations policies dashboard. See Entity Exploring on page 77 for instructions on how to view and organize that information in a variety of ways.

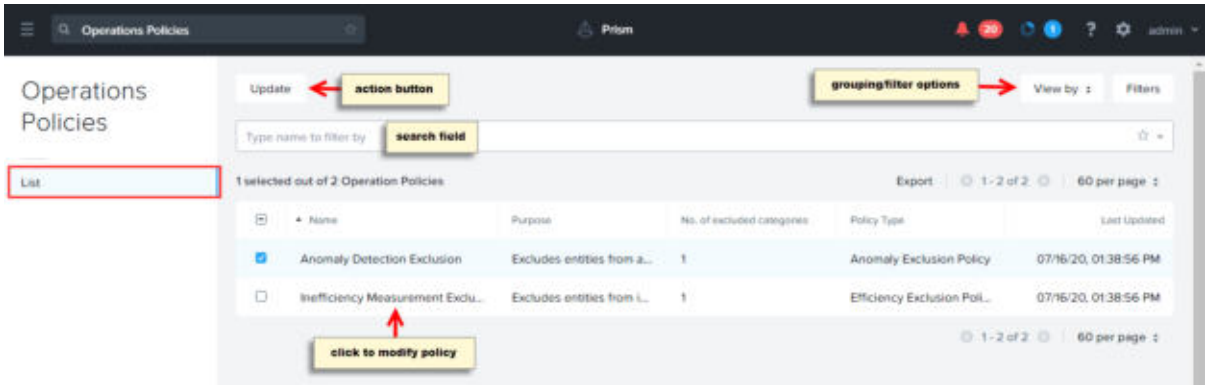

Figure 361: Operations Policies Dashboard

The following table describes the fields that appear in the operations policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 239: Operations Policy Fields

Parameter	Description	Values
Name	Displays the application name for the instance. Click the name to display the Application Info window. The window displays the application name, VM name, TCP ports, and UDP ports.	(app name)
Purpose	Describes the purpose of the policy.	(VM name)
No. of Excluded Categories	Displays the number of categories the policy applies to (see Category Management on page 413).	(integer)

Parameter	Description	Values
Policy Type	Describes what type of policy this is.	(type name)
Last Updated	Displays the date and time when the policy was last updated.	(date and time)

You can filter the list either by entering a string in the search field or by clicking the **Filters** button, checking the **Name** box, selecting a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and entering a string in the field.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

To modify a policy, select that policy and then select **Update** from the **Actions** pull-down menu (see Modifying an Operations Policy on page 418).

Modifying an Operations Policy

You can customize the list of categories to which an operations policy applies.

About this task

No categories are assigned to the Anomaly Detection Exclusion and Inefficiency Measurement Exclusion polices by default. To add (or delete) one or more categories to an operations policy, do the following:

Procedure

1. Go to the operations policy dashboard, select the desired policy, and then select **Update** from the **Actions** pull-down menu (see Operations Policies View on page 417).
2. In the policy page, click the **Add** link.

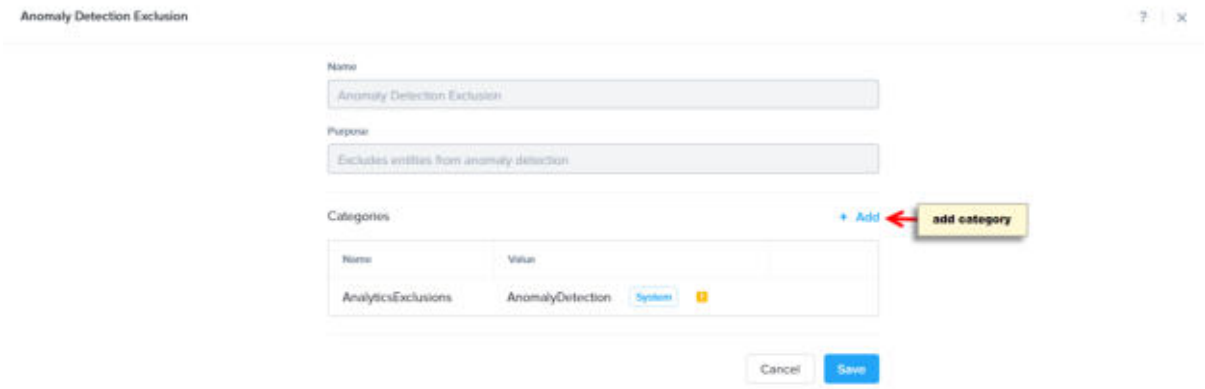

Figure 362: Anomaly Detection Exclusion Page (add category)

3. In the Add window, enter a category name and then click the plus icon. Repeat this step for all categories you want to add.

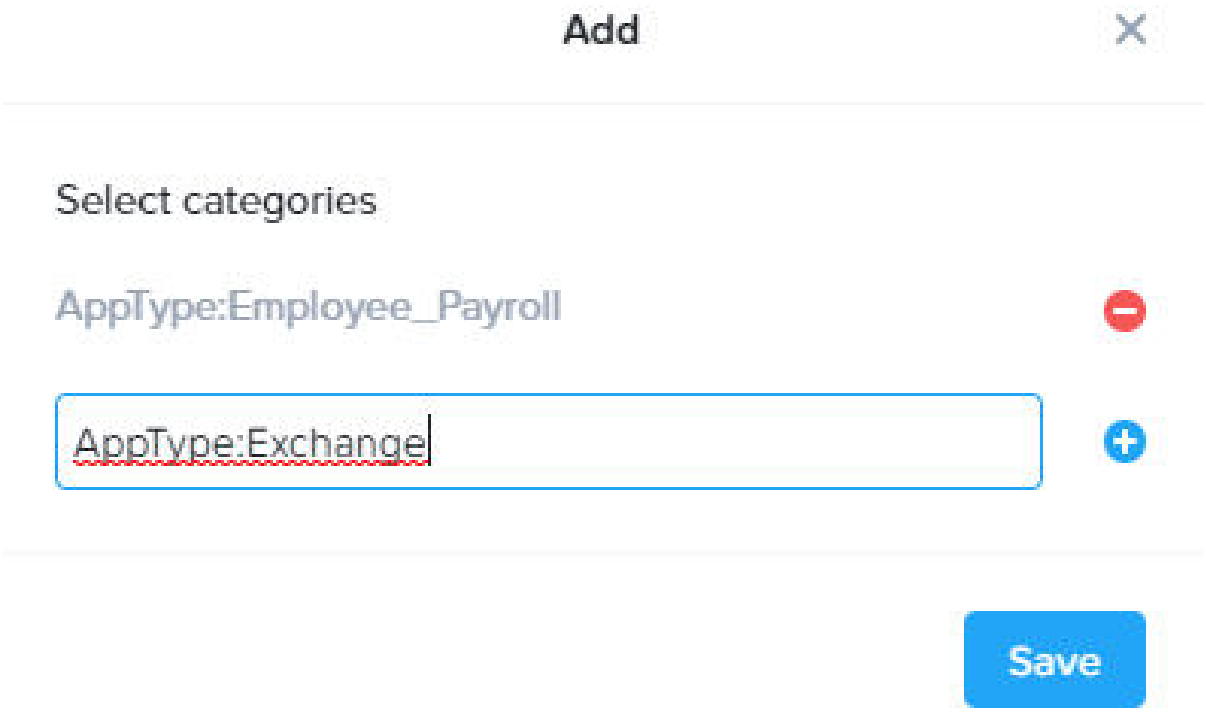

Figure 363: Add Window

4. Click the **Save** button.

   The policy is now enabled for entities in the selected categories.

5. To delete a category from the list, click the garbage pail icon for that category and then click the **Save** button.

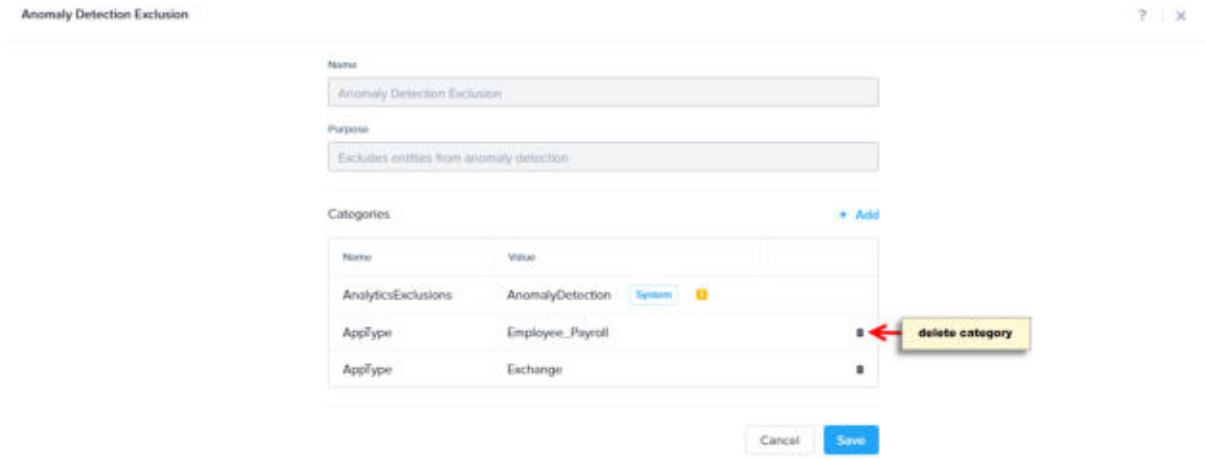

Figure 364: Anomaly Detection Exclusion Page (delete category)

Enabling an Operations Policy

You can enable (or disable) an operations policy for selected VMs.

About this task

To disable anomaly detection or inefficiency measurement for selected VMs or for all the VMs in selected clusters, do the following:

Procedure

1. Do one of the following:

   » Go to the List tab of the VMs dashboard, select the target VMs, and then select **Disable Efficiency Measurement** or **Disable Anomaly Detection** from the **Actions** pull-down menu (see VMs Summary View on page 82).

   » Go to the List tab of the clusters dashboard, select the target clusters, and then select **Disable Efficiency Measurement** or **Disable Anomaly Detection** from the **Actions** pull-down menu (see Clusters Summary View on page 146).

   This immediately disables the specified operation (enables the exclusion policy) for the target VMs.

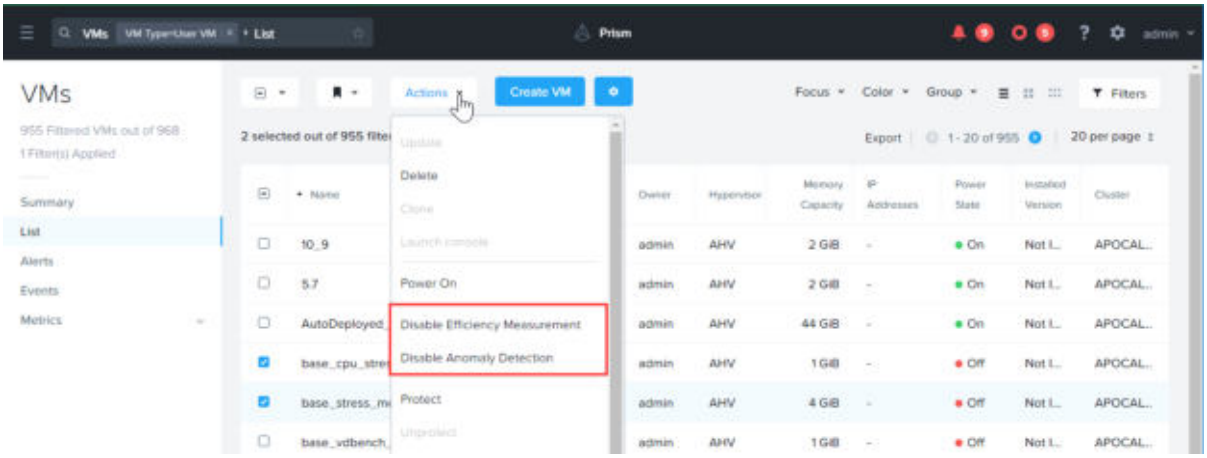

Figure 365: VM Actions Menu

2. To re-enable the operation, select the target VMs or clusters, and then select **Enable Efficiency Measurement** or **Enable Anomaly Detection** from the **Actions** pull-down menu.

   This immediately enables the specified operation (disables the exclusion policy) for the target VMs.

Resource Planning

You can review and analyze current and potential resource needs through the planning dashboard.

To access the planning dashboard, perform the following.

1. Click the hamburger icon.
2. Select **Operations** from the entities menu. (See Entities Menu on page 13).
3. Select **Planning** from the **Operations** menu.

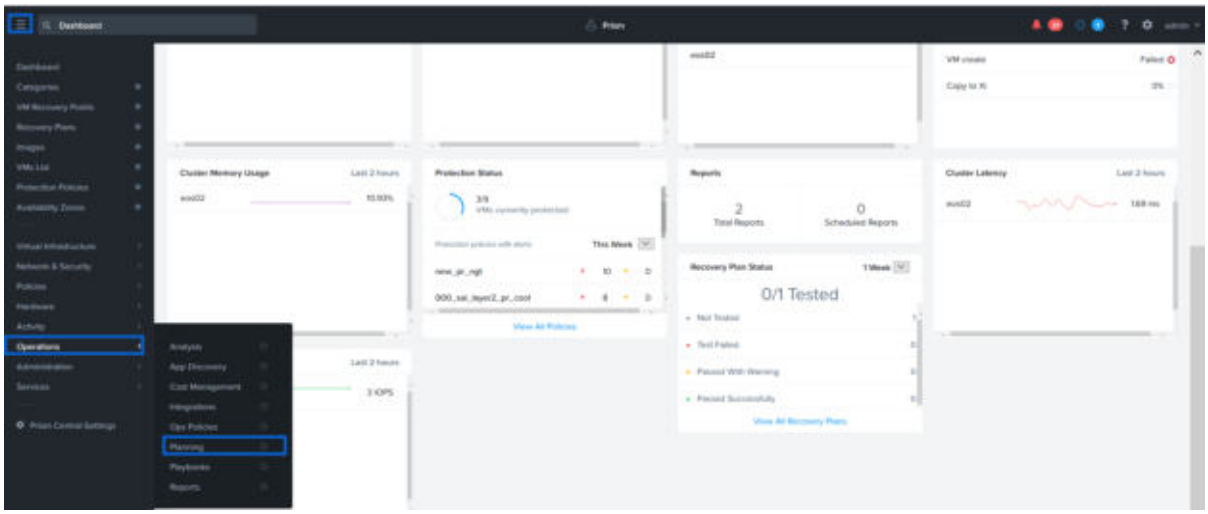

Figure 366: Planning

The planning dashboard includes two tabs:

> Note: You must have a Prism Pro license to use the resource planning tools.

- The **Capacity Runway** tab allows you to view current resource runway information across the registered Nutanix and Non-Nutanix clusters (see Capacity Runway View on page 561).
- The **Scenarios** tab allows you to create "what if"scenarios to assess the future resource requirements for potential work loads that you specify (see Scenarios View on page 563).
- All resource capacity in a cluster is considered by default when estimating the runway. To reserve some capacity before doing any runway analysis, see Updating Capacity Configurations on page 581.

Prism Central includes machine-learning capabilities that analyze resource usage over time and provide tools to monitor resource consumption, identify abnormal behavior, and guide resource planning (see Behavioral Learning Tools on page 582).

Capacity Runway View

Click the **Capacity Runway** tab in the planning dashboard (see Resource Planning on page 560) to view summary resource runway information for the registered Nutanix and Non-Nutanix clusters and to access detailed runway information about each cluster.

> Note:
>
> - Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
> - Capacity planning checks the past 90 days of data for monthly seasonality, it also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
> - For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 540 for more information.

- Click a column header to order the rows by that value (alphabetically or numerically as appropriate).
- Click the cluster name to display the **Capacity** tab of the details page for that cluster (see Cluster Details View on page 153).

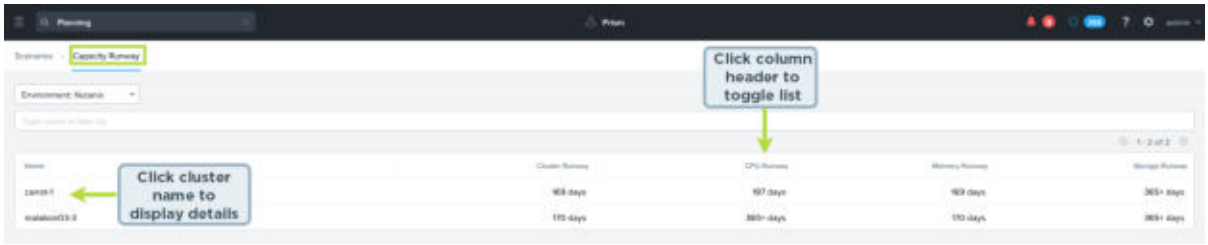

Figure 367: Capacity Runway View (Nutanix)

Figure 368: Capacity Runway View (Non-Nutanix)

The following table describes the fields that appear in the Capacity Runway view.

Table 240: Capacity Runway Tab Fields

Parameter	Description	Values
Name	Displays the cluster name. Clicking the name displays the **Capacity** tab of the details page for that cluster. This tab displays current, historical, and projected usage information. (see Cluster Details View on page 153).	(cluster name)
Cluster Runway	Displays the overall runway for the cluster. The "cluster" runway is the shortest CPU, memory, or storage runway; that is, how long will it be before one of these resources runs out of capacity? Note: The "runway" is the number of days the existing CPU, memory, or storage capacity can handle the workload based on the current workload consumption model.	xxx [days]
CPU Runway	Displays the cluster runway for CPU usage.	xxx [days]
Memory Runway	Displays the cluster runway for memory usage.	xxx [days]
Storage Runway	Displays the cluster runway for storage usage.	xxx [days]

Scenarios View

Click the **Scenarios** tab in the planning dashboard (see Resource Planning on page 560) to create "what if" scenarios to assess the future resource requirements for potential work loads that you specify. The Scenarios view displays a list of the previously created and saved scenarios.

> Note: If there are no saved scenarios, a "Welcome to capacity planning" message appears with a **Get Started** button. Clicking the **Get Started** button is the same as clicking the **New Scenario** button.

- Click the **New Scenario** button (upper right) to create a new scenario (see Creating a Scenario on page 564).
- Click a scenario name to display the details for that scenario.
- Click the appropriate pencil icon to modify that scenario, or click the X icon to delete that scenario.
- Enter a string in the search field to filter the list for scenarios with that string in the name.

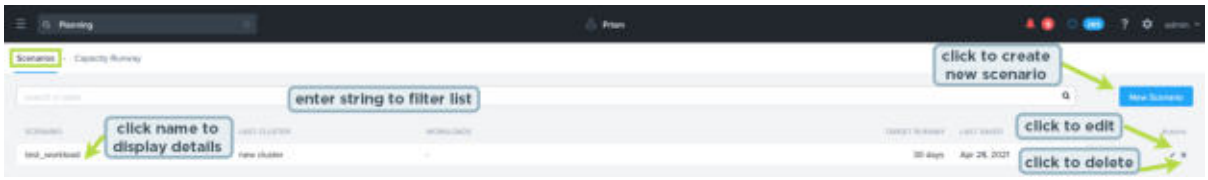

Figure 369: Scenarios View

The following table describes the fields that appear in the Scenarios view.

Table 241: Scenarios Tab Fields

Parameter	Description	Values
Scenario	Displays the scenario name. Clicking on the name displays the details of that scenario (see Creating a Scenario on page 564).	(scenario name)
Last Cluster	Displays the name of the cluster to which the scenario applies. (The column title is "Last Cluster" because you can change the target cluster when modifying a scenario.)	(cluster name)
Workloads	Displays the workloads included in the scenario. This field lists workloads added to the scenario. In the figure above, the "tom" scenario lists "sql" in this field indicating an SQL workload was added to the scenario. A blank field indicates either there is no workload (for a scenario based on a new cluster) or only the existing (current) workload is applied (for a scenario based on an existing cluster).	(blank),

Parameter	Description	Values
Target Runway	Displays the target runway duration for this scenario, that is how many days am I planning for. This is specified when the scenario is created or modified ("Target" parameter).	xxx [days]
Last Saved	Displays the time and date for the last (most recent) time the scenario was saved.	(time and date)
Actions	• To edit the scenario, click the pencil icon (see Modifying a Scenario on page 580). • To delete the scenario, click the X icon.	(n/a)

Creating a Scenario

About this task

Anticipating future resource needs can be a challenging task. To address this task, Prism Central provides an option to create "what if" scenarios that assess the resource requirements for possible future work loads. This allows you to evaluate questions like

- How many new VMs can the current cluster support?
- If I need a new database server in a month, does the cluster have sufficient resources to handle that increased load?
- If I create a new cluster for a given set of workloads, what kind of cluster do I need?
- If I remove a set of VMs or nodes, how will my cluster look?

You can create various "what if" scenarios to answer these and other questions. The answers are derived by applying industry standard consumption patterns to the hypothetical workloads and current consumption patterns for existing workloads. To create a "what if" scenario, do the following:

Procedure

1. Go to the **Scenarios** tab of the planning dashboard (see Scenarios View on page 563) and click the **New Scenario** button.
   The Scenario configuration page appears.

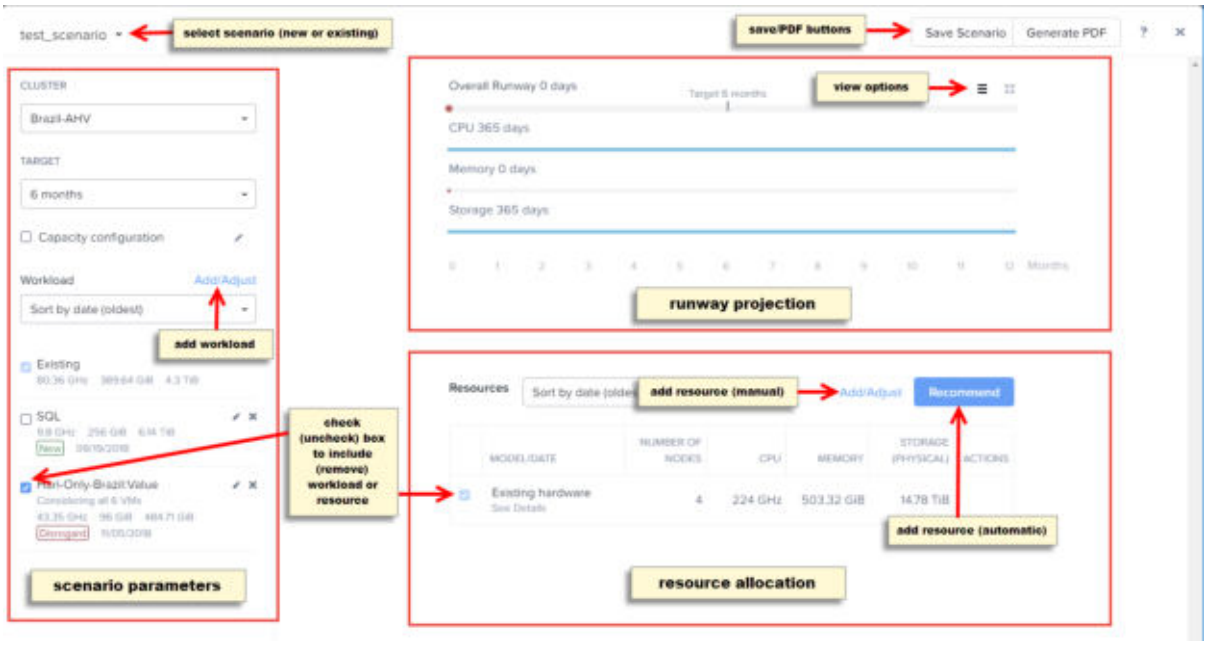

Figure 370: Scenario Configuration Page (example)

2. In the top left field, select **New Scenario** (to start with a blank scenario configuration) or the name of an existing scenario (to start with that scenario configuration) from the pull-down list and do the following in the indicated fields:

   Note: If there are no existing scenarios (see Scenarios View on page 563), the only option is **New Scenario**.

   a. **Cluster**: Select either **New Cluster** or an existing cluster from the pull-down list.

      Selecting an existing cluster means you start with the resources and workloads currently in that cluster; selecting **New Cluster** means you start with a blank slate (no resources or workloads).

      Note: You can select a cluster at any time to see the runway for that cluster using this scenario.

   b. **Target**: Select a runway duration (1 to 12 months) from the pull-down list.

   c. **Vendor Type** (new cluster only): For a new cluster, select the hardware vendor of choice (Nutanix, Dell, Lenovo, Cisco, IBM) from the pull-down menu.

   d. **Capacity configuration**: Check this box to reserve some capacity prior to analyzing the scenario.

      All the storage, CPU, and memory capacity in the cluster is used by default. However, you have the option to reserve (withhold) some capacity before doing the analysis. To specify

the capacity to reserve, click the pencil icon and do the following in the Edit Cluster Configuration window:

1. **Replication Factor**: Select the desired replication factor (2 or 3) from the pull-down list.
2. **CPU Reservation**: Specify the percentage of CPU capacity to keep in reserve. (The default is 0.)
3. **Memory Reservation**: Specify the percentage of memory capacity to keep in reserve. (The default is 0.)
4. **Storage Reservation**: Specify the percentage of storage capacity to keep in reserve. (The default is 0.)
5. **CPU Overcommit Ratio for New VM Worloads**: Specify the ratio of virtual CPU to physical CPU. (The default is 1.)
6. **Memory Overcommit Ratio for New VM Worloads**: Specify the ratio of virtual memory to physical memory. (The default is 1.)
7. Click the **Save Configuration** button to save the configuration and close the window and then check the **Capacity configuration** box (if it is not already checked) to apply this configuration.

Figure 371: Edit Cluster Configuration Window

> Note: Continue to the next step (e) to add a workload to the scenario. To remove the impact of selected VMs in an existing workload from the scenario, skip to the following step (f).

e. To add a workload, click the **Add/Adjust** link in the Workload section, which opens the Add/Adjust Workload window, and then click the **New Workload** radio button.

   This displays fields for adding a workload. In the **Workload** field, select the desired workload from the pull-down list, enter appropriate information in the remaining fields, and then click the **Save** button to add that workload to the scenario. You can choose from

a variety of workloads. The configuration fields vary depending on the selected workload. The following table describes the configuration fields for each workload.

> Note: If you want to include multiple workloads in the scenario, repeat this step as many times as needed.

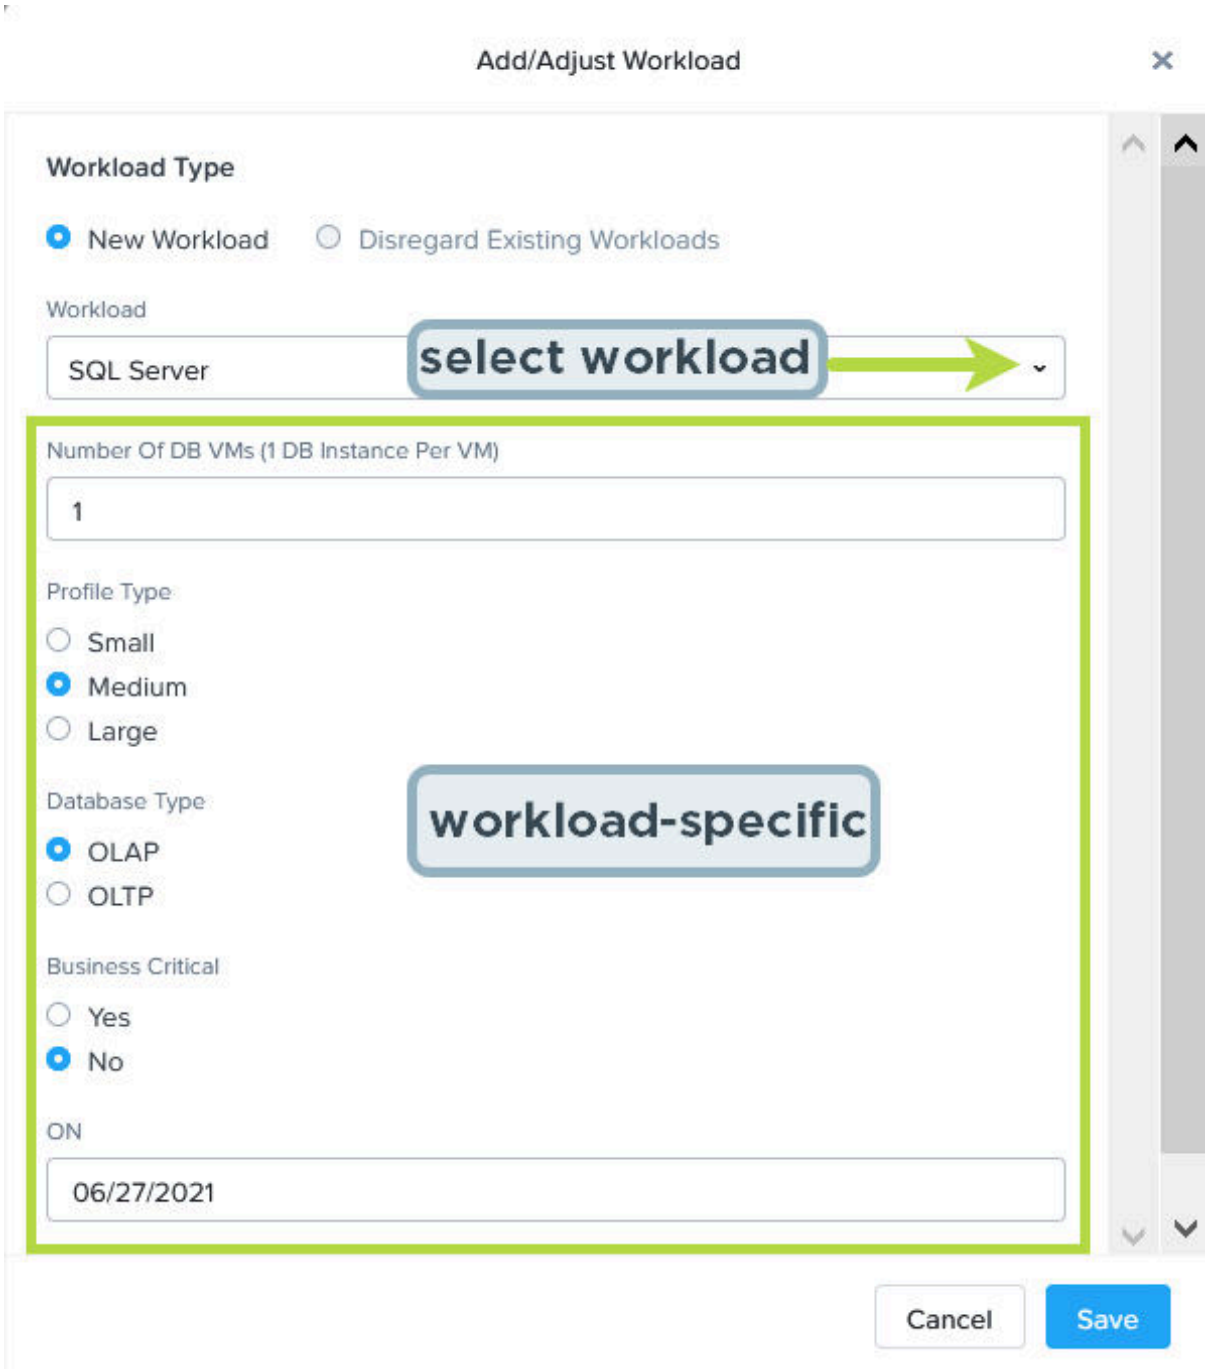

Figure 372: New Workload Fields (SQL Server example)

Table 242: Workload Fields

Workload	Field	Description	Values
(all)	Workload	Select the desired workload type from the pull-down menu.	SQL Server, VMs, VDI, Splunk, Xen App, Change in Demand
SQL Server (This emulates the workload for a specified set of SQL databases.)	Number of Databases	Enter the number of databases to support. The default is 1.	(number)
	Profile Type	Select (click the radio button for) the database size. The default is medium.	Small, Medium, Large
	Database Type	Select the database type. There are two options: • on-line transaction processing (OLTP) is characterized by a large number of short on-line transactions. • on-line analytical processing (OLAP) characterized by a relatively low volume of transactions. This is the default.	OLAP, OLTP
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)

Workload	Field	Description	Values
VMs (This emulates the workload for a specified set of VMs.)	Based On	Select (click the radio button for) the source for the VM profile: Note: If the list of available profiles or VMs is long, you can filter the list by entering a string in the "filter by" field. • **VM Profiles**: Select this option to use an existing profile. Select (click the line for) the desired profile from the provided list in the following table. • **Existing VMs**: Select this option to use an existing VM as the profile. Select the desired VM from the provided list. • **New Profile**: Select this option to create a profile from scratch. For a new profile, enter a name, the number of vCPUs (4 by default), and the amount of memory (4 GiB by default) and HDD (20 GiB by default) capacity in the displayed fields.	VM Profiles, Existing VMs, New Profile
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
	Number of VMs	Enter the number of VMs to add.	(number)
VDI (This emulates a virtual desktop infrastructure workload for a specified number of users.)	Vendor	Select (click the radio button for) the vendor product, either View from VMWare (default) or XenDesktop from Citrix.	View, XenDesktop
	User Type	Select the target user type: task (default), knowledge, power, or developer.	task worker, knowledge worker, power user, developer
	Provision Type	Select how to provision the VMs. The selection of provisioning types vary depending on the selected vendor: • View: full clones, V2V/P2V, linked clones • XenDesktop: full clones, V2V/P2V, provisioning services, machine creation services	full clones, V2V/P2V, linked clones, provisioning services, machine creation services

Workload	Field	Description	Values
	Number of users	Enter the number of users.	(number)
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
Splunk (This emulates a Splunk workload for a specified configuration.)	Daily Average Index Size	Enter a value (number of events) for the daily average index size. The default is 500.	(number)
	Hot Retention Days	Enter the number of days to retain hot data. The default is 7.	(number)
	Cold Retention Days	Enter the number of days to retain cold data. The default is 60.	(number)
	Search Users	Enter the number of search users. The default is 5.	(number)
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
Xen App (This emulates a XenApp workload for a specified configuration.)	Vendor	Select (click the radio button for) the vendor, either Citrix (default) or Microsoft	Citrix (XenApp), Microsoft
	Operating System	Select the operating system to run. The default is Windows 2012R2.	Windows 2012R2, Windows 2008R2
	Provision Type	Select how to provision XenApp. The choice of provisioning types vary depending on the selected vendor: • Citrix: PVS, MCS • Microsoft: VM clone	PVS, MCS, VM clone
	System Data	Enter the amount of system data (in GB), that is the space consumed by each RDSH Server image. The default is 100.	xxx [GB]
	Number of Concurrent Users	Enter the number of concurrent users. The default is 1000.	(number)
	MCS Diff Per VM	Enter a size (in GB) for the MCS difference disk per VM. The default is 20.	xxx [GB]

Workload	Field	Description	Values
	User Profile Data	Enter a size per user (in MB) for the user profile data, which includes the Windows profile and registry settings. The default is 20. The size is typically in the 20-100 MB range.	xxx [MB]
	PVS Write Cache Size Per VM	Enter a size (in GB) for the PVS write cache per VM. The default is 15.	xxx [GB]
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
Change in Demand	Increase/ Decrease	Specify how to change the current workload, that is select (from the pull-down list) whether to increase or decrease the workload.	Increase, Decrease
	By Percentage	Enter the amount you want to change the workload as a percentage of the current workload. Enter the value as an integer (no percent mark).	(1-unlimited)
	On	Enter a date when the workload change should take effect by clicking in the field and then selecting the target date from the displayed calendar. You can make the change effective immediately or at any date in the future.	(date)

Adding a workload creates demand reflected in the Runway chart. Red indicates one or more of the resources are insufficient to support the scenario; blue indicates current resources are sufficient. In the case of a new cluster, the chart is red initially because no resources have been added yet. A scenario based on an existing cluster may or not be red for the applied workload(s) running on the current resources. See Analyzing a Scenario on page 577 for more information.

f. To remove the impact of selected VMs from an existing workload, click the **Add/Adjust** link in the Workload section and then click the **Disregard Existing Workloads** radio button (in the Add/Adjust Workload window).

Select the category of VMs you want excluded from the workload in the **VMs To Disregard** field, enter the starting date in the **Disregard From** field, and then click the **Save** button. Note the following:

- If the selected category contains VMs from multiple clusters, only those VMs from the current cluster are considered for removal.
- If a VM is a member of multiple categories, it is removed just once even when both categories are selected. For example, if VM1 is a member of both Category1 and Category2 and both categories are marked for removal, VM1 is considered only for the

first workload listed for removal, which in this case applies to Category1 (and ignored for Category2).

Note: To use this option, you must first create a category to select (see Category Management on page 413).

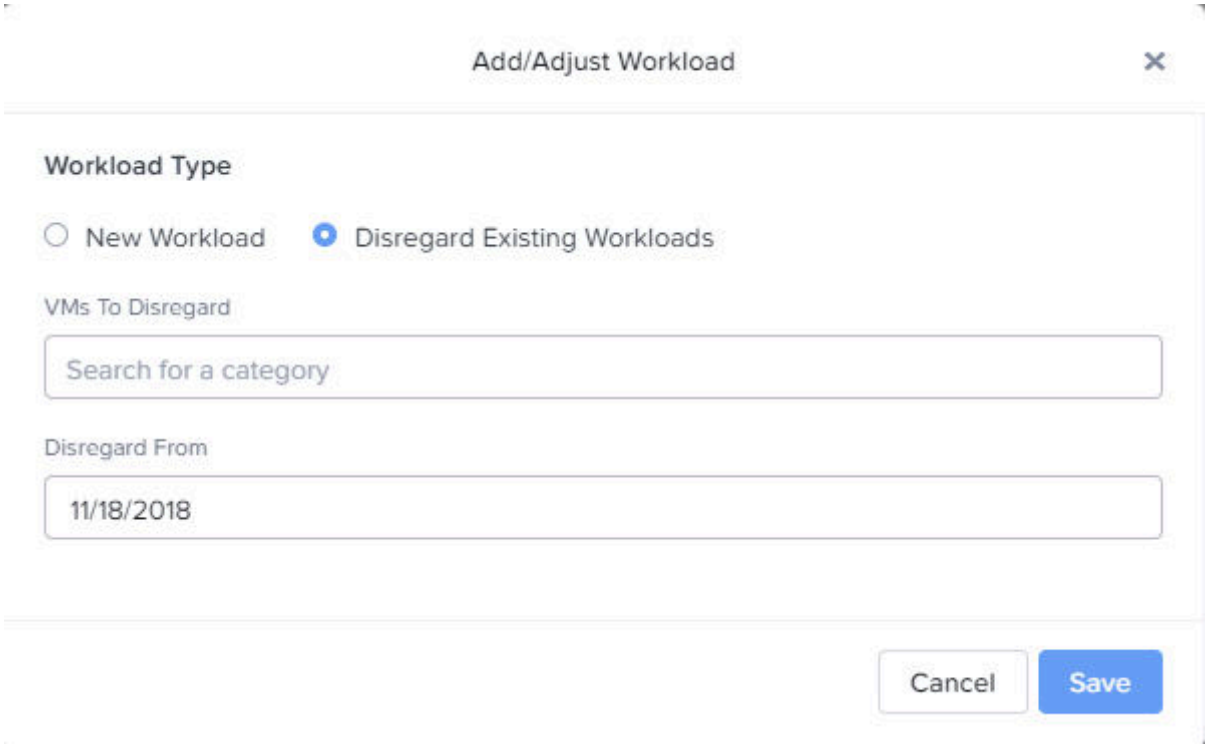

Figure 373: Disregard Existing Workloads Fields

3. To add a resource, do one of the following:

   - Click the **Recommend** button in the Resources section of the screen. The system automatically analyzes the scenario and then adds one or more models to the Resources

list. The Runway chart turns all blue indicating sufficient resources are now available. (The **Recommend** button is active only when at least one resource is red in the Runway chart.)

- Click the **Add/Adjust** button in the Resources section and then click the **New Nodes** radio button in the Add/Adjust Resource window to explicitly add a resource. Enter appropriate information in the following fields and then click the **Save** button.

  - **Model**: Select the model type from the pull-down list. The models match the selected vendor.

    > Note: The options for the following fields are adjusted to match the selected model type.

  - **CPU**: Select the amount of CPU (in GHz) from the pull-down list.
  - **Memory**: Select the amount of memory (in GB) from the pull-down list.
  - **HDD**: Select the amount of HDD capacity (in GB) from the pull-down list.
  - **SSD**: Select the amount of SSD capacity (in GB) from the pull-down list.
  - **# of Nodes**: Enter the number of nodes to include (1-4 depending on the model type).
  - **On**: Enter the starting date for this resource. Clicking in the field displays a calendar from which you can select the date.

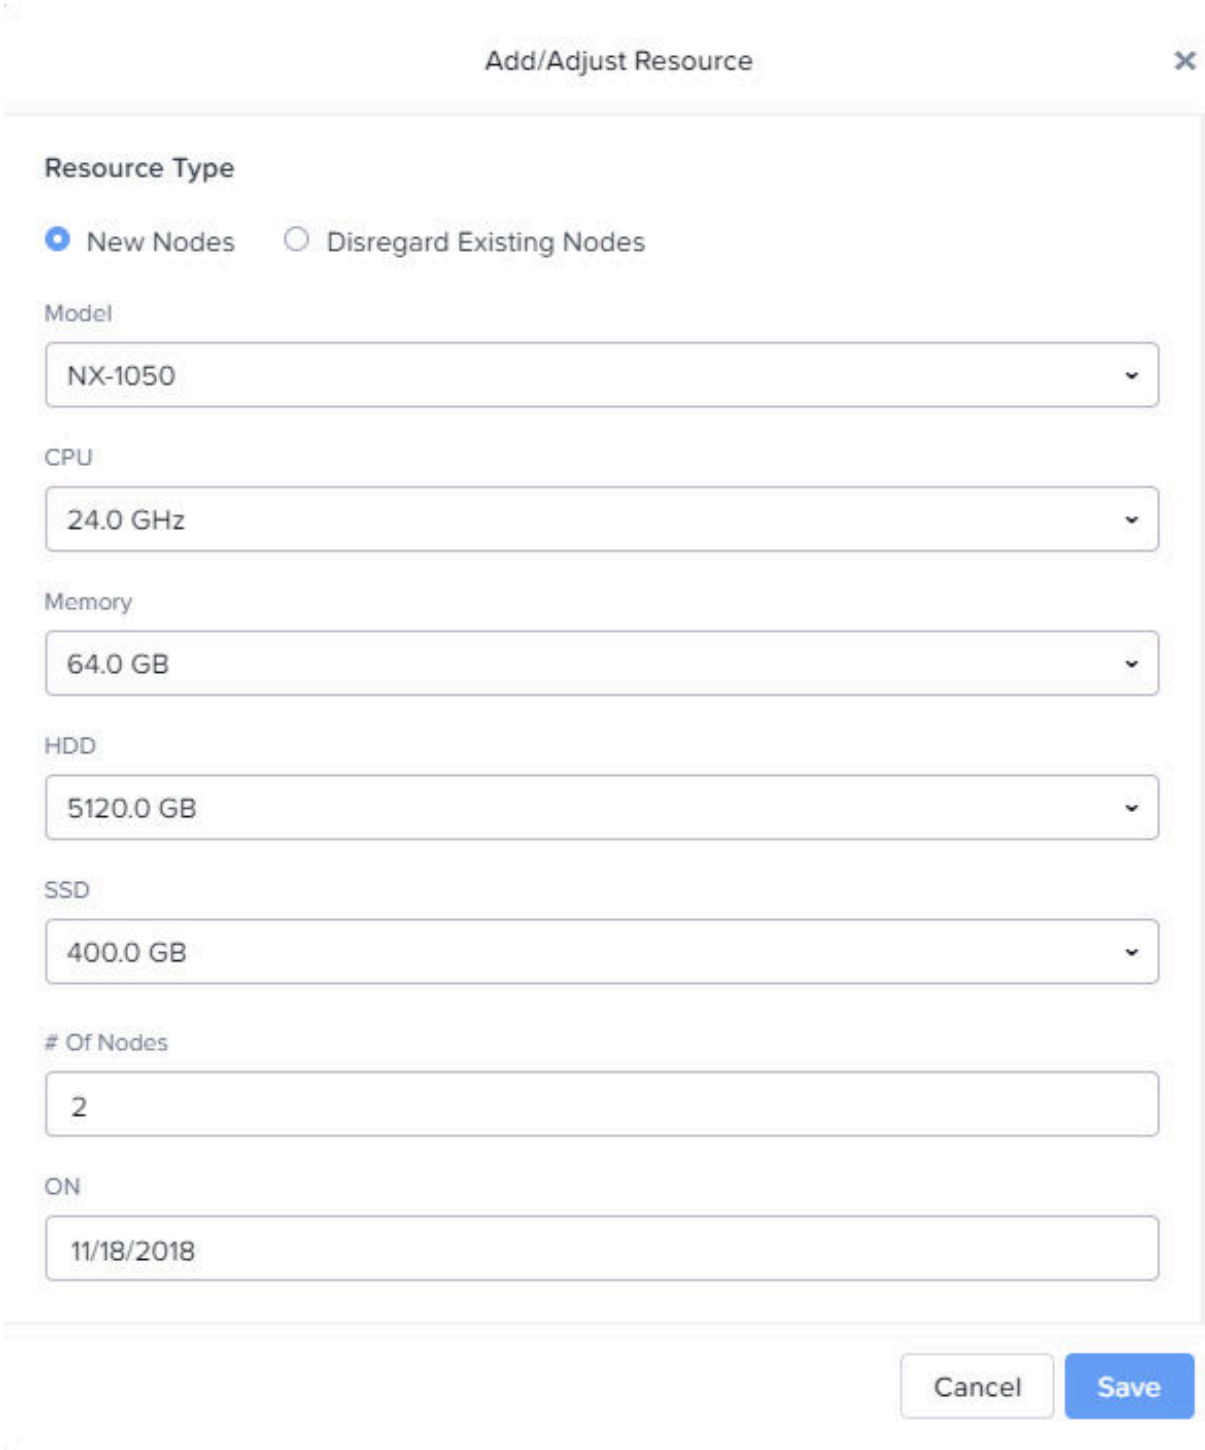

Figure 374: New Nodes Fields

The Runway chart is updated accordingly. However, unlike with the **Recommend** button which adds a resource(s) guaranteed to satisfy the target runway, the resource you add explicitly may or may not satisfy the target runway completely. If it does not, one or more of the resources will remain red in the chart.

Note: You can make changes at any time to the allocated resources and any of the configuration fields to try various scenarios.

4. To remove one or more nodes from a resource, click the **Add/Adjust** button in the Resources section, click the **Disregard Existing Nodes** radio button in the Add/Adjust Resource window, select the node(s) to disregard from the displayed list, and then click the **Save** button.

   The total, disregarded, and effective CPU, memory, and storage are displayed at the bottom. These values are adjusted as you select (unselect) each node to be removed from the scenario. In addition, various stability checks are evaluated when a node is removed including whether this results in too many lost nodes on the same day, system capacity falls below system usages, resources are insufficient to maintain the RF level, and the erasure coding ratio can no longer be maintained (when erasure coding is enabled). When one of the checks fails, an appropriate warning or error message appears.

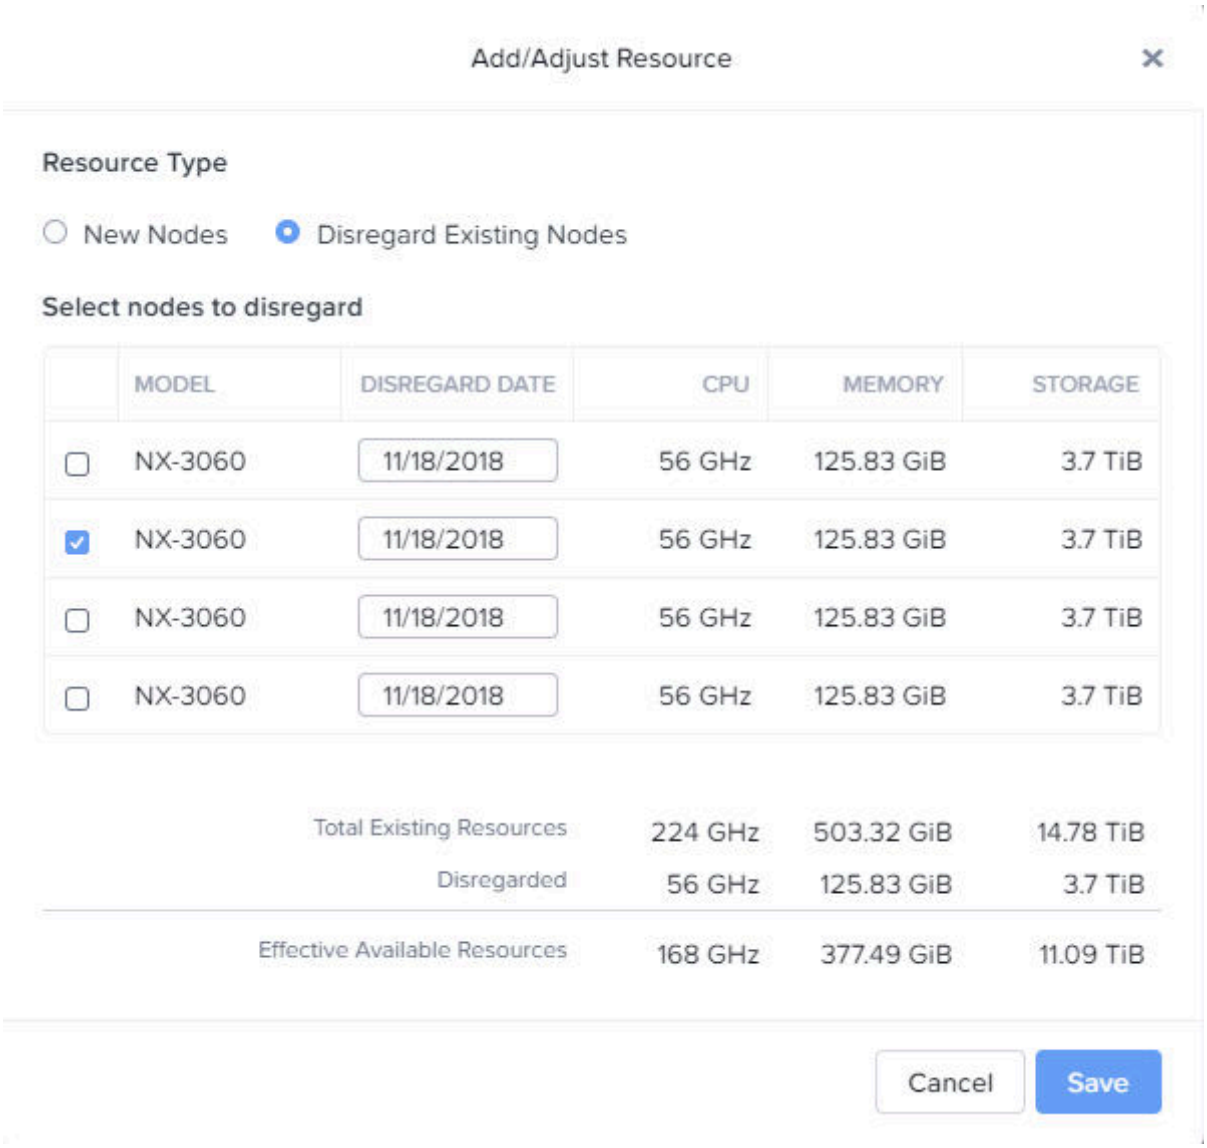

Figure 375: Disregard Existing Nodes Fields

5. To save a scenario, click the **Save Scenario** button (upper right). The saved scenario appears in the Scenarios view list (seeScenarios View on page 563).

6. To download a PDF version of the scenario, click the **Generate PDF** button (upper right).

   If the target runway in the scenario is not satisfied (one or more resources are red), the PDF includes a list of node recommendations that satisfy the runway target.

7. To close the scenario screen, click the X icon (upper left).

Analyzing a Scenario

The scenario screen provides a simple interface to configure and analyze a variety of scenarios.

Example Analysis

To illustrate this process, consider the following example. In this case an existing cluster consists of a NX8150-G4 with four nodes. It has total CPU, memory, and storage capacity that easily handles the current workload.

For this example let's make the target runway six months and add a VDI workload for 100 full clone users that will start in two months. When that workload is added, the runway (previously blue for all) shows that while the current CPU capacity is adequate for the added workload (blue), the memory and storage capacity (red) will not be when the workload starts in two months.

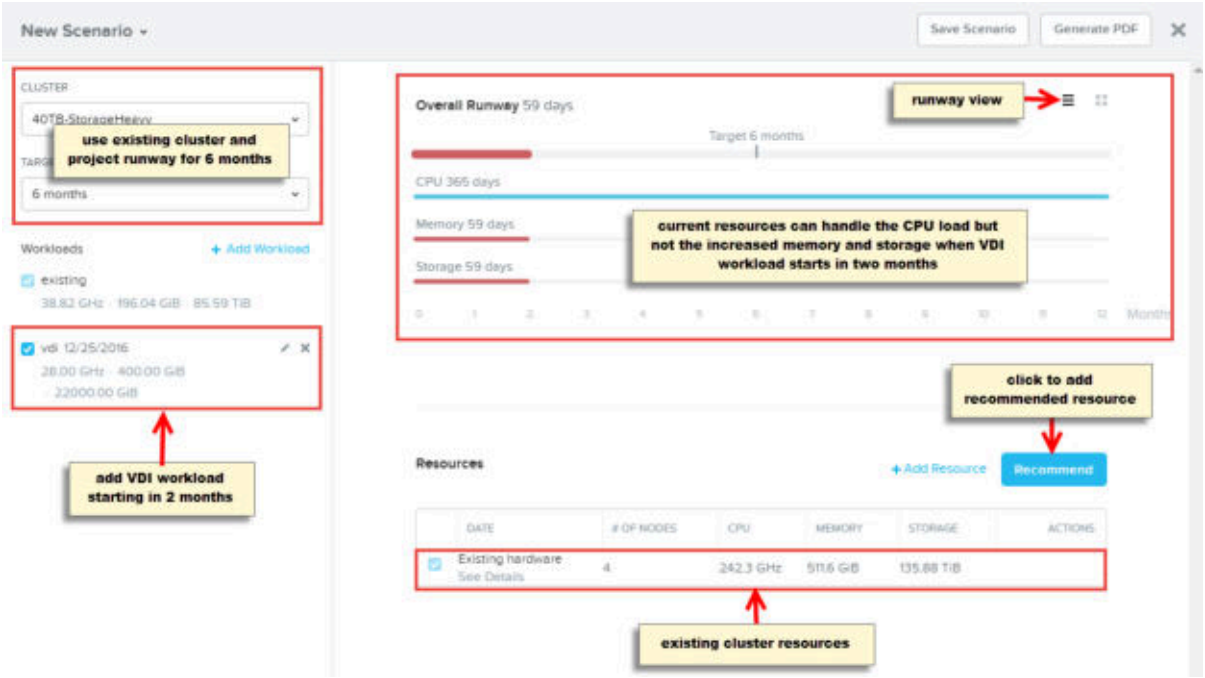

Figure 376: Scenario Screen (with added workload)

In order to meet the VDI workload demand, more resource capacity is required. You can experiment with how to address this deficit by clicking the **+ Add Resource** button and adding a model of your choice, but the simplest approach is to let the planning algorithm determine what to add by clicking the **Recommend** button. The result in this example is a recommendation to add a two node NX-1065-G5. Adding it turns the overall runway blue for all resources.

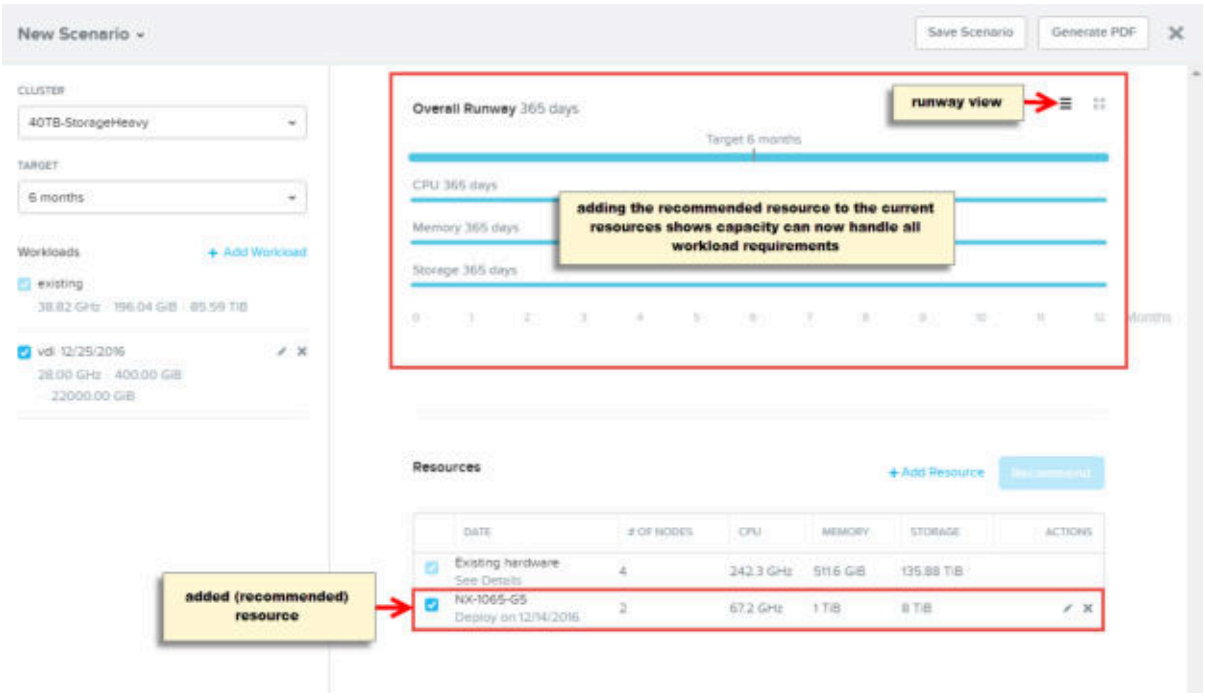

Figure 377: Scenario Screen (with added workload and recommended resource)

You can evaluate a variety of scenarios without leaving the configuration screen. For example, you can change the target runway duration, change the demand by adding or deleting workloads (check/uncheck the box for a workload or resource to add/remove from a scenario without deleting it), modify the workload parameters, or change the allocated resources. In each case the projected runway is updated immediately to reflect the changes you make. This provides a quick and robust method for testing out a variety of scenarios.

Viewing Options

The runway view, which you display by clicking the runway view icon, provides a view of the projected runway overall and by each resource (CPU, memory, and storage). You can also view the usage of each resource by clicking the usage view icon. The following three figures display the storage, CPU, and memory usage views for the example before adding the recommended resource.

- The first two months of each graph (from the left) shows the current workload usage.
- The red dot indicates when the VDI workload begins and the corresponding impact on the resource usage.
- The dotted line indicates the current capacity for each resource. In this case the graphs show that the added workload does not exceed the current CPU capacity but does exceed the storage capacity (by a little) and the memory capacity (by a lot).
- You can use the cursor to move the select point line anywhere on the graph to display the projected usage at that point in time.

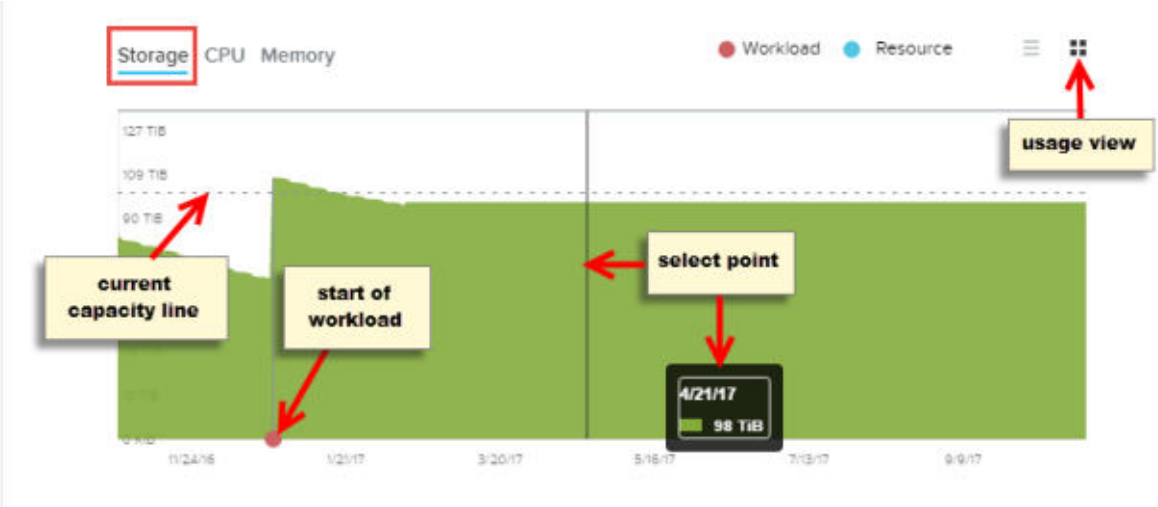

Figure 378: Storage Usage View

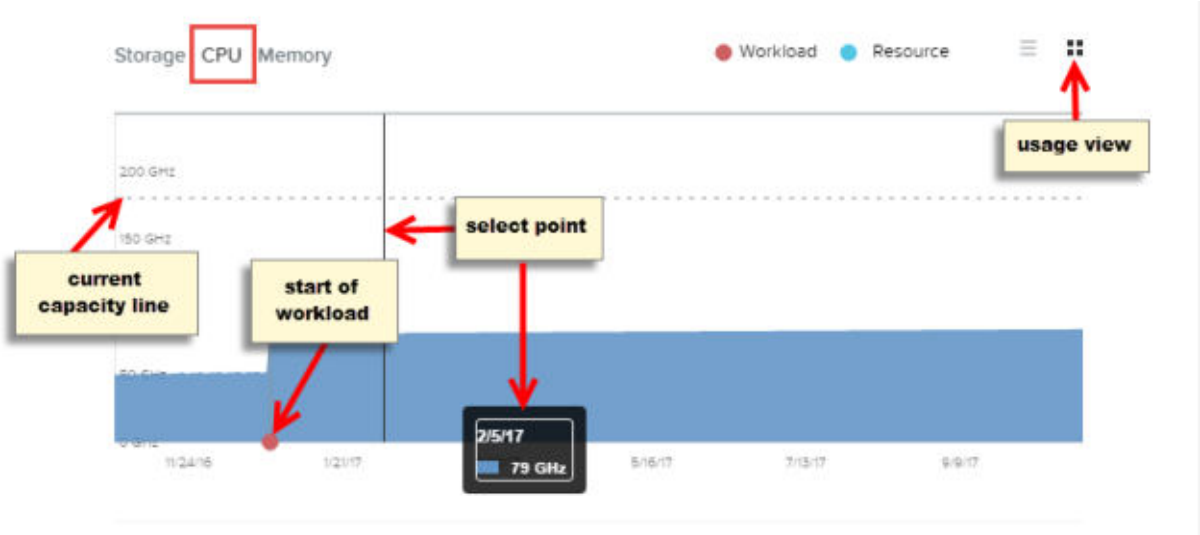

Figure 379: CPU Usage View

Figure 380: Memory Usage View

Modifying a Scenario

About this task

To modify an existing planning scenario, do the following:

Procedure

1. In the Scenarios view (see Scenarios View on page 563), click the pencil icon for the target scenario to display the screen for that scenario.

   Note: The following steps are brief. See Creating a Scenario on page 564 for details and screen shots.

2. To change the target cluster (or start with a new cluster), select the desired cluster from the **Cluster** field pull-down list.
3. To change the runway duration, select the desired length (1-12 months) from the **Target** field pull-down list.
4. To change the reserved capacity, check (or uncheck) the **Capacity configuration** box and click the associated pencil icon to modify the parameters as desired.
5. To change the workload, do one or more of the following:

   - To add a new workload, click the **Add/Adjust** link to open the Add/Adjust Workload window, click the click the **New Workload** radio button, select the desired workload from the pull-down list in the **Workload** field, enter appropriate information in the remaining fields, and then click the **Save** button.
   - To remove the impact of selected VMs from an existing workload, click the **Add/Adjust** link to open the Add/Adjust Workload window, click the click the **Disregard Existing Workloads** radio button, select the category of VMs you want excluded from the

workload in the **VMs To Disregard** field, enter the starting date in the **Disregard From** field, and then click the **Save** button.

- To modify an existing workload, click the associated pencil icon to open the Add/Adjust Workload window to that workload and make the desired changes.
- To delete an existing workload, click the associated X icon.

6. To change the resources, do one or more of the following:

   - To add a new (recommended) resource automatically, click the **Recommend** button.

     Note: The **Recommend** button is active only when at least one resource is red in the runway chart.

   - To add a new (specified) resource manually, click the **Add/Adjust** button, click the **New Nodes** radio button in the Add/Adjust Resource window, enter appropriate information in the displayed fields, and then click the **Save** button.
   - To remove a resource, click the **Add/Adjust** button, click the **Disregard Existing Nodes** radio button in the Add/Adjust Resource window, select the node(s) to remove from the displayed list, and then click the **Save** button.
   - To modify an existing resource, click the associated pencil icon to open the Add Resource window to that resource and make the desired changes.
   - To delete an existing resource, click the associated X icon.

7. To save the changed scenario, click the **Save Scenario** button.

Updating Capacity Configurations

About this task

The Prism Central planning tools (see Resource Planning on page 560) assume by default that all the storage, CPU, and memory capacity in the cluster should be considered when estimating the runway. However, you have the option to reserve some capacity before doing the runway analysis. To reserve storage, CPU, or memory capacity, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Capacity Configurations** in the **Settings** page.
   The Update Capacity Configurations window appears.

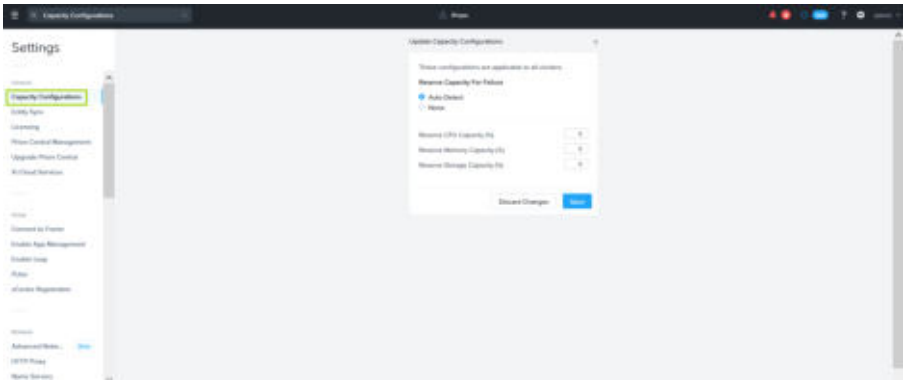

Figure 381: Update Capacity Configurations Window

2. Do the following in the indicated fields:

   a. **Reserve Capacity For Failure**: Click the appropriate radio button to specify whether runway analysis should account for the impact of a node failure.

      - Click **Auto Detect** to account for a node failure in the runway estimates. In this case the reserved capacity percentages (following three fields) refer to the amount of CPU, memory, or storage resources in the cluster minus the largest single node. This means the runway values are computed with enough cushion to account for a single node failure.
      - Click **None** to ignore failure scenarios in the runway estimates. In this case no capacity is reserved for a potential node failure.

   b. **Reserve CPU Capacity (%)**: Enter the percentage of cluster CPU to reserve (0-100).

      Note: The default is 0% of the cluster total for CPU, memory, and storage capacity. Settings a higher value in one or more of these three fields means leaving out that amount of the total cluster CPU, memory, or storage capacity from the runway analysis. For example, setting this field to "10" means10% of the total available CPU capacity is subtracted before doing the runway analysis.

   c. **Reserve Memory Capacity (%)**: Enter the percentage of cluster memory to reserve (0-100).

   d. **Reserve Storage Capacity (%)**: Enter the percentage of cluster storage to reserve (0-100).

3. When all the entries are correct, click the **Save** button to apply the values and close the window.

Behavioral Learning Tools

Prism Central includes machine-learning capabilities that analyze resource usage over time and provide tools to monitor resource consumption, identify abnormal behavior, and guide resource planning. These tools include

- VM "right sizing" where VMs are analyzed and those that exhibit inefficient profiles are identified.
- Anomaly detection to record when performance or resource usage is outside an expected range based on learned VM baseline behavior.
- "Smart" alerts that trigger when specified anomalies are recorded.
- Reports that summarize cluster efficiency.

VM Right Sizing

It is useful to look at the profile of your VMs when analyzing problems in a cluster or assessing future resource needs. This can help you identify VMs that are not optimally configured such as ones that consume too many resources, are constrained, are over provisioned, or are inactive.

The right sizing feature identifies inefficient VMs that fit one of the profiles described in the following table.

Table 243: Inefficient VM Profiles

Type	Description
Bully VM	A "bully" VM is one that consumes too many resources and causes other VMs to starve. A VM is considered a bully when it exhibits one or more of the following conditions for over an hour: • CPU ready time > 5% • Memory swap rate > 0 Kbps • Host I/O Stargate CPU usage > 85% One or more bully VMs might cause cluster performance to degrade. Identifying bully VMs can help in analyzing whether one or more of those VMs are misbehaving or need additional resources.
Constrained VM	A "constrained" VM is one that does not have enough resources for the demand and can lead to performance bottlenecks. A VM is considered constrained when it exhibits one or more of the following baseline values, based on the past 21 days: • CPU usage > 90% (moderate), 95% (severe) • CPU ready time > 5% , 10% • Memory usage > 90%, 95% • Memory swap rate > 0 Kbps (no moderate value) One or more constrained VMs might cause a performance bottleneck. To provide adequate host resources, resize (increase) the constrained VMs.
Over-provisioned VM	An "over-provisioned" VM is the opposite of a constrained VM, meaning it is a VM that is over-sized and wasting resources which are not needed. A VM is considered over-provisioned when it exhibits one or more of the following baseline values, based on the past 21 days: • CPU usage < 50% (moderate) or < 20% (severe) and CPU ready time < 5% • Memory usage < 50% (moderate) or < 20% (severe) and memory swap rate = 0 Kbps To prevent host resource wastage, resize (decrease) the over-provisioned VMs.
Inactive VM	A VM is "inactive" in either of the following states: • "Dead" VM: A VM is considered dead when it has been powered off for at least 21 days. • "Zombie" VM: A VM is considered a zombie when it is powered on but does fewer than 30 read or write I/Os (total) and receives or transfers fewer than 1000 bytes per day for the past 21 days. To prevent host resource wastage, delete dead and zombie VMs.

You can view information about inefficient VMs from the **VM Efficiency** and **Impacted Cluster** widgets on the home or a custom dashboard (see Main Dashboard on page 60 and Adding Dashboard Widgets on page 70) and from the **Capacity** tab of a cluster details page (see Cluster Details View on page 153).

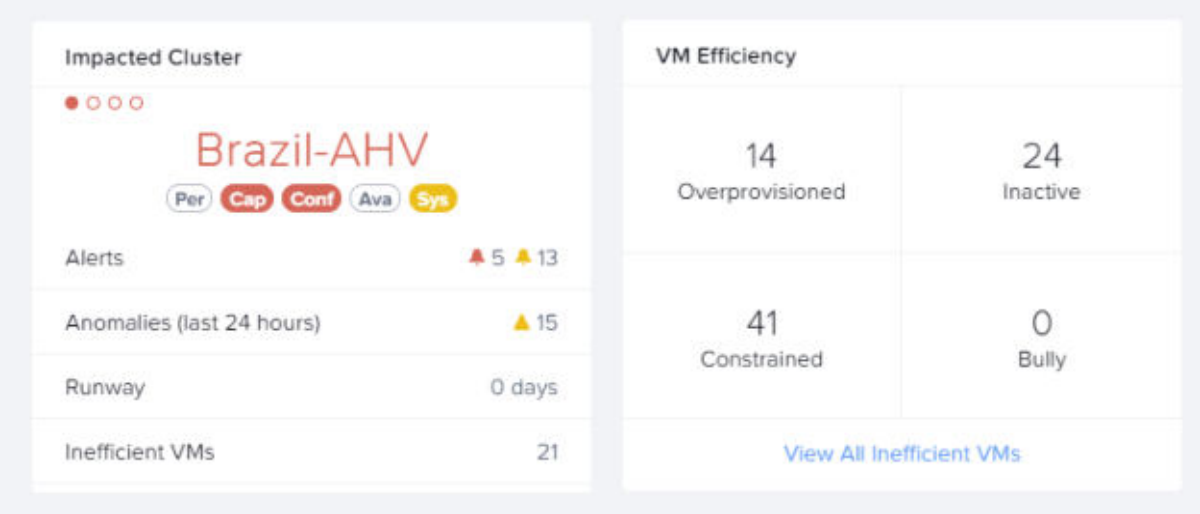

Figure 382: Impacted Cluster and VM Efficiency Widgets

You can also use a search filter from the VMs Summary View on page 82 to identify inefficient VMs.

Figure 383: Filtering VMs

Anomaly Detection

The system predicts a normal behavior band for various metrics based on historical data. The anomaly detection module monitors a predefined set of metrics on a daily basis and publishes baseline values for each of the metrics.

- Twenty-seven metrics are monitored for VMs, hosts, and clusters.

- Data for each metric from the past 21 days is recorded and analyzed, a normal behavior band is established, and predictions for the next 2 days are formulated.
- The behavior bands and predictions are adjusted accordingly when time period or trend patterns are observed, for example low CPU on weekends or increasing CPU usage.

The anomaly detection module measures usage every five minutes and compares that usage with the predicted values. If the observed value is outside the band, it flags that value as an anomaly. Each anomaly is recorded as an event (see Event Details on page 222). Anomalies appears as outliers in the behavioral anomaly event details screen and the VM details Metrics tab.

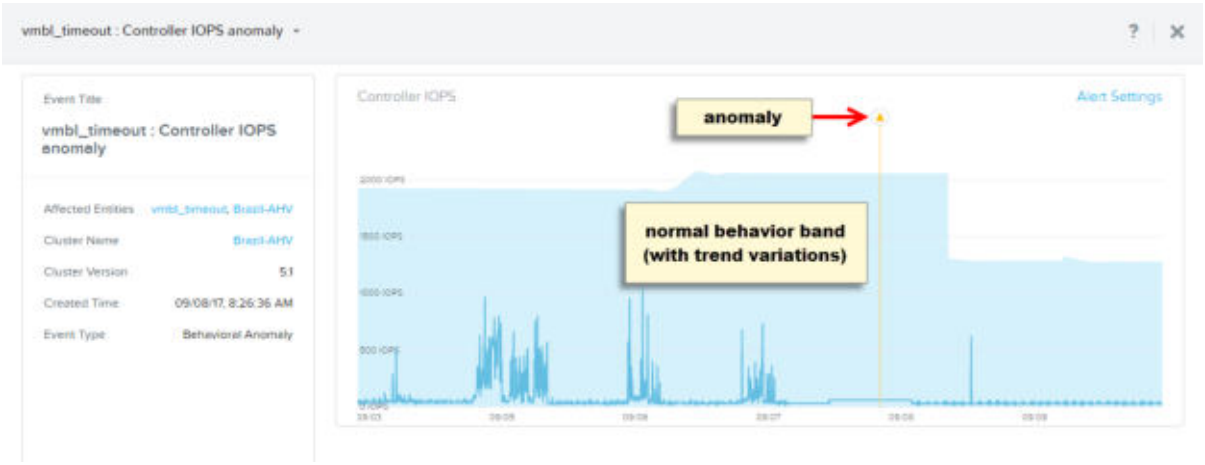

Figure 384: Anomaly Event Details Screen

Smart Alerts

You can create custom policies to generate alerts when behavioral anomalies occur. You can generate a critical or warning alert when a behavioral anomaly occurs for the following conditions:

- a specified VM, host, or cluster
- a specified metric
- every time or only when the anomaly resides within a certain range

See Creating Custom Alert Policies on page 230 for instructions on how to create a custom anomaly-based alert.

Cluster Efficiency Reports

One of the default reports (see Reports View on page 194) is the **Cluster Efficiency Summary** report. This report provides cluster statistics (host and VM counts), performance and usage statistics (average and peak CPU and memory usage), runway metrics, and information about inefficient VMs (bully, constrained, over-provisioned, and inactive). You can run this report to get a quick view of how efficiently the cluster is performing.

Task Automation

You can automate routine administrative tasks through Prism Central by using the X-Play feature.

Overview

The X-Play feature allows you to automate certain administrative tasks.

> Note: X-Play is a Prism Pro feature. If you disable the Prism Pro features, you can still create playbooks and actions, but you can't enable or run the playbooks..

X-Play is an easy to use automation tool that helps you to automate routine administrative tasks, and auto-remediate issues that may occur in your system. You can achieve this automation by creating **Playbooks**.

Playbook allows you to define a trigger that results in the execution of an action or a series of actions. A trigger may be an event that occurs in the system, such as an alert or a request made by you. The resultant actions that you configure can be VM actions, communication actions, alert actions, or report actions.

To get started, you can clone one of the predefined playbooks or create a playbook based on your requirement. You can create playbooks based on predefined alert triggers or you can create a manual trigger (see Creating Playbooks using Triggers on page 596).

You can view the instances of playbook execution using the Plays on page 586 menu. You can see a snapshot view of the executed plays from the **Plays** widget in the **Dashboard**. Additionally, you can view the plays executed in the last 24 hours in the **Impacted Cluster** widget.

Roles

You can perform the following X-Play operations if you have the Prism Admin or the Super Admin role.

- Create, read, update, delete, enable, and run playbooks.
- Create, read, update, or delete action template in the action gallery.
- View executed plays.

If you have the Prism Viewer role, you can only view the playbooks, plays, and action gallery. All other roles do not have permissions to the playbook, play, or action gallery.

Action Gallery

Action gallery provides you a wide range of actions that can be used in Playbooks.

The actions available within the action gallery serve the most common IT automation use cases, see the Playbook Actions on page 587 table for details on the available system actions.

To view the action gallery, go to **Operations** > **Playbooks** > **Action Gallery**.

You can clone system actions from the action gallery, to create custom actions to use in your Playbooks.

> Note: You can not clone actions with any password filled in the action parameter.

You can update or delete any cloned actions later.

Plays

Plays are the execution instances of playbooks that you can use to monitor the status of all the playbooks that you have executed.

To view **Plays**, go to **Operations** > **Plays**. Alternatively, you can view plays from the **Plays** tab in the playbook details page.

You can view the following information within plays.

- Execution start time and execution end time of each action
- Result status
- Any inputs/outputs provided to the action
- Error messages
- Execution Flow (for branch actions)
- Autopilot-enabled playbook statistics and attempt details (see Configuring Autopilot for Playbook on page 599).

You can also select one of the following actions for the paused plays from the **Actions** menu.

- **Resume** - allows you to resume any play that is paused.
- **Abort** - allows you to prevent a paused play from resuming.

Playbook Actions

This section provides information on the playbook actions.

You can switch between the **Tiles** or **List** views by clicking the respective icons.

> Note: The VM actions configured in playbooks have certain prerequisites based on the hypervisor (AHV or ESXi) and the AOS version, see Prerequisites for VM Actions in Playbooks on page 594 for details.

Table 244: Playbook Actions

Action	Action Description	Configurable attributes
Conditional Actions		
Branch Action	Switch execution based on a matching criteria.	See Using Branch Action (Conditional Execution) on page 604 for details.
Data Manipulation Actions		

Action	Action Description	Configurable attributes

String Parser	Parse data from a string. You can use this action to parse response from any other actions like **REST API** or **SSH Action**.	• **String to Parse** - Enter the string.
		• **Format** - Select one of the following options.
		  • **JSON** - Enter the field to retrieve as JSONPath.
		  • **XML** - Enter the node/ tag/text to retrieve as XPath in the **XPath** field. Optionally, enter the comma separated list of namespaces as key=value in the XML Namespaces field.
		  • **Regular Expression** - Enter the string to parse as regular expression

VM Actions		
Power Off VM	Power off a VM.	• Select the VM from the **Target VM** drop-down menu. • Select the **Type of Power Off Action** (Guest Shutdown, ACPI, or Power Off). Note: Guest shutdown requires NGT to be installed on the VM.
Power On VM	Power on a VM.	Select the VM from the **Target VM** drop-down menu.
VM Add Memory	Add memory for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Add Absolute Memory** and **Absolute Maximum**.

Action	Action Description	Configurable attributes
VM Add CPU	Add CPU for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **vCPUs to Add** and **Maximum Number of vCPUs**. See the Prerequisites for VM Actions in Playbooks on page 594 for details.
VM Reduce Memory	Reduce memory for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the value for **Reduce Memory by** and **Minimum Limit**.
VM Reduce CPU	Reduce CPU for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **vCPUs to Remove** and **Minimum Number of vCPUs**.
VM Snapshot	Captures the state of a VM and saves it.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Time To Live** in number of days.
	Captures the state of a VM and saves it.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Time To Live** in number of days.
VM Add Disk	Add vDisk to a VM.	• Select the VM from the **Target VM** drop-down menu. • Select the **Bus Type** and enter the **Disk Size** in GiB.

Action	Action Description	Configurable attributes
VM Expand Disk	Expand vDisk for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Expand Disk By** in GiB. • Ente the **Maximum Disk Size** in GiB. • Choose from: • **Specific Disk** - Select the **Bus Type** and enter the **Index** value for the selected disk type. • **Most Used Disk**
Communication Actions		
Email	Send an email.	Enter the **Recipient**, **Subject**, **Message**, and **Attachment**. Note: The **Attachment** field is applicable only if you have selected **Generate Forecast Report** as a previous action. Note: This action requires the SMTP server to be configured. See *Configuring an SMTP Server* in the *Prism Central Web Console Guide*.
Slack	Send a Slack message.	Enter the **Slack Token**, **Channel ID**, and **Message**. Note: You must integrate Slack to automate receiving messages in Slack, see *Integrating Slack* for details.
Alert Actions		
Acknowledge Alert	Mark an Alert as acknowledged.	• Select the target Alert from the **Target Alert** drop-down menu.

Action	Action Description	Configurable attributes
Resolve Alert	Mark an Alert as resolved.	Select the alert from the **Target Alert** drop-down menu.
Send to ServiceNow	Send an Alert to one of the following ServiceNow target. • ServiceNow Instance	Enter the **ServiceNow Instance Name**, **Username**, and **Password** for the ServiceNow instance.
	• Mid Server	Enter the **Mid Server Address**, **Mid Server Port**, **Username**, and **Password**.
Send Alert to Webhook	Send a JSON Alert Object to a REST API endpoint.	Enter the **Alert**, **Method**, and **URL** for the outgoing webhook. Optionally, select the **Username**, **Password**, and **Request Headers**.
Report Actions		
Generate Forecast Report	Generate a just in time forecast report.	Select the cluster for which you want to generate the report and enter the days for the runway period.
Generate Report	Generate a custom report.	• Select the Report Type as **New Report**. You can also select any of previously created reports to be generated. • Enter the **Report Instance Name**. • Optionally, enter a **Description**. • Select the **Time Period for Report**.
Wait Actions		

Action	Action Description	Configurable attributes
Wait for Some Time	Wait for a specified number of minutes.	Enter the time period in Minutes in the Wait for field. Select the action to be taken at the end of the entered time period. • Resume - automatically resume the play once the wait time is up. • Stop - requires the play to be manually resumed. Note: If a play is not resumed before the wait time is up, then it will be aborted. Optionally, check to enable **Check trigger validity when the Playbook is resumed**. If this option is enabled, the system checks if the alert that caused it to run is still valid or if it has been resolved already.
Wait until Day of Month	Wait until the next occurrence of a specified day and time of the month.	Select the date in the **Wait until** field and set the **Time**. Select the action to be taken at the end of the entered time period. • Resume - automatically resume the play once the wait time is up. • Stop - requires the play to be manually resumed. If a play is not resumed before the wait time is up, then it will be aborted. Optionally, check to enable **Check trigger validity when the Playbook is resumed**. If this option is enabled, the system checks if the alert that caused it to run is still valid or if it has been resolved already.

Action	Action Description	Configurable attributes
Wait until Day Of Week	Wait until the next occurrence of a specified day and time of the week.	Select the day in the **Wait until** field and set the **Time**. Select the action to be taken at the end of the entered time period. • Resume - automatically resume the play once the wait time is up. • Stop - requires the play to be manually resumed. Note: If a play is not resumed before the wait time is up, then it will be aborted. Optionally, check to enable **Check trigger validity when the Playbook is resumed**. If this option is enabled, the system checks if the alert that caused it to run is still valid or if it has been resolved already.
Execution Actions		
IP Address PowerShell	Remotely run a PowerShell script for a given VM IP address.	• Enter the **VM IP Address**, **Username**, **Password**, and the **Path to Script**. • Enable or disable **HTTPS**.
IP Address SSH	Remotely run an ssh command for a given VM IP address.	Enter the **VM IP Address**, **Username**, **Password**, and the **Command to Run**. Note: SSH action has a timeout of 5 minutes. The execution of ssh command should complete within 5 minutes or it should be executed as a daemon. Otherwise, the action will be marked as failed even if command successfully executes after the timeout.

Action	Action Description	Configurable attributes
REST API	Make a call to a REST API endpoint.	• Select the **Method**. • Enter the **URL**, **Username**, **Password**, **Request Body**, and **Request Headers**.
VM PowerShell	Remotely execute PowerShell script for a given VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Username**, **Password**, and the **Path to Script**. • Enable or disable **HTTPS**.
VM SSH	Remotely execute an ssh command for a given VM	• Select the VM from the **Target VM** drop-down menu. • Enter the **Username**, **Password**, and the **Command to Run**.

Prerequisites for VM Actions in Playbooks

The VM actions configured in playbooks have certain prerequisites based on the hypervisor (AHV or ESXi) and the AOS version. Refer to the following table to understand the prerequisites and other dependencies for VM actions defined in playbooks.

> Note: VM actions are not supported for Hyper-V.

Table 245: VM Action Prerequisites

Action	Hypervisor	Prerequisites	Notes
VM Add CPU (Hot Add)	AHV	AOS 5.6 and above	The VM has a guest operating system that supports CPU hot add functionality.
	ESXi	None	The VM has a guest operating system that supports CPU hot add functionality.
VM Reduce CPU	AHV	VM must be powered off	
	ESXi	VM must be powered off	

Action	Hypervisor	Prerequisites	Notes
VM Add Memory	AHV	AOS 5.6 and above (Hot Add)	The VM has a guest operating system that supports Memory hot add functionality.
	ESXi	None	• The VM has a guest operating system that supports Memory hot add functionality. • You cannot increase the memory to above 3 GB due to vSphere restrictions.
VM Reduce Memory	AHV	VM must be powered off	
	ESXi	VM must be powered off	
VM Power On	AHV	None	
	ESXi	None	
VM Power Off (Guest shutdown or ACPI shutdown)	AHV	Nutanix Guest Tools (optional power functions)	Nutanix Guest Tools if you select Guest Shutdown.
	ESXi	VMware guest tools (optional power functions)	
VM Snapshot	AHV	None	• The VM snapshot action creates snapshots for AHV version lower than 5.10 and recovery points for AHV version 5.10 and above. • Maximum Time to Live value allowed is 365 days. • Snapshot can not be created for VMs which are protected by protection domain of legacy DR.

Action	Hypervisor	Prerequisites	Notes
	ESXi	None	• The VM snapshot action creates recovery points for all ESXi VMs. • Maximum Time to Live value allowed is 365 days. • Snapshot can not be created for VMs which are protected by protection domain of legacy DR.
VM Add Disk	AHV	AHV does not support hot add of SATA disks.	
	ESXi	PCI disk type is not supported.	
VM Expand Disk	AHV		
	ESXi	• PCI disk type is not supported. • AOS 5.15 or above	

Creating Playbooks using Triggers

You can create playbooks based on the following triggers.

- Alert Trigger - The **Alert** trigger allows you to select an alert policy as a trigger. When an alert is generated from this policy, the action configured in the playbook is triggered.
- Alerts Matching Criteria - The **Alerts Matching Criteria** trigger allows you to trigger playbook on any alert matching specified criteria.
- Event - The **Event** trigger allows you to trigger playbook when an event of a certain type is generated.
- Manual Trigger - The **Manual** trigger allows you to select an entity type (VM, cluster, or host) to manually trigger the playbook from the **Explore** view for the selected entity type.
- Time - The **Time** trigger allows you to trigger playbook at the specified time.
- Webhook - The **Webhook** trigger allows triggering the playbook from a POST to the webhook URL.

The configured playbooks are available in the **Entities** menu (see Entities Menu on page 13).

Creating Playbooks Using Alert

About this task

To create a playbook using a predefined alert, do the following.

> Note: Alert as trigger allows you to select only a specific alert that can be used in a playbook. To select multiple alerts based on a matching criteria, see Creating Playbooks Using Alerts Matching Criteria on page 599.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Alert** option.
4. In the search text field, enter the name of an existing alert policy.
5. Configure the filter criteria to target the most relevant alerts and the entities associated with the alert.
6. Select the target entity type.

   - All Applicable Clusters - The alerts from all the applicable clusters AND matching the alert criteria are used as trigger.
   - Clusters in Categories - The alerts from only the clusters (or entity type) within the specified categories AND matching the alert criteria are used as trigger.
   - Specific Clusters - The alerts from only the specified clusters AND matching the alert criteria are used as trigger.

7. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.

   - Branch
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   The **Parameters** option allows you to select certain inputs based on the **Trigger** that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.

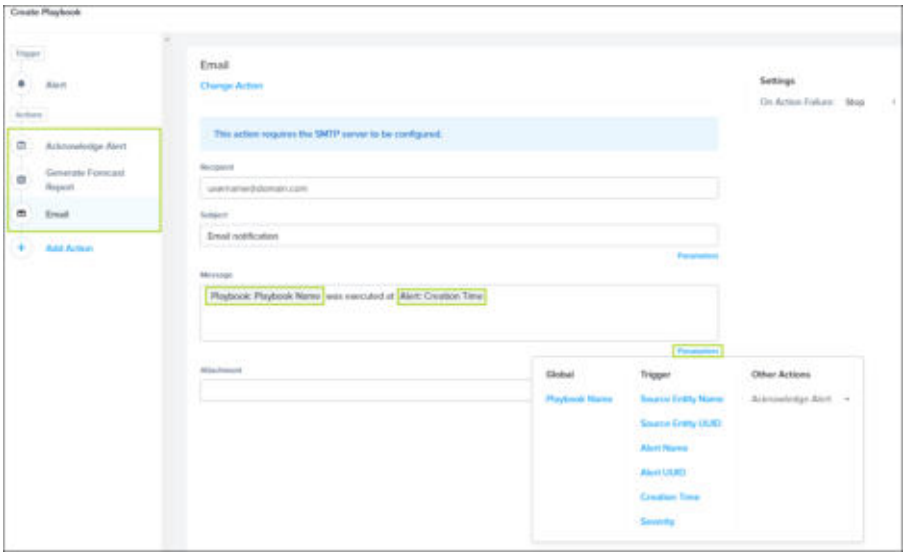

Figure 385: Parameters

To configure an action, set the action attributes based on your requirement. Refer to the Playbook Actions on page 587 table for action attributes and instructions on how to set the attributes. For each action, you must define the one of the following setting that you want to do in the event of action failure.

- Stop - Stop any further execution of the playbook when the action fails.
- Continue - Continue the execution of playbook when the action fails.

> Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action. You can also add actions before or after a particular action.

8. (Optional) Configure Autopilot for playbook.

   The Autopilot feature allows you to control the execution playbooks and automate the optimization of plays by adding specific metrics like key performance indicator (KPI) and optimal threshold for your automation tasks, see Configuring Autopilot for Playbook on page 599.

9. To save the playbook, do the following.
   a. Click **Save and Close**.
   b. Specify the **Name** and **Description**.
   c. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   d. Click **Save**.

Configuring Autopilot for Playbook

The Autopilot feature allows you to control the execution playbooks and automate the optimization of plays by adding specific metrics like key performance indicator (KPI) and optimal threshold for your automation tasks. You can use the autopilot to runs playbooks in the autopilot mode until the specified KPI is m new change

About this task

To configure autopilot for new or existing playbooks, do the following.

Before you begin

> Note: The autopilot feature is only applicable for playbooks that have metric-based alerts as triggers.

Procedure

1. Click **Set Autopilot**.
2. Select the **KPI Metric** from the drop-down menu. For example, for a **VM Memory Constrained** alert, select **Memory Usage (%)** as the KPI metric.
3. In **KPI Target Range**, select the Range option and specify the KPI target range value. Or, you can select **<** (less than) or **>** (greater than) options and specify a fixed value for the KPI target.
4. Enter the **Monitoring Duration** in minutes.
5. For **Max # Attempts**, enter the number of attempts that the playbook must run in the autopilot mode to achieve the KPI target.
6. (Optional) Check **Wait for approval for every attempt** to get an alert (INFO level) for every execution attempt of the playbook.
   If you select this option, you will have to resume the playbook after each run from the **Plays** page. The playbook remains in a suspended state until it is manually resumed after each run.
7. Click the **Autopilot Status** toggle-button to enable or disable the Autopilot. Click Save.

What to do next

You can view the details of the autopilot enabled playbooks from the Plays on page 586 page.

Creating Playbooks Using Alerts Matching Criteria

About this task

To create a playbook using alerts matching criteria as trigger, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Alerts Matching Criteria** option.
4. Select the following for either **All Alert Policies** or **Specific Alert Policies** from the following options.

   - **Impact Type** - Choose **Any** to select all impact types; or choose a single value or a combination of values from **Capacity**, **Performance**, **Configuration**, **Availability**, and **System Indicator** impact types.
   - **Severity** - Choose **Any** to select alerts having any severity levels; or choose a single value or combination of values from **Critical**, **Warning**, and **Info** severity levels.
   - **Cluster** - Choose **All Clusters** for alerts from from all clusters or **Specific Clusters** for alerts from only the specified clusters.

   For **Specific Alert Policies**, the available options are **Policies**, **Severity**, and **Cluster**.

   The actions defined in this playbook will be available for the selected entity type in the entities view.
5. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.

   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

6. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

Creating Playbooks Using Event

About this task

To create a playbook using event, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Event** option.

   > Note: VM related event triggers are applicable only to VMs running on AHV cluster. Event triggers for VM events for ESXi or HyperV VMs are not supported.

4. Select an event type from the events drop-down menu.
5. Select the target entity type by choosing one of the following options.
   - All Applicable VMs - The selected event type for all the applicable VMs is used as the trigger.
   - VMs in Categories - The selected event type for the VMs in the specified categories is used as the trigger.
   - Specific VMs - The selected event type for the specified VMs is used as the trigger.
6. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

7. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

Creating Playbooks Using Time

About this task

To create a playbook using a time trigger, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Time** option.
4. Select the **Trigger Action(s)** from the following occurrence type.
   - **Once** - Specify the day, time, and timezone.
   - **At a recurring interval** - Specify the **Recurrence** interval from daily, weekly, or monthly along with the required interval details.
5. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

6. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

What to do next

Once you create and enable a playbook with time trigger, you can go to the list view and click on the playbook to see the next time it will be executed.

Creating Playbooks Using Manual Triggers

About this task

To create a playbook using manual trigger for a VM, host, or cluster, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.

2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Manual** option.
4. Select the **Entity Type** from the following options.
   - VM
   - Host
   - Cluster

   The actions defined in this Playbook will be available for the selected entity type in the entities view.
5. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

6. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

Creating Playbooks Using Webhook

About this task

To create a playbook using a webhook as trigger do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Webhook** option.

4. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.

   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

5. To save the playbook, do the following.

   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

   After saving, open this playbook to view the instructions for using this trigger.

6. Once saved, open the playbook to view and copy the webhook parameters.

   Using the parameters (**API Method**, **URL**, **ID**, and **Body**) and values shown to create an incoming webhook for the external application to trigger playbook.

Using Branch Action (Conditional Execution)

You can create an X-Play playbook with conditional branch action for a trigger. Branch action allows you to switch execution based on a matching criteria.

About this task

To use a branch action, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Select and configure one of the following trigger.

   - Alert Trigger
   - Alerts Matching Criteria
   - Event
   - Manual Trigger
   - Time
   - Webhook

3. Click **Add Action** and select **Branch**.
4. Define condition for the branch.
   - Branch - Click the pencil icon to add a description for the conditional branch.
   - Condition - Select either **If**, **Else**, or **Else If** depending on the condition context that you want to set.

     Note:
     - You can create multiple branch actions for a trigger.
     - Once the playbook is run, you can view the **Execution Flow** of your playbook. To view the execution flow of any previously run playbook, go to **Operations** > **Plays** and click on the playbook name. To view the complete execution flow, click **View Full Playbook**.
   - Operand - Click **Parameters** to select a **Global** or **Trigger** entity.
   - Operator - Select the required operator from the drop-down menu. You can choose **regexp** for pattern matching.
   - Value - Enter the value for the operator, or click **Parameters** to select a **Global** or **Trigger** entity.

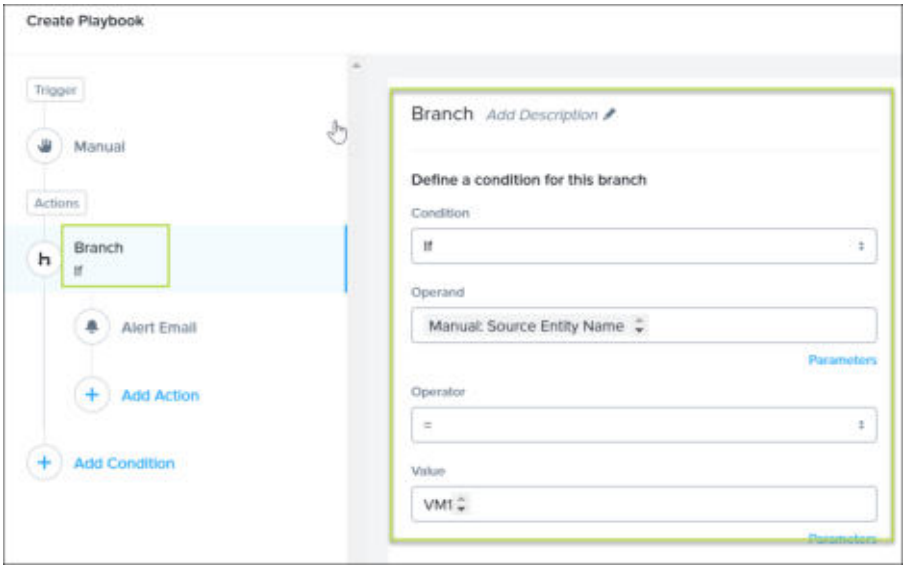

Figure 386: Branch Action

5. Click **Add Action** to view the list of all the available actions on the deck.

   To configure an action, set the action attributes based on your requirement. Refer to the Playbook Actions on page 587 table for action attributes and instructions on how to set the attributes.

   Note: For each action branch condition, you must add at least one action.

6. Create the execution flow using more branch actions (**Else-If** or **Else** options) and adding required actions for each condition.

   You can also configure nested branch actions as required. To delete any added branch action, click the three-dots icon select **Delete**.

7. To save the playbook, do the following.

   a. Click **Save and Close**.

   b. Specify the **Name** and **Description**.

   c. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.

   d. Click **Save**.

Running a Playbook (Manual Trigger)

About this task

You can run a playbook that has a manual trigger associated with the playbook.

Procedure

1. Click the hamburger icon from the main menu and go to **Virtual Infrastructure** > **VMs**.

   To run playbooks with manual trigger for other entities like **Clusters** or **Hosts**, go to **Hardware** > **Clusters** or **Hardware** > **Hosts**.

2. Select a VM, host, or cluster from the list and click **Actions**.

3. From the Actions drop-down menu, select **Run Playbook**.

4. In the **Run Playbook** dialog box, select the playbook that you want to run for the selected entity.

Configuring Manual Parameters

About this task

Parameters are the values that are dynamically filled in at **Playbook** runtime. You can either choose the parameters using the **Parameters** option at the time of creating **Playbooks**, or enter the parameters manually. Configuring manual parameters is useful if you want to use a parameter while cloning an action in the **Action Gallery**.

> Note: This task demonstrates the usage of configuring parameters manually with a scenario on cloning the default **Email**action.

Procedure

1. Go to **Operations** > **Action Gallery**.

2. Select **Email** from the **Actions** list.

3. Click **Actions** and select **Clone**.

4. Edit the **Name** of the **Action**.

5. In the **Recipient** field, enter the recipient email address.

6. In the **Subject** field, enter `Playbook {{playbook.playbook_name}} addressed Alert {{trigger[0].alert_entity_info.name}}`.

   Here, `{{playbook.playbook_name}}` and `{{trigger[0].alert_entity_info.name}}` are the available parameters. See Parameters for Triggers and Actions on page 609 for the complete list of parameters.

7. In the **Message** field, enter `Alert, {{trigger[0].alert_entity_info.name}}, has been addressed. Action has been taken on {{trigger[0].source_entity_info.name}}`.

8. Click **Copy** to complete.

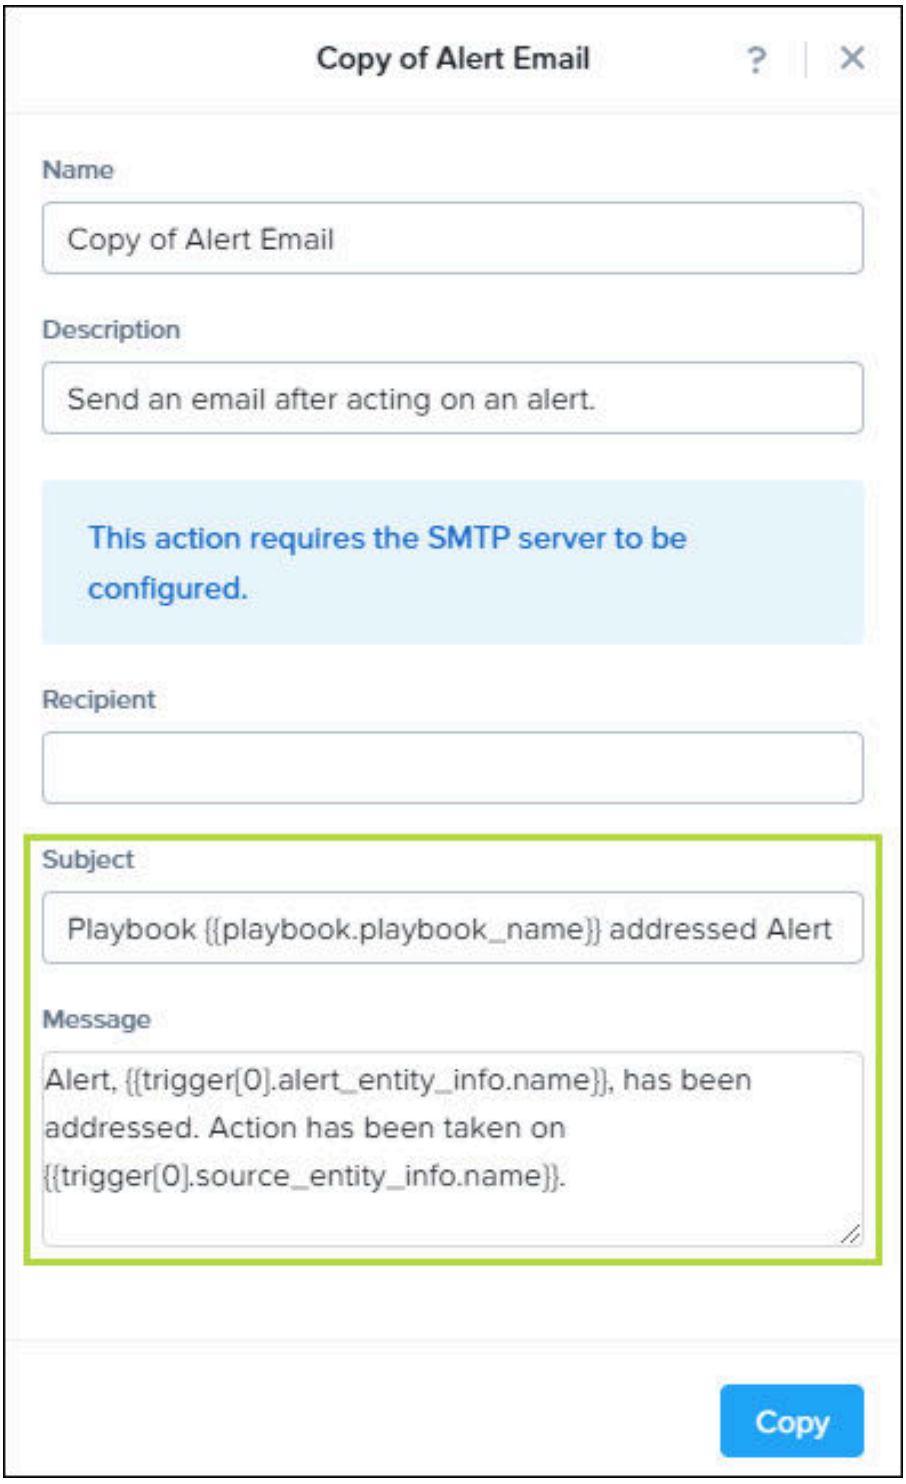

Figure 387: Configuring parameters manually during Clone operation

What to do next

The parameters are dynamically recognized when the cloned action is selected while configuring a **Playbook**.

Note:

- Not all fields accept parameters. The fields with a **Parameters** hyperlink in the Playbook creation workflow accept parameters as inputs.
- If you enter a parameter string into a field that does not accept a parameter, it is read as a string and no value is filled in dynamically.
- Validation errors arise if a parameter of incorrect data type is supplied to a field that is expecting another data type. See Parameters for Triggers and Actions on page 609 for the complete list of supported parameters, accepted data type and parameter format.

Parameters for Triggers and Actions

Refer to the following table for the complete list of parameters available for triggers and actions.

Table 246: Parameters Table

Entity	Parameter	
Playbook		
Playbook Name (string)	{{playbook.playbook_name}}	
Alert Trigger		
Source Entity Name (string)	{{trigger[0].source_entity_info.name}}	
Source Entity UUID (string)	{{trigger[0].source_entity_info.uuid}}	
Source Entity (object)	{{trigger[0].source_entity_info}}	
Alert Name (string)	{{trigger[0].alert_entity_info.name}}	
Alert UUID (string)	{{trigger[0].alert_entity_info.uuid}}	
Alert (object)	{{trigger[0].alert_entity_info}}	
Creation Time (long)	{{trigger[0].creation_time}}	
Severity (string)	{{trigger[0].severity}}	
Manual Trigger		
Source Entity Name (string)	{{trigger[0].source_entity_info.name}}	
Source Entity UUID (string)	{{trigger[0].source_entity_info.uuid}}	
Source Entity (object)	{{trigger[0].source_entity_info}}	
Playbook UUID (string)	{{trigger[0].action_rule_uuid}}	
Action		
Generate Forecast Report	Report PDF (file)	{{action[index].report}
	Report Name (String)	{{action[index].report.fileName}}
	Report Data (bytes)	{{action[index].report.fileData}}
IP Address Powershell	Return Code (long)	{{action[index].ret_code}}
	Standard Output (string)	{{action[index].stdout}}

Entity	Parameter	
	Standard Error (string)	{{action[index].stderr}}
IP Address SSH	Return Code (long)	{{action[index].status_code}}
	Standard Output (string)	{{action[index].response_headers}}
	Standard Error (string)	{{action[index].response_body}}
Slack	Channel Id (string)	{{action[index].channel_id}}
	Message Timestamp (string)	{action[index].timestamp}}
VM Add CPU	Resultant vCPU Count (long)	{{action[index].result_vcpu_count}}
VM Reduce Memory	Resultant Memory (double)	{{action[index].total_memory}}
VM Snapshot	Recovery Point UUID (string)	{{action[index].vm_recovery_point_entity.
	Recovery Point Name (string)	{{action[index].vm_recovery_point_entity.
	Recovery Point (object)	{{action[index].vm_recovery_point_entity}
VM SSH	Return Code (long)	{{action[index].ret_code}}
	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
Wait for Some Time	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}
Wait Until Day of Month	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (String)	{{action[index].planned_resume_time}}
Wait Until Day of Week	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}

Exporting or Importing Playbooks

You can export a playbook for backup reasons or import any previously exported playbook. If you export a playbook using a Prism Central instance, you can import that playbook to another Prism Central instance that is running the same or newer version only.

About this task

To export a playbook, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Select one or more playbooks by clicking the check-box against the playbook name
3. Click **Actions** > **Export**.
4. Optionally edit the default **Filename** and click **Export**.
   The playbook backup is downloaded in the .pbk format.

About this task

To import a playbook, do the following

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook** and click **Import**.
2. Click **Browse** to select the previously exported playbook backup file and click **Import**.

Integrating Slack

If you do not have a Slack account, create your account on the Slack website and set up a new workspace.

About this task

To integrate Slack for playbook actions, do the following.

> Note: You can follow these steps in an already existing workspace, however you will require permission from the workspace administrator to complete the installation steps.

Procedure

1. Go to https://api.slack.com/apps.
2. To create a Slack application, click **Create New App**.
3. In the **Create a Slack App** dialog box, do the following.
   1. Enter a name for the application.
   2. Select a development Slack Workspace. You must select the workspace on which you have the authority of accepting messages.
4. Set the app permissions - In the **Basic Information** menu, do the following.
   1. Click **Add features and functionality**.
   2. Go to **Permissions** > **Scopes** and Add **chat:write** or **chat:write.public** permission according to where you expect the application to post messages.
5. To install the application in your workspace, do the following.
   1. In **Basic Information** menu, click **Install your app to your workspace**.
   2. Click **Install App** (or **Reinstall App** if you have configured it previously). This redirects you to the **My Workspace** authorization page.
   3. Click **Authorize**.
6. To retrieve the Slack authorization token, do the following.
   1. Go to the **Install App** menu.
   2. Copy the **OAuth Access Token**. You must paste the copied token each time you use the **Slack** action in a playbook.

ServiceNow Integration with Prism Central (X-Play Support)

ServiceNow Integration Overview

ServiceNow provides IT services management (ITSM), and IT operations management (ITOM) as a software service.

ServiceNow integration with Prism Central allows seamless incident management for Prism Central alerts using ServiceNow. Export of alerts from Prism Central is achieved by creating an

X-Play playbook in Prism Central. The X-Play Action (Send to ServiceNow) pushes the Prism Central alerts to ServiceNow to be processed by the ServiceNow event management pipeline.

> Note: To integrate Prism Central with ServiceNow, enable the **Event Management** plug-in on your ServiceNow instance .

Enabling Event Management Plug-in on ServiceNow (Production Instance)

The ServiceNow **Event Management** plugin requires a separate subscription and can be activated by installing it from the ServiceNow store. . See ServiceNow documentation at https://docs.servicenow.com to enable the plugin for your production instance.

Enabling Event Management Plug-in on ServiceNow (Developer Instance)

About this task

To enable X-Play support on the ServiceNow developer instance, do the following.

Procedure

1. Go to the ServiceNow Developer Portal and create an account.
2. Log on to your account and go to **Manage** > **Instance**.
3. On the instance page, request a ServiceNow instance.
4. Once instance is provisioned, select **Action** > **Activate Plugin** for your instance.
5. On the dialog box that appears, enable the **Event Management** plug-in.

Configuring Alert Playbook for ServiceNow Alert Management

About this task

About this task

You can create a **Playbook** to send Prism Central alerts to ServiceNow.

Procedure

1. In Prism Central, click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Create Playbook**, then click **Select** in the **Alerts Matching Criteria** (option) card.
   - **Alert**: Allows you to select an alert policy as a trigger. When an alert is generated from this policy, the action configured in the playbook is triggered. For more information, see Creating Playbooks Using Alert on page 597.
   - **Alerts Matching Criteria**: Allows you to trigger playbook on any alert matching specified criteria such as severity levels, impact type, and so on. For more information, see Creating Playbooks Using Alerts Matching Criteria on page 599.
3. Click **Add Action**. Click **Select** in the **Send to ServiceNow** in the actions deck.

4. Configure the ServiceNow action.

   Note: This action requires the *Event Management Plug-in* to be enabled on the ServiceNow instance.

   a. Click **Parameters** and select Alert Matching Criteria: Alert in the **Alert** dropdown menu.

   b. Enter any additional information in the **Additional Info (Optional)**

   c. Select the appropriate ServiceNow instance based on the ServiceNow deployment for the **Send to** parameter.

      Select **Cloud Instance**, **Private Instance** or **MID Server** .

      If you select **MID Server**, you need to provide the MID Server IP address and Port details in next sections

      Note: The option to send the alert to ServiceNow is a default action in X-Play that calls the ServiceNow Event Collection API when a Nutanix alert matches the criteria specified.

   d. Enter the **ServiceNow Instance Name** in case of **Cloud Instance** or **Private Instance**. Enter the **MID Server IP Address** and **MID Server Port** details in case you selected **MID Server** in previous step.

      Note: Do not enter the complete URL of your ServiceNow instance in the **ServiceNow Instance Name** field. If your instance URL is *https://foo.servicenow.com*, the **ServiceNow Instance Name** you must enter is *foo*.

   e. Enter the **Username** and **Password** for the ServiceNow instance.

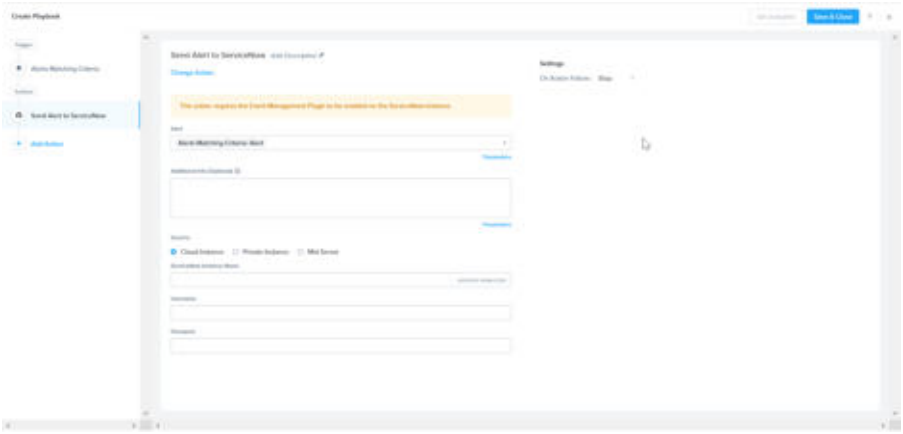

Figure 388: ServiceNow Action

5. To save the playbook, do the following.

   a. Click **Save & Close**.

   b. Specify the **Name** and provide a **Description**.

   c. Click the slider to enable the playbook. Prism Central executes only the playbooks that are enabled.

   d. Click **Save**.

   All Prism Central alerts of the specified type are forwarded to ServiceNow with the execution of the playbook created above.

What to do next

Log on to ServiceNow to view Nutanix alerts in the **Alert Intelligence All Alerts** menu after completing Importing Nutanix Event Rules to ServiceNow on page 614 and Activating the Create Incident Rule on page 614.

You can also create playbooks for actions that you want ServiceNow to initiate on Prism Central for the alerts sent to Service now. See .

Importing Nutanix Event Rules to ServiceNow

Ensure that you have imported Nutanix event rules to ServiceNow after you setup the alert playbook.

Before you begin

In ServiceNow, go to **Event Management** > **Event Rules** and check if the Nutanix event rules are already available.

The Nutanix event rules are provided in the Nutanix Event Rule Creation xml file. Download the *Nutanix Event Rule Creation* xml file from download.nutanix.com or contact Nutanix Support.

About this task

To import event rules using the *Nutanix Event Rule Creation* xml file, do the following.

Procedure

1. In ServiceNow, go to **System Update Sets** > **Retrieved Update Sets**.
2. On the right pane, click **Import Update Set from XML**.
   Import Update Set from XML is a clickable link on the right pane.
3. Browse and select the Nutanix Event Rule Creation xml file that you downloaded. Click **Upload**.
4. On the Update Sources page, click the uploaded Update Set to open update set details page.
5. Click **Preview Update Set**.
   If you find errors after the preview run is complete, click **Accept Remote Update** for each error. This resolves the error. Resolve all the errors.
6. After resolving all the errors, click **Commit Update Set**.

What to do next

Go to **Event Management** > **Event Rules**, Select **Source** in the search dropdown and type *Nutanix* in the **Search** box. The Nutanix Event Rules are now displayed.

Activating the Create Incident Rule

You need to activate the **Create Incident** rule in ServiceNow.

About this task

The incidence creation does not occur in ServiceNow until you set the active status of the **Create Incident** rule in Alert Management Rules in Event Management of the ServiceNow instance to true (Active status being false by default).

To change the Active status of the **Create Incident** to true, do the following:

Procedure

1. Go to **Event Management** > **Alert Management**. A list of Alert Management Rules is displayed.
2. Click the rule named **Create Incident**. Click (to check) the **Active** status checkbox below the rule name to set the **Create Incident** rule status to active.

Figure 389: Activate the Create Incident Rule

Configuring a Webhook Playbook for ServiceNow Action communications to Prism

X-Play now supports a webhook feature to invoke a playbook through REST calls. You can create webhook based playbooks for all the actions that you want ServiceNow to initiate on Prism Central for the alerts sent to ServiceNow.

About this task

Create webhook based playbooks using the **Webhook** trigger.

> Note: You can add one or more actions that you want SNOW to perform on Prism in one playbook.

To create a webhook based playbook, do the following.

Procedure

1. In Prism Central, click the hamburger icon from the main menu and go to **Operations** > **Playbook**.

2. Click **Create Playbook**, then click **Select** in the **Webhook** (option) card.

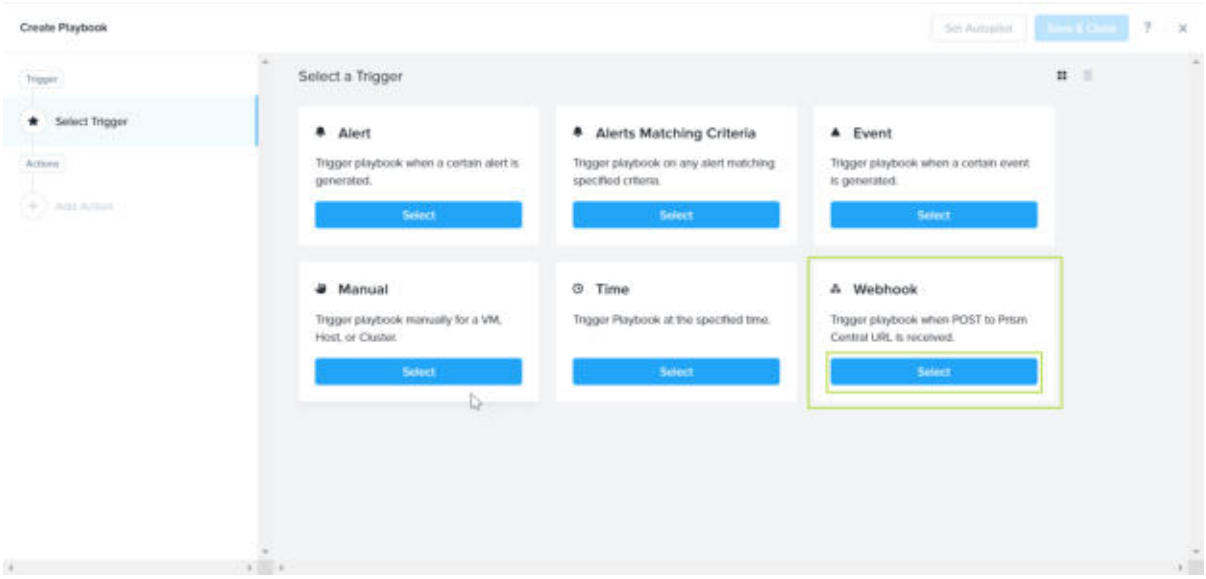

Figure 390: Select a Trigger

The Webhook page displays a message informing you to return to this playbook and review the instructions configured for this trigger.

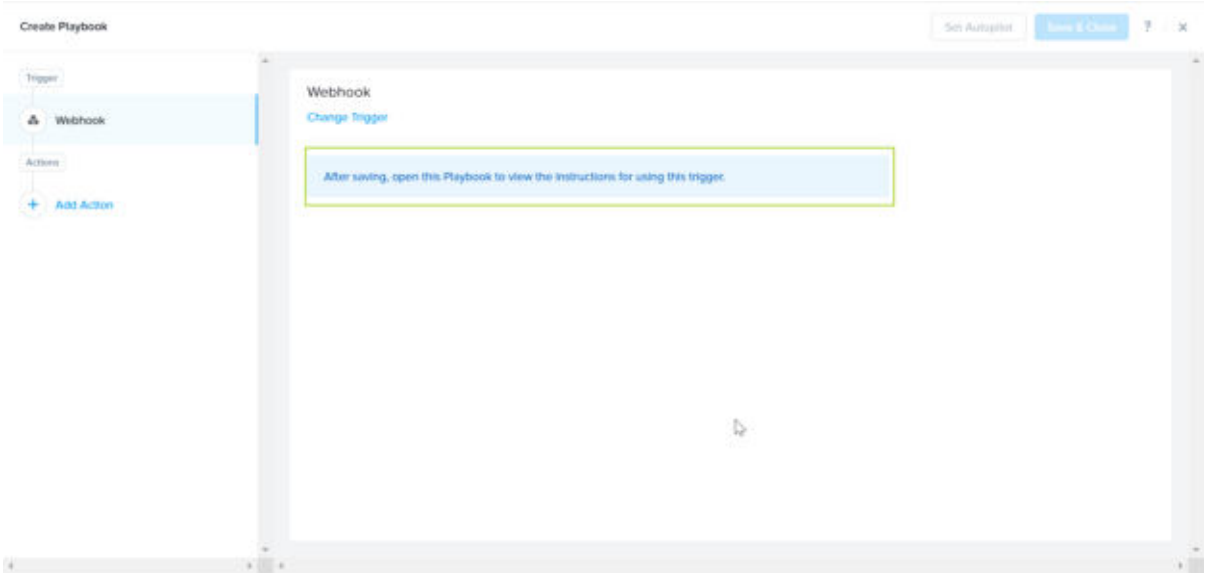

Figure 391: Webhook page

3. Click **Add Action**. Click **Select** in the appropriate action card(s) in the X-Play actions deck.

   You can select one or more actions for the playbook. After you complete configuring an X-Play action, click **Add Action** to add another action to the playbook.

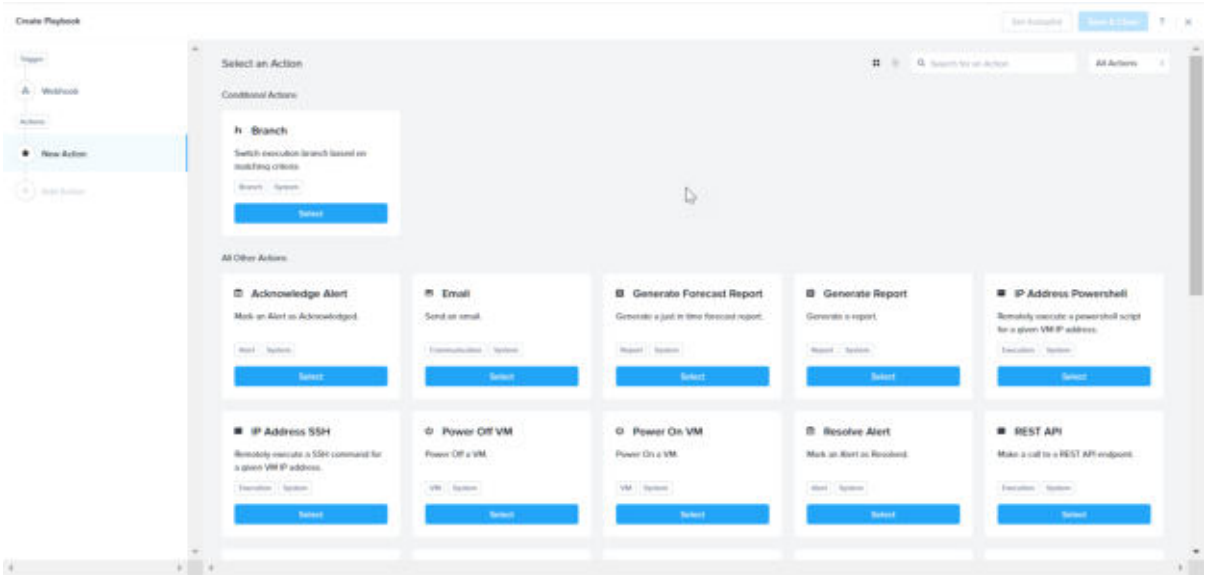

Figure 392: Select an Action page

   When you click **Select** on a X-Play action card, the action page opens. Provide all the necessary values for the parameters on the action page.

4. To save the playbook, do the following.

   a. Click **Save & Close**.

   b. Specify the **Name** and provide a **Description**.

   c. Click the slider to enable the playbook. Prism Central executes only the playbooks that are enabled.

   d. Click **Save**.

What to do next

Open the playbook you just created. Notice that the Webhook section is populated with the URL details and the body of the webhook. The URL of the webhook points to Prism Central; each webhook has a unique ID. Call this URL with the requisite parameters to trigger the playbook.

Reports Management

The report management feature provides you with an ability to configure and deliver the historical reports containing information about the infrastructure resources. This feature provides the operational insights of your infrastructure into your mailbox according to the schedule that you have configured.

A generated report can be divided into three main components.

- Style: Defines the way to represent the report, for example background color, logos, and so on.

- Representation of data: Defines the way that you want to represent the data. You can accomplish this by including different views (pre-defined and customizable).
- Data: Defines the actual data depending on your selection, for example information about different metrics; CPU usage, memory usage, IOPS bandwidth, VM count, host count, cluster count, license summary, and so on.

By default the **Reports** dashboard displays the **Cluster Efficiency Summary** and **Environment Summary** reports. The cluster efficiency summary report provides a detailed information about the predictive utilization and runway information of all the resources. The environment summary report provides information about the summary configuration of each cluster (host count, VM count, license information, CPU usage information, IOPS bandwidth, Controller VM I/O bandwidth) along with the licensing information that is registered to Prism Central.

> Note: You cannot modify or delete the custom (pre-defined or factory shipped) cluster efficiency summary and environment summary reports. However, you can clone the report and then modify the cloned reports. For more information about creating a custom reports, see Creating a New Report on page 619.

> Note: All the built-in roles except the **Prism Viewer** role can manage the reports (create, update, delete). Users with **Prism Viewer** privileges can only view the reports.

Internationalization in simplified Chinese and Japanese is supported. The generated report can have English, Chinese, and Japanese characters.

Some of the salient features of report management are as follows.

- Ability to generate and customize end-to-end reports.
- Ability to add different views to customize what data is displayed and how that data is represented.
- Ability to download the report in PDF and CSV format.
- Create a report definition from another report definition. You can use this feature to save time for a report that has a minor change from an existing one and copy from a shared read-only report from the other users.
- Customize the report definition. Ability to add a logo and make changes to the appearance of the report generated.
- Ability to send the reports through Email. Generated report in the PDF format or CSV format, or both, are sent as an attachment to the Email.
- Ability to schedule the report. Note that you can define only one schedule for a report definition.
- Ability to retain the reports for the specified period of time.
- Ability to check the report log details for the status of the report and any error message (if the report generation fails).
- Ability to create Role Based Access Control (RBAC) to define different roles and assign permissions accordingly.

Requirements

- License requirements: Prism Central should be running Prism Pro license.

Limitations

The report management feature in Prism Central has the following limitations.

- The report management feature currently only supports generation of reports in PDF format for 2000 entities. Nutanix recommends generating reports in CSV format for such scenarios.
- Generation of a report in CSV format supports a maximum of 25000 entities.
- Generation of a report in PDF format can only consist of a maximum 10 columns for a data table.
- The entire process of report generation fails if either of the report formats (PDF or CSV) fails to generate. This behavior is observed when both the formats are selected during report generation.

Alerts Generation

Alerts are generated during following failure scenarios.

- Generating a report
- Sending of a scheduled Email

Creating a New Report

Perform the following to create a new report.

Procedure

1. In the **Reports** dashboard (see Reports View on page 194), click **New Report**.
2. Add the views that you want to add to the report. For information about different views, see Adding Views to a Report on page 626.
3. Configure a schedule for the report. For more information about scheduling, see Scheduling a Report on page 643.
4. Configure the settings specific for this report. For more information about configuring report settings, Configuring Report Settings on page 645.

5. To save the report, do one of the following.

   - Click **Save**.

     This option saves the report and does not generate the report.

   - Click **Save and Run Now**

     .

     This option saves and generates the report. A **Run Report** window appears. For more information, see Generating a Report on page 620.

   The report appears in the **Report** dashboard.

   For example, in the following image, a new report is configured with different views (for example, bar chart, a metric summary, and a data table).

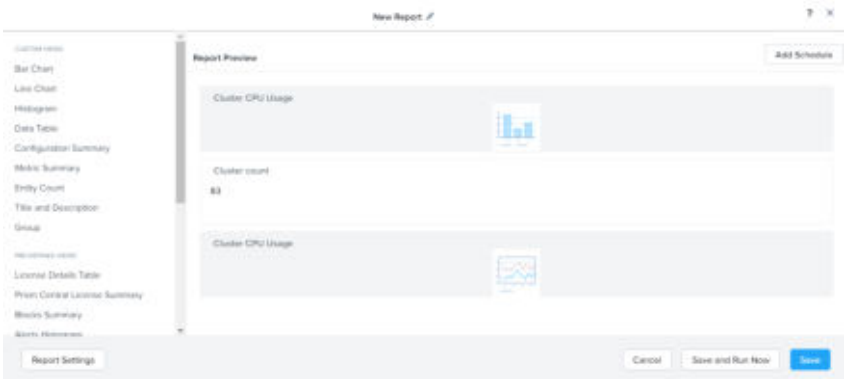

Figure 393: Create Report

Managing a Report

After you create a report, you can generate the report, edit the report, view instances of the report, clone the report, or delete the report.

- To generate a report, see Generating a Report on page 620.
- To edit a report, see Editing a Report on page 621.
- To clone a report, see Cloning a Report on page 622.
- To delete a report, see Cloning a Report on page 622.

Generating a Report

You can generate a report instance while creating a new report or on an existing report.

About this task

To generate a report, do the following.

Procedure

1. In the **Report** dashboard (see Reports View on page 194), click the check box against the report on the list.
2. From the **Actions** drop-down menu, select **Run**.

3. Enter the name of the report (new instance of the report) in the **Report Instance Name** field.
4. Select the duration of data that you want to display in the report from the **Report Time Period** drop-down menu.
   You can display the data for the last 24 hours, last week, or last month. You also have an option to select the date range.

   > Note: If you select **Custom Date Range** from the **Report Time Period**, you need to select the date range from the **From** and **To** option. If you do not specify the date and time period, the data is selected from the configured report.

   You can display the data for the last 24 hours, last week, or last month. You also have an option to select the date range.
5. Under **Report Format**, select the format of the report that you want to download later.
   - If you select either **PDF** or **CSV** check box, you can download the report in either of the two formats that you select.
   - If you select both **PDF** and **CSV** check box, you can select the recipient format by clicking **PDF** or **CSV** check box or both.

     > Note:
     >
     > - With recipient format, the email recipient can download the report in the format that you have selected.
     > - If you do not select any format of the report, by default you can download the report in only the PDF format. Also, by default, the email recipient receives the report in the PDF format.
6. Under **Email Report**, enter recipients (comma-separated) in the **Additional Recipients** text box.

   > Note: Email is sent to the recipients specified in this field along with the recipients that were configured during the report configuration. For more information about configuring Email settings, see Configuring Report Settings on page 645.
7. Click **Run**.
   A report is generated.

Editing a Report

You can edit an existing report. You can modify views and change other report settings in the Edit mode.

About this task

To edit a report, do the following.

Procedure

1. In the **Report** dashboard (see Reports View on page 194), click the check box against the report that is already generated.
2. From the **Actions** drop-down menu, select **Edit**.

3. Edit the report by updating the desired fields as described on the Creating a New Report on page 619 topic.
   To update the views that are added to the report, perform the following procedure.

   - 1. Select the check box of the view that you want to edit and click the drop-down menu in the right.
        You can edit the view, move the view into a group (if group already exists), or remove the view.
     2. To edit the view, click **Edit**.
     3. To move the view inside a group view, click the name of the group view. The view is moved into the group view.

        Note: This option only appears if you have a group view that is already configured and if you have adhered to hierarchical nature of the group view. For more information, see Adding a Group View on page 639.

     4. To remove the view from a group, select the view and click **Remove from Group**.

4. After you edit the required fields, do one of the following.

   - Click **Save and Run** to save and generate the report.
   - Click **Save** to save the report.

   The report is edited and saved.

Cloning a Report

Cloning a report creates a new copy of the report and at the same time you can add new views to the report.

About this task

To clone a report, do the following.

Procedure

1. In the **Report** dashboard (see Reports View on page 194), click the check box against the report that is already generated.
2. From the **Actions** drop-down menu, select **Clone**.
   A **Report Preview** window appears.
3. Update the desired fields.
   You can also add new views to the report. See the Creating a New Report on page 619 to add views and create a new instance of the report.
   The report is cloned.

Deleting a Report

Deleting a report removes all the generated report instances along with the selected report configuration.

About this task

To delete a report, do the following.

Procedure

1. In the **Report** dashboard (see Reports View on page 194), click the check box against the report that is already generated.
2. From the **Actions** drop-down menu, select **Delete**.
   The report is deleted.

Viewing Report Instances

Perform the following procedure to view the generated reports.

Before you begin

Ensure that you have generated at least one instance of the report that you are viewing. Otherwise, no instance found message is displayed.

Procedure

1. In the **Report** dashboard (see Reports View on page 194), click the report you want to view.
   All the instances of the report with timestamp information is displayed.

   Note: An instance of the report is generated only if you have manually generated the report or the report has been generated according to the configured schedule.

2. Select the instance of the report on which you want to perform the operations.

   Note: You cannot view multiple reports at the same time.

   You can perform the following operations on the report under the **Actions** drop-down menu.

Figure 394: Viewing Instance Information

Note: If the report generation fails then the Status column displays the status as Failed and if you hover on the field, you can view the reason for the failure.

a. **Resend Report**:

   - Select this option to send the report again to all the Email recipients that you have configured.
   - Select the format of the report (**PDF** or **CSV** or both) before you send an email, under **Report Format**.

   You also have an option to send the report to additional recipients by clicking **+ Add Additional Recipients** link. If you do not add any additional recipients, the report is sent to the recipients that are already configured in the **Report Settings** flow.

b. **Delete**: Select this option to delete the selected report instance.

Example Report - Environment Summary

A generated environment summary report has the following information.

Sample Environment Summary Report

- Stylistic information: The first page of the report is used to represent this kind of information. The stylistic information includes the name of the report, time it was generated, any logo that you have configured, header and footer information (for example, report name and page number).

Figure 395: Stylistic Information

- Metrics information: Metrics summary displays the cluster details information (for example, cluster usage, performance details, licensing information, etc).

  Following sample images displays the detailed information about the cluster efficiency summary information.

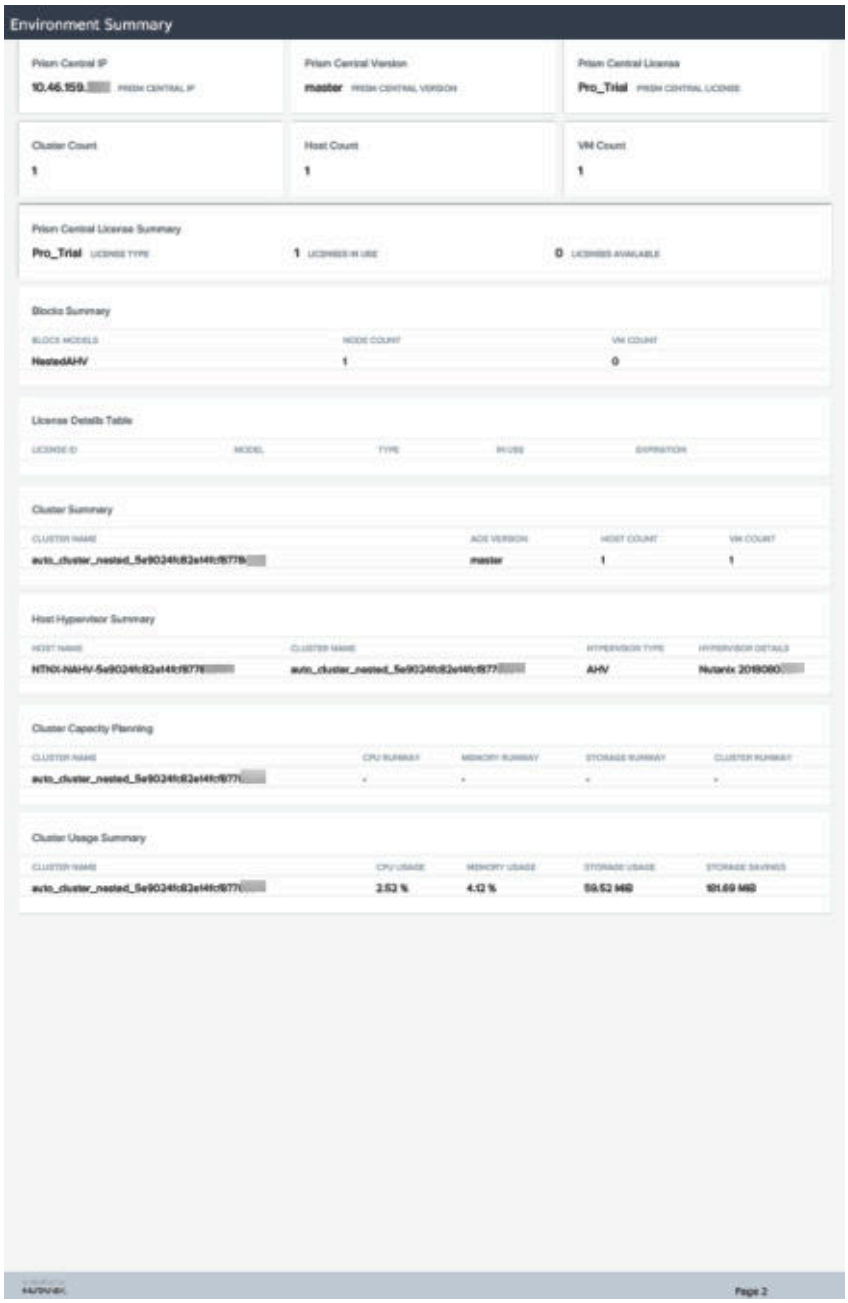

Figure 396: Detailed Metrics Information (Image 1)

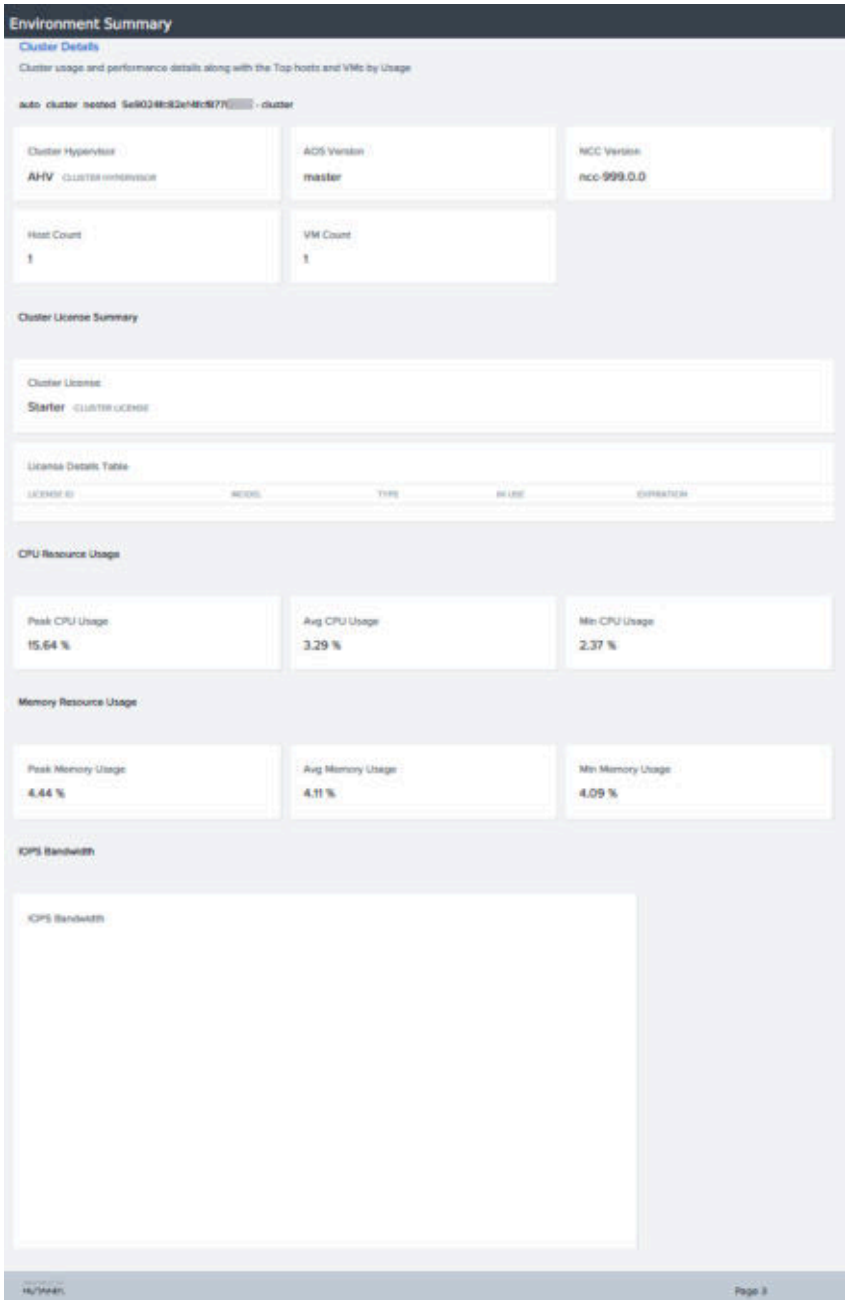

Figure 397: Detailed Metrics Information (Image 2)

Views in the Reports

The report management feature provides you different views that you can add to the report. You can also customize these views during the addition process itself. These views are displayed under **Custom Views** pane. You cannot customize some views, but these views can be directly included in the reports. These views are displayed under **Pre-Defined Views** pane.

Adding Views to a Report

Perform the following procedure to add views to the report.

Addition of Views

You can add following views to a report.

- To add a bar chart, line chart, or histogram view, see Adding a Bar Chart or a Line Chart or a Histogram View to a Report on page 627.
- To add a data table view, see Adding a Data Table View to a Report on page 630.
- To add a configuration summary view, see Adding Configuration Summary View on page 633.
- To add a metric summary view, see Adding Metric Summary View on page 635.
- To add an entity count view, see Adding Entity Count View on page 637.
- To add a title and description view, see Adding Title and Description View on page 639.
- To add a group view, see Adding a Group View on page 639.
- To add a pre-defined view, see Adding a Pre-Defined View on page 642.

Note:

- After you add the views to a report, the report displays the dummy data for the view that you have added. Only when you generate the report, the data for the view is updated and you can download the report to view its information.
- Collecting some VM data requires that NGT is installed. As a result reports include such data only from NGT-enabled VMs.

Adding a Bar Chart or a Line Chart or a Histogram View to a Report

You can use the bar chart, line chart, or a histogram view to represent the historical data in a graphical format. Perform the following procedure to add these views to the report.

Procedure

1. In the **New Report** wizard, select the type of view (bar chart, line chart, or a histogram view) that you want to add from the **Custom View** pane.

2. Configure the bar chart, line chart, or a histogram view.
   Following image displays the bar chart view.

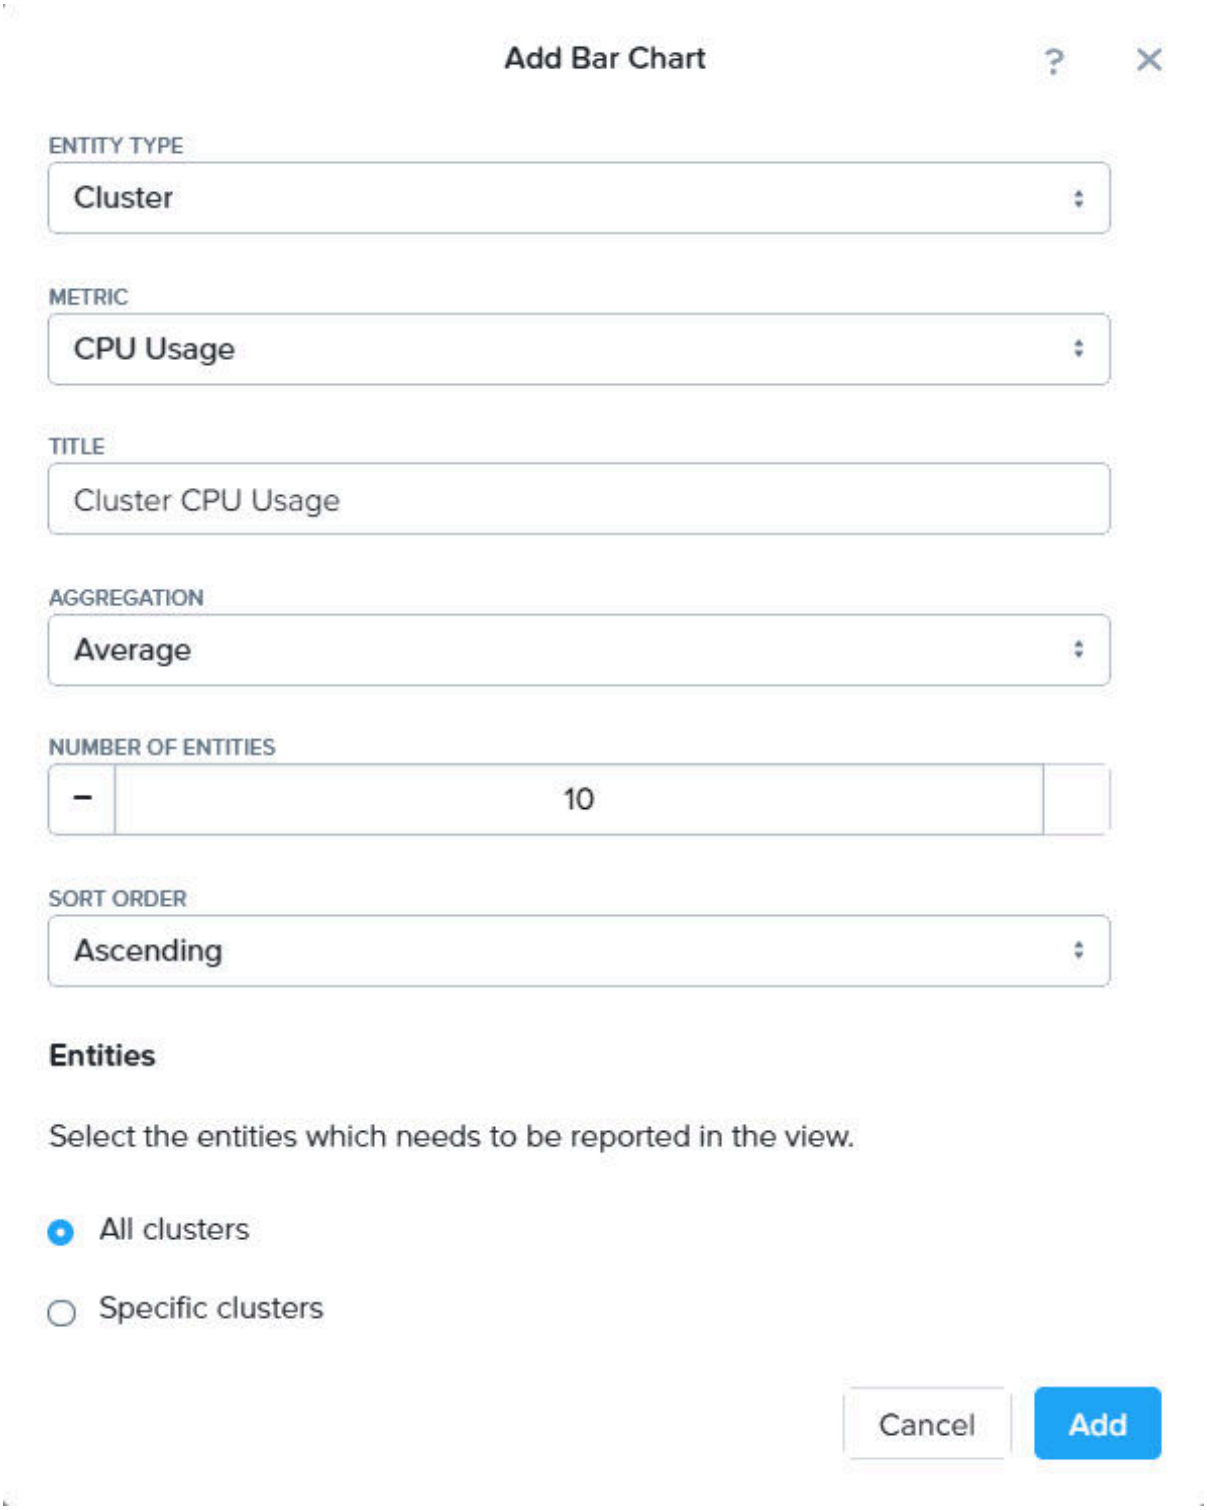

Figure 398: Bar Chart View

a. Select the entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) from the **Entity Type** drop-down menu for which you want to configure the view.

b. Select a specific metric of the entity type.

   The metric selection depends on what you have selected in the **Entity Type** drop-down menu. For example, if you select **Cluster** as an entity type, you can select one of the following metric.

   CPU Usage, **Memory Usage**, **Controller IOPS**, **IO Bandwidth**, **IO Latency**, **Free (Physical)**, **Storage Usage**, **Storage Capacity**, **Controller Read I/O Bandwidth**, **Controller Write I/O Bandwidth**, **Controller Read I/O Latency**, **Controller Write I/O Latency**, **Storage Logical Usage**, **Controller Read IOPS**, **Controller Write IOPS**.

c. Enter the title of the bar chart, line chart, or histogram view in the **Title** field.

d. Select the aggregation from the **Aggregation** menu.

   For the metrics that has time-series data, you can define the aggregation field. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric and select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the report.

   You can define the aggregation value of **Sum**, **Max**, **Min**, **Average**, **Last**, or **Count**.

   > Note: Aggregation is not supported for Line Chart view.

e. Enter the number of entities that you want to include in the report from the **Number of Entities** text box.

   > Note: This field does not appear for the Histogram view.

f. Select the **Limit entities in the chart** option to limit the number of entities that you want to include in the view and how you represent the entities (**Ascending** or **Descending**).

g. (Only for Histogram view) Enter number of buckets in the **Number of Objects** field.

   You can define the frequency distribution for all the entities in the entity type for the selected metric by using this field. For example, for a selected metric if the minimum value is 10 and maximum value is 100 and you have entered number of buckets as 5 then you will have 18 buckets configured (100-10/5). Depending on the value of the metric, different metrics will fall in the appropriate buckets.

h. Select the entities that you want to add.

   - **All Entities**: Select this option if you want to include all the entities that you have selected in the view
   - **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

3. Click **Add**.
   The configured view gets added to the report.

Adding a Data Table View to a Report

The data table view provides you with the point in time data of a particular entity. To include a data table view to the report, perform the following procedure.

Procedure

1. In the **New Report** wizard, select the **Data Table** view from the **Custom View** pane.

2. Configure the Data Table view.

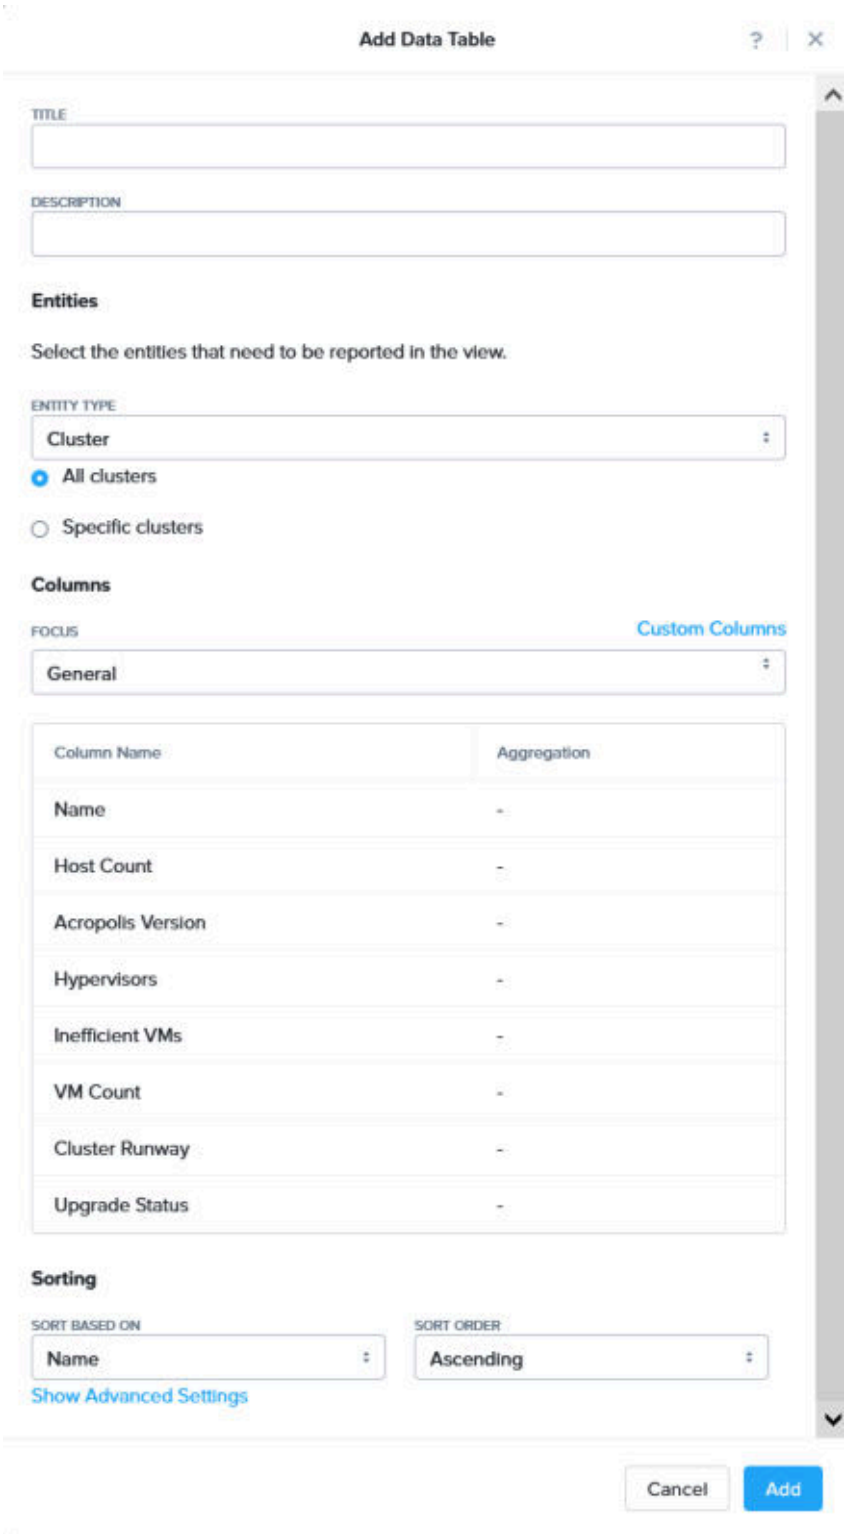

Figure 399: Data Table View

a. Enter the title of the data table view in the **Title** field.

b. Enter the description of the view in the **Description** field.

c. Select the entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) for which you want to configure the view.

d. Select the entities that you want add.

   » **All Entities**: Select this option if you want to include all the entities that you have selected in the view.

   » **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

e. For an entity that you have selected in the **Entity Type** drop-down menu, you can select a list of metrics and each of these metrics are displayed as columns in the table.

   The metric columns are further divided in to **General** or **Performance** depending on the entity that you have selected. Depending on your selection, pre-defined columns are displayed that will be included as part of the view. However, you can create your own custom column.

f. To create your own custom column, click the **Custom Columns** link.

   Select the columns that you want to add by click the **+** icon. You can select maximum of 10 metrics. Depending on your selection the focus area is updated.

   > Note: If any metric has time-series data, you can define the aggregation parameter also. If you do not define the aggregation, by default Average is configured for the metric.

g. Define the sorting by selecting the entity and order on which sorting should be performed by selecting the entity from the **Sort Based On** drop-down menu and **Sort Order** drop-down menu.

h. To configure advanced setting, click the **Show Advanced Settings** link.

i. Select the **Limit number of rows in the table** check box to limit the number of rows that you want to include in the view and enter the value in the text box that is provided.

j. Select the **Group Data** check box to select the entities based on which the data in the view should be grouped and the views should be repeated.

   » **All Entities**: Select this option if you want to perform the grouping for all the entities.

   » **Specific Entities**: Select this option if you want to perform the grouping for selected entities depending on your selection from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

   For example, if you want to perform grouping for all the VMs at the cluster level and you have three clusters registered, you will get three tables and each table will have VMs for that particular cluster. You can define multiple such views. The entire set of views as a group will be repeated for each of the registered cluster.

3. Click **Add**.
   The configured view gets added to the report.

Adding Configuration Summary View

You can add this view to include the static cluster configuration information, for example, cluster IP address, hardware model, license information, etc.

Procedure

1. In the **New Report** wizard, select the **Configuration Summary** view from the **Custom View** pane.

2. Configure the Configuration Summary view.

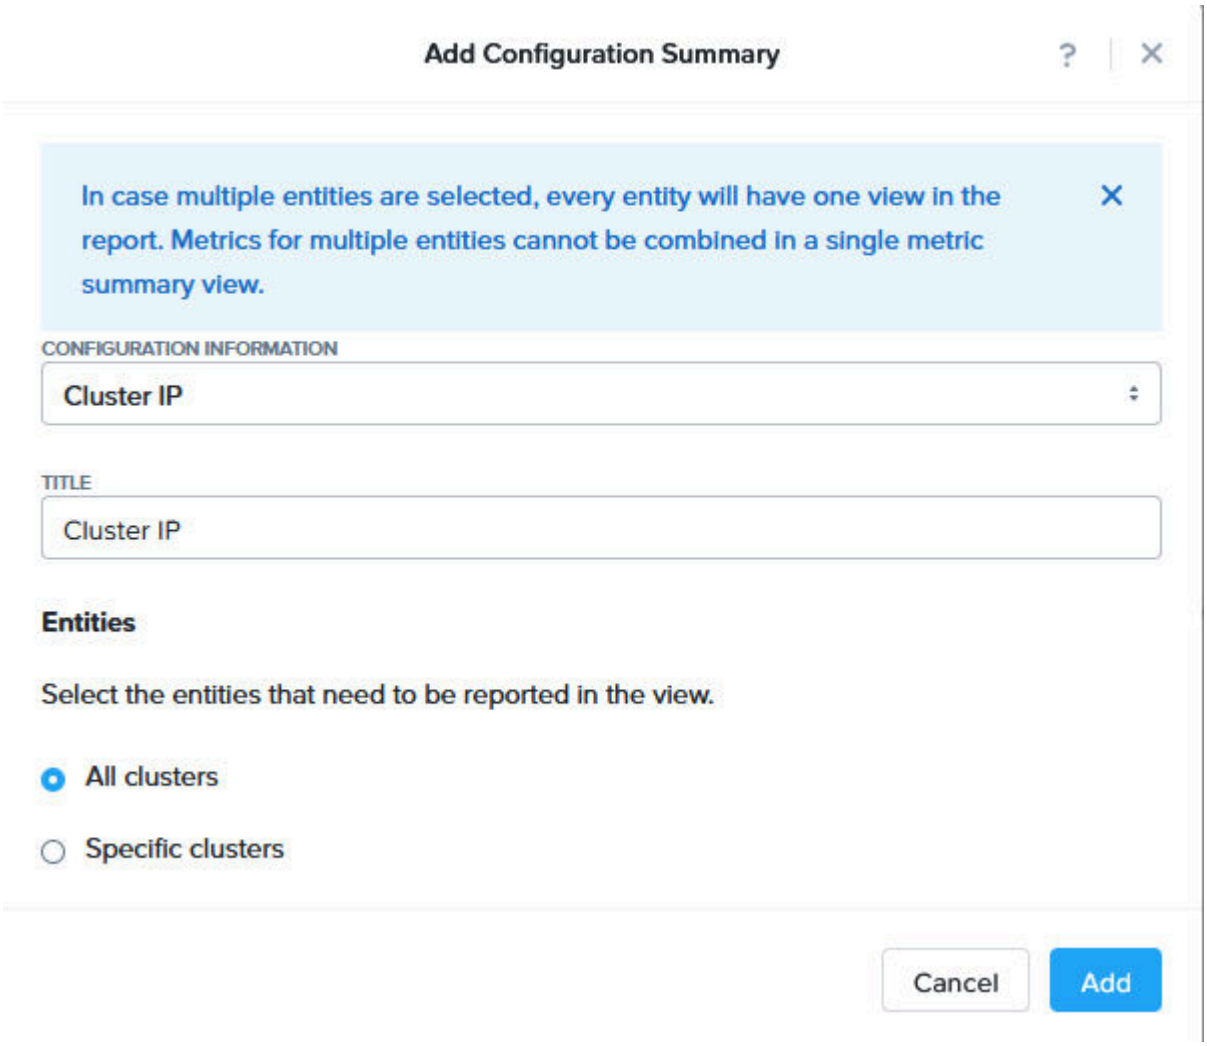

Figure 400: Configuration Summary View

a. Select the configuration information that you want to include from the **Configuration Information** drop-down menu.

   You can include the configuration information for the **Cluster IP**, **Prism Central IP**, **PC Version**, **Cluster Version**, **Cluster Hypervisor**, **Number of Blocks**, **Hardware Model**, **Cluster License**, or **Prism Central License**.

   Specific information is included depending on the cluster information that you have added. For example, if you have selected **Cluster IP**, the IP addresses of all the clusters that are registered is included in the view.

b. Enter the title of the report in the **Title** field.

c. Select the entities that you want to to add.

   » **All Clusters**: Select this option if you want to include this view for all the clusters that are registered.

» **Specific Clusters**: Select this option to define rules on a particular cluster. For example, if you want to include this view for AHV, you can select **Hypervisors** option and define the rule **Equal to** > **AHV** . You can specify multiple such rules by clicking the **+** icon.

3. Click **Add**.
   The configured view gets added to the report.

Adding Metric Summary View

You can add the metric summary view to get information on a metric for a defined aggregation value.

Procedure

1. In the **New Report** wizard, select the **Metric Summary** view from the **Custom View** pane.

2. Configure the Metric Summary view.

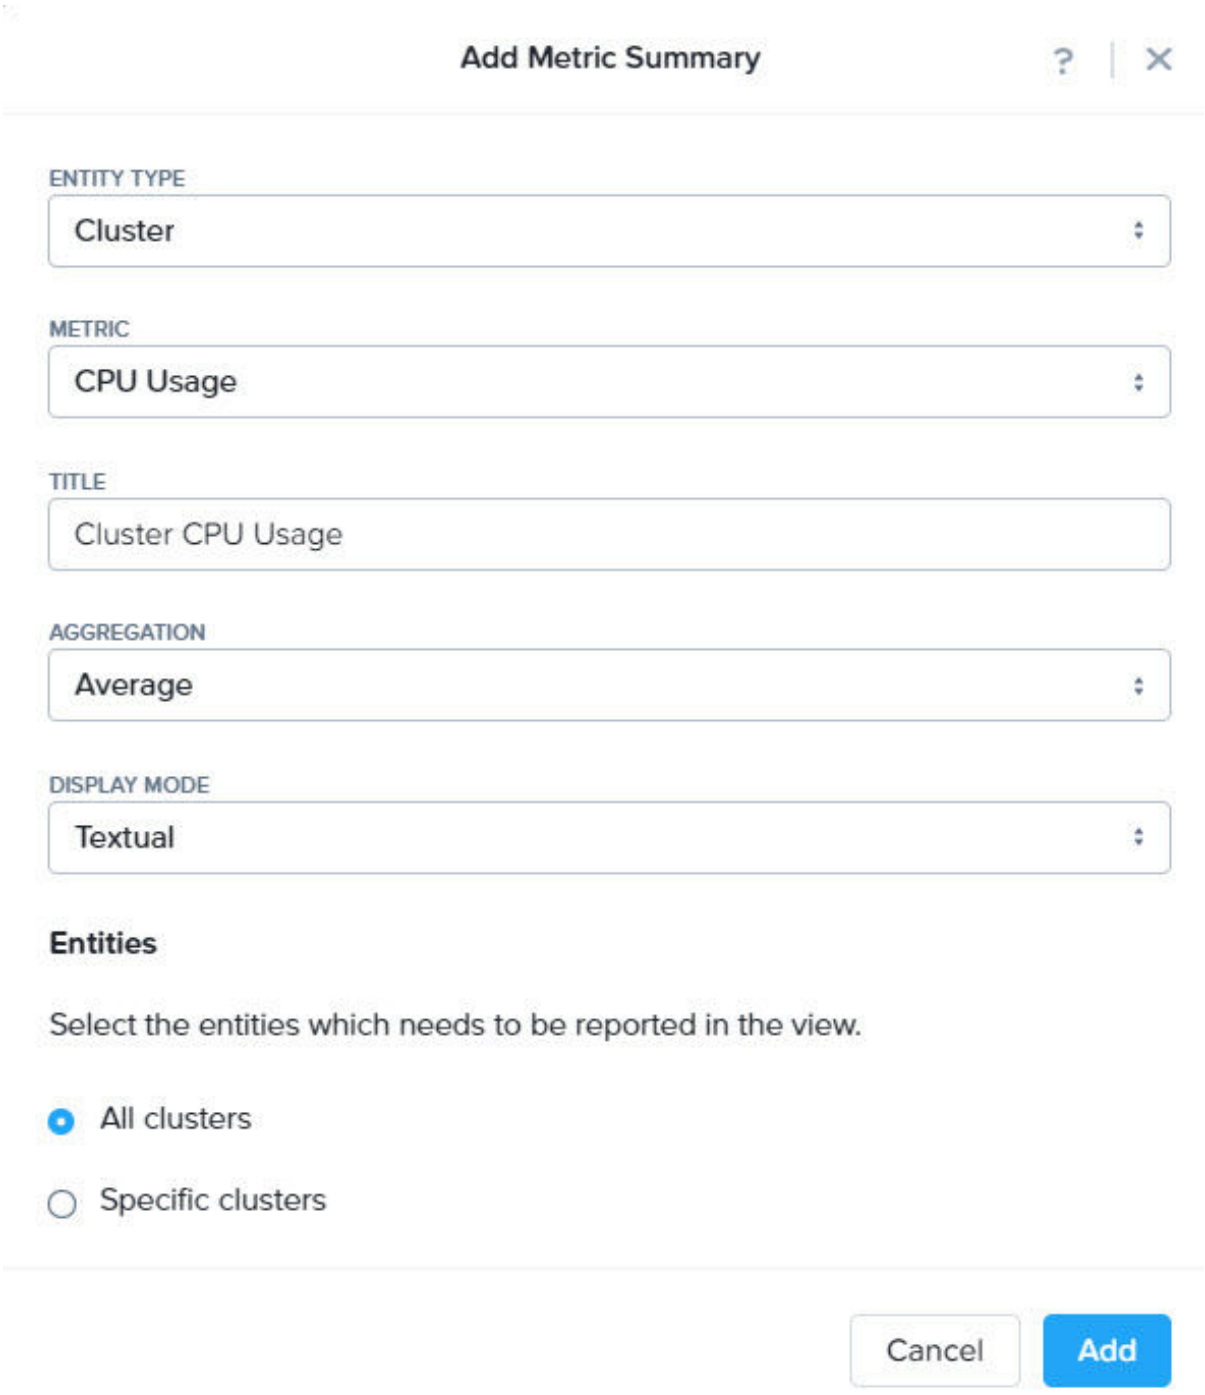

Figure 401: Metric Summary View

a. Select the entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) for which you want to define the metric summary.

b. For an entity that you have selected in the **Entity Type** drop-down menu, you can select a metric for which you want to add the information in the metric summary view.

c. Enter the title of the view in the **Title** field.

d. Select the aggregation from the aggregation **drop-down** menu.

   For the metrics that has time-series data, you can define the aggregation. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the view.

   You can define the aggregation value of **Sum**, **Max**, **Min**, **Average**, **Last**, or **Count**.

e. Select the mode of display from the **Display Mode** drop-down menu.

   You can select either **Textual** or **Graphical (Gauge Chart)**. The graphical chart is only displayed for consumption related metrics, for example, CPU usage.

f. Select the entities that you want to add.

   » **All clusters**: Select this option if you want to include all the entities that you have selected in the view.

   » **Specific clusters**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **Specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

3. Click **Add**.
   The configured view gets added to the report.

Adding Entity Count View

You can add the entity count view to get the number of entities for all the clusters or a particular cluster.

Procedure

1. In the **New Report** wizard, select the **Entity Count** view from the **Custom View** pane.

2. Configure the Entity Count view.

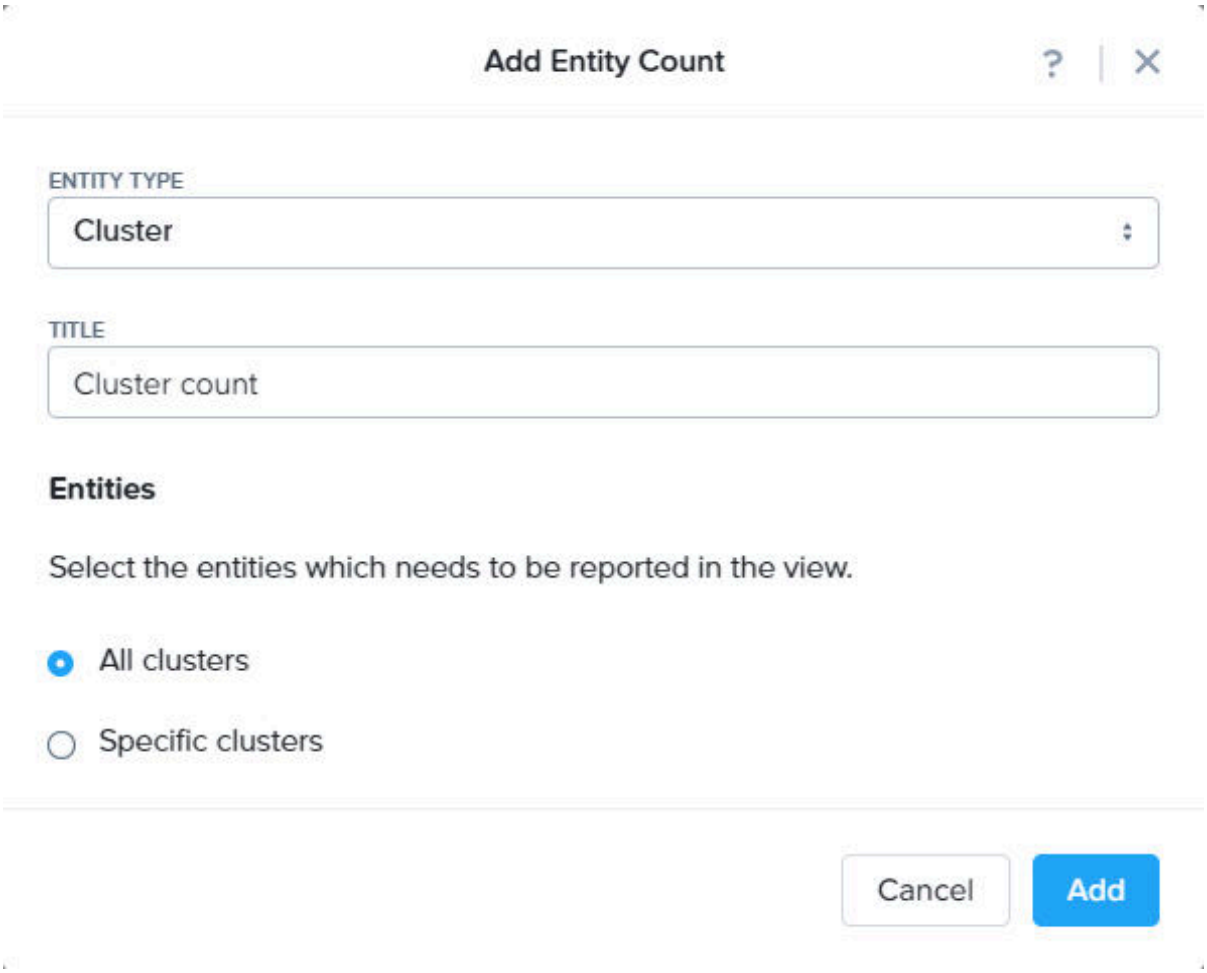

Figure 402: Entity Count View

   a. Select the entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) for which you want to get the count information.
   b. Enter the title of the view in the **Title** field.
   c. Select the entities that you want add.
      - **All clusters**: Select this option if you want to include all the entities that you have selected in the view.
      - **Specific clusters**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **Specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

3. Click **Add**.
   The configured view gets added to the report.

Adding Title and Description View

You can add the title and description view to add any miscellaneous information, for example, disclaimer etc.

Procedure

1. In the **New Report** wizard, select the **Title and Description** view from the **Custom View** pane.
2. Configure the Title and Description view.
   a. Enter the title of the view in the **Title** field.
   b. Enter the description in the **Description** field.
3. Click **Add**.
   The configured view gets added to the report.

Adding a Group View

Perform the following procedure to add a group view.

About this task

Grouping is always performed in an hierarchical manner. For example, if you have created a group for the cluster entity, you can add all the views to this group. However, if you have created the group for the VM entity, you cannot add a view to this group that you have created at the cluster level.

> Note: This view takes the full page width.

Procedure

1. Select the **Group** view from the **Custom View** pane in the new report wizard.

2. Configure the Group view.

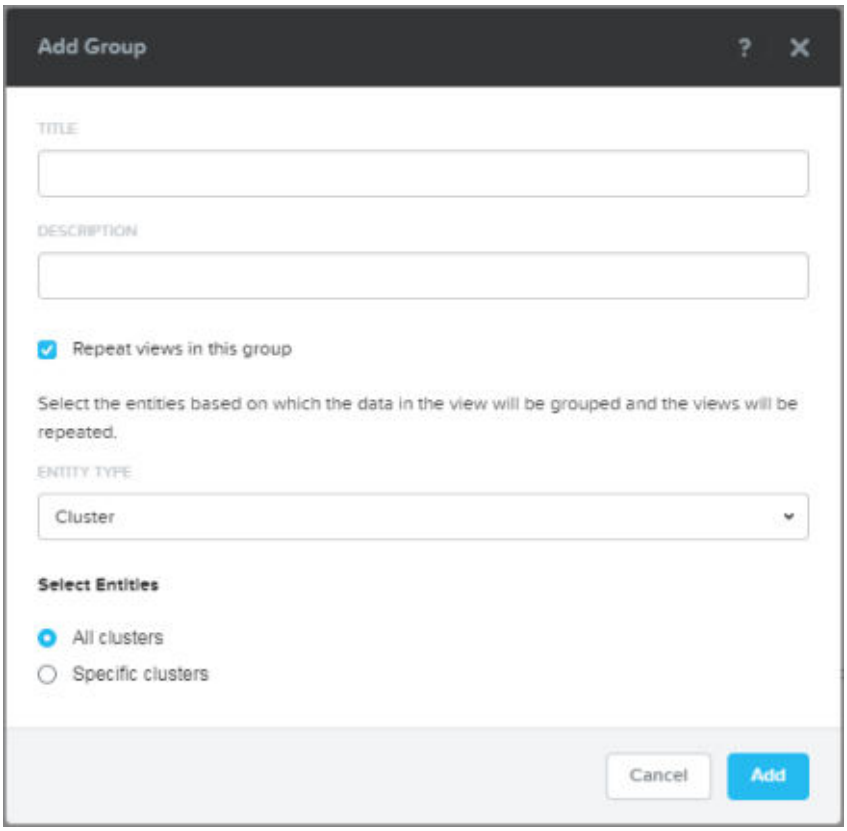

Figure 403: Group View

a. Enter the title of the data table view in the **Title** field..

b. Enter the description of the view in the **Description** field.

c. Select the **Repeat views in this group** check box to select the entities based on which the data in the view should be grouped and the views should be repeated.

d. Select the entity type from the **Entity Type** drop-down menu.

   » **All Entities**: Select this option if you want to perform the grouping for all the entities.

   » **Specific Entities**: Select this option if you want to perform the grouping for selected entities depending on your selection from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

3. Click **Add**.
The configured view gets added to the report. Depending on the entity type that you have configured, you can create grouping of the views. For example, if you have created the group view for the entity type VM, you can add all the views to this except the view with entity type cluster.

What to do next

You can add only the group view by following this procedure. If you want to add any views to this group, see Adding Views into a Group View on page 641

Adding Views into a Group View

Perform the following procedure to add views to a already created group view.

About this task

You can add a view to a group in multiple ways. The addition of views depends on whether you have created the group view or not.

Procedure

1. If you have created the group view, perform the following procedure. For more information on creating the group view, see Adding a Group View on page 639

   a. Select the check box of the view that you want to add to the group view, and click the drop-down menu in the extreme right of the view.

      The group to which you can add this view is displayed in the drop-down menu. You can perform the group only in an hierarchical manner. For example, if you have created group for the cluster entity, you can add all the views to this group. However, if you have created the group for the VM entity, you cannot add a view to this group that you have created at the cluster level.

   b. Select the group that you want to move this view.

      For example, in the following figure, the bar chart view with title Cluster CPU Usage can be moved to the group named **Cluster Group**.

Figure 404: Adding Views to Group view

      The view is moved in the group view.

2. If you have not created the group view, perform the following procedure.

   a. Select the check box of the view that you want to add to the group view, and click **Create Group** button at the top.

   b. Create the group by following the procedure Adding a Group View on page 639.
      The view is added to the newly created group.

   Depending on the grouping, the data in the view is grouped and the views are repeated. For example, if you want to perform grouping for all the VMs at the cluster level and you have three cluster registered, you will get three tables and each table will have VMs for that particular cluster. You can define multiple such views. The entire set of views as a group is repeated for the each cluster.

Adding a Pre-Defined View

In addition to adding and configuring views according to your requirements, you an also add some pre-defined views. You cannot modify these views.

About this task

As part of the pre-defined views, you can add the following views.

> Note: You can only include the **License Details Table** view as part of the group view.

Procedure

Select the view that you want to add from the **Pre-Defined Views** pane.

You can add the both detailed and summary licensing information and block summary information.

- **License Details Table:** Provides detailed information about the Prism Central licenses and includes information about license ID, model that are using the licenses, type of licenses, displays whether the licenses are used or not, and displays when the licenses are going to expire.
- **Prism Central License Summary:** Provides summary information of the Prism Central licenses and includes information about the type of licenses, number of used licenses, and number of licenses that are available.
- **Blocks Summary:** Provides summarized information about the different block models, number of nodes, and number of VMs running in the Prism Central.
- **Alerts Histogram:** Provides distribution of alerts over a specific time interval as configured during the execution of the report. X-axis plots the time period of the report as configured in the **Time Period of Report** field. Y-axis plots the number of alerts generated in the specific time period.
- **CPU Runway:** Provides visualized information about the historical CPU usage, predicted CPU usage, and predicted runway of the cluster.
- **Memory Runway:** Provides visualized information about the historical memory usage, predicted memory usage, and predicted runway of the cluster.
- **Storage Runway:** Provides visualized information about the historical storage usage, predicted storage usage, and predicted runway of the cluster.
- **List of Inactive VMs:** Provides detailed information about the inactive VMs on the cluster; including VM name, efficiency status, reasons for being marked as inactive and potential resource gains that can be accomplished by deleting inactive VMs or resizing over-provisioned VMs.
- **List of vCenter Inactive VMs:** Provides detailed information about the inactive VMs on the Non-Nutanix cluster; including VM name, efficiency status, reasons for being marked as inactive and potential resource gains that can be accomplished by deleting inactive VMs or resizing over-provisioned VMs.
- **List of Constrained VMs:** Provides detailed information about the constrained VMs on the cluster (that require more resources to function properly); including the VM name, efficiency status and the reasons for being marked as constrained.
- **List of vCenter Constrained VMs:** Provides detailed information about the constrained VMs on the Non-Nutanix cluster (that require more resources to function properly); including the VM name, efficiency status and the reasons for being marked as constrained.

- **List of Bully VMs:** Provides detailed information about the VMs whose activities are affecting other VMs that is not directly related to their sizes; including VM name, efficiency status and the reasons for being marked as bully.
- **List of vCenter Bully VMs:** Provides detailed information about the vCenter registered VMs whose activities are affecting other VMs that is not directly related to their sizes; including VM name, efficiency status and the reasons for being marked as bully.
- **List of Overprovisioned VMs:** Provides detailed information about the over-provisioned VMs on the cluster (that have more resources than they deserve); including the VM name, efficiency status, reasons for being marked as over-provisioned and potential resource gains that can be accomplished by deleting the inactive VMs or resizing the over-provisioned VMs.
- **List of vCenter Overprovisioned VMs:** Provides detailed information about the over-provisioned vCenter registered VMs on the cluster (that have more resources than they deserve); including the VM name, efficiency status, reasons for being marked as over-provisioned and potential resource gains that can be accomplished by deleting the inactive VMs or resizing the over-provisioned VMs.
- **Potential CPU Reclaim:** Provides summarized information about the potential CPU resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Memory Reclaim:** Provides summarized information about the potential memory resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Storage Reclaim:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.

Figure 405: Pre-Defined Views

Scheduling a Report

After you add all the views to the reports, you have an option to run the report at a custom generated time. This scheduling is applicable only for the specific report for which you are configuring the schedule.

Procedure

1. In the **New Report** wizard, click **Add Schedule**.

2. Select when you want to run the report from the **Report Runs** drop-down menu.
   You an run the report on daily, weekly, monthly, or yearly basis.

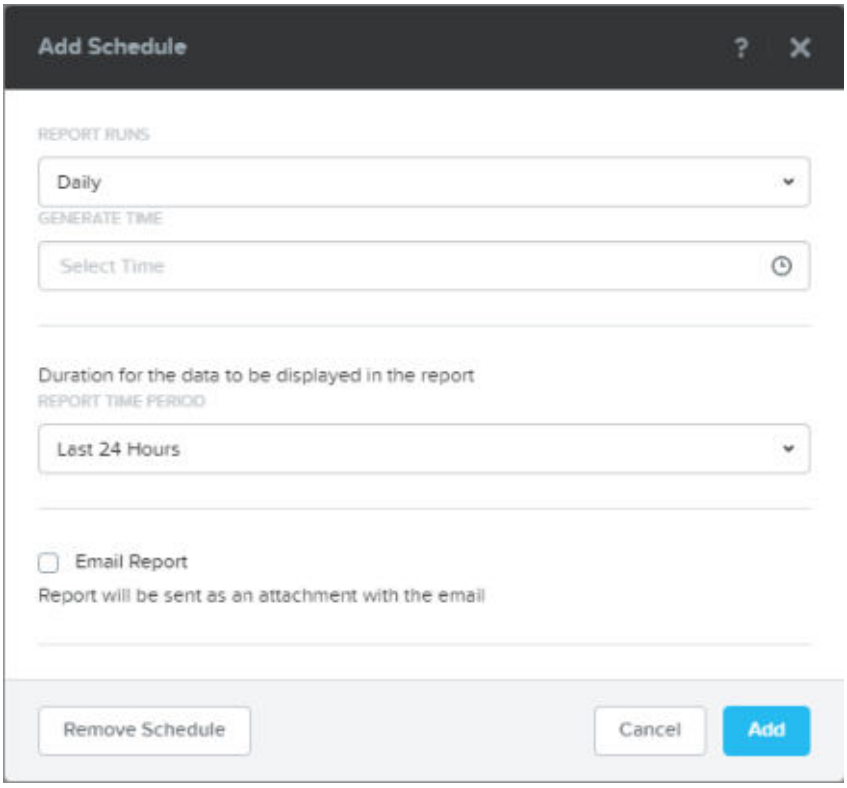

Figure 406: Report Scheduling

3. Select the time when you want to repeat the report.
   For example, if you select **Yearly** from the **Report Runs** drop-down menu, you can repeat the report by selecting the month when you want to run the report from the **Every** drop-down menu, date of the month from the **On** drop-down menu, time of the day when you want to run the report from the **Generated Time** drop-down menu.

4. Select the duration for the data that you want to display in the report from the **Report Time Period** drop-down menu.
   You can display the data for last 24 hours, last week, or last month. You also have an option to select a custom data range.

   a. If you select a **Custom Date Range** from the **Report Time Period**, you need to select the date range from the **From** and **To** option.

   If you a select a date range from 10 a.m. (Monday) to 10 a.m. (Tuesday), the data for this time is displayed in the report. For example, in a line chart, the X-axis will display the time range from 10 a.m. (Monday) to 10 a.m. (Tuesday) and the Y-axis will display the variation in data.

5. Select the **Email Report** check box and enter recipients (comma-separated) in the **Email Recipients** text box.
   Only the specific recipients that you have entered in the text box receives the report.

6. Click **Add**.
   The schedule for the report gets added.

7. After report is saved, you can also edit or delete the schedule by clicking **Edit Schedule**.
   a. To edit the schedule, make the desired changes to the schedule and click **Add**.
   b. To delete the schedule, click **Remove Schedule**.

Downloading Reports

You can download a report in the PDF format or CSV format or both.

About this task

- A PDF generated report displays a graphical representation of the infrastructure resources.
- A CSV generated report contains data of views separated by commas. You can use this CSV file to load CSV data into other systems. Data of views that you add to a report is populated to the CSV format report.
- The CSV report that you download is in the form of a .zip file. The .zip file contains CSV generated files for each view.
- For each group, a folder is created containing CSV files of each view of that group.
- If a view is repeated over a cluster entity or VM entity, the view has a CSV file for every repeating cluster. For example, a view named *View* that is repeating over clusters *Cluster_1*, *Cluster_2*, *Cluster_3*, the name of the CSV file is displayed as follows:
  - *1_1_View_Cluster_1.csv*
  - *1_2_View_Cluster_2.csv*
  - *1_3_View_Cluster_3.csv*
- If a group has a view repeating over a cluster or VM entity, the folder of the group consists a folder for every cluster and that cluster contains a CSV file for the view. For example, a group named Group having view named *View* on clusters *Cluster_1*, *Cluster_2*, *Cluster_3*, the CSV folder structure is displayed as follows:
  - *Group* / *Cluster_1* / *View.csv*
  - *Group* / *Cluster_2* / *View.csv*
  - *Group* / *Cluster_3* / *View.csv*

To download a report, do the following.

Procedure

1. On the **Reports** dashboard (see Reports View on page 194), click a report you want to download.
2. From the list of reports, click **PDF** or **CSV** under the **Download** column next to the report you want to download.

   Note: You can set the output format of the report, see Creating a New Report on page 619.

Configuring Report Settings

You can configure the report settings by configuring its appearance, email settings, and retention policy. Perform the following procedure to configure the report settings.

About this task

You can configure the report settings for individual reports or for all the reports that you have generated. Depending on where you configure the settings, the reports settings are applied.

> Note: If you apply the report settings both the global level (for all the reports) or for report level (when you are creating a new report), the setting applied at the report level takes precedence.

Procedure

1. Go to the report settings page.
   - » To configure the report settings at the global level, click **Report Settings** on the **Reports** dashboard (see Reports View on page 194).
   - » To configure the report settings at the report level, in the **New Report** wizard, click **Report Settings**.

2. Configure the appearance of the report.

   The appearance of the report is divided into two parts; cover page settings and content page settings. Configuration of logo and the background color is included in the cover page settings. Configuration of color of the header is included in the content page settings.

   > Note: The copyright information will not be displayed in the footer section, therefore the ability to customize copyright information is not allowed.

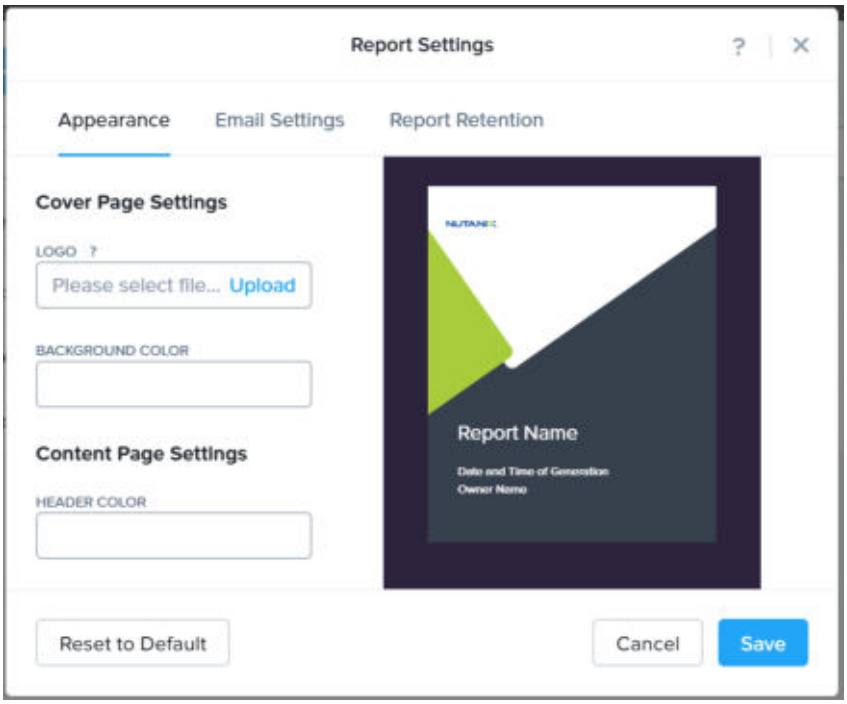

Figure 407: Report Settings Configuration

   a. To upload the logo for the report, click the **Upload** button.

      > Note: The logo size cannot be more than 1 MB.

   b. To change the background color of the report, select the background color from the **Background Color** field.

   c. To change the header color of the report, select the header color from the **Header Color** field.

3. Configure the Email settings.

   a. Enter the subject of the Email in the **Prepend Content to the Email Subject** text box.

      You can view the preview of the subject that you are adding in the Preview pane and is prepended to the subject *subject_email* Prism Report *report_name*.

   b. Enter the body of the Email in the **Append Content to the Email Body** text box.

      You can view the preview of the body that you are adding in the Preview pane and is appended to the body Please find attached the generated report *email_body*.

   c. Enter the recipients (comma-separated) of the report in the **Email Recipients** text box.

      The recipients that you have configured will get the report with the subject and body of the Email that you have entered.

4. Configure the retention policy of the report.

   a. Click the **Define Report Retention** check box.

      You can define the retention policy by entering the number of instances of the report that you want to retain or the duration of time that you want to retain the report.

   b. To configure the retention policy in instances form, click the **Number of Instances** option and enter the desired number of instances.

      > Note: You cannot retain more than 25 instances of a report. If you have 25 instances of a report and you generate a new instance of the same report, the first instance that was generated gets deleted.

   c. To configure the retention policy in duration form, click the **Time Duration** option and enter the desired number and duration in terms of **Days**, **Weeks**, or **Months** from the drop-down menu.

      For the time-based retention, you can retain a report for maximum of 3 months.

      > Note: If you do not define any retention policy, by default 10 instances of a report are retained.

5. Configure the Report Format of the report.

   a. Click the **PDF** check box or **CSV** check box or both to set a report format.

      A PDF generated report displays a graphical representation of infrastructure resources. A CSV generated report contains information separated by commas. You can use this CSV file to load CSV data into other systems.

   b. If you select the PDF format and CSV format in **Report Format**, you can select the **Recipient Format**. Select the format of report in which you want to send an e-mail.

      > Note: You cannot set the report format at the global level. You can only set the report format at the report level while creating a new report or editing an existing report.

6. Click **Save**.
   The report configuration is saved and next time you generate the report, the report generation follows the configured setting. If you apply the report settings both the global level (for all the reports) or for report level (when you are creating a new report), the setting applied at the report level takes precedence.

7. (Optional) Click **Reset to Default** to discard all the changes that you have made in the Report Settings page.

Role Based Access Control

Role-Based Access Control (RBAC) restricts Reports Management access to authorized users. For any role you are assigned, privileges are given to you to modify certain configurations in Reports Management. The following table shows the roles and certain privileges given to that role.

Table 247: Roles and Permissions

Access/ Permissior	Consumer	Developer	Operator	Prism Admin	Prism Viewer	Project Admin	Self-Service Admin	Super Admin
Create Common Report Config				#		#	#	#
Create Report Config				#		#	#	#
Create Report Instance				#		#	#	#
Delete Common Report Config				#		#	#	#
Delete Report Config				#		#	#	#
Delete Report Instance				#		#	#	#
Notify Report Instance	#	#	#	#	#	#	#	#
Update Common Report Config				#		#	#	#
Update Report Config				#		#	#	#
View Common Report Config	#	#	#	#	#	#	#	#
View Report Config	#	#	#	#		#	#	#

Access/ Permissior	Consumer	Developer	Operator	Prism Admin	Prism Viewer	Project Admin	Self-Service Admin	Super Admin
View Report Instance	#	#	#	#	#	#	#	#

13

SERVICES ENABLEMENT

You can enable selected services through Prism Central.

- Enabling Calm on page 651
- Enabling Files on page 652
- Enabling Foundation Central on page 652
- Enabling Karbon on page 652
- Enabling Objects on page 653

Enabling Calm

About this task

You can select, provision, and manage your business applications across all your infrastructure for both private and public clouds through the Nutanix Calm feature. Nutanix Calm provides automated application life cycle management, custom blueprints for the setup and management of enterprise applications, a marketplace to publish the blueprints to end users, and automated hybrid cloud management to provision your hybrid cloud architecture. See the *Nutanix Calm Administration and Operations Guide* for information about configuring and using Calm.

To enable Calm, do the following:

Procedure

1. Do one of the following:
   - » Click the collapse menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Calm** from the entities menu (see Entities Menu on page 13).
   - » Select **Enable App Management** from the Settings menu (see Settings Menu on page 17).
2. In the **Enable App Management** page, do the following:
   a. Check the **Enable App Management** box to enable application management.
   b. Check the **Enable Nutanix seeded blueprints** box to load your application store with a set of pre-defined blueprints.

      While loading the pre-defined blueprints is optional, it is recommended as it provides blueprints for many common use cases.
   c. Click the **Save** button.

Enabling Files

About this task

Nutanix Files provides file services to clients for file sharing across user work stations from a centralized and protected location.

To enable Files, do the following:

Procedure

1. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Files** from the entities menu (see Entities Menu on page 13).
2. In the Nutanix Files page, click the **Enable Files** button.

   After Files is enabled, the Files dashboard appears. See the *Nutanix Files Guide* for information about configuring and using Files.

Enabling Foundation Central

About this task

Foundation Central can manage several Foundation instances from a single pane of glass, allowing you to create clusters of remote nodes without needing to configure each of them individually.

To enable Foundation Central, do the following:

Procedure

1. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Foundation Central** from the entities menu (see Entities Menu on page 13).
2. In the Nutanix Foundation Central page, click the **Enable Foundation Central** button.

   After Foundation Central is enabled, the Foundation Central dashboard appears. See the *Foundation Central Guide* for information about configuring and using Foundation Central.

Enabling Karbon

About this task

Nutanix Karbon is a curated turnkey offering that provides simplified provisioning and operations of Kubernetes clusters. Kubernetes is an open-source container orchestration system for deploying and managing container-based applications.

To enable Karbon, do the following:

Procedure

1. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Karbon** from the entities menu (see Entities Menu on page 13).
2. In the Karbon page, click the **Enable Karbon** button.

   After Karbon is enabled, you can click the **here** link to go to the Karbon console. See the *Nutanix Karbon Guide* for information about configuring and using Karbon.

Enabling Objects

Nutanix Objects is a feature that allows you to create and manage object stores in a cluster.

About this task

To enable Objects, do the following:

Procedure

1. Click the collapse menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Objects** from the entities menu (see Entities Menu on page 13).
2. In the Objects page, click the **Enable** button to enable Objects.

   After Objects is enabled, the Object Stores dashboard appears. See the *Nutanix Objects Guide* for information about configuring and using object stores.

14

CUSTOMER SUPPORT SERVICES

Nutanix provides customer support services in several ways.

- Nutanix customer support can monitor your clusters and provide assistance when problems occur through the Pulse mechanism (see Pulse Health Monitoring on page 654 and Configuring Pulse on page 656).
- Nutanix customer support maintains a portal that you can access to request assistance, download various product updates, and view documentation (see Accessing the Nutanix Support Portal on page 666).
- If you need help, you can create a support ticket directly from Prism Central (see Creating a Support Case on page 659).

Pulse Health Monitoring

The feature known as Pulse provides diagnostic system data about Prism Central to Nutanix Support to deliver proactive, context-aware support for Nutanix solutions. Pulse unobtrusively collects this information with no effect on system performance. Pulse shares only basic system-level information necessary for monitoring the health and status of Prism Central. Information includes system alerts, current Nutanix software and hypervisor versions, and cumulative data about monitored clusters.

When Pulse is enabled, it sends a summary email of the cluster configuration to a Nutanix Support server daily by default. Pulse also collects the most important data like system-level statistics and configuration information more frequently to automatically detect issues and help make troubleshooting easier. With this information, Nutanix Support can apply advanced analytics to optimize your implementation and to address potential problems.

Note: Pulse sends messages through ports 80/8443/443 or through your mail server (see Configuring an SMTP Server on page 297). For the complete list of required ports, see Port Reference.

Note: When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse.

- To enable Pulse, click the **Continue** button in the message and follow the prompts.
- To continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

Pulse Transport Methods

Configure one of the following Pulse transport methods (in order of preference):

- [first choice] Enable Pulse and use Prism Central as a proxy for the Pulse data transmitted by each node (for clusters registered with Prism Central). Advantages: The configuration is

automatic (as described in Prism Central Proxy for Pulse Data on page 659), and no new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

- [second choice] Enable Pulse and configure an HTTP proxy server (see Configuring an HTTP Proxy on page 290). Advantage: No new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.
- [third choice] Enable Pulse and configure your firewall. Enable Pulse by using each Controller VM IP address in each managed cluster. See "Configuring Pulse" and "Pulse Access Requirements" in the *Prism Web Console Guide*. Disadvantage: New firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

Remote Diagnostics

Remote Diagnostics is a service that enables Nutanix Support to request granular diagnostic information from Pulse-enabled clusters. Pulse streams configuration data, metrics, alerts, events, and select logs back to Nutanix Support, but this information is aggregate in nature. The Pulse information is designed to provide a high-level state representation of the cluster. When the aggregate data is not detailed enough to diagnose a specific issue, Nutanix Support often needs to collect more diagnostic data from the cluster. Remote Diagnostics allows Nutanix Support to remotely collect the following data (and only the following data):

- Nutanix services logs
- Custom gflags being set for any Nutanix service
- Activity traces for Nutanix services
- Hypervisor logs
- Hypervisor config
- Cluster configuration
- System statistics like memory usage
- Nutanix NCC health check reports

Each time Remote Diagnostics triggers a collection, an entry is added to the audit trail for the cluster. There are always two entries, the start (initiation) and finish (termination) of the diagnostics collection.

Remote Diagnostics is enabled by default for every cluster where Pulse is enabled. If your security policy (or other consideration) does not allow Nutanix Support cluster access for remote diagnostics collection, you can disable Remote Diagnostics without turning off Pulse. Nutanix Support will still provide seamless and proactive support based on the Pulse data.

- To check the Remote Diagnostics status, SSH into a Controller VM and enter the following command:

  ```
  nutanix@cvm$ zkcat /appliance/logical/nusights/collectors/kCommand/override_config
  ```

- To disable Remote Diagnostics, SSH into a Controller VM and enter the following command:

  ```
  nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --enable=false --reason="text"
  ```

  The --reason argument is optional. Use it to enter a text string describing the reason for disabling Pulse.

- To enable Remote Diagnostics, SSH into a Controller VM and enter the following command:

  ```
  nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --enable=true --reason="text"
  ```

  The --reason argument is optional. Use it to enter a text string describing the reason for enabling Pulse.

Configuring Pulse

About this task

- See Pulse Health Monitoring on page 654 for Pulse configuration recommendations.
- Pulse sends messages through ports 80/8443/443, or if this is not allowed, through your mail server (see Configuring an SMTP Server on page 297). For the complete list of required ports, see Port Reference.
- When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse. To enable Pulse, click the **Continue** button in the message and follow the prompts; to continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

You can enable (or disable) Pulse at any time. To configure Pulse, do the following:

Procedure

1. Go to the Settings menu (see Settings Menu on page 17) and select **Pulse**.

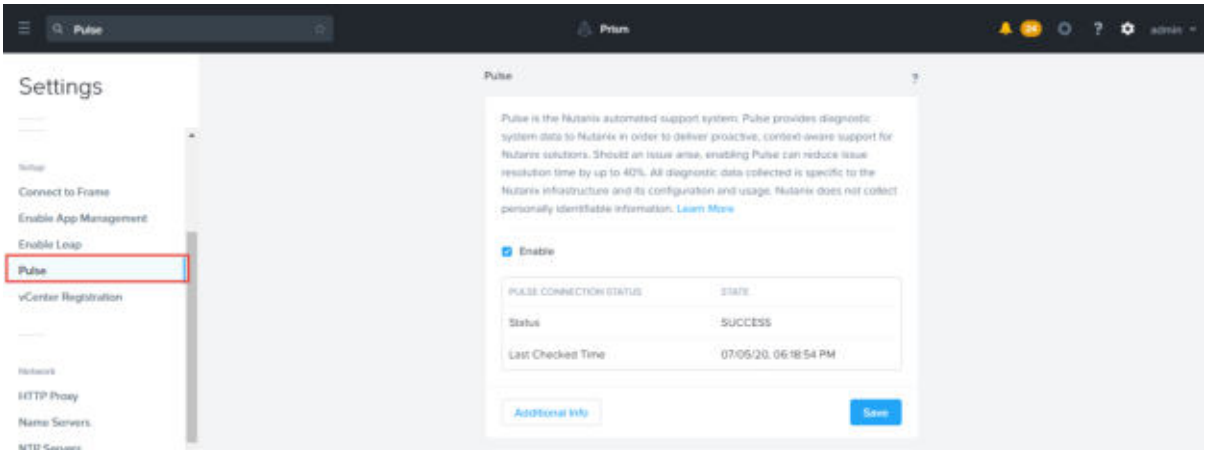

Figure 408: Pulse Window

2. To enable (disable) this feature, check (uncheck) the **Enable** box. Nutanix recommends enabling Pulse to allow Nutanix Support to receive system data and deliver proactive and context-aware support.

   When you enable Pulse, Pulse periodically sends data for troubleshooting purposes to Nutanix Insights tool and support team. Pulse is the underlying technology that securely transmits system-level diagnostic data to the Insights platform, enabling predictive health and context-aware support automation workflows. Nutanix Insights is an integrated service that utilizes this data to augment product support, reducing customer case volume and expediting issue resolution time.

   Nutanix does not share any data that Pulse sends, with any third parties unless permitted by your agreement with Nutanix or by the Nutanix Privacy Statement. Certain Nutanix

products require Pulse enablement for functionality and/or features. See the Nutanix Privacy Statement and applicable product documentation for more details.

You may elect to obfuscate (mask) entity names and IP addresses which are not obfuscated or masked by default. Entity names and IP addresses are not linked to individuals. These relate to Nutanix entities such as cluster names. (See KB 9326 for more information.)

- To check if partial scrubbing is enabled, use the following command:

  ```
  nutanix@cvm$ curl -H "Content-Type: application/json" -X GET -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
  ```

  If the output contains the string `"identificationInfoScrubbingLevel":"PARTIAL"` or `"identificationInfoScrubbingLevel":"AUTO"`, that means partial scrubbing is enabled. If the output contains the string `"identificationInfoScrubbingLevel":"ALL"`, that means `identificationInfoScrubbingLevel` is still set to `ALL` and the **Open In Prism** button on the Insights Portal remains deactivated.

- To obfuscate (mask) entity names and IP addresses, and update the PII scrub level for specific Prism Element clusters or Prism Central cluster, use the following command:

  ```
  nutanix@cvm$ curl -H "Content-Type: application/json" -X PUT -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse
  ```

- To obfuscate (mask) entity names and IP addresses, and update the PII scrub level for all registered Prism Element clusters from Prism Central, use the following command on a Prism Central cluster:

  ```
  nutanix@cvm$ curl -H "Content-Type: application/json" -X PUT -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
  ```

3. Click the **Additional info** button to display additional email recipient fields.

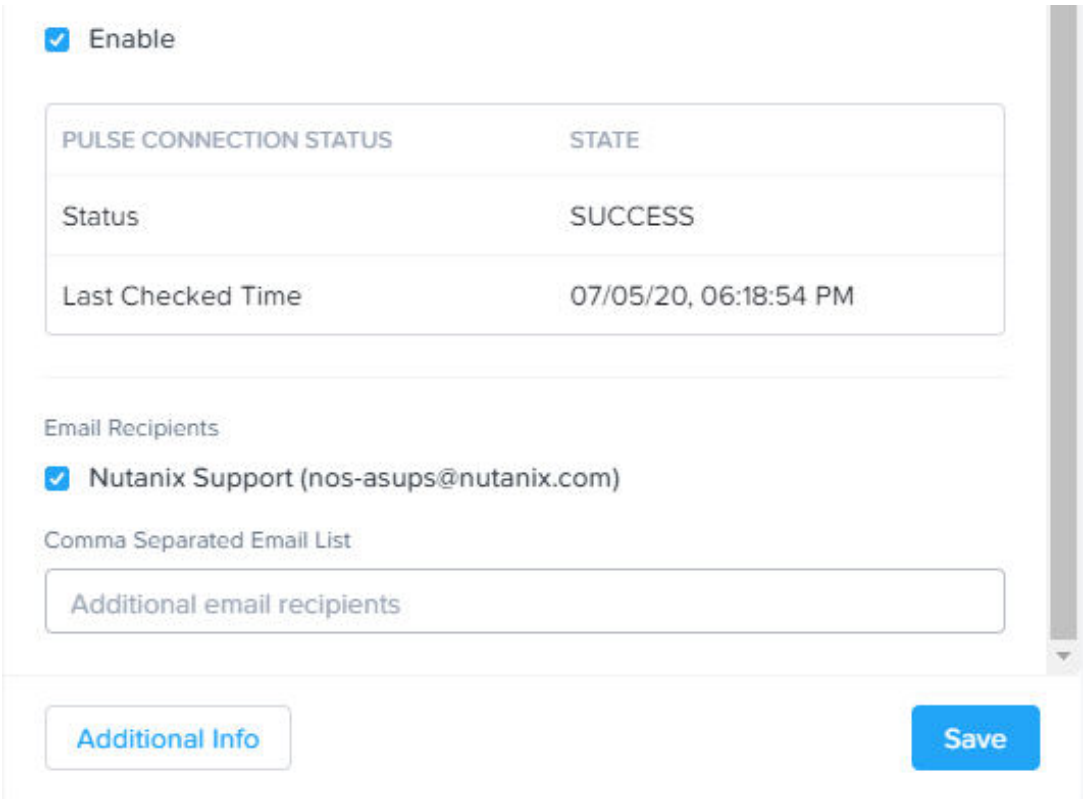

Figure 409: Pulse Additional Info

4. Check the **Pulse Connection Status** field for connection information.

   - **Status**: Displays the transport mechanism status, which is **Success** when the HTTP client can access the Pulse server successfully or **Failure** when it cannot (or **Unknown** when the status is not known).
   - **Last Checked Time**: The last time Pulse sent data and checked the connectivity with the Pulse server.

5. To add (remove) Nutanix customer support as a recipient of the cluster information, check (uncheck) the box next to **Nutanix Support (nos-asups@nutanix.com)** in the **Email Recipients** field.

   Cluster information is e-mailed automatically to Nutanix customer support at **nos-asups@nutanix.com** when this feature is enabled. If you do not want Nutanix customer support to receive the cluster information, you can uncheck this box at any time (or restart the e-mails by checking it again).

6. To send the cluster information to additional recipients, enter one or more e-mail addresses in the **Comma Separated Email List** field.

   In addition to (or instead of) sending the cluster information to Nutanix customer support, you can e-mail the information to yourself or others by entering recipient e-mail addresses in this field. Use a comma to separate multiple addresses.

7. Click the **Save** button to save the new setting and close the window.

Prism Central Proxy for Pulse Data

Prism Central can automatically act as a proxy for Pulse data transmitted by each node in a Prism Element cluster registered to that Prism Central instance.

How Do I Enable Prism Central Proxy for Pulse Data?

You do not have to explicitly enable this feature. It depends on your Prism Central and Prism Element configuration.

Pulse data from Prism Element nodes is automatically routed through Prism Central and then sent to Nutanix Support if you satisfy these requirements:

- You enable Pulse on Prism Central and each registered Prism Element cluster.
- Prism Central and each Prism Element cluster node are running NCC 3.5.2 or later.
- You have not configured an HTTP proxy on Prism Central or, if you have configured an HTTP proxy on Prism Central, it does not require basic authentication (a user name and password).
- You have not configured an HTTP proxy on Prism Element clusters registered to this Prism Central instance. If your Prism Element clusters are configured to use an HTTP proxy, the cluster nodes bypass the Prism Central Pulse proxy and transmit Pulse data to Nutanix Support.

If your Prism Central deployment is not available, the cluster nodes bypass the Prism Central Pulse proxy and transmit Pulse data to Nutanix Support.

For a Prism Central scale out deployment, each Prism Element node selects a Prism Central VM at random to act as its proxy.

Can I Use This Feature If I Have Configured an HTTP Proxy on Prism Central and Prism Element?

If you have configured an HTTP proxy on Prism Central, you can use this feature automatically if you satisfy these requirements:

- You have configured an HTTP proxy on Prism Central but it does not require basic authentication (a user name and password). Configuring an HTTP Proxy on page 290 describes setting up an HTTP proxy on Prism Central.
- Your Prism Element clusters are not configured to use a proxy. If your Prism Element clusters are configured to use a proxy, data is transmitted from each node to Nutanix support, bypassing the Prism Central Pulse proxy for pulse data.

Creating a Support Case

About this task

Nutanix customer support maintains a portal where you can get assistance by opening a support case and viewing the status of your open cases (see Accessing the Nutanix Support Portal on page 666). However, you can also create a support case directly from Prism Central. To create a support case or view information about your open cases, do the following:

> Note: A support portal connection is required before you can create a case (see Licensing on page 31). In addition, at least one cluster must be registered with Prism Central. Also, this feature is available only to customers who get support directly from Nutanix (such as NX and SX model customers), not customers who get initial support from third parties.

Procedure

1. Click the question mark icon in the main menu and select **Create Support Case** from the pull-down list.

   The **Create new support case** page appears. One of the following appears on this page:

   - If there are no open cases currently, a blank page with fields for creating a case appears.
   - If you have one or more open support cases, summary information about those cases appears (see Viewing Case Status on page 665). To create a case from this page, click

the **Create new support case** button (upper right). The blank page with fields for creating a case will appear.

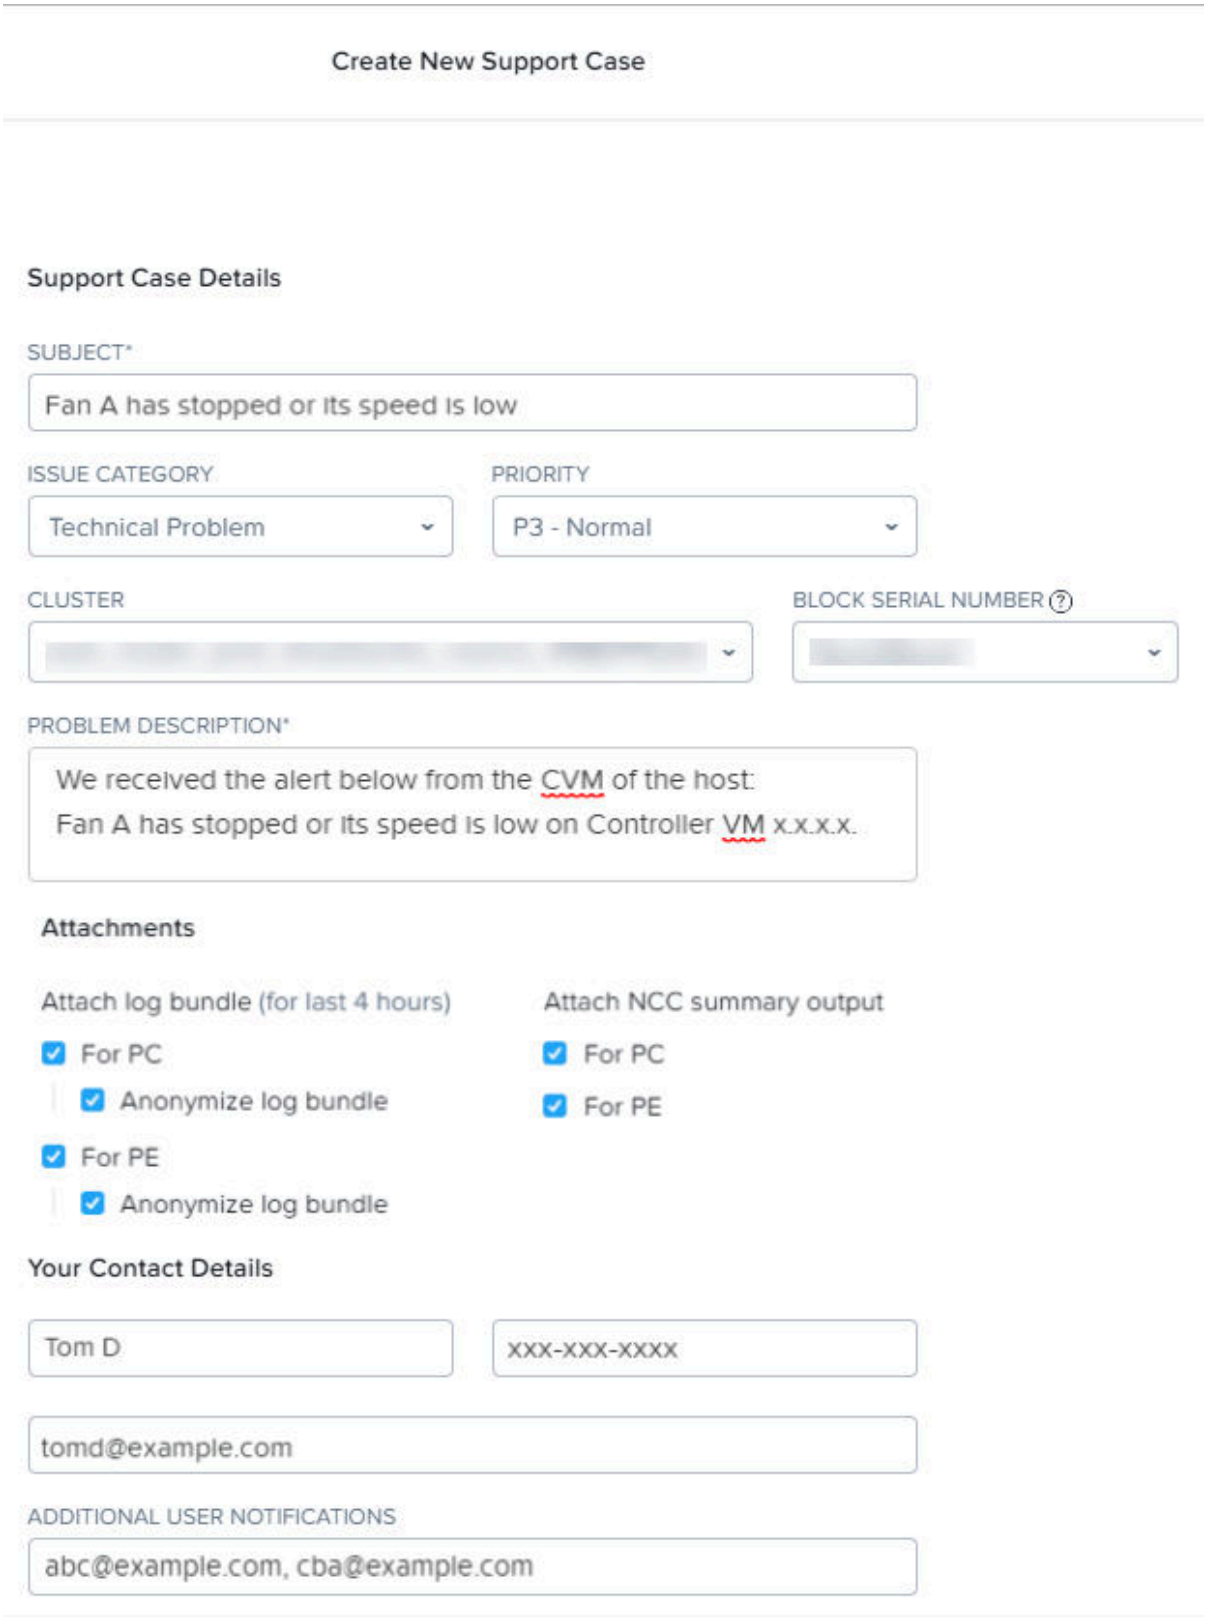

Figure 410: Create New Support Case Page

2. Do the following in the indicated fields:

   a. **Subject**: Enter a title that briefly describes the issue.

   b. **Issue Category**: Select the type of issue from the pull-down list.

      The types are **Technical Problem**, **Question**, **Comment/Feedback**, **Scheduled Upgrade/ Maintenance**, **RFE** (request for enhancement), **Licensing**, **Non Technical issue**, and **Technical Problem - Foundation**.

   c. **Priority**: Select the priority for this issue from the pull-down list.

      There are four priority levels based on the severity of the problem plus a request for enhancement option. Select the appropriate priority based on the following descriptions.

      - **P1 - Emergency**. System is not available and productivity has been halted. Product is unusable in its current state.
      - **P2 - Critical**. System is available but experiencing issues which have a direct impact on productivity. Major inconvenience.
      - **P3 - Normal**. System is having an occasional issue that has been identified as needing to be resolved, but the issue has not greatly affected productivity. Minor inconvenience.
      - **P4 - Low**. Questions about documentation, processes, or procedures. General requests about information.
      - **RFE - Request for Enhancement**. Feature requests for the product which would improve the experience or functionality for the customer.

   d. **Cluster**: Select the target cluster from the pull-down list.

   e. **Block Serial Number**: Select the serial number of the asset (node) in question from the pull-down list.

      If you need more information, see the "Installed Base" page on the Nutanix support portal for a list of your assets and the corresponding serial numbers.

   f. **Tell Us More**: Enter a description of your issue in the text box.

      Include any relevant details that might help Nutanix customer support analyze and resolve your issue.

   g. **Attach log bundle (for last 4 hours)**: Check this box to attach log files from the last four hours to the case. If you want the attached logs to be anonymized (personally identifiable information removed), check the **Anonymize log bundle** box.

      - **For PC**: Check this box to attach Prism Central log files from the last four hours to the case. If you want the attached logs to be anonymized (personally identifiable information removed), check the **Anonymize log bundle** box.
      - **For PE**: Check this box to attach Prism Element log files from the last four hours to the case. If you want the attached logs to be anonymized (personally identifiable information removed), check the **Anonymize log bundle** box.

      > Note: If the target cluster is not connected to the Internet, such as in a dark site, a log bundle or NCC summary report (following step) cannot be attached to the case through this method.

   h. **Attach NCC summary output**

- **For PC**: When this box is checked, NCC is run, and the summary output of Prism Central is attached. This ensures that the NCC (health checks) results reflect the current state.
- **For PE**: When this box is checked, NCC is run, and the summary output of Prism Element is attached. This ensures that the NCC (health checks) results reflect the current state.

i. **Your Contact Details**: Enter your name, phone number, and e-mail address in the indicated fields. Case updates are sent to your e-mail address. You can also have the case updates sent to (up to two) additional users by entering their e-mail addresses in the **Additional User Notification** field.

3. When the entered information is completed, click the **Submit** button.

   This submits the case to Nutanix customer support for review and response. The speed of response reflects the priority of the issue, where P1 and P2 issues have a higher priority and more rapid response than less critical issues, but Nutanix customer support strives to handle all cases as quickly as possible.

   A Prism Central task is started to create the case. If you selected to attach a log bundle and/or NCC summary output, a task is started for each selected item. The tasks run asynchronously, and you can check the progress (case creation, log collection, and health check tasks) through the tasks dashboard (see Tasks View on page 190). If a task encounters a problem, you can hover the cursor over the "Failed" status to see a brief explanation. Automatically generated comments (preceded by "Prism Central Automated Note:") are added to the case when the log/health checks collection starts and when the upload completes. When the tasks complete, the output (log bundle and NCC summary report) appear as attachments in the case summary (see Viewing Case Status on page 665).

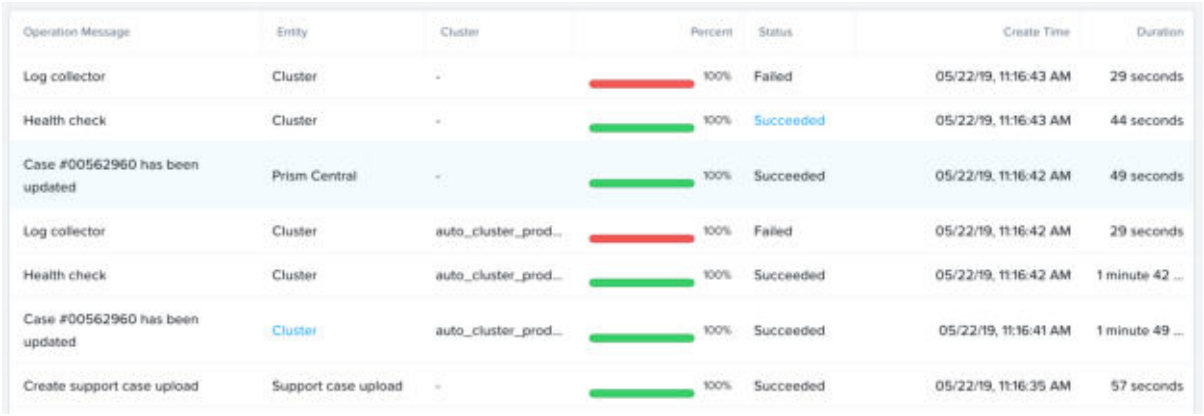

Figure 411: Tasks Dashboard for Prism Central (create case tasks)

Figure 412: Tasks Dashboard for Prism Element (create case tasks)

Viewing Case Status

You can view information about open support cases directly from Prism Central by clicking the question mark icon in the main menu and selecting **Create Support Case** from the pull-down list. This opens the **Create new support case** page, which displays summary information about open cases (if any). When you have one or more open cases, this page displays the following sections:

- **Open Cases** column (on the left). There is an entry (line) for each open case that includes the case name and number; click the desired case entry to displays summary information about that case in the main section of the page.
- Main section (in the middle). This section displays the creation date, description, attachments (if any), and other information related to the selected case. This is an abridged version of the full case details available from the Nutanix support portal (see below).
- **Key Insights** column (on the right). This section includes the following fields:
  - **Status**: Displays the current status of the case.
  - **Created By**: Displays the name of the user who created the case.
  - **Type**: Displays the type (category) of issue the user specified when creating the case.
  - **Case number**: Displays the number assigned to this case.
  - **Serial Number**: Displays the serial number of the asset (node) in question.
  - **Access Portal**: Includes a link to the Nutanix support portal. Clicking the **View Details in Portal** link opens the details page for that case in the support portal in a new tab or window. The details page in the support portal provides the full details and history of that case should you desire more information, and you can submit additional material to the case such as screen shots or other relevant information.

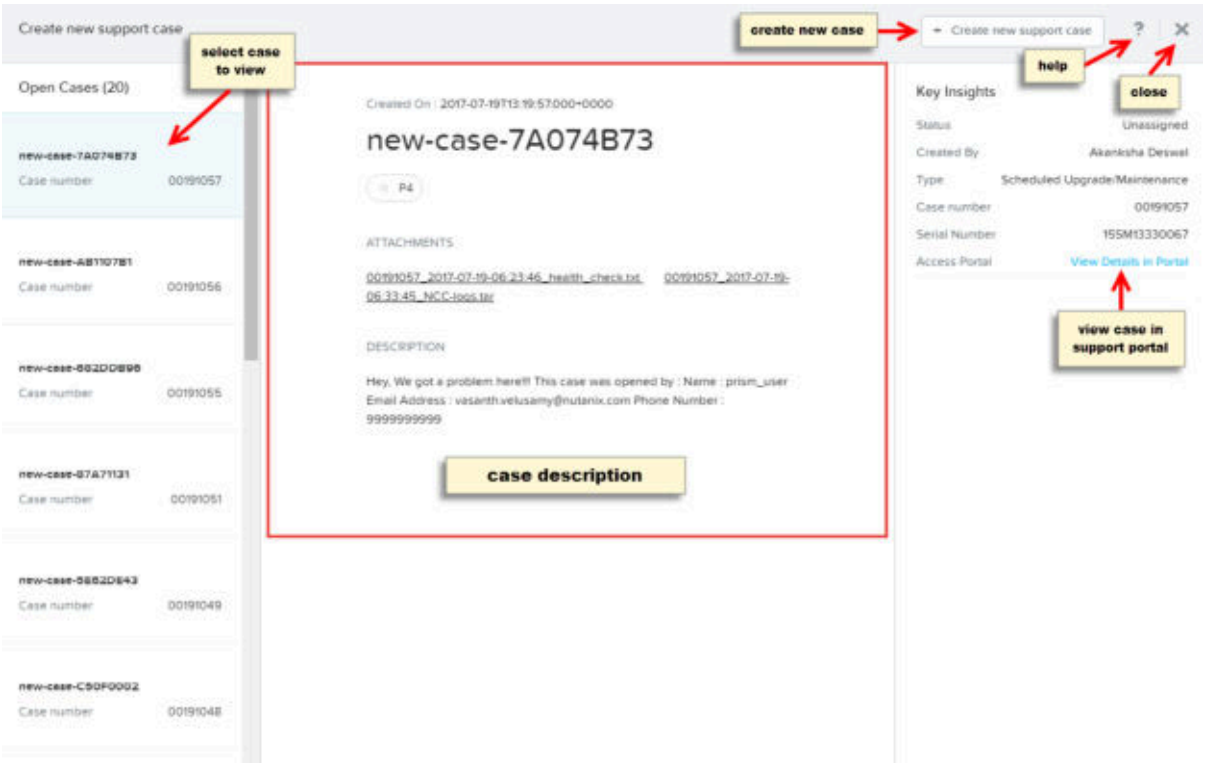

Figure 413: Create New Support Case Page (open cases)

Accessing the Nutanix Support Portal

About this task

Nutanix provides a variety of support services and materials through its support portal.

Procedure

1. To access the Nutanix support portal from Prism Central, select **Support Portal** from the user icon pull-down list of the main menu.

   The login screen for the Nutanix support portal appears in a new tab or window.

2. Enter your support account user name and password.
   The Nutanix support portal home page appears.

3. Select the desired service from the screen options.

   You can select an option from one of the main menu pull-down lists or search for a topic at the top of the screen, click one of the icons (Documentation, Open Case, View

Cases, Downloads) in the middle, or view one of the selections at the bottom such as an announcement or KB article. The following table lists the menu options.

> Note: Some options have restricted access and are not available to all users.

Table 248: Main Menu Options

Category	Option	Description
Documentation	Software Documentation	Displays a page from which you can view the Nutanix software manuals.
	Hardware Replacement Documentation	Displays a page from which you can view the Nutanix hardware replacement manuals.
	Knowledge Base	Displays a page from which you can view the knowledge base (KB) articles.
	Solutions Documentation	Displays a page from which you can view documents that describe how to implement the Nutanix platform to solve a variety of business applications.
	EOL Information	Displays a page from which you can view the end of life policy and bulletins.
	Field Advisories	Displays a page from which you can view field advisories.
	Training	Provides a link to the separate Nutanix training portal.
	Security Advisories	Displays a page from which you can view security advisories.
	AOS Upgrade Paths	Displays a page where you can see the supported AOS release upgrade paths.
	Compatibility Matrix	Displays a page from which you can view a compatibility matrix broken down (filtered) by hardware model, AOS version, hypervisor type and version, and feature version (NCC, Foundation, BMC/BIOS).
	Webinar Recordings	Displays a page with links to a selection of Nutanix training webinars.
Support	Open Case	Displays a form to create a support case.
	View Cases	Displays a page from which you can view your current support cases.
	.NEXT Forums	Provides a link to the (separate) Nutanix Next Community forum.
	Terms & Conditions	Displays a page from which you can view various warranty and terms and conditions documents.
Downloads	AOS (NOS)	Displays a page from which you can download AOS releases.

Category	Option	Description
	Acropolis File Services (AFS)	Displays a page from which you can download the Acropolis File Services.
	Acropolis Container Services (ACS)	Displays a page from which you can download the Acropolis Container Services.
	Hypervisor Details	Displays a page from which you can download Acropolis hypervisor versions. You can also download supporting files used when manually upgrading a hypervisor version (AHV, ESXi, or Hyper-V).
	Prism Central	Displays a page from which you can download the Prism Central installation bundle. There are separate bundles for installing on AHV, ESXi, or Hyper-V.
	Phoenix	Displays a page from which you can download Phoenix ISO files.
	Foundation	Displays a page from which you can download Foundation releases.
My Products	Installed Base	Displays a table of your installed Nutanix appliances, including the model type and serial number, location, and support coverage.

Category	Option	Description
	Licenses	Displays a table of your product licenses along with buttons to add or upgrade licenses for your clusters.

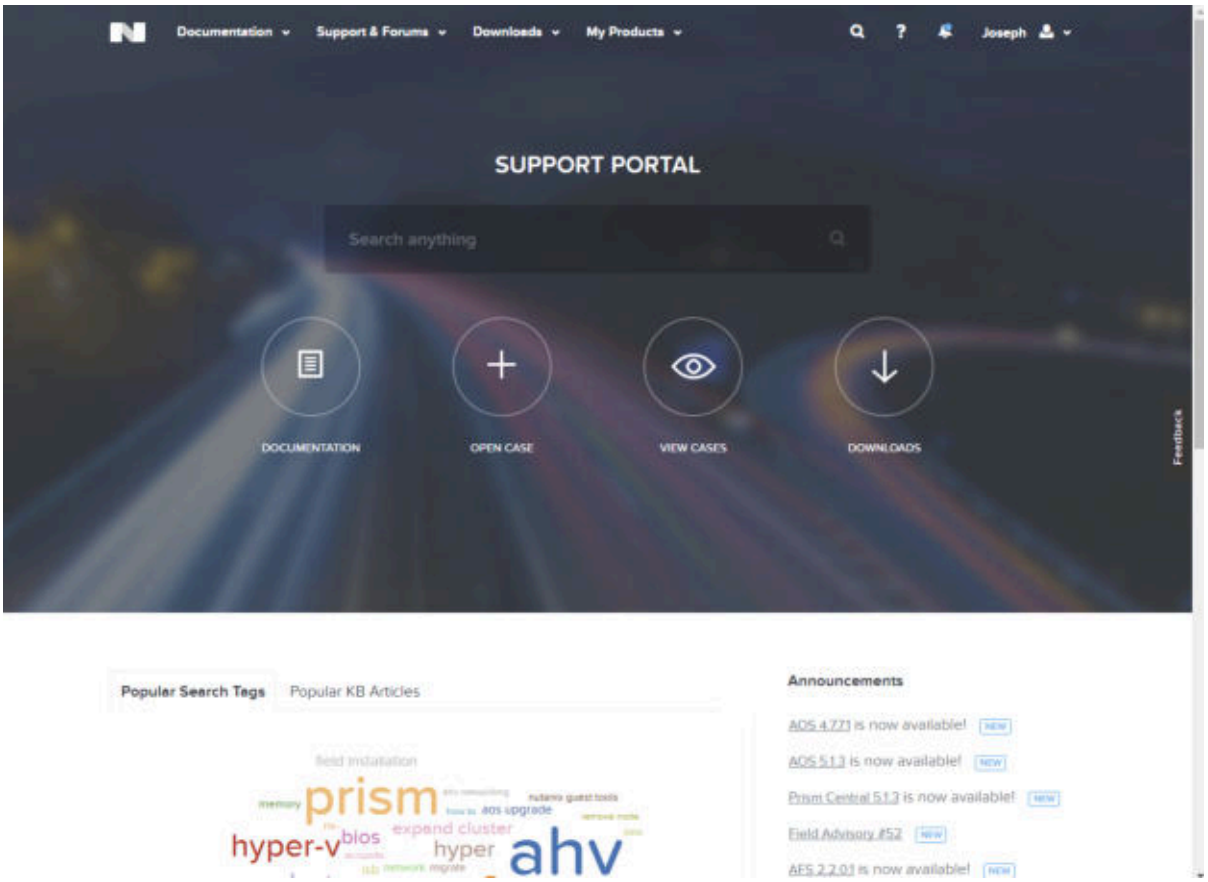

Figure 414: Nutanix Support Portal

Accessing the REST API Explorer

About this task

Nutanix provides a utility with the web console to help you get started with the REST API. The Explorer displays the parameters and format for the API calls that can be included in scripts. Sample API calls can be made to show the type of output you should expect to receive.

The v3 API can be viewed in the REST API Explorer.

Procedure

1. Log on to the web console, click the user icon in the upper-right corner of the web console, and click **REST API Explorer**.

   The REST API Explorer displays a list of the objects that can be managed by the API. Each line has four options:

   - **Show/Hide**: Expand or reduce the detail shown for the object
   - **List Operations**: Show all operations that can be run on this object
   - **Expand Operations**: Show the detailed view of the operations that can be run on this object

2. Find the line for the object you want to explore and click **Expand Operations**.
   For example, click **alerts**.
3. Click **GET** to show the details for this API call.
4. Click **Try it out!** to test the API call when used with your cluster.

15

HELP RESOURCES

There are several information sources that you can access at any time when you need help:

- Prism Central help documentation (see Accessing Online Help on page 671).
- Nutanix customer support portal (see Accessing the Nutanix Support Portal on page 666).
- Nutanix community forum (see Accessing the Nutanix Next Community on page 673).
- Glossary of terms (see Glossary on page 674).

Accessing Online Help

About this task

Prism Central includes online help documentation that you can access at any time.

Procedure

1. To open the online help, choose one of the following from the question mark icon pull-down list of the main menu:

   » Select **Help with this page** to display help documentation that describes the current screen.

   > Note: In a task window click the question mark icon in the upper right to display the help documentation for that window.

   » Select **Online Documentation** to display the Help Organization page.

   A context-sensitive help page or the Help Organization page appears in a new tab or window. (These pages are located on the Nutanix support portal.) The Help Organization page provides descriptions of the major help topics with links to the entry page for each

major topic. The display includes a breadcrumb at the top to navigate through the help pages.

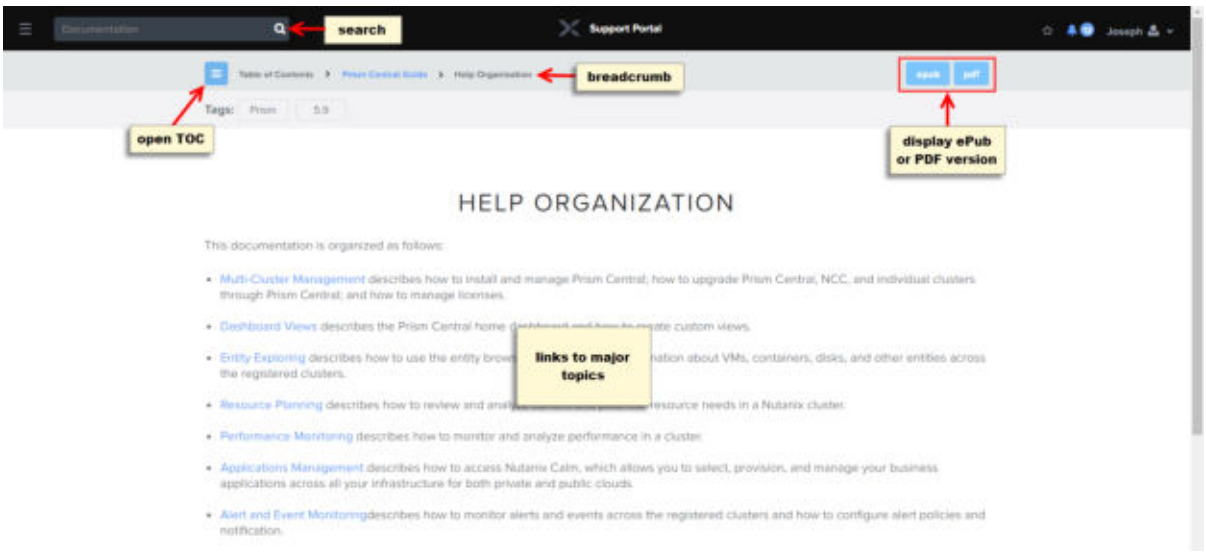

Figure 415: Help Organization Page

2. To select a topic from the table of contents, click the collapse menu icon (also know as a "hamburger" button) in the upper left.

   A table of contents pane appears on the left. Click a topic in the table of contents to display that topic.

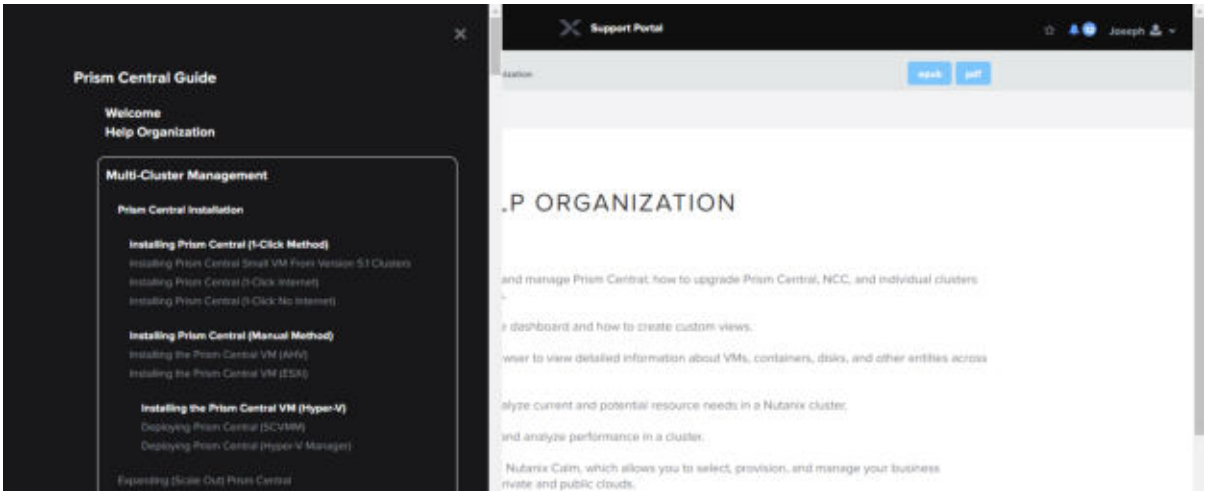

Figure 416: Table of Contents Pane

3. To display all the help contents as a single document, click the **epub** or **pdf** button in the upper right.

   You can view the *Prism Central Guide* in either ePUB or PDF format by selecting the appropriate button. If your browser does not support the selected format, you can download the PDF or ePUB file.

4. To search for a topic, click the magnifying glass icon in the main menu bar and enter a search string in the field.

   This searches not only the help contents, but also all the documentation, knowledge base articles, and solution briefs. Matching results appear below the search field. Click a topic from the search results to display that topic.

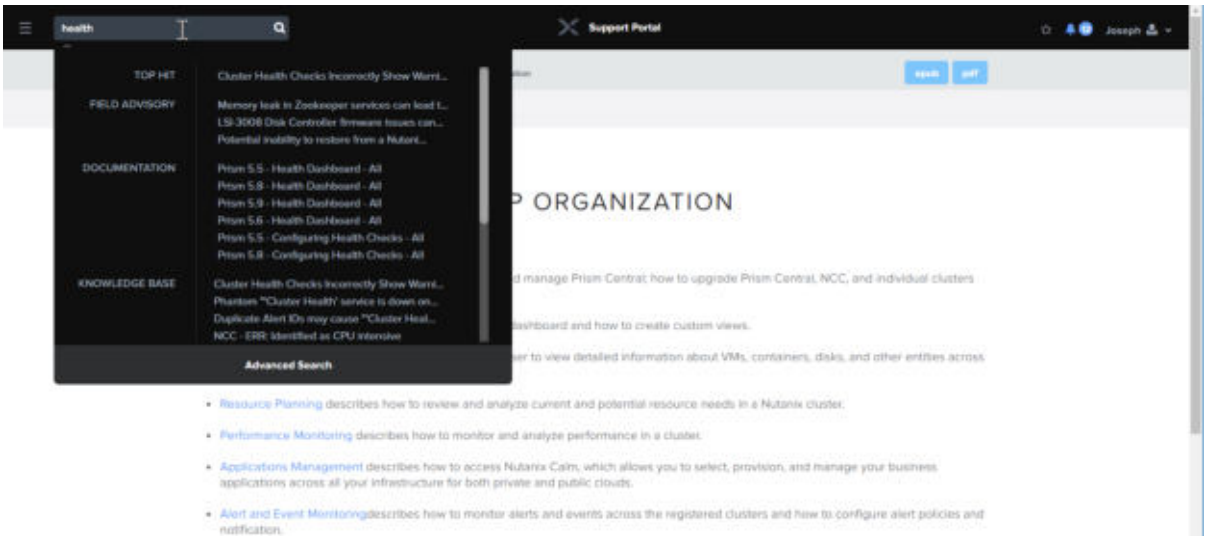

Figure 417: Search Results Example

Accessing the Nutanix Next Community

Nutanix maintains a community forum for customers and partners to facilitate a peer-to-peer exchange of ideas, tips, and information about Nutanix technologies and the rapidly changing landscape of data center IT.

Procedure

- To access the Nutanix next community forum from Prism Central, select **Nutanix Next Community** from the question mark icon pull-down list of the main menu (seeMain Menu on page 10.

  The Nutanix Next Community main page appears in a new tab or window. From this page you can search existing posts, ask questions, and provide comments.

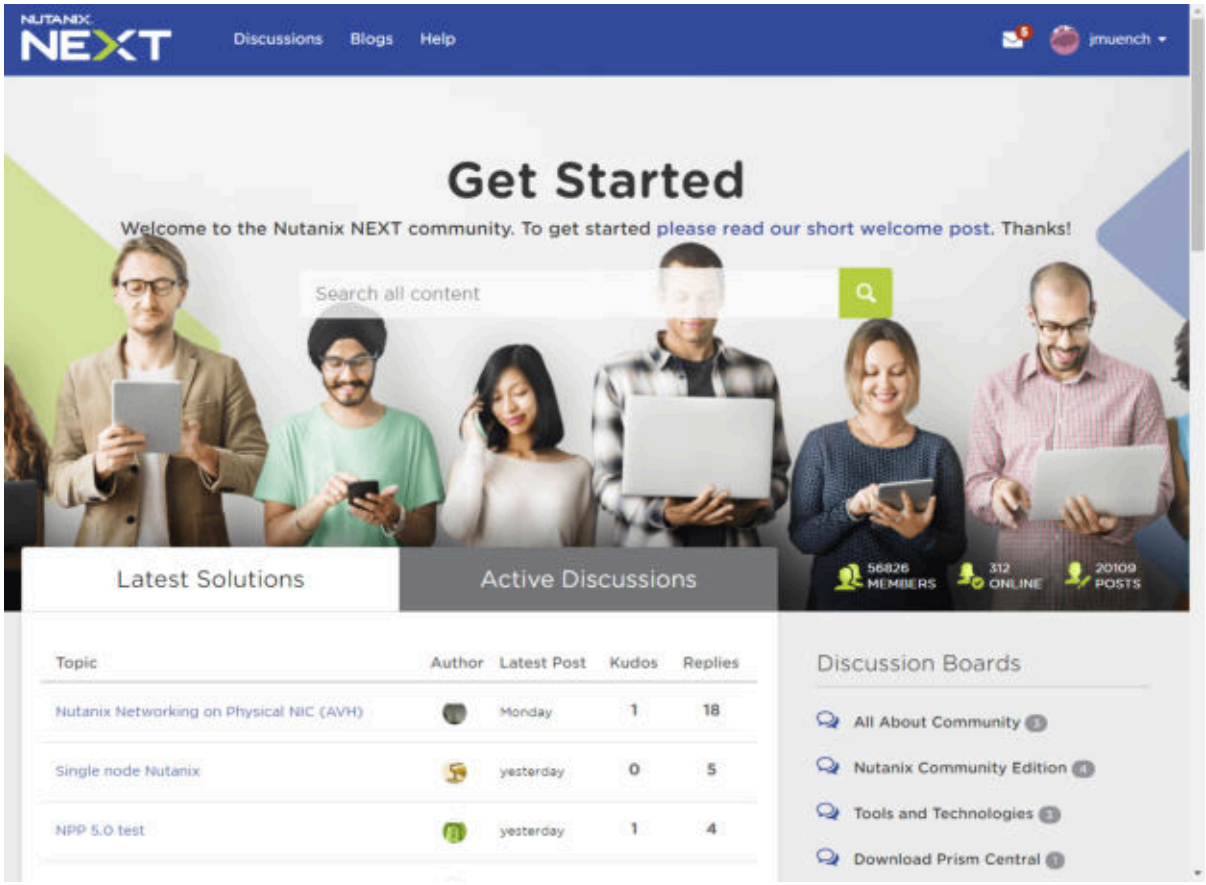

Figure 418: Next Community Screen

Glossary

aCLI

Acropolis command-line interface.

Acropolis

The Nutanix converged software fabric for virtualization and storage management. It consists of the Acropolis base software, Acropolis Distributed Storage Fabric, AHV, App Mobility Fabric, Prism, and Acropolis APIs.

Acropolis App Mobility Fabric

Provides virtualization management, volume management, and other distributed management functions for Acropolis.

AHV

Nutanix-specific hypervisor host providing core server virtualization and optimized for Nutanix cluster and guest VM use.

block

A set of Nutanix nodes contained in a single enclosure.

block fault tolerance

When certain conditions are met, redundant copies of any data required to serve I/O are placed on nodes that are not in the same block. In the case where multiple nodes in a block fail, the cluster can continue to run because cluster configuration data exists on other blocks. Guest VMs can continue to run because redundant copies of guest VM data and metadata exist on other blocks.

clone

A writeable copy of a vDisk.

Cloud Connect

A feature that enables you to back up and restore copies of virtual machines and files to and from an on-premise cluster and a Nutanix Controller VM located on the Amazon Web Service (AWS) cloud. The cloud-based cluster is managed as a remote site from the Prism Web Console or nCLI.

cluster

A group of nodes contained in one or more Nutanix blocks.

cold data

Data that did not have many rewrites or overwrites for a long time. For example, snapshots, file servers, archives, or backups.

compression

An optional storage container setting that enables one of two types of compression.

compression, inline

Data is compressed as it is written. This type of compression is recommended for workloads that perform batch processing.

compression, post-process

Data is compressed after it is written. The delay time between write and compression is configurable. Nutanix has recommends a delay value of 60. However, every workload has a different I/O profile, so you might require a different delay value otherwise. This type of compression is recommended for most workloads.

Consistency Group

A subset of VMs in a Protection Domain. All VMs within a Consistency Group for that Protection Domain will be snapshotted in a crash-consistent manner. For all VMs in a Consistency Group, a snapshot creates one snapshot for all VMs in the group.

Controller VM

A Nutanix VM that manages storage and other cluster functions on a node.

data resiliency

A measure of the types of failures a cluster can withstand; affected by block awareness and redundancy factor.

datastore

A logical storage container for files necessary for VM operations.

deduplication

The sharing of identical guest VM data on premium tiers (RAM and Flash) for improved performance or on capacity tiers (HDD) for storage space savings. Enabled by properties of a storage container or vDisk.

Distributed Storage Fabric

All storage functionality, including snapshots and clones, data protection, disaster recovery, data path redundancy, replication factors, deduplication, compression, erasure coding, and so on.

Enterprise Cloud Platform

The Nutanix solution that natively converges compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence

erasure coding

Optional algorithm included in the Acropolis base software to help reduce the storage used for fault tolerance. It helps to increase the effective or usable capacity on a cluster, depending on cluster size and data coldness.

failback

Planned failover initiated from recovery site.

failover

Moving VMs from a primary site to a recovery site.

failover, disaster

Failover when the primary site is down.

failover, planned

Failover when both sites are up.

guest VM

A VM running on a Nutanix cluster that executes a workload, such as VDI or Exchange, as opposed to a VM that is involved in cluster operations, such as a Controller VM.

host

An instance of the hypervisor that runs on a Nutanix node.

image service

A workflow in the Prism web console that enables a user to upload ISO or disk images (in ESXi or Hyper-V format) to a Nutanix AHV cluster by specifying a remote repository URL or by uploading a file from a local machine.

local replication

Multiple copies of data within a storage container. These copies exist for fault tolerance: if a physical disk fails, the cluster can recover data from another copy. The cluster manages the replicated data, and the copies are not visible to the user.

local snapshot

Snapshots stored on the same cluster where they originated.

nCLI

Nutanix command-line interface.

node

A physical server contained in a Nutanix block; runs a hypervisor host.

oplog

A write cache on a faster, more expensive storage tier.

Prism

Web-based management interface for managing Nutanix clusters.

Prism Central

Centralized management tool that runs as a separate VM configured as a single-node cluster to monitor and manage multiple clusters through a single web console.

Prism Element

A single cluster being managed by and available through the Prism Central web console.

Protection Domain

A group of VMs to be backed up locally on a cluster or replicated on the same schedule to one or more clusters. Protection Domains may be associated with remote sites. Protection Domain names must be unique across sites. A VM can be in at most one Protection Domain.

Protection Domain, active

A Protection Domain that manages live VMs and makes, replicates, and expires snapshots.

Protection Domain, inactive

A Protection Domain that receives snapshots from a remote cluster.

redundancy factor

The number of nodes plus 1 that the cluster can tolerate being down at one time. By default, Nutanix clusters have a redundancy factor of 2, which means that they can tolerate 1 node being down. They are configurable to redundancy factor 3 to enable tolerating 2 nodes being down.

remote replication, one-to-one

Replicates a production cluster with one idle cluster as a DR target.

remote replication, reciprocal

Cross replication within running (non-idle) production clusters.

remote site

A pair of clusters that can replicate data to each other. A configured remote site can also be located in the cloud (based on Amazon AWS, for example).

remote snapshot

A snapshot copied asynchronously from one cluster to another.

replication factor

The number of copies of data kept by a storage container. The default is 2. Storage Containers on clusters with redundancy factor of 3 can have replication factor of 3.

reserved capacity

A property of a storage container or vDisk that guarantees that a certain amount of storage space is available.

schedule

A property of a Protection Domain that specifies the intervals to take snapshots and how long the snapshots should be retained. A schedule optionally specifies which remote site to replicate to.

Self-Service Restore

Allows a user to restore a file within a virtual machine from the Nutanix protected snapshot with minimal Nutanix administrator intervention.

shadow clone

A cache of a vDisk on all the nodes in the cluster. When a vDisk is read by multiple VMs (such as the base image for a VDI clone pool), the cluster creates shadow clones of the vDisk.

snapshot

A read-only copy of the state and data of a VM at a point in time. Snapshots for a VM are crash consistent, which means that the VMDK on-disk images are consistent with a single point in time. That is, the snapshot represents the on-disk data as if the VM crashed. The snapshots are not, however, application consistent, meaning that application data is not quiesced at the time of snapshot, unless the Protection Domain is configured to use application-consistent snapshots.

storage container

A subset of available storage within a storage pool.

storage pool

A group of physical disks from one or more tiers.

storage replication adapter

A Nutanix-provided module that allows VMware Site Replication Manager (SRM) to use native remote replication.

tier

A type of physical storage in a Nutanix node.

vDisk

Data associated with a VM represented as a set of files on a datastore.

VM high availability

In virtualization management VM high availability, when a node becomes unavailable, VMs that are running on that node are restarted on another node in the same cluster.

VM mobility

The ability to export your existing VMs from one non-AHV cluster to an AHV cluster. This option requires that you install the Nutanix VM Mobility installer on all the VMs. The Nutanix VM Mobility installer deploys the drivers that are required at the destination AHV cluster. After you prepare the source VMs, they can be exported to the AHV cluster.

vStore

A separate mount point within a storage container which has its own NFS namespace. This namespace maps to a Protection Domain. Each vStore is exported as a device through the Nutanix SRA.

COPYRIGHT

Copyright 2022 Nutanix, Inc.

Nutanix, Inc.
1740 Technology Drive, Suite 150
San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

License

The provision of this software to you does not grant any licenses or other rights under any Microsoft patents with respect to anything other than the file server implementation portion of the binaries for this software, including no licenses or any other rights in any hardware or any devices or software that are used to communicate with or in connection with this software.

Conventions

Convention	Description
variable_value	The action depends on a value that is unique to your environment.
ncli> *command*	The commands are executed in the Nutanix nCLI.
user@host$ *command*	The commands are executed as a non-privileged user (such as nutanix) in the system shell.
root@host# *command*	The commands are executed as the root user in the vSphere or Acropolis host shell.
> *command*	The commands are executed in the Hyper-V host shell.
output	The information is displayed as output from a command or in a log file.

Version

Last modified: March 9, 2022 (2022-03-09T19:51:02+05:30)

BOOKLIST

Related Guides

Refer to the following guides for more information:

- *Acropolis Upgrade Guide*
- *Security Guide (AOS Security 5.18)*
- *Data Protection and Recovery with Prism Element*

Prism pc.2022.1

Prism Central Guide

March 22, 2022

Contents

Welcome viii

Help Organization ix

1. Prism Central Overview 10
- Main Menu 10
- Entities Menu 13
- Settings Menu 16
- Searching for Information 20
- Understanding Displayed Statistics 28

2. Prism Central Administration 30
- Installing or Upgrading 30
- Licensing 30
- Prism Central Management 30
  - Logging Into Prism Central 31
  - Logging Out of Prism Central 34
  - Managing Prism Central 34
  - Finding the Prism Central Version 43
  - Finding the AHV Version on Prism Central 44
  - Finding the AOS Version Using Prism Central 45
  - Register (Unregister) Cluster with Prism Central 46
  - Expanding (Scale Out) Prism Central 53
  - Shutdown (Start Up) Prism Central VM 56
  - Prism Central Disaster Recovery 57
  - IP Address Reconfiguration 60

3. Main Dashboard 64
- Main Dashboard Display 66
- Widgets on a Dashboard 68
- Creating a New Dashboard 73
- Adding Dashboard Widgets 74
- Modifying a Dashboard 79

4. Entity Exploring 81
- Compute and Storage Entities 86
  - VMs Summary View 86
  - VM Template Summary View 109
  - OVAs View 113
  - Images Summary View 117
  - Catalog Items View 121
  - Storage Containers Summary View 122
  - Storage Policies Summary View 133
  - Volume Groups Summary View 139

vCenter Datastores Summary View....149
Catalog Items View....156
Categories Summary View....157
OVAs View....160
Network and Security View....164
Subnets....165
Virtual Private Clouds....169
Floating IPs....174
Connectivity....174
Security Policies Summary View....182
Policies....185
Affinity Policies Summary View....186
NGT Policies View....187
Image Placement Policies Summary View....188
Bandwidth Throttling Policies Summary View....190
Security Policies Summary View....191
Data Protection and Recovery Entities....193
Protection Policies View....194
Recovery Plans View....195
VM Recovery Points View....196
Hardware Entities....197
Clusters Summary View....197
Hosts Summary View....217
Disks Summary View....227
GPUs Summary View....234
Activity Entities....237
Audits Summary View....238
Tasks View....241
Operations Entities....244
Reports View....245
Administration Entities....246
Users Summary View....246
Roles Summary View....250
Projects Summary View....254
Availability Zones View....259
LCM View....260
Services Entities....260

5. Alert and Event Monitoring....261
Alerts Summary View....261
Alert Details....269
Events Summary View....272
Event Details....274
Configuring Alert Emails....275
Configuring Alert Policies....278
Modifying System Alert Policies....279
Adding Custom Alert Policies....281
Viewing External Alert Policies....296
Stopping Alerts During a Maintenance Window....297
Prism Central Logs....299
Audit Log Events....302
Alerts/Health checks....304
Cluster....304
Controller VM....356
DR....367

License 426
Guest VM 427
Hardware 437
Node 439
Network 444
Nutanix Files 455
Other 481
Prism Central VM 493
Storage 493
System Indicator 504
VM 505

6. System Management 506
Configuring Prism Central When a Cluster Uses Proxy Servers 506
Who Needs to Use the Whitelist Method 507
Configuring an HTTP Proxy 508
Whitelisting Prism Central and Its Managed Clusters (nCLI) 513
Configuring Name Servers 514
Configuring NTP Servers 514
Configuring an SMTP Server 515
Configuring SNMP 516
Configuring Syslog Monitoring 522
Syslog Modules 523
Modifying UI Settings 526
Configuring a Banner Page 530
Internationalization (i18n) (Prism Central) 530
Localization (L10n) (Prism Central) 531
Changing the Language Settings (Prism Central) 531

7. Security and User Management 534

8. Compute and Storage (Cluster) Administration 535
VM Management 535
Creating a VM (AHV) 536
Creating a VM (Self Service) 546
Managing a VM (AHV) 551
Managing a VM (Self Service) 562
Adding Multiple vGPUs to the Same VM 570
Live Migration of vGPU-enabled VMs 572
Creating a VM (ESXi) 576
Managing a VM (ESXi) 581
Exporting a VM as an OVA 587
NGT Management in Prism Central 588
Storage Quality of Service (QoS) 595
VM Policy Management 598
VM Template Management 604
Limitations of VM Template Feature 604
Creating a VM Template 605
Deploying VM from a Template 606
Managing a VM Template 609
Storage Management 611
Storage Components 612
Storage Container Management 621

Creating a Volume Group 637
Modifying a Volume Group 641
Managing Volume Group Connections 642
Managing Volume Group Virtual Disks 646
Storage Policy Management 647
Storage Policy Based Encryption 649
Creating or Updating a Storage Policy 650
Managing Associations 653
Deleting a Storage Policy 655
Catalog Management 657
Adding a Catalog Item 658
Deleting a Catalog Item 660
Image Management 661
Requirements 662
Limitations 662
Creating an Image 662
Modifying an Image 678
Importing Images to Prism Central 679
Image Policy Management 680
OVA Management 688
Exporting a VM as an OVA 689
Uploading an OVA 690
Concatenating Upload using APIs 696
Deploying an OVA as VM 698
Downloading an OVA 704
Renaming an OVA 704
Deleting an OVAs 705
Performing Other Administrative Tasks 706
vCenter Server Integration 707
Registering vCenter Server (Prism Central) 709
Managing vCenter Server Registration Changes (Prism Central) 712

9. Prism Self Service Administration 713
Prism Self Service Overview 713
Configuring Prism Self Service 714
Migrating Prism Self Service 717
Project Management 718
Creating a Project 718
Modifying a Project 721

10. Network and Security Management 723
Network Connections 723
Configuring Network Connections 725
Creating or Updating Virtual Switch 728
Enabling LAG and LACP on the ToR Switch (AHV Only) 733
Modifying Network Connections 733

11. Category Management 735
Creating a Category 735
Modifying a Category 736
Assigning a Category 737

12. Operations Management....739
Performance Monitoring....739
Analysis Dashboard....739
Sessions Dashboard....765
Closing the Session....771
Switching Between Sessions....771
Application Discovery....771
Application Discovery View....772
Enabling Application Discovery....781
Connecting to the Cloud....787
Adding (Removing) Clusters to Monitor....788
Configuring vCenter Authentication....790
Discovering Applications....791
Publishing Applications....792
Creating Discovery Policies....792
Cost Management (Xi Beam)....795
Application Monitoring....796
Application Monitoring Summary View....796
Configuring Application Monitoring....809
Modifying Application Monitoring....812
Operations Policy Management....812
Operations Policies View....813
Modifying an Operations Policy....814
Enabling an Operations Policy....816
Resource Planning....816
Capacity Runway View....817
Scenarios View....819
Updating Capacity Configurations....837
Behavioral Learning Tools....838
Task Automation....841
Overview....842
Action Gallery....842
Plays....842
Playbook Actions....843
X-Play Integrations....854
Creating Playbooks using Triggers....865
Running a Playbook (Manual Trigger)....875
Configuring Manual Parameters....875
Exporting or Importing Playbooks....879
PagerDuty Integration with Prism....880
Reports Management....883
Creating a New Report....884
Managing a Report....885
Viewing Report Instances....888
Views in the Reports....898
Scheduling a Report....920
Downloading Reports....922
Configuring Report Settings....923
Role Based Access Control....926

13. Services Enablement....928
Enabling Calm....928
Enabling Files....929

Enabling Foundation Central....929
Enabling Karbon....929
Enabling Objects....930

14. Customer Support Services....931
Pulse Health Monitoring....931
Configuring Pulse....933
Prism Central Proxy for Pulse Data....936
Creating a Support Case....936
Viewing Case Status....942
Accessing the Nutanix Support Portal....943
Accessing the REST API Explorer....946

15. Help Resources....948
Accessing Online Help....948
Accessing the Nutanix Next Community....950
Glossary....951

Copyright....957
License....957
Conventions....957
Version....957

Booklist....958

WELCOME

You can monitor and manage entities across Nutanix clusters through a web console called Prism Central. This documentation describes how to use Prism Central.

- For a list and description of major topics, see Help Organization on page ix.
- For help with viewing this documentation, see Accessing Online Help on page 948.
- To access other Nutanix documents, see Nutanix support portal.

HELP ORGANIZATION

This documentation is organized as follows:

- Prism Central Overview on page 10 provides an overview of Prism Central including menu options and navigation tools.
- Prism Central Administration on page 30 describes how to install and manage Prism Central; how to upgrade Prism Central, NCC, and individual clusters through Prism Central; and how to manage licenses.
- Main Dashboard on page 64 describes the main (home) dashboard and how to customize it.
- Entity Exploring on page 81 describes how to use the entity browser to view detailed information about VMs, containers, disks, and other entities across the registered clusters.
- Resource Planning on page 816 describes how to review and analyze current and potential resource needs in a Nutanix cluster.
- Performance Monitoring on page 739 describes how to monitor and analyze performance in a cluster.
- Services Enablement on page 928 describes how to enable certain applications (Calm, Karbon, and Objects) through Prism Central.
- Alert and Event Monitoring on page 261 describes how to monitor alerts and events across the registered clusters and how to configure alert policies and notification.
- System Management on page 506 describes how to configure various system settings (NTP, SMTP, and so on) for Prism Central.
- Security and User Management on page 534 describes how to configure user authentication, manage local user accounts, control user access using roles, install SSL certificates, and control remote (SSH) access.
- Compute and Storage (Cluster) Administration on page 535 describes how to administer individual clusters through Prism Central such as how to create and manage VMs.
- Prism Self Service Administration on page 713 describes how to implement self-service features through Prism Central.
- Policies on page 185 describes how to create and apply policies to manage and control access and operations in a cluster.
- Reports Management on page 883 describes how to create and run reports about infrastructure resources.
- Task Automation on page 841 describes how to automate routine administrative tasks through Prism Central by using the X-play feature.
- Customer Support Services on page 931 describes how to configure Pulse (system diagnostics reporting feature), access the Nutanix customer support portal, and create a support ticket through Prism Central when you need help.
- Help Resources on page 948 describes how to use the online help to find what you need and how to access the Nutanix Next Community portal.

1

PRISM CENTRAL OVERVIEW

Nutanix provides an option to monitor and manage multiple clusters through a single web console. This multi-cluster view, known as Prism Central, is a centralized management tool that runs as a separate instance comprised of either a single VM or a set of VMs. Prism Central provides the following features:

- Single sign on for all registered clusters.
- Customizable main dashboard that displays summary information across the registered clusters (see Main Dashboard on page 64).
- Virtual infrastructure dashboards with drill-down options to manage VMs, storage containers, catalog items, images, categories, recoverable entities, and subnets across the registered clusters (Compute and Storage Entities on page 86).
- Dashboards to manage security policies, protection policies, recovery plans, Nutanix Guest Tools (NGT) policies, and image placement (see Policies on page 185).
- Hardware component dashboards with drill-down options to view detailed information about individual clusters, hosts, disks, and GPUs across the registered clusters (see Hardware Entities on page 197).
- Activity monitors for alerts, events, audits, and tasks (see Activity Entities on page 237).
- Tools to analyze system activity, plan for resource needs, create usage reports, and automate routine administrative tasks (see Operations Entities on page 244).
- Dashboards to manage inventory (LCM), projects, roles, users, and availability zones (see Administration Entities on page 246).
- Service connection pages to enable Calm, Karbon, and Buckets (see Services Entities on page 260).
- Settings menu from which you can configure Prism Central functions (see Settings Menu on page 16).

Main Menu

The main menu at the top of every screen provides access to the features of Prism Central. This section describes each of the main menu options.

Viewing Choices

Clicking the collapse menu button (also know as a "hamburger" button) in the far left of the main menu displays the entities menu (see Entities Menu on page 13). The Main dashboard is the first screen that appears after logging into Prism Central (see Main Dashboard on page 64).

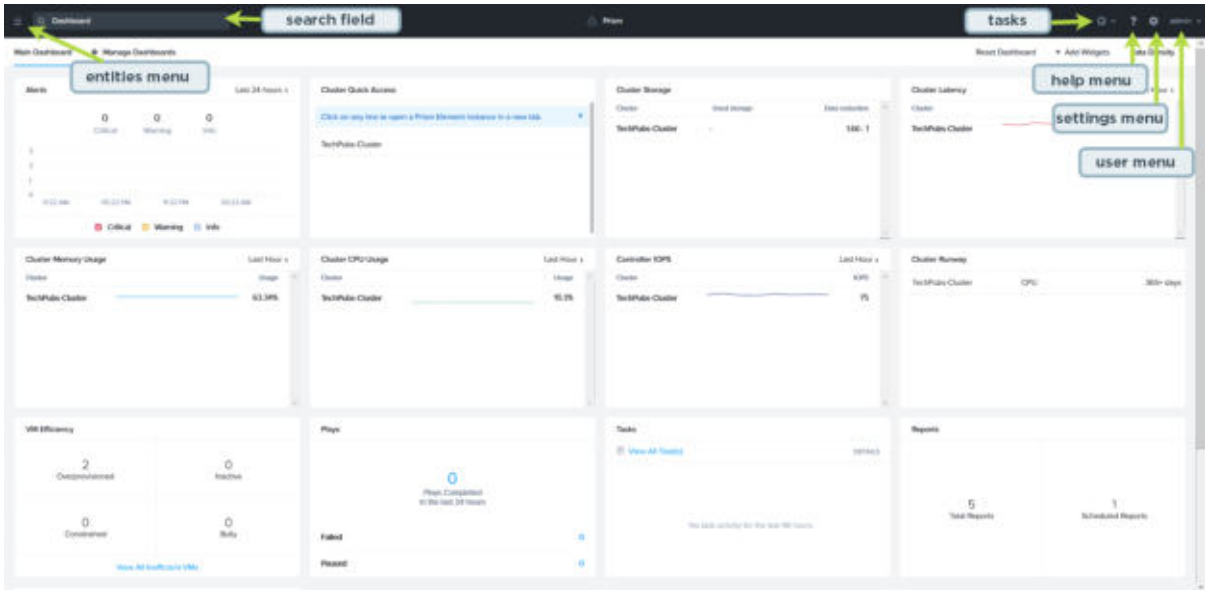

Figure 1: Prism Central Main Menu

Search, Alert, and Task Options

The main menu provides status information about active alerts and tasks plus access to advanced search capability:

- A field with a magnifying glass icon appears on the left of the main menu. Enter a string in this field to search for relevant content in Prism Central (see Searching for Information on page 20).
- An alerts icon appears on the right of the main menu when there are current critical (red), warning (yellow), or informational (gray) alert messages. The number of active alerts is displayed in the icon. Click the alert icon to display information about those alerts in the Alerts dashboard (see Alerts Summary View on page 261).
- A tasks icon appears to the right of the alerts when there are current tasks (running or completed within the last 48 hours). The icon is blue when all the tasks are running properly or have completed successfully. If one or more of those tasks do not complete successfully, the icon turns yellow (warning) or red (failure) . Click the tasks icon to see a drop-down list of the current tasks; click **View All Tasks** at the bottom of the list to display the tasks dashboard (see Tasks View on page 241).

Help Menu ("?" Icon)

A question mark icon appears on the right side of the main menu. Clicking the question mark displays a list of help resource options that you can select. The following table describes each option in the pull-down menu.

Table 1: Help Menu Options

Name	Description
Learn about search	Displays search guidelines (see Searching for Information on page 20).

Name	Description
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help on page 948).
Online Documentation	Opens the online help at the introduction page (see Accessing Online Help on page 948).
Support Portal	Opens a new browser tab (or window) at the Nutanix Support portal logon page (see Accessing the Nutanix Support Portal on page 943).
Nutanix Next Community	Opens a new browser tab (or window) at the Nutanix Next Community entry page (see Accessing the Nutanix Next Community on page 950). The portal is an online community site for customers and partners to exchange ideas, tips, and information about Nutanix technologies and related data center topics.
Create Support Case	Opens the **Create a new support case** page to view or create support cases with Nutanix customer support (see Creating a Support Case on page 936).

Settings Menu (Gear Icon)

A gear icon appears on the right side of the main menu. Clicking the gear icon launches the **Settings** page and displays a menu of tasks you can perform (see Settings Menu on page 16).

User Menu (<user_name> Icon)

A user icon appears on the far right side of the main menu with the current user logon name. Clicking the user icon displays a list of options to update your user account, log off from Prism Central, and other miscellaneous tasks. The following table describes each option in the pull-down menu.

Table 2: User Menu Options

Name	Description
Change Password	Opens the **Change Password** window to update your password (see “” ()).
Update Profile	Opens the **Update Profile** window to update your user name and email address (see “” ()).
Download Cmdlets Installer	Downloads the PowerShell installer for the Nutanix cmdlets. See the Nutanix *Command Reference* for information about installing the cmdlets locally and for cmdlet descriptions.
Download nCLI	Downloads the Nutanix command-line interface (nCLI) as a zip file to your local system. The download occurs immediately after clicking this option (no additional prompts). See the Nutanix *Command Reference* for information about installing the nCLI locally and for nCLI command descriptions.

Name	Description
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page (see Accessing the REST API Explorer on page 946).
About Nutanix	Opens the **About Nutanix** window that displays Nutanix operating system (AOS) and other version information (see Finding the Prism Central Version on page 43).
Nothing To Do?	Opens a game that is strictly for entertainment. To quit the game, click the "X" at the upper right of the screen.
Sign Out	Logs you out of Prism Central (see Logging Out of Prism Central on page 34).
Adjust Contrast (Chrome only)	Displays a contrast setting box at the bottom of the screen where you can set the Prism Central display to **Normal** (default) or **High** contrast. This option is available only when using the Chrome browser.

Entities Menu

Clicking the collapse menu ("hamburger") button in the main menu (see Main Menu on page 10) displays a menu of entities on the left. The following table describes each entity in the menu.

- The primary menu options are **Dashboard**, seven category names (**Compute & Storage**, **Network & Security**, **Hardware**, **Activity**, **Operations**, **Administration**, and **Services**), and **Global Settings**. Clicking a category name displays a submenu of entities in that category.
- Clicking the entity name displays a dashboard or other appropriate page for that entity.
- For quicker access to one or more submenu entities, you can promote a submenu option to the primary menu by bookmarking that entity (clicking the star icon to the right of the name). This displays that entity name below the **Dashboard** entry. You can remove the bookmarked entry by clicking the star again.
- The entity menu pane closes automatically when you click in the display area to the right. To keep the entity menu pane open all the time, click the **Lock Navigation Bar** button at the bottom. The button is a toggle; to unlock the navigation bar, click the button again.
- When the entities menu is displayed, the collapse menu button changes to an X button. Click the X button to hide the menu.

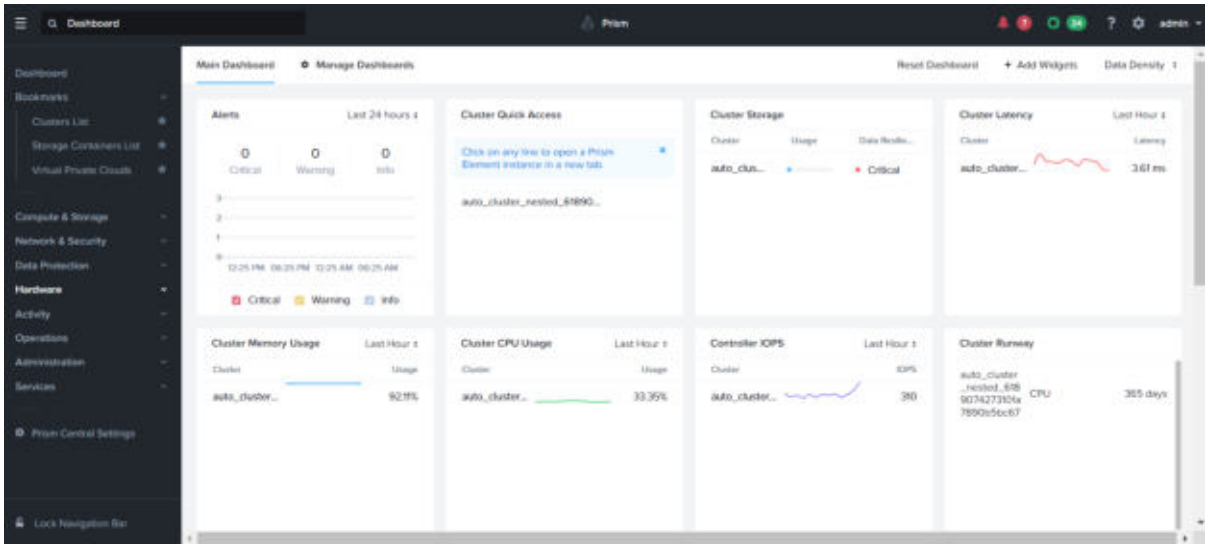

Figure 2: Entities Menu

Table 3: Entities Menu Options

Category	Name	Description
Dashboard		Displays the main dashboard (see Main Dashboard on page 64).
Compute & Storage	VMs	Displays the VMs dashboard (see VMs Summary View on page 86).
	Storage Containers	Displays the Storage Containers dashboard (see Storage Containers Summary View on page 122).
	Volume Groups	Displays the Volume Groups dashboard (see Volume Groups Summary View on page 139).
	vCenter Datastores	Displays the vCenter Datastores dashboard (see vCenter Datastores Summary View on page 149). This option appears when there is at least one registered vCenter instance.
	Catalog Items	Displays the Catalog dashboard (see Catalog Items View on page 121).
	Images	Displays the Images dashboard (see Images Summary View on page 117).
	OVAs	Displays the OVAs dashboard (see OVAs View on page 113).
Network & Security	Subnets	Displays the subnets dashboard (see Subnets View).
	Virtual Private Clouds	Displays the virtual private clouds dashboard (see Network and Security Entities).
	Floating IPs	Displays the floating IPs dashboard (see Network and Security Entities).

Category	Name	Description
	Connectivity	Displays the VPNs dashboard (see Network and Security Entities).
	Security Policies	Displays the Security Policies dashboard (see Security Policies Summary View on page 191).
Data Protection and Recovery	Protection Summary	Displays the protection summary dashboard.
	Protection Policies	Displays the Protection Policies dashboard (see Protection Policies View on page 194).
	Recovery Plans	Displays the Recovery Plans dashboard (see Recovery Plans View on page 195).
	VM Recovery Points	Displays the VM Recovery Points dashboard (see VM Recovery Points View on page 196).
Hardware	Clusters	Displays the Clusters dashboard (see Clusters Summary View on page 197).
	Hosts	Displays the Hosts dashboard (see Hosts Summary View on page 217).
	Disks	Displays the Disks dashboard (see Disks Summary View on page 227).
	GPUs	Displays the GPUs dashboard (see GPUs Summary View on page 234).
Activity	Alerts	Displays the Alerts dashboard (see Alerts Summary View on page 261)
	Events	Displays the Events dashboard (see Events Summary View on page 272).
	Audits	Displays the Audits dashboard (see Audits Summary View on page 238).
	Tasks	Displays the Tasks dashboard (see Tasks View on page 241).
Operations	Analysis	Displays the Analysis dashboard (see Analysis Dashboard on page 739).
	App Discovery	Displays the App Discovery dashboard or the enable wizard if application discovery is not enabled (see Application Discovery on page 771).
	Cost Management	Provides access to the Xi Beam site (see Cost Management (Xi Beam) on page 795).
	Integrations	Displays the Monitoring Integrations dashboard or the enable wizard if monitoring integrations is not enabled (see Application Monitoring on page 796).
	Ops Policies	Displays the Operations Policies dashboard (see Operations Policy Management on page 812).

Category	Name	Description
	Planning	Displays the capacity Planning dashboard (see Resource Planning on page 816).
	Playbooks	Displays the list of playbooks (see Task Automation on page 841).
	Reports	Displays the Reports dashboard (see Reports View on page 245).
Administration	Categories	Displays the Categories dashboard (see Categories Summary View on page 157).
	LCM	Displays the life cycle manage (LCM) dashboard (see LCM View on page 260).
	Projects	Displays the Projects dashboard (see Projects Summary View on page 254).
	Roles	Displays the Roles dashboard (see Roles Summary View on page 250).
	Users	Displays the Users dashboard (see Users Summary View on page 246).
	Availability Zones	Displays the Availability Zones dashboard (see Availability Zones View on page 259).
Services	Calm	Displays the Calm feature page or the enable wizard if Calm is not enabled(see Enabling Calm on page 928).
	Files	Displays the Files dashboard or the enable wizard if Files is not enabled (see Enabling Files on page 929).
	Foundation Central	Displays the Foundation Central dashboard or the enable wizard if Foundation Central is not enabled (Enabling Foundation Central on page 929).
	Karbon	Displays the Karbon feature page or the enable wizard if Karbon is not enabled (see Enabling Karbon on page 929).
	Objects	Displays the Objects feature page or the enable wizard if Objects is not enabled (see Enabling Objects on page 930).
Prism Central Settings		Displays the settings menu (see Settings Menu on page 16).

Settings Menu

Prism Central includes a Settings page from which you can configure a variety of system services. You can access the Settings page by doing either of the following:

- Click the gear icon in the main menu (see Main Menu on page 10).
- Click the collapse menu ("hamburger") button and then select **Prism Central Settings** from the entities menu (see Entities Menu on page 13).

The Settings page displays a menu of tasks (on the left) you can perform. Click the task to open the window or page for that option in the pane to the right. The following table describes each menu option.

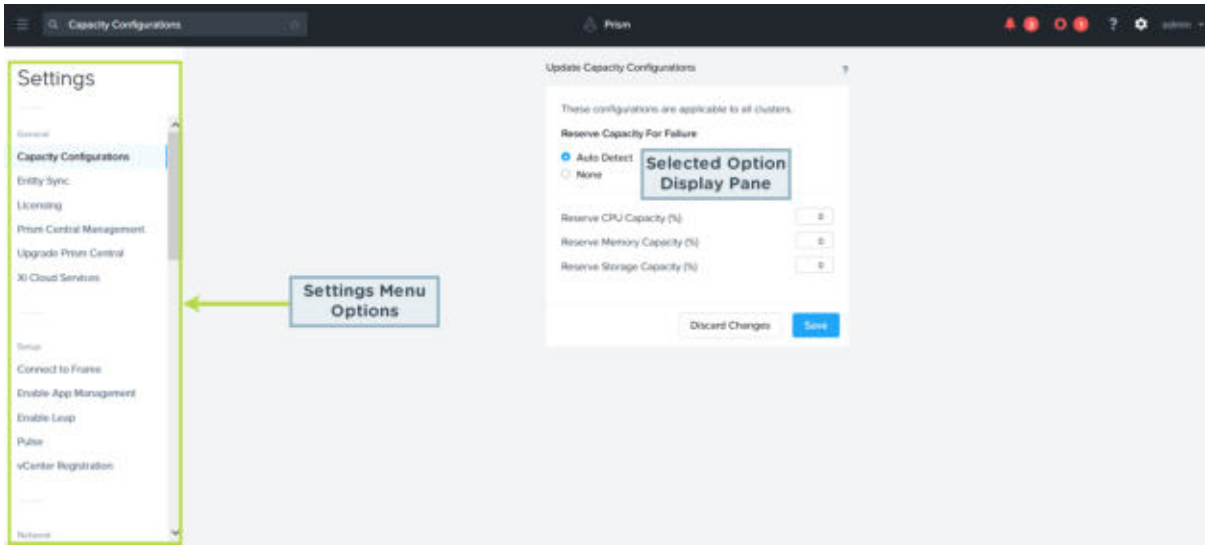

Figure 3: Settings Page

Table 4: Settings Menu

Category	Name	Description
General	Capacity Configurations	Opens the **Update Capacity Configurations** window, which allows you to configure sizing rules for cluster runway analysis and alerting (see Updating Capacity Configurations on page 837). This menu option does not appear if Prism Pro is disabled.
	Entity Sync	Opens the **Force Entity Sync** window, which synchronizes entities with the specified availability zones. See the *Leap Administration Guide* for more information.
	Licensing	Opens the **Licensing** page to install or update the Prism Central license that enables entitled Nutanix features (see Licensing on page 30).
	Migrate SSP	Opens the **Migrate Prism Self Service to Prism Central** window to migrate a pre-AOS 5.5 Prism Self Service (hosted in Prism Element) to Prism Central (see Migrating Prism Self Service on page 717).
	Prism Central Management	Opens the **Manage Prism Central** page for viewing information about Prism Central and optionally expanding (scale out) Prism Central across multiple VMs (see Managing Prism Central on page 34).

Category	Name	Description
	Upgrade Prism Central	Opens the **Upgrade Prism Central** window to upgrade the Prism Central VM to a newer version (see Installing or Upgrading on page 30).
	Xi Cloud Services	Opens the **Xi Cloud Services** windows, which allows you to log in to the Xi portal. See the *Xi Cloud Services Administration Guide* for more information.
Setup	Connect to Frame	Opens the **Connect to Frame** page to connect to Xi Frame. See the Xi Frame documentation for more information.
	Enable App Management	Opens the **Enable App Management** page to enable the Nutanix Calm features (see Enabling Calm on page 928).
	Enable Leap	Opens the **Leap** page to enable the Disaster Recovery as a Service (DRaaS) capability. See the *Leap Administration Guide* for more information.
	Pulse	Opens the **Pulse** page to enable (disable) the Nutanix automated support system (see Configuring Pulse on page 933).
	Self-Service Admin Management	Opens the Self-Service Admin Management page to modify the self-service configuration (see Configuring Prism Self Service on page 714). This menu option does not appear until a directory service is configured.
	vCenter Registration	Opens the **vCenter Registration** window to register (or unregister) clusters with vCenter (see Registering vCenter Server (Prism Central) on page 709).
Network	HTTP Proxy	Opens the **HTTP Proxies** window to configure an HTTP proxy to which the Nutanix software can connect (see Configuring an HTTP Proxy on page 508).
	Name Servers	Opens the **Name Servers** window to configure name servers (see Configuring Name Servers on page 514).
	NTP Servers	Opens the **NTP Servers** window to specify which NTP servers to access (see Configuring NTP Servers on page 514).
	SNMP	Opens the **SNMP Configuration** window to configure SNMP (see Configuring SNMP on page 516).
Flow	ID Based Security	Opens the **ID Based Security** page to add Active Directory domain services configurations and to import user groups for identity-based security policies. See the *Nutanix Flow Guide* for more information.

Category	Name	Description
	Microsegmentation	Opens the **Microsegmentation** page to enable this feature, which is disabled by default. Before you can configure and use application security policies, isolation environment policies, and quarantine policies, you must enable the feature. See the *Nutanix Flow Guide* for more information.
Security	Cluster Lockdown	Opens the **Cluster Lockdown** window, which allows you to delete (or add) public authentication keys used for SSH access into Prism Central. Removing all public keys locks down Prism Central from external access. For details, see “” () .
	SSL Certificate	Opens the **SSL Certificates** window to create a self-signed certificate. For details, see “” () .
User and Roles	Authentication	Opens the **Authentication Configuration** window to configure Prism Central authentication. For details, see “” () .
	Local User Management	Opens the **Local User Management** window. This window lists current users and allows you to add, update, and delete local user accounts. For details, see “” () .
	Role Mapping	Opens the **Role Mapping** window to configure role mappings that apply in the user authentication process. For details, see “” () .
Alerts and Notifications	Alert Email Configuration	Opens the **Alert Email Configuration** window, which allows you to configure rules and templates that apply to the alert messages from selected clusters (see Configuring Alert Emails on page 275).
	SMTP Server	Opens the **SMTP Server Settings** window to configure an SMTP server (see Configuring an SMTP Server on page 515).
	Syslog Server	Opens the **Syslog Servers** window to configure a syslog server (see Configuring Syslog Monitoring on page 522).
Appearance	Language Settings	Opens the **Language Settings** window to localize the Prism Central screen text (see Changing the Language Settings (Prism Central) on page 531).
	UI Settings	Opens the **UI Settings** window to configure Prism UI background themes, disable (or re-enable) the login screen background video and configure the automatic logout timeout (see Modifying UI Settings on page 526).
	Welcome Banner	Opens the **Edit Welcome Banner** window to create a welcome banner message that appears before users login to Prism Central (see Configuring a Banner Page on page 530).

Searching for Information

The main menu includes a search field on the left (see Main Menu on page 10) that allows you to find information about selected entities in various ways. If you need help at any time to navigate Prism Central or apply a search filter, you can click the question mark icon in the main menu and select **Learn about search**. This displays the Search Guidelines page that explains the search rules and options.

Search Basics

An entity is an object type such as a VM, cluster, node, security policy, project, report, event, alert, or audit. The search field is context-sensitive, which means it is populated automatically based on where you are in Prism Central. The first screen after logging on is the main dashboard, so "Dashboard" appears in the search field initially.

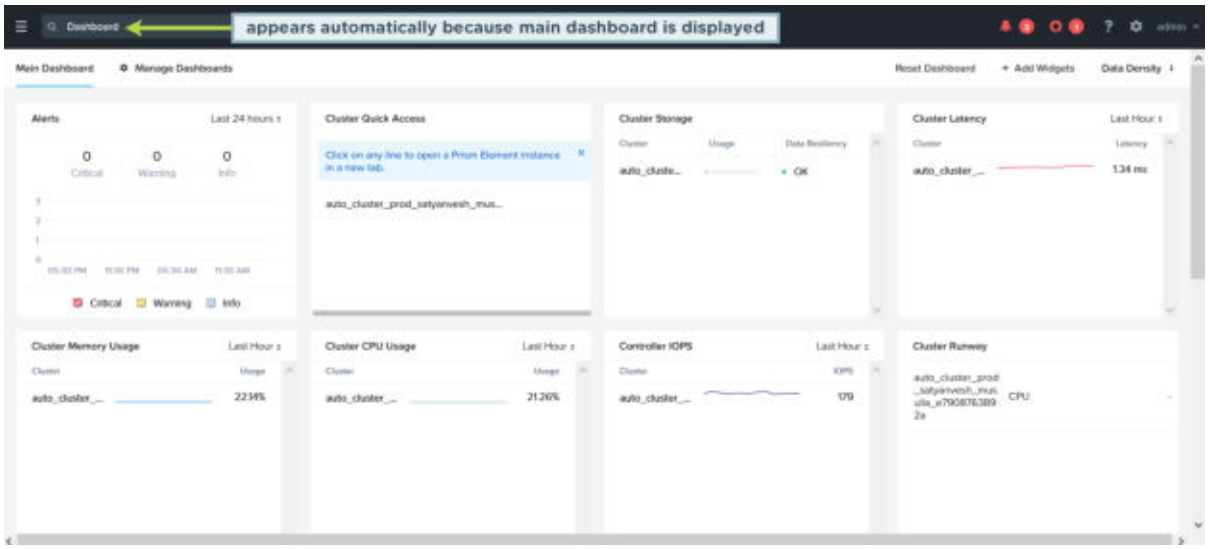

Figure 4: Search Field (initial screen)

As you go to different Prism Central screens, the search string changes automatically to match the current location, as illustrated in the following example from the clusters page.

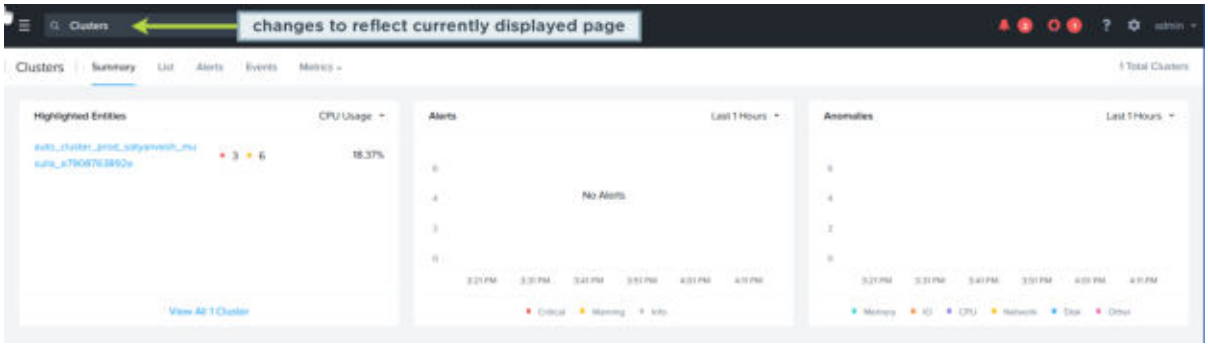

Figure 5: Search Field (from a different screen)

When you enter a string in the search field, a drop-down list appears with relevant matches across Prism Central. (Search strings are case insensitive.) For example, entering the string "VM" displays the following results:

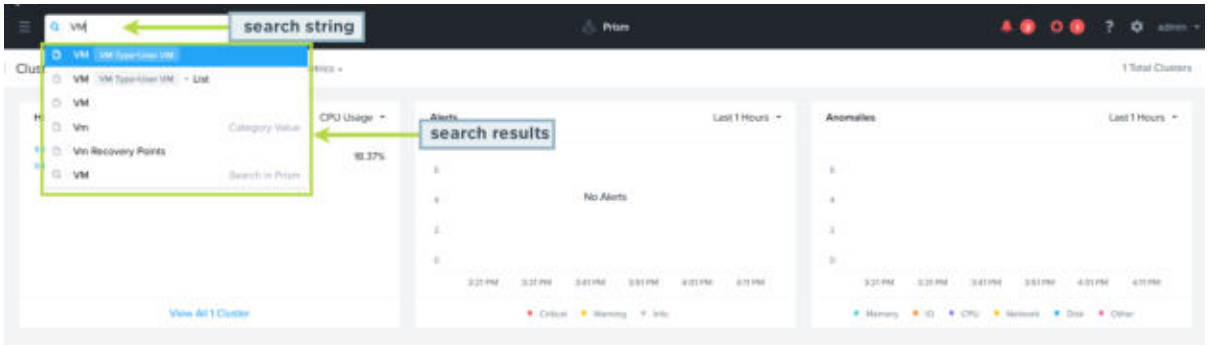

Figure 6: Search Results (for an entity type)

- "Category Value". Click this option to display the Category page with any VM-related entries.
- "VM Type=User VM". Click this option to display the VMs summary page filtered for user VMs.
- [blank]. Click this option to display the VMs summary page (unfiltered).
- "VM Type=User VM > List". Click this option to display the List tab of the user VMs page.
- "VM Type=User VM > Alerts". Click this option to display the Alerts tab of the user VMs page.
- "Search in Prism". Click this option to search across Prism Central for any information about the target entity. In the following example, two tables appear in the results, a list of the top VMs and a list of the top VM-related alerts. The top VMs list includes a link to the full list of VMs, and the top alerts list includes a link to the full list of alerts (which in this case is the same). Clicking the VM or alert name takes you to the details page for that VM or alert.

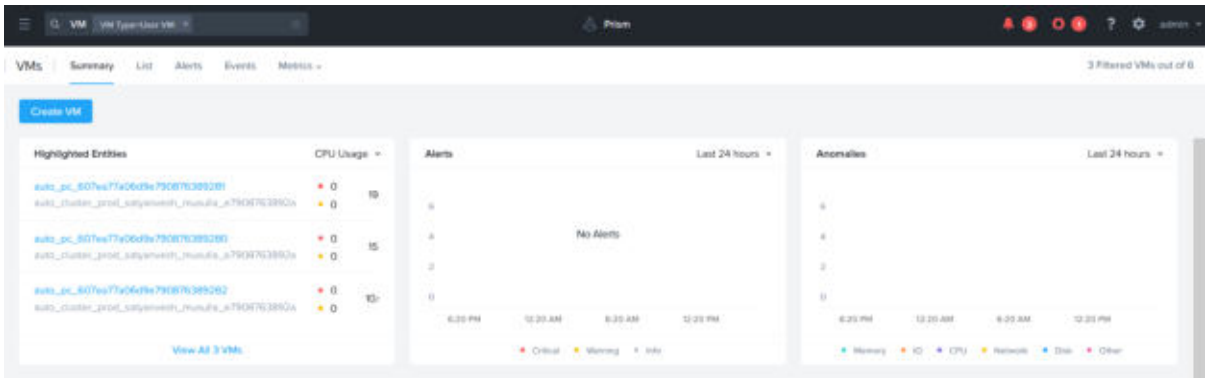

Figure 7: "Search in Prism" Results

Because search is context-sensitive, you can do a quick search from your current location without entering an explicit string. For example, if you are on the user VMs summary page, clicking in the search field (or clicking the forward slash [/] character) displays the search results for user VMs. In the following example, clicking the List, Alerts, or Events entry is the same as clicking those tabs on the page. Clicking "Memory Swap" displays the Memory Swap metrics page, and "Search in Prism" displays a results page for user VMs.

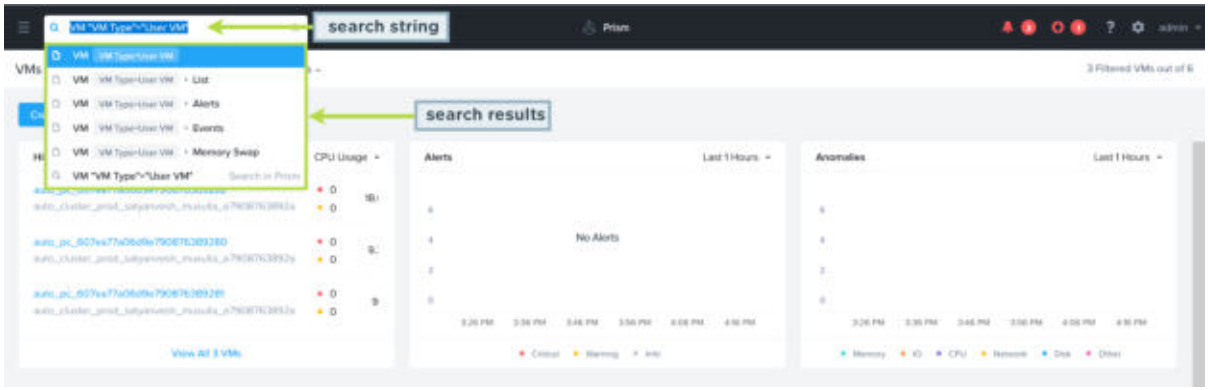

Figure 8: Search Results (for the current location)

Search provides intelligent responses based on whatever you enter in the field. For example, entering "scale out" returns a link to the Manage Prism Central page from which you can scale out Prism Central. If you enter a more generic string such as "version", the search engine returns results with links to pages across entities with version information.

Figure 9: Search Results (for generic string)

You can bookmark a filtered page (see Entity Exploring on page 81) by clicking the star icon in the search field. This adds a bookmark for that page (including any applied filters) to the entities menu (see Entities Menu on page 13).

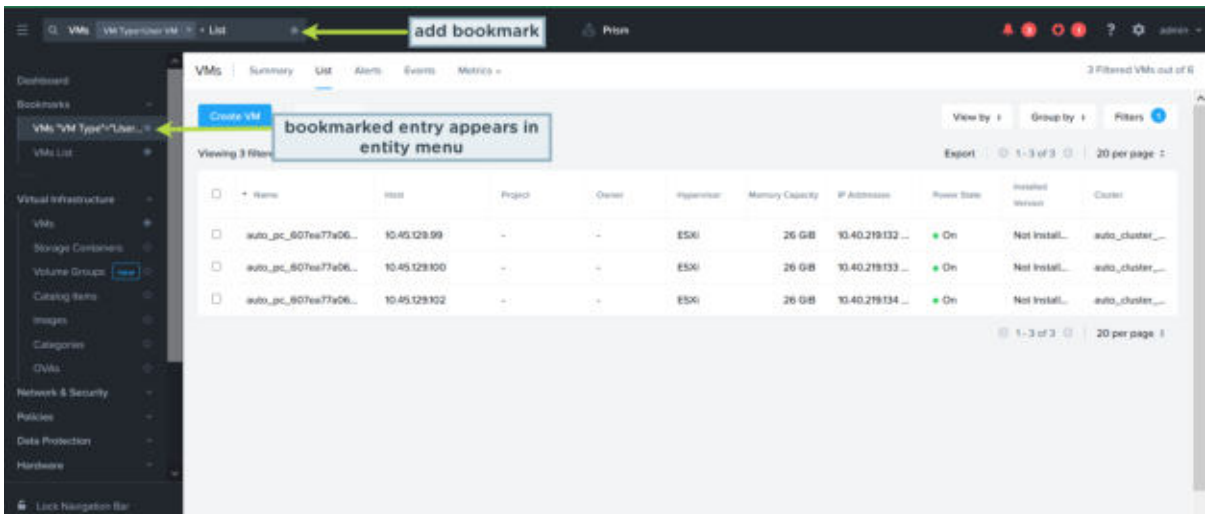

Figure 10: Bookmark Entry

On some pages, applied filters are not reflected in the main menu search field. For those pages, a local filters field appears. When you select one or more filters on such a page, the applied filters appear in the local filters field. You can do the following from this field:

- Remove a filter by clicking the X for that filter.
- Remove all filters by clicking **Clear** (on the right).
- Save the filter list by clicking the star icon. You can save a maximum of 20 filter lists per entity type.
- Use a saved filter list by selecting from the pull-down list.

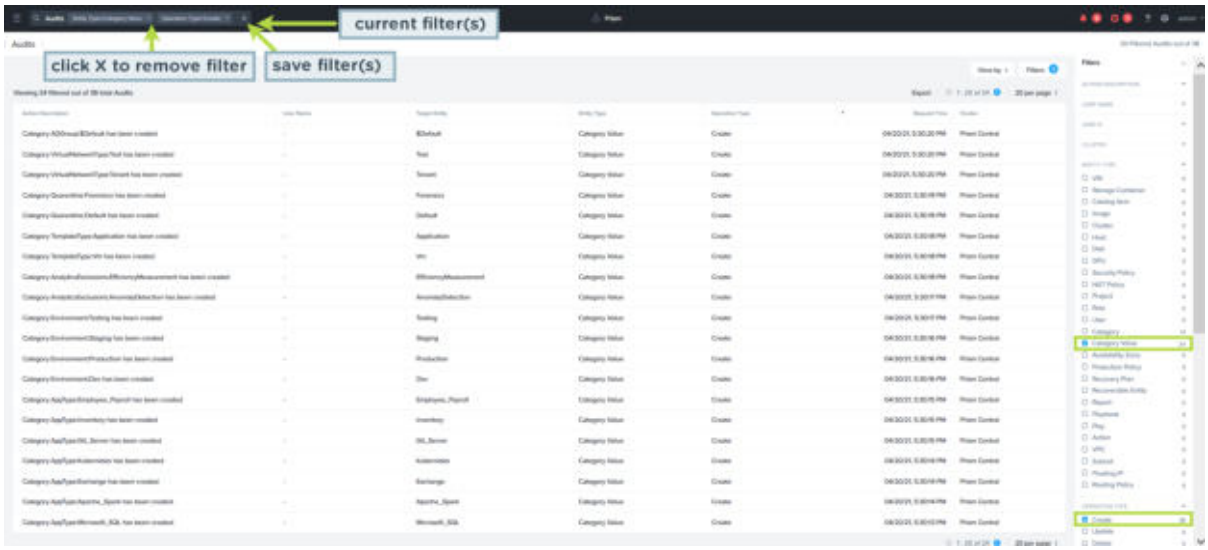

Figure 11: Local Filters Field

Query Syntax

The search field supports a range of simple and more complex query options in the following syntax forms:

- [entity|metric]: Enter an entity or metric type such as "`cluster`", "`vm`", or "`alerts`" to return information about the instances of that entity or metric across the registered clusters.
- <entity> <metric>: To refine an entity search, add a metric to the query. For example, "`vm iops`" returns IOPS information for the VMs.
- <entity> <attribute> <metric>: To further refine the search, add an attribute for the metric. For example, "`node failure alerts`" returns any failure alerts about the nodes.
- [metric|attribute] operator <value>: Enter an expression for a metric or attribute to return instances that match the expression. For example, "`"block model"=1050`" returns information about the NX-1050 nodes. See the following "Filter Expressions" section for the supported expression operators.
- <complex expression>: Combine the syntax forms to create complex queries. For example, "`clusters hypervisor = AHV "cpu usage" < 30`" returns information about clusters running the AHV hypervisor with CPU usage below 30%. Complex expressions have an implied AND so only instances that satisfy all conditions are returned.

- <action expression>: In addition to search queries, you can initiate certain actions from the query bar. For example, "`<vm_name> launch console`" launches a console window for the specified VM (if you are allowed to do so).

The following table describes the syntax rules for search queries.

Table 5: Syntax Rules

Rule	Example
RULES THAT GO TO WELL DEFINED PAGES	
entity type	vms
entity type + metric perspective (io, cpu, memory)	vms io
entity type + alerts	vm alerts
entity type + alerts + alert filters	vm alerts severity=critical
entity type + events	vm events
entity type + events + event filters	vm events classification=anomaly
entity type + filters (both metric and attribute)	vm “power state”=on
entity type + filters + metric perspective (io, cpu, memory)	vm “power state”=on io
entity type + filters + alerts	vm “power state”=on alerts
entity type + filters + alerts + (alert filters)	vm “power state”=on alerts severity=critical
entity type + filters + events	vm “power state”=on events
entity type + filters + events + event filters	vm “power state”=on events classification=anomaly
entity instance (name, ip address, disk serial, ...)	vm1, 10.1.3.4, BHTXSPWRM
entity instance + metric perspective (io, cpu, memory)	vm1 io
entity instance + alerts	vm1 alerts
entity instance + alerts + alert filters	vm1 alerts severity=critical
entity instance + events	vm1 events
entity instance + events + event filters	vm1 events classification=anomaly
entity instance + pages	vm1 nics, c1 capacity
parent instance + entity type	c1 vms
alert title search	disk bad alerts
page name search	analysis, tasks
RULES THAT ONLY GO TO SEARCH RESULT PAGES (exclusively)	

Rule	Example
entity type + metric	vm iops
entity type + attribute	vm power state
entity instance + metric	vm1 iops
entity instance + attribute	vm1 power state
entity type + filters + metric	vm “power state”=on iops
entity type + filters + attribute	vm “power state”=on hypervisor
help text search	upgrade cluster

Keywords

The following table lists the keywords or phrases you can use when formulating a query.

- The Object column identifies the type of object.
  - Entities: Lists the entities you can specify. The list is limited. For example, you cannot search for information about remote sites or network switches.
  - Fields: Lists the parameters (fields) you can specify.
  - Alerts: Lists the alert conditions you can specify.
  - Events: Lists the event conditions you can specify.
  - Actions: Lists the actions you can specify.
- The Entity column specifies the entities for which you can use that keyword in a query. For example, Alert queries can apply to any entity, but Fields has multiple rows broken down by entity type (such as cluster, VM, and disk) because there are select keywords that apply to each entity type.
- The Keywords column is a comma-separated list of the keywords or phrases you can use in a query for the specified object/entity type.

Table 6: Search Keywords

Object	Entity	Keywords
Entities	(n/a)	vm, cluster, node, container, disk
Fields	(not specified)	cpu usage, memory usage, disk usage, free physical storage, storage logical usage, saving ratio, savings, iops, read iops, write iops, io bandwidth, read io bandwidth, write io bandwidth, io latency, read io latency, write io latency, memory capacity, hypervisor
	Cluster	cluster name, ip address, version, number of hosts, cpu count, memory capacity, runway, storage runway, cpu runway, memory runway

Object	Entity	Keywords
	VM	vm name, ip address, host ip, virtual cpus count, power state, reserved memory, os, virtual hardware version
	Disk	serial, tier, mode, iops
	Container	container name, rf, on disk deduplication, perf-tier deduplication, compression, erasure coding
	Node	host name, ip address, service vm, hypervisor name, cpu capacity, cpu model, cpu sockets count, cpu cores count, cpu thread count, serial number, block model
Alerts	(any)	alert, alert title • Severity levels: critical, warning, info • Categories (impact types): capacity, performance, configuration, availability, system indicator • Dispositions types: resolved, unresolved, acknowledged, unacknowledged
Events	(any)	event type, title
Actions	VM	clone, migrate, delete, power on, power off, suspend, create vm, launch console, create network config, resume, snapshot, update, configure vm host affinity
	Cluster	launch prism element, unregister

Filter Expressions

You can use any of the following operators in an expression.

Table 7: Query Operators

Description	Operator
contains	~
does not contain	!~
starts with (“starts with” x)	=x*
ends with (“ends with” x)	=*x
equal to	=
not equal to	!=
greater than	>
greater than or equal to	>=

Description	Operator
less than	<
less than or equal to	<=
a metric range, for example CPU usage between 10 and 30	=[10 to 30]

The following implicit operator rules apply to expressions:

- In most cases, an AND is applied to all filters in an expression. For example, "hypervisor=AHV iops>100" means the hypervisor is AHV and IOPS are over 100.
- If the filters are on the same attribute, an OR is applied. For example, “hypervisor=AHV hypervisor=ESXi” means the hyperviosr is either AHV or ESXi.
- If the filters are on the same metric, an AND is applied to create a range. For example, "iops>100 iops<500" means IOPS in the range 100-500.
- If multiple range filters are defined in a single query, an OR is applied across the ranges. For example, "iops>100 iops<500 cpu>20 cpu<40" means IOPS in the 100-500 range or CPU in the 20-40 range.

Search Units

The data for search queries comes from the following parameters.

Table 8: Data Parameters

Data Type	Unit	Parameter
bytes usage related properties	GiBiBytes	memory_capacity_bytes, storage.usage_bytes, storage.free_bytes, storage.logical_usage_bytes, data_reduction.saved_bytes, storage.usage_bytes, storage.user_container_own_usage_bytes, data_reduction.saved_bytes, memory_size_bytes
runway properties	days	capacity.runway, capacity.storage_runway, capacity.cpu_runway, capacity.memory_runway
disk capacity related properties	TebiBytes	storage.capacity_bytes, storage.user_capacity_bytes
CPU/memory usage properties	percent	hypervisor_cpu_usage_ppm, hypervisor_memory_usage

Example Queries

Here are examples of various query types.

- Entity queries:

```
<cluster_name>
<cluster_ip_address>
```

```
<disk_serial#>
VMs "Power State"=On List
VMs Hypervisor=AHV List
powered on vms "memory capacity" > 32
```

- Performance queries:

```
clusters running out of cpu
clusters hypervisor = AHV "cpu usage" < 30
vm iops
```

- Alert queries:

```
node failure alerts
<cluster_name> alerts
<cluster_name> critical availability alerts
<alert title>
Alerts "Create Time"="08-Nov-2018 9:46 AM to 08-Nov-2018 10:46 AM" Severity=Critical
```

- Action queries:

```
<vm_name> launch console
<cluster_name> launch prism element
create vm
```

- Exploration queries:

```
clusters hypervisor=AHV
vm os=Linux
<cluster_name> vms
"block model"=1050
"cpu model"=Intel
containers Rf > 2
clusters version=4.6.2
hosts iops < 1000
<cluster_name> powered off vms
disks tier=ssd
vms "cpu usage"
vms “power state”
```

Understanding Displayed Statistics

Various statistics are displayed in the web console and Prism Central screens. There are three possible sources for a statistic:

> Note: Most displayed statistics appear in 30 second intervals. The values in the tables represent the most recent data point within the last 30 seconds. Prism Central collects the statistical data from each registered cluster, so the process of collecting that data could result in a longer lag time for some statistics displayed in Prism Central.

1. Hypervisor. When the hypervisor can provide usage statistics, those figures appear where appropriate. ESXi provides such statistics, but Hyper-V and AHV do not. Getting the statistics from ESXi means numbers displayed in the Prism display screens (both Prism Central and Prism Element) should match the corresponding ones in vCenter.

2. Controller (Stargate). When hypervisor statistics are unavailable or inappropriate, the Controller VM provides statistics from Stargate. Controller-reported statistics might differ from those reported by the hypervisor for the following reasons:

   - An NFS client might break up large I/O requests into smaller I/O units before issuing them to the NFS server, thus increasing the number of operations reported by the controller.
   - The hypervisor might read I/O operations from the cache in the hypervisor, operations not counted by the controller.
3. Disk (Stargate). Stargate can provide statistics from both the controller and disk perspective. The difference is that the controller perspective includes read I/O operations from memory and disk I/O operations, while the disk perspective includes just the disk I/O operations.

> Note: The distinction between hypervisor, controller, and disk statistics applies only to storage-related statistics such as IOPS, latency, and bandwidth.

Field labels in the Prism display screens help identify the information source:

- A field name that does not include either "Controller" or "Disk" indicates that statistic derives from the hypervisor (for example "IOPS").
- A field name that includes the word "Controller" indicates that statistic derives from the controller (for example "Controller IOPS").
- A field name that includes the word "Disk" indicates that statistic derives from the disk (for example "Disk IOPS").

The following table identifies the information source for various statistics. Overview, VM, and storage statistics come from either the hypervisor or controller. IFor VM statistics in a mixed ESXi/AHV cluster, the source depends on the hypervisor hosting that VM (hypervisor for ESXi-based VMs and controller for AHV-based VMs). Hardware statistics come from the disk. Metrics in the analysis page can come from any of the sources (hypervisor, controller, or disk) depending on the type of metric.

Table 9: Source for Displayed Statistics

Hypervisor Type	Overview, VM, and Storage	Hardware	Analysis
ESXi	hypervisor (controller for some storage stats)	disk	metric dependent
Hyper-V	controller	disk	metric dependent
AHV	controller	disk	metric dependent
Citrix Hypervisor	controller	disk	metric dependent
Mixed (ESXi + AHV)	hypervisor	disk	metric dependent
Prism Central	cluster dependent (hypervisor or controller)	disk	metric dependent

2

PRISM CENTRAL ADMINISTRATION

Administering Prism Central includes the following:

- Installing Prism Central and then upgrading Prism Central and registered clusters (see Installing or Upgrading on page 30)
- Managing licenses (see Licensing on page 30)
- Accessing Prism Central, managing the Prism Central instance (VM or set of VMs), and registering clusters with Prism Central (see Prism Central Management on page 30)

Installing or Upgrading

A Prism Central instance consists of either a single VM or a set of VMs that functions as a multi-cluster monitoring and control interface. You can install a Prism Central instance in either a Nutanix cluster or a non-Nutanix environment, although there are some limitations when you choose a non-Nutanix environment.

- To install or upgrade a Prism Central instance, see the "Prism Central Upgrade and Installation" topics in the *Acropolis Upgrade Guide*.

  > Note: Do not add any additional vDisks or NICs to a Prism Central VM above what is specified in the installation or upgrade instructions.

- To scale out an existing single VM Prism Central to multiple VMs, see Expanding (Scale Out) Prism Central on page 53.
- To upgrade NCC through Prism Central, see the "Upgrading NCC on Prism Central" topic the *Acropolis Upgrade Guide*.
- To upgrade one or more managed clusters through Prism Central, see the "Prism Central-Managed Clusters Upgrade" topic the *Acropolis Upgrade Guide*.
- If you use a proxy server, see Configuring Prism Central When a Cluster Uses Proxy Servers.

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the Nutanix License Manager Guide for information about how to license your cluster, apply add-on licenses, and manage your licenses in a variety of situations.

Prism Central Management

After installing Prism Central (see Installing or Upgrading on page 30), you can

- Log in and out through any supported browser (see Logging Into Prism Central on page 31 and Logging Out of Prism Central on page 34).

- Monitor the status of this Prism Central instance (see Managing Prism Central on page 34).
- Check version information (see Finding the Prism Central Version on page 43).
- Modify login page and automatic log off settings (see Modifying UI Settings on page 526).
- Register clusters with this Prism Central instance (see Register (Unregister) Cluster with Prism Central on page 46).
- Expand Prism Central from a single VM to a multi-VM instance (see Expanding (Scale Out) Prism Central on page 53).
- Shutdown or Start Up a Prism Central VM (see Shutdown (Start Up) Prism Central VM on page 56).
- Protect Prism Central to upto three AHV or ESXi clusters registered to it (see *Protecting Prism Central* in the *Prism Central Supplement* for pc.2021.7).
- Recover Prism Central from Prism Element running on the registered AHV or ESXi cluster (see *Recovering Prism Central (1-Click Disaster Recovery)* in the *Prism Central Supplement* for pc.2021.7).

Logging Into Prism Central

About this task

To log in to Prism Central, do the following. Knowledge base article KB 1661 lists default cluster credentials.

Procedure

1. Open a web browser, enter `http://`*`management_ip_addr`* in the address field, and press **Enter**. Replace *`management_ip_addr`* with the Prism Central VM IP address.

   > Note: The supported browsers are the current version and two major versions back of Firefox, Chrome, and Safari, plus Internet Explorer version 11 and Microsoft Edge.

   The browser redirects to the encrypted port (9440) and may display an SSL certificate warning. Acknowledge the warning and proceed to the site. If user authentication is enabled and the browser does not have the correct certificate, a denied access message may appear. For the complete list of required ports, see Port Reference.

2. If a welcome screen appears, read the message and then click the "Accept terms and conditions" bar at the bottom.

3. In the login screen, do one of the following: enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

   » Enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

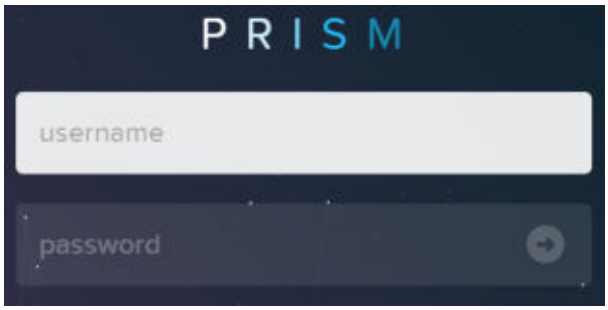

Figure 12: Login Screen

Note: If LDAP authentication is used, enter the user name in ***username@domain*** format; the ***domain\username*** format is not supported. The user principle name [UPN] attribute is used to find the user account in Active Directory, so your user name must be in that format.

   » If SAML authentication is enabled (see “” ()), enter your user name according to the configured identify provider in the indicated field. (This field appears on the login page only when SAML authentication is enabled.)

Figure 13: Login Screen (SAML)

Note: The login page includes background animation that is enabled by default. Click the **Freeze space time continuum!** link at the bottom right of the login screen to disable the animation (or the **Engage the warp drive!** link to enable the animation). To permanently disable (or enable) the animation, see Modifying UI Settings on page 526).

4. If you are logging in as an administrator for the first time (user name **admin** and default password **Nutanix/4u**), which requires that the default password be changed, enter a new

password in the **password** and **re-type password** fields and then press **Enter** or click the right arrow icon.

Figure 14: Login Screen (first admin login)

The password must meet the following complexity requirements:

- At least 8 characters long
- At least 1 lowercase letter
- At least 1 uppercase letter
- At least 1 number
- At least 1 special character
- At least 4 characters difference from the old password
- Should not be among the last 10 passwords

After you have successfully changed the password, the new password is synchronized across all Controller VMs and interfaces (Prism web console, nCLI, and SSH).

Note:

- The default password expiration age for the `admin` user is 60 days. You can configure the minimum and maximum password expiration days based on your security requirement.
  - `nutanix@cvm$ sudo chage -M <MAX-DAYS> admin`
  - `nutanix@cvm$ sudo chage -m <MIN-DAYS> admin`
- After you upgrade from an AOS earlier version and then attempt to log in to the Prism web console as the admin user, you are prompted to create a new admin user password.
- When you change the admin user password, update any applications and scripts using the admin user credentials for authentication. Nutanix recommends that you create a user assigned with the admin role instead of using the admin user for authentication. The *Prism Web Console Guide* describes authentication and roles.

5. If a license agreement screen appears (typically on the first login or if the EULA changed since the last login), which indicates the current EULA has not been acknowledged yet, do the following:

   a. Read the license agreement (on the left).

   b. Enter appropriate information in the **Name**, **Company**, and **Job Title** fields (on the right).

   c. Check the "I have read and agree to the terms ..." box.

   d. Click the **Accept** button.

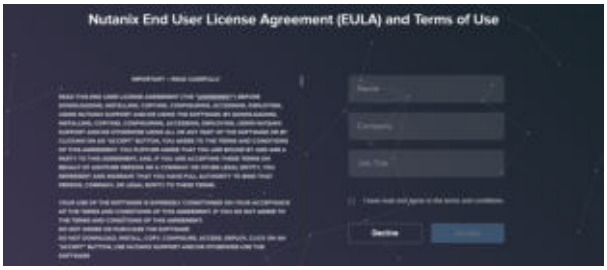

Figure 15: EULA Screen

Logging Out of Prism Central

Procedure

To log off from Prism Central, click the user icon in the main menu and then select the **Sign Out** option from the pull-down list. You are logged out immediately after selecting the option (no prompt or message).

Managing Prism Central

Prism Central consists of one or more VMs treated as a single instance that monitors and manages registered clusters. You can view information about the Prism Central VMs by clicking the gear icon in the main menu and then selecting **Prism Central Management** from the Settings menu (see Settings Menu on page 16). This displays the Manage Prism Central page.

> Note: This page is not available if Prism Central is hosted in a non-Nutanix environment, that is not hosted in a Nutanix cluster.

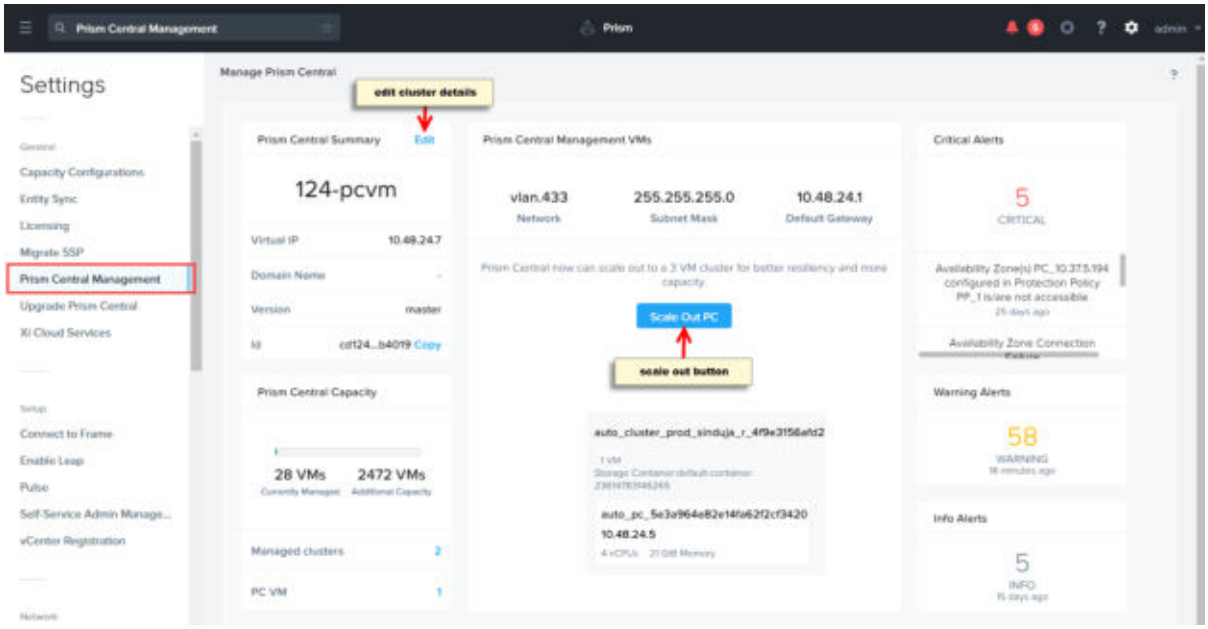

Figure 16: Manage Prism Central Page (1-VM instance)

The Manage Prism Central page contains the following:

- A "Prism Central Summary" pane on the upper left that displays fields for
  - Cluster name, which is "Unnamed" by default.
  - Domain name. A dash appears if a domain name is not specified.
  - Version. This is the Prism Central version number.
  - ID number. Click the copy link to copy the ID number so you can paste it elsewhere.
- A "Prism Central Capacity" pane on the lower left that displays fields for
  - Number of VMs monitored currently by this Prism Central instance.
  - Number of additional VMs (remaining capacity) this Prism Central instance can monitor.
  - Number of clusters registered to this Prism Central instance; click the number to display the clusters page (see Clusters Summary View on page 197).
  - Number of VMs that comprise this Prism Central instance; click the number to display the VMs page (see VMs Summary View on page 86).
- A "Prism Central VMs" pane in the middle that displays network address information (network name, subnet mask, and default gateway address) at the top and information below about each VM in the Prism Central instance (VM name and name of storage container in which it is located, IP address, number of vCPUs, and memory size). In addition, single VM instances include a **Scale Out PC** button to scale out this Prism Central instance (see Expanding (Scale Out) Prism Central on page 53).
- An "Alerts" pane on the right that displays a list of Prism Central-related alerts broken into Critical, Warning, and Info sections.

When the Prism Central instance consists of three VMs, the display changes slightly. The **Scale Out PC** button disappears, and a new **Add PCVM** button appears (top right). Prism Central instances are limited to a maximum of three VMs, but if you lose one of those VMs for any

reason, you can add a replacement by clicking the **Add PCVM** button. This displays the Add PC VMs page (similar to the **Scale Out PC** page) from which you can add the new VM.

> Note: Verify that the prerequisites in Expanding (Scale Out) Prism Central on page 53 are satisfied before adding a Prism Central VM.

To add or change the Prism Central domain name or virtual IP, click the **Edit** link in the "Prism Central Summary" pane. The Cluster Details window appears. Do the following:

1. In the **Cluster Name** field, enter a name for the cluster (default "Unnamed"). After naming the cluster, the field no longer appears when subsequently opening the window. Naming the cluster is optional.
2. In the **FQDN** field, enter the fully qualified domain name (FQDN) for the Prism Central cluster. For a multi-VM Prism Central, specifying a domain name allows balancing of requests across the Prism Central VMs to improve performance and resiliency. (This requires an administrator configure the domain name in the DNS server as the set of IP addresses associated with the Prism Central VMs.)

   When an FQDN is specified, it is the first choice for IP address resolution. Make sure the virtual IP (next field) is not mapped to multiple DNS entries as this can lead to resolution problems. If the FQDN fails to resolve in DNS, the existing Prism Central VM IP addresses are used.
3. In the **Virtual IP** field, enter an IP address that will be used as a virtual IP for the cluster. This is relevant if you have a multi-VM Prism Central.
4. When the fields are correct, click the **Update** button to save the changes and close the window.

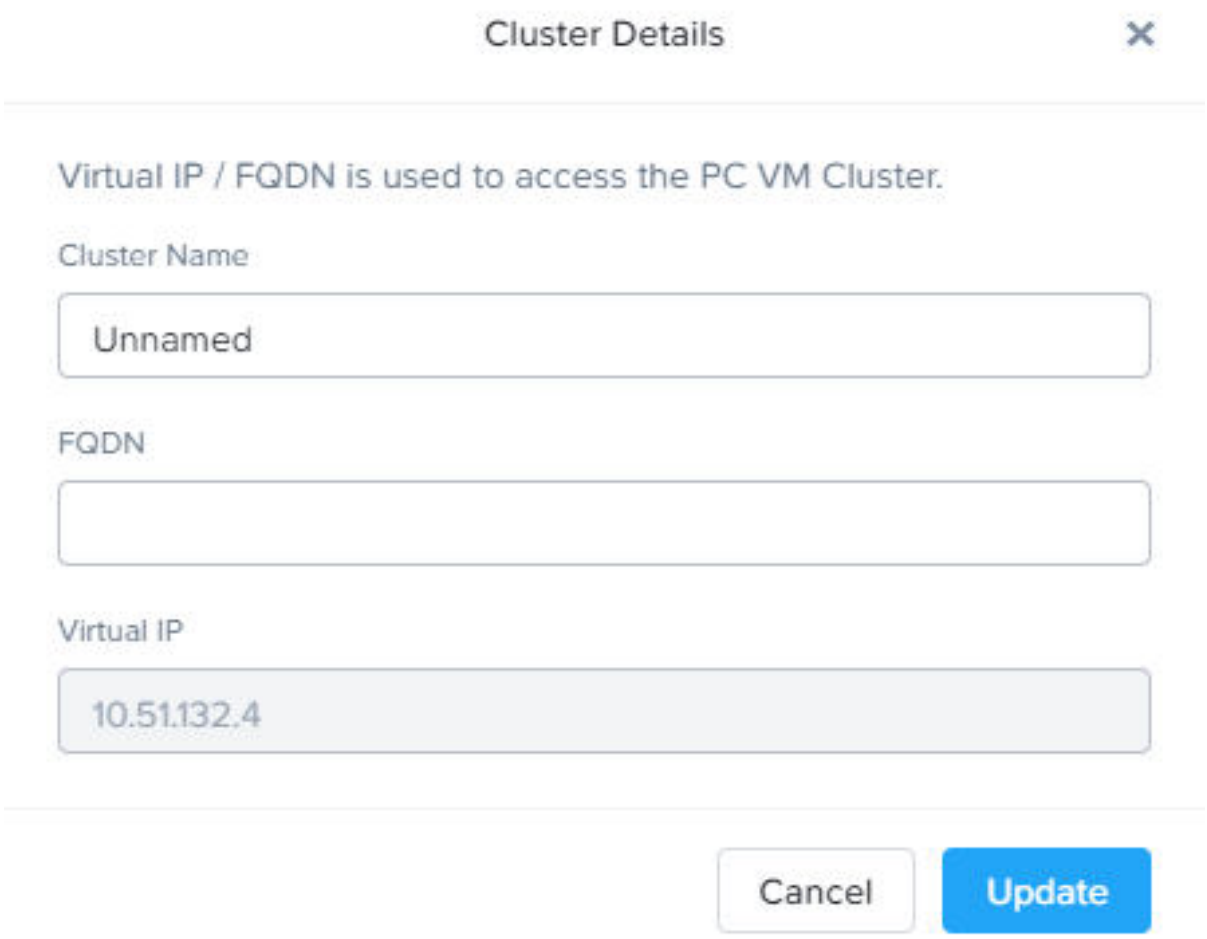

Figure 17: Cluster Details Window

Enabling Microservices Infrastructure

Some Prism Central services are packaged in containers as micro services, and the control plane for micro services platform (sometimes referred to as CMSP) enables the Microservices Infrastructure. The Microservices Infrastructure is disabled by default, but you can enable it at any time.

Before you begin

Address the following before enabling the Microservices Infrastructure:

- The Microservices Infrastructure is supported on clusters running AHV or ESXi only. For ESXi clusters you may need to enter your vCenter credentials (user name and password) and a network for deployment.
- The host cluster must be registered with this Prism Central instance (see Register (Unregister) Cluster with Prism Central on page 46).
- The host cluster virtual IP address and ISCSI data services IP address must be configured (see the "Modifying Cluster Details" section in the *Prism Web Console Guide*).
- Name servers and NTP servers must be configured in both the host cluster (Prism Element) and Prism Central.

- You may need to increase the Prism Central VM memory size before enabling the Microservices Infrastructure. The requirements are listed in KB 8932. If your Prism Central VM resources are below the recommended amount, manually increase the memory to the recommended amount.

About this task

To enable the Microservices Infrastructure, do the following:

CAUTION: Once enabled, the Microservices Infrastructure cannot be disabled. The following limitations apply after enabling the Microservices Infrastructure:

- Scaleout three-node PC VMs deployment: Small, large, or x-large PC VMs are supported. Reverting this deployment to a single PC VM deployment is not supported.
- Single PC VM deployment: Small, large, or x-large PC VM is supported. You can expand a single PCVM deployment to a scaleout three-node PC deployment.
- Changes to any of the following IP addresses are not supported: Prism Central virtual IP address, host cluster virtual IP address, host cluster iSCSI data services IP address
- Unregistering the host cluster from Prism Central is not supported. However, unregistering other clusters is still supported.
- Common Access Card (CAC) is not supported.
- Ensure that the IP addresses in subnet that you configure for microservices infrastructure do not conflict with the IP addresses in the management subnet.

  Note: Do not use the IP addresses in the subnet 10.100.0.0/24 for operational purposes like DNS.

Procedure

1. Go to the Manage Prism Central page (see Managing Prism Central on page 34).
2. In the **Prism Central on Microservices Infrastructure** box, click the **Enable Now** button.

Figure 18: Enable box (in Manage Prism Central page)

3. In the Enable Microservices Infrastructure window, do the following in the indicated fields:
   a. [first screen, no fields] Read the information and then click the **Continue** button.

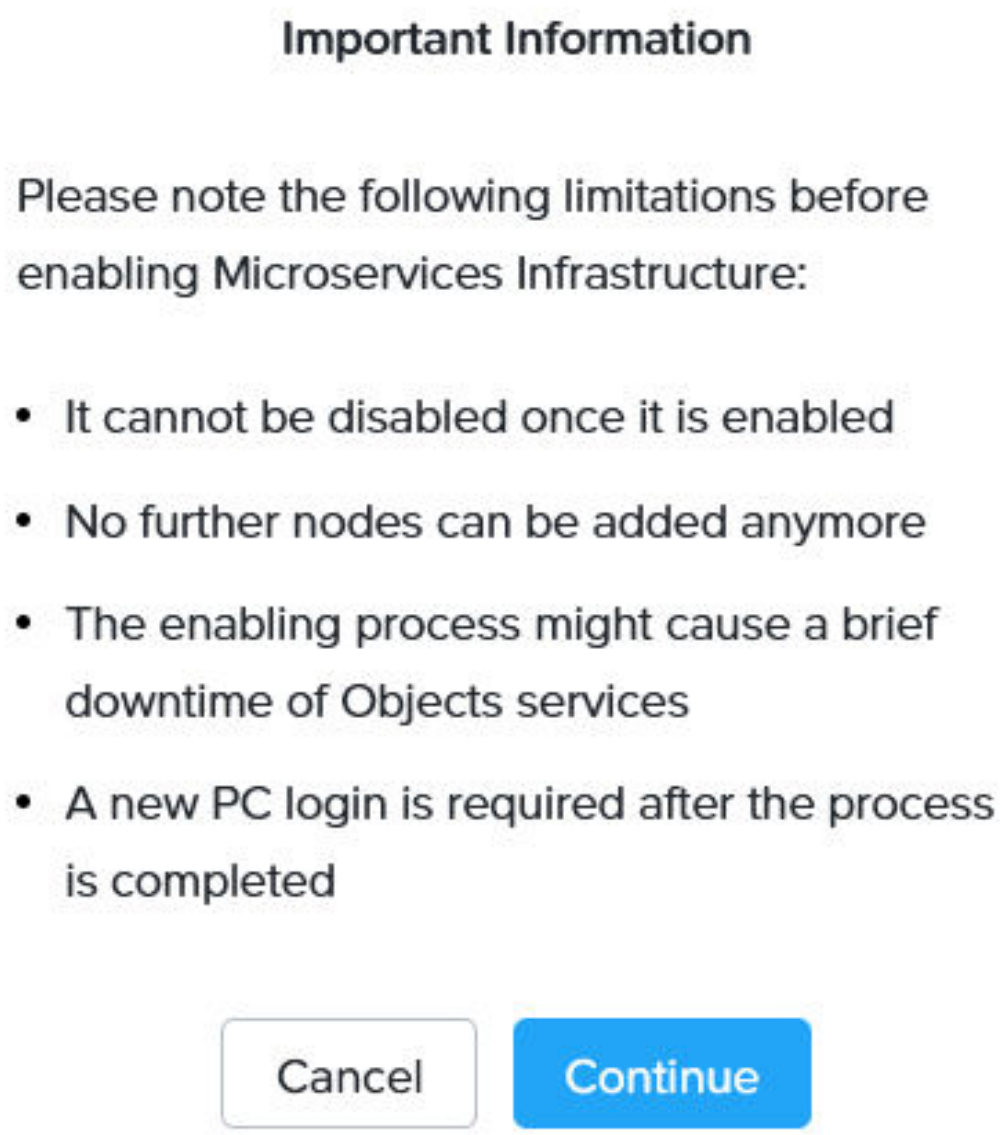

Figure 19:

The following sample window is displayed:

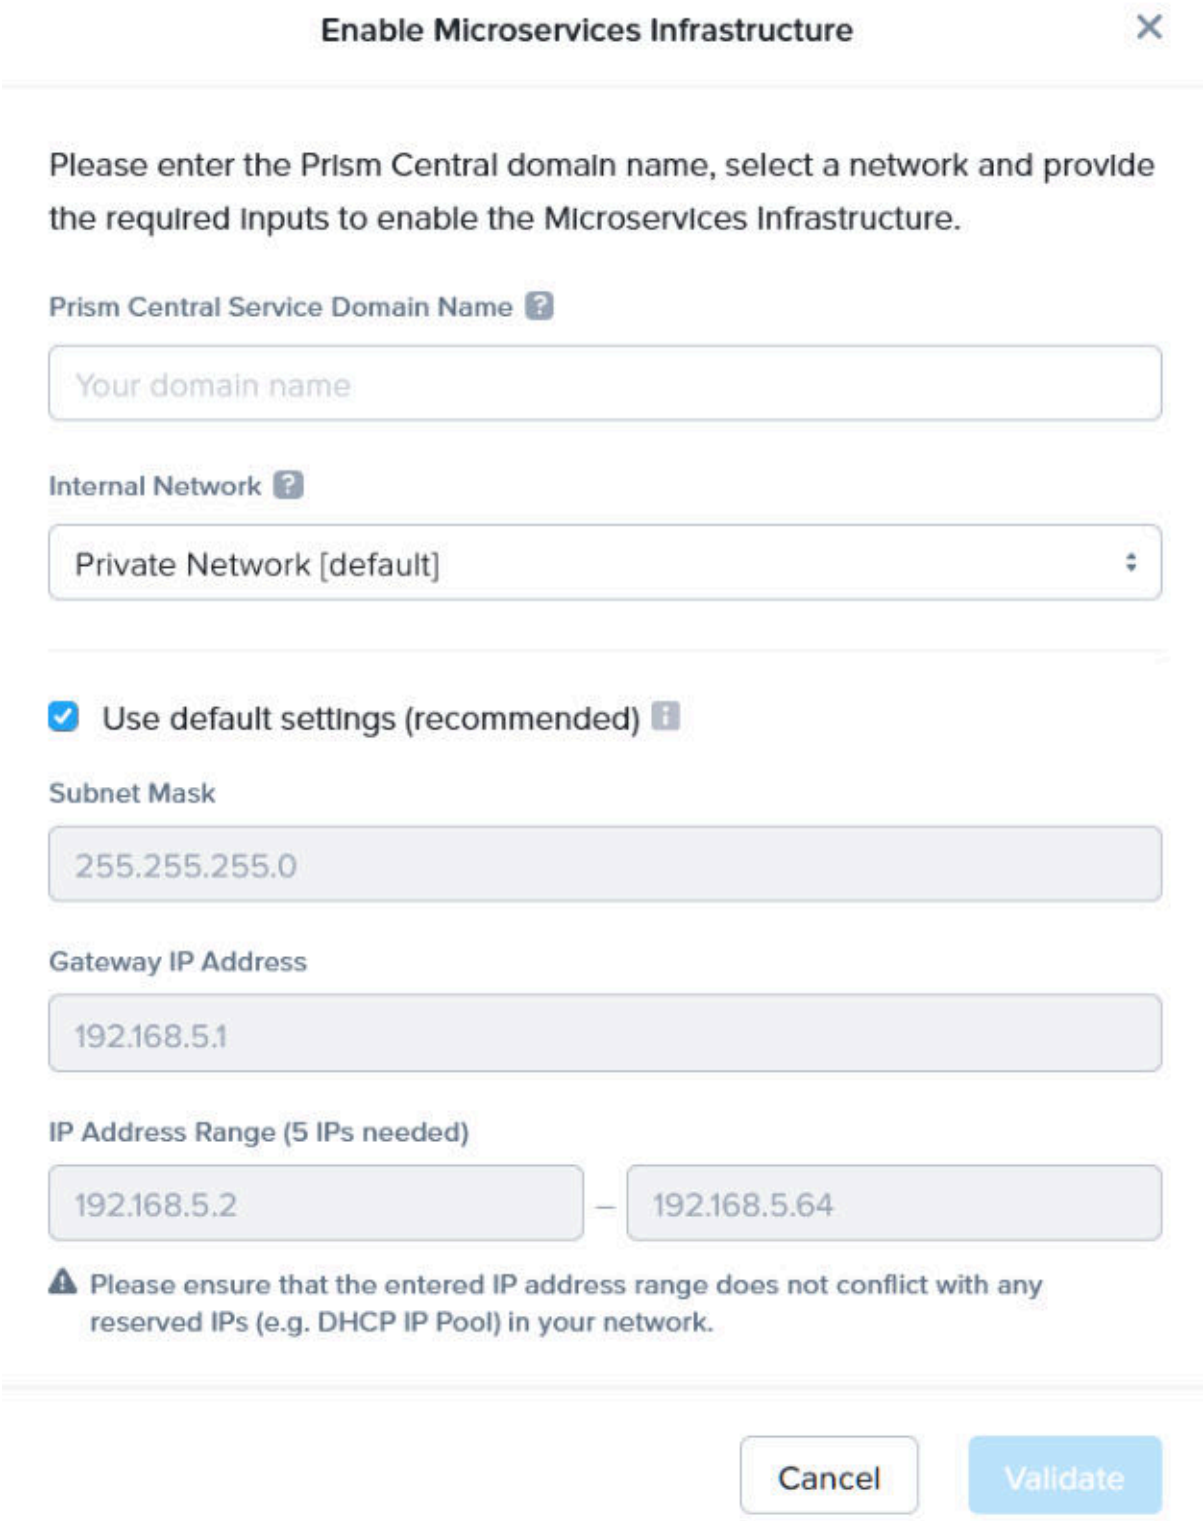

Figure 20:

b. **Prism Central Service Domain Name**: Enter a domain name for the service.

> Note: Ensure that the domain is not longer than 16 characters. Prism Central supports a maximum of 16 characters for domain name for enabling microservices infrastructure. If the length of the domain name exceeds 16 characters, enabling microservices infrastructure fails.

The domain name is used for exporting services under the Prism Central platform. The name must be unique and should be at least three level, for example test.nutanix.com. Do not use the name of an existing domain. (The parent domain's DNS server could redirect the DNS requests to the Prism Central service domain's DNS server to resolve the service names globally.) The IP address of the DNS server appears in the Prism Central management page.

> Note: When you configure microservices infrastructure, ensure that the DNS name you configure for CMSP does not end with test. Flow networking does not support test as a top level domain. For example, the following are valid domain configurations:
>
> - my.cluster.domain
> - my.test.cluster.test.domain
>
> However, the following are examples of domains that Flow networking does not support:
>
> - my.cluster.test
> - my.cluster.domain.test

c. **Internal Network**: Select the network to use for Prism Central micro services communication from the pull-down list.

   The default selection **Private Network [default]** is a pre-configured private VxLAN network. Instead, if you want microservices infrastructure to use a different network, you can select the network (managed or unmanaged) from the drop-down list. If the network you want microservices infrastructure to use does not appear in the list, you must first configure it (see Network Connections on page 723).

d. The **Use default settings (recommended)** checkbox is available only when you have retained the default selection (**Private Network [default]**) for **Internal Network** in the preceding step. This checkbox is checked by default. Do one of the following:

   » Retain the check mark for the **Use default settings (recommended)** checkbox and click **Validate**. (Go to step 5 on page 42.)

   Retaining the check mark for the **Use default settings (recommended)** checkbox allows Prism Central to use the **Private Network [default]** with the default values for **Subnet Mask**, **Gateway IP Address** and **IP Address Range**.

   » Clear the **Use default settings (recommended)** checkbox, if you want Prism Central to use the **Private Network [default]** setting with specific (non-default or custom) values for **Subnet Mask**, **Gateway IP Address** and **IP Address Range**.

You can configure the internal network for microservices infrastructure by entering the values for the **Subnet Mask**, **Gateway IP Address** and **IP Address Range** fields if you did one of the following:

- Selected a managed or unmanaged network other than **Private Network [default]** for **Internal Network** in step 3.c on page 41.

  If you selected a managed network in the step 3.c on page 41, the values in the **Subnet Mask**, **Gateway IP Address** and **IP Address Range** fields are already configured. If you selected an unmanaged network in the step 3.c on page 41, you must enter the necessary values in the respective fields.

- Cleared the **Use default settings (recommended)** checkbox with the **Private Network [default]** selection for **Internal Network** in step 3.d on page 41.

See the next step to configure the internal network for microservices infrastructure.

4. Enter the values for **Subnet Mask**, **Gateway IP Address** and **IP Address Range** fields.

   a. **Subnet Mask**: Enter the subnet mask.

   b. **Gateway IP Address**: Enter the IP address for the gateway.

   c. **IP Address Range**: Enter a range of IP addresses that the network can use.

      Enter a range of at least 10 available (unreserved) addresses for a 3-VM (scale out) Prism Central instance or at least 5 addresses for a 1-VM Prism Central instance. The addresses must be consecutive or sequential. For a managed network, the range of addresses for microservices infrastructure must be outside the range of reserved IP addresses (for example, DHCP IP Pool) in the selected network.

5. After you check that the values entered for all the fields are correct, click the **Validate** button.
   The system checks the settings. When the settings are confirmed, a `Validation was successful` message appears and the **Validate** button changes to an **Enable** button.

6. Click the **Enable** button.
   The enablement process begins, which takes some time (around 45 minutes). You can monitor progress from the tasks dashboard (see Tasks View on page 241). When the microservices infrastructure is enabled, the contents of the enable box in the Manage Prism Central page changes to display configuration information such as the domain name, the used IP addresses, and the service domain's DNS IP address.

What to do next

After the enablement tasks are completed, including the **IAM Migration & Bootstrap** task, log out of Prism Central. Wait at least 15 minutes before logging on to Prism Central.

Enabling Microservices Infrastructure (dark site)

Enabling the Microservices Infrastructure when there is no Internet access (dark site) requires some extra steps from the standard procedure.

About this task

To enable the Microservices Infrastructure at a dark site, do the following:

Procedure

1. From a device that has public Internet access, log on to the Nutanix support portal, select **Downloads** > **LCM** from the entity menu (hamburger icon), and then download the LCM Dark Site Bundle tar file.
2. Set up a local web server, upload `LCM Dark Site Bundle` to the server, and extract the files in a directory in the base of the web server.

   See the *Setting Up a Local Web Server* section in the *Life Cycle Manager Dark Site Guide* for more information. Set the following mime types in the LCM web server when you are using Windows Web Server: .json, .xz, .qcow2, .sign, .gz
3. In the Nutanix support portal **Downloads** page, select **Microservices Platform (MSP)**, download the **Upgrade (Darksite)** bundle to the web server base directory you used in step 2, and extract the files.
4. Now follow the standard procedure for enabling the Microservices Infrastructure through Prism Central (see Enabling Microservices Infrastructure on page 37).

Finding the Prism Central Version

About this task

You can check the version of your Prism Central instance using the Prism Central user interface or by running the following command.

Figure 21: Finding the Prism Central Version

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275778971001 on Brightcove.

```
nutanix@cvm$ ncli cluster info
```

Do the following to check the Prism Central version in the UI.

Procedure

1. Click the user icon in the main menu and then select the **About Nutanix** option from the pull-down list.

   An About Nutanix window appears that includes the AOS version number. It also includes a link to Nutanix patent information.

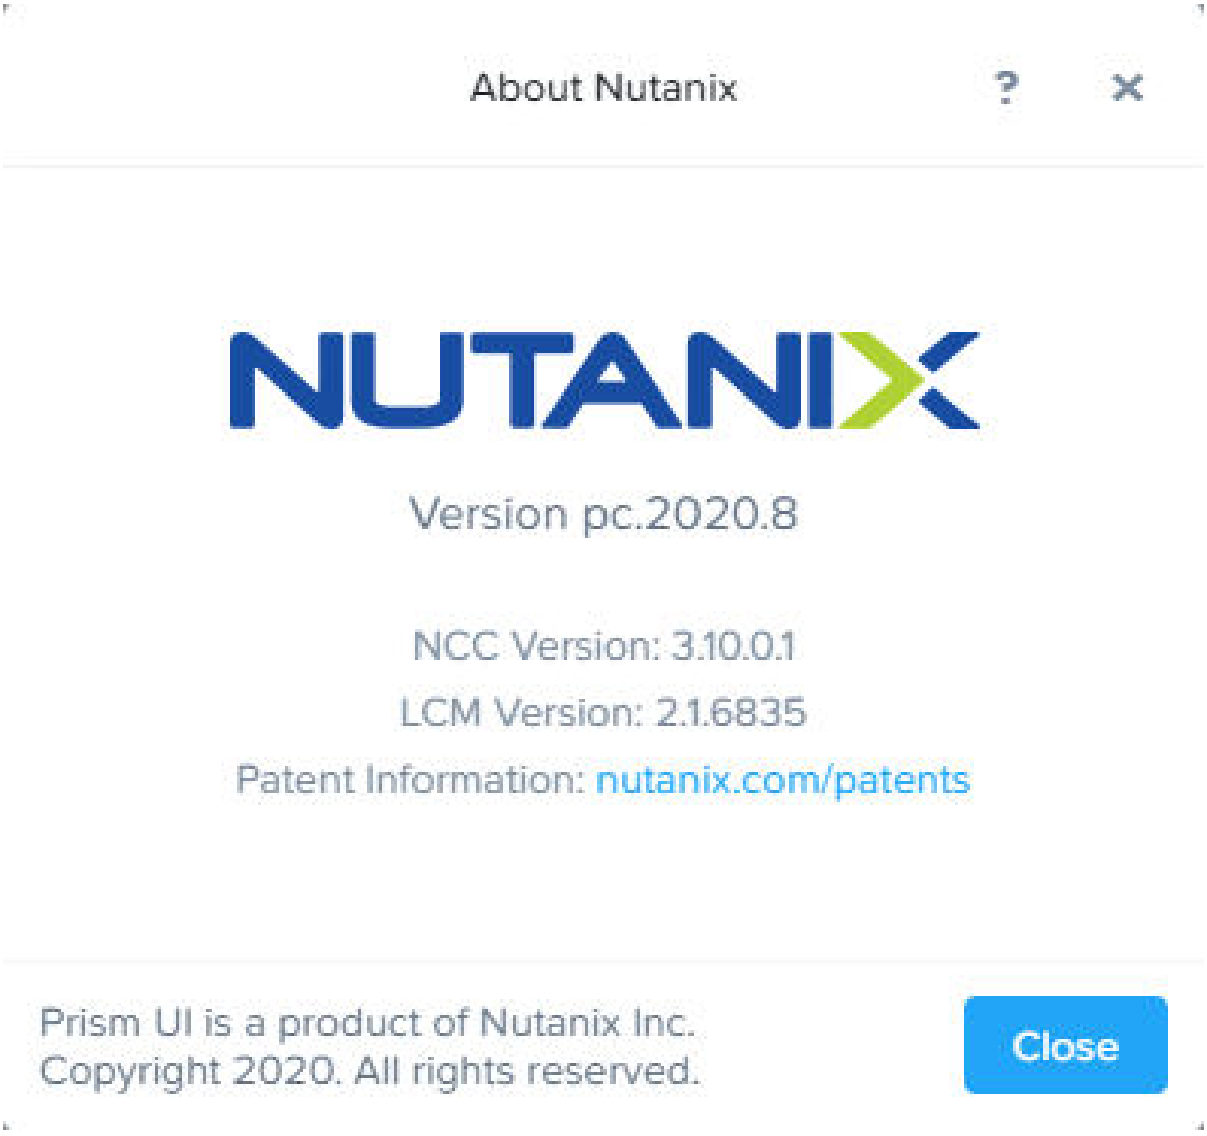

Figure 22: Example About Nutanix Window

2. Click **Close** to close the window.

Finding the AHV Version on Prism Central

You can see the installed AHV version in the Prism Central console.

About this task

To view the AHV version installed on any host in the clusters managed by the Prism Central, do the following.

Figure 23: Finding the AHV Version on Prism Central

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275777446001 on Brightcove.

Procedure

1. Log on to Prism Central.
2. In side bar, select **Hardware** > **Hosts** > **Summary** tab.
3. Click the host you want to see the hypervisor version for.
4. The Host detail view page displays the **Properties** widget that lists the `Hypervisor Version`.

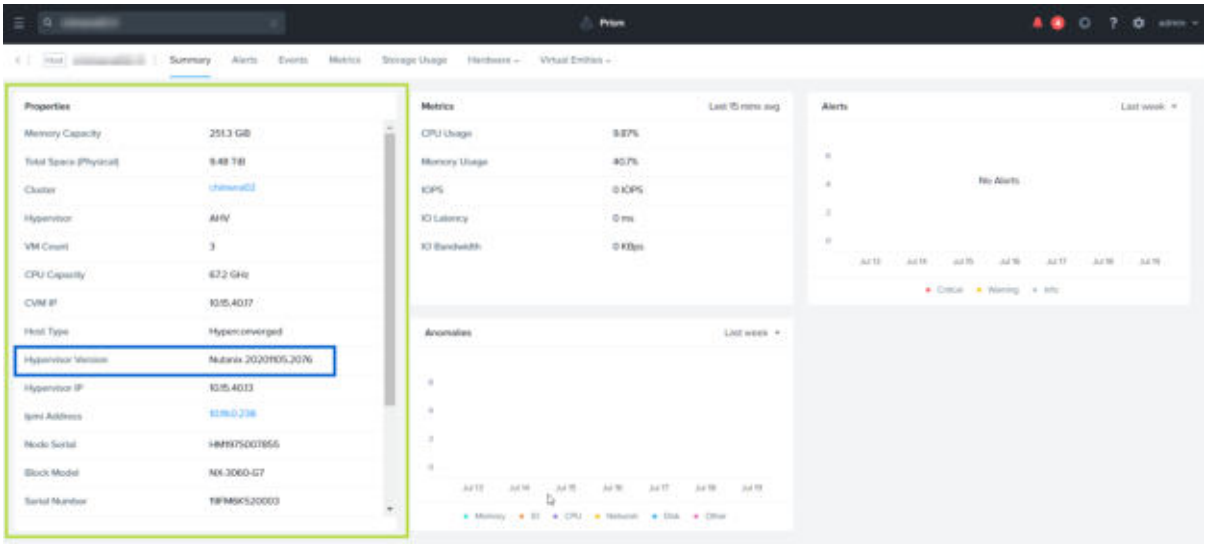

Figure 24: Hypervisor Version in Host Detail View

Finding the AOS Version Using Prism Central

To view the Nutanix AOS version running in the Prism Central, do the following:

Procedure

1. Log on to Prism Central.
2. Click the entities menu (hamburger menu) near the search field and then select **Hardware > Clusters** to view the cluster dashboard.

3. In the **List** tab (default), the AOS version of the cluster is listed against the cluster name.

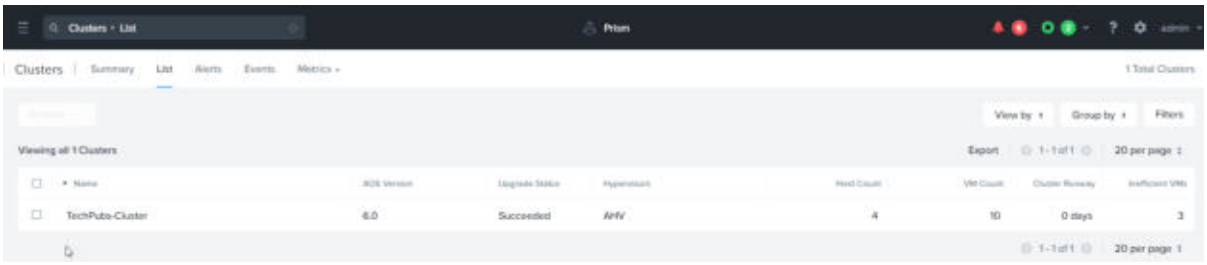

Figure 25: AOS Version in a Cluster

Register (Unregister) Cluster with Prism Central

Before you begin

- If you have never logged into Prism Central as the user **admin**, you need to log in and change the password before attempting to register a cluster with Prism Central.
- Do not enable client authentication in combination with ECDSA certificates on a registered cluster since it causes interference when communicating with Prism Central.
- Ports 9440 and 80 need to be open in both directions between the Prism Central VM and all the Controller VMs (and the cluster virtual IP address if configured) in each registered cluster. For the complete list of required ports, see Port Reference.
- A cluster can register with just one Prism Central instance at a time. To register with a different Prism Central instance, first unregister the cluster.

About this task

To register a cluster with Prism Central, do the following:

Procedure

1. Log in to the Prism web console on the target cluster as the user **admin**.
2. In the Health dashboard, from the **Actions** drop-down menu, select **Run Checks**. to run Nutanix Cluster Checks.

3. Do one of the following:

   » On the **Home** dashboard, click **Register or create new** from the **Prism Central** widget.

   » Click the gear icon and select **Prism Central Registration** from the **Settings** menu.

Figure 26: Home Dashboard

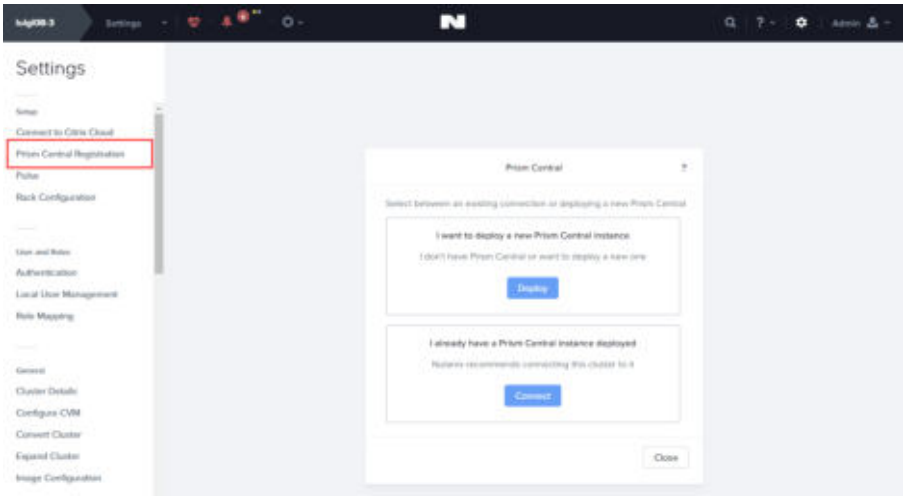

Figure 27: Settings Menu

4. In the first screen of the Prism Central dialog box, click the **Connect** button.

   This screen includes two options, **Deploy** and **Connect**. This procedure describes how to connect to an existing Prism Central instance. See *Installing Prism Central (1-Click Method)*

section in Acropolis Upgrade Guide for instructions on how to deploy a new Prism Central instance.

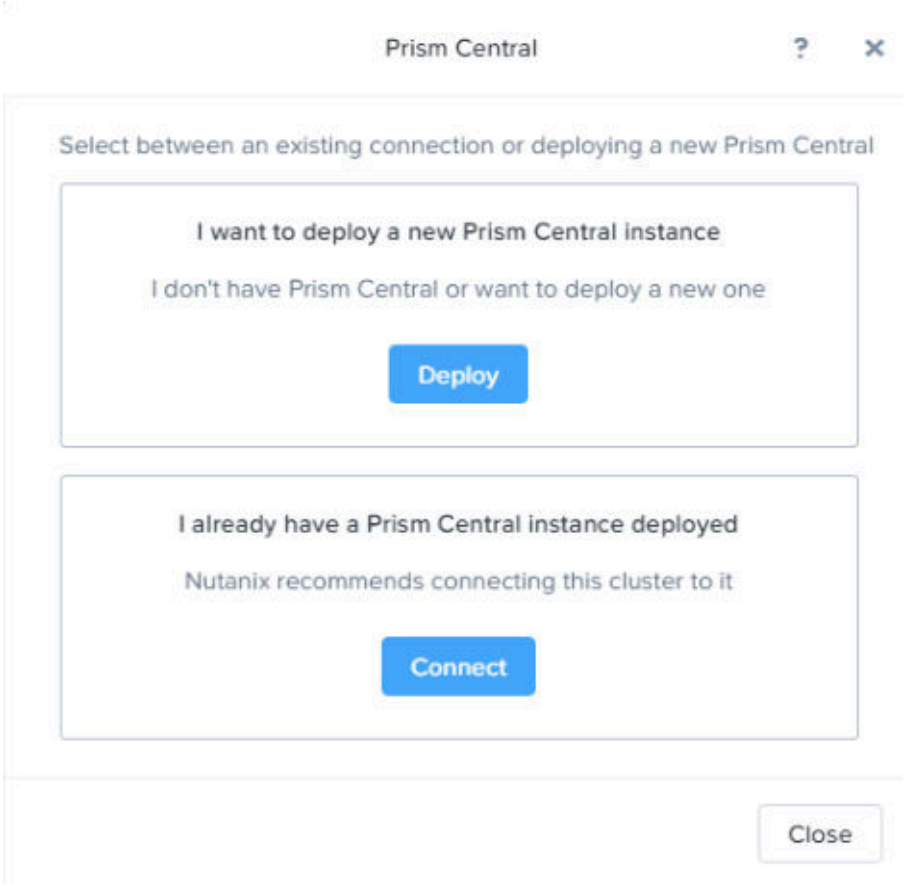

Figure 28: Prism Central Dialog Box (deploy/connect screen)

5. A screen appears outlining what services are available through Prism Element and Prism Central. After reviewing the message, click the **Next** button.

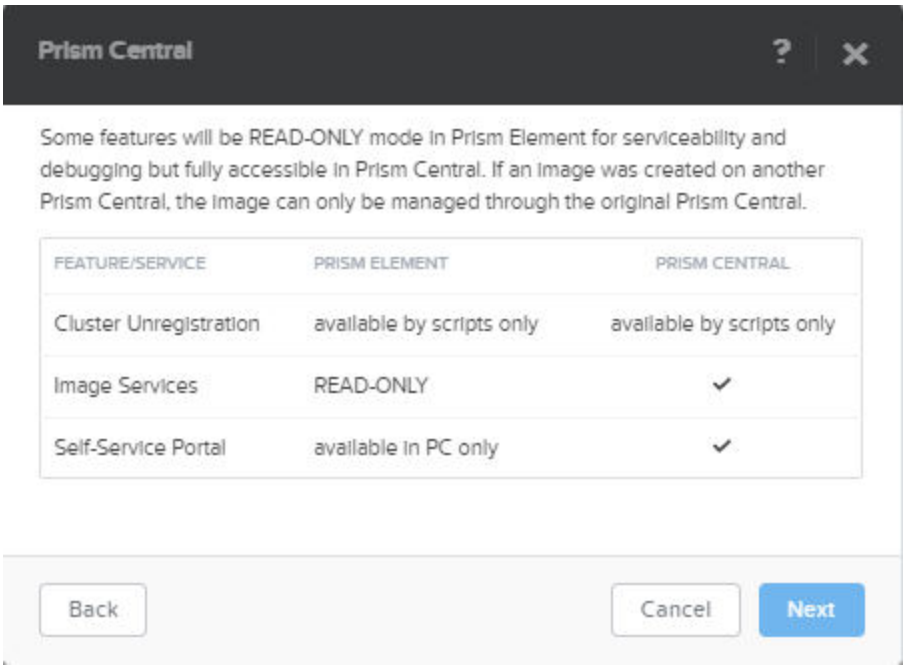

Figure 29: Prism Central Dialog Box (services screen)

6. In the registration screen, do the following in the indicated fields:
   a. **Prism Central IP**: Enter the IP address of the Prism Central VM.
   b. **Port**: The default port number is `80`. This is an optional field. For the complete list of required ports, see Port Reference.
   c. **Username**: Enter `admin` as the Prism Central user name.
   d. **Password**: Enter the Prism Central `admin` user password.

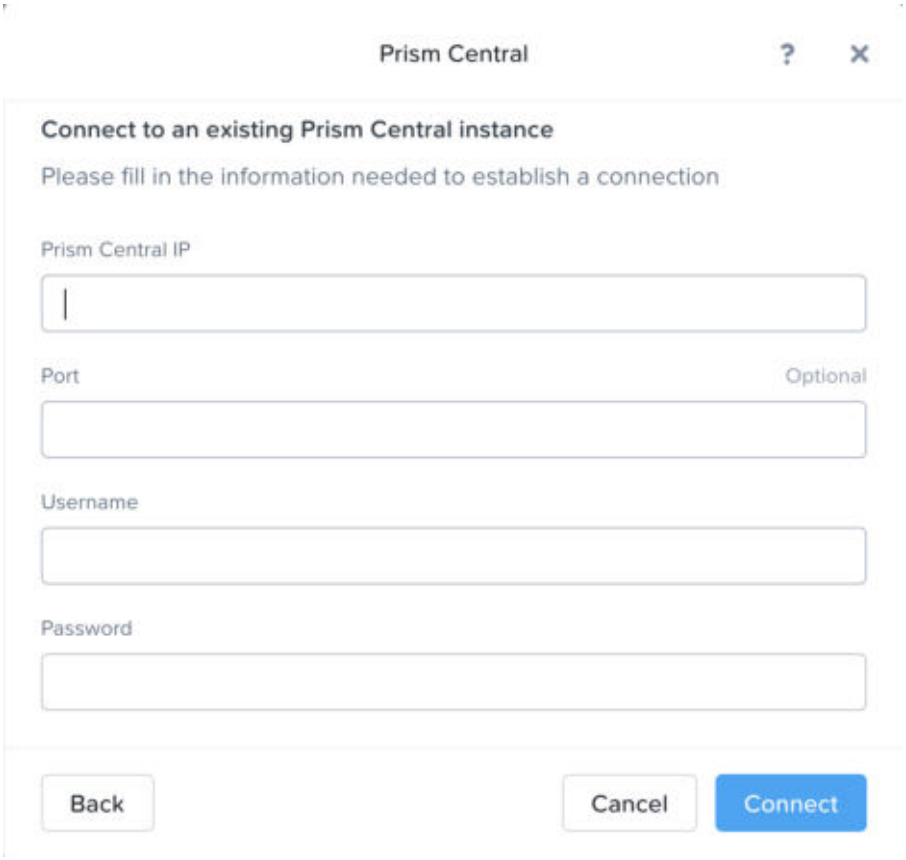

Figure 30: Prism Central Dialog Box (registration screen)

7. When all the fields are correct, click the **Connect** button to save the values and close the window.
   This registers the cluster on the specified Prism Central VM and allows the passing of information between the cluster and Prism Central.

> Note: The user credentials provided when registering a cluster (Prism Element) with Prism Central are only used once. After registration, modifying the admin password would not impact any communication between Prism Element and Prism Central.

Unregistering from Prism Central

About this task

To unregister a cluster from an existing Prism Central instance, do the following:

> Note: Unregistering a cluster through the Prism GUI is no longer available. This option was removed to reduce the risk of accidentally unregistering a cluster because several features

(including role-based access control, application management, micro-segmentation policies, and self-service capability) require Prism Central to run your clusters. If a cluster is unregistered from Prism Central, not only will these features not be available but the configuration for them may also be erased. Therefore, only the following procedure is available to unregister a cluster. See KB 4944 for additional details if you have enabled Self Service, Calm, or other special features in Prism Central.

Procedure

1. Log on to any Controller VM of the registered cluster through an SSH session.
2. Run the `cluster status` command and verify that all services are in a healthy state.

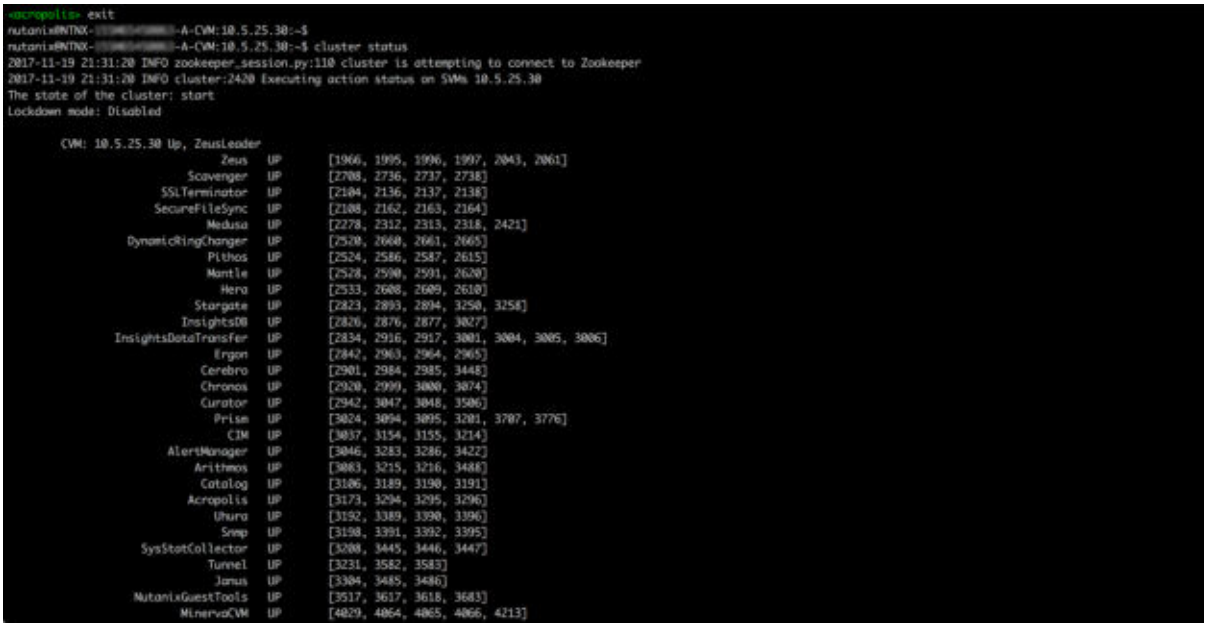

Figure 31: Cluster Status Command

3. Unregister the cluster from Prism Element.

```
nutanix@cvm$ ncli multicluster remove-from-multicluster external-ip-address-or-svm-ips=pc-name-or-ip username=pc-username password=pc-password force=true
```

Replace *pc-name-or-ip* with the Prism Central name or IP address and *pc-username* and *pc-password* with the login credentials for your Prism Central administrator account. This step can take some time (though typically just a few seconds). To check if the unregistration has completed, run the following command.

```
nutanix@cvm$ ncli multicluster get-cluster-state
```

4. Get the UUID for the cluster.

   ```
   nutanix@cvm$ ncli cluster info
   ```

   Find the "Cluster UUID" value in the displayed information.

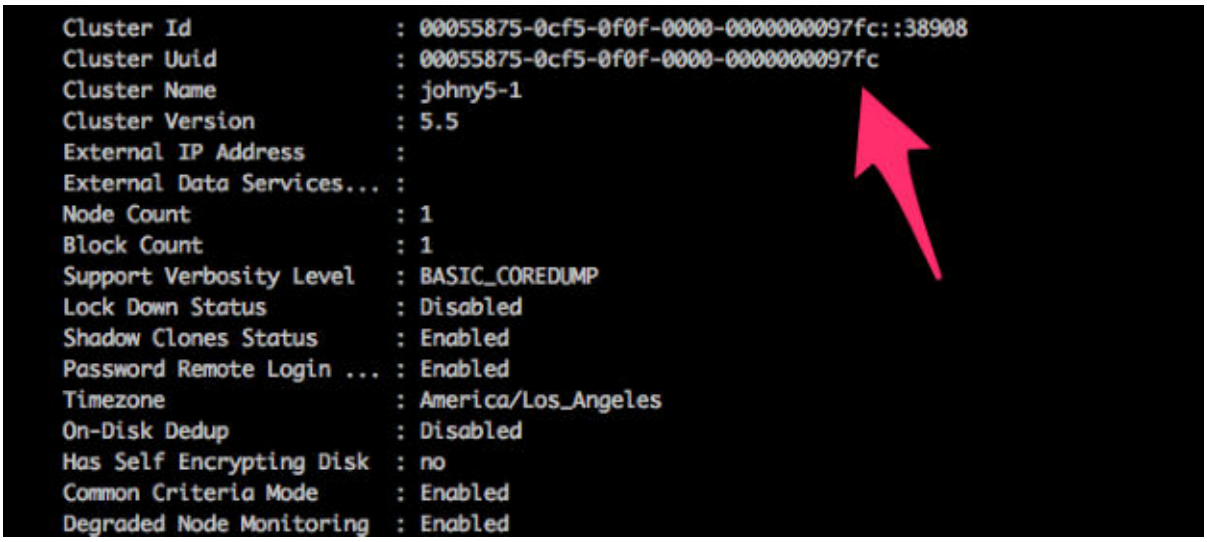

   Figure 32: Cluster Info Output

5. Log on to the Prism Central VM through an SSH session (as the "nutanix" user) and do the following:

   a. Run the unregistration clean-up script.

      ```
      [pcvm]$ python /home/nutanix/bin/unregistration_cleanup.py uuid
      ```

      Replace *uuid* with the value you obtained in step 4. This script removes all remaining registration information about that cluster and completes the unregistration process for the Prism Central VM.

   b. Get the UUID for Prism Central.

      ```
      [pcvm]$ ncli cluster info
      ```

      Find the "Cluster UUID" value in the displayed information (see step 4), which in this case is the UUID for Prism Central.

6. Go back to the Controller VM and run the **unregistration_cleanup.py** script to complete the unregistration process on the cluster.

   ```
   nutanix@cvm$ python /home/nutanix/bin/unregistration_cleanup.py uuid
   ```

   In this case the *uuid* is the Prism Central UUID obtained in step 5b.

What to do next

After performing these steps you can re-register the cluster with a new or re-created Prism Central instance.

If the clean up process does not complete successfully, try the following:

- Check the logs to indicate if there are any input errors when calling the script.
  The logs for the unregistration cleanup script can be found under **~/data/logs/unregistration_cleanup.log**.

- If errors occur during script execution, run the `cluster status` command and check that the cluster services are up and running. Rerun the script and check if it succeeds.

  Note: If Prism Self Service is enabled and the unregistered cluster is set as the default cluster, you must update the default project to select a new (registered) cluster. See the topic *Modifying A Project* in the Prism Central Guide.

Expanding (Scale Out) Prism Central

Before you begin

The following requirements must be met before you can expand Prism Central or add a Prism Central VM:

- The specified gateway must be reachable.
- No duplicate IP addresses can be used.
- The container used for deployment is mounted on the hypervisor hosts.
- When installing on an ESXi cluster:
  - vCenter and the ESXi cluster must be configured properly. See the *vSphere Administration Guide for Acropolis (using vSphere Web Client)* for details.
  - vCenter must be registered in Prism.
  - DRS must be enabled in vCenter.
  - vCenter is up and reachable during the deployment.

About this task

If Prism Central is just a single VM currently, you can expand it to three VMs. This increases both the capacity and resiliency of Prism Central (at the cost of maintaining two additional VMs). To expand this Prism Central instance across multiple VMs, do the following:

Note:

- Scale out Prism Central is supported on AHV and ESXi clusters only. If you are expanding a pre-AOS 5.17 Prism Central "large" VM that contains just a single vDisk, the upgrade process adds three new vDisks to each VM. This allows Prism Central to shard Cassandra metadata to improve performance. While the listed disk capacity of a Prism Central VM includes all the vDisks, the additional vDisks do not impact overall capacity. During the scale out you can monitor progress ("VM disk attach" tasks) from the Tasks dashboard.
- All scale out Prism Central VMs must run on the same cluster. For example, running two VMs in cluster_1 and one VM in cluster_2 is not supported.

Procedure

1. Click the gear icon and then select **Prism Central Management** from the Settings menu (see Settings Menu on page 16).
   The Manage Prism Central page appears. This page provides information about this Prism Central instance.

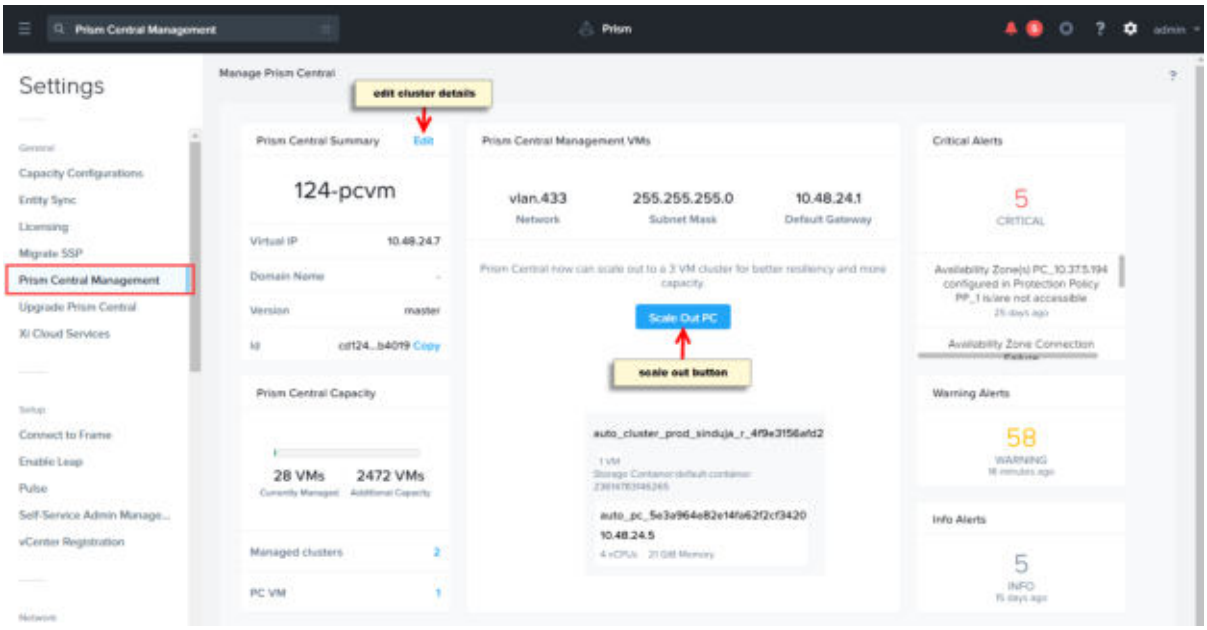

Figure 33: Manage Prism Central Page

2. To expand this Prism Central instance from one to three VMs, click the **Scale Out PC** button to display the Scale Out PC page and do the following:

   Note: A pop-up window appears explaining that scale out is a one-way process. Click the **Continue** button to display the Scale Out PC page. Once you scale out a Prism Central

instance from a single VM to multiple VMs, you cannot revert back. Deleting any of the Prism Central VMs may result in data loss.

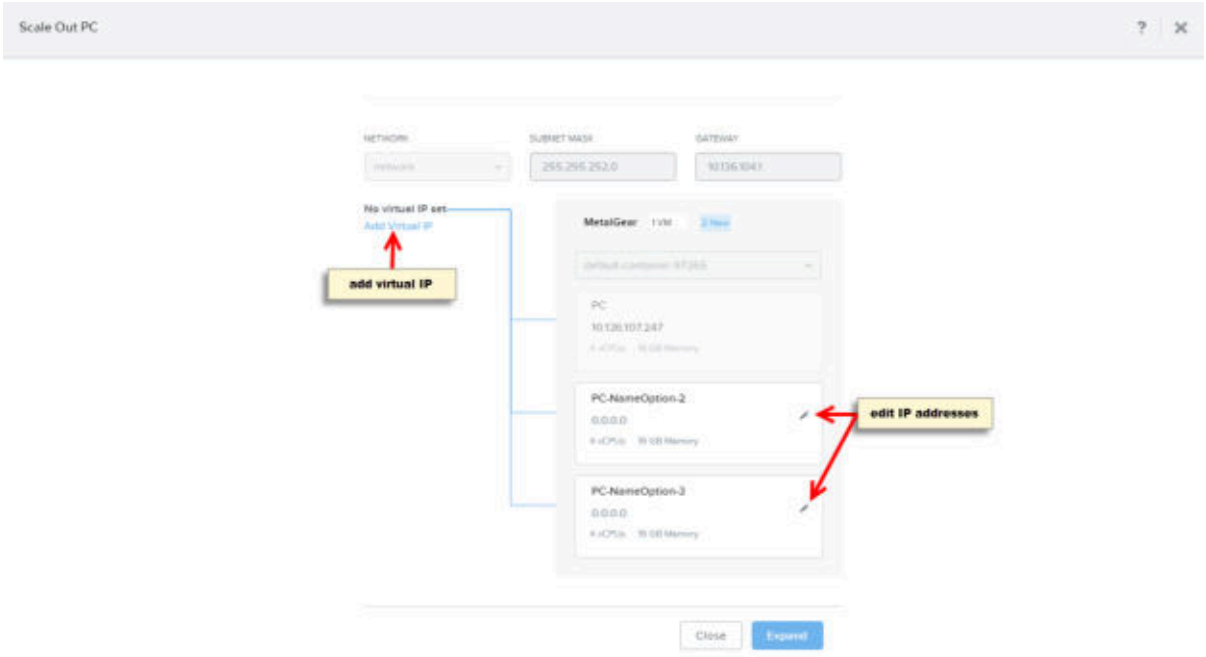

Figure 34: Scale Out PC Page

Note: The Prism Central VM image is deployed from the target cluster. If the required image cannot be accessed, typically because there is no Internet access (such as at a dark site), a pop-up error message appears. In this case you have the option to manually download and deploy the image as follows:

1. Log on to the Nutanix customer support portal, click **Downloads** > **Prism Central**, and click the download link for the target version to save the Prism Central binary .TAR and metadata .JSON files on your local media.
2. Log on (using SSH) to any Controller VM in the cluster specified in the error message and copy the Prism Central binary .TAR and metadata .JSON files to the Controller VM.
3. Run the following command to deploy the Prism Central image:

   ```
   nutanix@cvm$ ncli software upload software-type=PRISM_CENTRAL_DEPLOY file-path=file_path meta-file-path=metadata_file_path
   ```

   The *file_path* is the full (absolute) path to the .TAR file, and the *metadata_file_path* is the full path to the .JSON file. After this step completes, you can continue the scale out procedure.

a. Review the **Network**, **Subnet Mask**, and **Gateway** fields, which display the network name, subnet mask value, and gateway IP address applied to this Prism Central instance.

   The values in these three fields are read-only and cannot be changed.

   Note: The scale-out process uses the IP addresses from the network configured. If you used the default network configuration being the Private Network [default] while enabling

microservices infrastructure, Prism Central displays and uses Private Network [default] details in this section.

b. [optional] To specify a virtual IP address for Prism Central, click the **Add Virtual IP** link.

   This displays the Add PC Virtual IP window. A virtual IP can be used as a single point of access for Prism Central. Enter the IP address in the Virtual IP field and then click the **Update** button.

c. Specify IP addresses for the two new Prism Central VMs.

   1. Click the pencil icon for one of the new VMs. (The VM names are set automatically.) The **IP field** opens for editing. Enter the IP address and then click the green check mark icon to save that address.

Figure 35: Edit IP Field

   2. Repeat this step for the second new VM.

   The vCPU count and memory size in the new VMs are fixed and match the current Prism Central VM values; you cannot change these values.

d. When all the parameters are correct, click the **Expand** button.

   This starts the process of creating the new VMs and deploying this Prism Central instance. You can monitor progress from the Tasks page (see Tasks View on page 241).

What to do next

In some cases, especially when using Calm, a version mismatch with a linked product could occur after expanding Prism Central. To fix this situation, perform a life cycle manager (LCM) inventory after completing the Prism Central expansion procedure. See the LCM documentation for instructions on how to perform an LCM inventory.

Shutdown (Start Up) Prism Central VM

A Prism Central VM is managed like any other VM in a cluster through the Prism web console (Prism Element) for the cluster in which the Prism Central VM is running.

CAUTION: Prism Central supports features that could be damaged by shutting down the Prism Central VM at the wrong time. Do not shutdown a Prism Central VM until you are certain it can be done safely. If in doubt check with Nutanix customer service before attempting a shutdown.

- To shutdown (or start up) a Prism Central VM from the Prism web console, see the "Managing a VM (AHV)" or "Managing a VM (ESXi)" section in the *Prism Web Console Guide*.
- To shutdown (or start up) a Prism Central VM running on Hyper-V, see the appropriate Hyper-V documentation.

Prism Central Disaster Recovery

You can protect your Prism Central deployment (including scaleout deployment) and recover it when an event like a natural calamity, network outage, or power failure occurs and makes Prism Central unavailable. This feature enables you to automate continuous backup of your Prism Central VM (and various service configurations within it) to up to three AHV or ESXi clusters registered to Prism Central. The seamless 1-Click recovery restores important service configurations along with your Prism Central deployment.

The following services can be protected and data associated with them recovered.

- Leap
- Flow microsegmentation
- Prism Pro (AIOps)
- VM management
- Cluster management
- Identity and Access Management (IAM)

  > Note: IAM is recovered only after you enable microservices infrastructure.

- Categories
- Networking
- Licensing
- Reporting templates

Although the following services continue to run, they are not protected and therefore data associated with them not recovered.

- Flow networking (Virtualized Networking based Advanced Networking Controller - ANC)
- Calm
- LCM
- Karbon
- Objects
- Files
- Catalog
- Images
- VM templates

For Leap Recovery Jobs in progress, see this Note.

You can protect a Prism Central deployment either with the CLI-based Prism Central disaster recovery solution or the GUI-based Prism Central disaster recovery solution described in this document. For information about CLI-based Prism Central disaster recovery solution, contact Nutanix Support.

> Note: Nutanix does not recommend testing 1-Click Prism Central Disaster Recovery. Prism Central Disaster Recovery must be manually triggered only when the Prism Central VM becomes unavailable due to an unplanned event like a natural calamity, network outage, or power failure.

Protecting Prism Central

You can protect a Prism Central deployment to up to three clusters running AHV or ESXi. Clusters running Hyper-V are not supported.

Before you begin

Ensure that you meet the following requirements before you can protect your Prism Central.

- Prism Central must be running on version pc.2021.7 or newer.
- Prism Central must have NTP configuration to synchronize time between Prism Central and the registered clusters.
- The clusters used to protect (back up) Prism Central must be registered to Prism Central you want to protect.
- The clusters used to protect (back up) Prism Central must be running AOS 6.0 or newer.

About this task

Procedure

1. Log on to Prism Central.
2. Click the settings button (gear icon) at the top-right corner of the window.
3. Click **Prism Central Management** in the **General** section on the left pane.
4. Click **Protect Now** in the **Disaster Recovery** widget.
   The **Protect Now** option is available only when one or more AHV or ESXi clusters are registered to Prism Central.
   The Protect Prism Central window displays the services that will be protected and the services that will not be protected.
5. Click **Continue**, select one or more clusters to back up Prism Central, and then click **Proceed**. Clicking **Proceed** starts synchronizing Prism Central configuration data to the selected clusters. Creating the first backup on the selected clusters takes at least 30 minutes. After the first backup, the system keeps synchronizing Prism Central configuration data with the selected clusters every 30 minutes.
6. (optional) After the configuration, if you want to back up Prism Central to more clusters, click **+ Add Backup**. However, the number of clusters backing up Prism Central must not exceed three.
7. (optional) After the configuration, if you want to remove a cluster from backing up Prism Central, click **Remove**.

Recovering Prism Central (1-Click Disaster Recovery)

If you have protected your Prism Central deployment, you can use any registered AHV or ESXi cluster to recover Prism Central.

Before you begin

Ensure that you meet the following requirements before you can recover your Prism Central.

- The cluster used to recover Prism Central must be registered to the protected Prism Central.
- The ESXi cluster used to recover Prism Central must be registered to the vCenter Server.

- The cluster used to recover Prism Central must be running AOS 6.0 or newer.
- The cluster used to recover Prism Central must have ISCSI data service IP address for efficient recovery of Leap or Calm service configurations.

About this task

Procedure

1. Log on to any Prism Element registered to Prism Central you want to recover.
   The Prism Element dashboard shows Prism Central widget. Prism Central information (IP address and connection status) is available in the widget.
2. Click **Recover Now** in the Prism Central widget.
   The **Recover Now** option is available only when Prism Central is **Disconnected**.
   The Recover Prism Central window displays the services that will be recovered and the services that will not be recovered.
3. Click **Continue** and specify the following information.
   a. Select the cluster where you want to recover the Prism Central deployment.
   b. Verify the version of Prism Central that would recover on the selected cluster.
   c. Select the network where you want to recover and install Prism Central.
      The **Subnet Mask**, **Gateway**, and **DNS Address(s)** fields show the relevant information associated with the selected network.
   d. Enter details (name, IP address) for the Prism Central deployment you want to recover and click **Save**.
   e. Review the summary and click **Recover**.

   Prism Central recovers in at least 1 hour depending upon the configuration data it hosts. The recovered Prism Central shows all the guest VMs, disks, and metrics in at least 30–40 minutes more after the recovery. Do not perform any actions on the recovered Prism Central until all the recovery tasks are completed on the cluster. You can see the status of recovery and the related processes in the **Tasks** window.

What to do next

Consider the following after Prism Central recovery.

- Use the newly recovered Prism Central only.
  If the old Prism Central becomes available, shut down or delete the old Prism Central deployment because running the old Prism Central can cause data corruption.
- Reset the credentials.
  Prism Central recovers with the default credentials. Nutanix recommends changing the default credentials.
- Reconfigure the proxy server.
  If the old Prism Central had a proxy server, reconfigure the proxy server so that the recovered Prism Central maps to the correct IP address.

- Reconfigure the fully qualified domain name (FQDN).

  If the old Prism Central had an FQDN, reconfigure the FQDN so that the recovered Prism Central maps to the correct IP address.

- Leap Recovery Plan Jobs (RPJ) in progress: Perform the steps mentioned in KB-10962.

  If the old Prism Central had a failover task running (Leap) or protection policy with guest VMs protected with Synchronous replication schedule, perform the steps mentioned in KB-10962 to ensure that all the failover tasks stuck in the running state are terminated and a script is executed for efficient recovery of Prism Central.

- Re-enable Microservices Infrastructure.

  If the old Prism Central had Microservices infrastructure enabled, see *Enabling Microservices Infrastructure* in the *Prism Central Guide* to enable Microservices infrastructure on the recovered Prism Central.

IP Address Reconfiguration

In certain cases, you might need to reconfigure the IP address and gateway of a Prism Central VM. You can use this procedure to reconfigure the IP address and gateway of single Prism Central instances as well as Prism Central VMs in a Scale out Prism Central (clustered PC VMs or PC cluster).

Preparing to Reconfigure the IP Address and Gateway of PC VMs

Perform the tasks described in this topic before you start the IP address reconfiguration procedure.

About this task

You must perform the following tasks while the Prism Central VMs are still on the existing IP addresses.

Procedure

1. Coordinate Prism Central downtime, because the features and functionality of Prism Central will be unavailable for the entire duration of the IP address reconfiguration procedure.
2. Create a table to map the existing IP addresses with the new IP addresses of the Prism Central VMs for your reference.
3. Use SSH to log on to any running Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.
4. Verify if the PC VM or PC cluster is in a stable state.

   ```
   nutanix@pcvm$ cluster status
   ```

   This command is valid even if you only have a single PC VM and do not have a PC cluster.
5. Run the NCC health checks to make sure that the PC cluster is in a healthy state.
6. If you want to reconfigure the IP addresses of the entities in the AOS clusters (such as IP addresses of CVMs and hypervisor hosts), on which the PC is hosted, reconfigure those IP addresses first before you reconfigure the IP addresses of the PC VMs.

7. Stop the PC VM or PC cluster.

```
nutanix@pcvm$ cluster stop
```

This command is valid even if you only have a single PC VM and do not have a PC cluster.

Wait to proceed until an output showing all the services as DOWN is displayed, except the Zeus and Scavenger services.

Reconfiguring the IP Address and Gateway of Prism Central VMs

Before you begin

Ensure that you have completed the tasks described in Preparing to Reconfigure the IP Address and Gateway of PC VMs on page 60.

About this task

The procedure to reconfigure the IP addresses of PC VMs differs depending on whether or not IP Address Management (IPAM) is enabled in an AHV cluster. The procedure to reconfigure the IP addresses of PC VMs in an ESXi cluster is the same as the procedure of an AHV cluster in which IPAM is disabled.

IP Address Management (IPAM) is a feature of AHV that allows it to assign IP addresses automatically to VMs by using DHCP. See *IP Address Management* section in the *AHV Administration Guide*.

> Note:
>
> In this document, "managed network" refers to an AHV network that has IPAM enabled, and "unmanaged network" refers to an AHV network that does not use IPAM.

Perform the following procedure to reconfigure the IP address of a single PC VM or PC VMs in a PC cluster.

> Note:
>
> Perform steps 3, 4, and 5 only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

Procedure

1. Run the external_ip_reconfig script.

```
nutanix@pcvm$ external_ip_reconfig
```

2. Follow the prompts to type the new netmask, gateway, and external IP addresses.

   A message similar to the following is displayed when the procedure begins:

   ```
   External IP reconfig started
   ```

   A message similar to the following is displayed if the procedure is completed successfully:

   ```
   External IP reconfig finished successfully. Restart all the CVMs and start the cluster.
   ```

   The message indicates that you must restart the CVMs. In this case, you are reconfiguring the IP addresses of the PC VMs, so you must restart the PC VMs and not the CVMs.

   > Note:
   >
   > If your PC VMs are running in a managed network of an AHV cluster, do not start the PC VMs yet but proceed to step 3.
   >
   > If your PC VMs are in an unmanaged network of an AHV cluster or are in an ESXi cluster, skip steps 3, 4, and 5, and proceed to step 6 directly.

   A message similar to the following is displayed if the procedure fails:

   ```
   External IP reconfig Failed
   ```

   If the procedure fails, run the following command to check the log for troubleshooting purposes:

   ```
   nutanix@pcvm$ cat ~/data/logs/ip_reconfig.log
   ```

3. (Optional) Use SSH to log onto any CVM of the AOS cluster that is hosting the PC cluster or PC VM.

   > Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

4. (Optional) Update the NIC attached to the PC VM with the new IP address of the PC VM.

   > Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   ```
   nutanix@cvm$ acli vm.nic_update pc-vm-name pc-vm-mac request_ip=true ip=pc-vm-ip
   ```

   Replace the variables with their appropriate values as follows:

   - *pc-vm-name*: Name of the Prism Central VM whose IP address you modified.
   - *pc-vm-mac*: MAC address of the Prism Central VM whose IP address you modified (press **Tab** after you type the name of the PC VM to automatically populate the MAC address).
   - *pc-vm-ip*: New IP address of the Prism Central VM.

5. (Optional) Log on to the Prism Element web console and modify the IP address range in the IP pool for the PC network to work with the modified IP address of Prism Central VMs. In the Prism Element web console, perform the following:

   > Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   a. In the **VM** dashboard, click the **Table** view.

   b. Select a **PC VM** and click **Update**.

   c. In the **Update VM** dialog box, scroll down to the **Network Adaptors (NIC)** section.

   d. In the **VLAN Name** column, note the names of the networks attached to your PC VM.

   e. Click **Close**.

   f. Click the gear icon in the top-right corner, and under **Settings**, click **Network Configuration**.

   g. In the **Network Configuration** dialog box, identify the network names you noted in step d.

   h. Select each network and click the pencil icon next to the network.

   i. Under **IP Address Pools**, modify the IP address range if the new IP address of the PC VM is in a different IP address range than the previous IP address range.

6. Restart each Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.

7. Once the PC VMs are powered on, use SSH to log on to each PC VM or the single PC VM if you do not have a PC cluster and verify if the output of the following verification script displays the new IP addresses:

   ```
   nutanix@pcvm$ external_ip_reconfig_verify.py
   ```

8. Start the PC cluster.

   ```
   nutanix@cvm$ cluster start
   ```

9. Perform the following verification steps:

   a. Log on to the Prism Element web console of every cluster attached to the PC VM whose IP address you changed.

   b. On the **Home** page, verify if the **PC Registration** status is displayed as **OK**. This might take a few minutes after the PC VM or PC cluster is powered on.

3

MAIN DASHBOARD

The main dashboard is the opening screen that appears after logging into Prism Central. It provides a dynamic summary view across the registered clusters. To view the main dashboard at any time, select **Dashboard** from the Entities menu (see *Entities Menu* section).

> Important: Nutanix recommends using Google Chrome web browser to access Prism Central. The display of dashboard and widgets is optimized for Google Chrome browser.

Main Screen Layout

In addition to the main menu options (see *Main Menu* section), the main dashboard includes a menu bar and a display section for a collection of widgets that appear as tiles with targeted information about the registered clusters in each tile. The menu bar includes the following options:

- **Main Dashboard** tab. Click this tab to display the main dashboard. Additional tabs appear for any custom dashboards you create.
- **Manage Dashboard** button. Click this button to create a custom dashboard, edit the name of a custom dashboard, or delete a custom dashboard (see Creating a New Dashboard on page 73).
- **Reset Dashboard** button. Click this button to reset the main dashboard to the default set of widgets.
- **Add Widget** button. Click this button to add a widget to the displayed dashboard (see Adding Dashboard Widgets on page 74).

- **Data Density** button. Click this button to select the density of data representation in the widgets. Select one of the following three options:
  - Light—The information is lightly packed in the widget with more spacing between the elements.
  - Default—The information is moderately packed with sufficient spacing between the elements. This scheme of data density is based on user research by Nutanix.
  - Dense—The information is densely packed with minimal spacing between the elements for essential legibility.

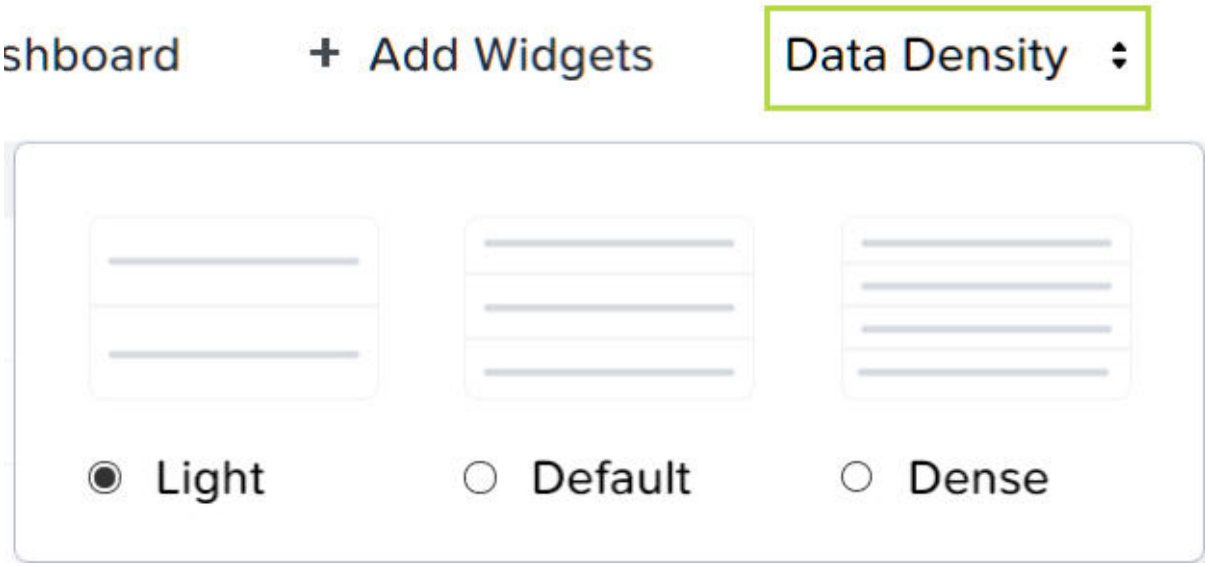

Figure 36: Data Density

Note: **Manage Dashboards** and **Add Widget** are Prism Pro license features. If Prism Pro is disabled, these buttons disappear, and you cannot customize or add dashboards. Any custom dashboards created before Prism Pro was disabled remain, but clicking the **Reset All** button (which replaces the **Reset Dashboard** button) deletes all custom dashboards and returns the main dashboard to the default state.

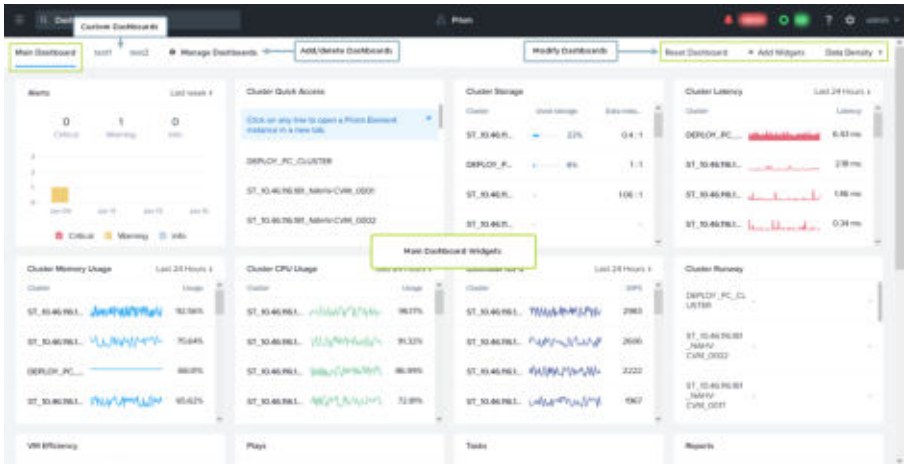

Figure 37: Main Dashboard

Main Dashboard Display

The main dashboard displays thirteen information tiles (widgets) by default. You can customize this view by adding or deleting widgets as desired (see Modifying a Dashboard on page 79).

The widgets are lined up in rows of four widgets per row. The default widgets are described in this section as they appear from left to right, top to bottom on the Main Dashboard page. A sample view of the Main Dashboard with the default widgets is displayed in the figure.

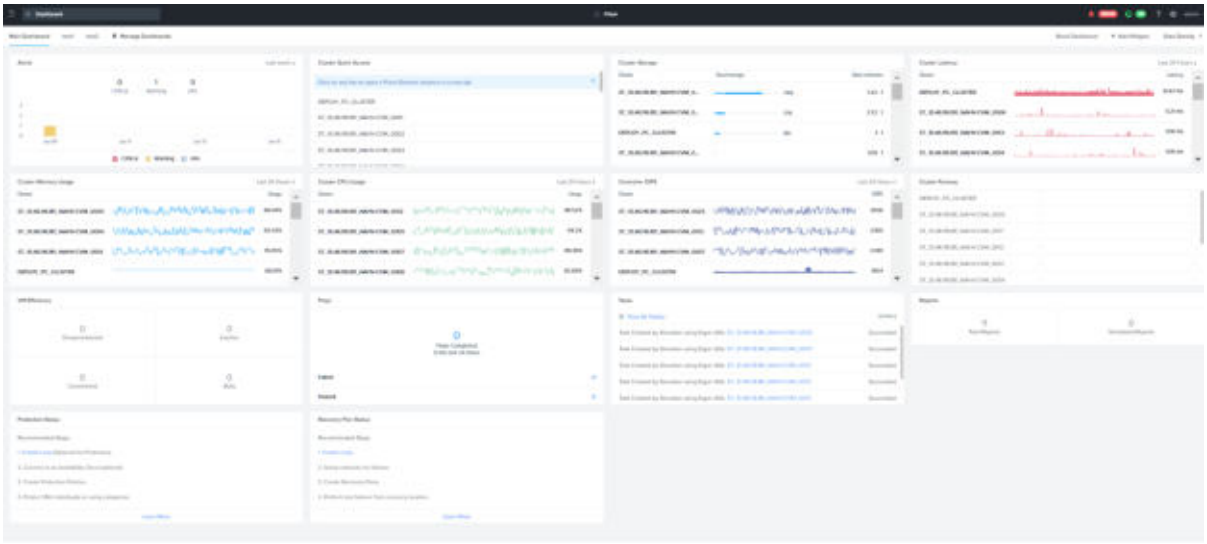

Figure 38: Main Dashboard with Default Widgets

The following figure displays the sample view of a widget with a period (time) selection:

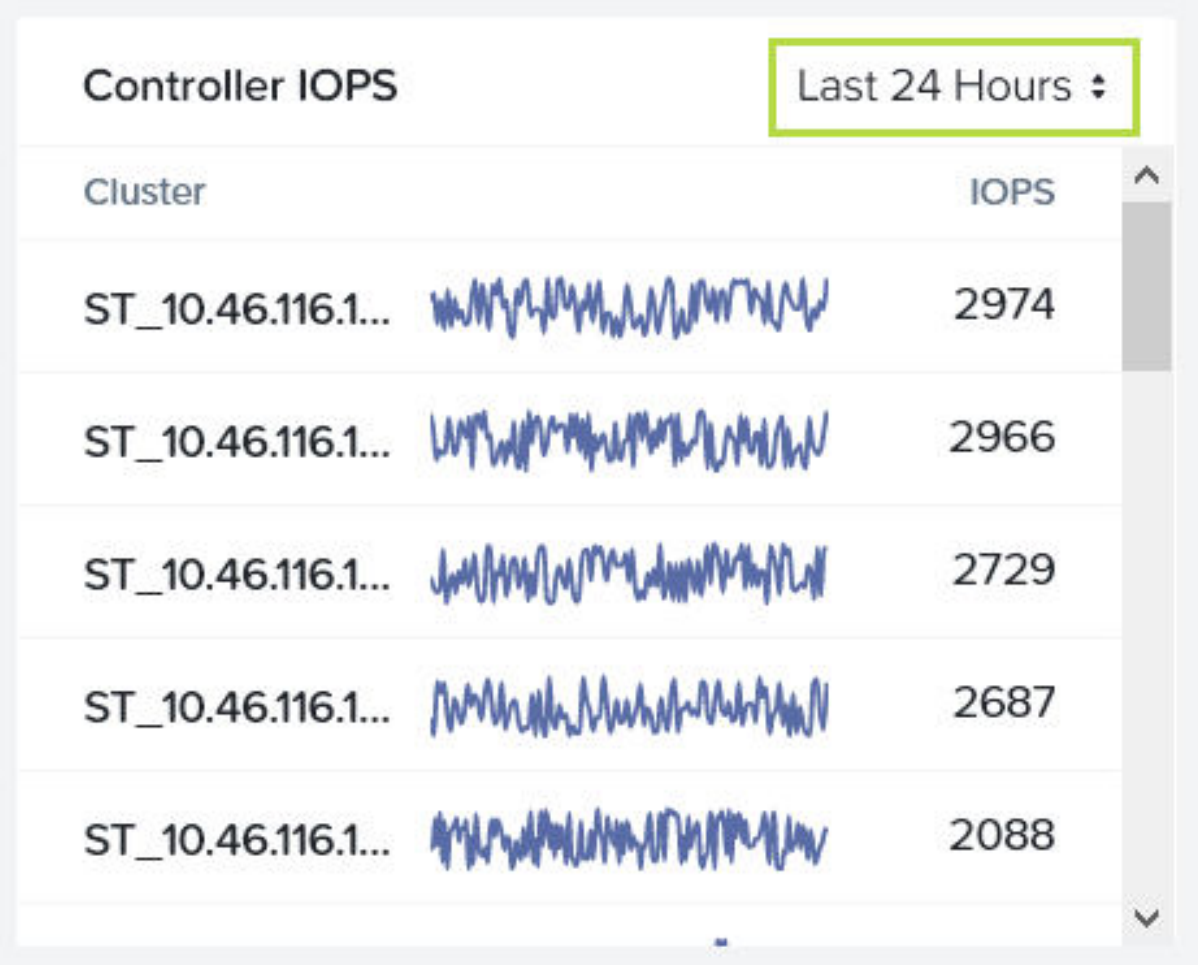

Figure 39: Sample Widget

- **Alerts**. This widget displays colored bar graph representing the alerts raised in the period selected. You can select the period on the drop down list at the top right corner of the widget. Each category of alerts is represented as a colored bar. Click any bar on the graph to see a list of the alerts represented in the bar. Click the **View All Alerts** link below the list to see all the alerts on the Alerts page.

  The Alerts categorized as **Critical**, **Warning** and **Info** are represented as red, yellow and blue bars respectively in the graphs.

- **Cluster Quick Access**. This widget displays a list of the clusters managed by the Prism Central. It allows you to open the Prism Element for a registered cluster in a new tab or window by simply clicking the line for that cluster. The cluster name appears on the left of each line while current health and alert icons (as appropriate) appear on the right.

- **Cluster Storage**. This widget displays storage and resiliency information for the highest usage clusters. Each line includes the cluster name, a usage column with a bar that visually indicates how much capacity is used currently, and a data resiliency column that displays the current status (critical, warning, OK, or unknown). The cluster list is ordered by data resiliency status with critical clusters at the top of the list. Click the cluster name to see the cluster details page. Hovering the cursor over the usage bar displays a box with usage details including data reduction ratio, usage percentage, used space, total space, fault domain (disk, node, or rack), and fault tolerance level (0, 1, or 2). The data reduction ratio indicates the data savings due to data reduction techniques such as deduplication,

compression, and erasure coding. (A 1:1 ratio indicates none of these data reduction techniques are in use currently.)

- **Cluster Latency**. This widget displays the total (read and write) IO latency average for the highest latency clusters. Clicking on a cluster name displays the summary page for that cluster.
- **Cluster Memory Usage**. This widget displays the percentage of total memory in use currently for the highest usage clusters. Clicking on a cluster name displays the summary page for that cluster.
- **Cluster CPU Usage**. This widget displays the percentage of total CPU in use currently for the highest usage clusters (or all clusters if there are fewer than five). Clicking on a cluster name displays the summary page for that cluster.
- **Controller IOPS**. This widget displays the total (read and write) controller IOPS for the highest volume clusters. Clicking on a cluster name displays the summary page for that cluster (see). The IOPS number comes from the controller when the hypervisor is AHV or Hyper-V and from the hypervisor when the hypervisor is ESXi.
- **Cluster Runway**. This widget alerts you to potential storage, CPU, or memory resource constraints across the clusters and provides an estimated runway (time remaining) before the resources are maxed out based on current usage. Clicking on a cluster name displays the capacity page for that cluster.
- **VM Efficiency**. This widget displays the number of VMs that are considered inefficient broken down by category: overprovisioned, inactive, constrained, and bully. It includes a link to the VMs dashboard for details about these VMs (see *VM Summary View*). See *Behavioral Learning Tools* in the *Prism Central Guide* for more information about VM efficiency.
- **Plays**. This widget displays a list of the Plays running on the cluster. The list is categorised to completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the Plays dashboard page.
- **Tasks**. This widget displays a list of recent tasks with the current status of each task. Click on the View All Tasks(s) link to view the Tasks page.
- **Reports**. This widget displays a table that lists the number of total and scheduled reports with a link to the reports dashboard (see *Reports Management* in the *Prism Central Guide*).
- **Protection Status**. This widget displays the protection plan status if one has been created. If not, it displays a set of recommended steps to create a protection plan.
- **Recovery Plan Status**. This widget displays the recovery plan status if one has been created. If not, it displays a set of recommended steps to create a recovery plan.

Widgets on a Dashboard

Resize the Widgets

You can vertically resize of a widget by clicking and dragging down the resize handle at the bottom of the widget. You cannot resize a widget horizontally.

The following figure provides an example of the resize handle.

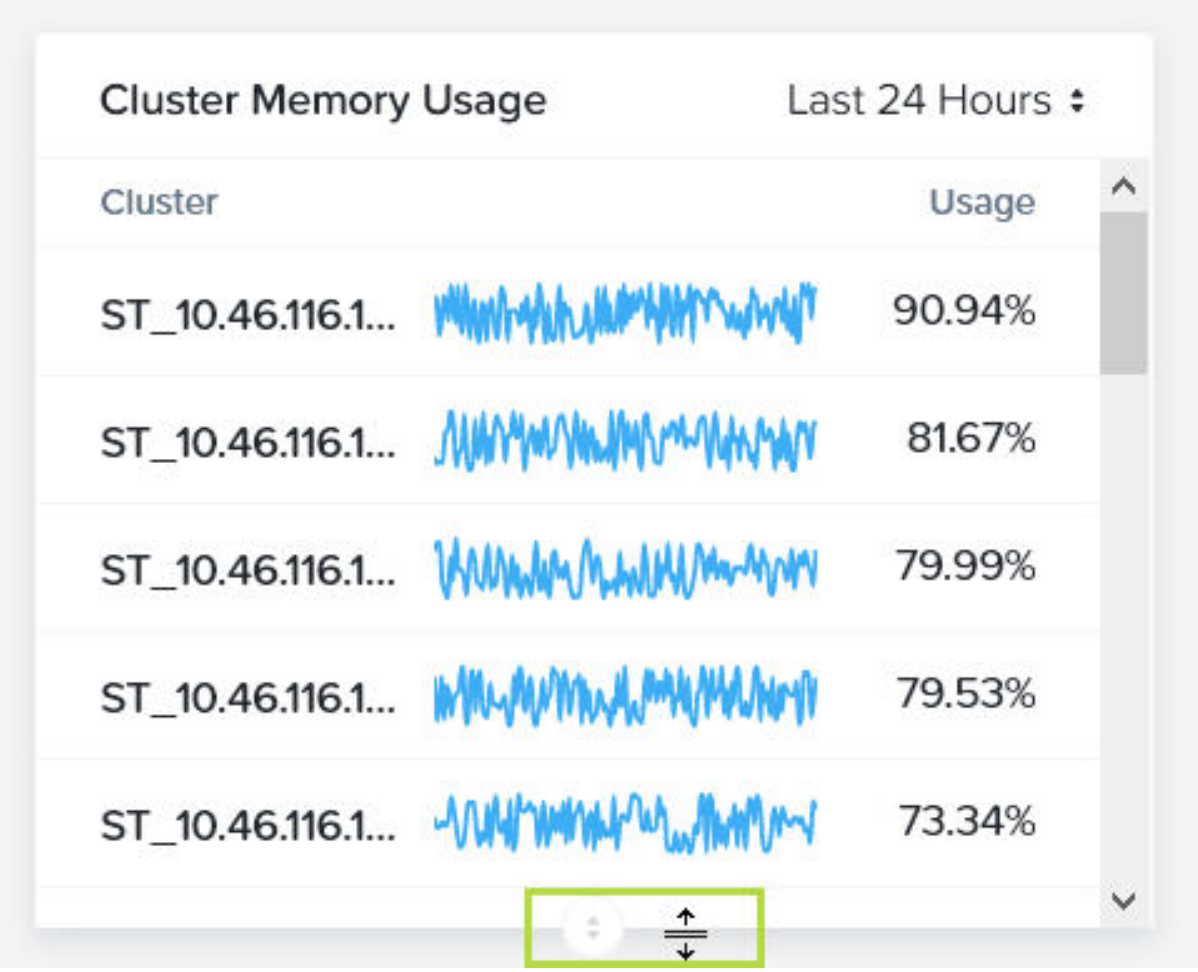

Figure 40: Vertically Resizing the Widget

When you click and drag up or down the resize handle to resize a widget, the widget placed below the resized widget shifts upward or downward.

You can resize a widget in multiples of half the default length of the widget, such as half, one-and-half, twice, two-and-half times the default length. Therefore, you can reduce the length to half or expand it to one-and-half times or twice the default length. Even if you attempt to resize it in any other variation of the default length, the widget auto-sizes itself to the nearest half length.

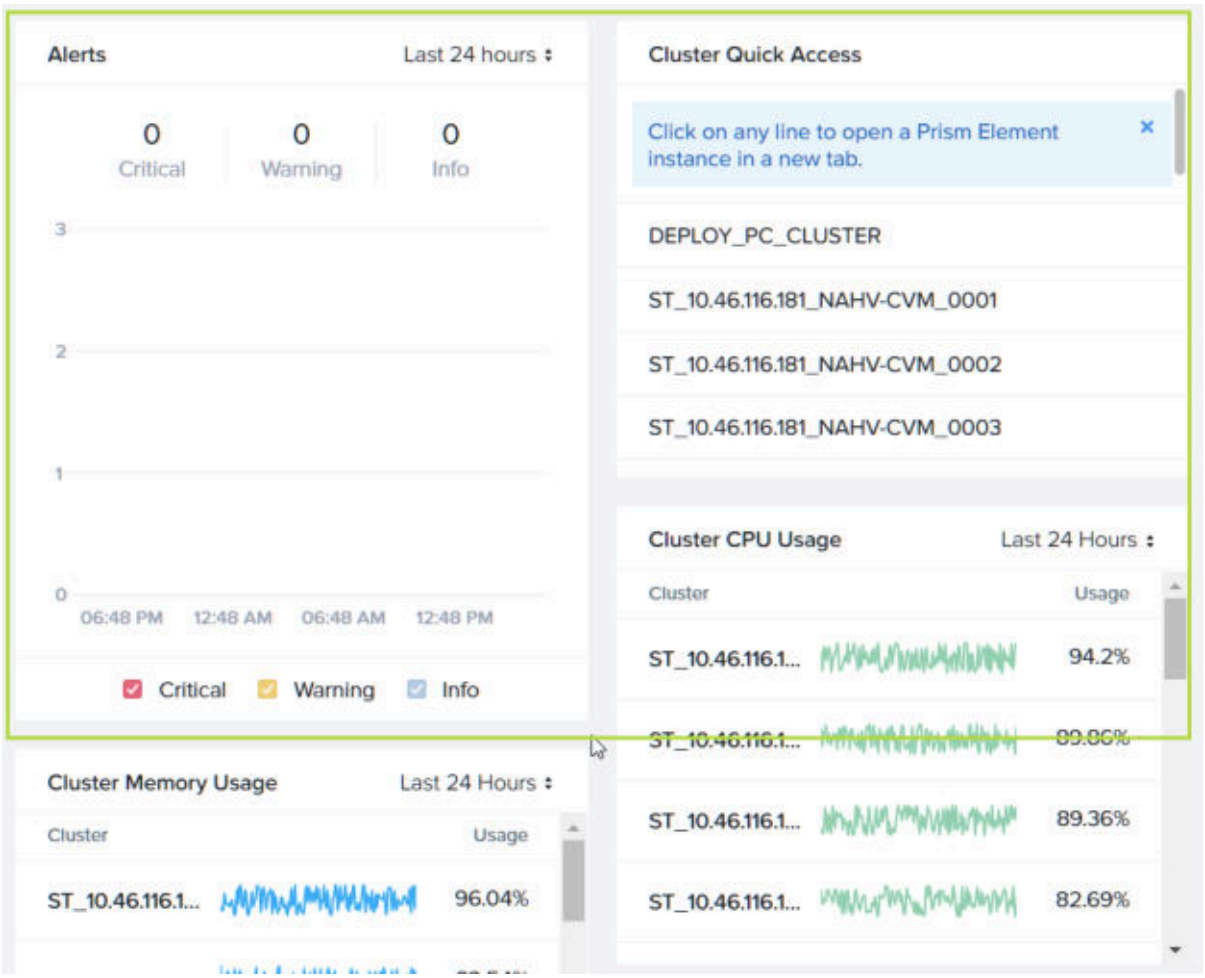

Figure 41: Widget Resize to One-and-half Times the Default Length

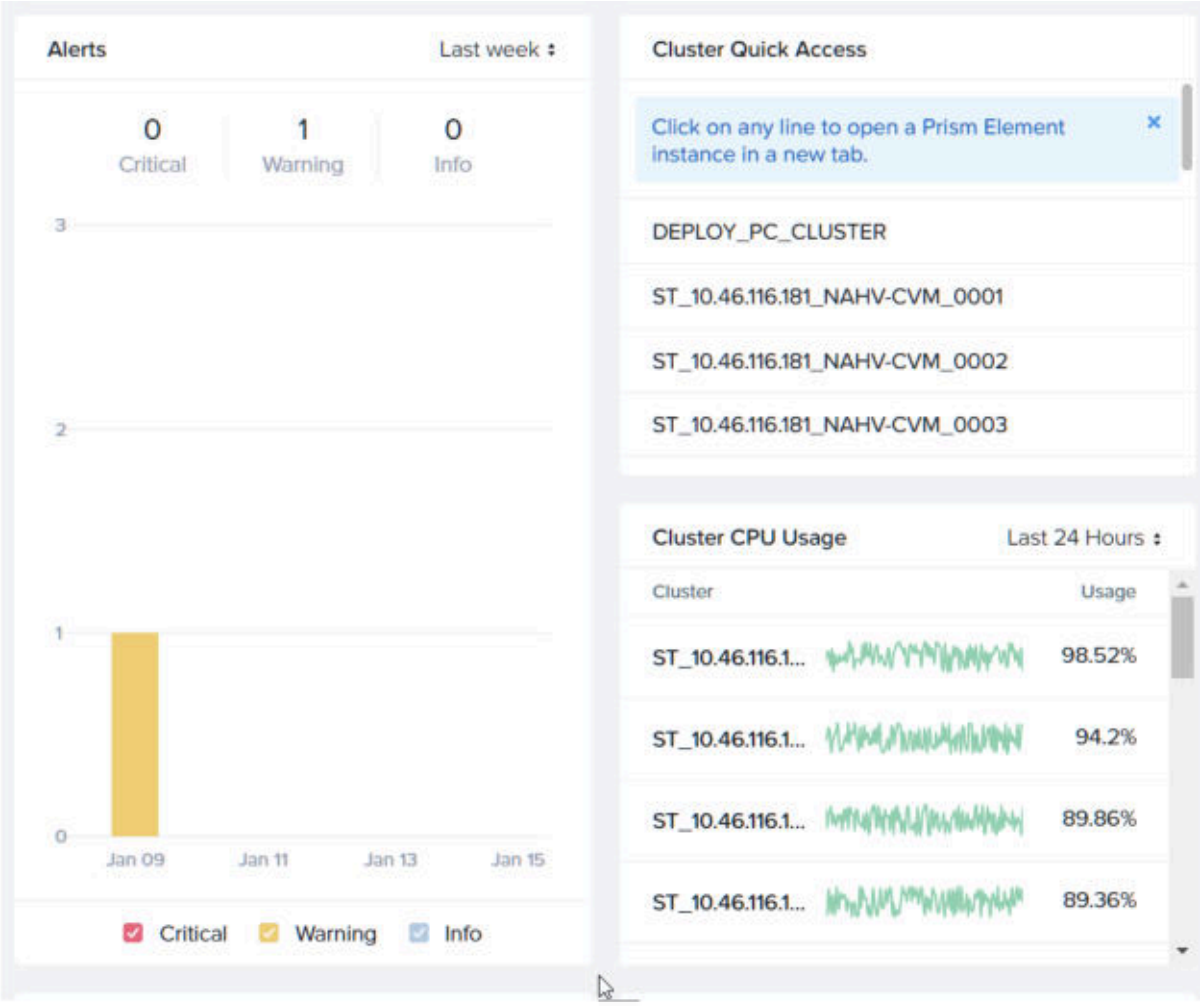

Figure 42: Widget Resize to Twice the Default Length

Drill-down or Expanding Data Elements in a Widget

You can expand or drill-down the data presented in the widgets by clicking on the data element. The following conditions are applicable:

- Expanding a data element in a widget is not possible in all widgets.
- In widgets that provide a list, clicking the list item or row opens an entity page in Prism Central or another application such as Prism Web Console in a new browser tab. For example, in the Cluster Quick Access widget, clicking the cluster name in the list opens the Prism Element Web Console for that cluster in a new browser tab. In the Cluster Storage widget, clicking a cluster name opens the cluster details view for that cluster.
- In widgets that provide numbers, clicking a number opens the respective entity page in Prism Central. For example, in the Reports widget, clicking the **Total Reports** or the **Scheduled Reports** number opens the Reports page. Similarly, clicking

- If the widget presents charts, click on the chart to display an expanded or drill-down view of the chart. Do not click on the entity (such as cluster name). If you do so, instead of displaying the drill-down view of the chart, the entity details page is displayed.

  The figure displays an example of expanding a chart in a widget.

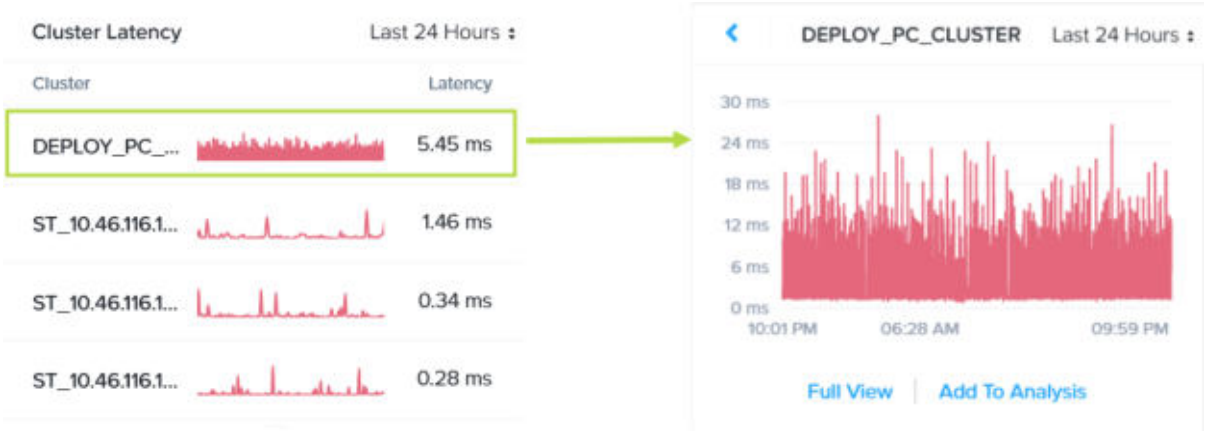

Figure 43: Widget Data Drill-down

At the bottom of the drill-down view, you can click **Full View** to open the details page of the entity(cluster or VM). Click **Add To Analysis** to add the chart to the Analysis dashboard (See *Analysis Dashboard* in *Prism Central Guide*.

Hover Details in Drill-down Views

After you click on a chart in a widget, the drill-down view of the chart is displayed. You can hover on any poin on the time chart to display the parameters at that point.

In the same example as presented in the Drill-down or Expanding Data Elements in a Widget on page 71, if you hover over the chart, the following sample details are displayed:

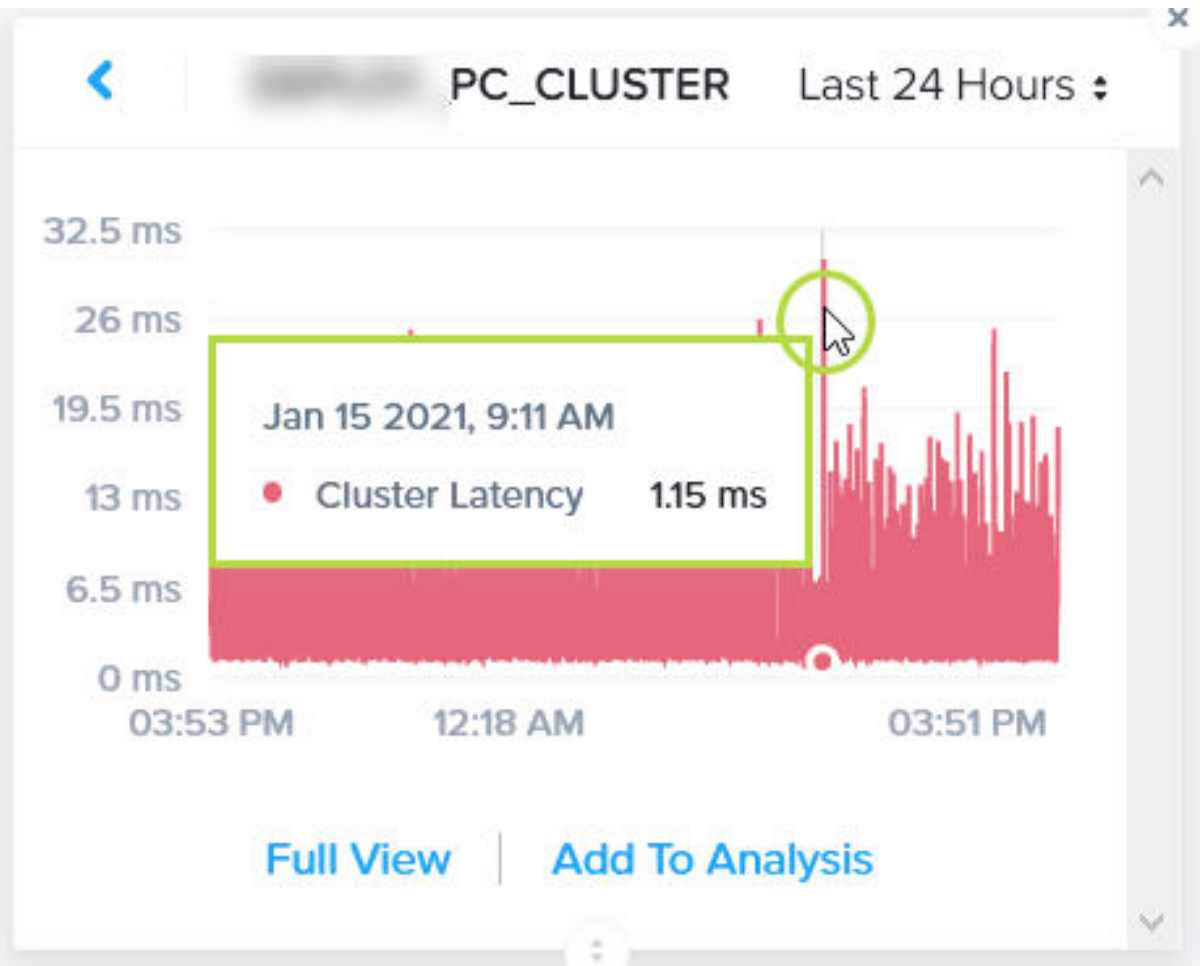

Figure 44: Hover Details

Creating a New Dashboard

About this task

The Main dashboard provides a default view into the clusters, but you can add custom views by creating custom dashboards. To create a custom dashboard, do the following:

> Note: Creating a dashboard requires a Prism Pro license. The **Manage Dashboards** button does not appear if Prism Pro is disabled.

Procedure

1. In the Main dashboard (see Main Dashboard on page 64), click the **Manage Dashboards** button.

2. In the Manage Dashboards window (see Modifying a Dashboard on page 79), do the following:

   a. Click the **New Dashboard** button.

   b. Enter a name for the dashboard in the displayed field.

   c. Click the **Save** button.

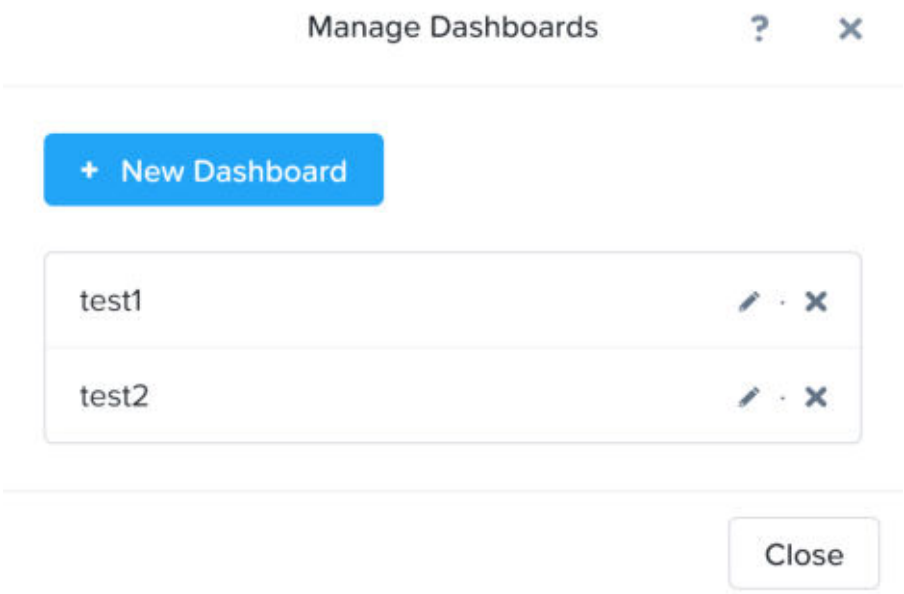

Figure 45: Manage Dashboards Windows (new dashboard)

3. Click the **Close** button to close the Manage Dashboards window.
   A tab for the new dashboard appears next to the **Main Dashboard** (or other custom dashboard) tab.

4. Click the tab for the new dashboard (to display that dashboard) and then click the **Add Widgets** button.
   The new dashboard is empty (no widgets) initially.

5. In the Add Widgets screen, select a widget and add it to the dashboard (see Adding Dashboard Widgets on page 74).
   Repeat this step until all the desired widgets have been added to the dashboard.

Adding Dashboard Widgets

Prism Central provides a selection of widgets that can be added to the home or a custom dashboard. Each widget provides information about a specific resource or usage topic that is displayed in a tile on the dashboard.

To add a widget to a dashboard, do the following:

> Note: Adding a widget to a dashboard requires a Prism Pro license. The **Add Widgets** button does not appear if Prism Pro is disabled.

1. Go to the target (main or custom) dashboard and click the **Add Widgets** button on the far right.

2. In the **Add Widgets** page, select the desired widget in the left column. (To find a particular widget, enter the name in the search field.) A preview of the selected widget appears in the middle column.
3. Enter values for the configurable parameters in the right column. The parameters vary by widget; some widgets do not have any configurable parameters.
4. Click the **Add to Dashboard** (or **Add & Return to Dashboard**) button to add the widget to the dashboard.

Figure 46: Add Widget Screen

The following table describes the widgets that you can add to a dashboard.

Table 10: Widget Parameters

Widget	Description	Parameter	Values
Custom Widgets			
Custom Alerts Widget	Displays a custom list of alerts.	Widget Name - Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster - Select the cluster(s) to monitor from the pull-down list. The default is all clusters.	All Clusters, <cluster name>
		Choose Size - Select the size (width/height dimensions) of the widget. The default is 1 x 1.	1 x 1, 2 x 1, 2 x 2, 2 x 3, 2 x 4, 4 x 4
		Severity. The default is critical.	Critical, Warning, Info

Widget	Description	Parameter	Values
		Entity. Select the entity of interest. The default is cluster.	Cluster, Storage, VM, Hardware, DR
Top Lists Widget	Displays a list of the top performers for a selected metric.	Widget Name - Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster - Select the cluster(s) to monitor from the pull-down list. The default is all clusters.	All Clusters, <cluster name>
		Choose Size - Select the size (width/height dimensions) of the widget. The default is 2 x 1.	2 x 1, 2 x 2, 2 x 3
		Entity. Select the entity of interest. The default is host.	host, VM
		Metric. Select the metric of interest. The default is IOPS.	IOPS, Memory Usage, CPU Usage, Bandwidth, Latency
Custom Chart Widget	Displays a performance graph for a selected entity and metric pair.	Widget Name - Enter a name for the widget. The default is "New Chart Widget".	(name)
		Entity Type. The default is host.	Host, Disk, Storage Pool, Storage Container, Virtual Machine, Virtual Disk
		Entity. (There is no default value.)	(entity name)
		Metric. (There is no default value.)	(metric name)
Cluster Info Widget	Displays cluster summary information about alerts, anomalies, runway, and inefficient VMs.	Widget Name - Enter a name for the widget. The default is "New Widget".	(n/a)
		Select a Cluster. Select the target cluster from the pull-down list.	(registered cluster names)
Compute & Storage Widgets - This widget appears on the Main dashboard by default (see Main Dashboard on page 64).			

Widget	Description	Parameter	Values
VM Efficiency	Displays a four-cell table that lists the number of overprovisioned, inactive, constrained, and bully VMs in the cluster with links to the details.	(no customizable parameters)	(n/a)
Hardware Widgets - All these widgets (except Performance) appear on the Main dashboard by default (see Main Dashboard on page 64).			
Cluster CPU Usage	Displays the percentage of total CPU in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Latency	Displays the total (read and write) IO latency average for the highest latency clusters.	(no customizable parameters)	(n/a)
Cluster Memory Usage	Displays the percentage of total memory in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Quick Access	Displays a list of registered clusters. Health and alert icons appear for each cluster. Clicking the cluster line opens Prism (element) for that cluster in a separate tab or window.	(no customizable parameters)	(n/a)
Cluster Runway	Displays storage, CPU, and memory runway estimates (time remaining before the resource reaches capacity). See the *Capacity Runway Tab* section in *Resource Planning* for more runway information.	(no customizable parameters)	(n/a)
Cluster Storage	Displays storage statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Controller IOPS	Displays the total (read and write) controller IOPS for the highest volume clusters.	(no customizable parameters)	(n/a)
Impacted Cluster	Displays information about any clusters that are impacted (performance, capacity, or other potential issues) and may need attention.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Performance	Displays latency, bandwidth, and IOPS statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Alerts	Displays colored bar graph representing the alerts raised in the period selected. Clicking the **View All Alerts** link below the list to see all the alerts on the Alerts page.	(no customizable parameters)	(n/a)
Activity Widgets - The Tasks widget appears on the Main dashboard by default (see Main Dashboard on page 64).			
Tasks	Displays a list of recent tasks with the current status of each task and a link to the Tasks dashboard.	(no customizable parameters)	(n/a)
Discovered Apps	Displays the number of applications discovered in a specified set of monitored clusters. Click the **App instances** number to display the application discovery dashboard. Click the **Identified** or **Unidentified** number to display the application discovery dashboard filtered for that condition. This widget displays the discovered apps information only if you enable App Discovery.	(no customizable parameters)	(n/a)
Operations Widgets - These widgets appear on the Main dashboard by default (see Main Dashboard on page 64).			
Plays	Displays a list of the Plays running on the cluster. The list is categorised to completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the Plays dashboard page.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Reports	Displays a table that lists the number of total and scheduled reports with a link to the Reports dashboard (see).	(no customizable parameters)	(n/a)
Disaster Recovery Widgets - These widgets appear on the Main dashboard by default (see Main Dashboard on page 64).			
Protection Status	Displays the protection plan status if one has been created. If not, it displays a set of recommended steps to create a protection plan.	(no customizable parameters)	(n/a)
Recovery Plan Status	Displays the recovery plan status if one has been created. If not, it displays a set of recommended steps to create a recovery plan.	(no customizable parameters)	(n/a)

Modifying a Dashboard

About this task

The Main dashboard provides a default view into the registered clusters, but you can customize that view at any time. To modify the Main dashboard or any other dashboard you create, do the following:

> Note: Customizing a dashboard requires a Prism Pro license. The **Add Widgets** and **Manage Dashboards** buttons do not appear if Prism Pro is disabled.

Procedure

1. Go to the Main dashboard (see Main Dashboard on page 64) and click the tab of the dashboard to modify (main or previously added custom dashboard).
2. To add a widget to the displayed (main or custom) dashboard, click the **Add Widgets** button, select a widget from the Add Widgets screen, and add it to the dashboard (see Adding Dashboard Widgets on page 74).
3. To delete a widget, click the X icon in the upper right of the displayed tile.

   A prompt appears to verify the delete; click the **OK** button. The widget then disappears from the dashboard.
4. To reset the Main dashboard to the default set of widgets (after you have previously added or deleted widgets), click the **Reset Dashboard** button.

   A prompt appears to verify the reset; click the **OK** button. The main dashboard returns to its default view.

5. To rename a custom dashboard, do the following:
   a. Click the **Manage Dashboards** button.
   b. Click the pencil icon for that dashboard.
   c. Enter a new name in the displayed field.
   d. Click the **Save** button.

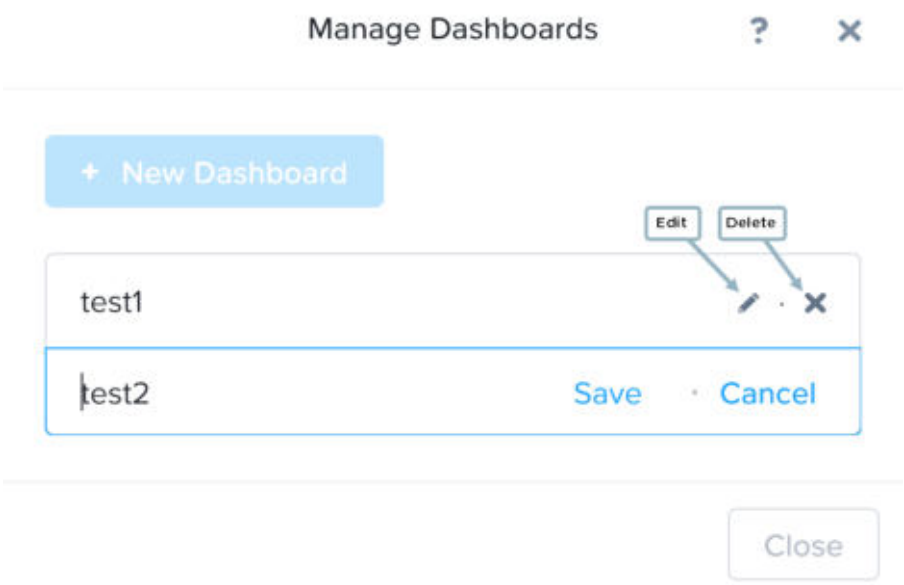

Figure 47: Manage Dashboards Window

6. To delete a custom dashboard, click the **Manage Dashboards** button and then click the X icon for that dashboard.
   A prompt appears to verify the delete; click the **OK** button. The tab for that dashboard disappears from the screen.

4

ENTITY EXPLORING

The entities menu (see Entities Menu on page 13) provides access to dashboards that allow you to view statistics about a range of entities across clusters and to organize that information in a variety of ways. An entity is an object type such as a VM, cluster, security policy, project, or report. Many of the entity dashboards include common elements in a summary view with drill-down options to view detailed information about individual entities of that type.

Note: This section describes common dashboard elements and how you can customize the view. See the appropriate "<entity> Summary View" section for a description of the specific fields and options applicable to each entity view. The Entities menu is organized into group categories, and the "<entity> Summary View" sections are located under one of the following group category headings:

- Compute and Storage Entities on page 86 (VMs, OVAs, images, catalog items, storage containers, storage policies, volume groups, vCenter datastores)
- Network and Security Entities (subnets, virtual private clouds, floating IPs, VPNs, security policies)
- Policies on page 185 (NGT policies, image placement policies)
- Data Protection and Recovery Entities on page 193 (protection summary, protection policies, recovery plans, VM recovery points)
- Hardware Entities on page 197 (clusters, hosts, disks, GPUs)
- Activity Entities on page 237 (alerts, events, tasks, audits)
- Operations Entities on page 244 (analysis, cost management, app discovery, monitoring integrations, operations policies, planning, playbooks, reports)
- Administration Entities on page 246 (categories, LCM, projects, roles, users, availability zones)
- Services Entities on page 260 (Calm, Files, Foundation Central, Karbon, Objects)

These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

Screen Layout

Common display elements in many dashboards include the following:

- A selection pane on the left that lists the entity type and number at the top.
  - For hardware and virtual infrastructure entity types like clusters and VMs, the pane includes five tabs: Summary, List, Alerts, Events, Metrics. Click the desired tab to display that content in the main section of the screen.
  - For most other entity types, the pane is blank (no tabs), and the relevant content appears automatically in the main section of the screen.
  - For some entity types, the pane does not appear.
- When a list appears in the main section, the display may include:
  - A menu bar with (left to right) selection options, action options (if any), grouping options, and display options for the list, and a **Filters** button on the far right that when clicked displays the **Filters** pane (hidden by default).
  - A query field that identifies the filters (if any) being applied to the list. This field displays all filter options that are currently in use. It also allows for basic filtering on the entity name.
  - A table (list) of entities. What is included in the table is determined by the grouping, display, and filter options. Entities are displayed 20 per page by default (use the scroll bar to scroll a page), but you can specify a different number per page (10 to 60) from the pull-down list above the table (click "*X - XX of XXX*" to display list).
  - A download icon (just above the table on the right); click the icon to download the table contents in CSV format. You can download a maximum of 1000 rows.

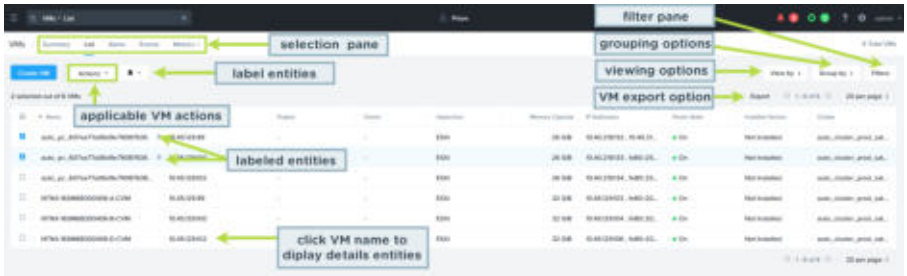

Figure 48: Dashboard Layout (example)

Filter Options

When the main section displays a list of entities, you can filter the list by clicking the far right icon in the menu bar to display the **Filters** pane. This pane includes a set of fields that vary according to the type of entity. Select the desired field values to filter the list on those values. An entry appears in the search field (see Searching for Information on page 20) for each value you select. You can save a filter by clicking the start icon in the search field, or you can remove a filter by clicking the X for that value.

Numeric filters have **To/From** fields to specify a range. These fields can take numeric values along with units. For example, the filter adjusts the scale accordingly when you type in "`10 K`" or "`100 M`".

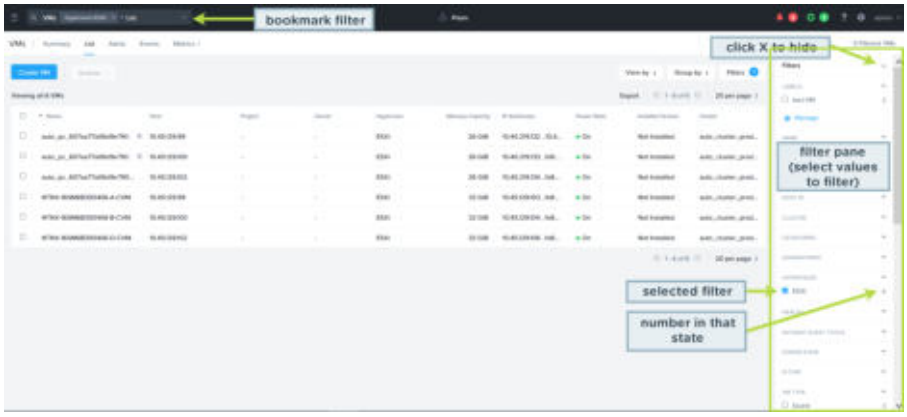

Figure 49: Filter Pane

Grouping Options

There are three (or four) drop-down menus on the right of the menu bar that allow you to organize the entity information in several ways.

- The **View by** menu allows you to specify the type of displayed information.

  Note: A general focus is available for all entities, but the other options are available only when appropriate for that entity.

  - Select **General** to display a set of general information parameters, as illustrated in the Explore Dashboard figure.
  - Select **Performance** to display a set of performance-specific parameters.
  - Select **Efficient** to display a set of efficiency-related parameters. "Efficiency" is determined through the VM behavioral learning engine (see Behavioral Learning Tools on page 838).
  - Select **GPU** to display a set of GPU-specific parameters.
  - Select **Data Protection** to display a set of data protection-related parameters.

  In addition to the preset **View by** displays, you can create one or more custom displays as follows:

  1. Select the **Add Custom** button at the bottom of the **Focus** menu.
  2. In the **<entity> Column** window, do the following:
     1. Enter a name for the custom display in the first (top) field.
     2. Select (click the blue + circle for) each entity property in the left column you want in your display. The selected entity properties appear in the right column. You can filter the entity-property list by entering a string in the search field above the left column.

You can select a maximum of 10 columns in one custom display including the Name column.

3. When the entity-property list in the right column is correct (you can remove an entity property from this list by clicking the red - circle for that entity), click the **Save** button.

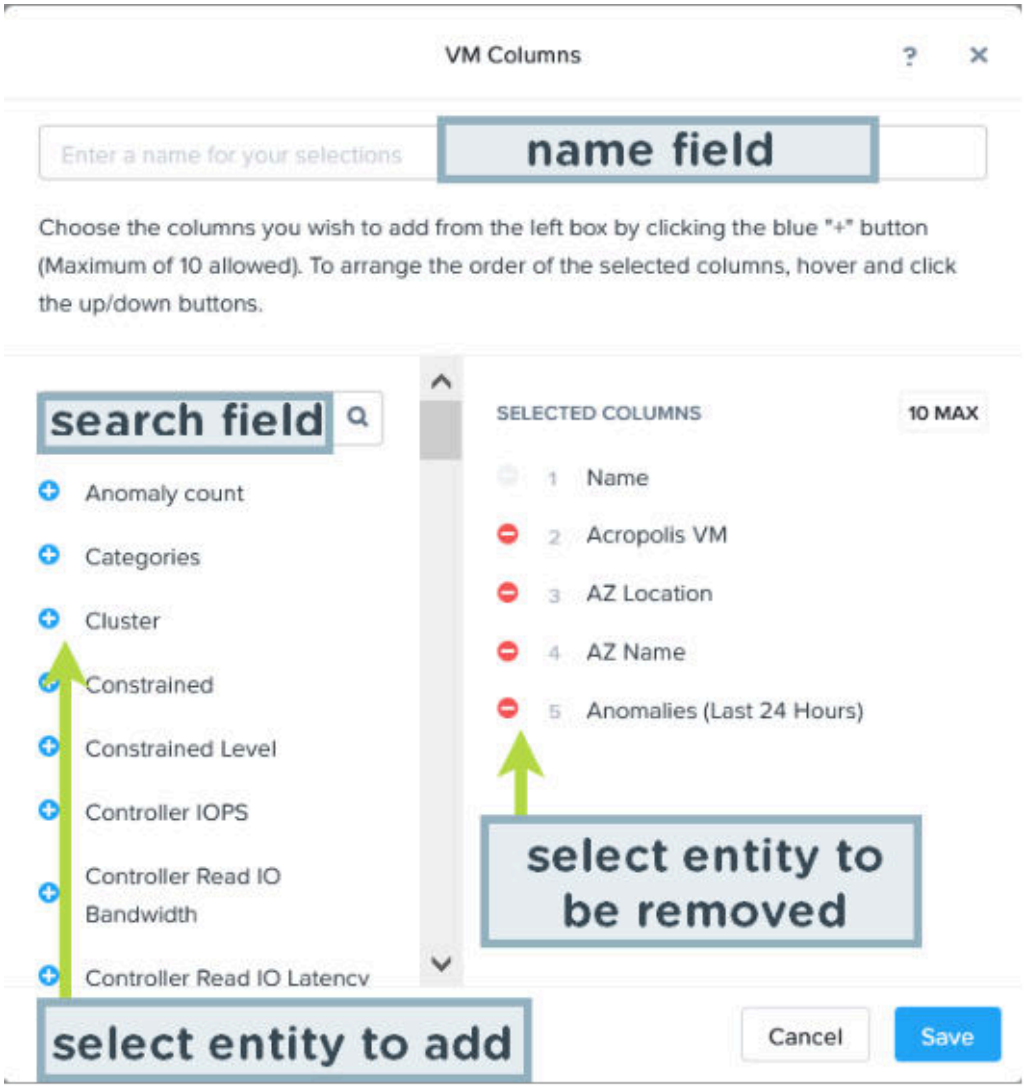

Figure 50: VM Column Window (example)

The **<entity> Column** window disappears and the new custom display appears (by name) in the **Focus** menu. To view that display at any time, select it from the menu.

> Note: Only the user that created a custom display can view that custom display. Other users and the administrative user cannot view that custom display.

- The **Color** menu allows you to color code the entries based on a selected parameter, as illustrated in the Circles View figure.
- The **Group** menu allows you to group the entries based on a selected parameter, as illustrated in the Circles View figure.

- The **Sort** menu allows you to sort the entries based on a selected parameter. This menu appears for the Tiles and Circles views only; ordering in the Grid (tabular) view is accomplished by clicking on a column header.

Action (and Selection) Options

The dashboard provides options to administer and categorize the entities. Actions are applied to selected entities that are currently visible. You can select a set of entities (rows) by selecting the first one and then using <Shift> select to select the last one.

- While some administrative tasks for a cluster must be done through Prism Element (see the *Prism Web Console Guide*), you can perform other administrative tasks directly from Prism Central. Action buttons for such tasks appear in the menu bar when an entity type is selected. For example, buttons for creating a VM and configuring the network appear in the VMs view. (No buttons appear when there are no relevant actions available for the selected entity type.)
- When one or more entities are selected from the list, an **Actions** menu appears in the menu bar. The menu includes actions that you can do to the selected entities. For example, the VMs menu includes actions such as power on or clone which will be executed on all the selected VMs. (Only currently valid actions are available; other menu options are inactive and appear gray.)
- When one or more entities are selected from the list, a label icon appears (next to the **Actions** menu). Labels allow you to create custom groupings for the entities. To add a label, click the icon and enter the label name in the field. The label is applied to all the selected entities in the list.

  > Note: Labels are available for VMs and clusters only.

- Use the selection menu to
  - Select all entries in the list by checking the box or choosing **Select all** from the pull-down list.
  - Unselect all by unchecking the box or choosing **Clear selections** from the pull-down list.
  - Show just the selected entries by choosing **Show selected entities** from the pull-down list.

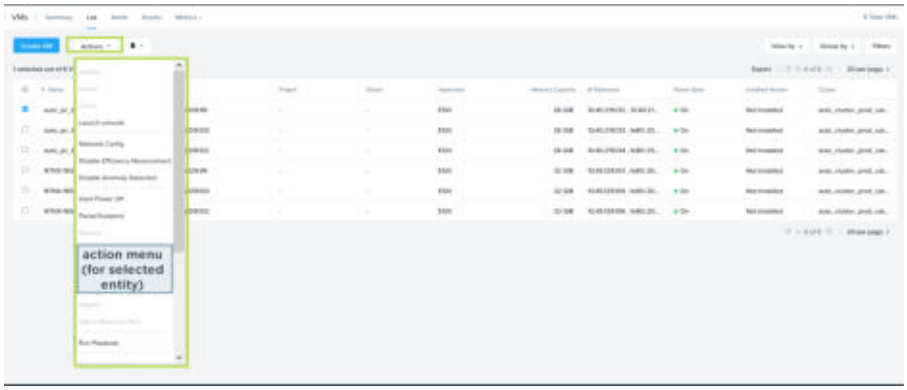

Figure 51: Action Options

Compute and Storage Entities

You can access dashboards for the following virtual infrastructure entities from the **Compute and Storage** category of the entities menu (see Entities Menu on page 13):

- VMs (see VMs Summary View on page 86)
- OVAs (see OVAs View on page 113)
- Images (see Images Summary View on page 117)
- Catalog Items (see Catalog Items View on page 121)
- Storage Containers (see Storage Containers Summary View on page 122)
- Volume Groups (see Volume Groups Summary View on page 139)
- vCenter Datastores (see vCenter Datastores Summary View on page 149)

> Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

VMs Summary View

To access the VMs dashboard, select **Compute & Storage** > **VMs** from the entities menu (see Entities Menu on page 13). The VMs dashboard summary view displays information about VMs across the registered clusters and allows you to access detailed information about each VM. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> Note: This section describes the information and options that appear in the VMs dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.
> - See VM Management on page 535 for information about creating and managing VMs.

Summary Tab

Clicking the **Summary** tab displays the following four widgets:

- **Suggested**: Displays a list of the VMs with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **Memory Usage**, and **IOPS**. Click the **View All XX VMs** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of VM-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click the graph, which then displays a list of those alerts. Clicking an alert displays the details page for that alert.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.
- **VM Efficiency**: Displays the number of VMs that are considered inefficient broken down by category (overprovisioned, inactive, constrained, and bully). This widget is the same that

appears in the main dashboard (see Main Dashboard on page 64). See Behavioral Learning Tools on page 838 for more information about VM efficiency.

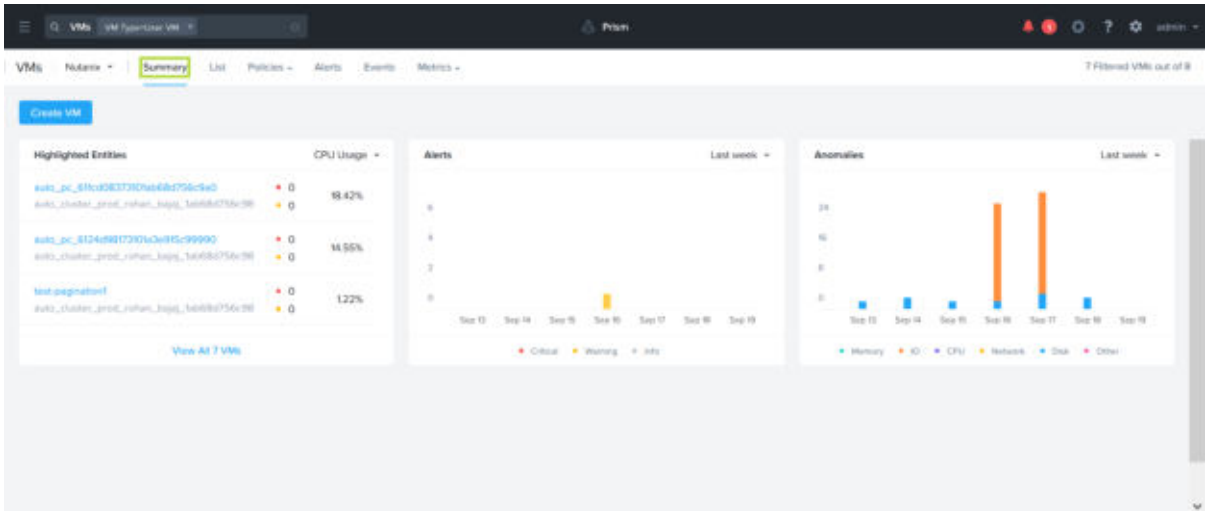

Figure 52: VMs Summary Tab (Nutanix)

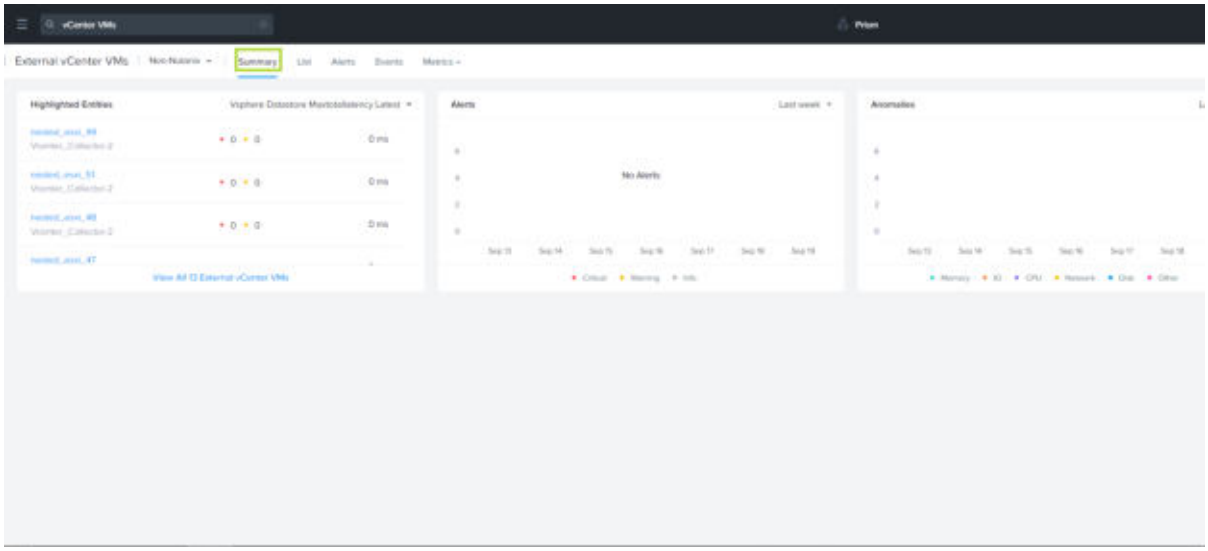

Figure 53: VMs Summary Tab (Non-Nutanix)

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the VMs across the registered clusters. The following table describes the fields that appear in the VMs list. The fields vary based on the **Focus** menu selection, which is **General**, **Performance**, **Efficiency**, or **GPU**. A dash (-) is displayed in a field when a value is not available or applicable.

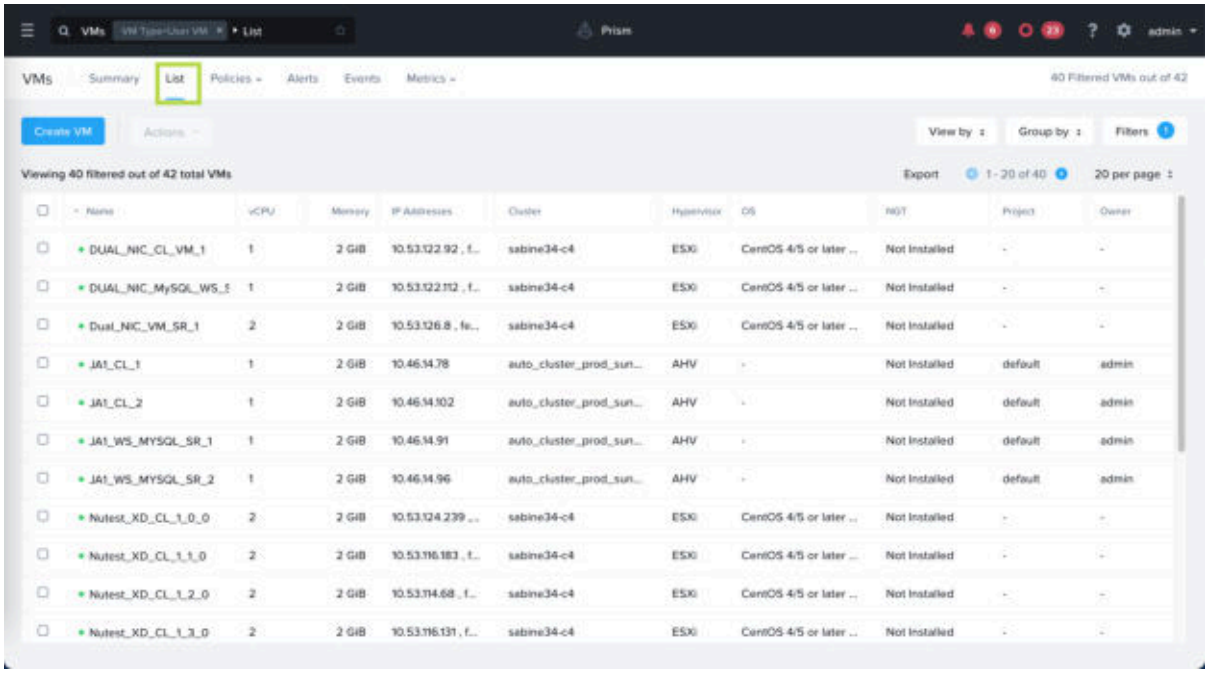

Figure 54: VMs List Tab

Table 11: VMs List Fields

Parameter	Description	Values
"General" Focus fields		
Name	Displays the VM name. Clicking the name displays the details page for that VM (see VM Details View on page 96).	(VM name)
vCPU	Displays the virtual CPU count of the VM.	(vCPU count)
Memory	Displays the total amount of memory available to this VM.	xxx [MB|GiB]
IP address	Displays one or more VM IP addresses.	(IP address)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Hypervisor	Displays the hypervisor type on which the VM is running.	AHV, ESX, or Hyper-V
OS	Displays the guest operating system of the VM.	(OS)
NGT	Displays whether Nutanix Guest Tools is installed on the VM.	Installed, Not Installed
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
"Performance" Focus Fields		

Parameter	Description	Values
Name	Displays the VM name.	(VM name)
Memory usage	Displays the percentage of allocated memory capacity currently being used by this VM.	0 -100%
Controller Read IOPS	Displays read I/O operations per second (IOPS) for this VM.	(number)
Controller Write IOPS	Displays write I/O operations per second for this VM.	(number)
Controller I/O Bandwidth	Displays I/O bandwidth used per second for this VM.	xxx [MBps|KBps]
Controller I/O Latency	Displays the average I/O latency for this VM.	xxx [ms]
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"Efficiency" Focus Fields		
Name	Displays the VM name.	(VM name)
Efficiency	Displays the assessed efficiency of the VM. If a VM is performing in an expected range, the efficiency is listed as "Good". If not, the type of inefficiency is displayed as determined by the VM behavioral learning engine (see Behavioral Learning Tools on page 838).	Good, Constrained, Overprovisioned, Inactive, Bully; ("NA" if insufficient baseline data for categorization, which is typically 21 days)
Efficiency Detail	Displays why a VM is considered inefficient. A dash (-) appears for a "Good" VM.	(text message)
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"GPU" Focus Fields		
Name	Displays the VM name.	(VM name)
GPU Configuration	Displays the GPU board (in case of Passthrough mode), vGPU software, type and number of vGPU instances in brackets (in case you configured the vGPU mode).	(GPU board/software-type name) Example: Passthrough - Nvidia Tesla M60 Example: vGPU - Nvidia GRID M60-8Q (2)
GPU Type	Displays the GPU configuration made as Passthrough or vGPU or if you did not configure GPU then "-"	[Passthrough|vGPU|-]

Parameter	Description	Values
GPU Usage	Displays the percentage of GPU capacity being used by the VM.	(percentage)
GPU Framebuffer Usage	Displays the percentage of GPU framebuffer (RAM) capacity being used by the VM.	(percentage)

You can filter the VMs list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range. Click the plus (+) sign to the right of the parameter name to expand the fields for that parameter; click the minus (-) sign to contract those fields.

Table 12: Filter Pane Fields

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the pull-down list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)
Name	Filters on the VM name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the VM name condition/string. Note: In this and the following two fields, the condition menu options are **Contains**, **Doesn't contain**, **Starts with**, **Ends with**, and **Equal to**.	(VM name string)
Host	Filters on the host name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the host name condition/string.	(host name string)
Cluster	Filters on the cluster name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the cluster name condition/string.	(cluster name string)
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a pull-down list appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of VMs tagged to each selected category is displayed on the right of the line.	(category name)

Parameter	Description	Values
Hypervisor	Filters on the hypervisor type. Check one or more of the boxes to filter on those hypervisors. The number of VMs currently on each hypervisor type is displayed on the right of the line.	AHV, ESX, HyperV
Health	Filters on the VM health state (good, warning, or critical). Select one or more states to return a list of VMs in that state(s). The number of VMs currently in each state is displayed on the right of the line.	Critical, Warning, Good
Power State	Filters on the VM power state. Select one or more states to return a list of VMs in that state(s). The number of VMs currently in each state is displayed on the right of the line.	On, Off, Suspended, Paused, Unknown
VM Type	Filters on the type of VM. Select either user VM or Controller VM (or both). The number of VMs of each type is displayed on the right of the line.	User VM, AHV Controller VM
Memory Usage	Filters on the amount of memory capacity being used. Check the box for the desired percentage range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of VMs utilizing memory in that range.	([xx] to [yy]% range)
Read IOPS	Filters on the read IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of VMs with read IOPS in that range.	([xx] to [yy] range)
Write IOPS	Filters on the write IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of VMs with write IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of VMs with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of VMs with average I/O latency in that range.	([xx] to [yy] range)

Parameter	Description	Values
Over Provisioned	Filters for over-provisioned VMs. Check the box(es) for the desired type (high and moderate). Note: This and the following two fields filter based on a VM efficiency algorithm, which is part of the VM behavioral learning capabilities (see Behavioral Learning Tools on page 838).	High, Moderate
Constrained	Filters for constrained VMs. Check the box(es) for the desired type (high and moderate).	High, Moderate
Efficiency	Filters for certain VM profiles. Check the boxes for the desired profile types. There is one for efficient VMs (good) and four for inefficient VMs (bully, over-provisioned, constrained, inactive).	Bully, Over Provisioned, Constrained, Inactive VM, Good
GPU Configuration	Filters for GPU configuration information such as model name. Enter the GPU configuration information in the field and then check the box. As you type a pull-down list appears to help you select the correct configuration information.	(configuration info)
GPU Type	Filters for GPU operational mode. Check the box for one or more of the GPU types.	vGPU, Passthrough, Passthrough(Compute)
GPU Usage	Filters on the amount of GPU capacity being used. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of GPUs in that range.	([xx] to [yy]% range)
GPU Framebuffer Usage	Filters on the amount of GPU framebuffer (RAM) capacity being used. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of GPUs in that range.	([xx] to [yy]% range)
vGPU Guest driver Version	Filters on the guest driver version. Enter the guest driver version number in the field.	(guest driver version number)

Export VM information. You can export the table containing the list of VMs and their information to a file in a CSV format by clicking the export icon on the right.

You can group the VM list in the following ways:

- The **Group** pull-down menu allows you to group the VM entries by cluster, hypervisor, power state, virtual CPU count, or health state. (You can only choose one.)

- When you select one or more VMs, the **Label** icon appears (to the left of the **Actions** menu). To assign the selected VMs a group label, click the **Label** icon and do one of the following:

Figure 55: Label Icon

  - To assign an existing label (if there are existing labels), select the desired label and then click **Apply changes**.
  - To assign a new label, enter a meaningful label in the blank text field and then click **Create new label**.

  You can now perform any actions available on this labeled group from the **Actions** menu (see following section). You can edit or delete the label by clicking the **Label** icon and selecting **Manage Labels**.

The VMs view includes two action buttons:

- To create a VM, click the **Create VM** button. You can create a VM on a cluster running AHV or ESXi (see Creating a VM (AHV) on page 536 or Creating a VM (ESXi) on page 576).
- To configure the network, click the **Network Config** button (see Configuring Network Connections on page 725).

The **Actions** menu appears when one or more VMs are selected. It includes the following actions:

- These actions can be applied to multiple VMs: **Delete**, **Power on**, **Power off**, **Pause/Suspend**, **Resume**, **Protect**, **UnprotectManage Categories**, **Quarantine VMs**, **Unquarantine VMs**, **Enable NGT**, **Disable NGT**, **Manage Ownership**
- These actions can be applied to only one VM at a time: **Update**, **Clone**, **Create VM Template**, **Launch console**, **Snapshot,**, **Migrate**, **Configure VM Host Affinity**, **Add to Catalog**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the current state of the selected VM(s). See Managing a VM (AHV) on page 551 or Managing a VM (ESXi) on page 581 for instructions on how to perform these actions.

Policies Tab

The **Policies** tab displays **Affinity Policies** (see Affinity Policies Defined in Prism Central on page 599) and **NGT Policies** (see NGT Policies on page 602).

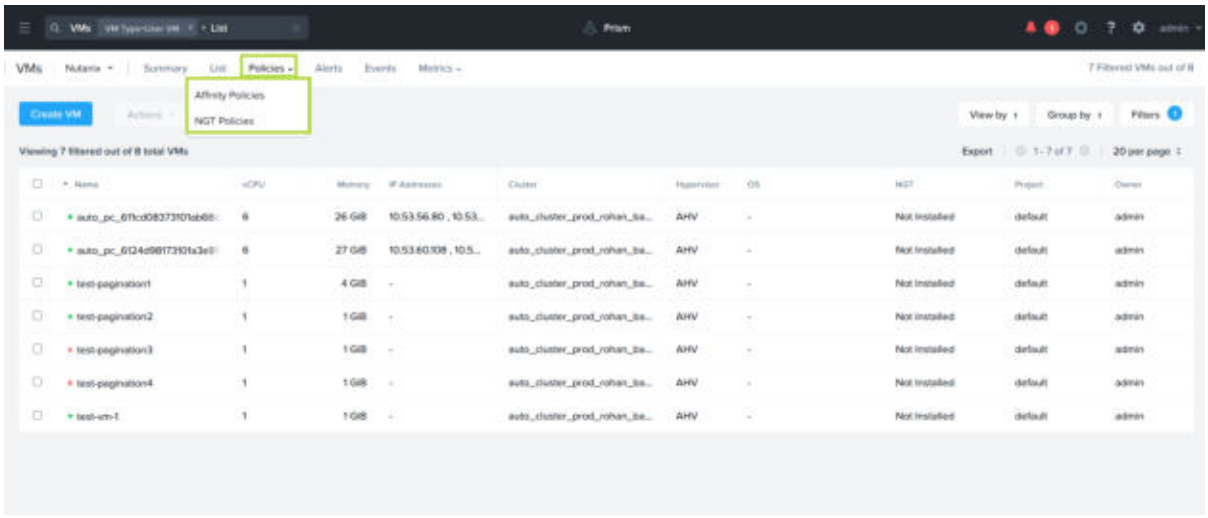

Figure 56: VMs Policies Tab

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just VM-related alerts across the registered clusters (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just VM-related events across the registered clusters (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the VMs. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant performance information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

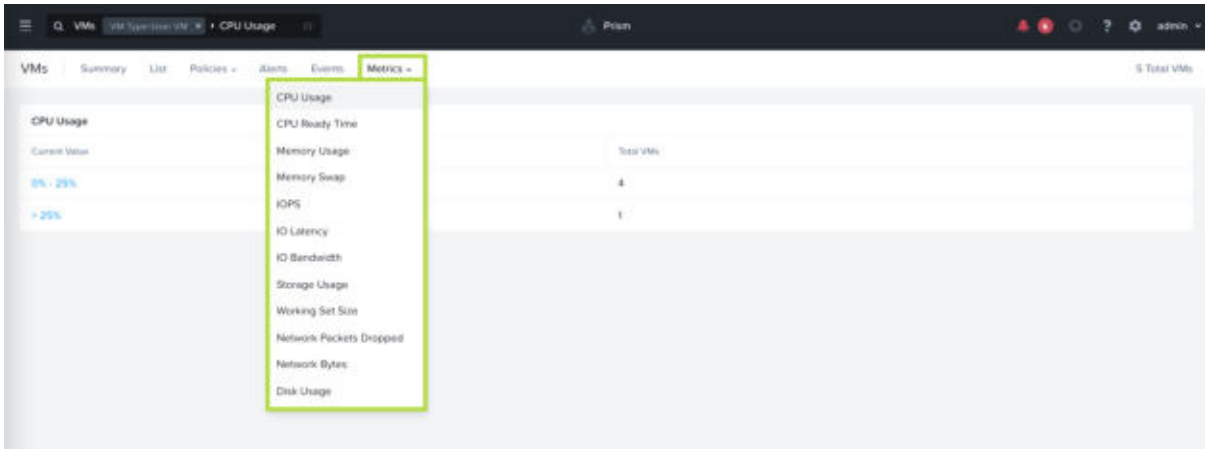

Figure 57: VMs Metrics Tab

Table 13: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total VMs (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking a percentile interval displays the Summary tab filtered to just those VMs. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
CPU Ready Time	Displays a CPU ready time percentage usage table.
Memory Usage	Displays a memory percentage usage table.
Memory Swap	Displays memory swap-out and swap-in rate tables.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate tables.
Usage	Displays total, snapshot, and shared storage size tables.
Working Set Size	Displays total, snapshot, and shared working set size tables. Working set is the amount of memory that a VM requires in a given time interval.
Network packets dropped	Displays tables for the number of transmitted and received packets dropped.
Network bytes	Displays tables for the amount of transmitted and received bytes (in GiBs).

Metric	Description
Disk Usage	Displays a disk usage table listing current values and total VMs (number). The current values are split into percentile intervals (for example, l3.68% - 4.1%, more than 4.93%). Clicking a percentile interval displays the Summary tab filtered to just those VMs. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.

VM Details View

To access the details page for a VM, go to the VMs **List** tab (see VMs Summary View on page 86) and click the VM name. You can also access the details page by clicking the VM name wherever that name appears, such as in a dashboard widget or search result.

The VM name and the following set of tabs appear on the left: Summary, Console, Data Protection, Alerts, Events, Metrics, NICs, Disks, Snapshots, Categories, and Apps & Relationships. Click a tab to display that information on the right. (Click the **Back to VMs** link to return to the VMs summary view.)

Note: VirtIO must be installed in a VM for AHV to display correct VM memory statistics.

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the VM (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.
- A **Storage Properties - Associations** widget that displays the storage properties of the VM as defined in the storage policy associated with the VM and the other entities like Virtual

Disks, Storage Containers, Categories and Storage Policies associated with the VM. See the Storage Properties Fields and the Associations Fields tables.

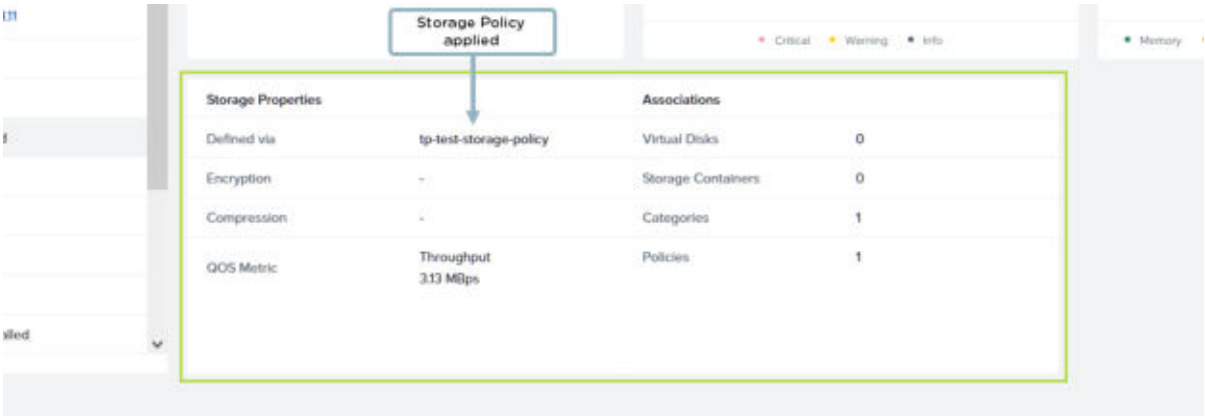

Figure 58: VM Storage Properties - Associations Widget

The widget displays the actual values for Encryption and Compression that is applied to the VM. For the values that are displayed in the Storage Policy dashboard as **Inherited from Cluster** also the **Storage Properties Association** widget displays the actual inherited values.

- Action button (above the widgets). Click the appropriate button to run that administrative action on the VM. The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM. See Managing a VM (AHV) on page 551 for instructions on how to perform each action.

  Note: You can perform administrative actions on VMs in Acropolis-managed clusters only. Unavailable actions are grayed out.

- 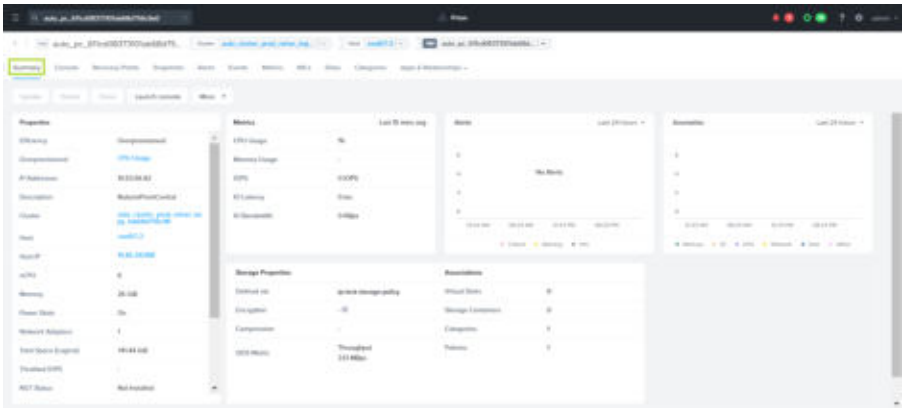

Figure 59: VM Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned.

Table 14: VM Properties Fields

Parameter	Description	Values
Efficiency	Displays the efficiency state for this VM. If the efficiency is not good, an additional field may appear that specifies the problem. For example, if the VM is constrained, a Constrained field appears that identifies the constrained resource such as the CPU or memory. Note: The Anomalies, Efficiency, Over provisioned, and Constrained parameters relate to the VM behavioral learning feature. See Behavioral Learning Tools on page 838 for a description of each parameter.	Bully, Over Provisioned, Constrained, Inactive VM, Good
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Host	Displays the host name. This field may be blank if the VM is powered off and a host is not assigned.	(host name)
Host IP	Displays the host IP address.	(IP address)
Virtual CPU Count	Displays the number of virtual CPUs assigned to this VM.	(number)
Memory Capacity	Displays the amount of memory available to this VM.	xxx [MB|GB]
IP Addresses	Displays one or more IP addresses assigned to the VM.	(IP address)
Power State	Displays whether the VM is powered on or powered off	On, Off
Network Adapters	Displays the number of network adapters available to this VM.	(# of adapter ports)
Disk Capacity	Displays the total disk capacity available to this VM.	xxx [GB|TB]
NGT Status		
Services Enabled	Displays the services enabled for this VM.	Calm, Karbon
NGT Installed Version		
NGT Cluster Version		
(the following fields appear when the VM is allocated to a GPU)		
GPU Type	Displays the GPU operational mode. If it is vGPU, the following fields also appear.	vGPU, Passthrough, None
GPU Configuration	Displays the vGPU profile used.	(vGPU profile name)

Parameter	Description	Values
Framebuffer	Displays the size of the GPU framebuffer (RAM).	xxx GiB
Virtual Slice	Displays the virtual slice applied. The "virtual slice" reflects the approximate amount of physical GPU resources that the vGPU can receive. Note: The Virtual Slice and vGPU Guest Driver Version fields do not appear for passthrough GPUs.	(slice amount)
vGPU Guest Driver Version	Displays the version number of the vGPU guest driver.	(version number)

The following table describes the fields in the **Storage Properties** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. For more information about the parameters, seeStorage Policy Details View on page 135.

Table 15: Storage Properties Fields

Parameter	Description	Values
Defined via	Displays the name of the storage policy that the VM is associated with.	Storage policy name
Encryption	Displays the status of encryption that the storage policy applies to the VM.	Enabled or Inherit from Cluster
Compression	Displays the status and type of compression that the storage policy applies to the VM.	Inline, Post Process, Off or Inherit from Cluster
QOS Metric	Displays the type of QOS metric that the storage policy applies to the VM with the throttled throughput value.	Throttled Throughput value in IOPS or MBps.

The following table describes the fields in the **Associations** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned.

Table 16: Associations Fields

Parameter	Description	Values
Virtual Disks	Displays the number of virtual disks associated with the VM.	(number)

Parameter	Description	Values
Storage Containers	Displays the number of storage containers associated with the VM.	(number)
Categories	Displays the number of categories associated with the VM.	(number)
Policies	Displays the number of policies like storage policies and Image Placement policies that are associated with the VM.	(number)

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget allows quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective pull-down menus.

Example: Displaying Relationship between Entities

The following image indicates that VM auto_pc_611cd08373101ab68d756c9a0 resides on host cool-07-2 and cluster auto_cluster_prod_rohan_bajaj_1ab68d756c98.

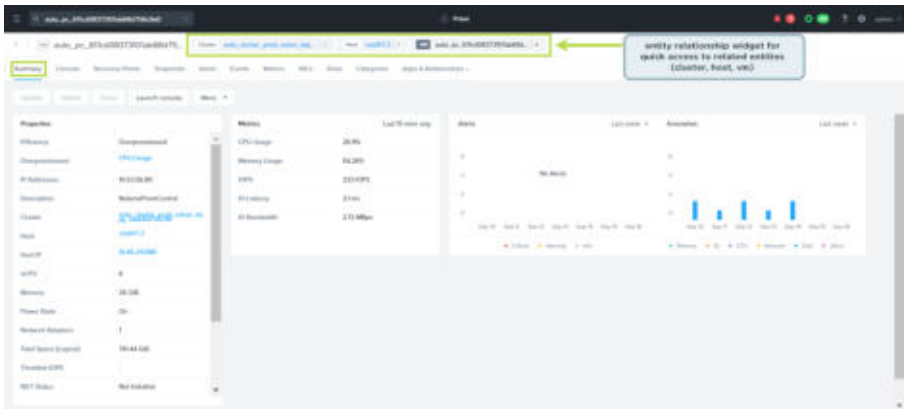

Figure 60: VM Summary: Entity Relationship Widget

Example: Viewing VMs on a Cluster

Click the **VM** pull-down menu to view the list of VM instances on the selected cluster. Alternatively, you search the VM instance name residing on the target cluster.

Figure 61: VMs on a Cluster

Example: Viewing VM Instances on a Host

Click the **VM** pull-down menu to view the list of VM instances on the selected host. Alternatively, you can search the VM instance name belonging to the target host.

Figure 62: VMs on a Host

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Console Tab

The **Console** tab displays the VM console screen. There are three icons above the console display (on the right).

- Click the left (three small boxes) icon to send a **Control-Alt-Delete** command to the console.
- Click the middle (camera) icon to take a screen shot of the console display.
- Click the right (box with an arrow in it) icon to open the console in a new window.

Recovery Points

The **Recovery Points** tab displays a list of recovery points (backup snapshots) when backups have been enabled. (The list is blank if there are no snapshots available.) The total number of recovery points and the latest and oldest recovery points are listed on the left. A list of all recovery points appears in a table on the right with the create time, location, expiry time, and recovery point type provided for each recovery point.

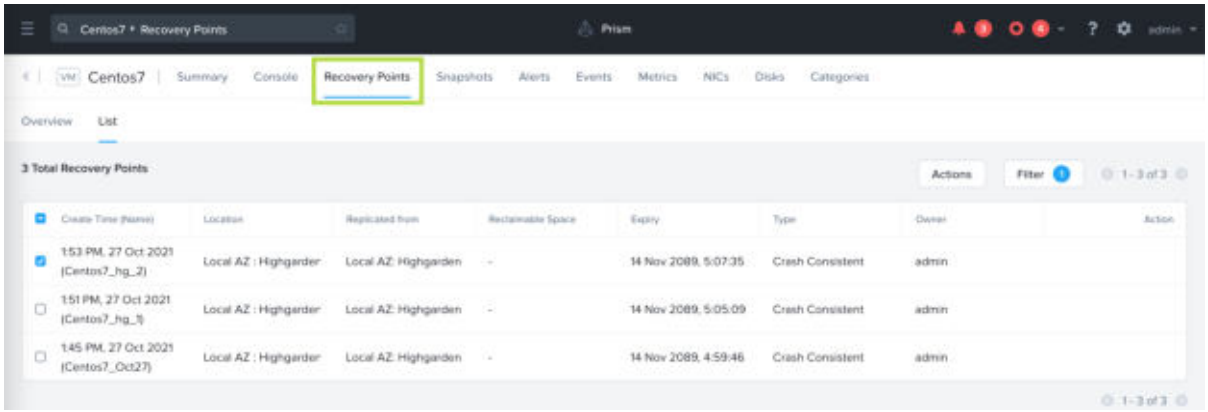

Figure 63: Recovery Points Tab

You can restore or replicate a VM from a recovery point. Select the recovery point and click the required action from the **Actions** menu:

- Click **Replicate** to replicate a VM from the selected recovery point either locally or remotely in a state of a chosen recovery point.
- Click **Restore** to restore the VM from the selected recovery point.
- Click **Delete** to delete the recovery point.

Snapshots Tab

The **Snapshots** tab displays information in tabular form about backup snapshots of the VM taking from Prism Element. Each line represent a snapshot, and the following information is displayed for each snapshot:

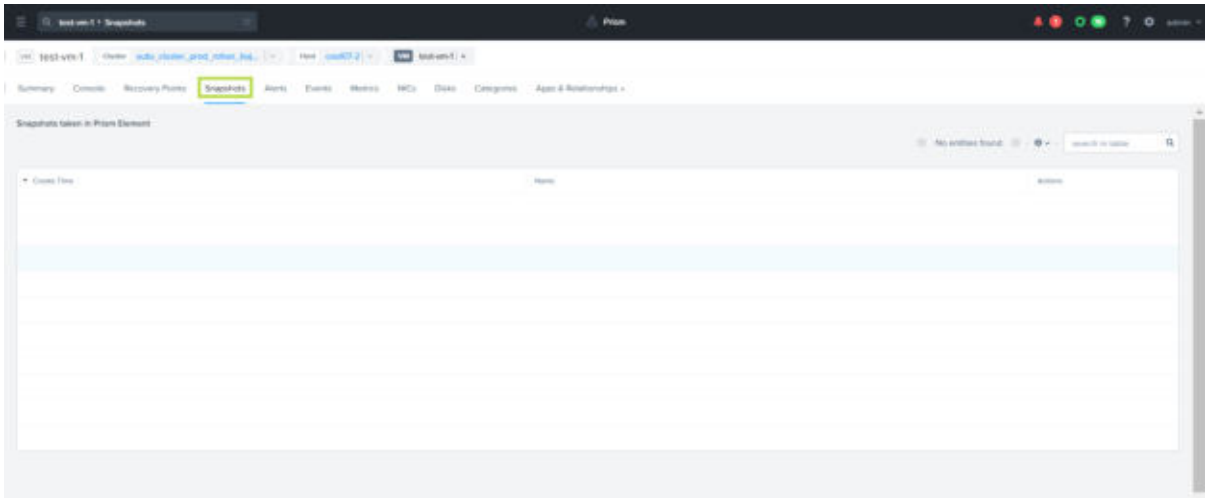

Figure 64: Snapshot Tab

- **Create Time**. Displays the time the backup snapshot was created (completed).
- **Name**. Displays a name for the backup if one was created.
- **Action**. Displays four action links:
  - Click the **Details** link to open a window that displays the snapshot details.

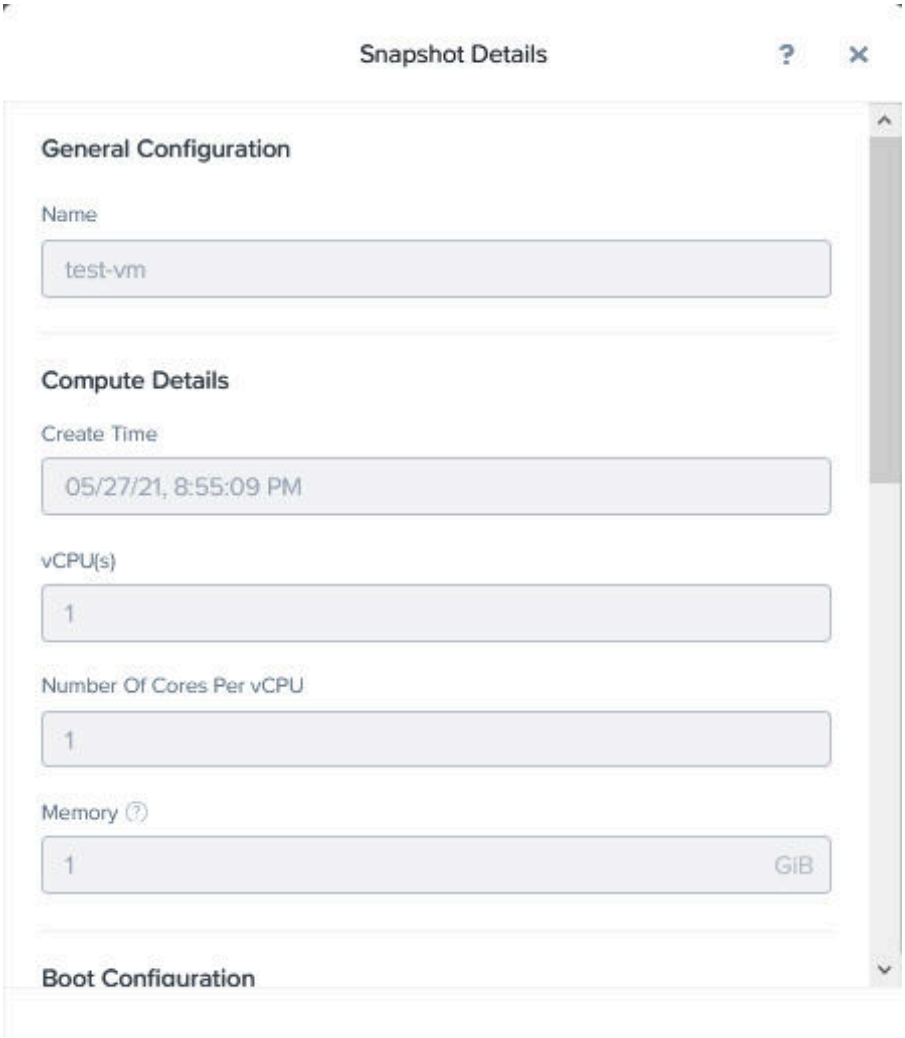

Figure 65: Snapshot Details Window

  - Click the **Clone** link to clone a VM from the snapshot.
  - Click the **Restore** link to restore the VM from the snapshot. This restores the VM back to the state of the selected snapshot.
  - Click the **Delete** link to delete the snapshot.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this VM (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this VM (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the VM. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 838). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check one or more appropriate boxes to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 282).

The following table describes the available metrics. (Some of these metrics are not available on all hypervisors.)

Table 17: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the VM (0–100%).
CPU Ready Time	Displays the current, high, and low percentage of CPU wait time (0–100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the VM (0–100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
IO Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
Usage	Displays separate graphs for current, snapshot, and shared storage usage (in GiBs) by the VM.

Metric	Description
Working Set Size	Displays separate graphs for total, write, and read storage usage (in GiBs) for the VM working set size.
Network Packets Dropped	Displays separate graphs for the number of transmitted and received packets dropped.
Network Bytes	Displays separate graphs for the amount of transmitted and received bytes (in GiBs).

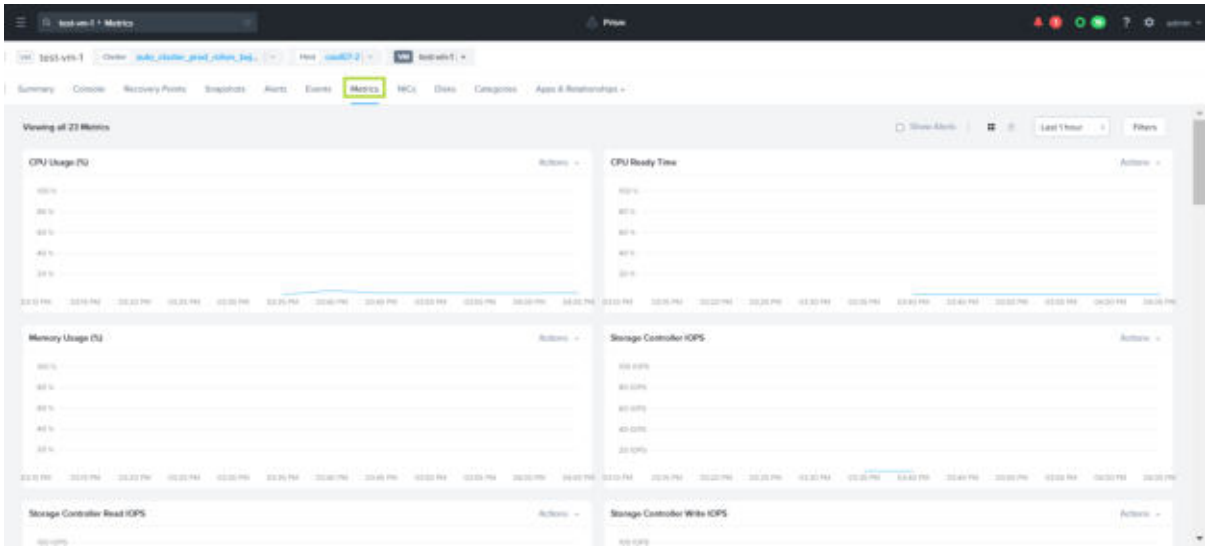

Figure 66: Metrics Tab: CPU Usage

NICs Tab

The **NICs** tab displays information in tabular form about the virtual NICs in the VM. Each line represent a virtual NIC, and the following table describes the fields.

Table 18: NIC Fields

Parameter	Description	Values
VLAN ID	Displays the VLAN name for this NIC.	(VLAN ID
MAC Address	Displays the virtual NIC MAC address.	(MAC address)
Network Connection State	Displays whether the NIC is connected to the network currently.	Connected, Disconnected
Requested IP Address	Displays the virtual NIC IP address	(IP address)
Action	Displays the available actions you can execute on this NIC.	Delete

Disks Tab

The **Disks** tab displays information in tabular form about the virtual disks in the VM. Each line represents a virtual disk, and includes the following fields.

- **Disk Address**: Displays the disk address (such as ide.0 or scsi.1).
- **Capacity**: Displays the disk capacity (in MiB or GiB).

Categories Tab

The **Categories** tab displays the categories and image placement policies associated with the VM. Each VM can have a one-to-many relationship with categories and the categories can have a many-to-one relationship with image placement policies.

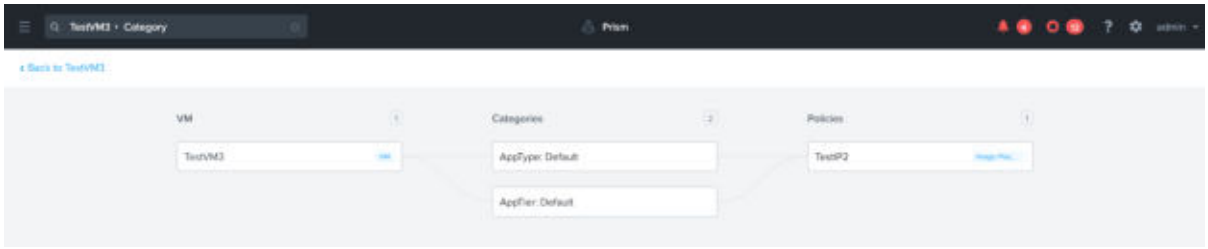

Figure 67: Categories Tab

For more information about categories management and image placement policies, see Category Management on page 735 and Image Placement Policies on page 680.

Apps & Relationships Tab

The **Apps & Relationships** tab displays a pull-down menu with two menu options: **Discovered Apps** and **App Relationships**.

The **Discovered Apps** menu option displays (up to) the top 20 applications that are communicating with other entities (usually client VMs). The discovered applications are sorted based on the number of clients communicating with the application. This menu option provides similar information as the application discovery dashboard, except it is filtered to display the details only for this VM (see Application Discovery View on page 772).

The **App Relationships** option displays three tabs: Visualization, Incoming List, and Outgoing List. All the three tabs display (up to) the top 20 results with respect to the incoming or outgoing client communication.

Visualization: Displays a visual representation of the VM and the applications running on it, along with its incoming and outgoing communication with other entities. Data for last 24 hours is fetched to display the information on the visualization tab. You can hover on an entity to see the applications it is communicating with. Hover on an application to see the application's incoming communication with external entities and outgoing communications of the VM on which the application is running.

> Note: In case of dual NIC VMs, click and expand the App listing to view the specific IP that is being used by the application.

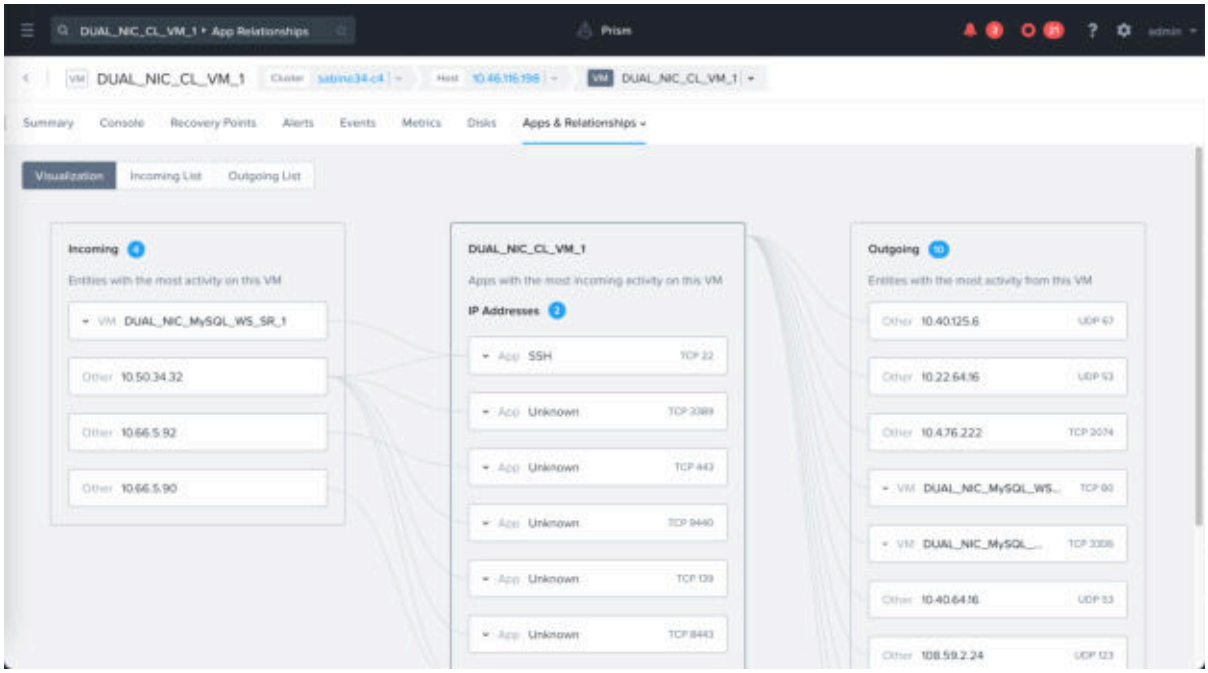

Figure 68: Visualization Tab

Incoming List: Displays a list of entities having incoming communication with this VM. Each line represents an incoming communication and includes the following fields.

Table 19: Incoming List Fields

Parameter	Description	Values
Source Entity Name	Displays the name of the communicating entity. If the entity is not identified by Prism Central (through registered PE or through vCenter using monitoring configurations feature), the name is displayed as unknown.	(VM name), Unknown
Source IP Address	Displays the IP Address of the communicating entity.	(IP Address)
Destination App	Displays the name of the application that is being communicated to. If the application is not an identified one, the Destination App is displayed as Unknown.	(application name), Unknown
Destination IP Address	Displays the IP Address of the VM.	(IP address)
Destination Port	Displays the port used by the application for communication.	(port number)
Destination App Identification Status	Indicates whether the application is an identified application type. An identified application is one that is a known type.	Identified, Unidentified

You can filter this list by opening the Filter pane to select a filter option. The following table describes the filter options for **Incoming List**. You can apply filters across multiple parameters.

The filters fetch the latest discovery and communication data and accordingly display the result set, to the maximum of 20 results.

Table 20: Filter Pane Fields

Parameter	Description	Values
Source Entity Name	Filters on the source entity name. Select a condition from the pull-down list (**Contains**, **Equal to**, **Not equal to**, **Doesn't contain**, **Starts with**, **Ends with**) and enter a string in the field.	(source entity name string)
Entity Type	Filters on the entity type. A VM belongs to a cluster managed by Prism Central or vCenter (only if using monitoring configurations feature). An Unknown entity is one that is not identified by Prism Central.	VM, Unknown
Source IP Address	Filters on the source IP address.(same options as Source Entity Name).	(source IP address string)
Destination App	Filters on the destination app (same options as Source Entity Name).	(Destination App string)
Destination IP Address	Filters on the destination IP address (same options as Source Entity Name).	(destination IP address string)
TCP Port	Filters on the TCP port number.	(TCP port number)
UDP Port	Filters on the UDP port number.	(UDP port number)
Time Range	Filters on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

Outgoing List: Displays a list of entities having outgoing communication from this VM. Each line represents an outgoing communication and includes the following fields.

Table 21: Outgoing List Fields

Parameter	Description	Values
Destination Entity Name	Displays the name of the entity this VM is communicating to. If the entity is not identified by Prism Central (through registered PE or through vCenter using monitoring configurations feature), the name is displayed as unknown.	(entity name), Unknown
Destination IP Address	Displays the IP Address of the entity this VM is communicating to.	(IP address)
Destination Port	Displays the port on which this VM is communicating with the destination entity.	(port number)

Parameter	Description	Values
Source IP Address	Displays the IP Address of the VM.	(IP address)
Source App Identification Status	Indicates whether the communicating application is an identified application type. An identified application is one that is a known type.	Identified, Unidentified

You can filter this list by opening the Filter pane to select a filter option. The following table describes the filter options for **Outgoing List**. You can apply filters across multiple parameters. The filters fetch the latest discovery and communication data and accordingly display the result set, to the maximum of 20 results.

Table 22: Filter Pane Fields

Parameter	Description	Values
Destination Entity Name	Filters on the destination entity name. Select a condition from the pull-down list (**Contains**, **Equal to**, **Not equal to**, **Doesn't contain**, **Starts with**, **Ends with**) and enter a string in the field.	(destination entity name string)
Entity Type	Filters on the entity type. A VM belongs to a cluster managed by Prism Central or vCenter (only if using monitoring configurations feature). An Unknown entity is one that is not identified by Prism Central.	VM, Unknown
Destination IP Address	Filters on the destination IP address. (same options as Destination Entity Name).	(destination IP address string)
TCP Port	Filters on the TCP port number.	(TCP port number)
UDP Port	Filters on the UDP port number.	(UDP port number)
Source IP Address	Filters on the source IP address (same options as Destination Entity Name).	(source IP address string)
Time Range	Filters on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

VM Template Summary View

To access the templates dashboard, select **Compute & Storage** > **Templates** from the entities menu (see Entities Menu on page 13). The templates dashboard summary view displays information about templates across the registered clusters and allows you to access detailed information about each template.

> Note: This section describes the information and options that appear in the templates dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.

- See Creating a VM Template on page 605 for information about creating the templates.
- See Managing a VM Template on page 609 for information about managing the templates.

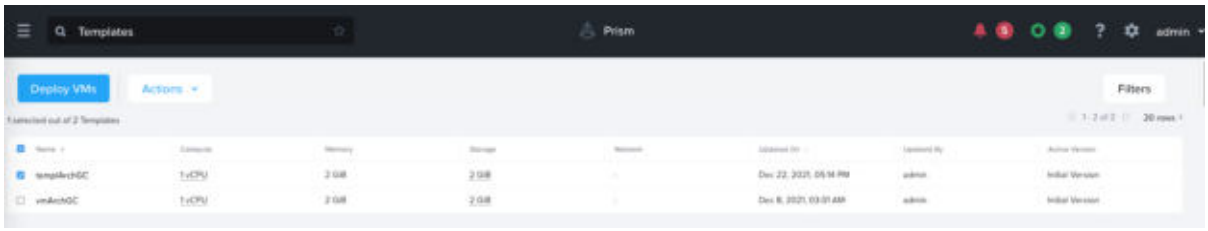

Figure 69: Templates View

Table 23: Template fields

Parameter	Description	Values
Name	Displays the template name. Clicking the name displays the details page for that template (see VM Template Details View on page 111).	(Template name)
Compute	Displays the virtual CPU count of the VM that is deployed using this template.	(vCPU count)
Memory	Displays the total amount of memory available to the VM that is deployed using this template.	xxx [MB|GiB]
Storage	Displays total amount of storage available to the VM that is deployed using this template.	xxx [MB|GiB]
Network	Displays the NIC count of the VM that is deployed using this template.	(NIC count)
Updated On	Displays the last date and time when the template has been updated.	(date and time)
Updated By	Displays the user who most recently updated the template.	(user)
Active Version	Displays the name of the active version of the template. An active version is the version of the template that by default gets deployed when you click the **Deploy VMs** button.	(version)

You can filter the template list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range. Click the plus (+) sign to the right of the parameter name to expand the fields for that parameter; click the minus (-) sign to contract those fields.

Table 24: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the template name. Select a condition from the pull-down list and enter a string in the field. It will return a list of templates that satisfy the template name condition/string. In this field, the condition menu options are **Contains**, **Doesn't contain**, **Starts with**, **Ends with**, and **Equal to**.	(template name string)

You can sort the list of templates by clicking on **Name** and **Updated On** columns.

The templates view includes an action button **Deploy VMs** along with the **Actions** menu. To deploy VMs using the VM template, click the **Deploy VMs** button. By default the active version of the template gets deployed.

The **Actions** menu appears when one or more templates are selected. It includes the following actions:

- Actions that can be applied to multiple templates: **Delete Template**.
- Actions that can be applied to only one template at a time: **Update Guest OS**, **Complete Guest OS Update**, **Cancel Guest OS Update**, **Update Configuration**, and **Delete Template**.

The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the selected template. See Managing a VM Template on page 609 for instructions on how to perform these actions.

VM Template Details View

To access the details page of a template, go to the templates view (see VMs Summary View on page 86) and click the template name. You can also access the details page by clicking the template name wherever that name appears, such as in a dashboard widget or search result.

The templates details page includes two views: **Summary** view and **Versions** view.

Summary View

The **Summary** view includes an action button **Deploy VMs** along with the **Actions** menu. There are only two options available in the **Actions** menu under this view: **Set as Active** and **Delete**. You can select a version and set it as active, or delete a selected version. You cannot delete an active version of the template.

The following set of widgets appears on the details page: **Template**, **Active Version**, **VM Properties**, and **VM Resources**. Click the **Back to Templates** link to return to the template summary view.

- **Template**: The widget displays the following:
  - **Available Versions**: Number of versions available
  - **Description**: User-provided description of the template
  - **Last Updated by**: User who has last updated the template
  - **Last Updated on**: Date and time of last update to the template

- **Active Version Configuration**: The widget displays the following:
  - **Version Name**: Name of the currently active version of the template
  - **Change Notes**: User-provided change notes for the active version
- **VM Properties**: The widget displays the following:
  - **Compute**: Compute resources of the VM such as CPU, core per CPU, and memory
  - **Boot Type**: Boot type of the VM such as legacy or UEFI
  - **Guest OS**: Guest OS of the VM. If the source VM of the template does not have NGT installed on it, then this field is displayed as unknown.
  - **Guest Customization**: Status of the guest customization application (applied/not applied)
  - **NGT Status**: Status of the NGT installation
- **VM resources**: The widget displays the following:
  - **GPU**: Number of GPU assigned to the VM
  - **Disks**: Number of disks and total disk capacity assigned to the VM
  - **Network**: Number of NICs assigned to the VM

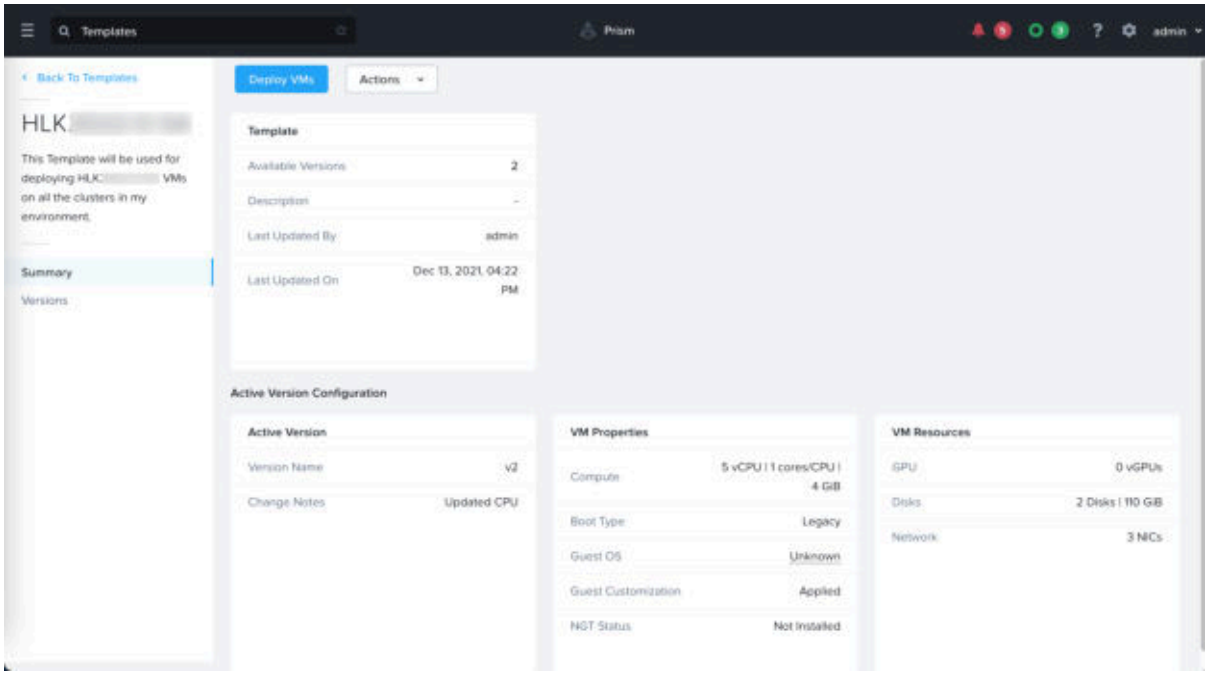

Figure 70: Summary View of a Template

Versions View

The **Versions** view includes an action button **Deploy VMs** along with the **Actions** menu. There are only two options available in the **Actions** menu under this view: **Set as Active** and **Delete**. You can select a version and set it as active, or delete a selected version. You cannot delete an active version of the template.

Table 25: Versions Fields

Parameter	Description	Values
Name	Displays the version name. This column also indicates if the version is an active version.	(template name)
Compute	Displays the virtual CPU count of the VM that gets deployed using this version.	(vCPU count)
Memory	Displays the total amount of memory available to the VM that gets deployed using this version.	xxx [MB|GiB]
Storage	Displays total amount of storage available to the VM that gets deployed using this version.	xxx [MB|GiB]
Network	Displays the NIC count of the VM that gets deployed using this version.	(NIC count)
Added On	Displays the last date and time when the version was added.	(date and time)
Added By	Displays the user who added the version.	(user)
Notes	Displays the user-specified change note for the version.	(version)

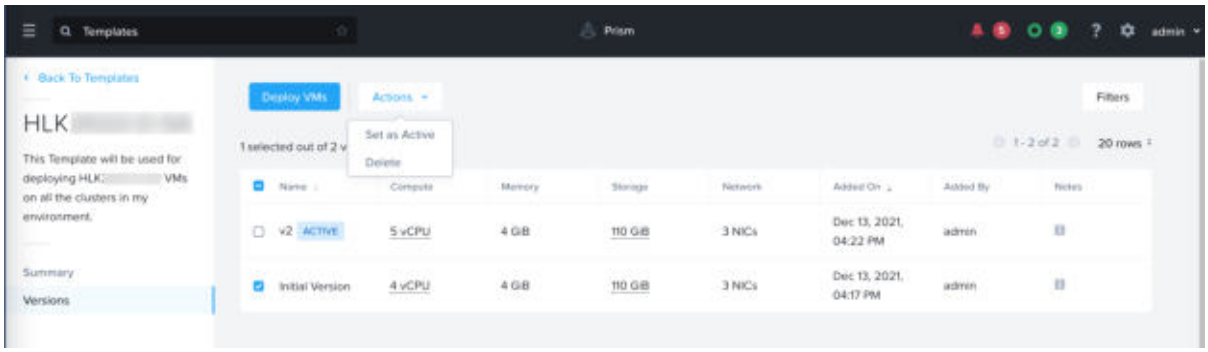

Figure 71: Versions View of a Template

You can sort the list of versions by clicking on **Name** and **Added On** columns.

OVAs View

To access the OVAs dashboard, select **Compute & Storage** > **OVAs** from the entities menu (see Entities Menu on page 13). The OVAs dashboard allows you to view information about the OVAs available in the

Note: This section describes the information and options that appear in the OVAs dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- See OVA Management on page 688 for information about the OVA operations.

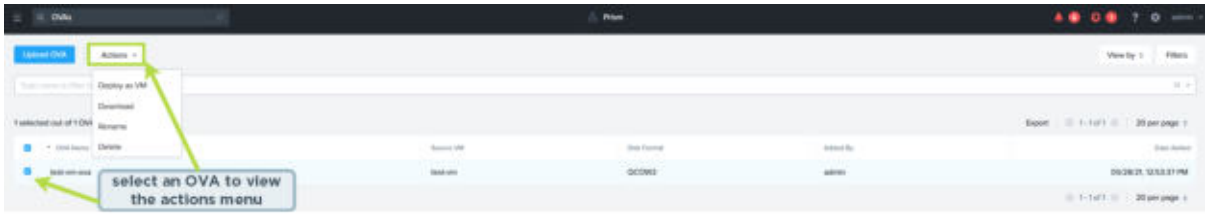

Figure 72: OVAs Dashboard

The following table describes the fields that appear in the catalog items list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 26: OVA List Fields

Parameter	Description	Values
OVA Name	Displays the name of the OVA file.	(file name)
Source VM	Displays the name of source VM that was exported as an OVA. The alert indicators for the task status are displayed beside the source VM for which the task is run. See the Alert indicator table for more information.	(file name)
Disk Format	Displays the format of the disks in the OVA.	[QCOW2|VMDK]
Added by	Indicates which user added this OVA.	(text string)
Date added	Displays the date and time when the OVA was added.	(date in mm/dd/yyyy format, time in hr:min:sec [AM|PM] format) Example: 07/30/20, 3:00:20 PM

Table 27: OVA Alerts

OVA Alerts	Status Message sample	Indicates...
	`Incorrect file uploaded. Multiple OVA files present in OVA.`	That the upload process failed during validation. You must fix the issue in the OVA file or upload the correct, compatible OVA file.

Figure 73:

	`Upload for OVA file has not started. Please start OVA file upload using Resume Upload option.`	That there is an interruption in one of the three sub-tasks in the upload process. You can resume the upload. See the *Resume Upload* action procedure.
	`Uploaded OVA file has not been verified. Please start verification using Resume Upload option.`	

A sample of the error hover message is as follows:

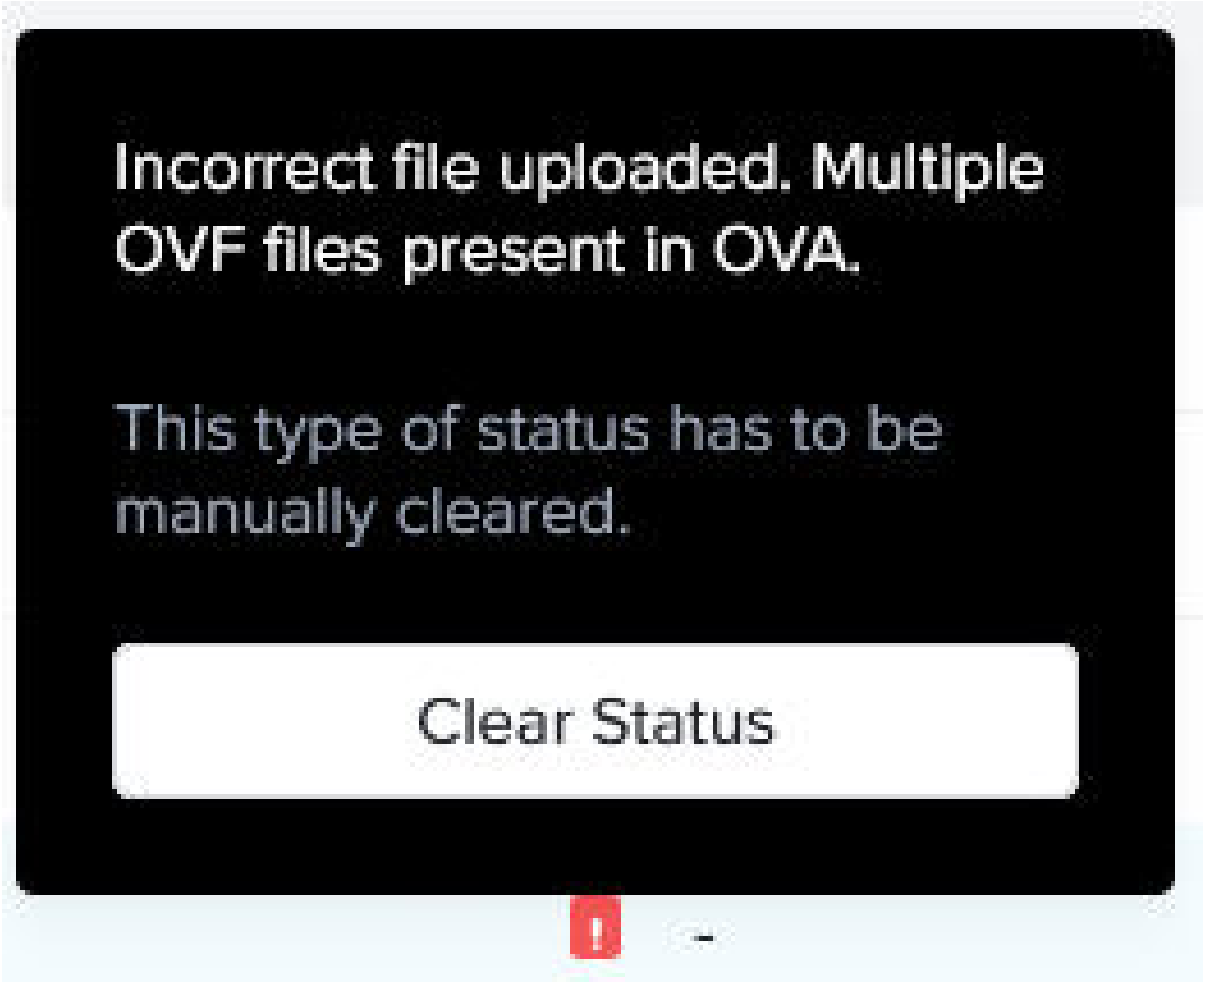

Figure 76: Sample error hover message

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

For each filter, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.

Table 28: Filter Pane Fields

Parameter	Description	Values
OVA Name	Filters on the OVA file name. It returns a list of OVAs that satisfy the name condition/string.	(OVA file name string)
Source VM	Filters on the Source VM name. It returns a list of OVAs for which the Source VM field satisfies the condition/string.	(Source VM name string)

Parameter	Description	Values
Disk Format	Filters on the disk format. Check the box(es) for the OVAs of the required disk format.	VMDK, QCOW2

The **Actions** menu appears when a OVA is selected. It allows you to:

- Deploy an OVA file as a VM.
- Download an OVA file to your local machine.
- Rename an OVA file.
- Delete an OVA file.
- Resume upload when an upload is interrupted and shows an error.

> Note: The resume upload action is not available in case the concatenate process of the upload is interrupted. In such a case, you need to run the Concatenate API from the REST API Explorer to resume the upload.

Images Summary View

To access the images dashboard, select **Compute & Storage** > **Images** from the **Entities** menu (see Entities Menu on page 13). The images dashboard allows you to view information about images available from this Prism Central instance.

> Note: This section describes the information and options that appear in the images dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See Image Management on page 661 for information about adding and managing images through Prism Central.
> - See Image Placement Policies on page 680 for information on AHV affinity policies.
> - See Bandwidth Throttling Policies on page 685 for information on how to create and apply NGT-related policies.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the Images.

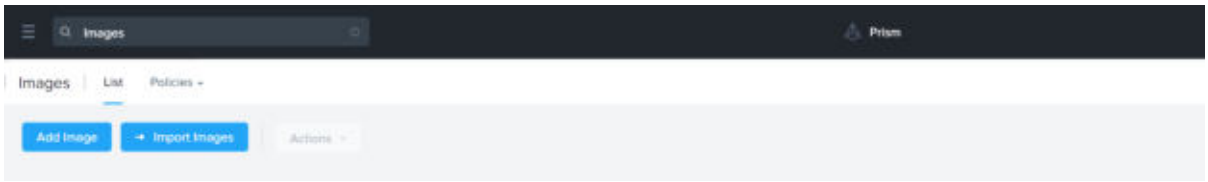

Figure 77: Images List View

The following table describes the fields that appear in the images list. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 29: Image List Fields

Parameter	Description	Values
Name	Displays the image name. Clicking the name displays the details page for that image (see Image Details View on page 119).	(name)
Description	Indicates which user uploaded this image.	(text string)
Type	Displays the image type.	ISO, Disk
Size	Displays the image size.	xxx [MB|GB]
Creator	Displays who created the image.	(name)

You can filter the images list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 30: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of storage containers that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It will return a list of images whose description field satisfy the condition/string.	(description string)
Type	Filters on the image type. Check the box(es) for the desired image types.	Disk, ISO

You can group the images list in the following ways:

- The **Group** pull-down menu allows you to group by type (disk or ISO).
- The **View by** pull-down menu allows you to view the entries by categories, creator, description, size, and type.

The images dashboard includes the following action buttons:

- **Add Image** Click to add an image (see Creating an Image on page 662).
- **Import Images** Click to import images from registered clusters to Prism Central (see Importing Images to Prism Central on page 679.

The **Actions** menu appears when one or more images are selected. It includes delete, update, and add image to catalog options (see Modifying an Image on page 678). The available

actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Policies Tab

The **Policies** tab displays the **Placement Policies** (see Image Placement Policies on page 680) and **Bandwidth Throttling Policies** (see Bandwidth Throttling Policies on page 685).

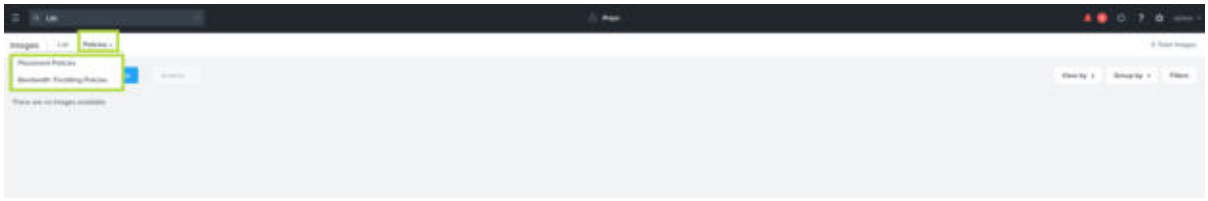

Figure 78: Images Policies Tab

Image Details View

You can view the detailed information about an image by going to the details page for that image.

To access the details page for an image, go to the images dashboard (see Images Summary View on page 117) and click the image name. An image details page includes the following:

- Image name (upper left). You can switch from one image to another by selecting a different image name from the pull-down list.
- Image information (middle left). The information includes the image description, type, size, and creator name (see Creating an Image on page 662).
- Tab options (lower left). Select the **Location** (default) or **Policies** tab to display that tab content to the right.
- Action buttons (upper right). When the window size is not large enough, some actions appear under a **More** menu button.
  - To delete the image, click the **Delete** button
  - To modify the name, description, or type, click the **Update Image** button (see Modifying an Image on page 678).
  - To add the image to the catalog, click the **Add Image to Catalog** button (see Adding a Catalog Item on page 658).
  - To assign the image to a category, click the **Manage Categories** button (see Assigning a Category on page 737).
  - To open a help page in a separate tab or window, click the question mark icon.
  - To close the details page, click the **X** icon.

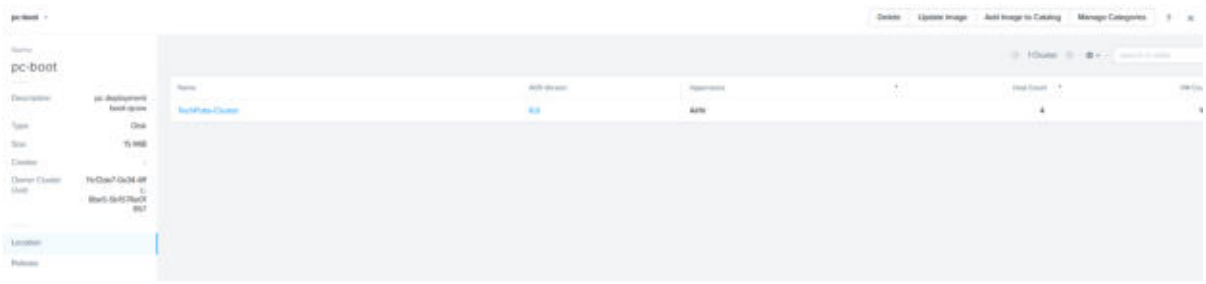

Figure 79: Image Details View (Location Tab)

Location Tab

The following table describes the information displayed in the **Location** tab, which appears by default when you first open the page.

Table 31: Location Fields

Parameter	Description	Values
Name	Displays the name of the cluster in which the image is stored. Each row represents another cluster that contains the image.	(cluster name)
AOS Version	Displays the AOS version running in the cluster.	(AOS version number)
Hypervisor	Displays the name of the hypervisor running in the cluster.	AHV, ESXi, Hyper-V
Host Count	Displays the number of hosts in the cluster.	(integer)
VM Count	Displays the number of VMs in the cluster	(integer)

Policies Tab

The **Policies** tab lists the policies that apply to the image (see Image Placement Policies on page 680). The following table describes the information displayed in the **Policies** tab.

Table 32: Policies Fields

Parameter	Description	Values
Policy Name	Displays the policy name.	(cluster name)
Status	Indicates whether the image is compliant according to the policy.	compliant, noncompliant
Details	Lists the clusters in which the policy is enforced.	(cluster names)

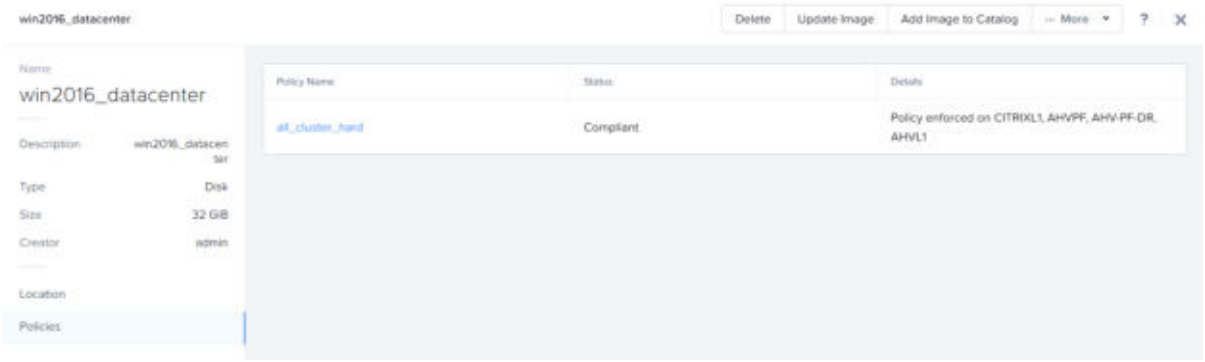

Figure 80: Image Details View (Policies Tab)

Catalog Items View

To access the catalog items dashboard, select **Compute & Storage** > **Catalog Items** from the **Entities** menu (see Entities Menu on page 13). The catalog items dashboard allows you to view information about the images and VMs in the catalog.

> Note: This section describes the information and options that appear in the catalog items dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See Catalog Management on page 657 for information about the catalog service and how to add (or delete) items.

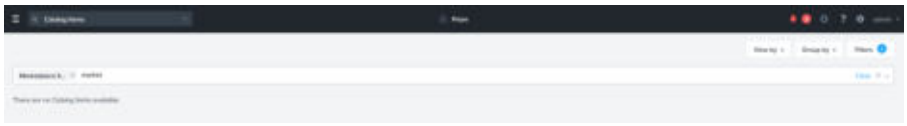

Figure 81: Catalog Items Dashboard

The following table describes the fields that appear in the catalog items list. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 33: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the item name.	(item name)
Type	Displays the type of item (image or VM).	Image, Vm
Description	Indicates which user added this item to the catalog.	(text string)

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 34: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of catalog items that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of catalog items whose description field satisfy the condition/string.	(description string)
Type	Filters on the item type. Check the box(es) for the desired item types.	Vm, Image

You can group the catalog items list in the following ways:

- The **Group** pull-down menu allows you to group by type (VM and image).
- The **View by** pull-down menu allows you to view by name, description, message, percent, reason, state, and type.

The **Actions** menu appears when one or more catalog items are selected. It allows you to delete the selected catalog items (see Deleting a Catalog Item on page 660).

Storage Containers Summary View

To access the storage containers dashboard, select **Compute & Storage** > **Storage Containers** from the entities menu (see Entities Menu on page 13). The storage containers dashboard summary view displays information about storage containers across the registered clusters and allows you to access detailed information about each storage container. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> Note: This section describes the information and options that appear in the storage containers dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities**: Displays a list of the storage containers with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that storage container. Click the **View all XX Storage Containers** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of storage container-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then

displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

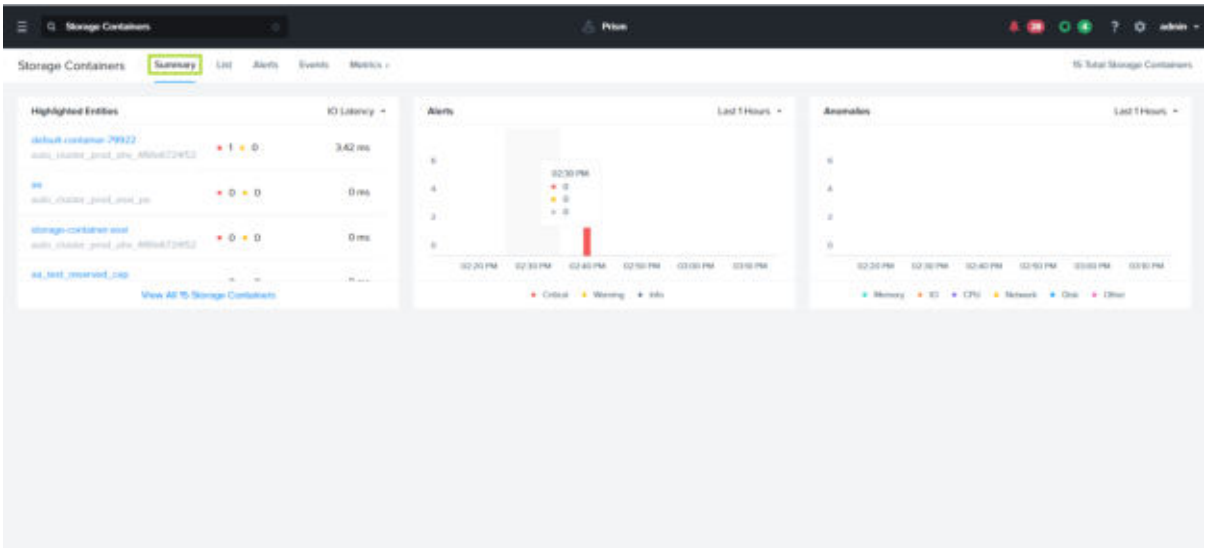

Figure 82: Storage Containers Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the storage containers across the registered clusters. The following table describes the fields that appear in the storage containers list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

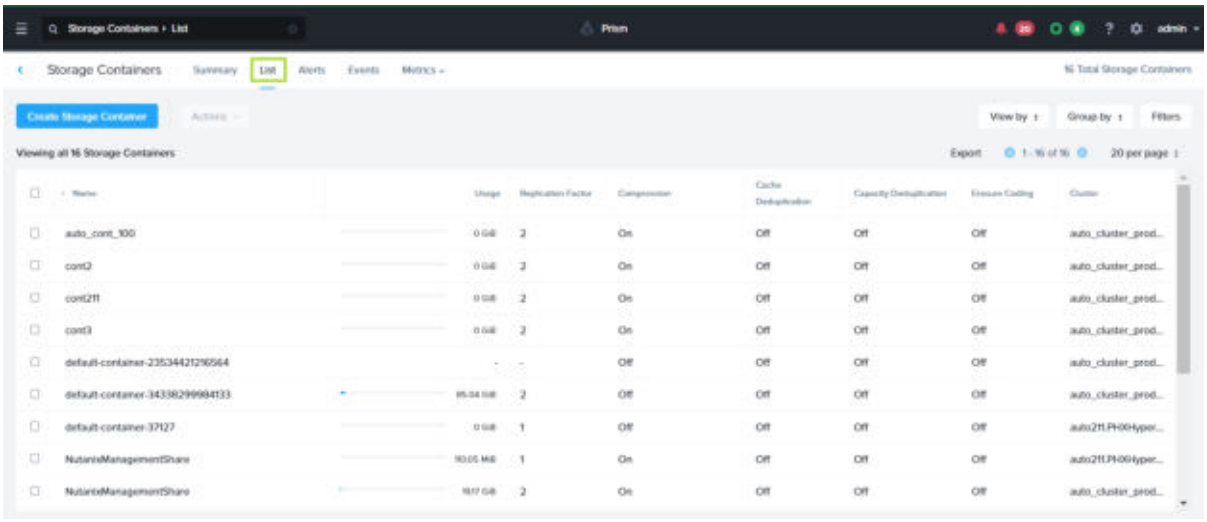

Figure 83: Storage Containers List Tab

Table 35: Storage Containers List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 128).	(name)
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the storage container. Hovering the cursor over the bar displays a window listing the used space, free space, and total space in the storage container.	xxx [GB|TB]
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether capacity deduplication is enabled, that is deduplication compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Cluster	Displays the name of the cluster in which the storage container resides. Click the name to display the cluster details (see Cluster Details View on page 205).	(cluster name)
"Performance" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 128).	(name)
Free Space	Displays the amount of free space available in the storage container.	xxx [GB|TB]

Parameter	Description	Values
Used Space	Displays the amount of used space in the storage container.	xxx [GB|TB]
Total Space	Displays the total amount of storage space in the storage container.	xxx [TB]
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
"Optimization" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 128).	(name)
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	*xx*:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GB|TB]
Compression Delay	Displays the time delay to perform compression on the data.	xx [m]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GB|TB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	*xx*:1

You can filter the storage containers list based on various parameter values. The following table describes the filter options available when you open the storage containers view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 36: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of storage containers that satisfy the name condition/string.	(storage container name string)
Cluster	Filters on the cluster name. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of storage containers in the clusters that satisfy the name condition/string.	(cluster name string)
Cache Deduplication	Filters on the cache deduplication setting. Select one or more settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
Replication Factor	Filters on the replication factor. Check the box for the desired replication factor (1, 2, 3). Only existing RF values appear. For example, if all containers are at RF 2, just a single box for RF 2 would appear. The number of storage containers currently in each setting is displayed on the right of the line.	1, 2, 3
Capacity Deduplication	Filters on the capacity deduplication setting. Select one or more settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
Erasure Coding	Filters on the erasure coding setting. Select one or both settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off
Free Space	Filters on the available storage space for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with available capacity in that range.	([xx] to [yy] GiB range)
Health	Filters on the storage container health state. Select one or more states to return a list of storage containers in those states. The number of storage containers currently in each state is displayed on the right of the line.	Critical, Warning, Good

Parameter	Description	Values
Used Space	Filters on the used storage space for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with used storage in that range.	([xx] to [yy] GiB range)
Total Space	Filters on the total available capacity for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with maximum capacity in that range.	([xx] to [yy] GiB range)
IOPS	Filters on the current IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of storage containers with IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of storage containers with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of storage containers with average I/O latency in that range.	([xx] to [yy] range)

You can group the storage containers list in the following ways:

- The **Group** pull-down menu allows you to group the host entries by replication factor, compression, erasure coding, cluster, or health setting. (You can only choose one.)
- The **View by** pull-down menu allows you to group the host entries by the information parameters (fields), which vary depending on whether you selected the **General**, **Performance**, or **Optimization** focus. (You can only choose one parameter.) You can also use the **+ Add custom** entry to create a custom view.

There are no action options available from the storage containers view (no action buttons and no **Actions** menu options when a storage container is selected).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just storage container-related alerts across the registered clusters (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just storage container-related events across the registered clusters (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the storage containers. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

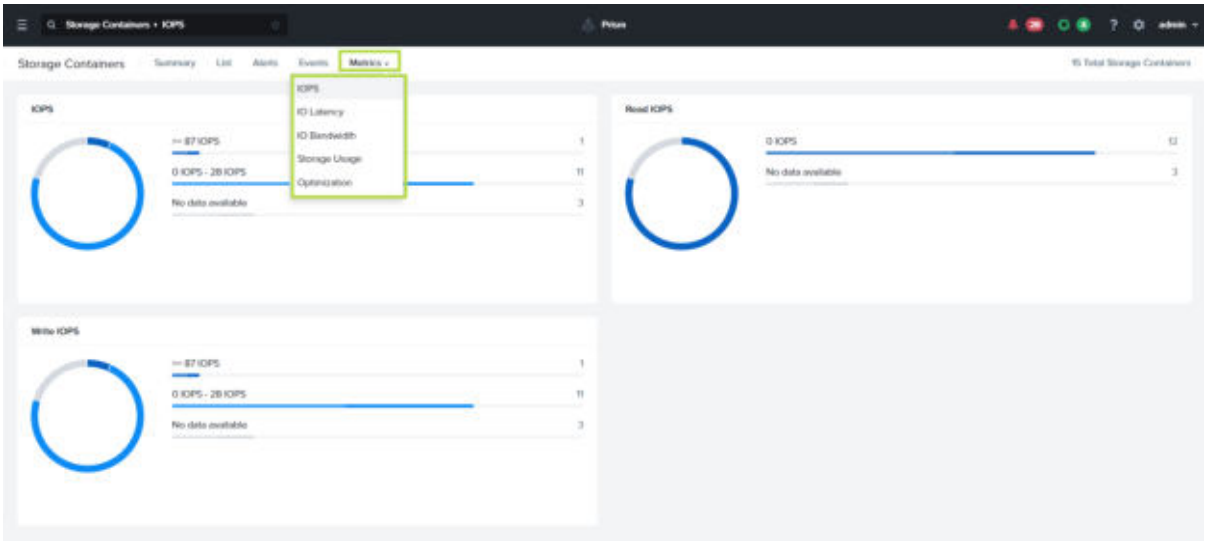

Figure 84: Storage Containers Metrics Tab

Table 37: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000). Note: The same format also applies to the other metrics in this table.
I/O Latency	Displays total, read, and write I/O latency rate graphs.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.
Storage Usage	Displays storage usage graph.
Optimization	Displays replication factor, data reduction ratio, and overall efficiency graphs.

Storage Container Details View

To access the details page for a storage container, go to the storage container dashboard **List** tab (see Storage Containers Summary View on page 122) and click the storage container name. You can also access the details page by clicking the storage container name wherever that name appears, such as in a dashboard widget or search result.

The storage container name and the following set of tabs appear at the top: **Summary**, **Alerts**, **Events**, **Metrics**, and **Usage**. Click a tab to display that information. (Click the arrow at the upper left to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the storage container (see following table).
- A **Usage** widget that displays usage data for the storage container (see following table).
- An **Optimization & Performance** widget that displays optimization and performance data for the storage container (see following table).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last week** (default), **Last 1 Hours**, or **Last 24 hours** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

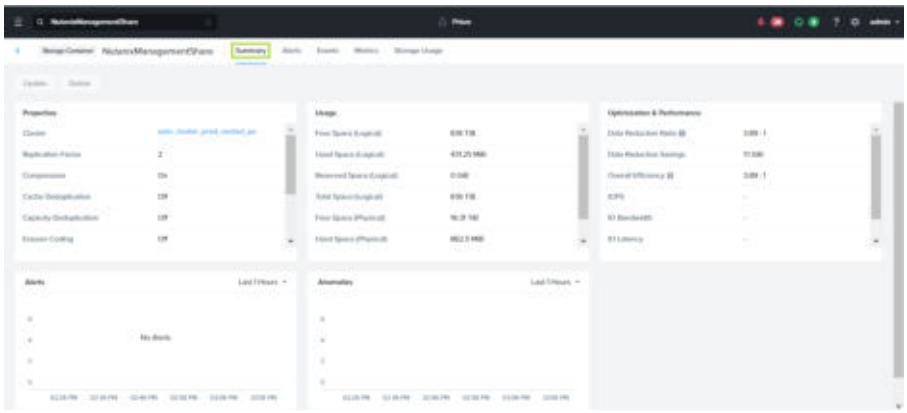

Figure 85: Storage Container Summary Tab

The following table describes the fields in the **Properties**, **Usage**, and **Optimization & Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 38: Storage Container Widget Fields

Parameter	Description	Values
"Properties" Widget		
Cluster	Displays the name of the cluster in which the storage container resides. Click the name to display the cluster details (see Cluster Details View on page 205).	(cluster name)
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]

Parameter	Description	Values
Compression	Displays whether compression is enabled.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether capacity deduplication is enabled, that is deduplication compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Filesystem Whitelists	Displays the IP addresses for file systems that are whitelisted.	[IP addresses]
Compression Delay	Displays the delay (number of minutes) before data changes are compressed. A zero value indicates that compression is immediate (not delayed).	xx min
Thick Provision	Displays the reserved storage capacity for thick provisioned VMs.	xxx [GB|TB]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GB|TB]
"Usage" Widget		
Free Space (Logical)	Displays the amount of logical (effective) free space available in the storage container. Note: Logical space accounts for the storage container replication factor. RF1 means the logical and physical spaces are the same, RF2 means the logical space is half the physical space, and RF3 means the logical space is a third of the physical space.	xxx [GB|TB]
Used Space (Logical)	Displays the amount of logical used space in the storage container.	xxx [GB|TB]
Reserved Space (Logical)	Displays the amount of logical reserved space in the storage container.	xxx [GB|TB]
Total Space (Logical)	Displays the amount of logical total space in the storage container.	xxx [GB|TB]

Parameter	Description	Values
Free Space (Physical)	Displays the amount of physical free space available in the storage container.	xxx [GB|TB]
Used Space (Physical)	Displays the amount of physical used space in the storage container.	xxx [GB|TB]
Total Space (Physical)	Displays the amount of physical total space in the storage container.	xxx [GB|TB]
"Optimization & Performance" Widget		
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	*xx*:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GB|TB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	*xx*:1
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
Erasure Coding	Displays whether erasure coding is enabled for the storage container.	[On|Off]

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this storage container (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this storage container (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the storage container. Click the **Metrics** tab to display graphs for nine metrics: **Storage Controller IOPS**, **Storage Controller Read IOPS**, **Storage Controller Write IOPS Storage Controller Latency**, **Storage Controller Read Latency**, **Storage Controller Write Latency**, **Storage Controller IO Bandwidth**, **Storage Controller Read IO Bandwidth**, **Storage Controller Write IO Bandwidth**. Each graph is a rolling time interval performance or usage monitor. The baseline range (when established) appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 838). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week).
- Click the **Actions** pull-down menu (upper right of any graph) to monitor the metric performance in greater detail (see Analysis Dashboard on page 739):
  - Click **Add to Analysis** to add an active monitor for the metric on the Analysis dashboard.
  - Click **Select Analysis Session** to assign the metric to a target session.
- Click the filter button to display the filter panel; check (uncheck) the desired boxes to display (hide) those metrics.
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 282).

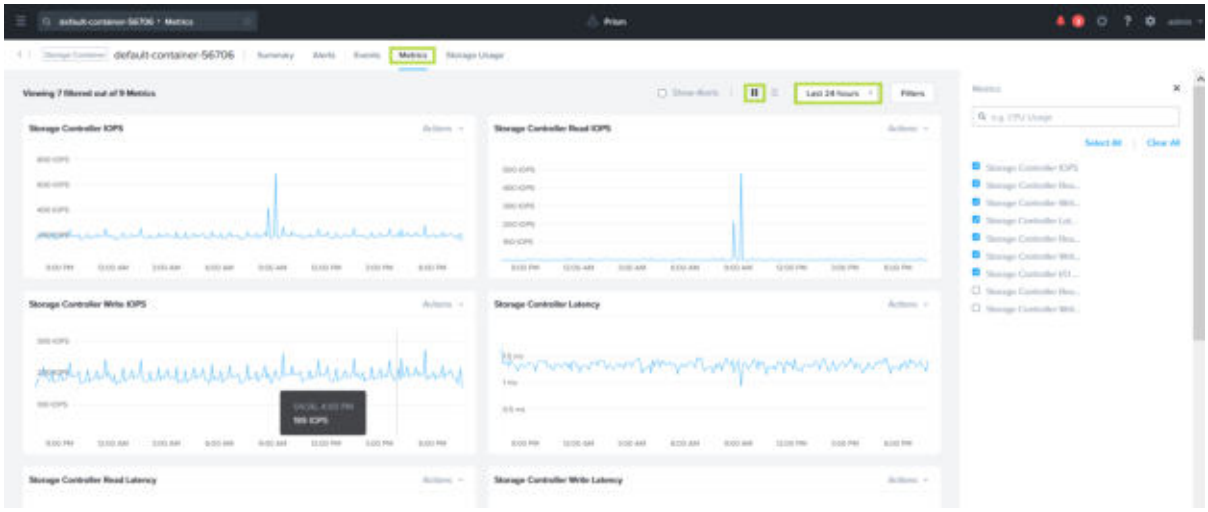

Figure 86: Storage Container Metrics Tab

Storage Usage Tab

The **Storage Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of storage container storage usage that can vary from one to several hours depending on activity moving from

right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of container storage space used by each disk tier (SSD and DAS-SATA).

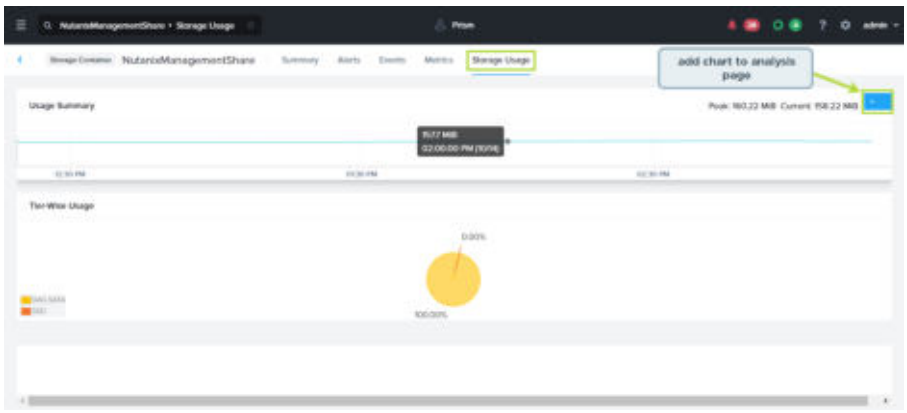

Figure 87: Storage Container Usage Tab

Storage Policies Summary View

To access the Storage Policies dashboard, select **Compute & Storage** > **Storage Policies** from the entities menu (see Entities Menu on page 13).

> Note: Storage Policies dashboard is not available if you log on as a view-only user.

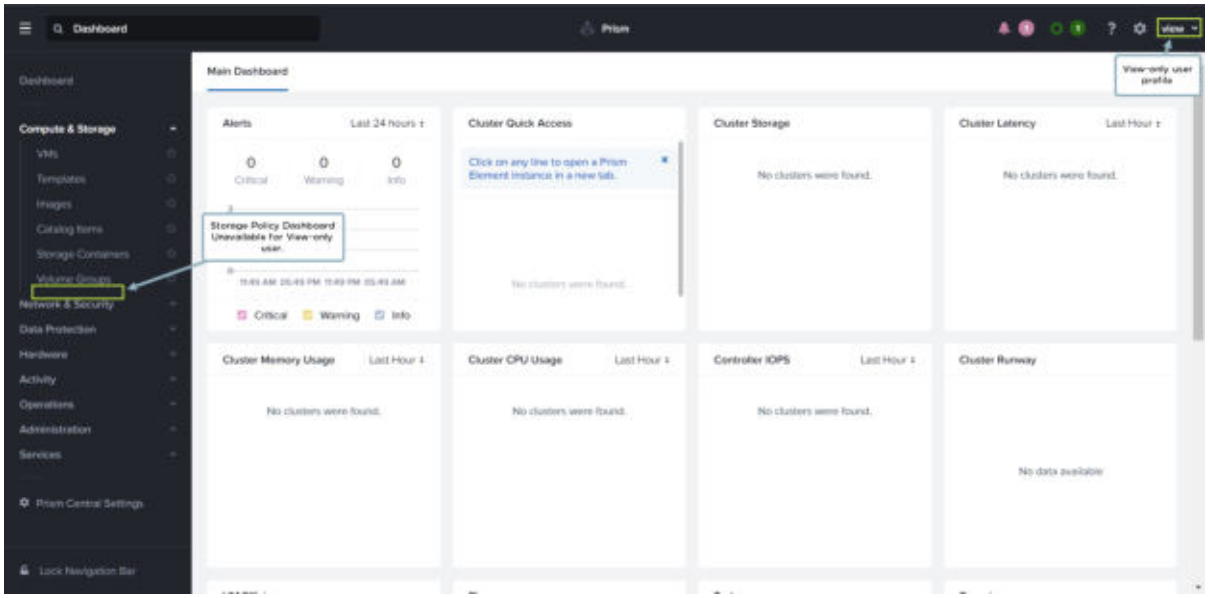

Figure 88:

The Storage Policies dashboard summary view displays information about the storage policies configured in Prism Central for the registered clusters and allows you to access detailed information about each storage policy. The dashboard has **List** tab with a display area below the selected tab that displays a list of the storage policies.

Note: This section describes the information and options that appear in the Storage Policies dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.
- See Storage Policy Management on page 647 for information about storage policies, creating and managing storage policies.

List Tab

The List tab appears by default when you first open the Storage Policies dashboard, displays a list of the security policies across the registered clusters.

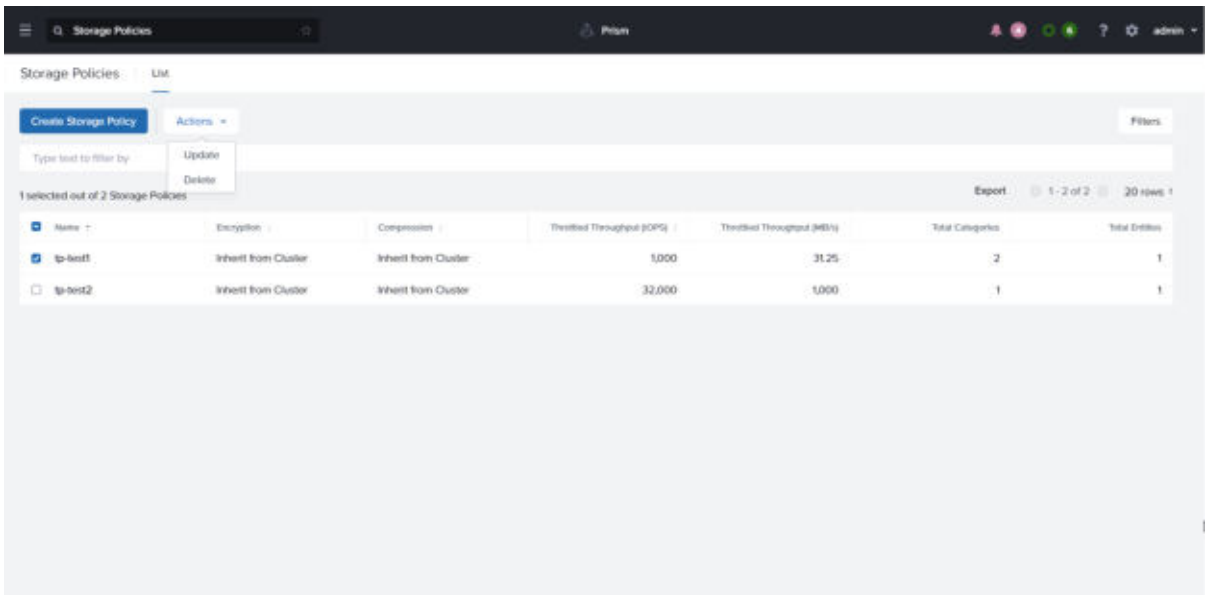

Figure 89: Storage Policy List Tab

The List tab displays the following:

- **Create Storage Policy**button: Click this button to open the Create Storage Policy page. See Creating or Updating a Storage Policy on page 650 for more information about creating a storage policy.
- **Actions** drop down menu: This menu includes two actions - **Update** and **Delete**.

  Click **Update** to open the Update Storage Policy page. The Update Storage Policy has the same fields as the Create Storage Policy page. See Creating or Updating a Storage Policy on page 650 for more information about creating a storage policy.

  Click **Delete** to delete a storage policy.

  Note: If there are any categories or entities like VMs that are associated with the storage policy, Prism Central displays a message box that says the following:

  ```
  This will remove the enforcement of the storage policy ‘tp-test1’ on all 1
  associated entities across 2 Categories. Any policy overridden by tp-test1
  ```

```
might get enforced on the entities.
```

The message box also provides a check box that seeks your confirmation before you proceed to delete the storage policy without dissociating the associated categories and entities from the policy:

```
I understand the impact of this action.
```

You can preview the changes by clicking the **Preview Changes** link. The Change Preview dialog box provides details of the categories and entities that are affected by the deletion.

The following table describes the fields that appear in the security policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 39: Storage Policies List Fields

Parameter	Description	Values
Name	Displays the name of the Storage Policy. Clicking the name displays the details page for that storage policy.	(Name of storage policy)
Encryption	Displays the encryption mode selected.	Enabled or Inherit from Cluster
Compression	Displays the compression mode selected.	Inline or Post Process (if enabled), Disabled or Inherit from Cluster
Throttled Throughput (IOPS)	Displays the throttled throughput value in terms of IOPS.	(Integer number)
Throttled Throughput (MBps)	Displays the throttled throughput value in MBps.	(Integer number)
Total Categories	Displays the total number of Categories that are associated with the storage policy.	(Integer number)
Total Entities	Displays the total number of entities that the storage policy is applied to. Note: A single storage policy can be applied to multiple entities. Multiple storage policies cannot be applied to a single entity.	(Integer number)

For more information about Storage Policies, see Storage Policy Management on page 647.

Storage Policy Details View

To access the details page for a storage policy, go to the storage policy **List** tab (see Storage Policy Management on page 647) and click the name of the storage policy you want to view the details of. You can also access the details page by clicking the name of the storage policy wherever that name appears, such as in a dashboard widget or search result.

The storage policy details page includes the following tabs:

- **Summary**—displayed by default when you open the storage policy details page. This tab provides information about the storage policy summarized into widgets.
- **Categories**—displays a list of Categories associated with the storage policy.
- **Entities**—displays the list of entities including VMs associated with the storage policy.

Summary Tab

The **Summary** tab displays the following:

- **Update** button: Click this button to open the Update Storage Policy page and modify the storage policy. This page has the same parameters as the Create Storage Policy page. See Creating or Updating a Storage Policy on page 650 for more information about creating a storage policy.
- **Delete** button: Click this button to delete the storage policy. When you delete a storage policy, the categories and entities associated with the policy are governed by any other policy that is applied to the category or entity.

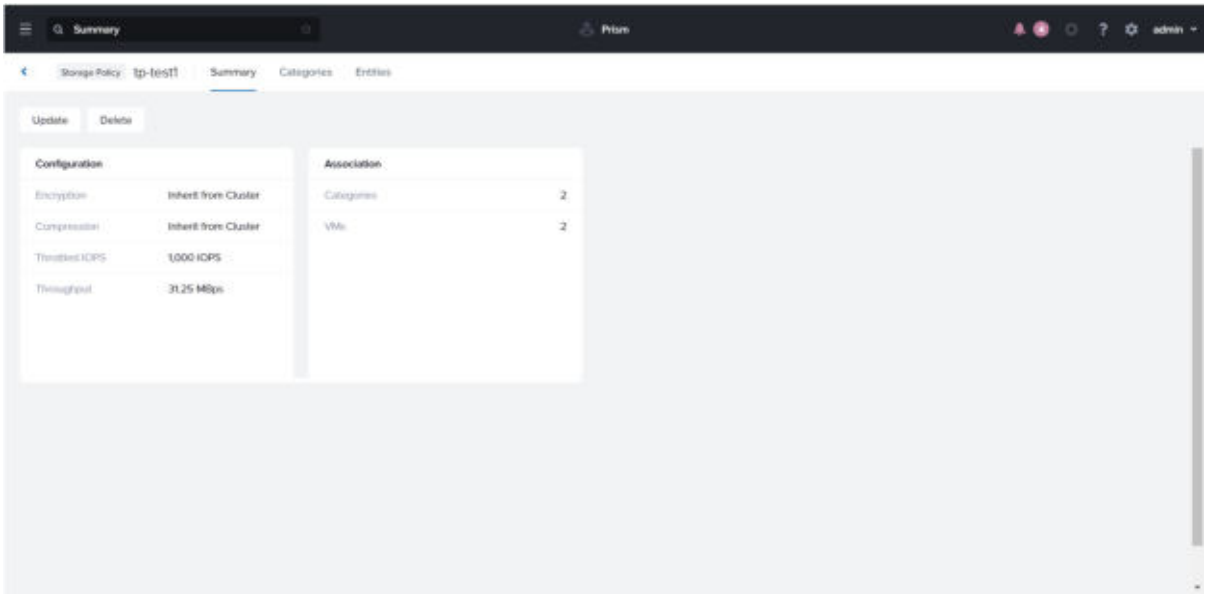

Figure 90: Storage Policy Details Summary Tab

The **Summary** tab displays information about the storage policy in following widgets:

- **Configuration**: This widget provides the following information.

Table 40: Configuration Properties

Properties	Description
Encryption	Indicates whether the storage policy involves encryption of the data. The possible values are: • **Enabled** where the encryption is enabled. • **Inherit from Cluster** where the entity inherits the encryption configuration from the cluster.
Compression	Indicates whether the storage policy involves compression of the data. The possible values are: • **Inline** • **Post Process** • **Off** to disable compression. • **Inherit from Cluster** where the entity inherits the compression configuration from the container.
Throttled IOPS	Provides the throttled throughput value in IOPS.
Throughput	Provides the throttled throughput value in MBps.

- **Association**: This widget provides the following information.

Table 41: Association Properties

Properties	Description
Categories	Indicates the number of categories associated with the storage policy.
VMs	Indicates the number of entities or VMs associated with the storage policy.

Categories Tab

The Categories tab displays the following:

- **Manage Associations** button: Click this button to add or delete category associations.
- The list of categories associated with the storage policy. The following table provides the details of the categories associated.

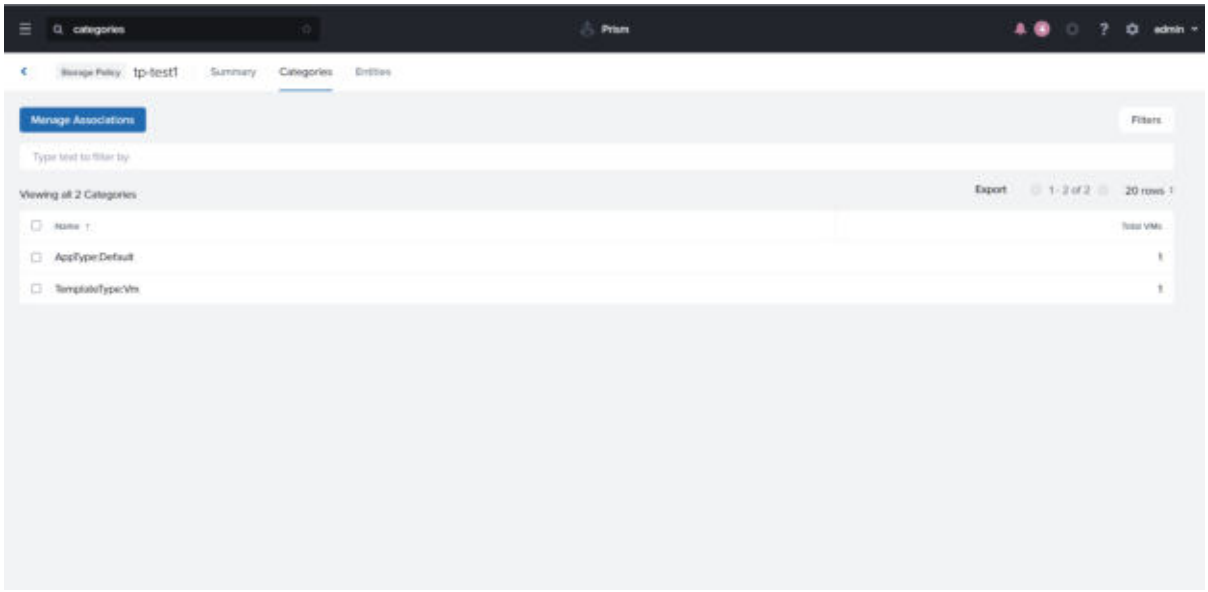

Figure 91: Storage Policy Details Categories Tab

Table 42: Categories List Fields

Properties	Description
Name	Displays the name of the category.
Total VMs	Displays the total number of VMs associated with the category.

Entities Tab

The Entities tab displays the following:

- **Manage Categories** button: Click this button to add or delete category associations. You can only associate entities to a storage policy if the entities are associated with one or more categories.

  Note: Storage policies use categories to associate with entities like VMs. For more information about the requirements and limitations for storage policies, see Storage Policy Management on page 647.

- The list of entities (VMs) associated with the storage policy. The following table provides the details of the entities associated.

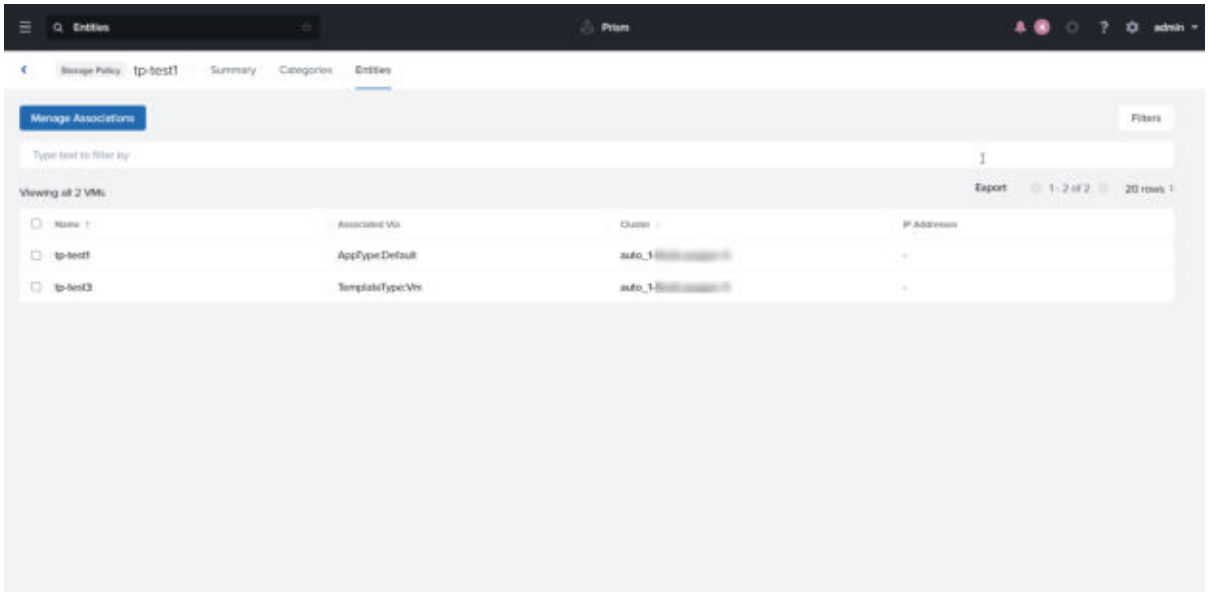

Figure 92: Storage Policies Details Entities Tab

Table 43: Entity List Fields

Properties	Description
Name	Displays the name of the category.
Associated via	Displays the categories that the entity and the storage policy are concurrently associated with. The storage policy is associated with the entity because the storage policy is associated with this category that is associated with the entity. For example, StoragePolicy1 is associated with CategoryA which is in turn associated with VM1. The category displayed in **Associated via** is CategoryA. The other categories associated with VM1 but not associated with StoragePolicy1 are not displayed here.
Cluster	Displays the name of the cluster that the entity is configured on.
IP Addresses	Displays the IP addresses configured on the entity.

Volume Groups Summary View

To access the volume groups dashboard, select **Compute & Storage** > **Volume Groups** from the entities menu (see Entities Menu on page 13). The volume groups dashboard summary view displays information about volume groups across the registered clusters and allows you to access detailed information about each volume group. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, and **Metrics**) with a display area below the selected tab.

Note:

- This section describes the information and options that appear in the volume groups dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.

- See Storage Management on page 611 for information about creating and managing volume groups.
- Some volume group details and actions are unavailable for clusters running a pre-6.0 AOS version.

Summary Tab

Clicking the **Summary** tab displays the following two widgets:

- **Highlighted Entities**: Displays a list of the volume groups with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that volume group. Click the **View all XX Volume Groups** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of storage container-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.

Figure 93: Volume Groups Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of volume groups across the registered clusters. The following table describes the fields that appear in the volume groups list. A dash (-) is displayed in a field when a value is not available or applicable.

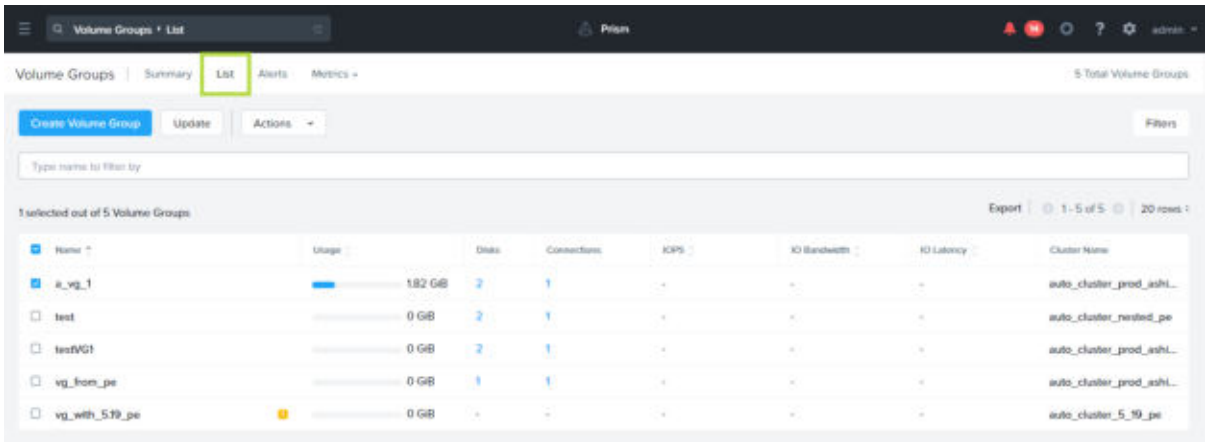

Figure 94: Volume Groups List Tab

Table 44: Volume Groups List Fields

Parameter	Description	Values
Name	Displays the name of the volume group. Click the name to display the details page for that volume group (see Volume Group Details View on page 143). An exclamation point icon appears if the cluster is running a pre-6.0 AOS version, which indicates the available information and actions are limited.	(name)
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the volume group. Hovering the cursor over the bar displays a window listing the used space, free space, and total space in the volume group.	xxx [GB|TB]
Disks	Displays the number of virtual disks in the volume group. Clicking the number displays the **Virtual Disks** tab of the details view (see Volume Group Details View on page 143).	(integer)
Connections	Displays the number of external client connections to the volume group. Clicking the number displays the **Connections** tab of the details view (see Volume Group Details View on page 143).	(integer)
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Cluster Name	Displays the name of the cluster in which the volume group resides. Click the name to display the cluster details (see Cluster Details View on page 205).	(cluster name)

You can filter the volume groups list based on various parameter values. The following table describes the filter options available when you open the volume groups view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 45: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the volume group name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of volume groups that satisfy the name condition/string.	(volume group name string)
Usage	Filters on the used storage space for a volume group. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of volume groups with used storage in that range.	([xx] to [yy] GiB range)
Cluster Name	Filters on the cluster name. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of volume groups in the clusters that satisfy the name condition/string.	(cluster name string)
IOPS	Filters on the current IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It returns a list of volume groups with IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It returns a list of volume groups with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of volume groups with average I/O latency in that range.	([xx] to [yy] range)

The volume groups view includes two action buttons and an actions menu:

- To create a volume group, click the **Create Volume Group** button (see Creating a Volume Group on page 637).
- To update or delete a volume group, click the **Update** or **Delete** button (see Modifying a Volume Group on page 641).
- To manage external client connections for a volume group, select the volume group and then select **Manage Connections** from the **Actions** menu (see Managing Volume Group Connections on page 642).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just volume group-related alerts across the registered clusters (see Alerts Summary View on page 261).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the volume groups. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

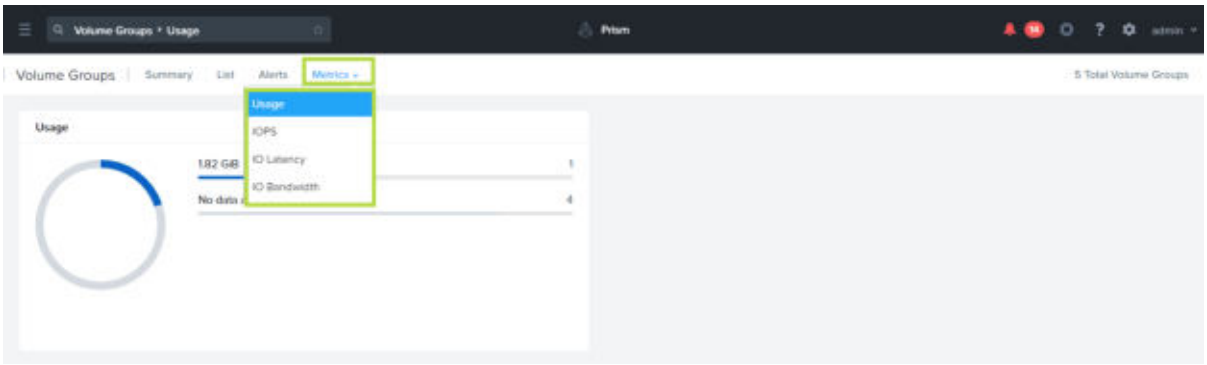

Figure 95: Volume Groups Metrics Tab

Table 46: Metrics Tab Fields

Metric	Description
Usage	Displays volume group usage graph.
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000). Note: The same format also applies to the other metrics in this table.
I/O Latency	Displays total, read, and write I/O latency rate graphs.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.

Volume Group Details View

To access the details page for a volume group, go to the volume groups dashboard **List** tab (see Volume Groups Summary View on page 139) and click the volume group name. You can also access the details page by clicking the volume group name wherever that name appears, such as in a dashboard widget or search result.

The volume group name and the following set of tabs appear at the top: **Summary**, **Virtual Disks**, **Connections**, **Alerts**, and **Metrics**. Click a tab to display that information. (Click the arrow at the upper left to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the storage container (see following table).

- A **Usage and Performance** widget that displays usage data for the storage container (see following table).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last week** (default), **Last 1 Hours**, or **Last 24 hours** from the pull-down menu.

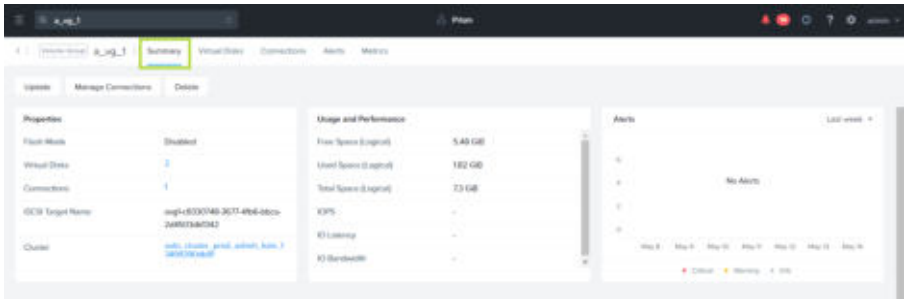

Figure 96: Volume Group Summary Tab

The following table describes the fields in the **Properties** and **Usage and Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 47: Volume Group Widget Fields

Parameter	Description	Values
"Properties" Widget		
Flash Mode	Displays whether Flash Mode is enabled for the volume group.	[Enabled, Disabled]
Virtual Disks	Displays the number of virtual disks in the volume group. Clicking the number displays the **Virtual Disks** tab.	(integer)
Connections	Displays the number of external client connections to the volume group. Clicking the number displays the **Connections** tab.	(integer)
iSCSI Target Name	Displays the name of the iSCSI target.	(ISCSI name)
Cluster	Displays the name of the cluster in which the volume group resides. Click the name to display the cluster details (see Cluster Details View on page 205).	(cluster name)
"Usage and Performance" Widget		

Parameter	Description	Values
Free Space (Logical)	Displays the amount of logical (effective) free space available in the volume group. Note: Logical space accounts for the replication factor. RF1 means that the logical and physical spaces are the same, RF2 means that the logical space is half the physical space, and RF3 means that the logical space is a third of the physical space.	xxx [GB|TB]
Used Space (Logical)	Displays the amount of logical used space in the volume group.	xxx [GB|TB]
Total Space (Logical)	Displays the amount of logical total space in the volume group.	xxx [GB|TB]
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]

The **Summary** tab includes three action buttons:

- To update or delete the volume group, click the **Update** or **Delete** button (see Modifying a Volume Group on page 641).
- To manage external client connections for the volume group, click the **Manage Connections** button (see Managing Volume Group Connections on page 642).

Virtual Disks Tab

Clicking the **Virtual Disks** tab displays a list of virtual disks for the volume group. The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable.

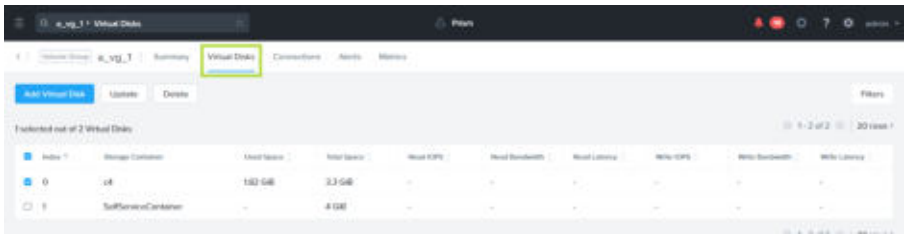

Figure 97: Volume Group Virtual Disks Tab

Table 48: Virtual Disks Fields

Parameter	Description	Values
Index	Displays the index number assigned to the virtual disk.	(integer)
Storage Container	Displays the name of the storage container in which the virtual disk is located.	(container name)
Used Space	Displays the amount of used space in the virtual disk.	xxx [GB|TB]
Total Space	Displays the total amount of storage space in the virtual disk.	xxx [GB|TB]
Read IOPS	Displays the current read I/O operations per second (IOPS) for the virtual disk.	[0 - unlimited]
Read Latency	Displays the average read I/O latency for the virtual disk.	xxx [ms]
Read Bandwidth	Displays read I/O bandwidth used per second for the virtual disk.	xxx [MBps|KBps]
Write IOPS	Displays the current write IOPS for the virtual disk.	[0 - unlimited]
Write Latency	Displays the average write I/O latency for the virtual disk.	xxx [ms]
Write Bandwidth	Displays write I/O bandwidth used per second for the virtual disk.	xxx [MBps|KBps]

You can filter the virtual disks list based on various parameter values. The following table describes the filter options available when you open the virtual disks view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 49: Filter Pane Fields

Parameter	Description	Values
storage container	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of volume groups that satisfy the name condition/string.	(storage container name string)
[Total|Used] Space	Filters on the total or used storage space for a virtual disk. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of virtual disks with total or used storage in that range.	([xx] to [yy] GiB range)

Parameter	Description	Values
[Read|Write] IOPS	Filters on the current read or write IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It returns a list of virtual disks with IOPS in that range.	([xx] to [yy] range)
[Read|Write] Bandwidth	Filters on the read or write I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It returns a list of virtual disks with I/O bandwidth usage in that range.	([xx] to [yy] range)
[Read|Write] Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of virtual disks with average I/O latency in that range.	([xx] to [yy] range)

The **Virtual Disks** tab includes three action buttons:

- To add a virtual disk to the volume group, click the **Add Virtual Disk** button (see Managing Volume Group Virtual Disks on page 646).
- To update or delete a virtual disk, select the target disk and then click the **Update** or **Delete** button (see Managing Volume Group Virtual Disks on page 646).

Connections Tab

Clicking the **Connections** tab displays a list of external client connections to the volume group. The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable.

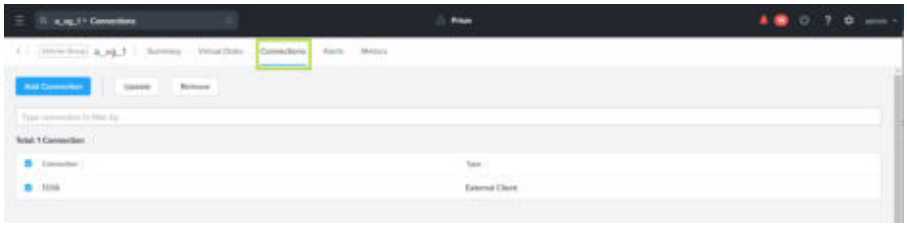

Figure 98: Volume Group Connections Tab

Table 50: Connections Fields

Parameter	Description	Values
Connection	Displays the IP address or IQN of the external client.	(IP or IQN address)
Type	Displays the type of connection.	external client

The **Connections** tab includes three action buttons:

- To add an external client connection, click the **Add Connection** button (see Managing Volume Group Connections on page 642).
- To update or remove a connection, click the **Update** or **Remove** button (see Managing Volume Group Connections on page 642).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this volume group (see Alerts Summary View on page 261).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the volume group. Click the **Metrics** tab to display graphs for 10 metrics: **Usage**, **Controller IOPS**, **Controller Read IOPS**, **Controller Write IOPS Controller Average I/O Latency**, **Controller Average Read I/O Latency**, **Controller Average Write I/O Latency**, **Controller I/O Bandwidth**, **Controller Read I/O Bandwidth**, **Controller Write I/O Bandwidth**. Each graph is a rolling time interval performance or usage monitor. The baseline range (when established) appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 838). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- Click the **Actions** pull-down menu (upper right of any graph) to monitor the metric performance in greater detail (see Analysis Dashboard on page 739):
  - Click **Add to Analysis** to add an active monitor for the metric on the Analysis dashboard.
  - Click **Select Analysis Session** to assign the metric to a target session.
- Click the filter button to display the filter panel; check (uncheck) the desired boxes to display (hide) those metrics.
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 282).

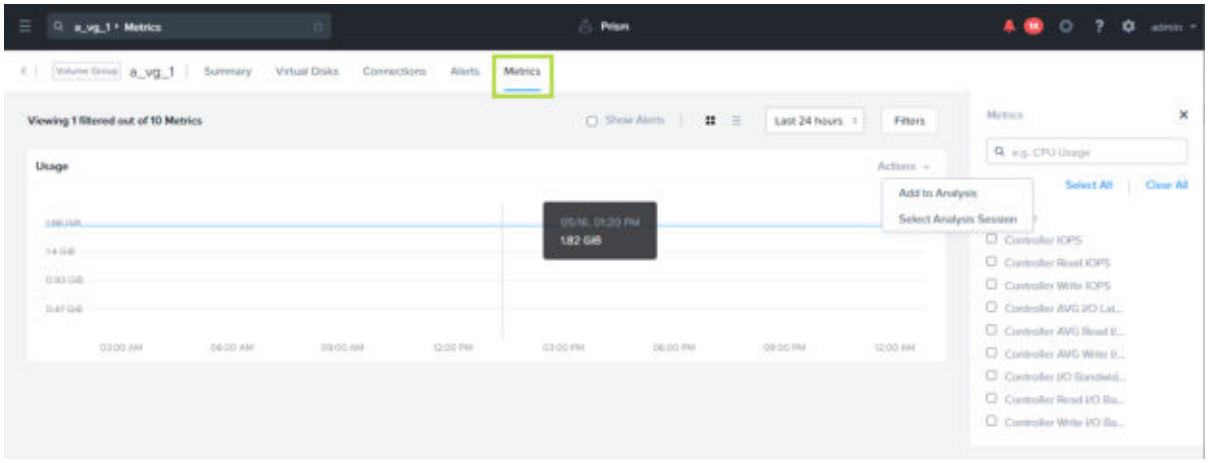

Figure 99: Volume Group Metrics Tab

vCenter Datastores Summary View

To access the vCenter datastores dashboard, select **Compute & Storage** > **vCenter Datastores** from the entities menu (see Entities Menu on page 13). The vCenter datastores dashboard summary view displays information about datastores across the registered vCenter instances and allows you to access detailed information about each datastore. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below for the selected tab.

Note:

- This section describes the information and options that appear in the vCenter datastores dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- The vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 796 for more information.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities**: Displays a list of the datastores with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The only option is **Vsphere Datastore Numberreadaveraged Average**. Click the **View All XX vCenter Datastores** link at the bottom to display the **List** tab (following section).
- **Alerts**: Displays a list of datastore-related alerts that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

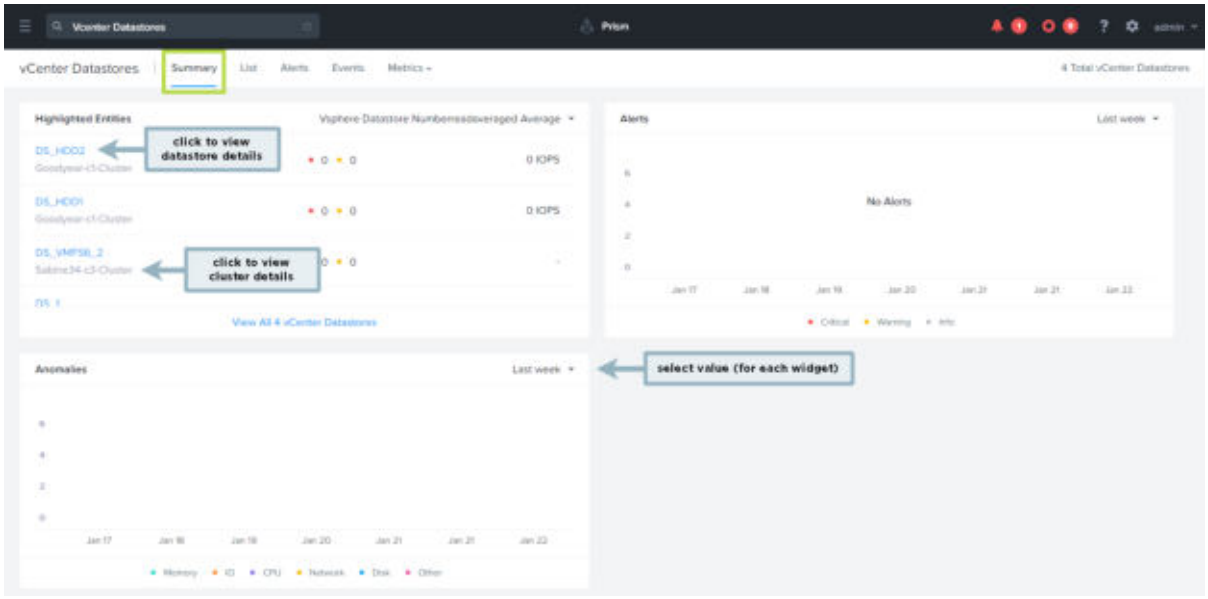

Figure 100: Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the registered vCenter datastores. The following table describes the fields that appear in the datastores list. A dash (-) is displayed in a field when a value is not available or applicable.

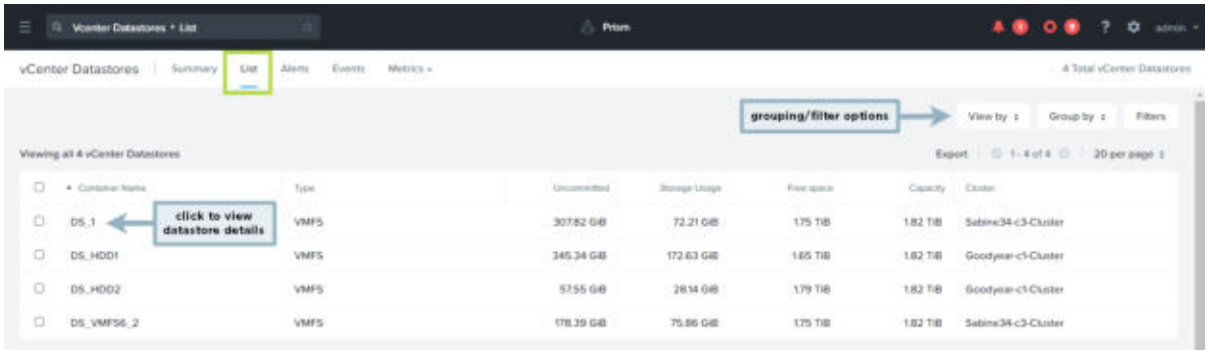

Figure 101: List Tab

Table 51: Datastores Fields

Parameter	Description	Values
Container Name	Displays the name of the datastore container. Click the name to display the details for that container (see vCenter Datastore Details View on page 152).	(container name)
Type	Displays the file system type.	VMFS, NFS
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB|TiB]

Parameter	Description	Values
Storage Usage	Displays the amount of used storage space.	xx [GiB|TiB]
Free Space	Displays the amount of free storage space.	xx [GiB|TiB]
Capacity	Displays the total amount of storage capacity.	xx [GiB|TiB]
Cluster	Displays the name of the cluster in which the datastore resides.	(cluster name)

You can filter the datastores list based on a variety of parameter values. The following table describes the filter options available when you open the datastores view **Filters** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 52: Filter Pane Fields

Parameter	Description	Values
Container Name	Filters on the container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of containers that satisfy the name condition/string.	(name string)
Cluster	Filters on the cluster name. Enter a string in the field to display a list of clusters that satisfy the name string.	(cluster name)
Storage Usage	Filters on the amount of used storage space. Check the box(es) for the desired range or enter a range in the "**from <low> to <high> TiB**" field. The number of containers in each range is displayed on the right of the line.	([xx] to [yy] range)
Capacity	Filters on the amount of total storage capacity.	([xx] to [yy] range)
Free Space	Filters on the amount of free storage space.	([xx] to [yy] range)
Uncommitted	Filters on the amount of uncommitted storage space.	([xx] to [yy] range)
Type	Filters on file system type.	VMFS, NFS

To download the table contents in CSV format, click the **Export** link.

You can group the datastore containers list in the following ways:

- The **View by** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.
- The **Group by** pull-down menu allows you to group the datastore entries by cluster or type. (You can only choose one.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just datastore-related alerts across the registered clusters (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just datastore-related events across the registered clusters (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the datastores. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant performance information. The following table describes the available metrics.

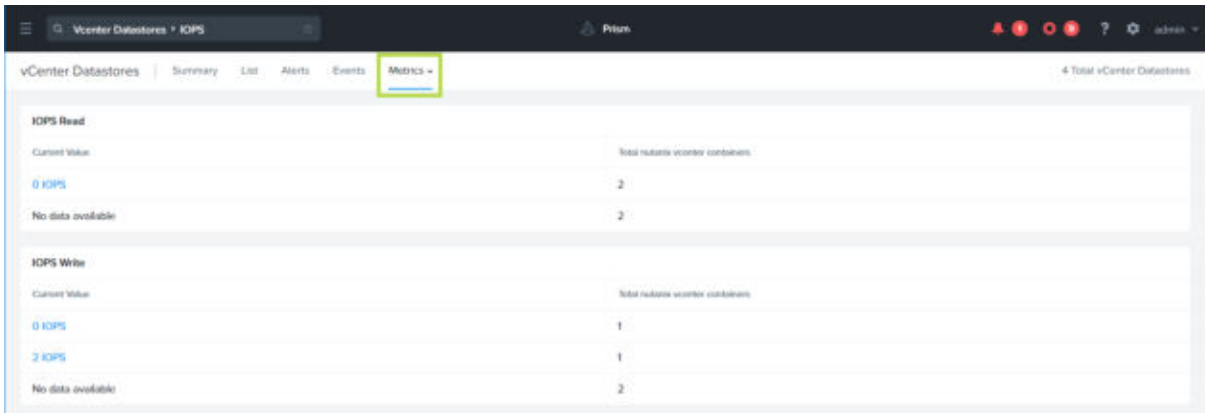

Figure 102: Metrics Tab (IOPS)

Table 53: Metrics Tab Fields

Metric	Description
IOPS	Displays the number of read and write IOPS for the listed number of containers. (Clicking an IOPS value redisplays the Summary tab filtered for those containers.)
Disk Usage	Displays the disk usage amount for the listed number of containers. (Clicking the usage value redisplays the Summary tab filtered for those containers.)
Disk Capacity	Displays the disk capacity for the listed number of containers. (Clicking the capacity value redisplays the Summary tab filtered for those containers.)
Disk Provisioned	Displays the provisioned disk space for the listed number of containers. (Clicking the provisioned value redisplays the Summary tab filtered for those containers.)

vCenter Datastore Details View

To access the details page for a vCenter datastore, go to the vCenter datastores **Summary** tab (see vCenter Datastores Summary View on page 149) and click the datastore name. You can

also access the details page by clicking the datastore name wherever that name appears, such as in a dashboard widget or search result.

The datastore name and the following set of tabs appear at the top: Summary, Alerts, Events, and Metrics. Click a tab to display that information below the tabs. (Click the **Back to vCenter Datastores** link to return to the datastores summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the VM (see following table).
- A **Metrics** widget that displays average metrics for the past 15 minutes (IOPS reads and writes, and disk usage, disk capacity, and disk provisioned amounts).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
- 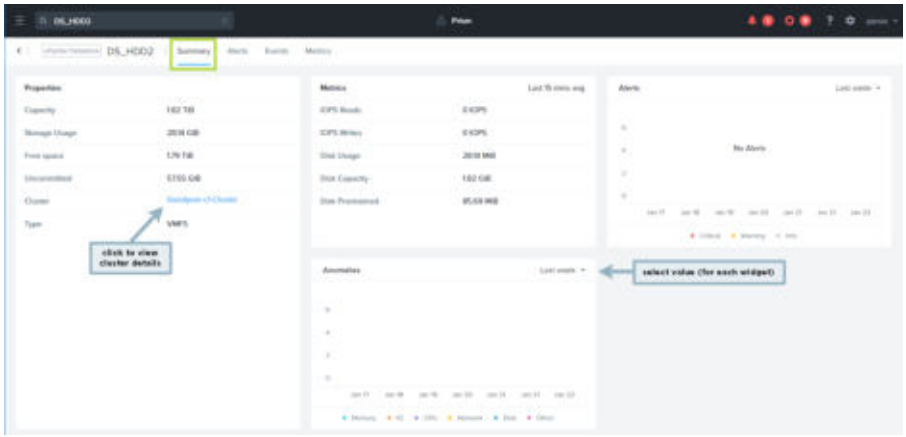

Figure 103: Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 54: Datastore Properties Fields

Parameter	Description	Values
Capacity	Displays the total amount of storage capacity.	xx [GiB|TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB|TiB]
Free Space	Displays the amount of free storage space.	xx [GiB|TiB]

Parameter	Description	Values
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB|TiB]
Cluster	Displays the name of the cluster in which the datastore resides. Click on the cluster name to displays details about that cluster.	(cluster name)
Type	Displays the file system type.	VMFS, NFS

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this vCenter datastore (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this vCenter datastore (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab displays a set of graphs for various metrics. Each graph is a rolling time interval performance or usage monitor for that metric.

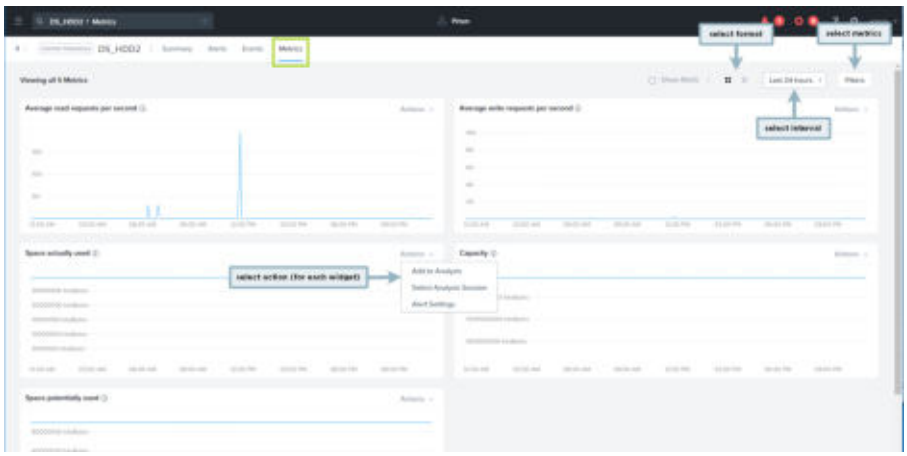

Figure 104: Metrics Tab

- Each graph includes a pull-down actions menu with some or all of the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard on page 739).
  - Select **Select Analysis Session**, which opens the **Select A Session** window, to add the metric to an existing or new analysis session. Choose an option from the pull-down list and then click the Add to Session button.

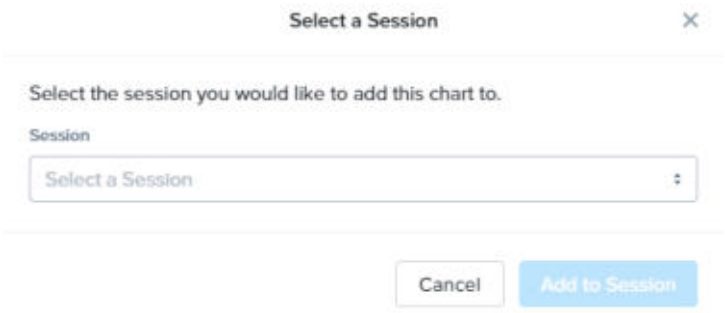

Figure 105: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 282).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired metrics.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to vCenter Datastores** link to return to the vCenter datastores dashboard.

The following table describes the metric graphs.

Table 55: Metric Performance Graphs

Metric	Description
Average read requests per second	Average number of read commands issued per second to the datastore during the collection interval.
Average write requests per second	Average number of write commands issued per second to the datastore during the collection interval.
Space actually used	Amount of space actually used by the virtual machine or the datastore. May be less than the amount provisioned at any given time, depending on whether the virtual machine is powered-off, whether snapshots have been created or not, and other such factors. Available from datastore and virtual machine target entities.
Capacity	Configured size of the datastore. Available from a datastore entity only.

Metric	Description
Space potentially used.	Amount of storage set aside for use by a datastore or a virtual machine. Files on the datastore and the virtual machine can expand to this size, but not beyond it. Available from datastore and virtual machine target entities.

Catalog Items View

To access the catalog items dashboard, select **Compute & Storage** > **Catalog Items** from the **Entities** menu (see Entities Menu on page 13). The catalog items dashboard allows you to view information about the images and VMs in the catalog.

> Note: This section describes the information and options that appear in the catalog items dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See Catalog Management on page 657 for information about the catalog service and how to add (or delete) items.

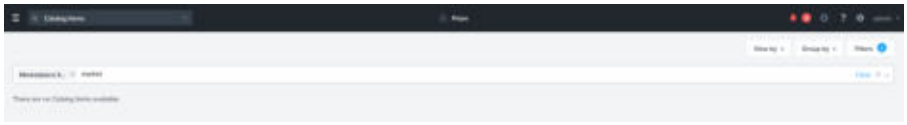

Figure 106: Catalog Items Dashboard

The following table describes the fields that appear in the catalog items list. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 56: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the item name.	(item name)
Type	Displays the type of item (image or VM).	Image, Vm
Description	Indicates which user added this item to the catalog.	(text string)

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 57: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of catalog items that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of catalog items whose description field satisfy the condition/string.	(description string)
Type	Filters on the item type. Check the box(es) for the desired item types.	Vm, Image

You can group the catalog items list in the following ways:

- The **Group** pull-down menu allows you to group by type (VM and image).
- The **View by** pull-down menu allows you to view by name, description, message, percent, reason, state, and type.

The **Actions** menu appears when one or more catalog items are selected. It allows you to delete the selected catalog items (see Deleting a Catalog Item on page 660).

Categories Summary View

To access the categories dashboard, select **Compute & Storage** > **Categories** from the entities menu (see Entities Menu on page 13). The categories dashboard allows you to view summary information about existing categories and access detailed information about each category.

> Note: This section describes the information and options that appear in the categories dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See Category Management on page 735 for information on how to create, modify, and apply categories.

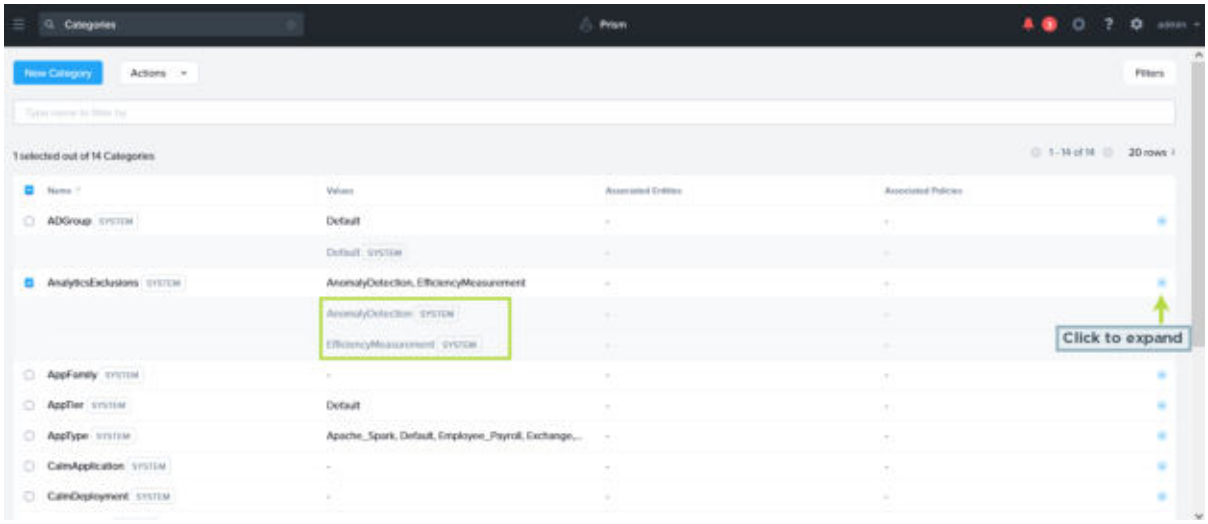

Figure 107: Categories Dashboard

The following table describes the fields that appear in the categories list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 58: Categories List Fields

Parameter	Description	Values
Name	Displays the name of the category.	(category name)
Value	Displays the values defined for the category. Click **Show more** (right of line) to see a line for each value. (Click**Show fewer** to collapse the list.) You may see an icon with a "showing X of Y" message to indicate there are additional values; click the icon to display the full list.	(value names)
Assigned Entities	Lists the number of entities assigned to this category.	(number)
Assigned Policies	Lists the number of policies assigned to this category.	(number)

You can filter the category list based on several parameter values. The following table describes the options available when you click the **Filter** button, which displays the Categories view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 59: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the category name. Select a condition from the pull-down list and enter a string in the field. It will return a list of categories that satisfy the name condition/ string. Note: In this and the Value field, the condition menu options are **Contains**, **Does not contain**, **Starts with**, **Ends with**, and **Equal to**.	(name string)
Entities	Filters on the entity type. Check the box for one or more entity types.	VM, Host, App, Blueprint, Cluster, Image, Marketplace Item, Reports, Subnet
Policies	Filters on the policy type. Check the box for one or more entity types.	Security Policy, Affinity Policy, Image Placement Policy, NGT Policy, Protection Policy, QoS Policy

There is a **New Category** action button to create a new category (see Creating a Category on page 735). The **Actions** menu appears when one or more categories are selected and includes the following options:

- This action can be applied to multiple categories: **Delete**
- This action can be applied to only one category at a time: **Update**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the selected category. See Modifying a Category on page 736 for instructions on how to perform these actions.

Category Details View

You can view detailed policy, value, and entity associations for a category by going to the details page for that category. To access the details page for a category, go to the categories dashboard (see Categories Summary View on page 157) and click the category name.

A category details page includes the following:

- Category name (upper left). You can switch from one category to another by selecting a different category name from the pull-down list in the upper left of the screen.

- Action buttons (upper right).
  - Click the **Update** button to update the category definition (see Modifying a Category on page 736); click the **Delete** button to delete the category. The button is grayed out if that action is not allowed. For example, you cannot delete system categories.
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.
- **Policies** (left), **Values** (middle), and **Entities** (right) columns that list the values defined for the category plus the policies and entities associated with that category. Placing the cursor over a policy, value, or entity displays lines that graphically indicate the association among the three parameters. For example, in the following figure placing the cursor over the **Forensics** value shows that it is associated with the **Quarantine-Forensics** security policy, but it is not associated with any entities currently.

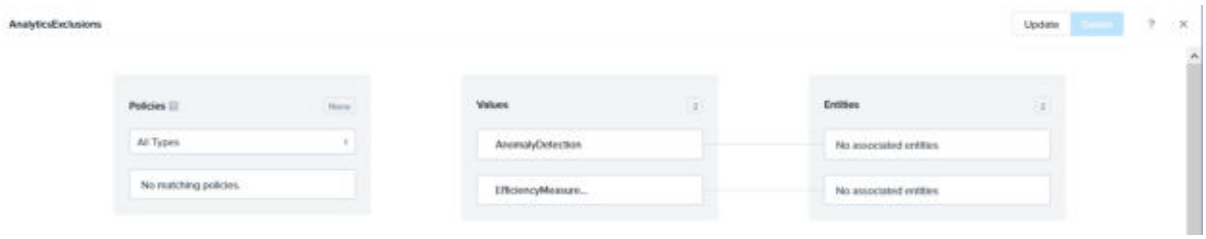

Figure 108: Category Details View

OVAs View

To access the OVAs dashboard, select **Compute & Storage** > **OVAs** from the entities menu (see Entities Menu on page 13). The OVAs dashboard allows you to view information about the OVAs available in the

> Note: This section describes the information and options that appear in the OVAs dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See OVA Management on page 688 for information about the OVA operations.

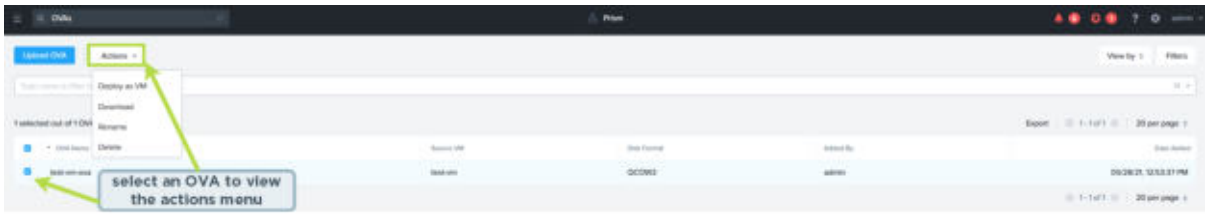

Figure 109: OVAs Dashboard

The following table describes the fields that appear in the catalog items list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 60: OVA List Fields

Parameter	Description	Values
OVA Name	Displays the name of the OVA file.	(file name)
Source VM	Displays the name of source VM that was exported as an OVA. The alert indicators for the task status are displayed beside the source VM for which the task is run. See the Alert indicator table for more information.	(file name)
Disk Format	Displays the format of the disks in the OVA.	[QCOW2|VMDK]
Added by	Indicates which user added this OVA.	(text string)
Date added	Displays the date and time when the OVA was added.	(date in mm/dd/yyyy format, time in hr:min:sec [AM|PM] format) Example: 07/30/20, 3:00:20 PM

Table 61: OVA Alerts

OVA Alerts	Status Message sample	Indicates...
	`Incorrect file uploaded. Multiple OVA files present in OVA.`	That the upload process failed during validation. You must fix the issue in the OVA file or upload the correct, compatible OVA file.

Figure 110:

	`Upload for OVA file has not started. Please start OVA file upload using Resume Upload option.`	That there is an interruption in one of the three sub-tasks in the upload process. You can resume the upload. See the *Resume Upload* action procedure.
	`Uploaded OVA file has not been verified. Please start verification using Resume Upload option.`	

A sample of the error hover message is as follows:

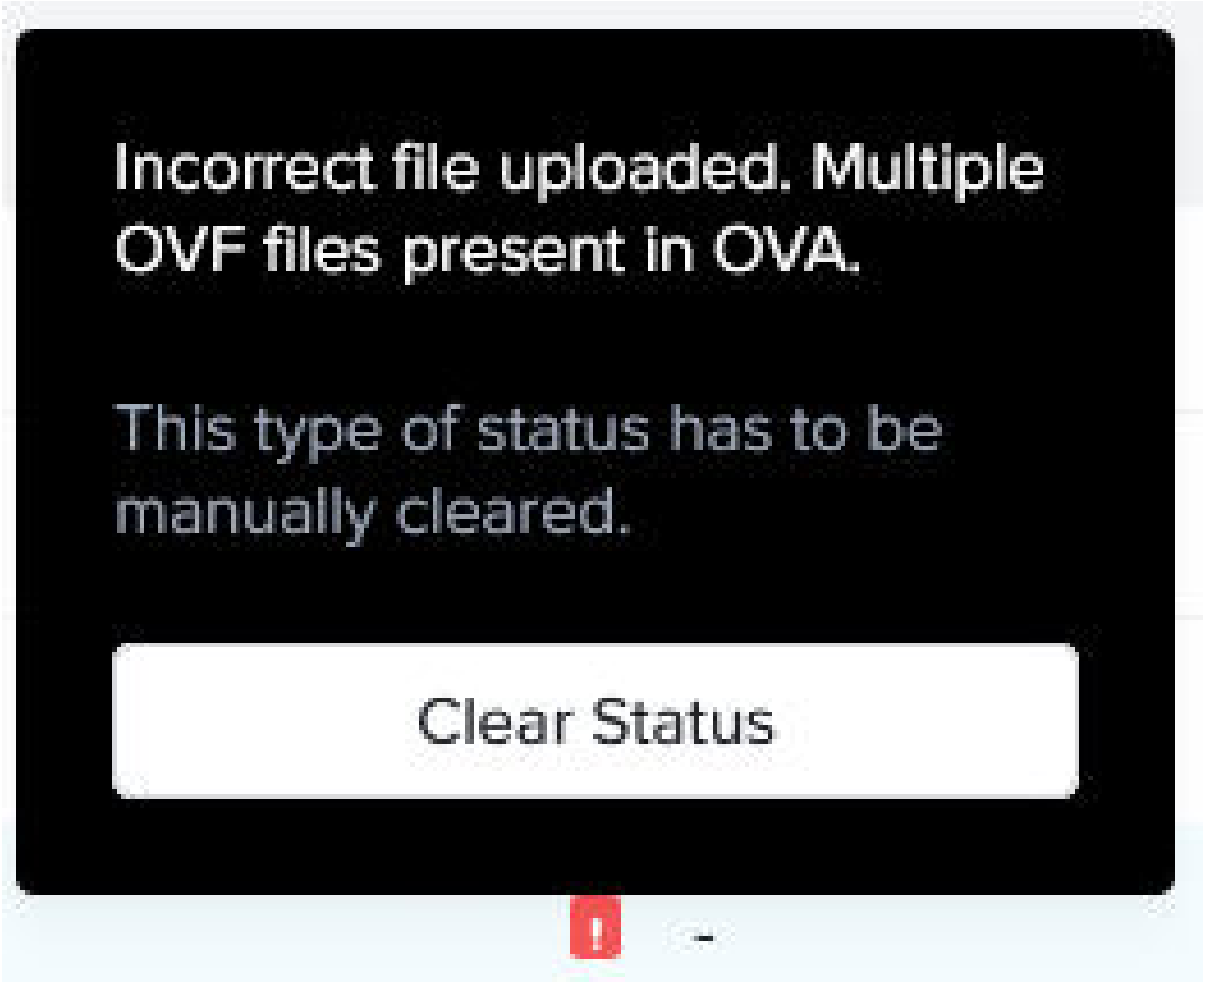

Figure 113: Sample error hover message

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

For each filter, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.

Table 62: Filter Pane Fields

Parameter	Description	Values
OVA Name	Filters on the OVA file name. It returns a list of OVAs that satisfy the name condition/string.	(OVA file name string)
Source VM	Filters on the Source VM name. It returns a list of OVAs for which the Source VM field satisfies the condition/string.	(Source VM name string)

Parameter	Description	Values
Disk Format	Filters on the disk format. Check the box(es) for the OVAs of the required disk format.	VMDK, QCOW2

The **Actions** menu appears when a OVA is selected. It allows you to:

- Deploy an OVA file as a VM.
- Download an OVA file to your local machine.
- Rename an OVA file.
- Delete an OVA file.
- Resume upload when an upload is interrupted and shows an error.

> Note: The resume upload action is not available in case the concatenate process of the upload is interrupted. In such a case, you need to run the Concatenate API from the REST API Explorer to resume the upload.

Network and Security View

The **Network and Security** category in the **Entities Menu** expands on-click to display the following networking and security entities that are configured for the registered clusters:

> Note: In

1. Subnets: This dashboard displays the subnets and the operations you can perform on subnets.
2. Virtual Private Clouds: This dashboard displays the VPCs and the operations you can perform on VPCs.
3. Floating IPs: This dashboard displays a list of floating IP addresses that you are using in the network. It allows you to request for floating IP addresses from the free pool of I addresses available to the clusters managed by the Prism Central instance.
4. Connectivity: This dashboard allows you to manage the following networking capabilities:
    - Gateways: This tab provides a list of network Gateways you have created and configured, and the operations you can perform on the network Gateways. You can check and upgrade the Gateway bundle in **Administration** > **LCM** > **Inventory**.
    - VPN Connections: This tab provides a list of VPN connections you have created and configured, and the operations you can perform on VPN connections.
    - Subnet Extensions: This tab provides a list of subnets that you have extended at the Layer 2 level using VPN (point-to-point over Nutanix VPN) or VTEP (point-to-multi-point including third party).
5. Security Policies: This dashboard provides a list of security policies you configured using Flow Segmentation. For information about Security Policies, see the *Flow Microsegmentation Guide*.

See "Network Connections" section for information on how to configure network connections.

Subnets (Overlay IP subnets), Virtual private clouds, floating IPs, and Connectivity are Flow Networking features. These features support flexible app-driven networking that focuses on VMs and applications instead of virtual LANs and network addresses. Flow Networking powers network virtualization to offer a seamless network experience with enhanced security. It is disabled by default. It is a software-defined network virtualization solution providing overlay capabilities for the on-premises AHV clusters.

Security policies drives the Flow Segmentation features for secure communications. See *Flow Microsegmentation Guide*.

Subnets

Manage subnets in the **List** view of **Subnets** dashboard in the **Network and Security** section.

To access the Subnets dashboard, select **Subnets** from the entities menu in Prism Central. The Subnets dashboard allows you to view information about the subnets configured for the registered clusters.

> Note: This section describes the information and options that appear in the Network and Security dashboard. See "" () for instructions on how to view and organize that information in a variety of ways.

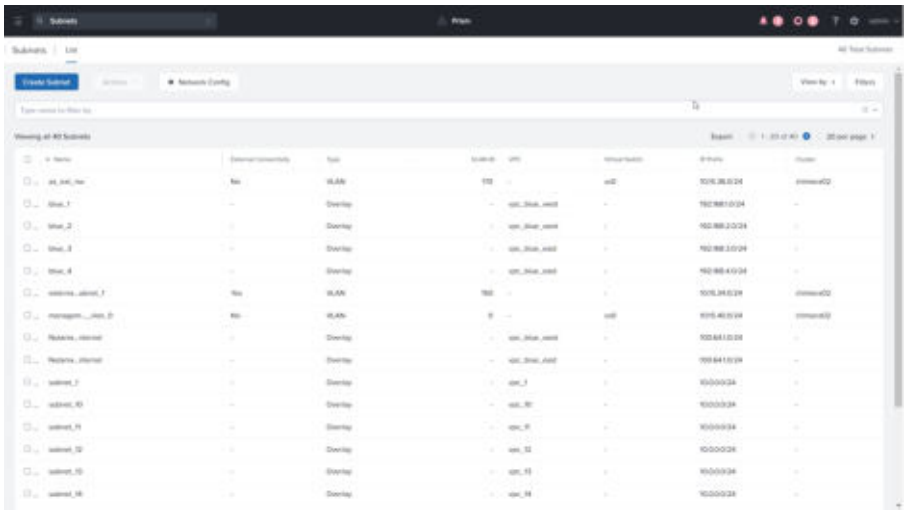

Figure 114: Subnets Dashboard

The following table describes the fields that appear in the subnets list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 63: Subnets Dashboard Fields

Parameter	Description	Values
Name	Displays the subnet name.	(subnet name)
External Connectivity	Displays whether or not the subnet has external connectivity configured.	(Yes/No)
Type	Displays the subnet type.	VLAN
VLAN ID	Displays the VLAN identification number.	(ID number)
VPC	Displays the name of the VPC that the Subnet is used in.	(Name of VPC)

Parameter	Description	Values
Virtual Switch	Displays the virtual switch that is configured for the VLAN you selected. The default value is the default virtual switch `vs0`. Note: The virtual switch name is displayed only if you add a VLAN ID in the VLAN ID field.	(virtual switch name)
IP Prefix	Displays the IPv4 Address of the network with the prefix.	(IPv4 Address/Prefix)
Cluster	Displays the name of the cluster for which this subnet is configured.	(cluster name)
Hypervisor	Displays the hypervisor that the subnet is hosted on.	(Hypervisor)

To filter the list by network name, enter a string in the filter field. (Ignore the **Filters** pane as it is blank.)

To view focused fields in the List, select the focus parameter from the **Focus** drop down list. You can create your own customised focus parameters by selecting **Add custom** from the drop down list and selecting the necessary fields after providing a **Name**, in the Subnet Columns.

There is a **Network Config** action button to configure a new network (see “” ()). The **Actions** menu appears when one or more networks are selected and includes a **Manage Categories** option (see “” ()).

Go to the Subnets list view by clicking **Network and Security** > **Subnets** on the left side-bar.

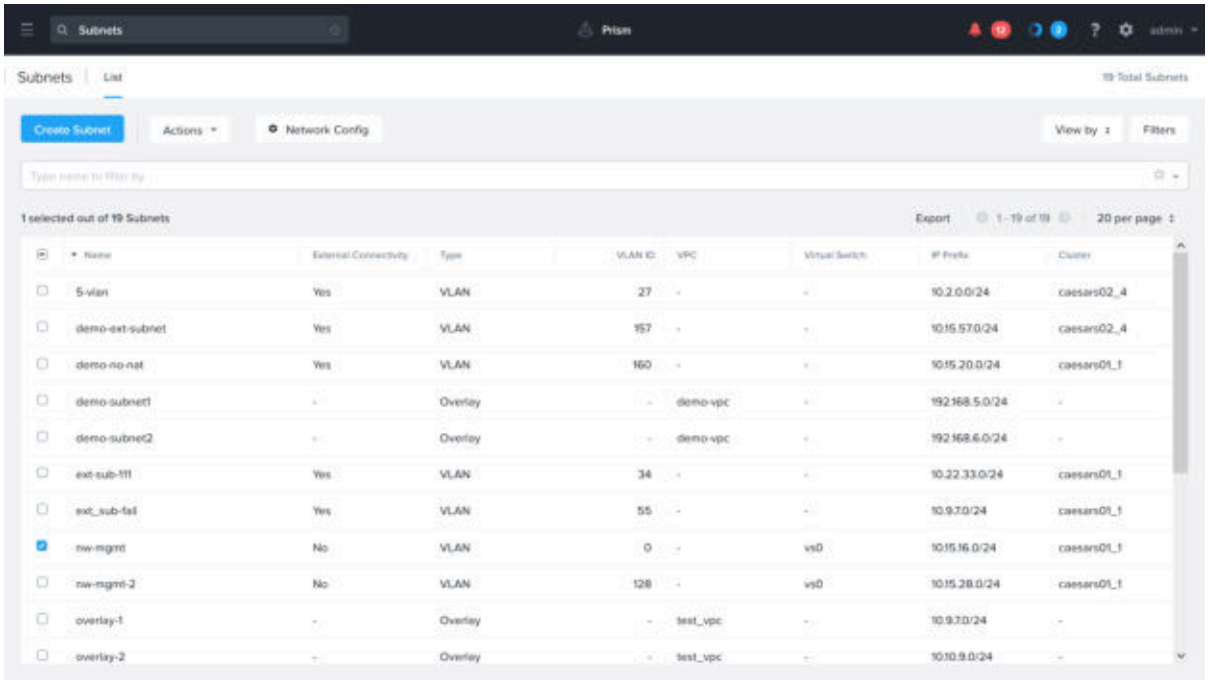

Figure 115: Subnets Page

To view or select actions you can perform on a subnet, select the subnet and click the **Actions** dropdown.

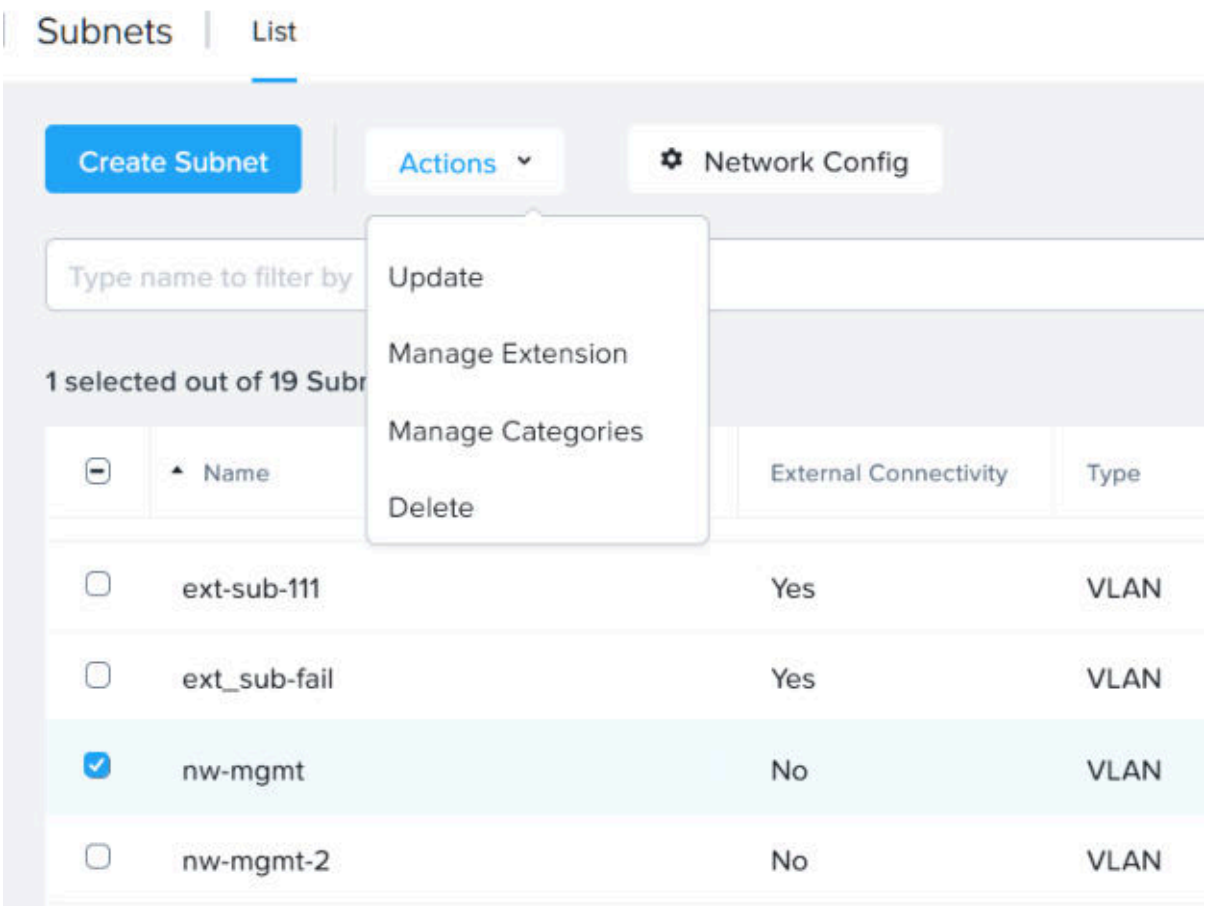

Figure 116: Subnet Actions

Table 64: Subnet Actions

Action	Description
Update	Click this action to update the selected subnet. see Updating a Subnet.
Manage Extension	Click this action to create a subnet extension. A subnet extension allows VMs to communicate over the same broadcast domain to a remote Xi availability zone (in case of Xi-Leap based disaster recovery) via the extension.
Manage Categories	Click this action to associate the subnet with a category or change the categories that the subnet is associated with.
Delete	Click this action to delete the selected subnet. See Deleting Subnets, Policies or Routes.

You can also filter the list of subnets by clicking the **Filters** option and selecting the filtering parameters.

Subnet Summary View

View the details of a subnet listed on the Subnets page.

To view the details of a subnet, click the name of the subnet on the subnet list view.

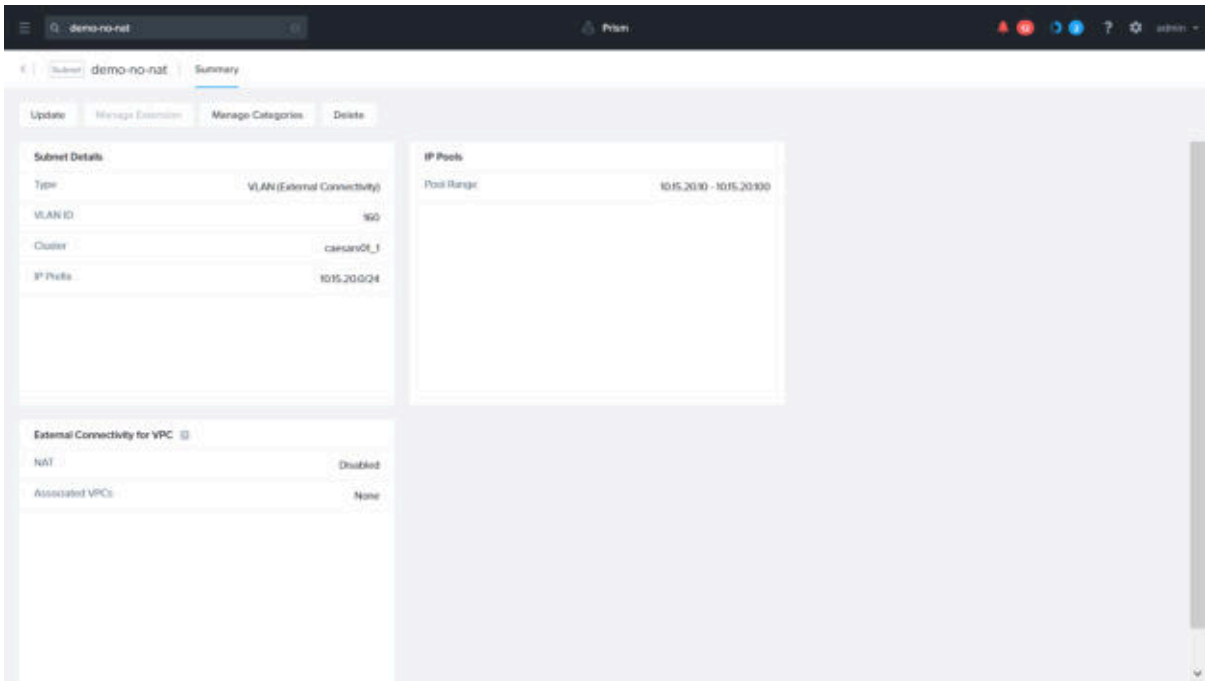

Figure 117: Subnet Summary Page

The Summary page provides buttons for the actions you can perform on the subnet, at the top of the page. Buttons for the following actions are available: **Update**, **Extend**, **Manage Categories**, and **Delete**.

The subnet Summary page has the following widgets:

Widget Name	Information provided
Subnet Details	Provides the following: • Type — Displays the type of network like VLAN or Overlay. • VLAN ID — Displays the VLAN ID. This parameter is displayed only for VLAN networks. • VPC — Displays the VPC name. This parameter is displayed only for Overlay networks. • Cluster — Displays the cluster that the VLAN network is configured on. This parameter is displayed only for VLAN networks. • IP Prefix — Displays the IP address prefix configured for the network. This parameter is displayed for both VLAN and Overlay networks.
IP Pool	Provides the IP address Pool Range assigned to the network.

Widget Name	Information provided
External Connectivity	Provides the following: • NAT — Displays whether NAT is enabled or disabled for VPCs connecting to the network. When you hover on the `Enabled`/`Disabled` status, the hover message displays details of VPCs connected to the external subnet. • Associated VPCs — Displays the VPCs associated with this external subnet.

Virtual Private Clouds

You can manage Virtual Private Clouds (VPCs) on the `Virtual Private Clouds` dashboard.

Go to the Virtual Private Clouds dashboard by clicking **Network and Security** > **Virtual Private Clouds** on the left side-bar.

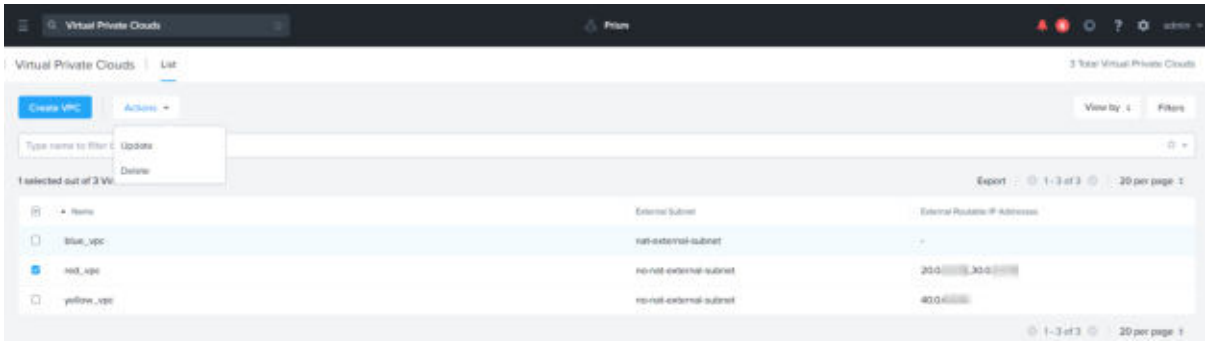

Figure 118: Virtual Private Clouds dashboard

You can configure the table columns for the VPC list table. The available column list includes `Externally Routable IP Addresses` that provides address space within the VPC that is reachable externally without NAT.. For the list of columns that you can add to the list table, see Customizing the VPC List View on page 169.

> Note:
>
> Ensure that the externally routable IP addresses (subnets with external connectivity without NAT) for different VPCs do not overlap.
>
> Configure the routes for the external connectivity subnets with next hop as the Router or SNAT IP address. Also configure the routes on the router for the return traffic to reach the VPC. See **External Connectivity** panel in VPC Details View on page 172.

To view or select actions you can perform on a VPC, select the VPC and click the **Actions** drop down.

You can also filter the list of VPC by clicking the **Filters** option and selecting the filtering parameters.

Customizing the VPC List View

About this task

You can customize the columns in the table. Click the **View by** drop down and select **+ Add custom**.

In the Virtual Network Columns dialog box, do the following.

Procedure

1. Enter a name for the view.

2. Select the columns you want displayed in the table.

   During the column selection, the columns you select are moved under the **Selected Columns** list. The ***Name*** (of the VPC) column is the default column already selected. You can add a maximum of 10 columns (including the ***Name*** column) to the **Selected Column** list.

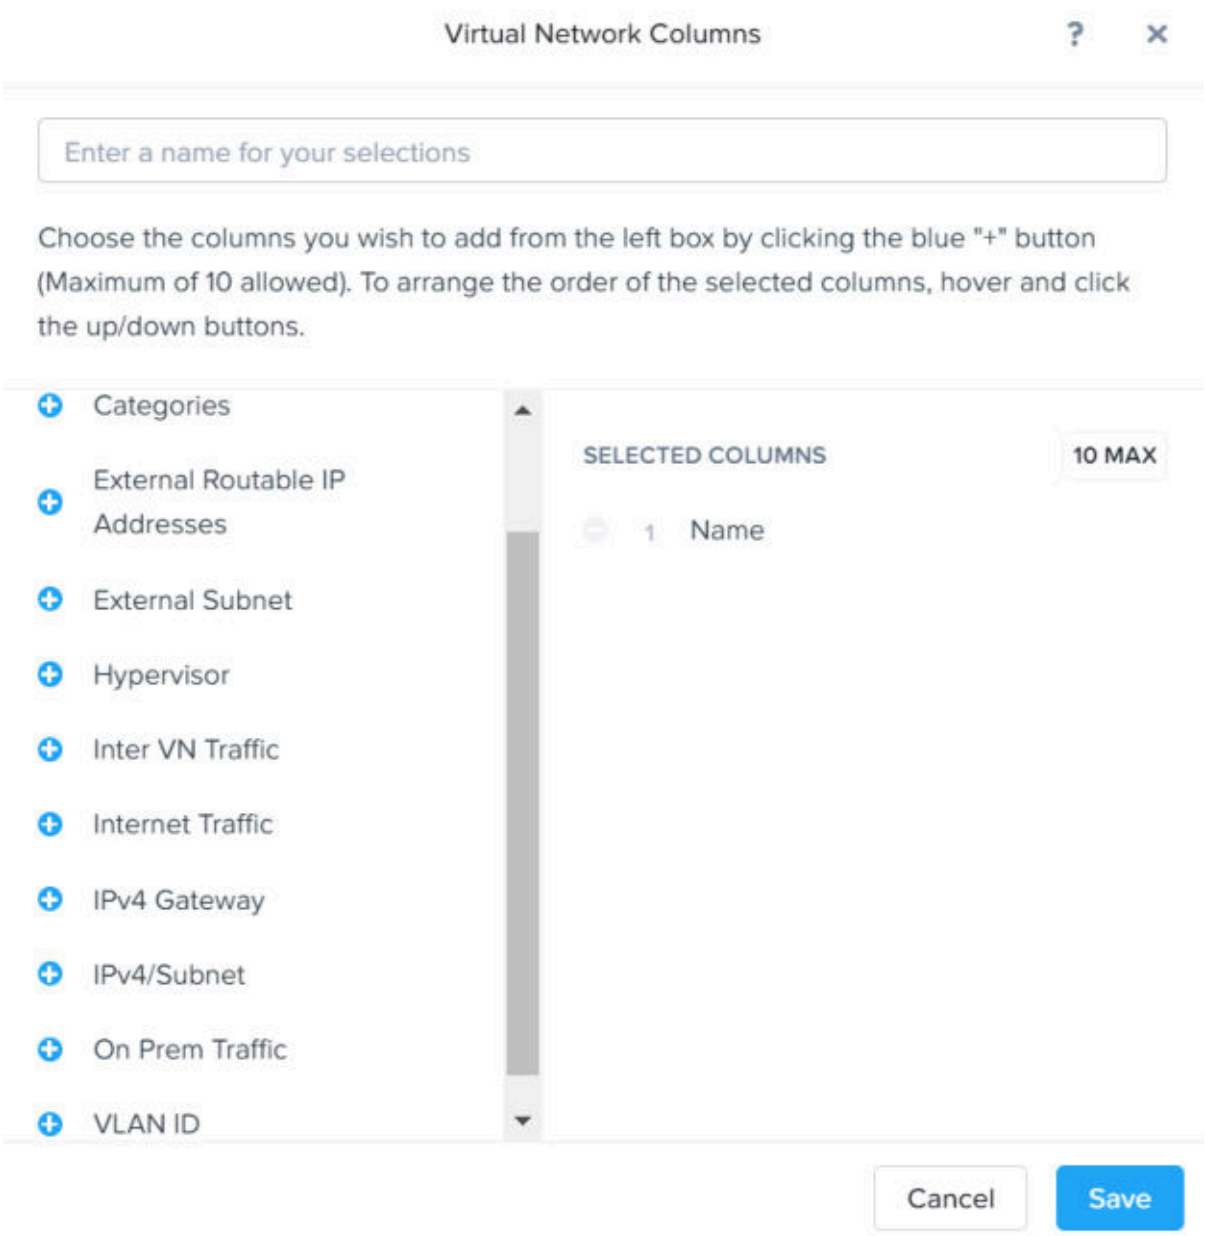

Figure 119: Customizing Columns in VPC View

   To arrange the order of the selected columns, hover on the column name and click the up or down arrow button as appropriate.

3. Click **Save**.

VPC Details View

To view the details of a VPC, click the name of the VPC on the VPC list view.

The VPC details view has the following tabs:

- **Summary**

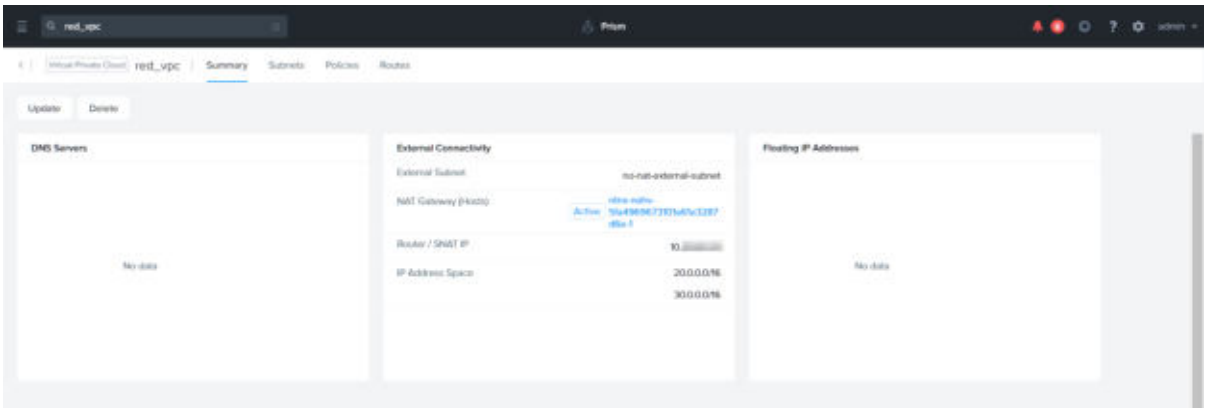

Figure 120: Summary Tab

The Summary tab provides the following panes:

- DNS Servers—Provides more information about the DNS Servers used by the VPC.
- External Connectivity—Provides the name of the external subnet, NAT Gateway host details, router/SNAT IP address and the IP address spaces or ranges configured for the VPC.
- Floating IP Addresses—Provides details of the floating IP addresses that the VPC uses.

- **Subnets**

Figure 121: Subnet Tab

The Subnet tab provides the following information for the subnets:

- Name—Displays the name of the subnet.
- IP Range—Displays the IP address range configured for the subnet.
- DHCP IP Pool—Displays the DHCP IP address pool configured for the subnet.
- Default Gateway IP—Displays the IP address used as the default gateway by the entities in the subnet.
- Actions—Displays the actionable links to **Edit** or **Delete** the subnet.

- **Policies**

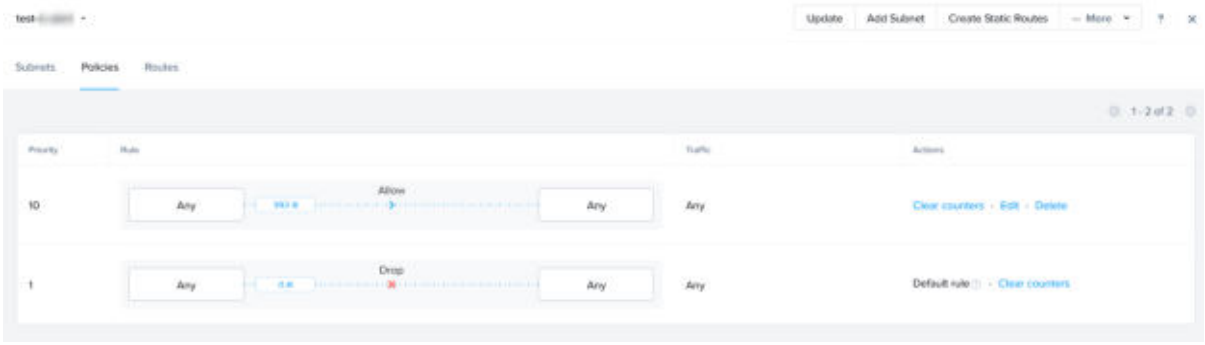

Figure 122: Policies Tab

The Policies tab maps the following information about the security-based traffic shaping policies you configure:

- Priority—The traffic priority.
- Rule—The **Allow** or **Deny** rule set for the priority.
- Traffic—The traffic type that the priority and rule should be applied to.
- Actions—Actions you can take on the policy. You can perform three actions: **Clear counters**, **Edit** the policy or **Delete** the policy.

- **Routes**

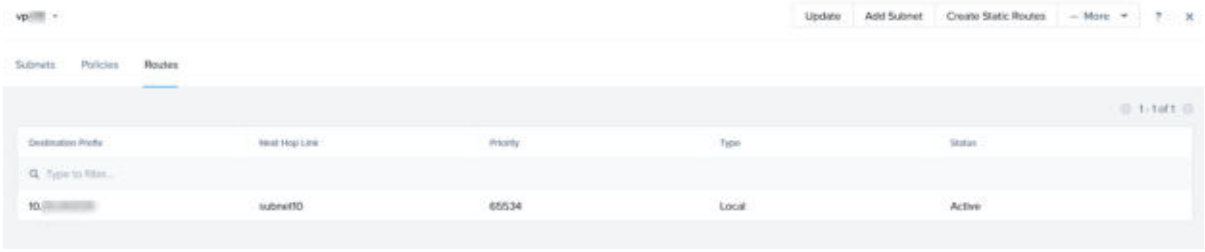

Figure 123: Routes Tab

The Router tab provides the following information about the routes:

The VPC details view has the following configuration options for the VPC:

- **Update**: Use this option to update the VPC. For more information, see Updating Virtual Private Cloud.
- **Add Subnet**: Use this option to add a subnet to the VPC. For more information, see Creating a Subnet.
- **Create Static Routes**: Use this option to create a static route. For more information, see Creating Static Routes.
- **Update Static Routes**: Use this option to update static route configurations that you already created. For more information, see Updating Static Routes.

- **Create Policy**: Use this option to create traffic policies in addition to the pre-configured default policy. When you create a VPC, there is one default policy that Advanced Networking creates for the VPC. This policy is pre-configured and cannot be edited.

  For more information, see Creating a Policy.

- **Clear All Counters**: Allows you to clear all the counters for the VPC.
- **Delete**: Allows you to delete the VPC. See Deleting a Virtual Private Cloud and Deleting Subnets, Policies or Routes.

Floating IPs

You can access floating IPs on the **`Floating IPs`** dashboard or list view in the **`Network and Security`** section.

For information about floating IP addresses and their role in Flow Networking, see *SNAT and Floating IP Address* in Essential Concepts.

Go to the Floating IPs dashboard by clicking **Network and Security** > **Floating IPs** on the left side-bar.

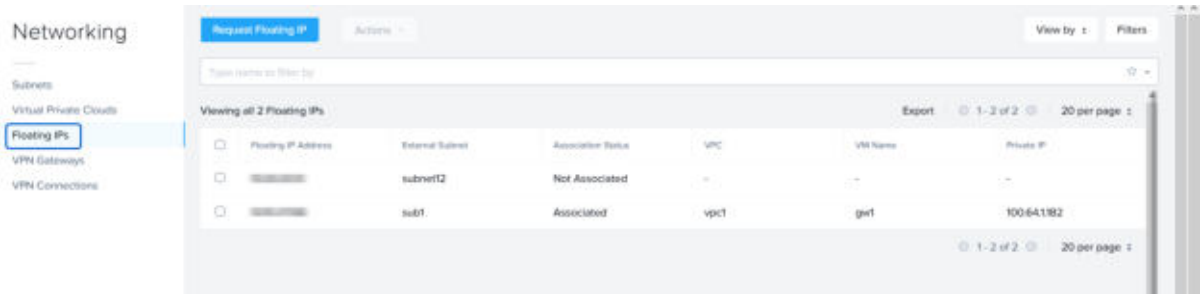

Figure 124: Floating IPs dashboard

To view or select actions you can perform on a floating IP address assigned, select the floating IP address and click the **Actions** drop down. The following actions are available for a selected floating IP address:

- Update—Assign or change the assignment of the floating IP address. You can assign the floating IP address to a IP address such as a private IP address in a VPC or the primary IP address of a VM or a secondary IP address created on a VM.
- Delete—Delete the floating IP address. The deleted IP address returns to the IP address pool as unused. Before you delete, ensure that it is not assigned to a private IP address or a VM. Change the assignment to **`None`** if it is already assigned, using the **Update** action.

  Note: Floating IP addresses are not reachable (Pings fail) unless you associate them to primary or secondary IP addresses of VMs. For information about assigning floating IP addresses to secondary IP addresses of VMs, see Assigning Secondary IP Addresses to Floating IPs.

To filter the list of floating IP address assignments, click the **Filters** option and select the appropriate filtering parameters.

To request floating IP addresses, see Requesting Floating IPs.

Connectivity

You can access network Gateways, VPN connections and subnet extensions on the **`Connectivity`** dashboard.

Click **Network & Security** > **Connectivity** to see the Connectivity dashboard.

The Connectivity dashboard opens on the Gateways tab. To see the VPN connections, click the VPN Connections tab. To see the subnets extended across AZs, click the Subnet Extensions tab.

Gateways Summary View

The Connectivity dashboard opens on the Gateways dashboard or summary view.

The **Gateway** dashboard provides a list of gateways created for the clusters managed by the Prism Central.

The **Gateways** dashboard provides a **Create Gateway** dropdown menu that lets you create a **Local** or a **Remote** gateway. You can create a local or remote gateway with **VPN** or **VTEP** service. See Creating a Network Gateway for more information.

You can select a gateway from the list (select the checkbox provided for the gateway) and then perform an action provided in the **Actions** dropdown list. The **Actions** dropdown list allows you to **Update** or **Delete** the selected gateway.

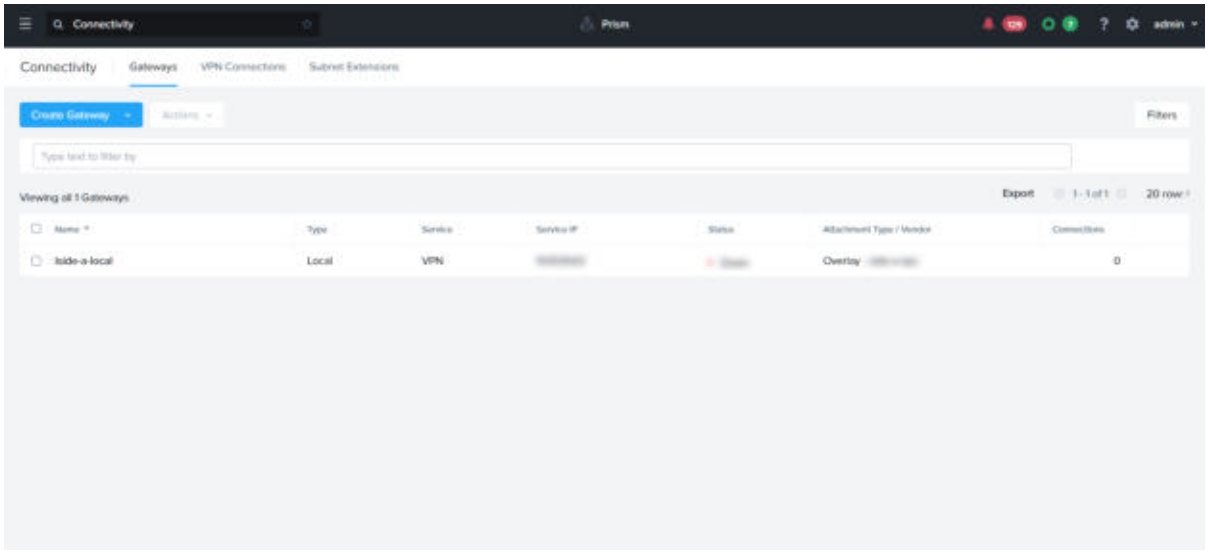

Figure 125: Gateways dashboard

The **Gateway** summary list view provides the following details about the gateway.

Table 65: Gateway List Fields

Parameter	Description	Values
Name	Displays the name of the gateway.	(Name of gateway)
Type	Displays the gateway type.	(Local or Remote)
Service	Displays the service that the gateway uses.	(VPN or VTEP)
Service IP	Displays the IP address used by the service.	(IP address)
Status	Displays the operational status of the gateway.	(Up or Down)

Parameter	Description	Values
Attachment Type/ Vendor	Displays the type of subnet associated with the gateway.	(VLAN or Overlay-VPC name)
Connections	Displays the number of service connections (such as VPN connections) configured and operational on the gateway.	(number)

You can click the name of a gateway to open the gateway details page that presents the information about the gateway in widgets.

Gateway Details View

You can click the name of a gateway in the **Gateway** dashboard list to open the gateway details page that presents the information about the gateway in widgets.

The gateway details page displays the name of the gateway on the top left corner.

- On the top right corner, the close button (X) allows you to close the details page.
- The **Update** button opens the **Update Gateway** page. See *Updating Gateways* for more information.
- The **Delete** button allows you to delete the gateway. See *Deleting Gateways* for more information.

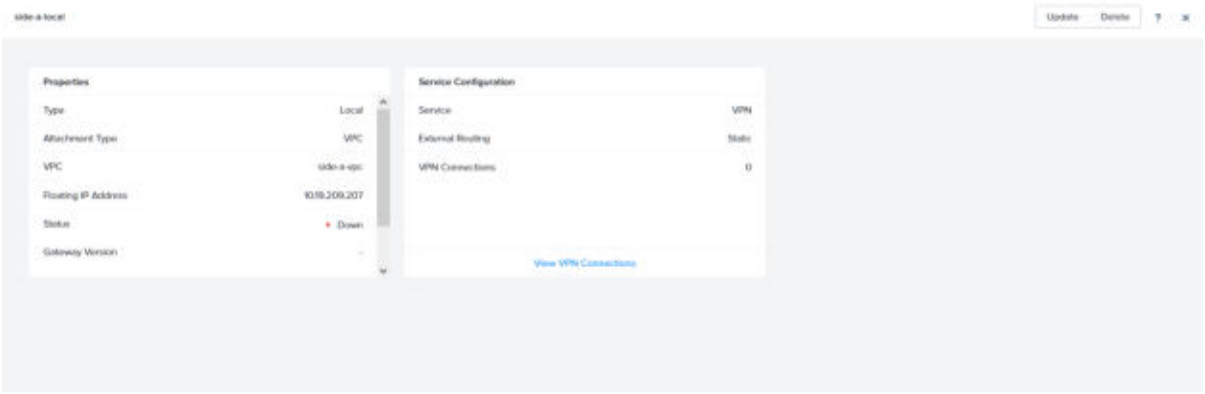

Figure 126: Gateway Details View

The details about the gateway are organized in widgets as follows:

Table 66: Gateway Details

Parameter	Description	Values
Properties widget		
Type	Displays the gateway type.	(Local or Remote)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)

Parameter	Description	Values
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address)
Status	Displays the operational status of the gateway.	(Up or Down)
Gateway Version	Displays the version of the Nutanix gateway appliance deployed.	(Version)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Service Configuration		
Service	Displays the service used by the gateway.	(VPN or VTEP)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
VPN Connections	Displays the total number of VPN connections associated with the gateway.	(Number - actionable link. Click the link to open the VPN connection details page for the associated VPN connection.)
View VPN Connections	Click this link to open the VPN Connections tab.	-

VPN Connections Summary View

The Connectivity dashboard allows you to open the VPN Connections dashboard or summary view.

VPN Connection: Represents the VPN IPSec tunnel established between local gateway and remote gateway. When you create a VPN connection, you need to select two gateways between which you want to create the VPN connection.

The **VPN Connections** dashboard provides a list of VPN connections created for the clusters managed by the Prism Central.

The **VPN Connections** dashboard provides a **Create VPN Connection** button that opens the **Create VPN Connection**. See Creating a VPN Connection for more information.

You can select a VPN connection from the list (select the checkbox provided for the VPN connection) and then perform an action provided in the **Actions** dropdown list. The **Actions** dropdown list allows you to **Update** or **Delete** the selected VPN connection.

The **VPN Connections** summary list view provides the following details about the VPN connection.

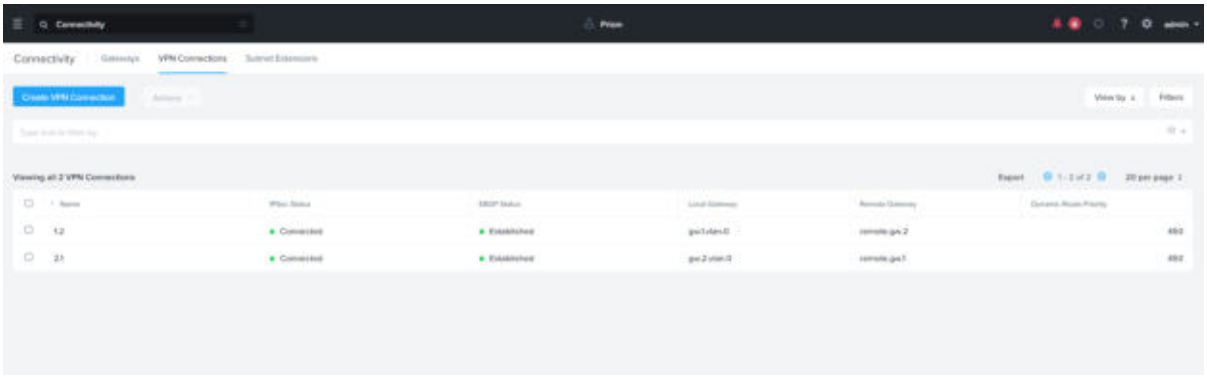

Figure 127: VPN Connections dashboard

Table 67: VPN Connections List Fields

Parameter	Description	Values
Name	Displays the name of the connection.	(gateway name)
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Local Gateway	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Remote Gateway	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Flow networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)

VPN Connections Details View

You can click the name of a VPN connection in the **VPN Connections** dashboard list to open the VPN connection details page that presents the information about the VPN connection in widgets.

The VPN connection details page displays the name of the VPN connection on the top left corner.

- On the top right corner, the close button (X) allows you to close the details page.
- The **Update** button opens the **Update VPN Connection** page. See Updating VPN Connection for more information.
- The **Delete** button allows you to delete the VPN connection. See Deleting a VPN Connection for more information.

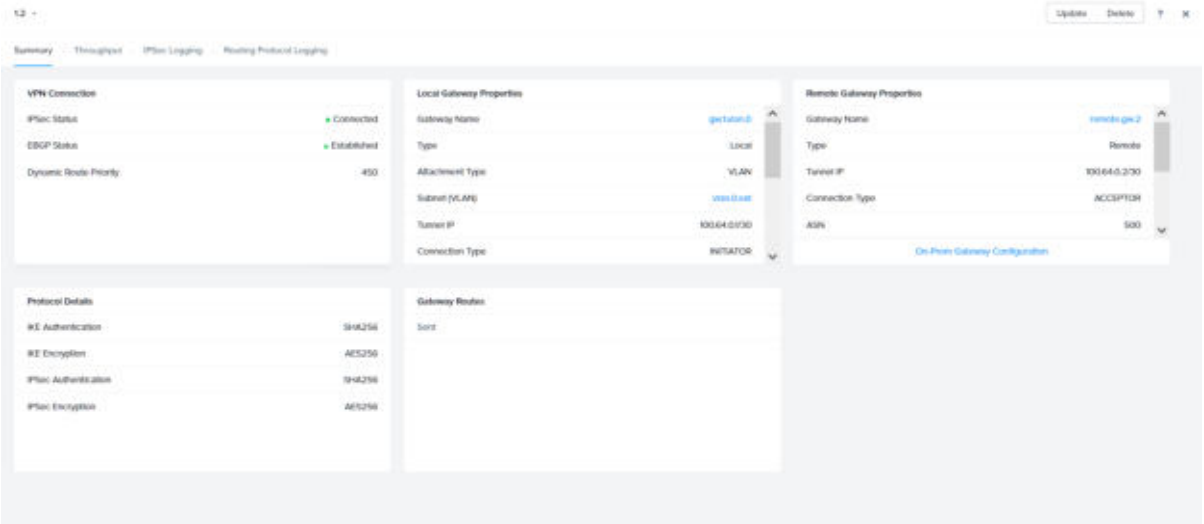

Figure 128: VPN Connection Details

The details about the VPN connection are organized in widgets as follows:

- Summary tab—See the *VPN Connection Summary Tab Details* table below.
- Throughput tab—See the *VPN Connection Throughput Tab Details* table below.
- IPSec Logging tab—Provides logs for the IPSec tunnel.
- Routing Protocol Logging tab—Provides logs for the routing protocol used in the VPN connection.

Table 68: VPN Connection Summary Tab Details

Parameter	Description	Values
VPN Connection widget		
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)

Parameter	Description	Values
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Flow networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)
Local Gateway Properties		
Gateway Name	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Type	Displays the type of gateway.	(Local)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Tunnel IP	Displays the Tunnel IP address of the local gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address that you assigned to the local gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	(Up or Down)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)

Parameter	Description	Values
Remote Gateway Properties		
Gateway Name	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Type	Displays the type of gateway.	(Remote)
Tunnel IP	Displays the Tunnel IP address of the remote gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
ASN	Displays the ASN of the EBGP route. This information is only displayed if you configured EBGP as the External Routing protocol.	(Number)
Vendor	Displays the name of the vendor of the gateway appliance at the remote site.	(Name of vendor of gateway appliance)
External IP	Displays the IP address assigned to remote the gateway.	(IP Address that you assigned to the remote gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	-
Protocol Details		
Service	Displays the service used by the gateway.	(VPN or VTEP)
Gateway Routes	Displays the status of the routes used by the gateways.	(Sent)

Viewing an Extended Subnet

The Connectivity dashboard opens on the Subnet Extensions dashboard or summary view.

About this task

Perform the following procedure to view an extended subnet at the on-prem site.

Procedure

1. Click the hamburger icon in the top-left corner of the **Dashboard**.
   The main feature list appears.

2. Click **Network & Security** > **Connectivity** > > **Subnet Extensions**.
   The Subnet Extensions tab displays a list of the extended subnets.

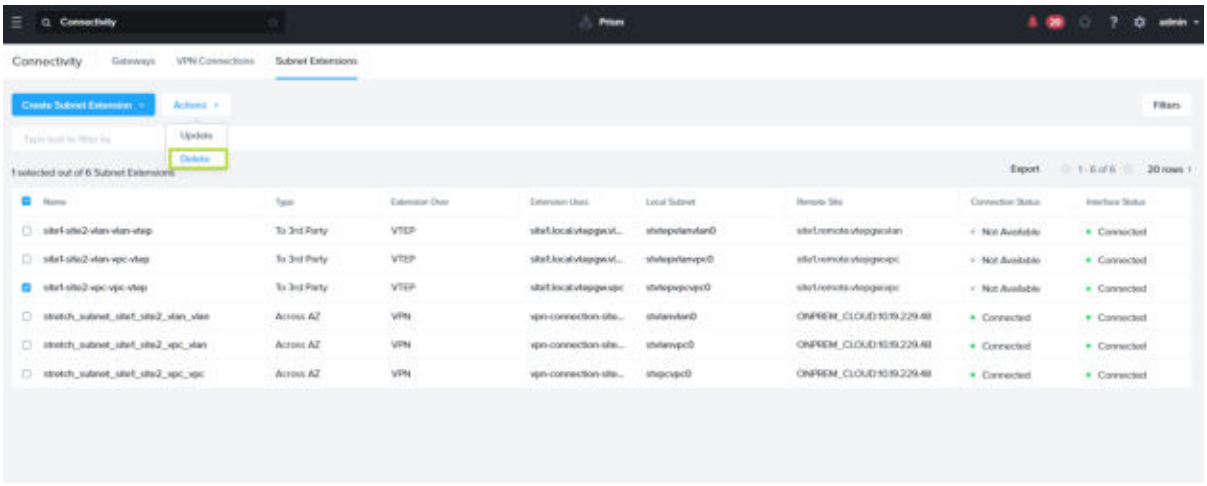

Figure 129: Extended Subnet

What to do next

You can also see which subnets are extended, in the Subnets dashboard. The **List** tab displays the list of configured subnets. The subnets that are extended have the **Extended Subnet** tag alongside their name.

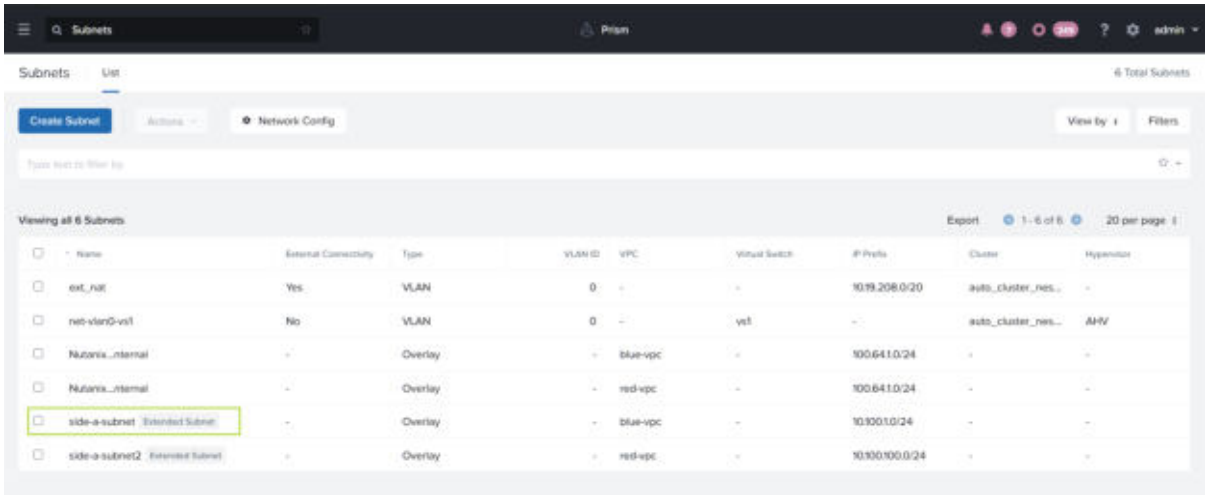

Figure 130: Subnet List View - Extended Subnets

Security Policies Summary View

To access the security policies dashboard, select **Policies** > **Security Policies** from the entities menu (see Entities Menu on page 13). The security policies dashboard allows you to view summary information about defined security policies.

Note: This section describes the information and options that appear in the security policies dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- See *Flow Microsegmentation Guide* for information about how to create and apply security policies.

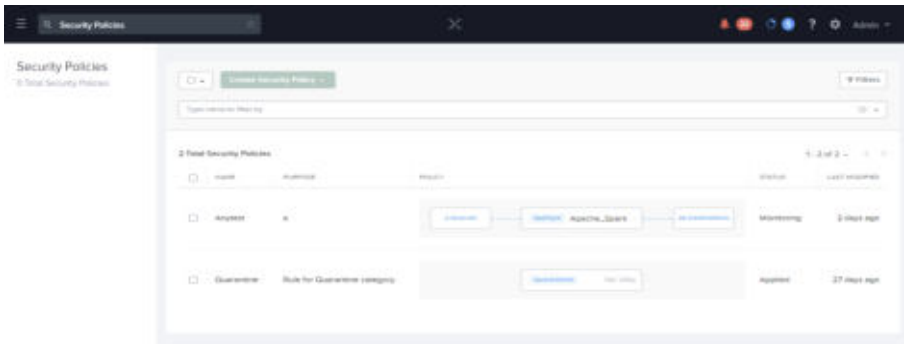

Figure 131: Security Policies Dashboard

The following table describes the fields that appear in the security policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 69: Security Policies List Fields

Parameter	Description	Values
Name	Displays the policy name. The policy is one of three types: application, quarantine, or isolation.	(name), Application, Quarantine, Isolation
Purpose	Describes (briefly) the policy's purpose.	(text string)
Policy	Displays (high level) what the policy does.	(boxed text)
Status	Displays the current status of the policy (either applied currently or in monitoring mode).	Applied, Monitoring
Last Modified	Displays the date the policy was last modified (or the creation date if the policy has never been modified).	(date)

You can filter the security polices list based on several parameter values. The following table describes the filter options available when you open the Security Policies view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 70: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of security policies that satisfy the name condition/string.	(policy name string)
Type	Filters on the policy type. Check the box for one or more of the policy types (application, quarantine, isolation). It will limit the list to just those policy types.	Application, Quarantine, Isolation
Status	Filters on the policy status. Check the box for applied or monitoring.	Applied, Monitoring

The security policies dashboard includes a **Create Security Policy** action button with a drop-down list to **Secure an Application** or **Isolation Environments**.

The **Actions** menu appears when one or more policies are selected. It includes options to update, apply, monitor, and delete. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Security Policy Details View

To access the details page for a security policy, click on the desired security policy name in the list (see Security Policies Summary View on page 182). The Security Policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The rule status appears below the name and indicates whether the policy is being applied currently or is in monitoring mode.
- Three columns appear that specify the Inbound policy (on the left), the affected entities (in the middle), and the Outbound policy (on the right).
- There are three action buttons (upper right).
  - Click the appropriate button to update, apply, monitor, or delete the policy (see *Nutanix Security Guide* for details). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

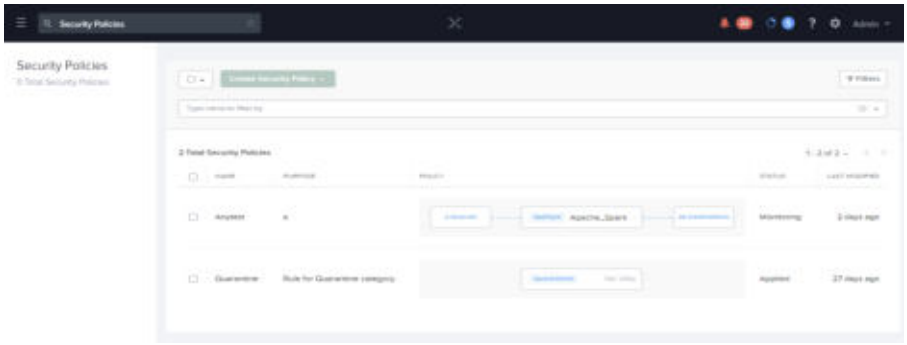

Figure 132: Security Policy Details View: Monitoring Rule Example

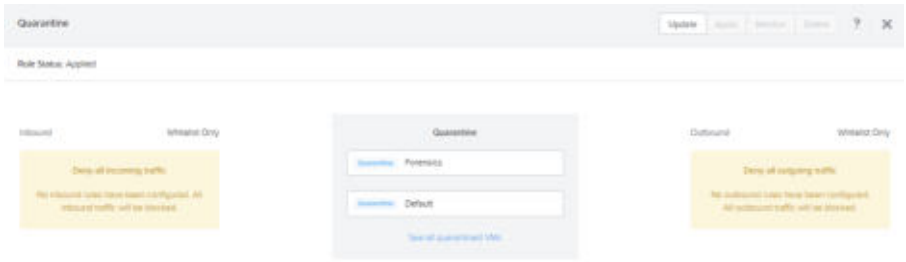

Figure 133: Security Policy Details View: Applied Rule Example

For more information about Security Policies, see the *Flow Microsegmentation Guide*.

Policies

You can access dashboards for the following policy types from the **VM**, **Images**, **Storage**, and **Network & Security** entities.

> Note: See the *Leap Administration Guide* for information about creating and applying protection policies and recovery plans.

VM Policies

- Affinity Policies (see Affinity Policies Summary View on page 186
- NGT Policies (see NGT Policies View on page 187)

Image Policies

- Image Placement Policies (see Image Placement Policies Summary View on page 188)
- Bandwidth Throttling Policies (see Bandwidth Throttling Policies Summary View on page 190)

Storage Policies

Storage Policies (see Storage Policies Summary View on page 133)

Security Policies

Storage Policies (see Security Policies Summary View on page 191)

Affinity Policies Summary View

To access the affinity policies dashboard, select **Compute & Storage** > **VMs** > **Policies** > **Affinity Policies** from the **Entities** menu (see Entities Menu on page 13).

> Note: This section describes the information and options that appear in the affinity policies dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.
> - See Affinity Policies Defined in Prism Central on page 599 for information on how to create or modify the affinity policies.

The affinity policies dashboard displays a list of current policies that includes the name and type for each policy.

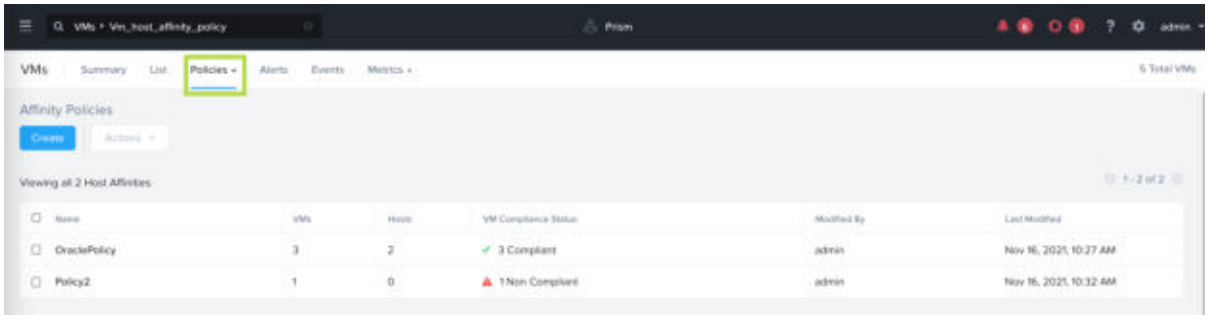

Figure 134: Affinity Policies Dashboard

The following table describes the fields that appear in the affinity policies list. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 71: Affinity Policies List Fields

Parameter	Description	Values
Name	Displays the policy name.	(name)
VMs	Displays the count of VMs associated with this policy.	(number of VMs)
Hosts	Displays the count of hosts associated with this policy.	(number of hosts)
VM Compliance Status	Displays the compliance status of the VMs associated with this policy. If the policy is being applied and the compliance status is not yet known, the status is displayed as Pending. If a VM is part of multiple VM-Host affinity policies, the oldest policy is applied on the VM. For rest of the policies, the VM is displayed as non-compliant.	(number of VMs Compliant/Non Compliant/Pending)

Parameter	Description	Values
Modified By	Displays the name of the user who modified the policy last time.	(user)
Last Modified	Displays the date and time when the policy was modified last time.	(date & time)

Affinity Policies Details View

To access the details page for an Affinity policy, click the desired policy name in the list (see Affinity Policies Summary View on page 186). The affinity policy details page includes the following tabs:

- **Summary**: On the **Summary** tab, you can view the **Overview**, **Associations**, and **Compliance Status** sections. All the three sections display information related to the policy.

  The **Summary** tab view also includes options to **Update**, **Delete**, and **Re-Enforce** the policy. If any of the VMs become non-compliant, you can use this option to re-enforce the policy after fixing the issue for the non-compliance.

- **Entities**: On the **Entities** tab, you can view the details of the **VMs** and **Hosts** entities that are associated with the policy. The details displayed on the **Hosts** tab include the host name, cluster name, and the category used to associate, while the details displayed on the **VMs** tab include the VM name, host name, cluster name, category used to associate, and VM compliance status. If a VM is non-compliant, the cause of non-compliance is also displayed along with the status.

Figure 135: Affinity Policies Details view

NGT Policies View

To access the Nutanix guest tools (NGT) policies dashboard, select **Compute & Storage** > **Policies** > **NGT Policies** from the dowp-down menu (see Entities Menu on page 13). The NGT policies dashboard allows you to view information about existing NGT policies and create new policies.

> Note: This section describes the information and options that appear in the NGT policies dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See NGT Policies on page 602 for information on how to create or modify NGT policies.

The NGT dashboard displays a list of current policies that includes the name and type for each policy.

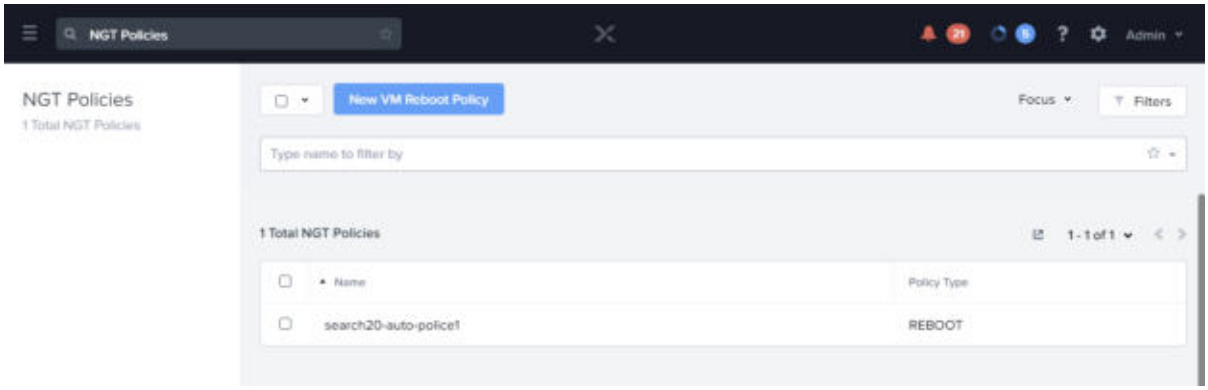

Figure 136: NGT Policies Dashboard

To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following parameters and values.

- **Name**: Enter name in search field.
- **Created By**: System, User

There is a **New VM Reboot Policy** action button to create a new reboot policy (see Creating a New Restart Policy on page 602). The **Actions** menu appears when one or more policies are selected and includes the following options:

- **Update** (see Modifying a Restart Policy on page 604)
- **Delete**

Image Placement Policies Summary View

The image placement policies dashboard allows you to view summary information about defined image placement policies.

To access the image placement policies dashboard, select **Images** > **Policies** > **Image Placement** from the **Entities** menu (see Entities Menu on page 13).

Note: This section describes the information and options that appear in the image placement policies dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- See Image Placement Policies on page 680 for information about how to create and apply image placement policies.

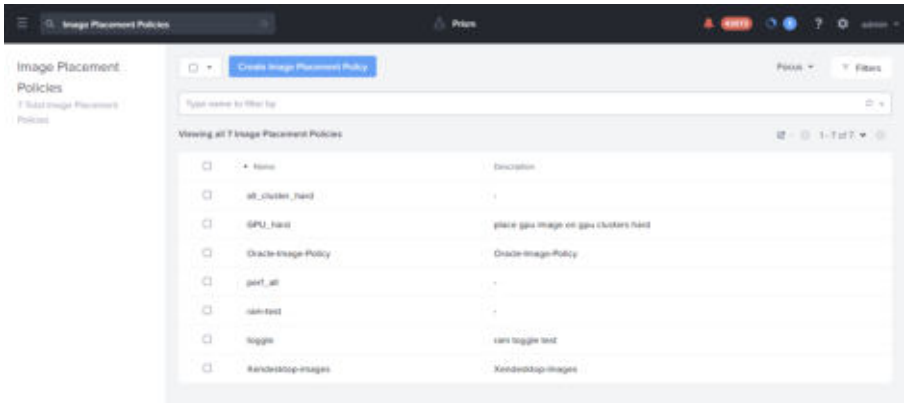

Figure 137: Image Placement Policies Dashboard

The following table describes the fields that appear in the image placement policies list. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 72: Image Placement Policies List Fields

Parameter	Description	Values
Name	Displays the policy name.	(name)
Description	Describes (briefly) the purpose of the policy.	(text string)

You can filter the image placement polices list by opening the **Filter** pane. To apply a filter, check the box and enter a text string in the **Contains** field for the Name or Description parameter. You can apply filters across both parameters.

The image placement policies dashboard includes a **Create Image Placement Policy** action button (see Configuring an Image Placement Policy on page 682).

The **Actions** menu appears when one or more policies are selected. It includes options to **Update** and **Delete**. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Image Placement Policy Details View

To access the details page for an image placement policy, click on the desired policy name in the list (see Image Placement Policies Summary View on page 188). The image placement policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The policy enforcement status appears below the name and indicates whether the enforcement is hard or soft.
- Two columns appear that specify the images in question (on the left) and the target clusters (on the right). Click **Expand all** (or **Collapse all**) to display (or hide) the list of images or clusters.

- There are four action buttons (upper right).
  - Click the **Update** or **Delete** button to update or delete the policy (see Updating an Image Placement Policy on page 682). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

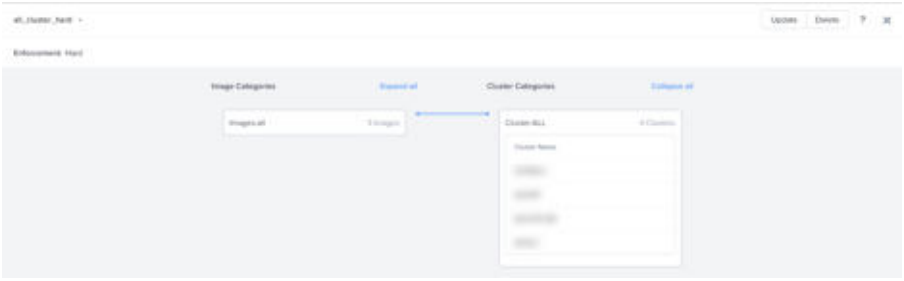

Figure 138: Image Placement Policy Details View

Bandwidth Throttling Policies Summary View

The bandwidth throttling policy allows you to limit the bandwidth consumed during image creation using the URL option in specific clusters. The bandwidth throttling policies dashboard allows you to view summary information about user-defined bandwidth throttling policies.

To access the bandwidth throttling policies dashboard, select **Images** > **Policies** > **Bandwidth Throttling Policies** from the **Entities** menu (see Entities Menu on page 13).

Note: This section describes the information and options that appear in the bandwidth throttling policies dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.
- See Bandwidth Throttling Policies on page 685 for information about how to create and apply bandwidth throttling policies.

Figure 139: Bandwidth Throttling Policies Dashboard

The following table describes the fields that appear in the image bandwidth throttling policies list. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 73: Bandwidth Throttling Policies List Fields

Parameter	Description
Policy Name	Displays the policy name.
Description	Describes (briefly) the purpose of the policy.
Bandwidth Throttle Limit	Describes the bandwidth throttle limit configured for the policy.
Last Updated On	Describes the date and time when the bandwidth throttling policy was created or updated.

Click the check box associated with the policy name to update or delete the bandwidth throttling policy.

Select specific policies and click **Actions** > **Update** to update the selected bandwidth throttling policies.

Select **Actions** > **Delete** to delete the selected bandwidth throttling policies.

You can filter the bandwidth throttling polices list by opening the **Filter** pane. To apply a filter, check the box and enter a text string in the **Contains** field for the name or description parameter. You can apply filters across both parameters.

The bandwidth throttling policies dashboard includes a **Create Policy** action button (see Creating Bandwidth Throttling Policy on page 686).

Security Policies Summary View

To access the security policies dashboard, select **Policies** > **Security Policies** from the entities menu (see Entities Menu on page 13). The security policies dashboard allows you to view summary information about defined security policies.

Note: This section describes the information and options that appear in the security policies dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- See *Flow Microsegmentation Guide* for information about how to create and apply security policies.

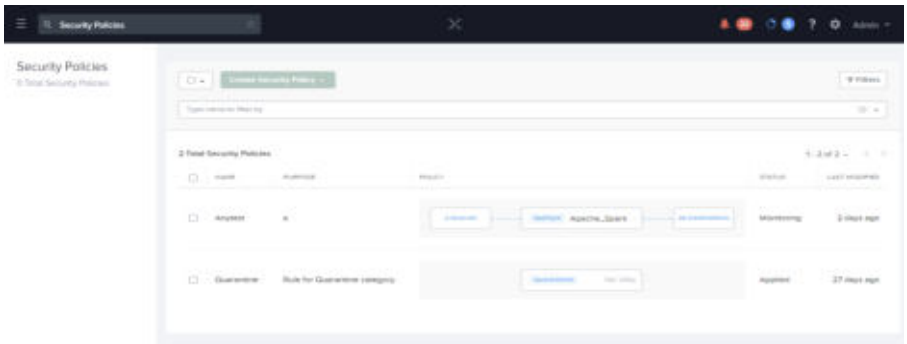

Figure 140: Security Policies Dashboard

The following table describes the fields that appear in the security policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 74: Security Policies List Fields

Parameter	Description	Values
Name	Displays the policy name. The policy is one of three types: application, quarantine, or isolation.	(name), Application, Quarantine, Isolation
Purpose	Describes (briefly) the policy's purpose.	(text string)
Policy	Displays (high level) what the policy does.	(boxed text)
Status	Displays the current status of the policy (either applied currently or in monitoring mode).	Applied, Monitoring
Last Modified	Displays the date the policy was last modified (or the creation date if the policy has never been modified).	(date)

You can filter the security polices list based on several parameter values. The following table describes the filter options available when you open the Security Policies view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 75: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of security policies that satisfy the name condition/string.	(policy name string)
Type	Filters on the policy type. Check the box for one or more of the policy types (application, quarantine, isolation). It will limit the list to just those policy types.	Application, Quarantine, Isolation
Status	Filters on the policy status. Check the box for applied or monitoring.	Applied, Monitoring

The security policies dashboard includes a **Create Security Policy** action button with a drop-down list to **Secure an Application** or **Isolation Environments**.

The **Actions** menu appears when one or more policies are selected. It includes options to update, apply, monitor, and delete. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Security Policy Details View

To access the details page for a security policy, click on the desired security policy name in the list (see Security Policies Summary View on page 191). The Security Policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The rule status appears below the name and indicates whether the policy is being applied currently or is in monitoring mode.
- Three columns appear that specify the Inbound policy (on the left), the affected entities (in the middle), and the Outbound policy (on the right).
- There are three action buttons (upper right).
  - Click the appropriate button to update, apply, monitor, or delete the policy (see “” () for details). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

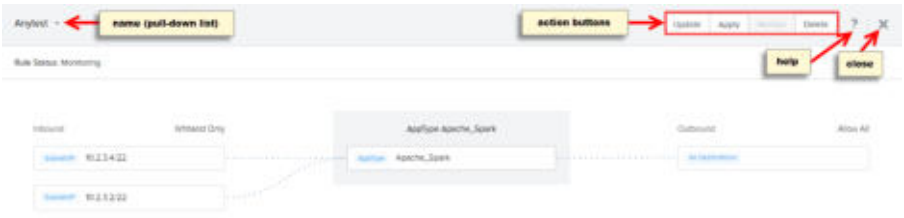

Figure 141: Security Policy Details View: Monitoring Rule Example

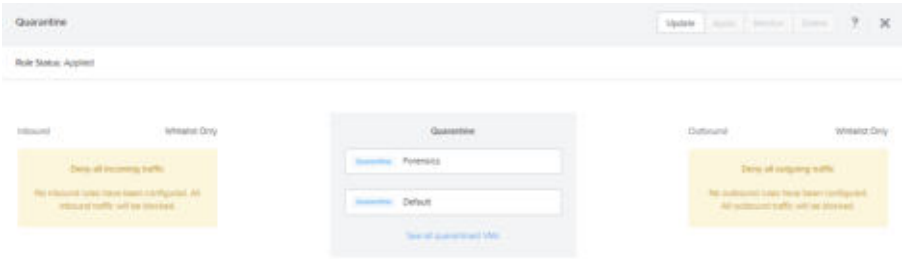

Figure 142: Security Policy Details View: Applied Rule Example

Data Protection and Recovery Entities

You can access dashboards for the following data protection types from the **Data Protection** category of the entities menu (see Entities Menu on page 13):

- Protection Summary
- Protection Policies (see Protection Policies View on page 194)
- Recovery Plans (see Recovery Plans View on page 195)

- VM Recovery Points (see VM Recovery Points View on page 196)

Protection Policies View

To access the protection policies dashboard, select **Policies** > **Protection Policies** from the entities menu (see Entities Menu on page 13). The protection policies dashboard allows you to view information about the current policies and create new policies.

> Note: This section describes the information and options that appear in the protection policies dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See the *Leap Administration Guide* for information about how to create and use protection policies.

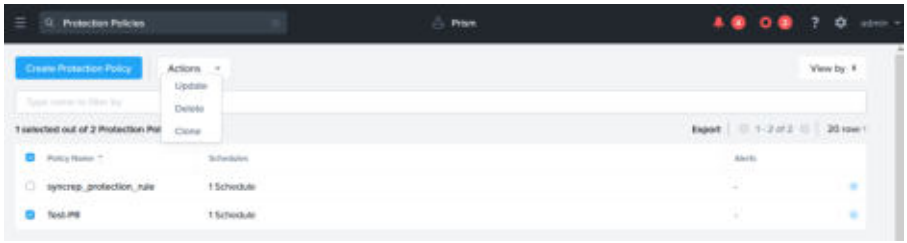

Figure 143: Protection Policies Dashboard

The following table describes the fields that appear in the protection policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 76: Protection Policies Fields

Parameter	Description	Values
Name	Displays the protection policy name.	(item name)
Source	Displays the replication source name for the protection policy	(source name)
Destination	Displays the replication destination name for the protection policy.	(destination name)
RPO	Displays the recovery point objective (RPO) for the protection policy	(time interval)
Remote Retention	Displays the number of retention points at the remote availability zone.	*xx* Recovery Points
Local Retention	Displays the number of retention points at the local availability zone.	*xx* Recovery Points

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

Click the **Create Protection Policy** button to create a new protection policy. The **Actions** menu appears when one or more protection policies are selected and includes the following options:

- **Update**: Update the protection policy.
- **Clone**: Clone the protection policy.
- **Delete**: Delete the protection policy.

Recovery Plans View

To access the recovery plans dashboard, select **Policies** > **Recovery Plans** from the entities menu (see Entities Menu on page 13). The recovery plans dashboard allows you to view information about the current plans and create new plans.

> Note: This section describes the information and options that appear in the recovery plans dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See the *Leap Administration Guide* for information about how to create and use recovery plans.

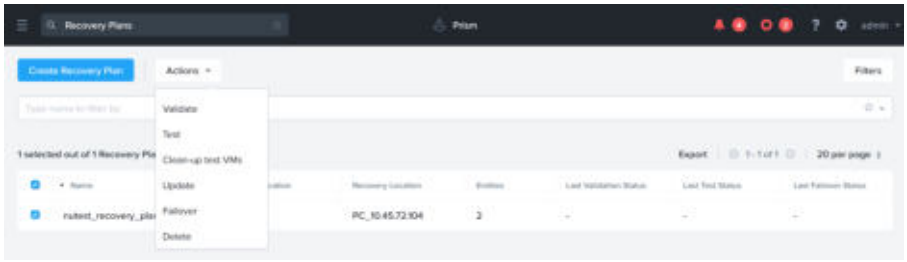

Figure 144: Recovery Plans Dashboard

The following table describes the fields that appear in the recovery plans list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 77: Recovery Plans Fields

Parameter	Description	Values
Name	Displays the recovery plan name.	(name)
Source	Displays the primary availability zone for the recovery plan.	(zone name)
Destination	Displays the recovery availability zone for the recovery plan.	(zone name)
Entities	Displays the number of VMs and snapshots associated with the recovery plan in the availability zone in which you are viewing the recovery plan.	(integer)

Parameter	Description	Values
Last Validation Status	Displays the status of the most recent validation of the recovery plan.	
Last Test Status	Displays the status of the most recent test performed on the recovery plan.	

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

Click the **Create Recovery Plan** button to create a new recovery plan. The **Actions** menu appears when one or more recovery plans are selected and includes the following options:

- **Validate**: Clone a recovery plan.
- **Test**: Test the recovery plan.
- **Update**: Update the recovery plan.
- **Failover**: Perform a failover operation.
- **Delete**: Delete the recovery plan.

VM Recovery Points View

To access the recoverable entities dashboard, select **Data Protection & Recovery** > **VM Recovery Points** from the entities menu (see Entities Menu on page 13). A recoverable entity is a VM that you can recover from a snapshot. The recoverable entities dashboard allows you to view information about the entities for which snapshots exist in the availability zone.

> Note: This section describes the information and options that appear in the recoverable entities dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See the *Xi Leap Administration Guide* for information about recoverable entities and how to use them.

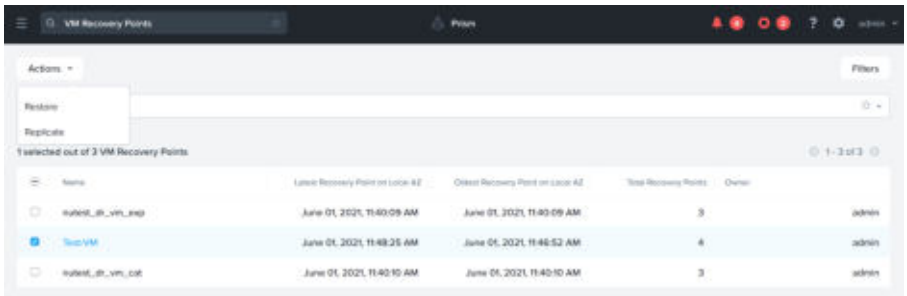

Figure 145: VM Recovery Points Dashboard

The following table describes the fields that appear in the recoverable entities list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 78: VM Recovery Points List Fields

Parameter	Description	Values
Name	Displays the recoverable entity name.	(item name)
Latest Recovery Point on Local AZ	Displays the time at which the most recent snapshot was taken. This is specific to the local availability zone.	(time and date)
Oldest Recovery Point on Local AZ	Displays the time at which the oldest available snapshot was taken. This is specific to the local availability zone.	(time and date)
Total Recovery Points.	Displays the total number of recovery points available locally.	(integer)
Owner	Displays the owner who created the recovery points	(user name)

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

The **Actions** menu appears when one or more recoverable entities are selected and includes the following options:

- **Restore**: Restore a VM entity from the snapshot taken at the recovery point.
- **Replicate**: Replicate the snapshot to a paired availability zone.

Hardware Entities

You can access dashboards for the following hardware components from the **Hardware** category of the entities menu (see Entities Menu on page 13):

- Clusters (see Clusters Summary View on page 197)
- Hosts (see Hosts Summary View on page 217)
- Disks (see Disks Summary View on page 227)
- GPUs (see GPUs Summary View on page 234)

Clusters Summary View

To access the clusters dashboard, select **Hardware** > **Clusters** from the entities menu (see Entities Menu on page 13). The clusters dashboard allows you to view summary information about registered clusters and access detailed information about each cluster. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> Note: This section describes the information and options that appear in the clusters dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities**: Displays a list of the clusters with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **IOPS**, and **Aggregate Hypervisor Memory Usage Pp**. Click the **View all XX Clusters** link at the bottom to display the **List** tab (following section).

- **Alerts**: Displays a list of cluster-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.

- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) , **Last 1 Hours**, or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

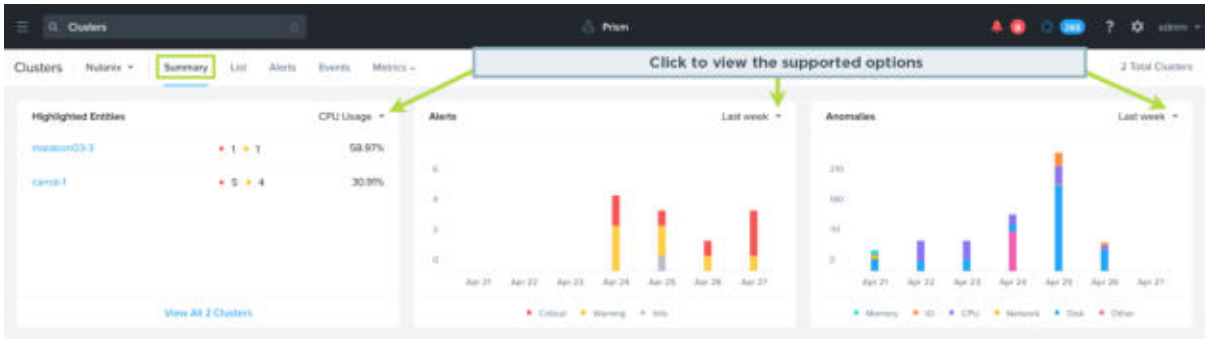

Figure 146: Clusters Summary Tab (Nutanix)

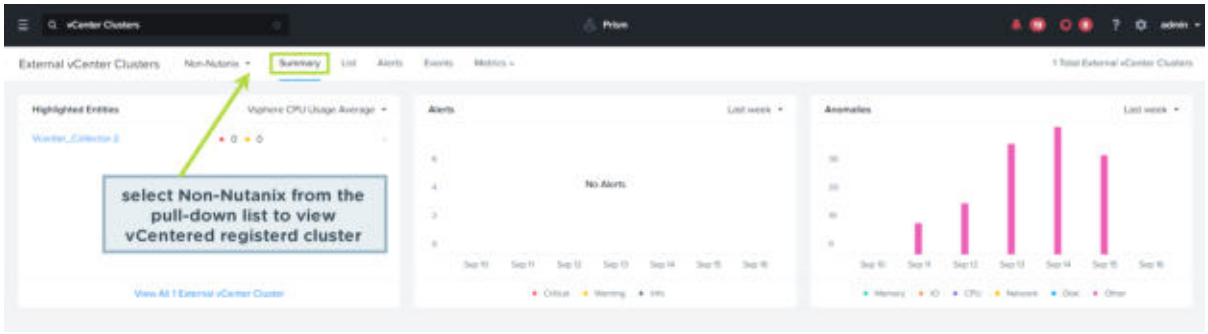

Figure 147: Clusters Summary Tab (Non-Nutanix)

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the registered clusters. The following table describes the fields that appear in the clusters list. The fields vary based on the **View by** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

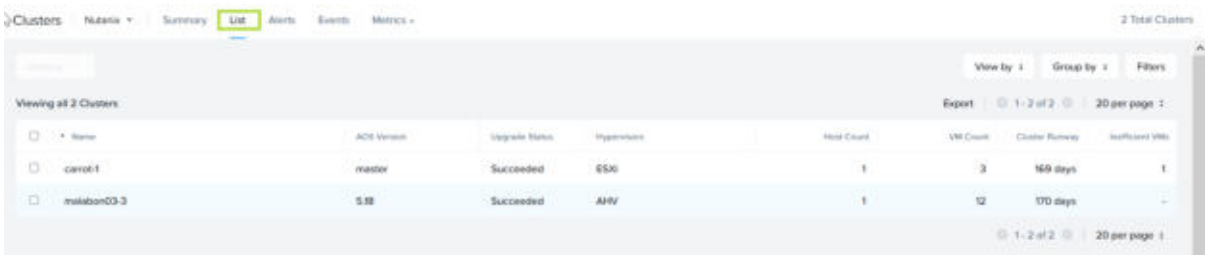

Figure 148: Clusters List Tab (Nutanix)

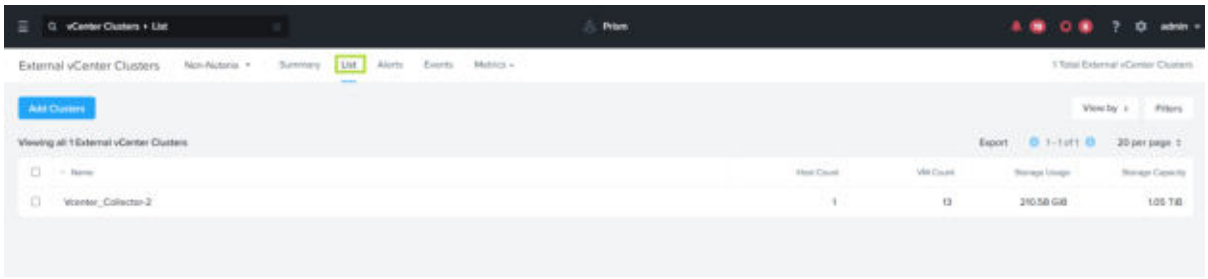

Figure 149: Clusters List Tab (Non-Nutanix)

Table 79: Clusters List Fields

Parameter	Description	Values
View by "General" Fields		
Name	Displays the cluster name. Clicking on the name displays the details page for that cluster (see Cluster Details View on page 205).	(cluster name)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Upgrade Status	Displays the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, Hyper-V
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)
VM Count	Displays the total number of VMs in the cluster (in any state).	(number of VMs)

Parameter	Description	Values
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 205 for more information).	(number of days)
Inefficient VMs	Displays the number of inefficient VMs in the cluster (see Behavioral Learning Tools on page 838 for more information).	(number)
View by "Performance" Fields		
Name	Displays the cluster name.	(cluster name)
CPU Usage	Displays the percentage of CPU capacity in the cluster currently being used.	0 -100%
Memory Usage	Displays the percentage of memory capacity in the cluster currently being used.	0 -100%
IOPS	Displays total (both read and write) I/O operations per second (IOPS) for this cluster.	(number)
IO Bandwidth	Displays total I/O bandwidth used per second in this cluster.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency in this cluster.	xxx [ms]
View by "Capacity" Fields		
Name	Displays the cluster name.	(cluster name)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 205 for more information).	(number of days)
CPU Runway	Displays the predicted CPU runway (time period) before the cluster requires additional resources.	(number of days)
Memory Runway	Displays the predicted memory runway (time period) before the cluster requires additional resources.	(number of days)
Storage Runway	Displays the predicted storage memory runway (time period) before the cluster requires additional resources.	(number of days)
View by "Encryption" Fields		
Name	Displays the cluster name.	(cluster name)

Parameter	Description	Values
Encryption Type	Displays the cluster encryption type.	Not Encrypted, Software, Hardware, Software and Hardware
Encryption Scope	Displays the encryption scope.	Cluster, Entity
KMS Type	Displays the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)

You can filter the cluster list based on a variety of parameter values. The following table describes the filter options available when you open the Clusters view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 80: Filter Pane Fields

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the pull-down list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)
Name	Filters on the cluster name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of clusters that satisfy the name condition/ string.	(cluster name string)
AOS Version	Filters on AOS version. Select one or more versions to return a list of clusters running those version(s). The number of clusters currently running each version is displayed on the right of the line.	(Acropolis version numbers across clusters currently)
Hypervisors	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, HyperV
Health	Filters on the cluster health state (good, warning, or critical). Select one or more states to return a list of clusters in that state(s). The number of clusters currently in each state is displayed on the right of the line.	Critical, Warning, Good

Parameter	Description	Values
CPU Usage	Filters on the amount of total CPU being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
IOPS	Filters on the total (both read and write) IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of clusters with total IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the total I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of clusters with total I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of clusters with average I/O latency in that range.	([xx] to [yy] range)
Upgrade Status	Filters on the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Categories	Filters on the category type. Search for the category name. For example, ADGroup: $Default	(category name string)
Encryption Type	Filters on the encryption type. Check the box for the desired encryption type. It will return a list of clusters with the corresponding encryption type applied.	Not Encrypted, Software, Hardware, Software and Hardware
Encryption Scope	Filters on the encryption scope. Check the box for the entities where encryption is applied on. It will return a list of clusters with the corresponding encryption scope applied.	Cluster, Entity
KMS Type	Filters on the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External

You can group the clusters list in the following ways:

- The **Group by** pull-down menu allows you to group the cluster entries by AOS version, host count, or health state. (You can only choose one.)
- When you select one or more clusters, the **Label** icon appears (to the left of the **Actions** menu). To assign the selected clusters a group label, click the **Label** icon and do one of the following:

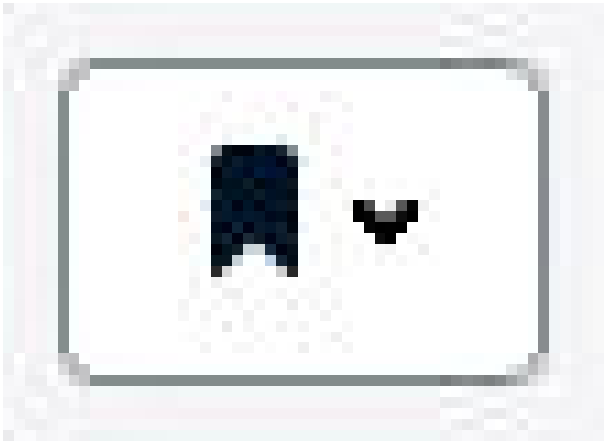

Figure 150: Label Icon

  - To assign an existing label (if there are existing labels), select the desired label and then click **Apply changes**.
  - To assign a new label, enter a meaningful label in the blank text field and then click **Create new label**.

  You can now perform any actions available on this labeled group from the **Actions** menu (see following section). You can edit or delete the label by clicking the **Label** icon and selecting **Manage Labels**.

The **Actions** menu appears when a cluster is selected. The pull-down list includes the following actions:

- Select **Launch Prism Element** to launch Prism element for that cluster in a separate tab or window (depending on your browser settings).

  Note: When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user drop-down list are **REST API Explorer**, **About Nutanix**, **Support Portal**, **Help**, **Nutanix Next Community**, and **Sign Out**.

-
- Select **Upgrade Software** to upgrade the AOS version on that cluster (see Installing or Upgrading on page 30).
- Select **Rack Configuration** to configure the rack awareness feature. This option appears only for clusters that satisfy the conditions for rack awareness. See the *Prism Web Console Guide* for instructions on configuring rack awareness.
- Select **Manage Categories** to configure cluster categories.
- Select **Disable Efficiency Measurement** to disable efficiency measurement for the cluster.

- Select **Disable Anomaly Detection** to disable anomaly detection for the cluster.
- Select **Run Playbook** to run or create a new playbook for the cluster entity type.
- Select **Download Encryption Keys** to download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just cluster-related alerts (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just cluster-related events (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the clusters. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information below the tab. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

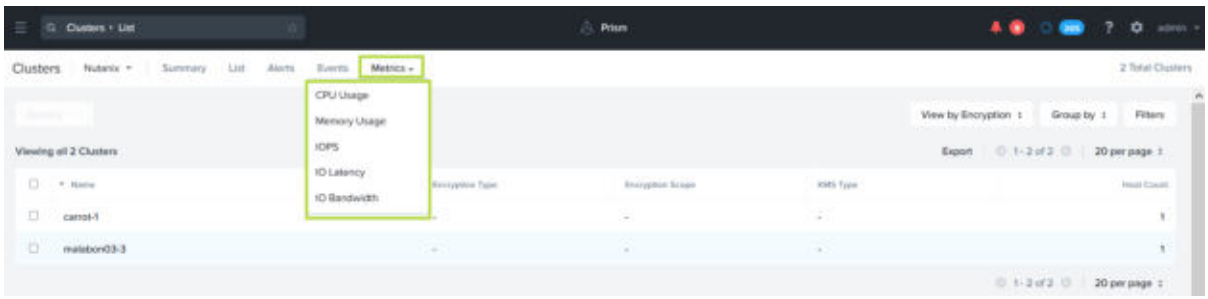

Figure 151: Clusters Metrics Tab

Table 81: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total clusters (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those clusters. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.

Metric	Description
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Cluster Details View

To access the details page for a cluster, go to the clusters dashboard **List** tab (see Clusters Summary View on page 197) and click the cluster name. You can also access the details page by clicking the cluster name wherever that name appears, such as in a dashboard widget or search result.

The cluster name and the following set of tabs appear on the top: **Summary**, **Alerts**, **Events**, **Metrics**, **Storage Usage**, **Hardware**, **Virtual Entities**, and **Capacity**. Click a tab to display that information below the tab. (Click the **Back to Clusters** link to return to the summary view.)

> Note:
>
> For Non-Nutanix clusters, the cluster name and the following set of tabs appear on the top: **Summary**, **Alerts**, **Events**, **Metrics**, and **Capacity**. Click a tab to display that information below the tab. (Click the **Back to Clusters** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the cluster (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.
- A **VM Efficiency** widget that displays the number of VMs, which are considered inefficient broken down by category (overprovisioned, inactive, constrained, and bully). See Behavioral Learning Tools on page 838 for more information about VM efficiency.
- Action buttons (only the applicable ones appear):
  - **Launch Prism Element**: Click this button to launch Prism Element for this cluster in a new tab or window.

    > Note: When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user drop-down list are

REST API Explorer, **About Nutanix**, **Support Portal**, **Help**, **Nutanix Next Community**, and **Sign Out**.

- **Upgrade Software**: Click this button to upgrade the AOS version on this cluster (see Installing or Upgrading on page 30).
- **Rack Configuration**: Click this button to configure the rack awareness feature (see the *Prism Web Console Guide*).
- Select **Manage Categories** to configure cluster categories (see Category Management on page 735).
- Select **Disable Efficiency Measurement** to disable efficiency measurement for the cluster.
- Select **Disable Anomaly Detection** to disable anomaly detection for the cluster.
- Select **Run Playbook** to run or create a new playbook for the cluster entity type (see Creating Playbooks using Triggers on page 865)..
- Select **Download Encryption Keys** to download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.

Figure 152: Cluster Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 82: Cluster Properties Fields

Parameter	Description	Values
Health	Displays the cluster health state (good, warning, or critical).	Critical, Warning, Good
Storage Usage	Displays the amount of storage used in the cluster	xxx [GiB|TiB]
Storage Capacity	Displays the total amount of storage capacity in this cluster.	xxx [GiB|TiB]
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see the "Capacity Tab" section below).	(number of days)
VM Count	Displays the number of VMs in the cluster.	(number of VMs)

Parameter	Description	Values
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of hosts)
Upgrade Status	Displays the status of the last (or current) upgrade attempt.	(status condition)
IP Address	Displays the virtual IP address for the cluster (if defined).	(IP address)
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, or Hyper-V

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget allows quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective pull-down menus.

Example: Viewing Cluster Instances

Click the **Cluster** pull-down menu to view the list of cluster instances. Alternatively, you search a cluster name.

Figure 153: Clusters

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this cluster (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this cluster (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the cluster. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric below the tab. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

> Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 838). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 282).

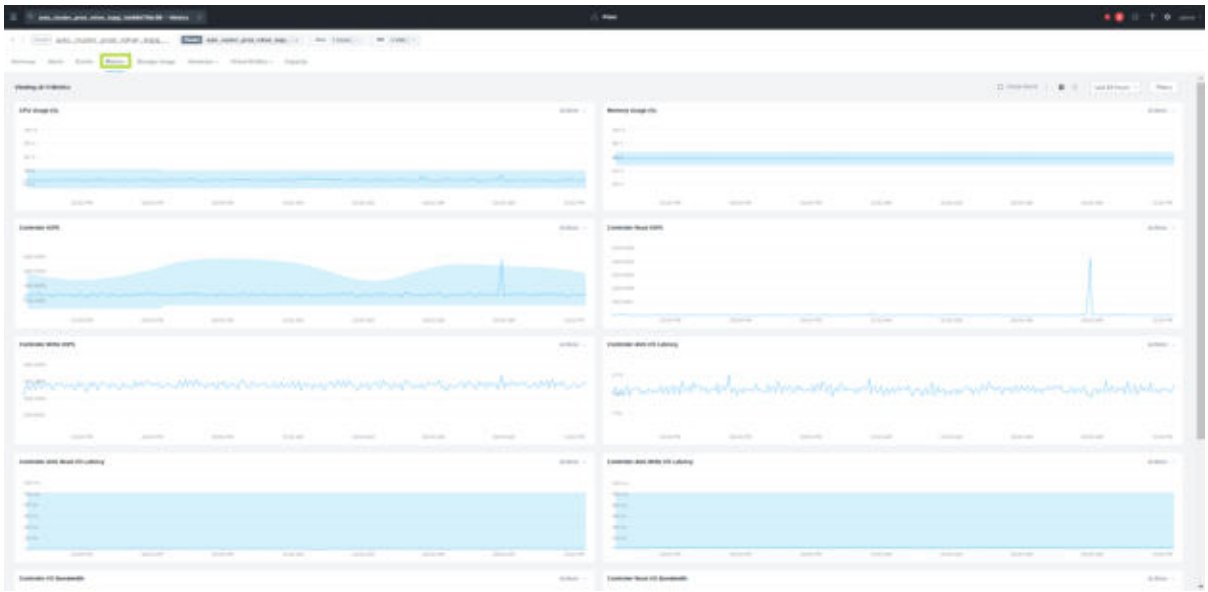

Figure 154: Cluster Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 83: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the cluster (0–100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the cluster (0–100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the cluster.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the cluster.
I/O Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the cluster.

Usage Tab

The **Usage** tab displays the following graphs:

- The **Cluster-wide Usage Summary** graph displays a rolling time interval monitor of total storage usage across the cluster that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of storage space used by each disk tier (SSD and DAS-SATA) across the cluster.

Figure 155: Cluster Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist in this cluster. Clicking a hardware entry such as **Hosts** or a virtual entity such as **Containers** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those in this cluster. See Hardware Entities on page 197 and Compute and Storage Entities on page 86 for more information.

Capacity Tab

The **Capacity** tab displays current and historical usage information and provides resource planning tools. It includes the following sections:

Note:

- The capacity planning feature requires a Prism Pro license. If Prism Pro is disabled, the **Capacity** tab is grayed out and not available.

- Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
- Capacity planning checks the past 90 days of data for monthly seasonality, it also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
- For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 796 for more information.

- A **<cluster_name> Runway** area (upper left) that displays the current cluster runway overall and broken down by CPU, memory, and storage. Runway refers to how long the cluster can continue to run normally, based on the current consumption rate, before the existing resources are used to capacity. The overall runway value is the same as the lowest value for any single resource (storage, CPU, or memory). For example, if the storage runway is 14 days while the CPU and memory runways are 45 days, the overall runway is 14 days (the storage value).

  Note: A plus sign in any runway value, for example the "89+" in the Cluster Capacity Tab figure, indicates that the predicted runway is more than that number of days, but the estimation stopped at that number.

  - Clicking **Storage Runway** displays a storage usage graph and table (to the right).
  - Clicking **CPU Runway** displays a CPU usage graph and table.
  - Clicking **Memory Runway** displays a memory graph and table.

- An **Optimize Resource** button (middle left). Clicking this button opens the **Optimize Resources** window, which identifies areas to improve cluster resource allocation or capacity. Recommendations might include

  - Table of "inefficient" (over-provisioned, inactive, constrained, and bully) VMs based on the VM behavioral learning engine (see Behavioral Learning Tools on page 838). Click the appropriate link to see a list of the VMs in that state.
  - An add node recommendation. The recommendation is based on the current (historical) CPU, memory, and storage usage demand across the cluster.

  Click the **Get Report** link to run the Cluster Efficiency report (see Reports Management on page 883).

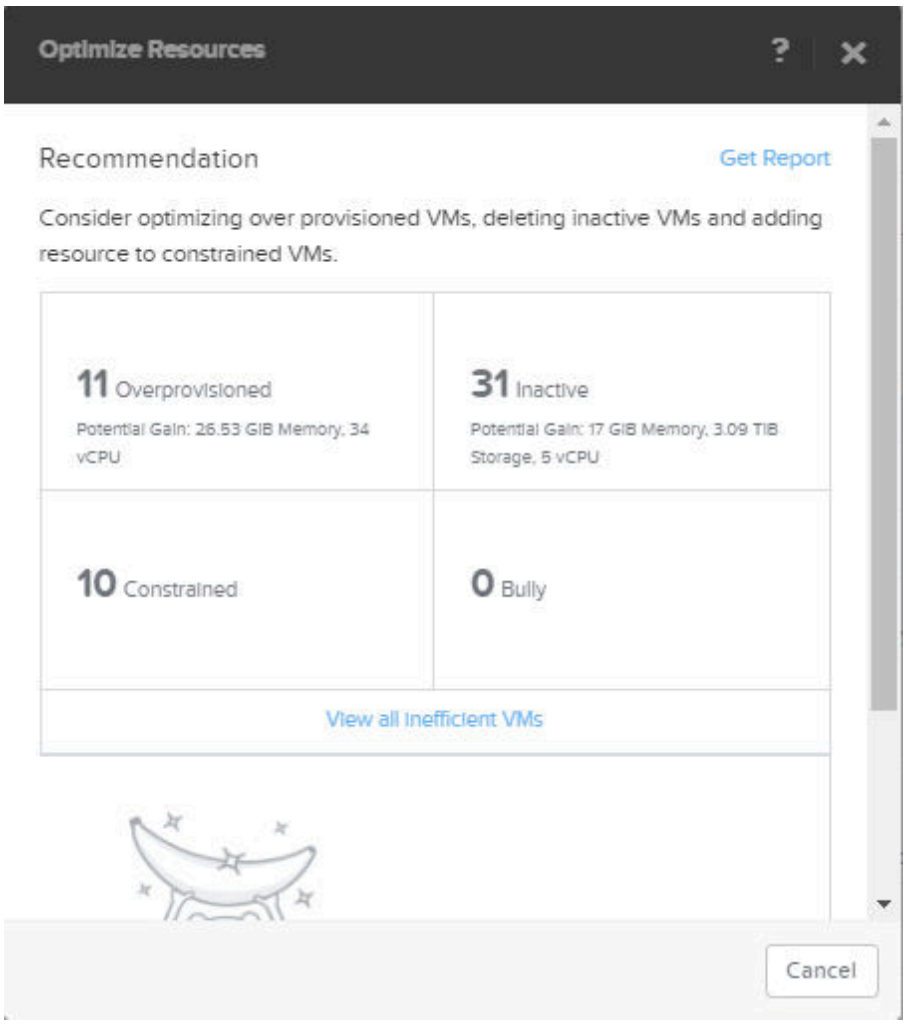

Figure 156: Optimize Resources Window

- A **Get Started** button (lower left). Clicking this button opens the resource planning scenario page. From this page you can create "what if" scenarios to test future workloads against current or hypothetical storage, CPU, and memory resource capacity (see Creating a Scenario on page 820).
- A usage history graph (upper right). The graph displays storage, CPU, or memory usage depending on which is selected in the left column. Placing the cursor anywhere on the horizontal axis displays the value(s) at that time. A (solid red) maximum capacity line

appears near the top of the graph, and an additional (dotted red) effective capacity line appears on the storage graph. A (vertical) "Today" line is displayed to indicate where the current date is located on the graph. When there is an alert or event, a bell icon appears either on the Today line or at the starting point of the alert. Clicking the bell icon (red for critical alert, orange for warning alert, or gray for event) displays more information about that alert or event.

- A table (lower right) that displays usage information. The information in the table varies depending on which runway (storage, CPU, or memory) is selected.

Figure 157: Cluster Capacity Tab (Nutanix Cluster)

Figure 158: Cluster Capacity Tab (Non-Nutanix Cluster)

Capacity Tab: Storage View

When **Storage Runway** is selected, the graph and table display storage usage information. The maximum and effective storage capacity values (in TiBs) for the cluster appear above the graph. Effective capacity of a cluster is computed based on the configured reservations (from the capacity configurations settings) and the replication factor for all the containers in the cluster.

The storage view includes two tabs:

- Click the **By Usage** button to see cluster-wide storage information (see previous figure).
- Click the **By Storage Container** button to see storage information per storage container. You can select all storage containers from the pull-down list (upper left of graph) or a single

storage container from the pull-down list or by clicking the storage container name in the table.

Figure 159: Storage Container View (Nutanix Cluster)

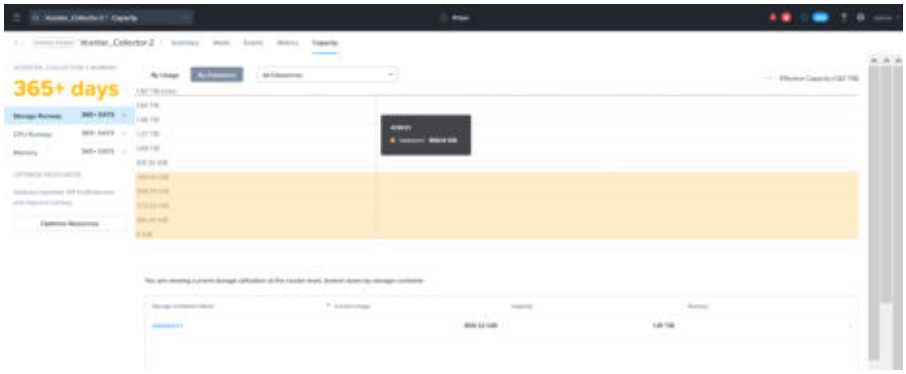

Figure 160: Storage Container View (Non-Nutanix Cluster)

The following table describes the fields in the storage table.

Table 84: Capacity Tab: Storage Fields

Parameter	Description	Values
"By Usage" tab (cluster-wide)		
Name	Displays a descriptive name for a type of storage. There are four types: • Live Usage: The amount of live storage. • Snapshot Usage: The amount of storage used for snapshots. • System Usage: The amount of storage for everything else (total usage - (reserved + snapshot + live)). System usage includes garbage, parity, and other miscellaneous uses. • Reserved Usage: The amount of storage reserved for thick provisioned VMs (ESXi).	Live Usage, Reserved Usage, Snapshot Usage, System Usage

Parameter	Description	Values
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB|TB]
"By Storage Container" tab: All Storage Containers		
Storage Container Name	Displays the name of the storage container.	(storage container name)
Current Usage	Displays the amount of storage space used by the storage container currently.	xxx [GB|TB]
Capacity	Displays the total amount of storage capacity available to the storage container.	xxx [TB]
"By Storage Container" tab: <storage container name>		
Storage Container Name	Displays one of the following: • If "All Storage Containers" is selected, the storage container names are displayed. Click a name to see the details for that storage container. • If a single storage container is selected, it displays the four storage types described for the Name field in the "By Usage" tab (Live Usage, Reserved Usage, Snapshot Usage, and System Usage).	(storage container name) OR Live Usage, Reserved Usage, Snapshot Usage, System Usage
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB|TB]

Capacity Tab: CPU View

When **CPU Runway** is selected, the graph and table display CPU usage information. The maximum and effective CPU capacity values (in GHz) for the cluster appear above the graph. The CPU view includes two tabs:

- Click the **Overall** button to see cluster-wide CPU information.
- Click the **By Host** button to see CPU information per host. You can select all hosts from the pull-down list (upper left of graph) or a single host from the pull-down list or by clicking the host name in the table.

The following table describes the fields in the CPU table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the CPU table.

Table 85: Capacity Tab: CPU Fields

Parameter	Description	Values
"Overall" tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)

Parameter	Description	Values
CPU Usage (%)	Displays the percent of CPU capacity used by the VM currently.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)
"By Host" tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average CPU Usage over last 1 hour	Displays the average host CPU usage during the last hour.	xxx [GHz]
Current CPU Capacity	Displays the CPU capacity of the host.	xxx [GHz]
Runway	Displays the current runway (estimated time interval that current CPU capacity is sufficient for the workload).	(number of days)
"By Host" tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)

Capacity Tab: Memory View

When **Memory Runway** is selected, the graph and table display memory usage information. The maximum and effective memory capacity values (in GiBs) for the cluster appear above the graph. The memory view includes two tabs:

- Click the **Overall** button to see cluster-wide memory information.
- Click the **By Host** button to see memory information per host. You can select all hosts from the pull-down list (upper left of graph) or a single host from the pull-down list or by clicking the host name in the table.

The following table describes the fields in the memory table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the memory table.

Table 86: Capacity Tab: Memory Fields

Parameter	Description	Values
"Overall" tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]

Parameter	Description	Values
Memory Usage (%)	Displays the percentage of memory capacity used by the VM.	(0–100%)
"By Host" tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average Memory Usage over last 1 hour	Displays the average host memory usage during the last hour.	xxx [GiB]
Runway	Displays the current runway (estimated time interval that current memory capacity is sufficient for the workload).	(number of days)
Current Memory Capacity	Displays the memory capacity of the host.	xxx [GiB]
"By Host" tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM currently.	(0–100%)

Metrics Tab

The **Metrics** tab allows you to view performance metrics for the cluster, both for actual past performance and for predicted future performance. The **Metrics** tab displays a section on the left with a list of metrics.

> Note: The predicted metrics and identified anomalies are based on sophisticated machine-learning capabilities. See Behavioral Learning Tools on page 838 for information about these capabilities and how they are used.

- Clicking a metric displays a graph on the right. (Some metrics have multiple graphs.) The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph. Placing the cursor anywhere on the horizontal axis displays the value at that time. To set the time interval (last 24 hours, last week, last 21 days), select the duration from the pull-down list on the right.

  > Note: The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph may not appear if less than 21 days of data is available.

- To display predicted performance, check the **Prediction** box. This displays projected performance for the next week.
- To create an alert for this cluster based on either behavioral anomalies or status thresholds, click the **Set Alerts** link above the graph (see Creating Custom Alert Policies on page 282).

The following table describes the available metrics.

> Note: Some of these metrics are not available on all hypervisors.

Table 87: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the VM (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the VM (0 - 100%).
I/O Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.

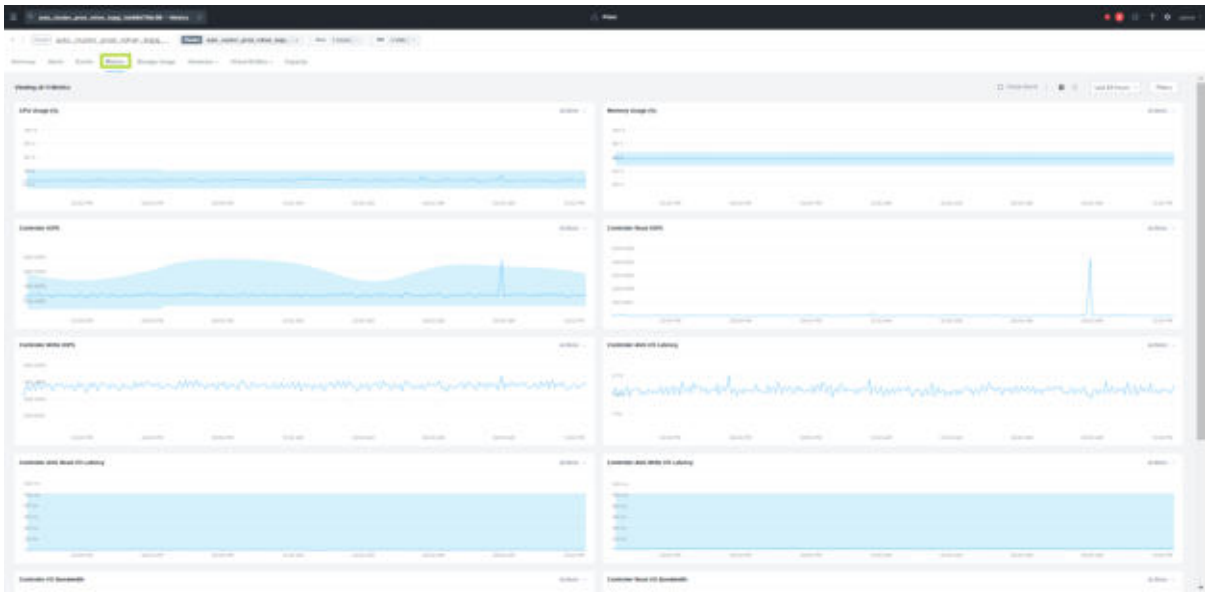

Figure 161: Metrics Tab: CPU Usage metric

Hosts Summary View

To access the hosts dashboard, select **Hardware** > **Hosts** from the entities menu (see Entities Menu on page 13). The hosts dashboard allows you to view summary information about hosts across the registered clusters and access detailed information about each host. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area to the right for the selected tab.

> Note: This section describes the information and options that appear in the Hosts dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Suggested**: Displays a list of the hosts with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **Memory Usage**, **IO Latency**, and **IOPS**. Click the **View all XX Hosts** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of host-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

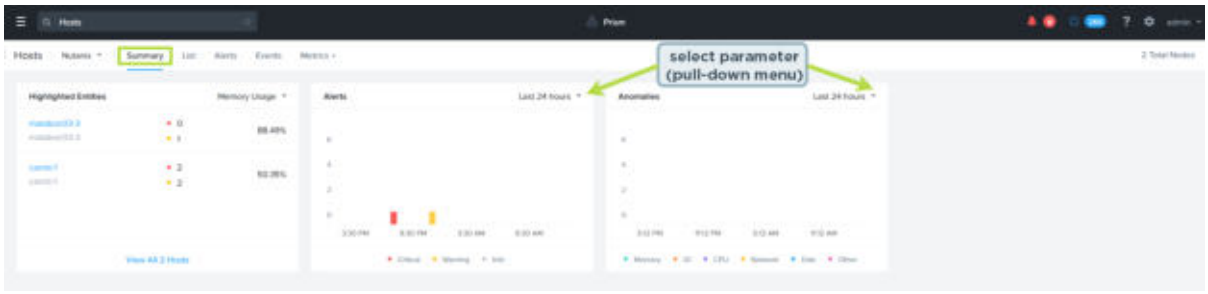

Figure 162: Hosts Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the hosts across the registered clusters. The following table describes the fields that appear in the hosts list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

Figure 163: Hosts List Tab

Table 88: Hosts List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the name of the host. Clicking on the name displays the details page for that host (see Host Details View on page 222).	(host name)

Parameter	Description	Values
Host IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Hypervisor	Displays the hypervisor type running on the host.	AHV, ESX, or Hyper-V
Memory Capacity	Displays the memory capacity of the host.	xxx [MB|GB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
"Performance" Focus Fields		
Name	Displays the name of the host.	(host name)
CPU Usage	Displays the percentage of CPU capacity currently being used by this host.	0 - 100%
Memory Usage	Displays the percentage of memory capacity currently being used by this host.	0 -100%
IOPS	Displays I/O operations per second (IOPS) for this host.	[0 - unlimited]
Disk IO Bandwidth	Displays I/O bandwidth used per second for this host.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this host.	xxx [ms]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)

You can filter the hosts list based on a variety of parameter values. The following table describes the filter options available when you open the Hosts view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 89: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the host name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of hosts that satisfy the name condition/ string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of hosts that reside in the clusters which satisfy the name string.	(cluster name string)

Parameter	Description	Values
Health	Filters on the host health state (good, warning, or critical). Select one or more states to return a list of hosts in that state(s). The number of hosts currently in each state is displayed on the right of the line.	Critical, Warning, Good
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a pull-down list appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of hosts tagged to each selected category is displayed on the right of the line.	(category name)
Hypervisor	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, HyperV
Memory Capacity	Filters on the host memory capacity. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It will return a list of hosts with memory capacity in that range.	([xx] to [yy] GiB range)
CPU Usage	Filters on the amount of CPU being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of hosts utilizing CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
GPUs	Filters for GPU configuration information such as model name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. As you type a pull-down list appears to help you select the correct configuration information. It will return a list of hosts that satisfy the GPU condition/string.	(configuration info)
IOPS	Filters on the IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of hosts with IOPS in that range.	([xx] to [yy] range)

Parameter	Description	Values
IO Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of hosts with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of hosts with average I/O latency in that range.	([xx] to [yy] range)

You can group the hosts list in the following ways:

- The **Group** pull-down menu allows you to group the host entries by cluster, hypervisor type, or health state. (You can only choose one.)
- The **View by** pull-down menu allows you to group the host entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

The **Actions** menu appears when one or more hosts are selected. It includes a **Manage Categories** option (see Assigning a Category on page 737).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just host-related alerts across the registered clusters (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just host-related events across the registered clusters (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

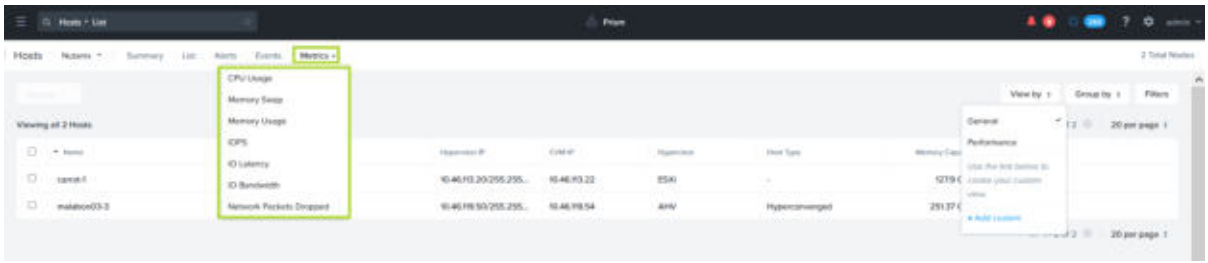

Figure 164: Hosts Metrics Tab

Table 90: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total hosts (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those hosts. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Swap	Displays memory swap-out and swap-in rate tables.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Host Details View

To access the details page for a host, go to the hosts dashboard **List** tab (see Hosts Summary View on page 217) and click the host name. You can also access the details page by clicking the host name wherever that name appear, such as in a dashboard widget or search result.

The host name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, **Usage**, **Hardware**, and **Virtual Entities**. Click a tab to display that information on the right. (Click the **Back to Hosts** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the host (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
- A **Manage Categories** button to manage (set) categories for the host (see Category Management on page 735).

Figure 165: Host Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 91: Host Properties Fields

Parameter	Description	Values
Memory Capacity	Displays the total memory capacity for this host.	xxx [MB|GB]
Disk Capacity	Displays the total amount of disk capacity on this host.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the host resides. Clicking the name displays the details page for that cluster (see Cluster Details View on page 205).	(cluster name)
Host IP	Displays the host IP address.	(IP address)
Hypervisor	Displays the hypervisor name.	(hypervisor name)
VM Count	Displays the number of VMs running on this host.	(number)
Block Model	Displays the block model number.	(model series number)
Serial Number	Displays the block serial number.	(block serial number)
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Host Type		
IPMI Address	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console.	(IP address)
Node Serial	Displays the node serial number. The node serial is a unique number passed through from the manufacturer. (The form can vary because it is determined by each manufacturer.)	(manufacturer serial number)

Parameter	Description	Values
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on the metadata disk.	[0 - 100%]
Oplog Disk Size	Displays the current size of the operations log. (The Oplog maintains a record of write requests in the cluster.) A portion of the metadata disk is reserved for the Oplog, and you can change the size through the nCLI.	xxx [GB]
Monitor Enabled	Displays whether the host is high availability (HA) protected. A **Yes** value means HA is active for this host. A **No** value means VMs on this host are not protected (will not be restarted on another host) if the host fails. Normally, this value should always be **Yes**. A **No** value is likely a sign of a problem situation that should be investigated.	[Yes|No]
Disks	Displays the number of disks in each storage tier in the host. Tier types vary depending on the Nutanix model type.	DAS-SATA: (number), SSD-SATA: (number), SSD-PCIe: (number)
GPUs	Displays the number and type of GPUs in the host. For example, if the host contains four Tesla M10 GPUs, this field displays "Tesla M10 (4)".	(GPU type and number)
Datastore(s)	Displays the names of any datastores.	(names)

Figure 166: Host Summary Tab

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget provides quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective pull-down menus.

Example: Viewing Hosts on a Cluster

Click the **Host** pull-down menu to view the list of host instances on the selected cluster. Alternatively, you can search the host instance name belonging to the target cluster.

Figure 167: Hosts

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this host (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this host (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the host. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 838). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).

- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 282).

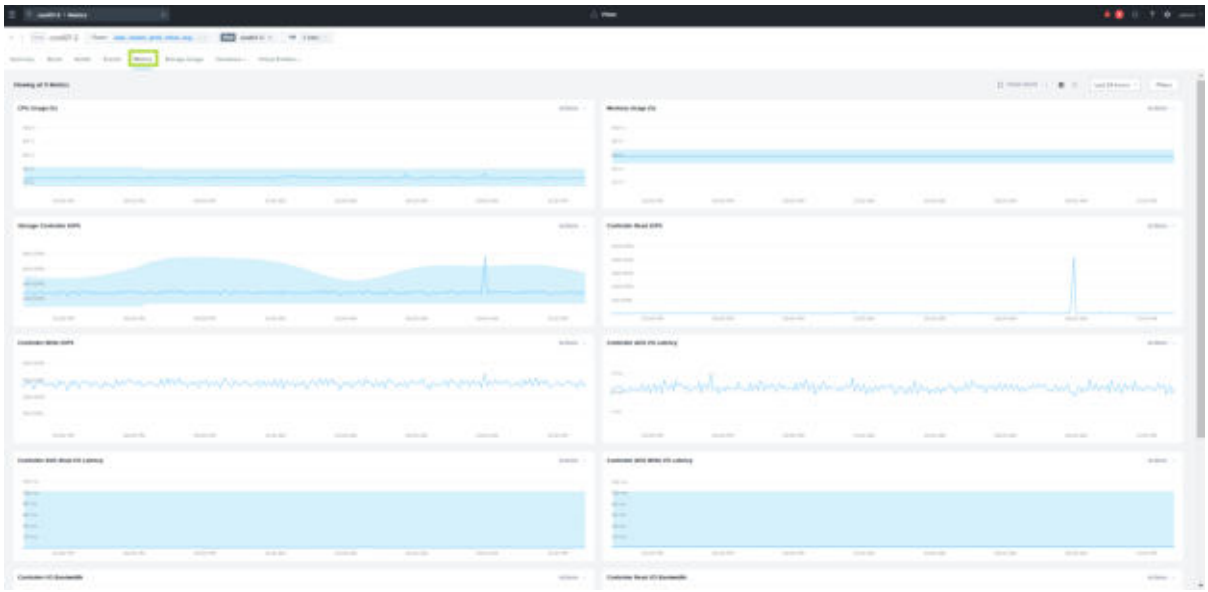

Figure 168: Host Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 92: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the host (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the host (0 - 100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the host.
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the host.
IO Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the host.

Usage Tab

The **Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of host storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more

in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of host storage space used by each disk tier (SSD and DAS-SATA).

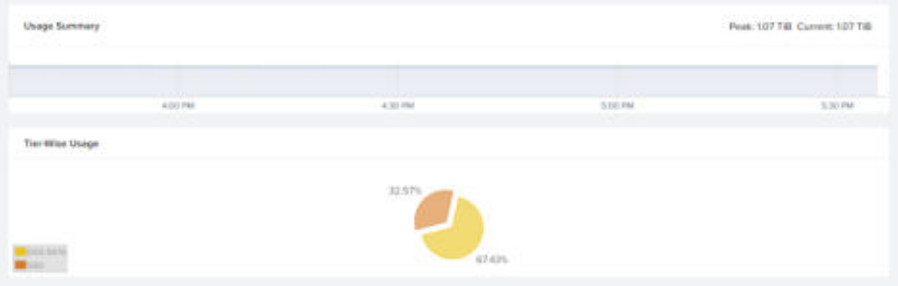

Figure 169: Host Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist on this host. Clicking a hardware entry such as **Disks** or a virtual entity such as **VMs** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those on this host. See Hardware Entities on page 197 and Compute and Storage Entities on page 86 for more information.

Disks Summary View

To access the disks dashboard, select **Hardware** > **Disks** from the entities menu (see Entities Menu on page 13). The disks dashboard allows you to view summary information about disks across the registered clusters and access detailed information about each disk.

> Note: This section describes the information and options that appear in the Disks dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

Clicking the **Summary** tab displays the following three widgets:

- **Suggested**: Displays a list of the disks with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Bandwidth**, **IOPS**, **IO Latency**, and **Disk Usage**. Click the **View all XX Disks** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of disk-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

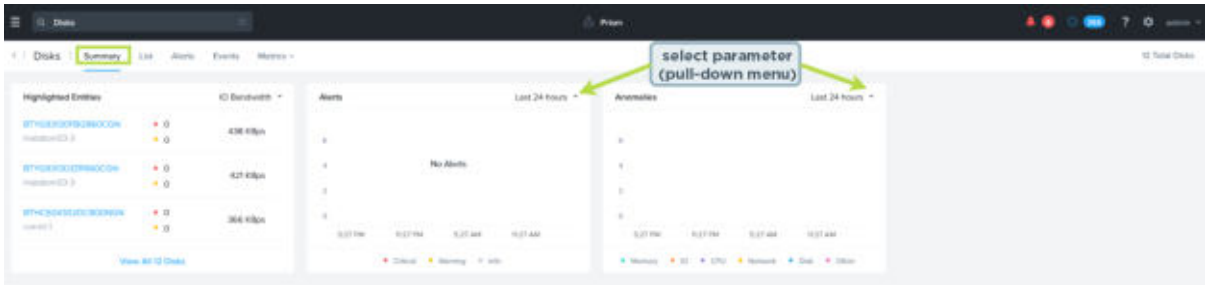

Figure 170: Disks Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the disks across the registered clusters. The following table describes the fields that appear in the disks list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

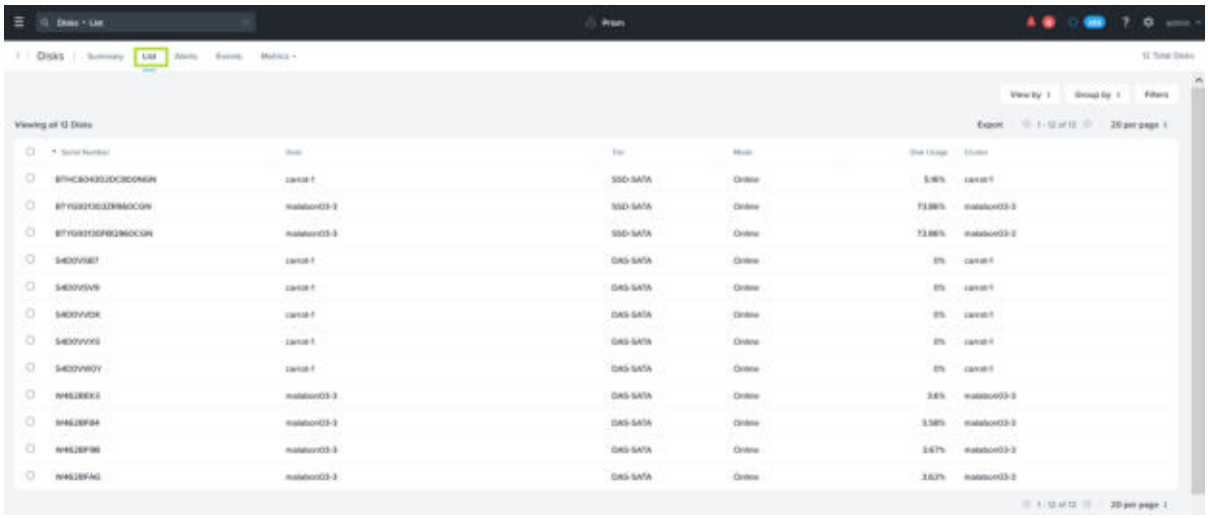

Figure 171: Disks List Tab

Table 93: Disks List Fields

Parameter	Description	Values
"General" Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Host	Displays the name of the host in which this disk resides.	(host name)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe | SSD-SATA | DAS-SATA]

Parameter	Description	Values
Mode	Displays the operating state of the disk.	online, offline
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB|TB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
"Performance" Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB|TB]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB|TB]
IOPS	Displays I/O operations per second (IOPS) for this disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for this disk.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this disk.	xxx [ms]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)

You can filter the disks list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 94: Filter Pane Fields

Parameter	Description	Values
Serial Number	Filters on the disk serial number. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of disks that satisfy the serial number condition/string.	(serial number string)
Host	Filters on the host name. Enter a string in the field. It will return a list of disks in the hosts that satisfy the host name condition/string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of disks in the clusters that satisfy the cluster name condition/string.	(cluster name string)

Parameter	Description	Values
Mode	Filters on whether the disk is online or offine. Check the box for one or both of these modes. The number of disks currently in each state is displayed on the right of the line.	Online, Offline
Tier	Filters on whether the disk is in the solid state (SSD-SATA) or hard disk (DAS-SATA) tier. Check the box for one or both of these modes. The number of disks currently in each tier is displayed on the right of the line.	DAS-SATA, SSD-SATA
Health	Filters on the disk health state (good, warning, or critical). Select one or more states to return a list of disks in that state(s). The number of disks currently in each state is displayed on the right of the line.	Critical, Warning, Good
Disk Usage	Filters on the used capacity. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of disks with used capacity in that range (0-100%).	([xx] to [yy]% range)
Disk Capacity	Filters on the total capacity. Enter an amount range in the "**from <low> to <high> GiB**" field. It will return a list of disks with total capacity in that range.	([xx] to [yy] GiB range)
IOPS	Filters on the IOPS. Enter a range in the "**from <low> to <high> iops**" field. It will return a list of disks with IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the I/O bandwidth used. Enter a range in the "**from <low> to <high> bps**" field. It will return a list of disks with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Enter a range in the "**from <low> to <high> ms**" field. It will return a list of disks with average I/O latency in that range.	([xx] to [yy] range)

You can group the disks list in the following ways:

- The **Color** pull-down menu allows you to color code the disk entries by tier type, mode, or health state. (You can only choose one.) A legend appears at the bottom to indicate what each color means in that grouping.
- The **Group** pull-down menu allows you to group the disk entries by host, tier type, mode, cluster, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** pull-down menu allows you to group the disk entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just disk-related alerts across the registered clusters (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just disk-related events across the registered clusters (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

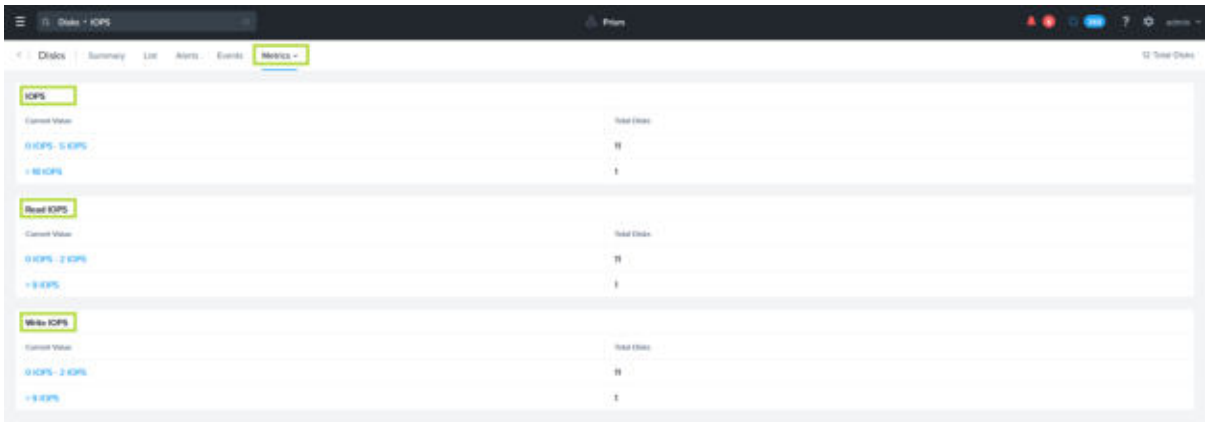

Figure 172: Disks Metrics Tab

Table 95: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS tables listing current values and total disks (number). The current values are split into intervals (for example, less than 20, 20-40, 40-60, more than 60). Clicking on an interval displays the Summary tab filtered to just those disk. Note: The same format also applies to the other metrics in this table.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Disk Details View

To access the details page for a disk, go to the disks dashboard **List** tab (see Disks Summary View on page 227) and click the disk name. You can also access the details page by clicking the disk name wherever that name appear, such as in a dashboard widget or search result.

The disk name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, and **Usage**. Click a tab to display that information on the right. (Click the **Back to Disks** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the disk (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

Figure 173: Disk Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 96: Disk Properties Fields

Parameter	Description	Values
Disk Usage	Displays the amount of used space on the drive.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)
Host	Displays the name of the host in which the disk resides.	(host name)
Host IP	Displays the IP address of the host.	(IP address)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe | SSD-SATA | DAS-SATA]
Mode	Displays whether the disk is currently online or offline.	[online|offline]

Parameter	Description	Values
Disk Capacity	Displays the total physical space on the drive.	xxx [GB|TB]
Status	Displays the operating status of the disk. Possible states include the following: • **Normal**. Disk is operating normally. • **Data migration initiated**. Data is being migrated to other disks. • **Marked for removal, data migration is in progress**. Data is being migrated in preparation to remove disk. • **Detachable**. Disk is not being used and can be removed.	Normal; Data migration initiated; Marked for removal, data migration is in progress; Detachable
Self Encryption Drive	Displays whether this is a self-encrypted drive.	Not Present, Present

Figure 174: Disk Summary Tab

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this disk (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this disk (see Events Summary View on page 272).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the disk. Click the **Metrics** tab and then the desired metric name (**IOPS**, **IO latency**, and **IO Bandwidth**) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph.

> Note: The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).

- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 282).

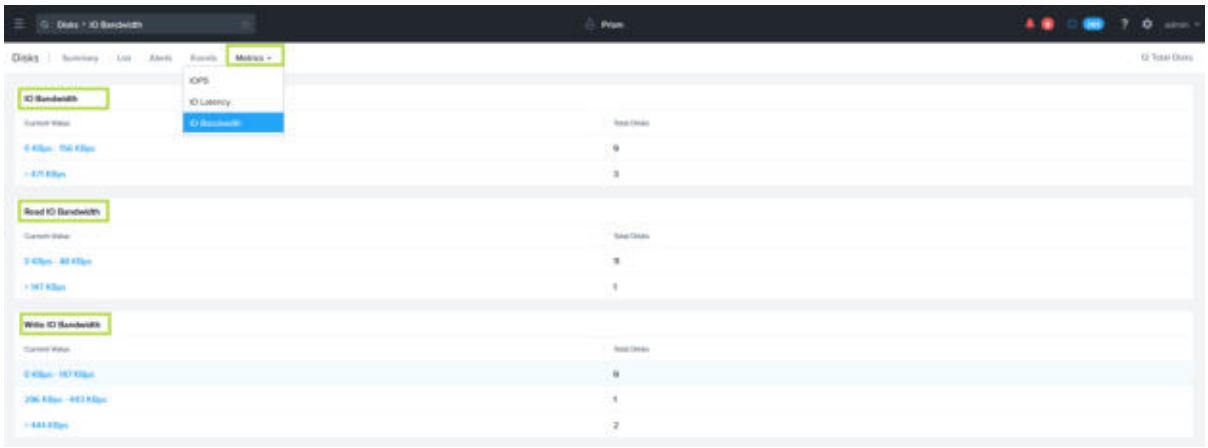

Figure 175: Disk Metrics Tab

Usage Tab

The **Usage** tab displays the following graph:

- The **Usage Summary** graph displays a rolling time interval monitor of disk storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

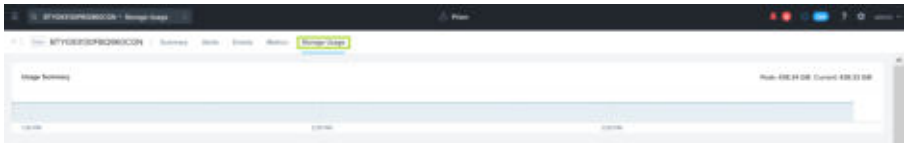

Figure 176: Disk Usage Tab

GPUs Summary View

To access the GPUs dashboard, select **Hardware** > **GPUs** from the entities menu (see Entities Menu on page 13). The GPUs dashboard allows you to view summary information about GPUs across the registered clusters and access detailed information about each GPU.

> Note: This section describes the information and options that appear in the GPUs dahsboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.

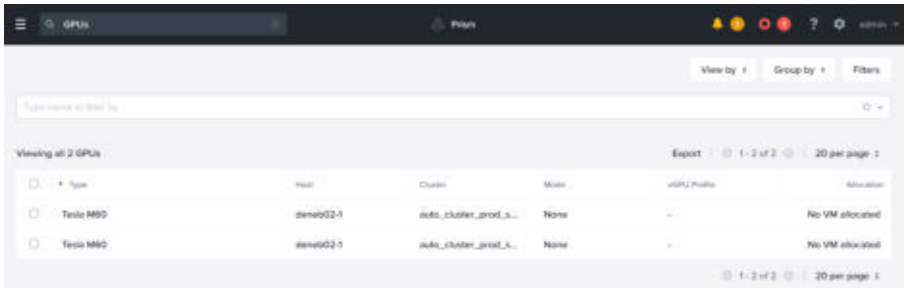

Figure 177: GPUs Dashboard

The following table describes the fields that appear in the GPUs list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 97: GPU List Fields

Parameter	Description	Values
Type	Displays the GPU model type.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster	Displays the name of the cluster in which the GPU resides.	(cluster name)
Mode	Displays the mode in which the GPU is operating.	None, vGPU, passthrough
Allocation	Indicates the number of VMs allocated to the GPU.	"No VM allocated", "*x* of *y* VMs allocated"

You can filter the GPUs list based on a several parameter values. The following table describes the filter options available when you open the GPUs view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 98: Filter Pane Fields

Parameter	Description	Values
Type	Filters on the GPU model type. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of GPUs that satisfy the type condition/string.	Tesla M10, Tesla M60, Tesla M60 compute
Host	Filters on the host name. Enter a string in the field. It will return a list of GPUs in the selected hosts.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of GPUs in the selected clusters.	(cluster name string)

Parameter	Description	Values
Mode	Filters on the GPU operation mode. Check the box for one or more of these modes. The number of GPUs currently in each mode is displayed on the right of the line.	None, vGPU, Passthrough

You can group the GPUs list in the following ways:

- The **Group** pull-down menu allows you to group the GPU entries by cluster, host, or mode. (You can only choose one.)
- The **Sort** pull-down menu allows you to group the GPU entries by type, cluster, mode, or allocation. (You can only choose one parameter.)

GPU Details View

To access the details page for a GPU, click on the desired GPU type entry in the list (see GPUs Summary View on page 234). You can also access the details page by clicking the GPU name wherever that name appear, such as in a dashboard widget or search result. The GPU details page includes **Summary** and **VMs** tabs.

Summary Tab

Clicking the **Summary** tab, which appears by default, displays the following:

- A section on the left that displays summary information about the GPU (see following table).
- A section of the right that displays GPU performance metrics. The graphs are rolling time interval performance monitors that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. This section includes the following graphs:
  - **GPU Usage**: Displays the percentage of GPU capacity being used.
  - **GPU Framebuffer Usage**: Displays the percentage of GPU framebuffer (RAM) capacity being used.
- Action button on the upper right (which appears regardless of which tab is selected):
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

The following table describes the GPU summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 99: GPU Summary Fields

Parameter	Description	Values
GPU Type	Displays the GPU type for this entry.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster Name	Displays the name of the cluster in which the GPU resides.	(cluster name)

Parameter	Description	Values
Host	Displays the name of the host in which the GPU resides.	(host name)
Mode	Displays the GPU operational mode. If it is vGPU, an additional field appears that displays the vGPU profile used.	None, vGPU, Passthrough
Allocation	Displays the number of VMs allocated to this GPU. This field does not appear if no VMs are allocated.	"*x* of *y* VMs allocated"
Framebuffer (RAM) Per GPU	Displays the framebuffer (RAM) size per GPU.	xx GiB
ID	Displays the GPU ID number.	(ID number)

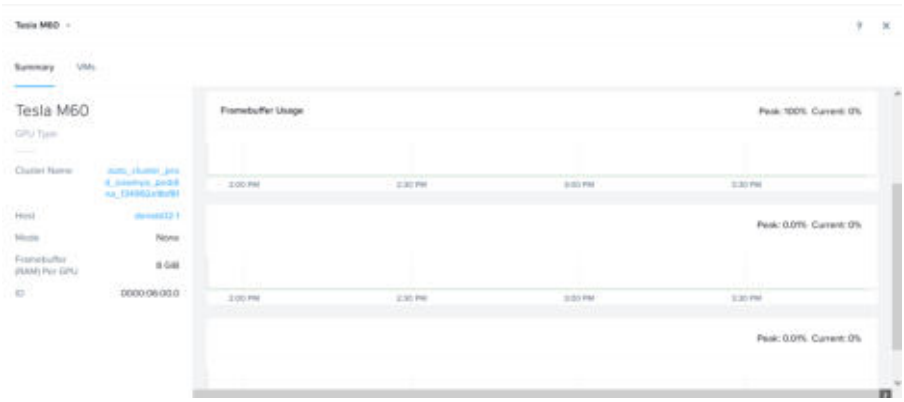

Figure 178: GPU Summary Tab

VMs Tab

Clicking the **VMs** tab displays a table of VMs allocated (attached) to the GPU. The table includes the following fields:

- **Name**: Displays the VM name. Click the name to display the details page for that VM (see VM Details View on page 96).
- **GPU Usage**: Displays the percentage of GPU capacity used by this VM.
- **GPU Framebuffer Usage**: Displays the percentage of GPU framebuffer (RAM) capacity used by this VM.

Activity Entities

You can access dashboards for the following activity monitors from the **Activity** category of the entities menu (see Entities Menu on page 13):

- Alerts (see Alerts Summary View on page 261)
- Events (see Events Summary View on page 272)
- Audits (see Audits Summary View on page 238)
- Tasks (see Tasks View on page 241)

Audits Summary View

To access the audits dashboard, select **Activity** > **Audits** from the entities menu (see Entities Menu on page 13). The audits dashboard allows you to view a list of actions performed across the registered clusters.

Note:

- This section describes the information and options that appear in the audits dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- Audit information appears only for those registered clusters running AOS 5.10 or later.
- Audit logs with default values are generated when updates to VMs are initiated, either by Prism Central **Self Service** users or by using Nutanix v3 API calls for the first time.
- The retention period for audit entries is four weeks by default.
- See Prism Central Logs on page 299 for all Prism Central related events. Additionally, you can refer to the Audit Log Events on page 302 for the complete list of captured audit events.

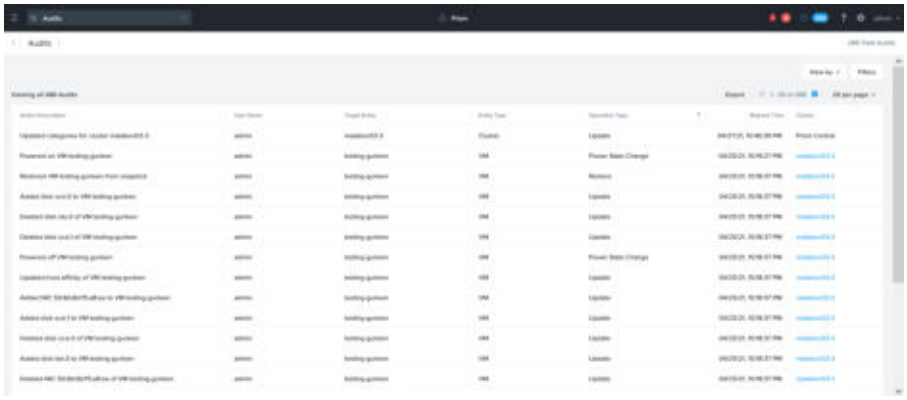

Figure 179: Audits Dashboard

The following table describes the fields that appear in the audits list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 100: Recovery Plans Fields

Parameter	Description	Values
Action Description	Describes the action taken such as "deleted VM *vm-name*" or "added disk *disk-name*"	(action description)
User Name	Displays the name of the user who requested the action.	(user name)
Target Entity	Displays the entity name. Click the name to go to the details page for that entity.	(entity name)

Parameter	Description	Values
Entity Type	Displays the entity type such as VM or host.	(entity type)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)

To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. The following table describes the filter options available. You can apply multiple filters.

Table 101: Filter Pane Fields

Parameter	Description	Values
User Name	Enter a name string in the field to filter for users who requested an action.	(user name)
Entity Type	Check the boxes of one or more entities to filter for actions on those entity types.	VM, Storage Container, Catalog Item, Image, Cluster, Host, Disk, GPU, Security Policy, NGT Policy, Project, Role, User, Category, Availability Zone, Protection Policy, Recovery Plan, Recoverable Entity, Report
Operation Type	Check the boxes of one or more operations to filter on those operations.	Create, Update, Delete, Power State Change
Request Time	Check an interval box to filter for actions that were requested during that time period. For the custom interval option (**from** *xxx* **to** *xxx*), click in each field and select a date from the pop-up calendar.	Last 1 hour, Last 24 hours, Last week, From *xxx* to *xxx*
Cluster	Enter a cluster name in the field to filter for actions in the cluster.	(cluster name)
User IP	Enter a user IP address and then click the Add button to filter for actions requested by that user. You can add multiple user IP addresses.	(IP address)

Audit Details View

To access the details page for an action, go to the audits dashboard (see Audits Summary View on page 238) and click the action description. The details page includes the following:

- The action description (upper left). You can switch from one action details page to another by selecting from the pull-down list.
- A section on the left that displays summary information about the action (see following table).
- A section on the right that displays a table of information specific to that action. The table shows the attributes that were changed during the action and the current value that is set for the attribute. The attributes vary depending on the specifics of the action (see following examples).

The following table describes the action summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 102: Action Summary Fields

Parameter	Description	Values
Action Description	Displays the action description.	(description)
User Name	Displays the name of the user who requested the action.	(user name)
Target Entity	Displays the name of the entity that was the action target. Click the entity name to display the details for that entity.	(entity name)
Affected Entities	Displays the names of the entities that were affected by the action. Click an entity name to displays the details page for that entity.	(one or more entity names)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
User IP	Displays the IP address of the user.	(IP address)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)
Status	Displays the status of the action.	Succeeded, Failed

In the first example the action is adding more categories for the cluster, so the details simply include the names of the attached categories.

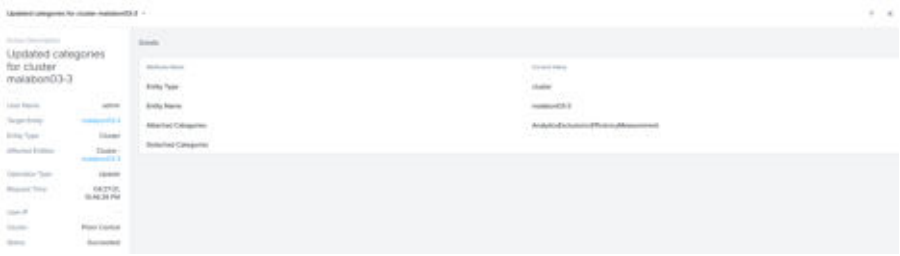

Figure 180: Audit Details Page (example 1)

In the second example the action is adding a disk, so the details include the VM name, disk address, size, and source disk UUID.

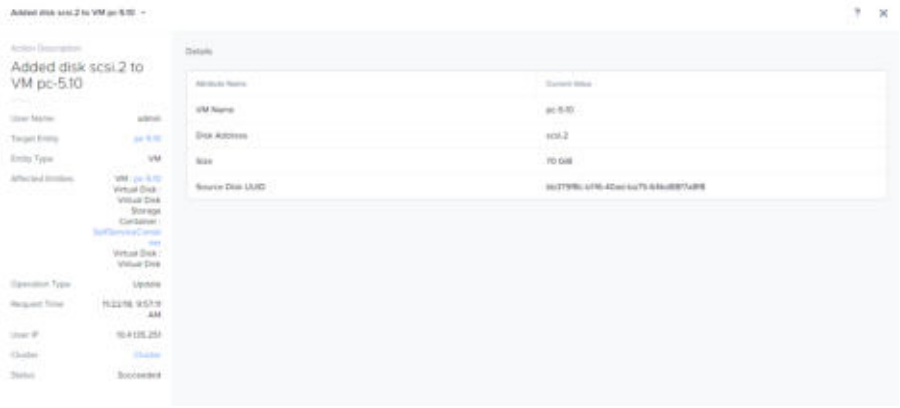

Figure 181: Audit Details Page (example 2)

In the third example the action is creating a VM, so the details include a variety of information about that VM.

Figure 182: Audit Details Page (example 3)

Tasks View

An icon appears in the main menu when there are current tasks (running or completed within the last 48 hours). The icon appears blue when a task runs normally, yellow when it generates a warning, or red when it fails. Clicking the icon displays a drop-down list of current tasks. You can see the Tasks dashboard, which displays information about all tasks across the registered clusters, by doing one of the following:

- Clicking the **View All Tasks** link at the bottom of the current tasks drop-down list.
- Clicking the **View All Tasks** link in the Tasks widget in the main dashboard (see Main Dashboard on page 64).
- Selecting **Activity** > **Tasks** in the entities menu (see Entities Menu on page 13).

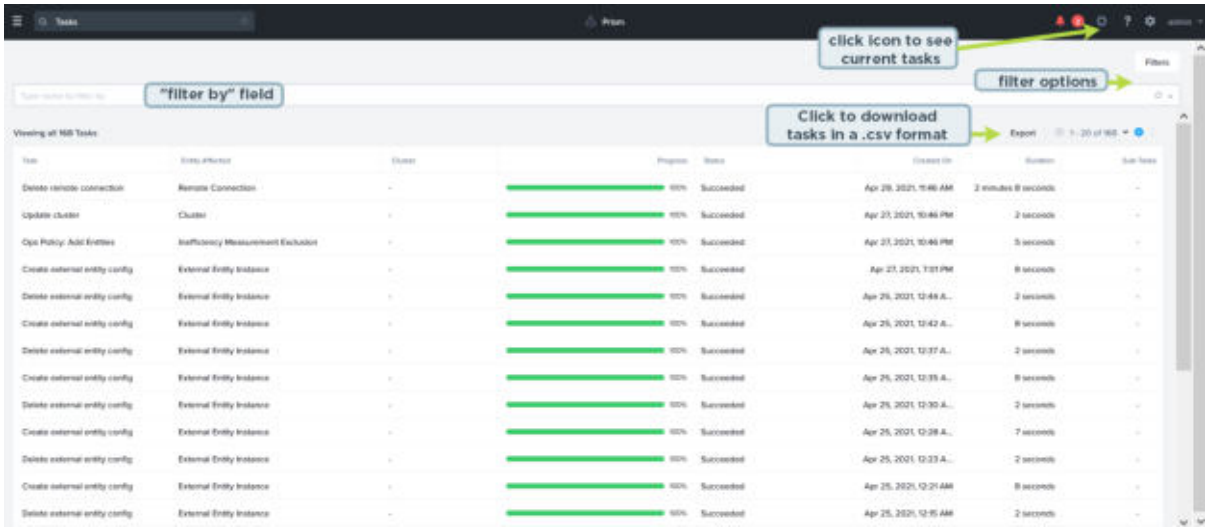

Figure 183: Tasks Screen

You can do the following in the Tasks dashboard:

- Filter the list by entering a name in the "filter by" field.
- Filter the list by clicking the **Filters** button. This displays a filter pane on the right of the screen. Click on one or more of the states to filter the list to just tasks in those states. The states are Aborted, Canceled, Failed, Queued, Running, Paused, Skipped, Succeeded, and

Scheduled. The number of tasks in each state currently appear to the right of the state name. Click the X icon (upper right) to close the filter pane.

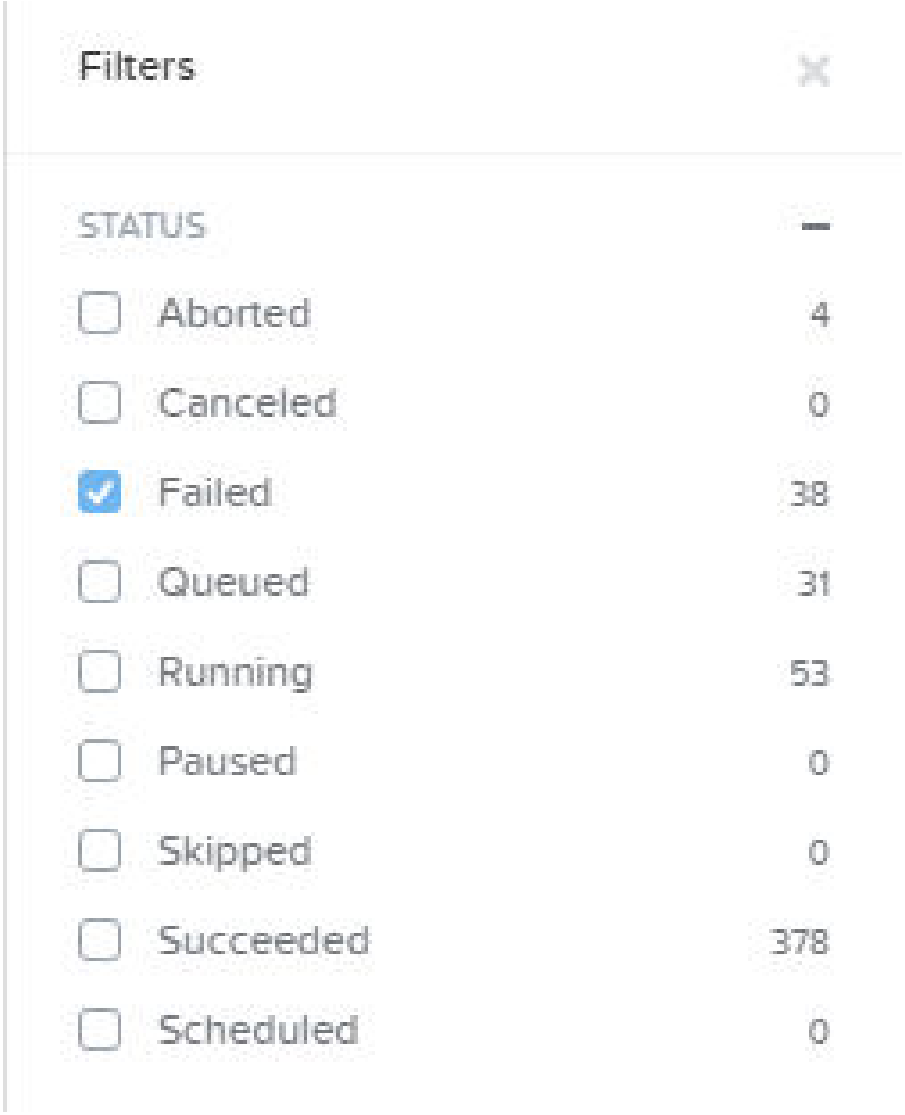

Figure 184: Filter Pane

- Download the task list in CSV format by clicking the download icon (on the right above the table).

The following table describes the fields in the tasks list. Each task appears in the list for a minimum of one hour after completion, but how long that task remains in the list depends on several factors. In general, the maximum duration is two weeks. However, tasks are rotated off the list as new tasks arrive, so a task might disappear from the list much sooner when activity is high. In some cases a task appears for longer than two weeks because the last task for each component is retained in the listing.

Table 103: Tasks List Fields

Parameter	Description	Values
Task	Specifies which type of operation the task is performing.	(name of operation such as "upgrade")
Entity Affected	Lists the entity name or describes the operation. If this is a live link, click it to either see more details or go to the entity details page.	(name or description)
Cluster	Specifies the cluster name(s) on which the task is running. A dash (-) appears when the task applies to Prism Central rather than one of the registered clusters.	(cluster names)
Progress	Indicates the current percentage complete for the task.	0%-100%
Status	Indicates the task status, which can be pending, running, completed, or failed.	Succeeded, Running,Failed, Queued
Created On	Displays when the task began.	xx [seconds| minutes| hours| days] ago
Duration	Displays how look the task has been running or took to complete.	xx [seconds| minutes| hours| days]
Sub-Tasks		

Operations Entities

You can access dashboards for the following tools from the **Operations** category of the entities menu (see Entities Menu on page 13):

- Analysis: Analyze cluster activity (see Analysis Dashboard on page 739)
- Cost Management: Analyze and manage costs (see Cost Management (Xi Beam) on page 795)
- App Discovery: Discover applications running in a cluster (see Application Discovery on page 771)
- Monitoring Configurations: Monitor select applications in a cluster (see Application Monitoring on page 796)
- Operations Policies: Manage operations policies (see Operations Policy Management on page 812)
- Planning: Plan for resource capacity and usage (see Resource Planning on page 816)
- Playbooks: Create playbooks for task automation (see Task Automation on page 841)
- Reports: Generate resource and activity reports for Nutanix and Non-Nutanix environments (see Reports View on page 245)

Reports View

The Reports dashboard allows you to create and run reports about the Nutanix and vCenter entities and metrics of your choice. To access the Reports dashboard, select **Operations** > **Reports** from the entities menu (see Entities Menu on page 13).

> Note: This section describes the information and options that appear in the Reports dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See Reports Management on page 883 for detailed information on how to create, configure, and run reports.

Figure 185: Reports Dashboard

The following table describes the fields that appear in the reports list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 104: Reports List Fields

Parameter	Description	Values
Report Name	Displays the name of the report. Note: The **SYSTEM** label appears against the report name for pre-defined reports.	(report name)
Owner	Displays the name of the user who created the report.	(user name)
Scheduled	Indicates whether this report is scheduled to run at specified times. • "Yes" means the report is scheduled to run automatically at specified times (see Frequency). • "No" means the report is not scheduled to run. To run the report you must either manually run it or update the report to specify a schedule.	Yes, No
Frequency	Displays when (how frequently) the report is scheduled to run.	Daily, Weekly, Monthly, Yearly

Parameter	Description	Values
Last Updated Time	Displays the last updated data and time of the report.	mm/dd/yy, hh:mm:ss AM/PM

You can filter the reports by report name. To apply a filter, click the **Filter** button to open the **Filter** pane. Enter a string in the report name field. It will return a list of reports that satisfy the name string.

You can group the reports by predefined status, that is the predefined reports appear as one group while the custom reports appear as a separate group.

The Reports view includes two action buttons: **+ New Report** and **Report Settings**.

- Click the **New Report** button to create a new custom report (see Creating a New Report on page 884).
- Click the **Report Settings** button to configure the report appearance, e-mail, and retention settings (see Configuring Report Settings on page 923).

The **Actions** menu appears when one or more reports are selected. It includes the following actions (see Managing a Report on page 885):

- This action can be applied to multiple reports: **Delete**
- These actions can be applied to only one report at a time: **Run**, **Edit**, **Clone**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the selected report.

Administration Entities

You can access dashboards to manage the following objects from the **Administration** category of the entities menu (see Entities Menu on page 13):

- Users (see Users Summary View on page 246)
- Roles (see Roles Summary View on page 250)
- Projects (see Projects Summary View on page 254)
- Availability Zones (see Availability Zones View on page 259)
- LCM (see LCM View on page 260)

Users Summary View

To access the users dashboard, select **Administration** > **Users** from the entities menu (see Entities Menu on page 13). The users dashboard allows you to view a list of local and Active Directory-specified users.

> Note: This section describes the information and options that appear in the users dashboard.
>
> - See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
> - See “” () for information about creating and maintaining local user accounts.

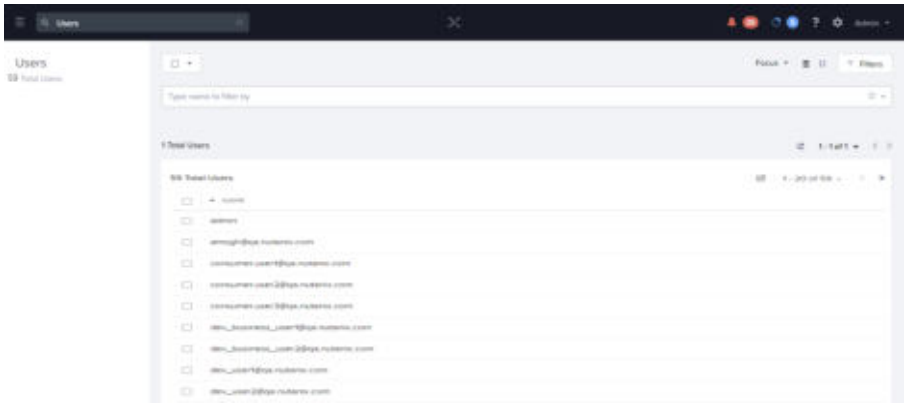

Figure 186: Users Dashboard

The Users view displays a list of local users and project-assigned Active Directory users. You can order the list alphabetically (A-Z or Z-A) by clicking (toggling) the **Name** column header. Clicking on a user name displays the details page for that user.

From the**Filter** pane you can filter the users list by name. Check the name box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field. It will return a list of users that satisfy the name condition/string.

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by name.

User Details View

To access the details page for a user, go to the users dashboard (see Users Summary View on page 246) and click the user name. The details page includes the following:

- The user name (upper left). You can switch from one user details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Trends**, **VM**, **Project**) to display that tab contents below.

Summary Tab

Clicking the **Summary** tab, which appears by default when you first open the page, displays the following:

- A section on the left that displays summary information about the user (see following table).
- A section of the right that displays vCPU, storage, and memory usage graphs for the top five VMs for this user. Place the cursor in a metric graph to see which VM is using that portion of the resource (vCPU, storage, or memory).

The following table describes the project summary information fields. A dash (-) is displayed in a field when a value is not available.

Table 105: User Summary Fields

Parameter	Description	Values
Display Name	Displays the name that appears in the Prism Central displays. This might be a truncated version of the name (next parameter) when that name is long.	(user display name)
Name	Displays the user name.	(user name)
Projects Count	Displays the number of projects to which this user is a member.	(number)
Total VMs	Displays the number of VMs assigned to this user.	(user name)
vCPU Usage	Displays the number of vCPUs used by this user.	(number)
Memory Usage	Displays the amount of memory used by this user.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this user.	xxx [GiB]

Figure 187: User Summary Tab

Trends Tab

The **Trends** tab displays the following graphs.

- The Total VMs graph displays a rolling time interval monitor of VMs owned by this user. Placing the cursor anywhere on the horizontal axis displays the number at that time.
- The **vCPU Usage** graph displays a rolling time interval monitor of user vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of user memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of user storage usage.

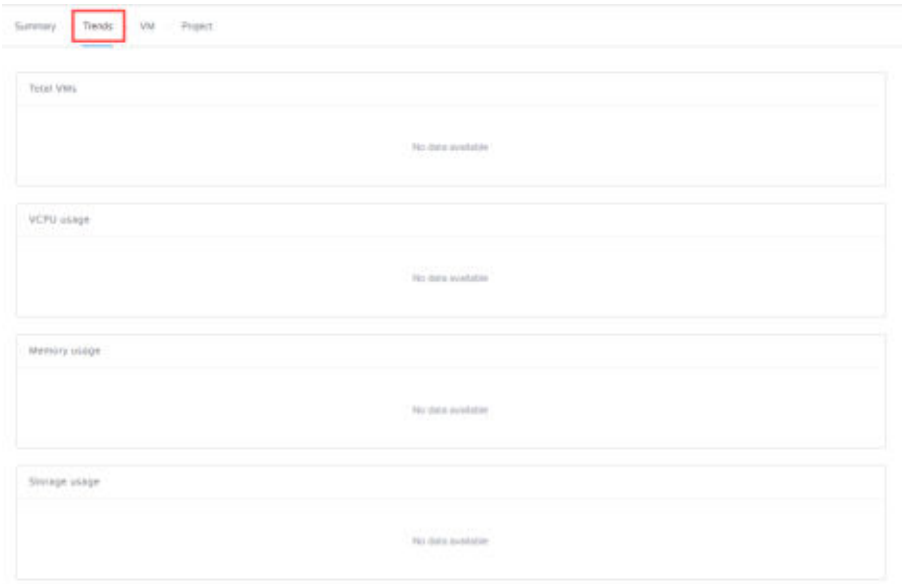

Figure 188: User Trends Tab

VM Tab

The **VM** tab displays a table of VMs owned by the user. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the VM information fields.

Table 106: User VM Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
Project Name	Displays the name of the project associated with this VM.	(project name)
Virtual CPU Count	Displays the number of vCPUs allocated to the VM.	(number)
Memory Capacity	Displays the amount of memory capacity allocated to the VM.	xxx [GiB]
Disk Capacity	Displays the amount of disk capacity allocated to the VM.	xxx [GiB]

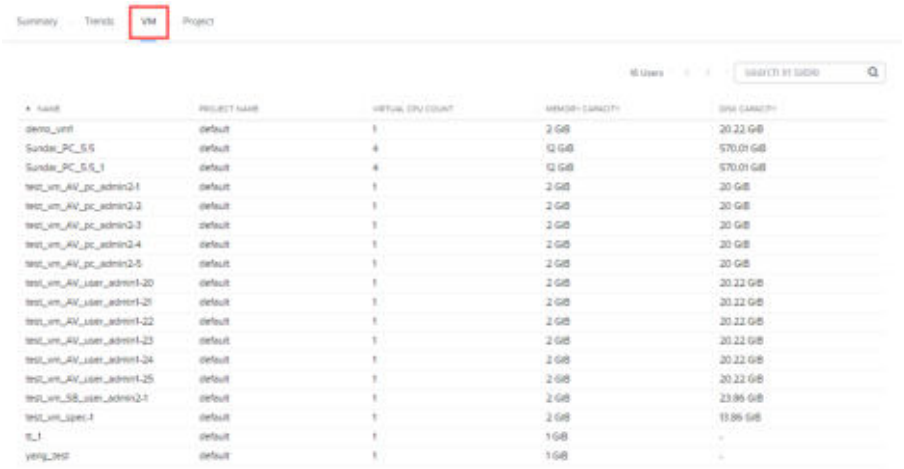

Figure 189: User VM Tab

Project Tab

The **Project** tab displays a table of projects to which the user is a member. You can filter the list by entering a string in the search field (upper right above the table). The following table describes the users information fields.

Table 107: User Project Fields

Parameter	Description	Values
Name	Displays the name of the project.	(project name)
vCPU Usage	Displays the number of vCPUs used by this project.	(number)
Memory Usage	Displays the amount of memory used by this project.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this project.	xxx [GiB]

Figure 190: User Project Tab

Roles Summary View

To access the roles dashboard, select **Administration** > **Roles** from the entities menu (see Entities Menu on page 13). The roles dashboard allows you to view summary information about built-in (default) and custom (created) roles and to access detailed information about each role.

Note: This section describes the information and options that appear in the roles dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize this information in a variety of ways.
- See “” () for information on how to create, manage, and apply roles.

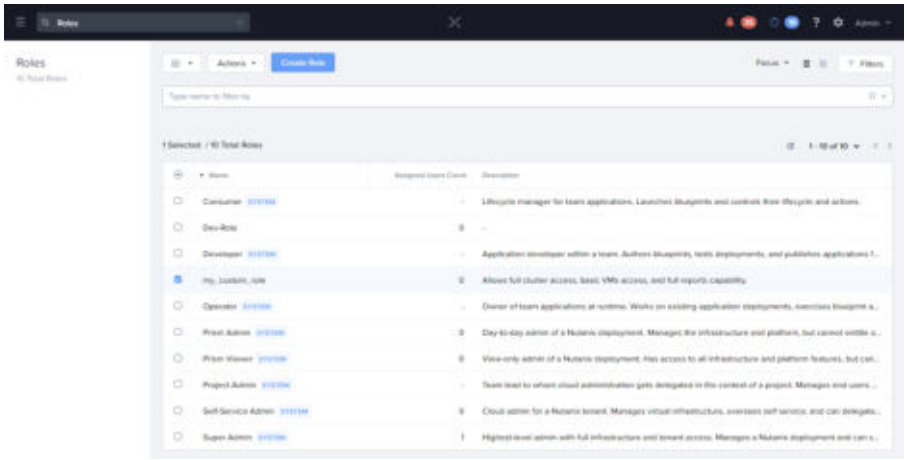

Figure 191: Roles Dashboard

The following table describes the fields that appear in the Roles list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 108: Roles List Fields

Parameter	Description	Values
Name	Displays the role name.	(role name)
Assigned Users Count	Displays the number of users assigned this role.	(number)
Description	Describes the role (if a description was provided when the role was created or updated).	(number)

You can filter the roles list based on several parameter values. The following table describes the filter options available when you open the roles dashboard **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 109: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the role name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of role that satisfy the name condition/ string.	(project name string)
Assigned Users Count	Filters on the number of assigned users. Enter a range in the "**from <low> to <high>**" field. It will return a list of roles that have an assigned user count within that range.	([xx] to [yy] range)

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by name or assigned user count. (You can only choose one parameter.)

The dashboard includes a **Create Role** button (see “” ()). The **Actions** menu appears when one or more roles are selected. It allows you to manage assignments for, duplicate, update, or delete a role (see Role Details View on page 252).

> Note: You cannot delete or modify the default roles.

Role Details View

To access the details page for a role, go to the roles dashboard (see Roles Summary View on page 250) and click the role name. The details page includes the following:

- The role name (upper left). You can switch from one role details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Users**, **User Groups**, **Role Assignment**) to display that tab contents below.
- Action buttons (upper right). The available actions appear in bold; other actions are grayed out.
  - **Manage Assignment**: Click this button to assign the role to users, groups, and entities (see “” ()). This button does not appear for certain default roles such as Prism Admin that you are not allowed to assign.
  - **Duplicate**: Click this button to duplicate the role. It opens the role configuration screen preconfigured with the same permissions as this role (see “” ()).
  - **Delete**: Click the button to delete the role. You cannot delete or modify the default roles.
  - **Update Role**: Click this button to modify the role permissions (see “” ()).
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

Summary Tab

Clicking the **Summary** tab, which appears by default when you first open the page, displays the following:

- A section on the left that displays summary information about the role (see following table).

- A section in the middle which displays a list of actions that the role is permitted to execute. Click the **Actions** header to display additional actions for those roles with more than a single page of permitted actions. You can filter the actions list by entering a string in the search field (upper right above the list).

The following table describes the role summary information fields. A dash (-) is displayed in a field when a value is not available.

Table 110: Role Summary Fields

Parameter	Description	Values
Name	Displays the role name.	(role name)
Assigned Users	Displays the number of users assigned this role.	(number)
Assigned User Groups	Displays the number of user groups assigned this role.	(number)

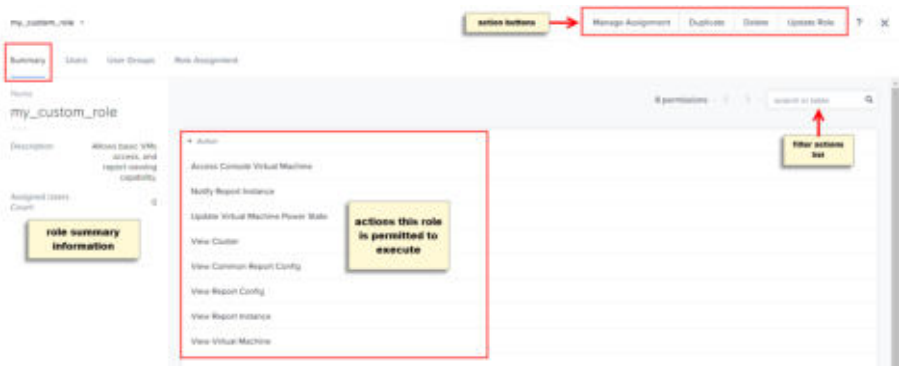

Figure 192: Role Summary Tab

Users Tab

The **Users** tab displays a list of users assigned to the role. You can filter the list by entering a string in the search field (upper right above the table).

User Groups Tab

The **User Groups** tab displays a list of user groups assigned to the role. You can filter the list by entering a string in the search field.

Roll Assignment Tab

The **Roll Assignment** tab displays the relationship between the users, groups, and entities assigned to this role. This tab does not appear for certain default roles such as Prism Admin that you are not allowed to assign.

Figure 193: Role Assignment Tab

Projects Summary View

To access the projects dashboard, select **Administration** > **Projects** from the entities menu (see Entities Menu on page 13). The projects dashboard allows you to view summary information about projects that have been created and access detailed information about each project.

Note: This section describes the information and options that appear in the projects view.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- See Project Management on page 718 for information on how to create and manage projects.

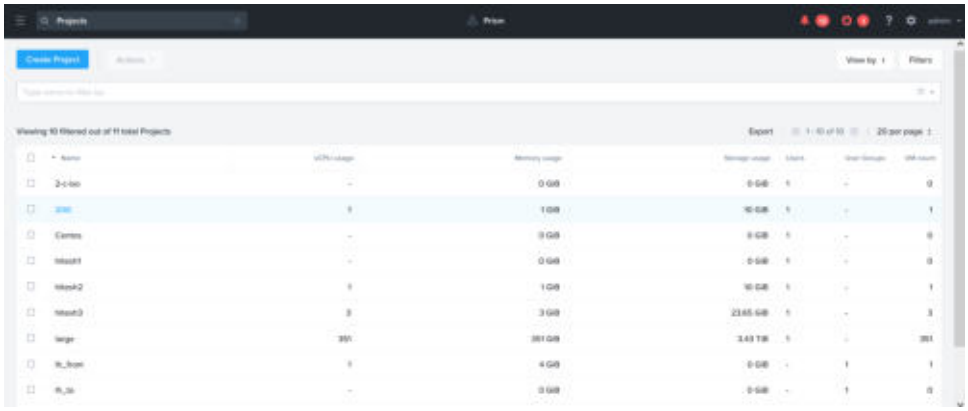

Figure 194: Projects Dashboard

The following table describes the fields that appear in the Projects list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 111: Projects List Fields

Parameter	Description	Values
Name	Displays the project name.	(project name)
vCPU Usage	Displays the number of vCPUs used by the project.	(number)
Memory Usage	Displays the amount of memory used by the project.	xxx [GiB]

Parameter	Description	Values
Storage Usage	Displays the amount of storage used by the project.	xxx [GiB]
Users	Displays the number of users in the project.	(number)
User Groups	Displays the number of user groups in the project.	(number)
VM Count	Displays the number of VMs in the project.	(number)

You can filter the projects list based on several parameter values. The following table describes the filter options available when you open the Projects view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 112: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the project name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of projects that satisfy the name condition/string.	(project name string)
vCPU Usage	Filters on the number of vCPUs. Enter a range in the "**from <low> to <high>**" field. It will return a list of projects within that range of vCPU usage.	([xx] to [yy] range)
Memory Usage	Filters on the amount of memory used. Enter a range in the "**from <low> to <high> GiB**" field. It will return a list of projects within that range of memory usage.	([xx] to [yy] GiB range)
Storage Usage	Filters on the amount of storage space used. Enter a range in the "**from <low> to <high> GiB**" field. It will return a list of projects within that range of storage usage.	([xx] to [yy] GiB range)
VM Count	Filters on the number of VMs. Enter a range in the "**from <low> to <high>**" field. It will return a list of projects within that range of total number of VMs.	([xx] to [yy] range)

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by name, vCPU usage, memory usage, storage usage, or VM count. (You can only choose one parameter.) This is the only grouping option available in the Projects view.

The Projects view includes a **Create Project** button (see Creating a Project on page 718). The **Actions** menu appears when one or more projects are selected. It allows you to update or delete the project (see Modifying a Project on page 721). You can update only one project at a time, but you can select multiple projects when deleting.

Note: You cannot delete the default project.

Project Details View

To access the details page for a project, go to the projects dashboard (see Projects Summary View on page 254) and click the project name. The details page includes the following:

- The project name (upper left). You can switch from one project details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Usage**, **VMs**, **Users**) to display that tab contents below.
- Action buttons (upper right). The available actions appear in bold; other actions are grayed out.
  - **Delete**: Click the button to delete the project. You cannot delete or modify the default project.
  - **Update Project**: Click this button to modify the role permissions (see *Modifying a Custom Role* in the Security Guide).
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

Summary Tab

Clicking the **Summary** tab, which appears by default when you first open the page, displays the following:

- A section on the left that displays summary information about the project (see following table).
- A section of the right that displays vCPU, storage, and memory usage graphs for the top five users and top five VMs in the project. Place the cursor in a metric graph to see which user or VM is using that portion of the resource (vCPU, storage, or memory).

The following table describes the project summary information fields. A dash (-) is displayed in any table field when a value is not available or applicable.

Table 113: Project Summary Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
VM Count	Displays the number of VMs in the project.	(number)
Users	Displays the number of users in the project.	(number)
User Groups	Displays the number of user groups in the project.	(number)
Networks	Displays the networks that the project is allowed to access.	(network names)
vCPU Usage	Displays the number of vCPUs used by the project.	(number)

Parameter	Description	Values
Memory Usage	Displays the amount of memory used by the project.	xxx [GiB]
Storage Usage	Displays the amount of storage used by the project.	xxx [GiB]
Description	Displays a project description, which is typically who created the project.	(description)

Figure 195: Project Summary Tab

Usage Tab

The **Usage** tab displays the following graphs. Set the graph time period by selecting the duration (1 day or 1 week) from the **Show** drop-down menu.

- The **vCPU Usage** graph displays a rolling time interval monitor of project vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of project memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of project storage usage.

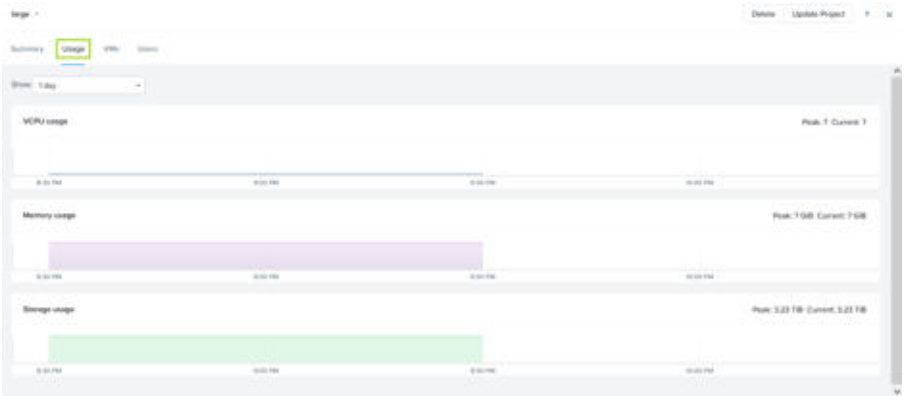

Figure 196: Project Usage Tab

VMs Tab

The **VMs** tab displays a table of VMs assigned to the project. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the VMs information fields.

Table 114: VMs Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
Power State	Displays whether the VM is currently on or off.	On (green circle), Off (red circle)
Owner	Displays the name of the user who owns this VM.	(user name)
Virtual CPU Count	Displays the number of vCPUs allocated to the VM.	(number)
Memory Capacity	Displays the amount of memory allocated to the VM.	xxx [GiB]
Disk Capacity	Displays the amount of disk space allocated to the VM.	xxx [GiB]

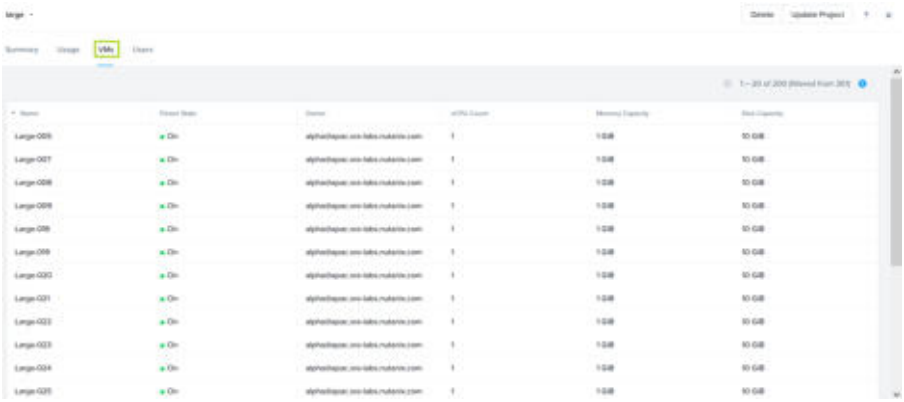

Figure 197: Project VMs Tab

Users Tab

The **Users** tab displays a table of users (individuals and group members) assigned to the project. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the users information fields.

Table 115: Users Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)

Parameter	Description	Values
Total VMs	Displays the number of VMs owned by this user.	(user name)
vCPU Usage	Displays the number of vCPUs used by this user.	(number)
Memory Usage	Displays the amount of memory used by this user.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this user.	xxx [GiB]

Figure 198: Project Users Tab

Availability Zones View

To access the availability zones dashboard, select **Administration** > **Availability Zones** from the entities menu (see Entities Menu on page 13). The availability zones dashboard allows you to view information about the availability zones connected to Prism Central currently and to connect to other zones.

Note: This section describes the information and options that appear in the availability zones dashboard.

- See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- See the *Xi Leap Administration Guide* for information about availability zones and how to connect to them.

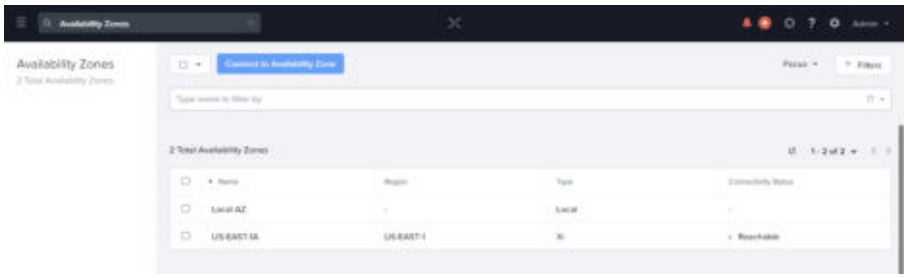

Figure 199: Availability Zones Dashboard

The following table describes the fields that appear in the availability zones list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 116: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the zone name.	(zone name)
Region	Displays the region in which the zone is located.	(region name)
Type	Displays the type or zone. The availability zone that you are logged in to is shown as a local availability zone. Availability zones that are backed by on-premises Prism Central instances are shown as type physical. Availability zones in Xi Cloud Services are shown as type Xi.	Local, Physical, Xi
Connectivity Status	Indicates whether Prism Central can access the zone currently.	Reachable, Unreachable

To filter the list based on the type of zone (local, physical, or Xi), click the **Filters** button and select the desired type(s) in the **Filters** pane.

Click the **Connect to Availability Zone** button to connect to a new availability zone. The **Actions** menu appears when one or more zones are selected. It allows you to disconnect from the selected zones.

LCM View

To access the LCM dashboard, select **Administration** > **LCM** from the entities menu (see Entities Menu on page 13). The life cycle manager (LCM) tracks software and firmware versions of the various components in a cluster. It allows you to view information about the current inventory and update the versions as needed. See the Life Cycle Manager Guide for more information.

Services Entities

You can enable the following services from the **Services** category of the entities menu (see Entities Menu on page 13):

- Calm (see Enabling Calm on page 928)
- Files (see Enabling Files on page 929)
- Foundation Central (see Enabling Foundation Central on page 929)
- Karbon (see Enabling Karbon on page 929)
- Objects (see Enabling Objects on page 930)

5

ALERT AND EVENT MONITORING

Prism Central provides several mechanisms to monitor and configure cluster alerts and events.

- The Alerts and Events dashboards, which you can access from the entities menu (see Entities Menu on page 13), allow you to view and monitor alert and event activity across the registered clusters.
  - The Alerts Summary View on page 261 displays a list of alerts that you filter in various ways. You can also drill down for detailed information and corrective actions for an alert (see Alert Details on page 269)
  - The Events Summary View on page 272 displays a list of events that you filter in various ways. You can also drill down for detailed information for an event (see Event Details on page 274).
- You can configure rules for who should receive email alerts and customize the alert messages (see Configuring Alert Emails on page 275).
- You can customize the list of events that generate an alert (see Configuring Alert Policies on page 278).
- You can review event log files when necessary (see Prism Central Logs on page 299).

> Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default. (Some DR events are retained for just 7 days by default.)

Alerts Summary View

The Alerts dashboard allows you to view summary information about alert messages across the registered clusters, access detailed information about each alert, and view alert policies from any source (user defined, system defined, or external defined).

To access the Alerts dashboard, select **Activity** > **Alerts** from the entities menu (see Entities Menu on page 13). The top pane includes **List** and **Alert Policies** tabs with **User Defined**, **System Defined**, and **External Defined** tabs under **Alert Policies**. Click the tab to display that content below the tab.

> Note: This section describes the information and options that appear in the Alerts dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways. See Alerts/Health checks on page 304 for a list of alert messages and corrective actions.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of active alerts across the registered clusters. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

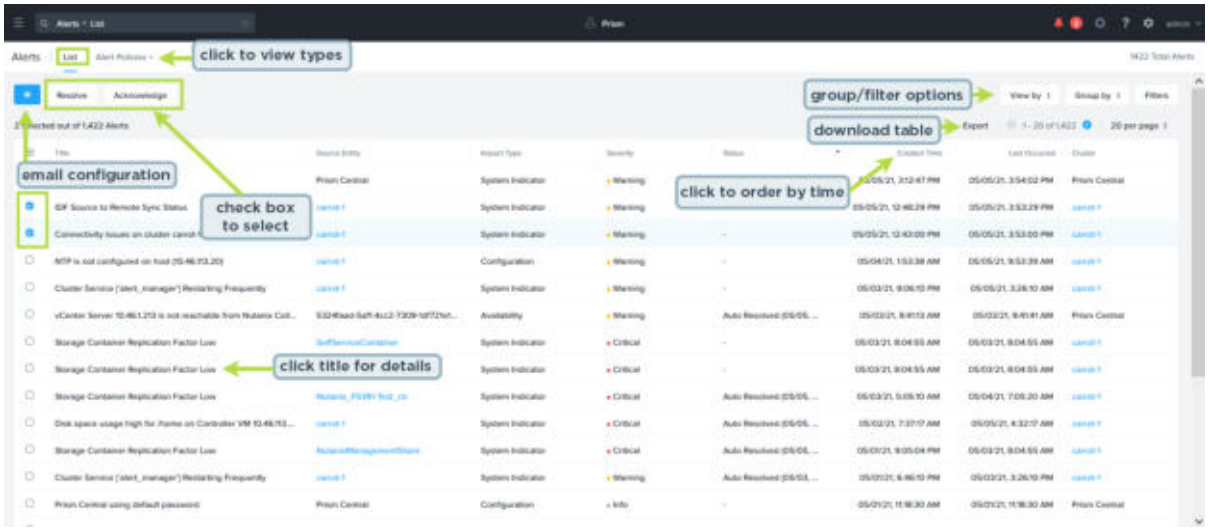

Figure 200: Alerts List Tab

- To order the alerts, click a column header, which toggles the alerts alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of alerts in CSV format, click **Export**. The maximum number of alerts you can export is 1000.
- To display the details page for an alert, click the alert title (see Alert Details on page 269).
- To enable alert emails and specify email addresses to which alerts should be sent, click the **Email Configuration** button (see Configuring Alert Emails on page 275).
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 81).
- To group the alerts, select an option from the **Group** pull-down menu. You can group the alerts by cluster, severity, or impact type. (You can only choose one.)
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Cluster: Enter name in search field.
  - Provider: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Severity: Critical, Warning, Info
  - Resolved: Yes (Manual), Yes (Auto), No
  - Acknowledged: Yes, No
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Create Time: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*
  - Last Occurred: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*
  - Title: Enter a string in the field.

- To select (or unselect) alerts in the list, do one of the following:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all** (or **Clear selections**) to select (unselect) all the alerts in the list.
- When one or more alerts are selected, **Resolve** and **Acknowledge** appear in the menu; click one or both to resolve and/or acknowledge the selected alerts. The user name and time appears in the appropriate field when an alert is acknowledged or resolved.

Table 117: Alerts List Fields

Parameter	Description	Values
(selection box)	Click this box to select the message for acknowledgement or resolution.	n/a
Title	Displays the name of the alert. Click the name to displays the details page for that alert.	(alert name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the alert applies. A comma separate list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Impact Type	Displays the category in which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Severity	Displays the severity level of this condition. There are three levels: **Critical** An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. **Warning** An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. **Info (Informational)** An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational

Parameter	Description	Values
Status	Indicates whether the alert has been resolved. Resolving an error means you set that error as fixed. (The alert may return if the condition is scanned again at a future point.) If you do not want to be notified about the condition again, turn off the alert for this condition (see Configuring Alert Policies on page 278). • A blank value means the alert is not resolved. • An "Acknowledged By" value means this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, check the box for that alert and then select the **Acknowledge** button. • An "Auto Resolved (*`date_time`*)" value means this alert was resolved automatically at the specified date and time. This requires that auto alert is enabled for this alert type. • A "Resolved By *`user`* (*`date_time`*)" value means this alert was resolved manually by the specified user at that date and time. To manually resolve an alert, check the box for that alert and then select the**Resolve** button.	(blank), Acknowledged By *`user`* (*`date_time`*), Auto Resolved (*`date_time`*), Resolved By *`user`* (*`date_time`*)
Create Time	Displays the date and time when the alert occurred.	(date and time)
Last Occurred	Displays the date and time when the alert last occurred (before this occurrence). If this is the first occurrence, both Create Time and Last Occurred display the same date and time.	(date and time)
Cluster	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 205).	(cluster name)

Alert Policies: User Defined Tab

Clicking the **User Defined** tab displays a list of the custom (user defined) alert policies. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

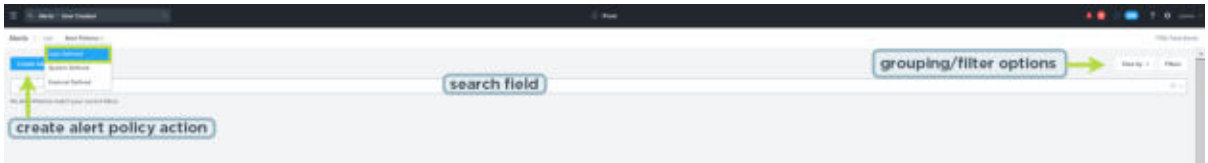

Figure 201: Alert Policies: User Defined Tab

- To create a custom policy, click the **Create Alert Policy** button (see Adding Custom Alert Policies on page 281).
- To view or modify a custom policy, click the policy name (see Modifying Custom Alert Policies on page 286).
- To delete a custom policy, select the policy and then select **Delete** from the **Actions** pull-down menu.
- To order the polices, click a column header, which toggles the list alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 81).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Auto Resolve: Enabled, Disabled
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 118: Alert User Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to displays the policy definition page (see Modifying Custom Alert Policies on page 286).	(policy name)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)

Parameter	Description	Values
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Auto Resolve	Displays whether auto resolved is enabled.	Enabled, Disabled
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Policies: System Defined Tab

Clicking the **System Defined** tab displays a list of the default (system defined) alert policies. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

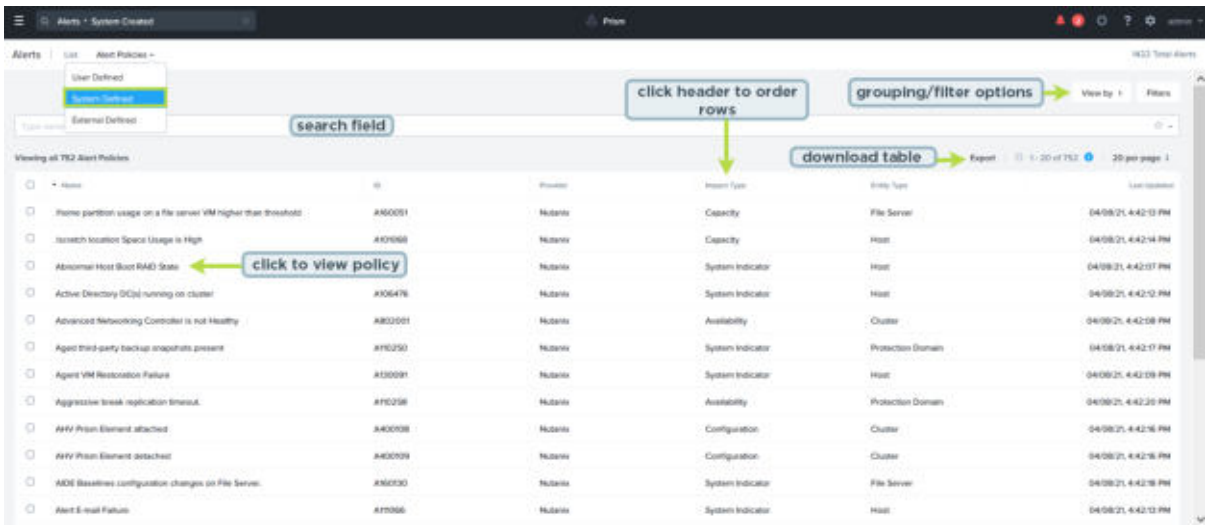

Figure 202: Alert Policies: System Defined Tab

- To view or modify a system policy, click the policy name (see Modifying System Alert Policies on page 279).
- To order the polices, click a column header, which toggles the list alphabetically, numerically, or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 81).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include

in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.

  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - ID: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator

- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 119: Alert System Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to display the Update Policy window (see Modifying System Alert Policies on page 279).	(policy name)
ID	Displays the identification number assigned to the policy.	(ID number)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Policies: External Defined Tab

Clicking the **External Defined** tab displays a list of the alert policies defined by external entities, typically an application through the API. The following figure is a sample view, and the table

describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

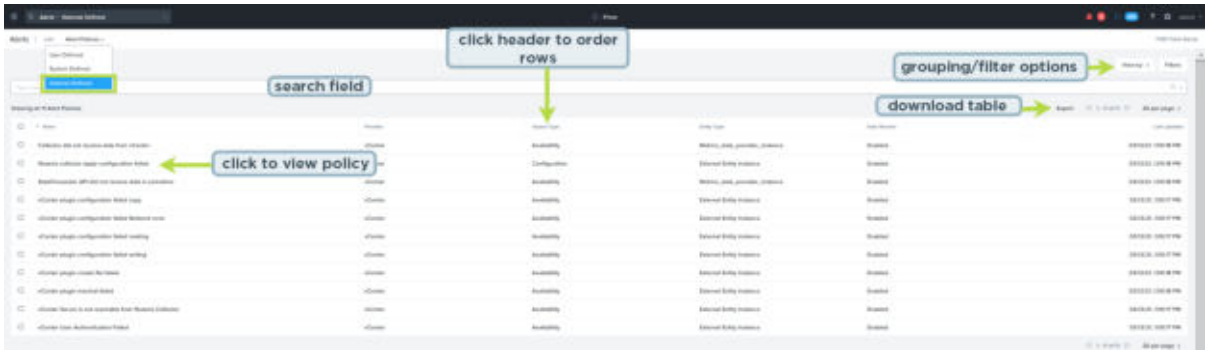

Figure 203: Alert Policies: External Defined Tab

- To view an external policy, click the policy name (see Viewing External Alert Policies on page 296).
- To order the polices, click a column header, which toggles the list alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 81).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Auto Resolve: Enabled, Disabled
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 120: Alert System Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to display the ***policy_name*** window (see Viewing External Alert Policies on page 296)..	(policy name)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Auto Resolve	Displays whether auto resolved is enabled.	Enabled, Disabled
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Details

Clicking on an alert message in the dashboard (see Alerts Summary View on page 261) or anywhere else the alert title appears, such as in a search list, displays detailed information about that alert. The alert details appear in the left column. Possible causes for the alert appear to the right. The most likely cause appears first with other possible causes (if any) appearing below in the order of likelihood. Each cause includes a Recommendation section that describes the recommended corrective action and in many cases a Details section that provides additional context and instructions.

- To switch among alert details pages, select the desired alert from the pull-down list in the upper left of the screen.
- To acknowledge or resolve the alert, click the **Resolve** or **Acknowledge** button in the upper right.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 121: Alert Details Fields

Parameter	Description	Values
Alert Title	Displays the title of the alert.	(alert title)

Parameter	Description	Values
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the alert. Clicking an entity link opens the detail page for that entity.	(entity names)
Severity	Displays the severity level of this condition. There are three levels: **Critical** An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. **Warning** An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. **Informational (Info)** An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational
Cluster Name	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 205).	(cluster name)
Cluster Version	Displays the AOS version running on the cluster.	(version number)
Created Time	Displays the date and time when the alert first occurred.	(time and date)
Last Occurred	Displays the date and time for the most recent occurrence of the alert.	(time and date)
Impact Type	Displays the impact type to which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Policy	Displays the name of the alert policy.	(alert policy name)
Status (for manually resolved alerts)	Displays the current alert status: acknowledged, resolved, or neither (blank).	(blank), Acknowledged, Resolved
Status (for auto resolved alerts)	Displays the current alert status: acknowledged, resolved, auto resolved, or neither (blank).	(blank), Acknowledged, Auto Resolved

Parameter	Description	Values
Acknowledged By	Indicates whether the alert has been acknowledged. • A blank (dash) value means the alert is not acknowledged. • An "N/A *date_time*)" entry indicates this alert was acknowledged automatically at the specified date and time. This requires that auto alert is enabled for this alert type (see Configuring Alert Policies on page 278). • A "*user* *date_time*" entry indicates this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, click the **Acknowledge** button (upper right). The field value changes from blank to the user name and date/time the alert was acknowledged.	(blank), N/A *date_time*), *user* *date_time*
Resolved By	Indicates whether the alert has been resolved. The options (blank, N/A, or user) are the same as for Acknowledged. To manually resolve an alert, click the **Resolve** button (upper right). The field value changes from blank to the user name and date/time the alert was resolved.	(dash), N/A *date_time*), *user* *date_time*
Description	Describes the reason for the alert.	(event description)

Figure 204: Alert Details Screen

Events Summary View

The Events dashboard summary view displays a list of event messages across the registered clusters. Event messages describe cluster actions such as adding a storage pool or taking a snapshot. Unlike alerts, event messages are simply informational without the need to acknowledge or resolve. The following figure is a sample view, and the table describes each field in this view. To access the Events dashboard, select **Activity** > **Events** from the entities menu (see Entities Menu on page 13).

> Note: This section describes the information and options that appear in the Events dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.

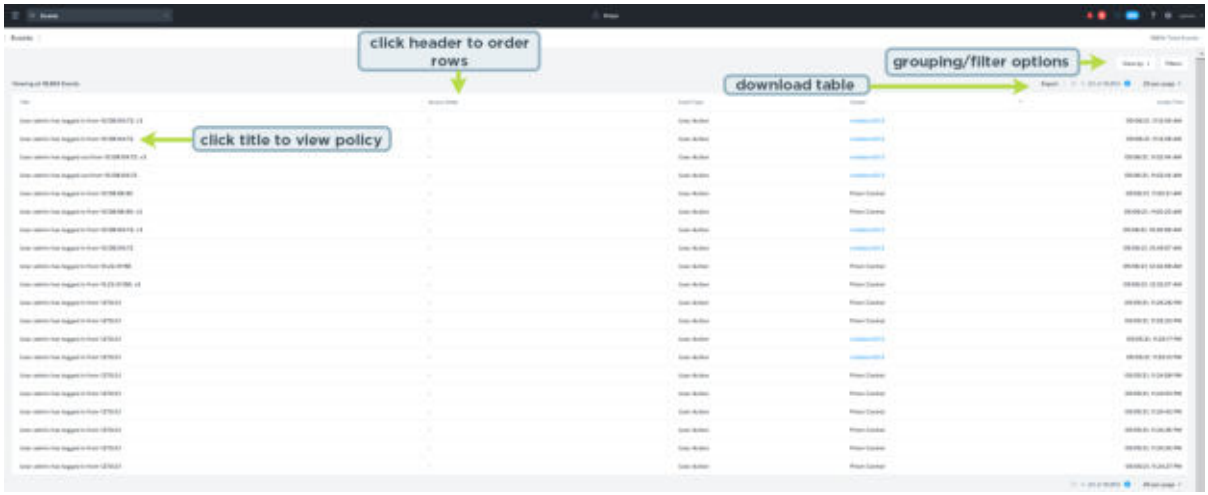

Figure 205: Events Dashboard

- To display the details page for an event, click the event title (see Event Details on page 274).
- To order the events chronologically, click the **Create Time** column header, which toggles the event list latest-to-earliest or earliest-to-latest.
- To download the table of alerts in CSV format, click the **Export** link.
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following event parameters and values.
  - Event Type: Behavioral Anomaly, System Action, User Action
  - Cluster: Enter name in search field.
  - Create Time: Last 1 hour, Last 24 hours, Last week, From *xxx* to *xxx*

Table 122: Events View Fields

Parameter	Description	Values
Title	Displays the name of the event. Click the name to displays the details page for that event (see Event Details on page 274).	(event name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the event applies. A comma separated list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)

Parameter	Description	Values
Event Type	Displays the category in which the event is classified.	System Action, User Action, Behavioral Anomaly (see Behavioral Learning Tools on page 838), DR
Cluster	Displays the name of the cluster in which the event was issued.	(cluster name)
Create Time	Displays the date and time when the event occurred.	(date and time)

Event Details

Clicking on an event message in the dashboard (see Events Summary View on page 272) displays detailed information about that event. The event details appear in the left column. Additional information, such as a description or performance graph, appears to the right.

Figure 206: Event Details Screen

- To switch among event details pages, select the desired event from the pull-down list in the upper left of the screen.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 123: Event Details Fields

Parameter	Description	Values
Event Title	Displays the title of the event.	(alert title)
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the event. Clicking an entity link opens the detail page for that entity.	(entity names)
Cluster Name	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 205).	(cluster name)
Cluster Version	Displays the AOS version running on the cluster.	(version number)

Parameter	Description	Values
Created Time	Displays the date and time when the event occurred.	(time and date)
Event Type	Displays the type of event. Three types are defined, and each event is assigned to one of these types.	Behavioral Anomaly, System Action, User Action

Configuring Alert Emails

About this task

Prism Central allows you to configure the alert messages sent by Prism Central. To configure alert settings, reporting rules, and message templates, do the following:

Note:

- Alert emails sent by Prism Central are in addition to any alert emails you might have configured on individual clusters through the Prism Element web console. You will receive email from both entities in this case. Prism Central alert emailing is not enabled by default; you must explicitly enable it and specify the recipients (Nutanix customer support and/or supplied email addresses). If you enable alerts through Prism Central and do not want to receive double email notifications for the same alert, disable customer email notification for those alerts on the individual clusters through Prism Element (but keep email notification for Nutanix customer support enabled).
- Prism Central requires an SMTP server to send alert email messages (see Configuring an SMTP Server on page 515).

Procedure

1. Do one of the following:

   » Click the gear icon and select **Alert Email Configuration** from the **Settings** menu (see Settings Menu on page 16).

   » In the **Alerts** dashboard, click the **Configure** button and select **Email Configuration** from the pull-down list.

   The Alert Email Configuration window appears.

2. Click the **Settings** tab and do the following:

   a. Enable Prism Central to sent alert emails.

      - Check the **Every Single Alert** box to send an email whenever the event occurs.
      - Check the **Daily Digest** box to send a cumulative (24 hour) list of alerts once a day. (If you click the **Skip empty digest email** box, an email is not sent for the days when no alerts occur.)

      You can check one or both boxes. If neither box is checked, no alert emails are sent by Prism Central.

   b. To send alert notifications to others, enter their email addresses in a comma separate list in the**Email Recipients** field.

   c. Check the **Send alert emails to Nutanix Support** box to send alert notifications to Nutanix customer support.

   d. Click the **Save** button to apply the changes and close the window.

      If you want to configure rules or an email message template, continue to the following steps before clicking the **Save** button.

   The **Tunnel Connection** section displays mail transport status information.

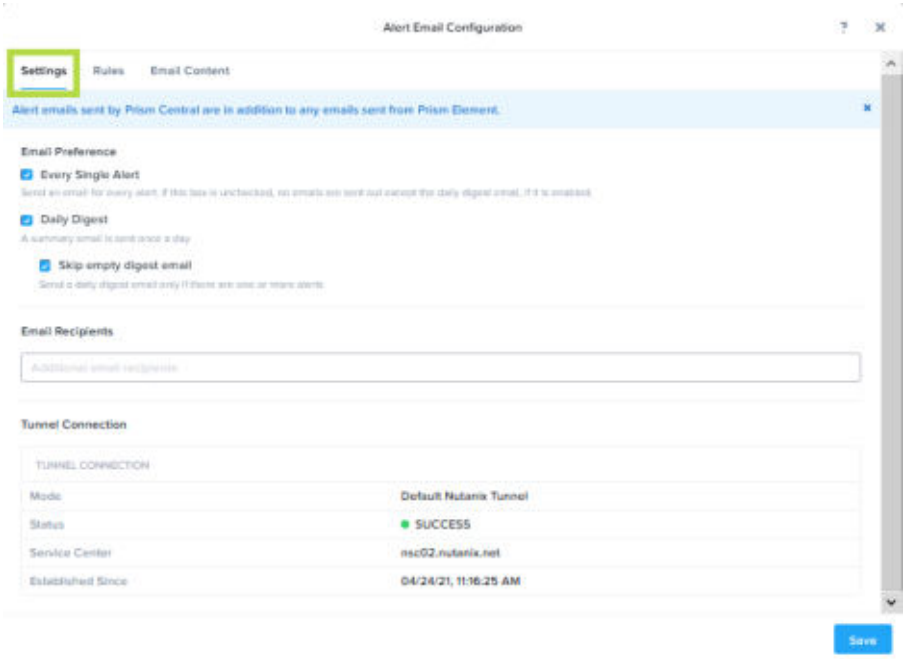

Figure 207: Alert Email Configuration Window (Settings tab)

3. To create a custom alert email rule, click the **Rules** tab and the **New Rule** button, and then configure the rule as follows:

   a. Specify the conditions for generating the alert:

      - **Alert Severity**: Select one or more of the severities from the pull-down list (Critical, Warning, Info, All).
      - **Impact Type**: Select one or more of the categories from the pull-down list (Availability, Capacity, Configuration, Performance, System Indicator, All).
      - **Cluster**: Select one or more of the clusters from the pull-down list (*`cluster_name`*, All).
      - **Alert Contains**: Enter a key phrase or word that should generate an email notification whenever the alert contains that phrase. For example, if you want to get an email notification when an alert contains the phrase "low memory", enter `low memory` in the field.

   b. Specify who should receive the alert email by entering recipient email addresses as a comma separated list in the **Send Email To** field.

   c. Click the **Save Rule** button to save the rule.

   d. Repeat these steps to create more custom rules.

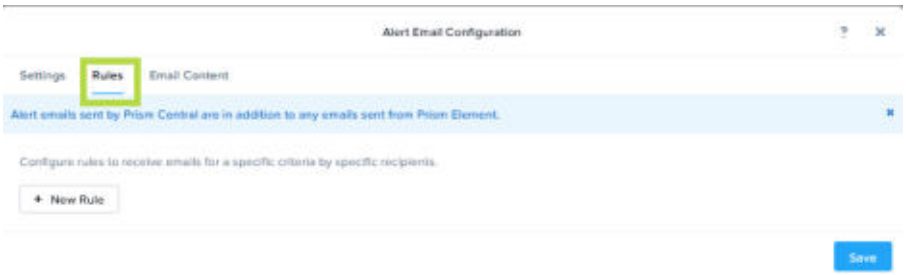

Figure 208: Alert Email Configuration Window (Rules tab)

4. To create a template for the email messages, click the **Email Content** tab and do the following:

   a. In the **Prepend Subject** field, enter the desired text.

      This text appears at the beginning of the subject field in each alert email. If the field is left blank, no prepended text appears in the subject.

   b. In the **Append Body** field, enter the desired text.

      This text appears at the end of the message body in each alert email. If the field is left blank, no appended text appears in the message body.

   c. Click the **Save** button to apply the changes and close the window.

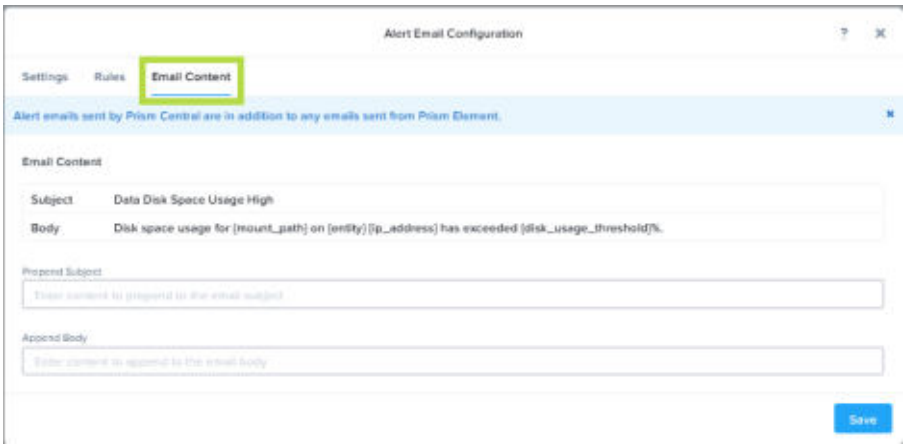

Figure 209: Alert Email Configuration Window (Template tab)

Configuring Alert Policies

About this task

The system monitors a variety of conditions and sends an alert whenever one of the alert conditions is detected (when alerting is enabled). There are default policies for these alerts, but you have the option to modify the default policies and add new policies.

> Note: To help ensure that Prism Central and each managed cluster are taking advantage of NCC features, ensure that:
>
> - Each node in your cluster is running the same NCC version.
> - Prism Central and each cluster managed by Prism Central are all running the same NCC version.

There are three categories of alert policies:

Procedure

- System Defined Alert Policies. You can view and customize policies for any of the system alerts (see Modifying System Alert Policies on page 279).
- User Defined Alert Policies. You have the option to create custom alert policies (see Adding Custom Alert Policies on page 281).

- External Defined Alert Policies. There is an API that external sources can use to define alert policies. These are policies you can view (but not modify) through the alerts dashboard (see Viewing External Alert Policies on page 296).

Modifying System Alert Policies

About this task

There are default policies for system alerts, but you can customize which system events should trigger an alert. To modify the system alert policies, do the following:

Procedure

1. Go to the **System Defined** tab of the alerts dashboard and click the name of the target system alert policy (see Alerts Summary View on page 261).
2. In the **Update Policy** window, do the following in the indicated fields:
   a. **Global Rule**: Uncheck (or check) the box next to the severity to disable (or re-enable) this as an alert.

      All the alerts are enabled (box checked) by default. In most cases this field includes just a single box with the word Critical, Warning, or Info indicating the severity level. Checking the box means this event will trigger an alert of that severity. Unchecking the box means an alert will not be issued when the event occurs. In some cases, such as in the example figure about disk space usage, the event can trigger two alerts, a warning alert when one threshold is reached (in this example 75%) and a critical alert when a second threshold is reached (90%). In these cases you can specify whether the alert should be triggered (check/uncheck the box) and at what threshold (enter a percentage in the box).

   b. **Auto Resolve These Alerts**: Uncheck (or check) the box to disable (or re-enable) automatic alert resolution.

      Automatic alert resolution is enabled for all alert types (where applicable) by default. When this is enabled, the system will automatically resolve alerts under certain conditions such as when the system recognizes that the error has been resolved or when the

initiating event has not reoccurred for 48 hours. (Automatic resolution is not allowed for some alert types, and this is noted in the policy window for those types.)

c. **Exceptions**: To exclude a cluster from this global rule, click the **Add Exceptions** link to open an exceptions line and then select the target cluster from the pull-down list.

   The same check boxes for severity and auto resolve appear. Create a custom rule for the cluster by checking (unchecking) the boxes that should apply to that cluster. Repeat for each cluster you want changed from the global rule.

d. Click the **Save** button to save the changes and return to the alerts dashboard.

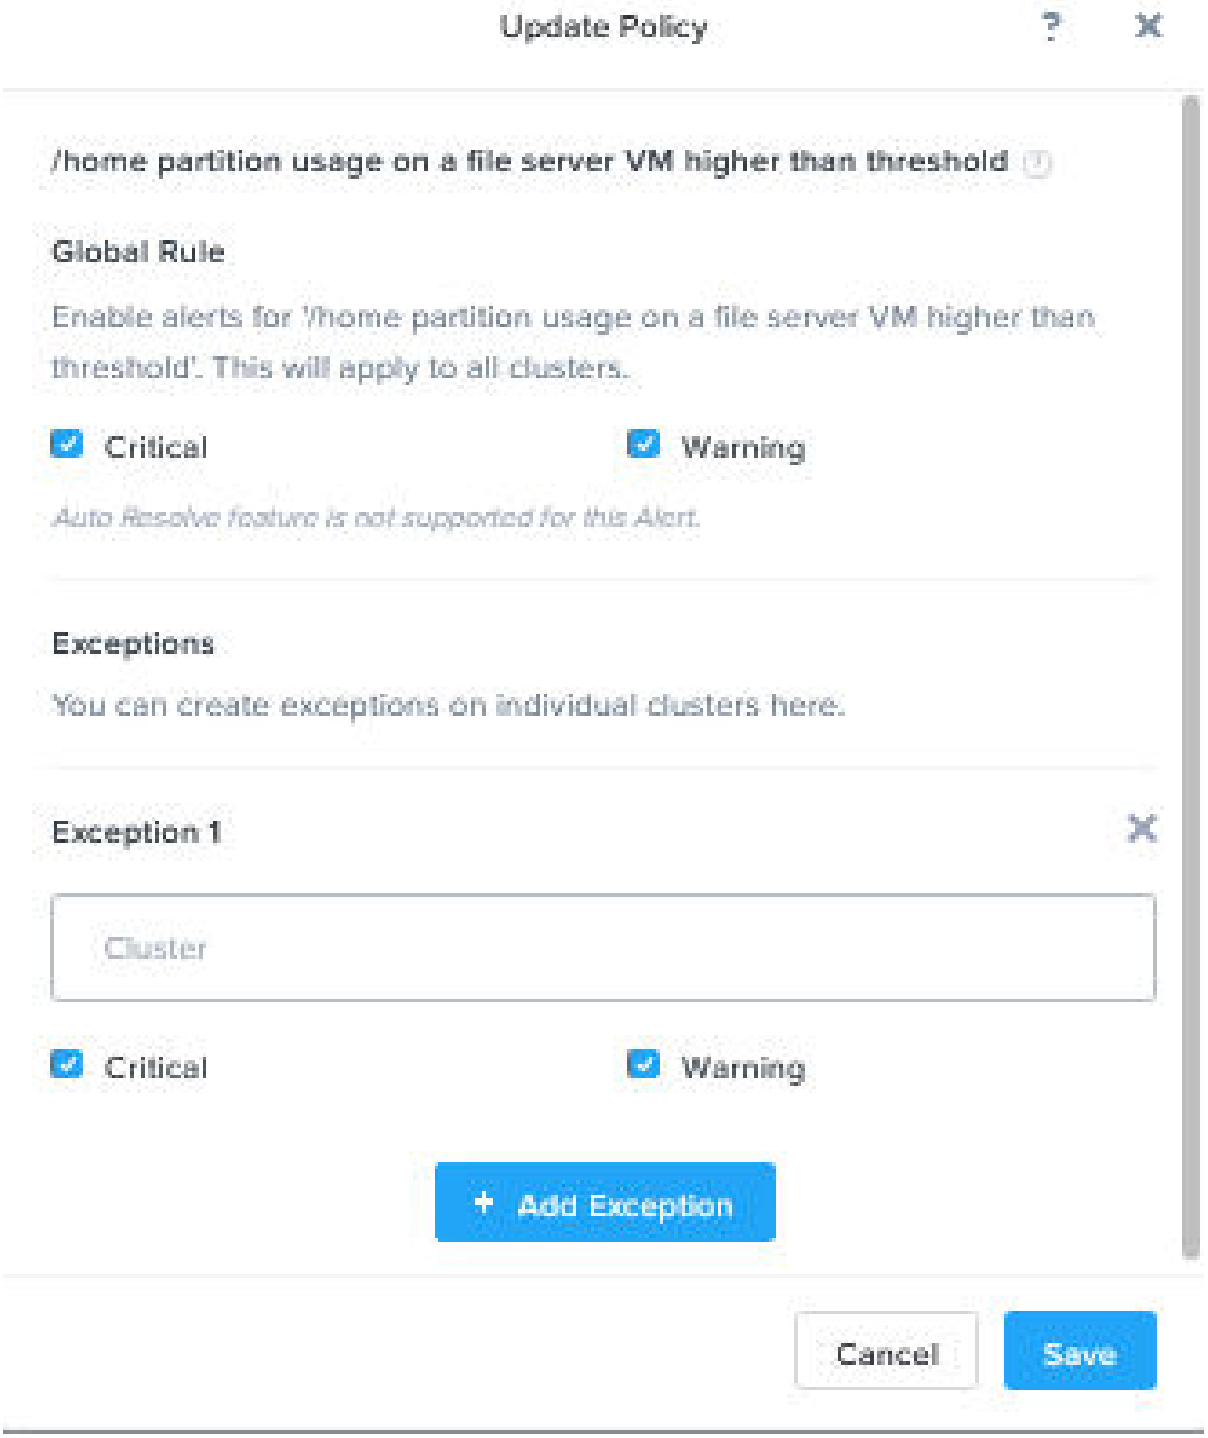

Figure 210: Update Policy Window (example)

Adding Custom Alert Policies

Prism Central provides a mechanism to define alert policies to monitor the entities with respect to the metrics that are critical to your virtual infrastructure. You can define alert policies to monitor entities such as VMs, hosts, and clusters. You can also define a single alert policy for either all VMs, all hosts, or all clusters that share some common criteria. You can create, update,

delete, enable, and disable the alert policies. Each enabled alert policy is executed every five minutes. You can also configure an alert policy so that Prism Central can automatically resolve the alert after the condition causing the alert is resolved.

For a complete description of the metrics available to create custom alert policies, see Alert Metrics on page 286.

Creating Custom Alert Policies

About this task

To create a custom alert policy, do the following.

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and click the **Create Alert Policy** button (see Alerts Summary View on page 261).
   The Create Alert Policy page appears.

Figure 211: Create Alert Policy page (example)

2. In the left pane, do the following in the indicated fields:
   a. **Entity Type**: Select the entity type (**Host**, **VM**, or **Cluster**) from the pull-down list.
   b. **Entity**: Select the target entity from the pull-down list. The choices allow you to specify a single entity or a group of entity. For example, if the entity type is **Host**, you can specify

All Hosts, **One Host**, or **All Hosts in a Cluster**. For all options except **All Hosts**, a search field appears to select the target host, category, or cluster.

> Note: The category option is available only if you select the **VM** entity type.

c. **Metric**: Select the desired metric from the pull-down list.

   The list of metrics varies depending on the entity type. When the metric is selected, more fields appear in the right pane to specify the conditions for the policy.

d. **Policy Name**: Enter a name for the policy.

   The system automatically generates a name for this field based on the policy criteria you have selected, but you can edit the name as desired.

e. **Description** (optional): Enter a description for the policy.

f. **Impact Type**: Select the type of impact the alert condition causes to the entity (**Performance**, **Capacity**, **Configuration**, or **Availability**) from the pull-down list. The default impact type is **Performance**.

g. **Auto resolve alerts**: Check this box to automatically resolve the alert generated by the policy after the condition causing the alert no longer persists (the condition must not occur again for 48 hours).

h. **Enable Policy**: Check this box to enable (or uncheck to disable) the policy.

3. In the right pane, do the following in the indicated fields:

   a. **<Metric Name>**: The top field displays an anomalies chart from the (machine learned) expected behavior for the specified metric. Hovering over the data displays additional information. Select the chart duration (last 12 hours, last 24 hours, or last week) from the pull-down list on the right.

      > Note: The following fields allow you to specify that an alert should be generated when the metric deviates from the expected behavior (**Behavioral Anomaly** fields), when the metric

crosses a set limit (**Static Threshold** fields), or both. See Behavioral Learning Tools on page 838 for more information about behavioral anomalies.

Note the following maximum limit of **Behavioral Anomaly** policies for VMs:

- For a small Prism Central VM (4 vCPUs, 16 GB memory) with more than 4000 VMs, you must not have more than 10 **Behavioral Anomaly** policies for VMs.
- For a large Prism Central VM (8 vCPUs, 32 GB memory) with more than 10,000 VMs, you must not have more than 5 **Behavioral Anomaly** policies for VMs.
- For a Prism Central VM in general, the total number of alert policies must not exceed 50.

If the number of policies exceeds the recommended maximum limit, the user-created alerts and behavioral anomaly alerts might be skipped.

b. **Every time there is an anomaly, alert**: Check this box to generate an alert when an anomaly occurs. Select **Critical** or **Warning** from the pull-down list to set this as a critical or warning level alert.

c. **Ignore all anomalies between**: Check this box to trigger the alert only when the anomaly occurs outside a certain range (percentages or units). Select the operator and value for the range.

d. **Alert critical if**: Check this box to generate a critical alert. Select the operator and value for the critical alert.

e. **Alert warning if**: Check this box to generate a warning alert. Select the operator and value for the warning alert. This option is only available if the behavioral anomaly threshold is not set.

f. **Trigger alert if condition persists for**: Select the duration from the pull-down list (0-30 minutes) for which the alert condition must persist before the system generates an alert.

4. Click **Save** to save the alert policy.

Overlapping Policies

If you are creating an alert policy specific to only an entity (a VM, host, or cluster), but you have earlier created a global policy with the same metric for all entities of that entity type (all VMs, all hosts, or all clusters), the policy that you are creating becomes an overlapping policy. For example, if you have created a policy (global-host-policy-1) that triggers a critical alert if the memory usage of any host exceeds 95%. However, you are now creating a policy (host1-policy) that triggers a critical alert if the memory usage (same metric) of host1 (a specific host) exceeds 90%, host1-policy is now an overlapping policy for global-host-policy-1.

When you are creating a new policy (see Creating Custom Alert Policies on page 282), if you select the same entity and metric as an existing policy but different values and you click **Save**, Prism Central displays a message that a similar policy exists as shown in the following screenshot:

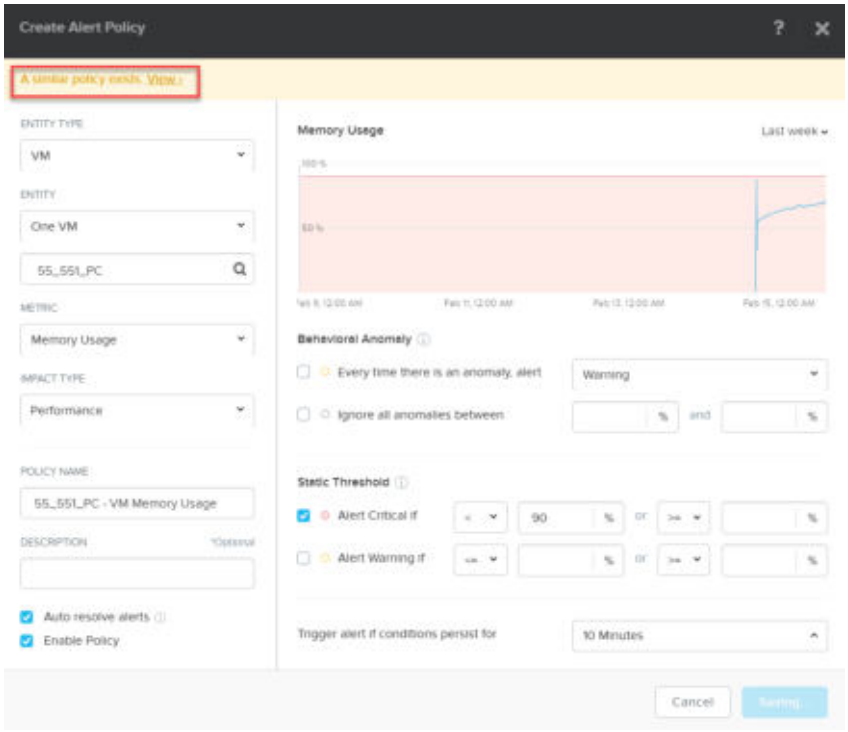

Figure 212: Overlapping Policy

Click **View** to display the overlapping policies and click **Save** to save the overlapping policy.

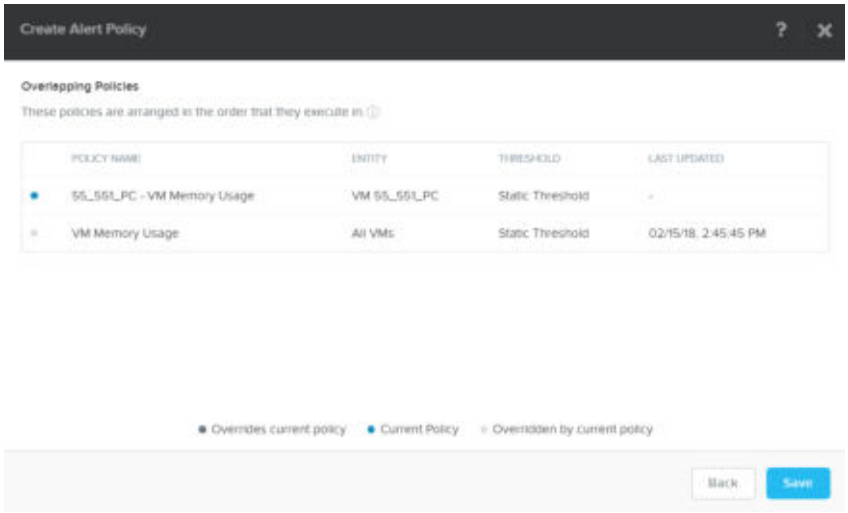

Figure 213: Overlapping Policies List

Precedence of Overlapping Polices

If an overlapping policy exists (policy that applies to a specific entity) for a global policy (policy that applies to all the entities of that entity type), the overlapping policy takes precedence over the global policy.

If multiple policies apply to an entity during policy execution (that is overlapping policies exist), the policy with the highest precedence for that entity is evaluated and all other policies that apply to that entity are ignored. Even if a policy is disabled, the policy might still overlap other

policies. If the policy that has the highest precedence is disabled, that policy is not applied and all other policies that apply to that entity are also ignored.

Policies are evaluated in the following order of precedence (highest precedence policy at the top):

1. Policy applied to a specific entity
2. Policy applied to an entity type in a category
3. Policy applied to an entity type in a cluster
4. Policy applied to all entities of an entity type

Precedence of Overlapping Polices in a Category

A category policy (policy that applies to all VMs in a category) takes precedence over a global policy (a policy that applies to all VMs in the cluster). However, a VM can belong to multiple categories. In this case, the category policy that is most recently updated takes precedence over any other category policy that applies to a VM.

For example, consider that you have created a global policy (global-policy-1) for all VMs in the cluster. You have VM-1 that belongs to both the categories: category-1 and category-2. You have created the category-1-pol policy for all VMs in category-1 and the category-2-pol policy for all VMs in category-2. In this case, all three policies apply to VM-1, but only one policy can be applied to a single VM. Because category-1-pol and category-2-pol are category policies, category-1-pol and category-2-pol take precedence over global-policy-1. However, you still have two policies (category-1-pol and category-2-pol) that apply to the same VM (VM-1). In this case, the policy that is updated most recently takes precedence over the other.

Modifying Custom Alert Policies

About this task

To modify a custom alert policy, do the following:

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and click the name of the target custom alert policy (see Alerts Summary View on page 261).
   The *`policy_name`* page appears. which includes the same fields as the Create Alert Policy page (see Creating Custom Alert Policies on page 282).
2. Make the desired changes in the appropriate fields and then click **Save**.

Deleting Custom Alert Policies

About this task

To delete a custom (user-created) alert policy, do the following:

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and select the target custom alert policy in the list (see Alerts Summary View on page 261).
2. Select **Delete** from the **Actions** pull-down menu.

Alert Metrics

The following table describes the VM metrics.

Table 124: VM Metrics

CPU Ready Time	
Metric	`hypervisor.cpu_ready_time_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the time a virtual machine waits to use the physical CPU out of the total CPU time allotted to the VM. AHV: Ratio of the VM wait time to the total CPU time allotted to the VM in percentage. ESXi: Value of `cpu.ready.summation` for a VM in percentage.
Supported Hypervisors	AHV and ESXi
CPU Usage	
Metric	`hypervisor_cpu_usage_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the total CPU cycles currently in use by the VM reported by the hypervisor. Note: This metric identical to the system created alert policy A3040. ESXi: Value of `cpu.usage.average` reported by ESXi for a VM. Hyper-V: Value of `AvgCpuUtilizationPercent` reported by Hyper-V for a VM.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Memory Usage	
Metric	`hypervisor_memory_usage_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the memory used by the VM out of the total configured memory reported by the hypervisor. Note: This metric is identical to the system created alert policy A3021. AHV: Memory usage is always 100% because AHV does not support overcommit. ESXi: Value of `mem.usage.average` reported by ESXi for a VM.
Supported Hypervisors	ESXi and AHV
Memory Balloon*	
Metric	`hypervisor_memory_balloon_reclaimed_bytes`
Rollup	Average

Unit	Bytes
Description	Total guest memory in bytes reclaimed through ballooning. Note: This metric is similar to the system created alert policy A3024.
Supported Hypervisors	ESXi
Memory Balloon Target*	
Metric	`hypervisor_memory_balloon_reclaim_target_bytes`
Rollup	Average
Unit	Bytes
Description	Target value in bytes for a VM's memory balloon.
Supported Hypervisors	ESXi
Memory Swap In Rate*	
Metric	`hypervisor_swap_in_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the memory swapped in from the disk to active memory in kilobytes per second.
Supported Hypervisors	ESXi
Memory Swap Out Rate*	
Metric	`hypervisor_swap_out_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the memory swapped out from active memory to the disk in kilobytes per second.
Supported Hypervisors	ESXi
Controller AVG Read I/O Latency	
Metric	`controller_avg_read_io_latency_usecs`
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the read I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller AVG Write I/O Latency	
Metric	`controller_avg_write_io_latency_usecs`
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the write I/O latency of the VM in microseconds reported by the controller.

Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller AVG I/O Latency	
Metric	`controller_avg_io_latency_usecs`
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the read and write I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Read IOPS	
Metric	`controller_num_read_iops`
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM reads in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Write IOPS	
Metric	`controller_num_write_iops`
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM writes in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller IOPS	
Metric	`controller_num_iops`
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM read and writes in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Read IO Bandwidth	
Metric	`controller_read_io_bandwidth_kBps`
Rollup	Derived from summary rollup type.
Unit	KBps
Description	Read IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Write IO Bandwidth	
Metric	`controller_write_io_bandwidth_kBps`
Rollup	Derived from summary rollup type.
Unit	KBps

Description	Write IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller IO Bandwidth	
Metric	`controller_io_bandwidth_kBps`
Rollup	Derived from summary rollup type.
Unit	KBps
Description	Read and write IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Network Transmit Packets Dropped	
Metric	`hypervisor_num_transmit_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network transmit packets dropped for the VM.
Supported Hypervisors	ESXi and AHV
Network Receive Packets Dropped	
Metric	`hypervisor_num_receive_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network receive packets dropped for the VM.
Supported Hypervisors	ESXi and AHV
Controller Disk Usage	
Metric	`controller_user_bytes`
Rollup	Average
Unit	bytes
Description	Disk usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
I/O working set size	
Metric	`controller.wss_3600s_union_MB`
Rollup	Average
Unit	MB
Description	Size of the I/O working set size in megabytes.
Supported Hypervisors	All hypervisors
Write I/O working set size	
Metric	`controller.wss_3600s_write_MB`
Rollup	Average
Unit	MB

Description	Size of the I/O working set size writes in megabytes.
Supported Hypervisors	All hypervisors
Read I/O working set size	
Metric	`controller.wss_3600s_read_MB`
Rollup	Average
Unit	MB
Description	Size of the I/O working set size reads in megabytes.
Supported Hypervisors	All hypervisors
Snapshot usage	
Metric	`controller.snapshot_usage_bytes`
Rollup	Average
Unit	bytes
Description	Snapshot usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
Shared Data	
Metric	`controller.shared_usage_bytes`
Rollup	Average
Unit	bytes
Description	Snapshot usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
Network Bytes Received	
Metric	`hypervisor_num_received_bytes`
Rollup	Summation
Unit	bytes
Description	Number of network bytes received reported by the hypervisor.
Supported Hypervisors	All hypervisors
Network Bytes Transmitted	
Metric	`hypervisor_num_transmitted_bytes`
Rollup	Summation
Unit	bytes
Description	Number of network bytes transmitted reported by the hypervisor.
Supported Hypervisors	All hypervisors
* Experimental only. Might not be supported in the future releases.	

Table 125: Host Metrics

The following table describes the host metrics.

CPU Usage	
Metric	`hypervisor_cpu_usage_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the total CPU cycles currently in use by the host. Note: This metric is identical to the system created alert policy A6414. AHV: Ratio between the sum of kernal, user, and iowait and the sum of kernal, user, iowait, and idle over all the CPUs. For hosts with hyperthreads enabled, the host CPU usage is 10% more than the total value. ESXi: Value of `cpu.usage.average` for a host in percentage. Hyper-V: Value of `LogicalProcessorTotalRuntimePercent` from the output of the command `Get-LocalHostPerformance | Convertto-Json -Depth 5 -Compress`.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Memory Swap In Rate*	
Metric	`hypervisor_swap_in_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the swap-in memory in kilobytes per second. Note: This metric is similar to the system created alert policy A6413. ESXi: Value of `mem.swapinRate.average` for the host.
Supported Hypervisors	ESXi
Memory Swap Out Rate*	
Metric	`hypervisor_swap_out_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the swap-out memory in kilobytes per second. Note: This metric is similar to the system created alert policy A6413.

	ESXi: Value of `mem.swapoutRate.average` for the host.
Supported Hypervisors	ESXi
Memory Usage	
Metric	`hypervisor_memory_usage_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Ratio of the total memory allocated to all VMs over the total memory size on the host. Note: This metric is identical to the system created alert policy A6414.
	AHV: Committed memory size from `/proc/meminfo` over the total memory size. Memory usage is always 100%.
	ESXi: Value of `mem.usage.average` for the host.
	Hyper-V: Amount of used memory over the total memory.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Network Transmit Packets Dropped	
Metric	`hypervisor_num_receive_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network transmit packets dropped.
	ESXi: Value of `net.droppedTx.summation` for the host.
Supported Hypervisors	ESXi and AHV
Network Receive Packets Dropped	
Metric	`hypervisor_num_receive_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network receive packets dropped.
	ESXi: Value of `net.droppedRx.summation` for the host.
Supported Hypervisors	ESXi and AHV
Controller AVG I/O Latency	
Metric	controller_avg_io_latency_usecs
Rollup	Average
Unit	Microseconds

Description	Average read and write I/O latency of the host in microseconds reported by the controller.
Controller AVG Read I/O Latency	
Metric	`controller_avg_read_io_latency_usecs`
Rollup	Average
Unit	Microseconds
Description	Average read I/O latency of the host in microseconds reported by the controller.
Controller AVG Write I/O Latency	
Metric	`controller_avg_write_io_latency_usecs`
Rollup	Average
Unit	Microseconds
Description	Average write I/O latency of the host in microseconds reported by the controller.
Controller Read IOPS	
Metric	`controller_num_read_iops`
Rollup	Average
Unit	IOPS
Description	Rate of the host reads in IOPS reported by the controller.
Controller Write IOPS	
Metric	`controller_num_write_iops`
Rollup	Average
Unit	IOPS
Description	Rate of the host writes in IOPS reported by the controller.
Controller IOPS	
Metric	`controller_num_iops`
Rollup	Average
Unit	IOPS
Description	Rate of the host reads and writes in IOPS reported by the controller.
* Experimental only. Might not be supported in the future releases.	

Table 126: Cluster Metrics

The following table describes the host metrics.

Controller Read IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller reads in IOPS.
Controller Write IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller writes in IOPS.
Controller IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller reads and writes in IOPS.
Controller AVG Read I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the read I/O latency of the controller in microseconds.
Controller AVG Write I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the write I/O latency of the controller in microseconds.
Controller AVG I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the read and write I/O latency of the controller in microseconds.
Controller I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Read and write IO bandwidth in kilobytes per second reported by the controller.
Controller Read I/O Bandwidth	
Rollup	Average
Unit	KBps

Description	Read IO bandwidth in kilobytes per second reported by the controller.
Controller Write I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Write IO bandwidth in kilobytes per second reported by the controller.
CPU Usage	
Rollup	Average
Unit	PPM
Description	CPU usage of the cluster in PPM reported by the hypervisor.
Memory Usage	
Rollup	Average
Unit	PPM
Description	Memory usage of the cluster in PPM reported by the hypervisor.

Viewing External Alert Policies

External alert policies are those created by applications.

Nutanix provides an API that applications can use to define alert policies. The **External Defined** tab displays a list of the application-defined alerts (see Alerts Summary View on page 261). You cannot modify an externally defined policy, but you can view the policy by clicking the policy name in the list.

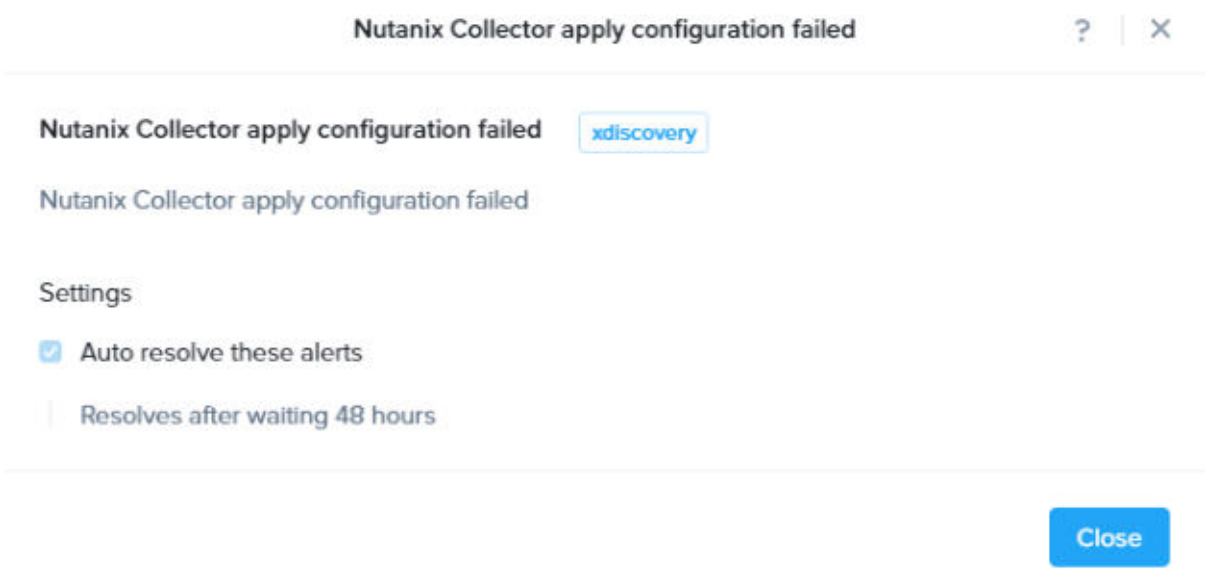

Figure 214: External Defined Policy Window (example 1)

Stopping Alerts During a Maintenance Window

You can stop alerts from being generated during a maintenance window by using Prism Central. You must log on to the Prism Central VM with SSH and create configurations to stop the alerts by using the nuclei command line interface.

About this task

You can stop the alerts of only the **cluster** entity type.

Perform the following procedure to create or modify a configuration to stop the alerts.

Procedure

1. Log on to the Prism Central VM with SSH.
2. Create a specifications file as shown in the following example to create or update a configuration.

```
{
 "spec": {
   "resources": {
     "scope_entity_list": [
       {
         "entity_type": "cluster",
         "uuid_list": ["123", "223"]
       }
     ],
     "schedule_list": [
       {
         "interval_multiple": 1,
         "duration_secs": 300,
         "end_time": "2017-08-20T23:47:45.141Z",
         "start_time": "2017-07-20T23:47:45.141Z",
         "interval_type": "DAILY"
       }
     ],
   }
},
"api_version": "3.0",
"metadata": {
   "kind": "blackout"
}
```

```
}
```

Specify the values for the following parameters:

- Under `scope_entity_list`, specify the scope of the entities for which you want to stop the alerts as follows:
  - entity_type. Specify the type of entities for which you want to stop the alerts. You can stop the alerts of only the **cluster** entity type in this release.
  - uuid_list. Specify the UUIDs of the entities for which you want to stop the alerts. For example, specify the UUIDs of the clusters for which you want to stop the alerts as shown in the example specification file.

    Note: If this field is empty, alerts for all entities of that entity type are stopped. For example, use this parameter if you want to stop the alerts for specific clusters. If you want to stop the alerts for all clusters, leave this field empty.

- Under `schedule_list`, define the schedule during which the alerts should be stopped as follows:
  - interval_type. Define the recurrence to stop the alerts (ONCE, DAILY, WEEKLY, or MONTHLY).
  - interval_multiple. Define the frequency to stop the alerts. This is related to interval_type. For example:
    - If interval_type=Daily and interval_multiple=1, the alert is stopped every day.
    - If interval_type=Daily and interval_multiple=2, the alert is stopped every alternate day.
  - start_time and end_time. Define start and end time to stop the alerts. The date time format is same as used in Nutanix REST API calls.

    ```
    date-fullyear(4digit) "-" date-month(2digit) "-" date-mday(2digit)
    "T" time-hour(2digit) ":" time-minute(2digit) ":" time-second(2digit)
    [.] "Z"  or ("+" / "-") time-hour(2digit) ":" time-minute(2digit)
    ```

  - duration_secs. Define the duration in seconds for which the alerts should be stopped.
  - day_of_week_list. Define the days of the week when the alerts should be stopped. You must include this parameter only if interval_type=WEEKLY.

3. Create a configuration to stop the alerts.

   ```
   $ nuclei --username admin --password password blackout.create spec_file=filepath
   ```

   Replace *password* with the password of the admin user account and *filepath* with the path of the specification file that you created.

4. List all the configurations that stop alerts.

   ```
   $ nuclei --username admin --password password blackout.list
   ```

   Replace *password* with the password of the admin user account.

   The output of this command lists the UUID of the each configuration that you have created.

5. Update a configuration.

   ```
   $ nuclei --username admin --password password blackout.put uuid spec_file=filepath
   ```

   Replace *password* with password of the admin user account, *filepath* with the path of the specification file that you created, and *uuid* with UUID of the configuration that you want to update (obtained in step 4).

6. Get a configuration.

   ```
   $ nuclei --username admin --password password blackout.get uuid
   ```

   Replace *password* with password of the admin user account and *uuid* with UUID of the configuration that you want to get (obtained in step 4).

Prism Central Logs

This section lists the logs that reside in a Prism Central VM. These logs record Prism Central-related events and actions. A Prism Central VM runs a subset of the services that run in a Controller VM plus a service about VMs (Metropolis) that is unique to Prism Central. Prism Central logs contain comparable information as the similarly named ones for a Controller VM. See the "logs" section in the *Acropolis Advanced Administration Guide* for more information.

> Note:
>
> - The timestamps for all Nutanix service logs are moved to UTC (in ISO 8601:2020-01-01 T00:00:00Z) from Prism version 5.18.
> - All operating system logs will not be moved to UTC, hence Nutanix recommends that you set the server local time to UTC.

Table 127: Location: **/home/nutanix/data/logs**

Log	Contents	Frequency
alert_manager.[out, ERROR, FATAL, INFO, WARNING]	Alert manager process output	
aplos.out	Aplos V3 gateway logs	
aplos_engine.out	Aplos V3 engine logs	
athena.[out, ERROR, INFO, FATAL]	Authentication and identity management activity	
cassandra_monitor.[out, ERROR, FATAL, INFO]	Cassandra database monitor process output	
catalina.out	Catalina/Tomcat for Prism process output	
connection_splicer.[out, ERROR, FATAL, INFO, WARNING]	Internal process connection status	
consolidated_audit.log	Audit logs for VM operations, see Audit Log Events on page 302 for details.	
cron_avahi_monitor.log	Avahi process status	

Log	Contents	Frequency
cron_time_check.log	Check time difference across Prism Central VMs when Prism Central is a multi-VM instance.	every 1 min
disk_usage.log	Disk and inode usage of all partitions on the Prism Central VM	every 1 min
dynamic_ring_changer.[out, ERROR, FATAL]	Metadata migration across nodes activity	
genesis.out	Nutanix software start process output	
hyperint_monitor.[out, ERROR, FATAL, INFO, WARNING]	Hypervisor integration activity	
karbon_core.out (formerly acs_controller.out)	Karbon controller output. This file applies only when Karbon is enabled (see Enabling Karbon on page 929).	
karbon_ui.out (formerly acs_ui.out)	Karbon user interface output. This file applies only when Karbon is enabled.	
metropolis.out	Internal Prism Central service about VMs (Metropolis) to validate cluster choice with scheduler, check out images from the catalog service, and forward fanout requests to Prism Element	
prism_gateway.[out, ERROR, FATAL, INFO]	Prism leader activity	
prism_monitor.[out, ERROR, FATAL, INFO]	Prism (web console, nCLI, REST API) monitor process output	
scavenger.out	Log and core file clean-up status	
search.[out, FATAL]		
search_monitor.[ERROR, INFO, WARNING]		
send-email.log	E-mail alerts sent from the Controller VM	every 1 min
ssh_tunnel.log	Connect status to nsc.nutanix.com for the remote support tunnel	
stats_aggregator.[out, ERROR, FATAL, INFO]	Statistics aggregator process output	
support-info.log	Daily automated support (ASUP) alerts	

Log	Contents	Frequency
using-gflags.log	gflags status	
zeus_config_printer.INFO	Contents of cluster configuration database	
zookeeper_monitor.[out, ERROR, INFO]	Cluster configuration and cluster state activity	

Table 128: Location: **/home/nutanix/data/logs/cassandra**

Log	Contents
system.log	Cassandra system activity

Table 129: Location:**/home/nutanix/data/logs/sysstats**

Log	Contents	Frequency	Command
df.info	Mounted filesystems	every 5 sec	df -h
disk_usage.INFO	Disk usage across disks	every 1 hour	du
interrupts.INFO	CPU interrupts	every 5 sec	
iostat.INFO	I/O activity for each physical disk	every 5 sec	sudo iostat
iotop.INFO	Current I/O in realtime	every 5 sec	sudo iotop
lsof.INFO	List of open files and processes that open them	every 1 min	sudo lsof
meminfo.INFO	Memory usage	every 5 sec	cat /proc/meminfo
metadata_disk_usage.INFO	Disk usage for metadata drives	every 5 sec	
mpstat.INFO	CPU activities per CPU	every 5 sec	mpstat
ntpq.INFO	NTP information	every 1 min	ntpq -pn
ping_gateway.INFO	Pings to the default gateway	every 5 sec	ping
ping_hosts.INFO	Pings to other Prism Central VMs (if Prism Central is a multi-VM instance)	every 1 min	ping
sar.INFO	Network bandwidth	every 5 sec	sar -n DEV, -n EDEV
top.INFO	Real-time CPU and memory activity	every 5 sec	top

Table 130: Location: **/home/nutanix/data/serviceability/alerts**

Log	Contents
num.processed	Alerts that have been processed

Table 131: Location: **/var/log**

Log	Contents
dmesg	OS start messages
kernel	OS kernel messages
messages	OS messages after starting

Audit Log Events

The following table provides the list of all the VM events that are audited and captured as part of the audit logs.

Table 132: Audit Event List

Event Type	Operation	Description
VmCreateAudit	Create	Created VM
VmReplicateAudit	Create	Replicated VM
VmCloneAudit	Create	Cloned VM
VmDeleteAudit	Delete	Deleted VM
VmRestoreAudit	Restore	Restored VM from snapshot
VmNicAddAudit	Update	Added NIC to VM
VmNicUpdateAudit	Update	Updated NIC of VM
VmNicDeleteAudit	Update	Deleted NIC of VM
VmDiskAddAudit	Update	Added disk to VM
VmDiskDeleteAudit	Update	Deleted disk of VM
VmPowerOnAudit	Power State Change	Powered on VM
VmPowerOffAudit	Power State Change	Powered off VM
VmResetAudit	Power State Change	Reset VM
VmPowerCycleAudit	Power State Change	Power cycled VM
VmPauseAudit	Power State Change	Paused VM
VmSuspendAudit	Power State Change	Suspended VM
VmResumeAudit	Power State Change	Resumed VM
VmGuestShutdownAudit	Power State Change	Guest shutdown on VM
VmGuestRebootAudit	Power State Change	Guest rebooted on VM

Event Type	Operation	Description
VmAcpiShutdownAudit	Power State Change	ACPI shutdown VM
VmAcpiRebootAudit	Power State Change	ACPI rebooted VM
VmPowerStateReconcileAudit	Power State Change	Reconciled VM power state
HAFailoverStartAudit	HA	Started HA failover on host
HAFailoverCompleteAudit	HA	Completed HA failover on host
EnteringMaintenanceModeAudit	Update	Entering maintenance mode on host
MaintenanceModeAudit	Update	Host in maintenance mode
ExitMaintenanceModeAudit	Update	Exited maintenance mode on host
VmHostAffinityUpdateAudit	Update	Updated host affinity of VM
VmDiskResizeAudit	Update	Resized disk of VM
VmDiskUpdateAudit	Update	Updated disk of VM
VmUpdateAudit	Update	Updated VM
VmExportAudit	Create	VM exported to OVA
OvaCreateAudit	Create	OVA created
OvaDeleteAudit	Delete	OVA deleted
OvaUpdateAudit	Update	OVA name updated
OvaConcatenateAudit	Update	OVA parts concatenated
VmMigrateAudit	Migrate	Migrated VM from one host to another host
VsCreateAudit	Create	Virtual switch created
VsUpdateAudit	Update	Virtual switch updated
VsDeleteAudit	Delete	Virtual switch deleted
MigrateBridgeToVirtualSwitchAudit	Migrate	Migrated bridge to virtual switch
IPAssignmentToVMAudit	Update	IP assigned to VM
IPUnassignmentToVMAudit	Update	IP unassigned from VM
SpanCreateAudit	Create	SPAN session created
SpanUpdateAudit	Update	SPAN session updated
SpanDeleteAudit	Delete	SPAN session deleted
VmHostAffinityPolicyCreateAudit	Create	VM Host Affinity Policy created
VmHostAffinityPolicyUpdateAudit	Update	VM Host Affinity Policy updated
VmHostAffinityPolicyDeleteAudit	Delete	VM Host Affinity Policy deleted

Event Type	Operation	Description
VmHostAffinityPolicyReEnforceAudit	Update	VM Host Affinity Policy Re-enforced
ImageCreateAudit	Create	Created Image
ImageDeleteAudit	Delete	Deleted Image
ImageUpdateAudit	Update	Updated Image

Alerts/Health checks

Cluster

Table 133: Cluster In Read-Only Mode [1058] [A1195]

Name	Cluster In Read-Only Mode
Description	Single-node cluster is in read-only mode.
Alert message	`Single node cluster with CVM service_vm_external_ip is running in read-only mode due to reason.`
Cause	One of the metadata SSDs has failed.
Impact	Native backups are disabled.
Resolution	Replace the failed SSD. Refer to the Nutanix documentation for instructions.
KB Article	8156
Severity	Warning

Table 134: CVM Time Difference High [3025] [A1017]

Name	Time Drift Check
Description	Checks whether CVM times match.
Alert message	`Time differs between Controller VMs.`
Cause	Timedrift between CVMs is high. NTP connectivity may be lost.
Impact	The cluster may experience downtime.
Resolution	Check NTP/network configuration or connectivity.
Time Drift Threshold	Warning Threshold: 3 seconds
KB Article	1512
Severity	Warning

Table 135: Upgrade bundle available [3031] [A3031]

Name	Upgrade Bundle Available

Description	Bundle available for upgrade.
Alert message	`Bundle bundle_name is available for upgrade_type upgrade.`
Cause	Newer bundle available for upgrade
Impact	Cluster is running on older version.
Resolution	Upgrade the cluster with newer bundle. Refer to the Nutanix documentation for instructions.
Severity	Info

Table 136: Cluster services are down [3032] [A3032]

Name	Cluster Services Down Check
Description	Check that no services are down in the cluster.
Alert message	`Cluster Service: service_name is down on the Controller VM ip_address.`
Cause	One or more services in the cluster are down.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to serve I/O requests.
Resolution	[If the reason for the service(s) being down is unknown, run "ncc health_checks run_all", perform log collection covering the period of this alert timestamp, and provide the results to Nutanix Support via case attachment for further investigation. If due to planned maintenance or triage, please ensure the service is started again using "cluster start" as soon as possible. Report any reoccurring issues.]
KB Article	1518
Severity	Critical

Table 137: Alert Manager Service Check [3033]

Name	Alert Manager Service Check
Description	Check if Alert Manager service is available.
Cause	This check usually indicates that the Alert Manager service is not healthy, but there could be other causes.
Impact	Cluster issues may stay unnoticed.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.
KB Article	3377

Table 138: Cluster Service Restarting Frequently [3034] [A3034]

Name	Cluster Services Status

Description	Check if services have restarted recently across the cluster.
Alert message	`There have been multiple service restarts of services across all Controller VM(s). Latest crash of these services have occurred at timestamps respectively.`
Cause	This alert indicates that one or more services in the cluster were restarted.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.
KB Article	3378
Severity	Warning

Table 139: Multiple Vcenters discovered [6415] [A6415]

Name	Multiple vCenter Servers Discovered
Description	Multiple vCenter servers discovered
Alert message	`Discovered more than one vCenter servers : 'vcenter_ips'`
Cause	Nodes in the current Nutanix cluster are managed by multiple vCenters.
Impact	You will not be able to perform VM Management related operations.
Resolution	(A) Make sure all the nodes of this Nutanix cluster are moved to any one vCenter. (B) Contact Nutanix support.
Severity	Warning

Table 140: Common port group between ESXi hosts is absent [6416] [A6416]

Name	No common port group among ESXi hosts
Description	Common port group between ESXi hosts is not present
Alert message	`'message'`
Cause	ESXi hosts does not have any common port group.
Impact	You will not be able to perform network related operations.
Resolution	Ensure that all the nodes of this cluster have at least one common port group or all the nodes are in a single datacenter. If you are still not able to resolve the issue, contact Nutanix support.
Severity	Warning

Table 141: Cassandra service status check [21009]

Name	Cassandra service status check
Description	Check that all Cassandra nodes are online and and functioning normally

Cause	Cassandra service is not running in one or more hosts.
Impact	Cluster resilience might degrade.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.
KB Article	1547

Table 142: Cassandra service is running out of memory [21011] [A21011]

Name	Cassandra memory usage
Description	Check if the Cassandra service is running out of memory.
Alert message	`Cassandra service is running out of memory. Contact Nutanix support for assistance.`
Cause	Cassandra has crossed memory threshold more than 5 times in the last 3 minutes.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.
KB Article	1576
Severity	Critical

Table 143: Cassandra Waiting For Disk Replacement [21014] [A21014]

Name	Cassandra Waiting For Disk Replacement
Description	Metadata disk not replaced for Disk Replace Op
Alert message	`Disk disk_id for node service_vm_external_ip not replaced. Node performing disk replace: ip_address_doing_ring_change`
Cause	Disk not replaced for Disk Replace Op.
Impact	Disk Replacement operation will not proceed to completion.
Resolution	Ensure that a replacement disk is present on the node undergoing disk replace. If it is and the alert persists, contact Nutanix support.
Severity	Critical

Table 144: Domain Fault Tolerance is low for Metadata [21016] [A21016]

Name	Cassandra Domain Fault Tolerance Warning
Description	Domain fault tolerance is low for metadata

Alert message	`domain` fault tolerance is reduced from desired fault tolerance(`desired_fault_tolerance`) to `current_fault_tolerance` for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Cassandra-domain-aware fixer has been triggered to re-arrange metadata and improve the domain fault tolerance. No action is required from your side.
Severity	Info

Table 145: Domain Fault Tolerance is low for Metadata [21017] [A21017]

Name	Cassandra Domain Fault Tolerance Error
Description	Domain fault tolerance is low for metadata
Alert message	`domain` fault tolerance is reduced from desired fault tolerance(`desired_fault_tolerance`) to `current_fault_tolerance` for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Please run Cassandra-domain-aware fixer operation using the nCLI command 'cluster cassandra-domain-aware-fixer-start'.
Severity	Warning

Table 146: Domain Fault Tolerance is low for Metadata [21018] [A21018]

Name	Cassandra Domain Fault Tolerance Critical
Description	Domain fault tolerance is low for metadata
Alert message	`domain` fault tolerance is reduced from desired fault tolerance(`desired_fault_tolerance`) to `current_fault_tolerance` for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Please add more nodes to reduce the skew in number of nodes in each domain. It is preferred to maintain almost equal number of nodes in all the domains to achieve best domain fault tolerance.
Severity	Warning

Table 147: Node detached from metadata ring [21019] [A21019]

Name	Metadata Drive Detached With Fixer Check
Description	Node detached from metadata ring

Alert message	`Metadata drive on CVM ip_address is now detached from ring due to reason.`
Cause	Node was removed from the metadata store as part of the ring fixer operation.
Impact	Cluster performance may be degraded till node gets added back.
Resolution	No action required. Node will get added back to the metadata store automatically.
Severity	Info

Table 148: Metadata service restarting frequently due to long GC pauses. [21022] [A21022]

Name	Cassandra Restarting Due to Long GC Pauses
Description	Metadata service restarting frequently due to long GC pauses.
Alert message	`Metadata service on node ip_address is restarting frequently due to long garbage collection pauses.`
Cause	Metadata Service is undergoing long garbage collection pauses. This could possibly be due to cluster overload.
Impact	Cluster performance may be significantly degraded.
Resolution	Engage Nutanix Support to investigate why the Metadata service is significantly overloaded.
Severity	Warning

Table 149: Cluster In Override Mode [101057] [A101057]

Name	Cluster In Override Mode
Description	Single-node cluster is in override mode.
Alert message	`Single node cluster with CVM service_vm_external_ip is running in override mode due to reason.`
Cause	One of the metadata SSDs has failed.
Impact	Cluster is not disk fault tolerant.
Resolution	Replace the failed SSD. Refer to the Nutanix documentation for instructions.
KB Article	8132
Severity	Warning

Table 150: Reset Current Passwords for admin and nutanix SSH accounts on {cluster_name} [101081] [A101081]

Name	Password Strength Status
Description	High-strength password policy has been enabled, please reset the password.

Alert message	`High-strength password was recently enabled, please reset current passwords for admin and nutanix SSH accounts on cluster_name - cluster_uuid.`
Cause	High-strength password policy has been enabled.
Impact	High-strength password policy has been enabled but low-strength passwords may still be in use on the cluster.
Resolution	Please reset the current passwords. Follow the Security Guide's Hardening Instructions on Nutanix Support Portal for the password requirements.
Severity	Warning

Table 151: Same timezone check [103085] [A103085]

Name	CVM same timezone check
Description	Check that all CVMs are in the same timezone
Alert message	`vm_types are not in the same timezone.`
Cause	Some of the Controller VMs in the cluster are not in the same timezone.
Impact	It is recommended to have all the CVMs in the same timezone to avoid potential issues.
Resolution	Make sure all Controller VMs are in the same time zone. If they are not, then set the timezone using KB 1050
KB Article	2044
Severity	Warning

Table 152: Cluster NCC version check [103086]

Name	Cluster NCC version check
Description	Check that NCC version is consistent on the cluster
Cause	NCC version is different across the CVMs.
Impact	Some NCC checks can not be executed on all CVMs.
Resolution	Upgrade NCC to the desired version to keep it consistent across CVMs.
KB Article	3640

Table 153: CVM IP address is unreachable [103088] [A103088]

Name	Inter-CVM connectivity check
Description	Check that all CVMs are reachable via ping.
Alert message	`CVM dest_ip is unreachable`
Cause	The Controller VM is down or there is a network connectivity issue.
Impact	Storage may be unavailable and/or performance issues may be observed.

Resolution	If the Controller VM does not respond to ping, turn it on. Ensure that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	1635
Severity	Warning

Table 154: Host same timezone check [103092]

Name	Host same timezone check
Description	Check that all hosts are in the same timezone
Cause	Not all hosts in the cluster are in the same timezone.
Impact	Ensure all hosts in the cluster are in the same timezone to help ensure all guest VMs share the same timezone and log messages are timestamped consistently.
Resolution	Ensure that all hosts are in the same timezone by following the instructions in KB 1050
KB Article	2044

Table 155: Checks if AOS version of the target site configured for synchronous replication is on AOS version >= 5.17 [110021] [A110021]

Name	Remote AOS Version Check
Description	Checks if the target site configured is on AOS version >= 5.17 to support Synchronous Replication
Alert message	`For protection rule: protection_rule_name, target cluster(s): remote_site_list have version older than 5.17`
Cause	Target site is on AOS version < 5.17
Impact	Synchronous Replication will not be supported.
Resolution	Upgrade the cluster on the target site to AOS version >= 5.17. It is recommended to have both the source and target sites at the same version.
KB Article	7364
Severity	Warning

Table 156: Round-trip travel (RTT) to target Prism Element(s) is greater than 5ms, the maximum value allowed for synchronous replications. [110022] [A110022]

Name	Remote Site Latency Check
Description	Check if the Round-trip travel (RTT) is less than 5ms, the maximum supported for Synchronous Replication
Alert message	`Round-trip travel (RTT) to target Prism Element(s) err_msg is greater than max_permitted_latency`

Cause	Target cluster is unreachable or Round-trip travel (RTT) to target cluster is greater than 5ms.
Impact	Synchronous Replication will be affected.
Resolution	Ensure that the target cluster is reachable and the RTT is less than 5ms or choose another target cluster.
KB Article	7366
Severity	Warning

Table 157: NearSync Replication of Protection Domain has not progressed. [110264] [A110264]

Name	Nearsync Replication Stuck
Description	Nearsync replication of the protection domain has not progressed.
Alert message	`NearSync replication of the protection domain protection_domain_name to the remote site remote_name for the snapshot snapshot_id has not made any progress in the last time_string.`
Cause	Network outage, saturated network or unavailability of services like Stargate, Cerebro on the remote site. Snapshots of the protection domain will not be available on the remote site.
Cause	SSD snapshot reserve on the remote site does not have enough free disk space.
Impact	Protection domain may transition out of NearSync.
Impact	Snapshots of the protection domain will not be available on the remote site.
Resolution	Check the network connection and status of the remote site.
Resolution	Ensure that the remote site has enough disk space available in the SSD snapshot reserve.
KB Article	10249
Severity	Warning

Table 158: Stale secondary synchronous replication configuration found. [110456] [A110456]

Name	Stale secondary Synchronous Replication configuration check
Description	Check if there are any stale secondary synchronous replication configurations present on Prism Central.
Alert message	`Stale secondary synchronous replication configuration found for vms err_msg.`
Cause	Stale secondary synchronous replication configuration left behind.
Impact	Stale entries could cause unwanted tasks to be generated and they cannot be cleaned up automatically on the secondary as it could affect the VM stretch state hence require a manual cleanup.
Resolution	Cleanup the stale secondary synchronous replication configuration left behind on Prism Central.

KB Article	9716
Severity	Warning

Table 159: Check if a Recovery Plan with Witness configured has asynchronously protected VMs or has VMs with Synchronous Replication paused [110458] [A110458]

Name	Witness configured Recovery Plan has asynchronously protected VMs and VMs with Synchronous Replication paused check
Description	Check if a Recovery Plan with Witness configured has asynchronously protected VMs or has VMs with Synchronous Replication paused
Alert message	`The Witness configured Recovery Plan 'recovery_plan_name' has asynchronously protected VMs or has VMs with Synchronous Replication paused`
Cause	Recovery Plan with Witness configured has asynchronously protected VMs or has VMs with Synchronous Replication paused
Impact	The entities will not be managed by Witness
Resolution	Ensure all the VMs that are part of the Witness configured Recovery Plan are Synchronously protected
KB Article	11000
Severity	Warning

Table 160: Stale primary synchronous replication parameters found. [110460] [A110460]

Name	Stale primary Synchronous Replication configuration check
Description	Check if there are any stale primary synchronous replication configurations present on Prism Central.
Alert message	`Stale primary synchronous replication parameters found for vms err_msg.`
Cause	Stale primary synchronous replication configuration left behind probably because of a restore operation.
Impact	Stale configurations could cause unwanted tasks to be generated and they cannot be cleaned up automatically on the primary as it could affect the VM stretch state hence require a manual cleanup.
Resolution	Delete the stale VM(s) on primary site.
KB Article	11002
Severity	Warning

Table 161: Cassandra nodes up [111002]

Name	Cassandra nodes up
Description	Check that all Cassandra nodes are running

Cause	One or more Cassandra nodes are down which may be due to several issues including network issue, disk related issue, CPU or resource crunch.
Impact	Meta data load is inconsistent.
Resolution	Cassandra will repair itself. If other Cassandra checks are failing, contact Nutanix support.
KB Article	1369

Table 162: Cassandra tokens consistent [111003]

Name	Cassandra tokens consistent
Description	Check that all Cassandra nodes are consistent
Cause	The tokens may be inconsistent because of manual intervention or because of bad configuration.
Impact	Meta data load is inconsistent.
Resolution	Follow instructions described in KB 1369.
KB Article	1369

Table 163: Name server configuration [111022]

Name	Name server configuration
Description	Check for DNS Server configuration
Cause	Name server is not configured on CVMs and hypervisor hosts or is not able to resolve queries.
Impact	Name resolution will not work.
Resolution	Verify if Name server is configured on CVMs and on hypervisor hosts and if it can resolves queries. Review KB 3005 for more details.
KB Article	3005

Table 164: HTTP proxy check [111023]

Name	HTTP proxy check
Description	Check if HTTP proxy is working
Cause	(A) No proxy configured. (B) Unable to connect to proxy on the port.
Impact	Cluster may not have access to Internet-based service such as 1-click auto and manual downloads, Pulse, LCM inventory, etc.
Resolution	Check the proxy server settings (network, port, user credentials, etc.)
KB Article	1708

Table 165: LDAP configuration [111024]

Name	LDAP configuration
Description	Check LDAP configuration.
Cause	LDAP not correctly configured in the cluster.
Impact	Directory users might not be able to log properly.
Resolution	Review KB 2997.
KB Article	2997

Table 166: Virtual IP Check [111027] [A111027]

Name	Virtual IP check
Description	Check if virtual IP is configured and reachable.
Alert message	*vm_type* **Virtual IP is configured but unreachable.**
Cause	Cluster virtual IP is not configured.
Cause	Cluster services are down or cluster is not started yet.
Impact	Nutanix features that use virtual IP address might be adversely affected.
Impact	
Resolution	Configure a valid virtual IP for the cluster.
Resolution	Verify that all cluster services are up.
KB Article	3708
Severity	Critical

Table 167: Advanced Vector Extensions (AVX) enabled [111035]

Name	Advanced Vector Extensions (AVX) enabled
Description	Checks AVX enabled on CPUs.
Cause	Intel's Advanced Vector Extensions (AVX) is disabled on the node.
Impact	If AVX is disabled, upgrade will fail.
Resolution	Enable AVX on the node. Review KB 8148.
KB Article	8148

Table 168: Prism Central using default password [111050] [A111050]

Name	PC using default password
Description	PC is using default password

Alert message	`Please update the default password on PC for user(s): users to harden the security on the clusters. Nutanix highly recommends not using default passwords for clusters.`
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration - Using default passwords could expose your system to unauthorized access.
Resolution	Ensure that Prism Central is using a custom password for improved security.
KB Article	6153
Severity	Critical

Table 169: Notifications dropped check [111052]

Name	Notifications dropped check
Description	Some alert notifications were dropped
Cause	One or more cluster services are not healthy.
Impact	Some of the alerts will not be raised. Cluster issues may stay unnoticed.
Resolution	Contact Nutanix support.
KB Article	4905

Table 170: Alert E-mail Failure [111066] [A111066]

Name	Alert E-mail Failure
Description	Failed to send alert Emails
Alert message	`Unable to send alert Email from the Controller VM service_vm_external_ip due to reason.`
Cause	Issue with the Email client
Impact	Alerts notification Emails will not be sent.
Resolution	Contact Nutanix Support.
KB Article	9937
Severity	Critical

Table 171: Failure to copy Image to Cluster [111072] [A111072]

Name	Image copy to Cluster check
Description	Detect if image copy to cluster is failing
Alert message	`Image could not be copied to Cluster`
Cause	Cluster may be unreachable from PC or Catalog service may be down on Cluster.
Impact	Image will not be available on the Cluster.

Resolution	Check Catalog service on cluster and connectivity to PC.
Severity	Warning

Table 172: SW Encryption Keys from {kms_name} are unavailable [111075] [A111075]

Name	SW Encryption Keys Unavailable
Description	Check to see if the KMS actually has the expected passwords for the container encryption keys.
Alert message	`SW encryption keys for one or more containers are not available from external key manager kms_name.`
Cause	If a 3rd party KMS gets restored from a backup or other incorrect management operation, it may have missing keys.
Impact	Cluster will have data unavailability if rebooted.
Resolution	Contact Nutanix and 3rd party key manager support.
KB Article	8223
Severity	Critical

Table 173: High number of tasks in the cluster [111082]

Name	High number of tasks in the cluster
Description	High number of tasks in the cluster
Cause	High number of tasks in the cluster
Impact	Operation on entities in the cluster might not progress.
Resolution	Contact Nutanix support for help.
KB Article	9866

Table 174: Cannot Remove Password Protected Disk(s) [130003] [A1106]

Name	Password Protected Disk Status
Description	Cannot remove password protected disk(s).
Alert message	`Cannot remove password protected disks disk_id_list from the configuration as they cannot be cleaned.`
Cause	The drives are offline, the drives are removed, or the node is down.
Impact	The disks cannot be removed from the configuration.
Resolution	Ensure the disks are accessible.
Severity	Warning

Table 175: Automatic Addition Of Node To Metadata Ring Disabled [130004] [A1079]

Name	Node Auto Add Check
Description	Automatic addition of the node to the metadata ring is disabled.
Alert message	`Automatic addition of node to metadata ring has been disabled on CVM service_vm_id with IP address ip_address.`
Cause	This node has been removed from the metadata store after being automatically added recently. Automatic addition of this node to the metadata store has now been disabled.
Impact	Cluster performance may be significantly degraded.
Impact	In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Automatic addition will be re-enabled when the node is manually added to the metadata store. If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. Otherwise, contact Nutanix support.
Severity	Warning

Table 176: Node Detached From Metadata Ring [130005] [A1055]

Name	Node Detached From Ring Check
Description	The metadata drive is detached from the ring.
Alert message	`Node ip_address is now detached from metadata ring due to reason.`
Cause	Either a metadata drive has failed, the node was down for an extended period of time, or an unexpected subsystem fault was encountered, so the node was removed from the metadata store.
Impact	Cluster performance may be significantly degraded.
Impact	In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	(A) If the metadata drive has failed, replace the metadata drive as soon as possible. Refer to the Nutanix documentation for instructions. (B) If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. For assistance, contact Nutanix support.
Severity	Warning

Table 177: Metadata Dynamic Ring Change Operation Stuck [130006] [A1117]

Name	Metadata DynRingChangeOp Status
Description	Dynamic Ring Change operation is not making progress.
Alert message	`operation for node ip_address is not progressing. Node performing the operation: ip_address_doing_ring_change.`

Cause	Node addition or removal is not making progress. A node in the cluster is unhealthy.
Impact	Node addition and node removal operations will not proceed to completion.
Resolution	If the alert persists, contact Nutanix Support.
KB Article	8195
Severity	Critical

Table 178: Metadata Dynamic Ring Change Operation Too Slow [130007] [A1116]

Name	Metadata DynRingChangeOp Slow Check
Description	Dynamic Ring Change operation too slow.
Alert message	`operation` for node `ip_address` is progressing slowly. Total elapsed time: `elapsed_time_mins` min(s). Node performing the operation: `ip_address_doing_ring_change`.
Cause	A node in the cluster is unhealthy, or the cluster contains an extraordinarily large amount of data.
Impact	Cluster performance may be degraded.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.
Severity	Warning

Table 179: Node Forwarding Metadata Requests [130008] [A1037]

Name	Node In Forwarding Mode Check
Description	Node Forwarding Metadata Requests
Alert message	Metadata service on node `ip_address` is running in forwarding mode due to `reason`.
Cause	Either a metadata drive has failed, node removal has been initiated, or an unexpected subsystem fault has been encountered.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
KB Article	8266
Severity	Critical

Table 180: Large Metadata Size Detected [130009] [A1119]

Name	Metadata Size

Description	Large metadata size detected.
Alert message	`Node ip_address contains a large amount of metadata.`
Cause	A node in the cluster contains a large amount of metadata and has exceeded thresholds.
Impact	Cluster performance may be degraded.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.
Severity	Warning

Table 181: Node Marked To Be Auto Added To Metadata Ring [130010] [A1078]

Name	Node Marked To Be Auto Added Check
Description	Node Marked To Be Auto Added To Metadata Ring
Alert message	`Node service_vm_id with IP address ip_address is marked to be added to metadata ring. reason.`
Cause	Node was detached from metadata ring earlier. It is now stable and will be added back to the metadata store.
Impact	The metadata service will start to use the node.
Resolution	No action is necessary.
Severity	Info

Table 182: Node Marked To Be Detached From Metadata Ring [130011] [A1054]

Name	Node Marked To Be Detached Check
Description	The node is marked to be detached from the metadata ring.
Alert message	`Node ip_address is marked to be detached from metadata ring due to reason.`
Cause	Either a metadata drive has failed, the node was down for an extended period of time, or an unexpected subsystem fault was encountered, so the node is marked to be removed from the metadata store.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	(A) If the metadata drive has failed, replace the metadata drive as soon as possible. Refer to the Nutanix documentation for instructions. (B) If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. For assistance, contact Nutanix support.
KB Article	8408
Severity	Critical

Table 183: Metadata Ring Imbalance [130012] [A1072]

Name	Metadata Imbalance Check
Description	Cassandra metadata imbalance.
Alert message	`Metadata ring imbalance for num_partitions partitions. The partitions are partition_list.`
Cause	One or more nodes have a disproportionately larger token range size. This imbalance may cause performance bottlenecks on the node(s) affected.
Impact	Node performance may be significantly degraded.
Resolution	Execute cassandra ring skew fix operation using the NCLI command 'cluster cassandra-token-range-skew-fixer-start'. UVMs' performance might be affected for a small amount of time.
Severity	Warning

Table 184: Curator Job Running Too Long [130015] [A1120]

Name	Curator Job Status
Description	The Curator job has been running for too long.
Alert message	`Curator job name with id execution_id has been running for a long time i.e. elapsed_time_secs seconds.`
Cause	Various
Impact	Background cluster maintenance tasks might get affected in the future. The root cause should be addressed as soon as possible.
Resolution	Contact Nutanix support if this issue persists.
Severity	Info

Table 185: Curator Scan Failure [130016] [A1081]

Name	Curator Scan Status
Description	Curator scan failure.
Alert message	`Curator scan_type scans have repeatedly failed to complete.`
Cause	Various
Impact	Background cluster maintenance tasks are not occurring. While cluster operation will not be immediately affected, the root cause should be addressed as soon as possible.
Resolution	Contact Nutanix support.
KB Article	3786
Severity	Critical

Table 186: Duplicate IP Address Detected [130017] [A1096]

Name	IP Configuration
Description	Duplicate IP address is detected.
Alert message	`Detected IP address conflict for cluster virtual IP address duplicate_ip. IP found on NIC with MAC address mac_address.`
Cause	The cluster virtual IP address may be configured on another host.
Impact	The cluster is not accessible through the virtual IP address. Other services that use the virtual IP address, such as remote replication or SCCM, may be affected.
Resolution	Either configure a different cluster virtual IP address or unconfigure the IP address from the other host.
Severity	Warning

Table 187: Duplicate Remote Cluster ID [130018] [A1038]

Name	Duplicate Remote Cluster ID Check
Description	Duplicate remote cluster ID.
Alert message	`Remote site already exists with the name 'conflicting_remote_name'.`
Cause	Two remote sites with different names or different IP addresses have same cluster ID. This can happen in two cases: (a) A remote cluster is added twice under two different names (through different IP addresses) or (b) Two clusters have the same cluster ID.
Impact	Protected data is not replicated to the remote site.
Resolution	In case (a) remove the duplicate remote site. In case (b) verify that the both clusters have the same cluster ID and contact Nutanix support.
Severity	Warning

Table 188: IP Address Not Hosted [130023] [A1097]

Name	Virtual IP Configuration
Description	IP address is not hosted.
Alert message	`Unable to host virtual IP address ip_address.`
Cause	The network configuration is incorrect.
Impact	The cluster is not accessible through the virtual IP address. Other services that use the virtual IP address, such as remote replication or SCCM, may be affected.
Resolution	Contact Nutanix support.
Severity	Critical

Table 189: SMTP Error [130060] [A1080]

Name	SMTP Configuration
Description	Failed to send Pulse and/or alert emails to Nutanix support with the current configuration.
Alert message	`Failed to send email from Controller VM ip_address due to error: error.`
Cause	SMTP server may not be configured.
Cause	The SMTP server is not reachable from the Controller VM.
Cause	Firewall is preventing emails from being forwarded through the SMTP server.
Impact	Nutanix support cannot proactively monitor cluster health and offer assistance before issues become acute.
Resolution	Configure the SMTP server.
Resolution	SMTP server should be reachable.
Resolution	Ensure that the firewalls allow the cluster to send email messages through the Nutanix SMTP servers (nsc01.nutanix.net and nsc02.nutanix.net) on port 80.
KB Article	1638
Severity	Warning

Table 190: Remote Site Latency Is High. [130075] [A1168]

Name	RemoteSiteHighLatency
Description	Latency to a Metro remote site is high.
Alert message	`Latency to remote site 'remote_name' is more than 'remote_site_high_latency_threshold_ms' ms , for 'sustained_high_latency_secs' seconds due to reason.`
Cause	Network issues or CVM resource constraints
Impact	Replication will fail.
Resolution	Check if remote site is reachable and resolve any network issues. Monitor CVM CPU and memory usage
Severity	Warning

Table 191: Possible degraded Node [130087] [A130087]

Name	Node Degradation Status
Description	Services on one of the nodes in the cluster are in a degraded state and possibly not making progress.
Alert message	`Possible degraded node ip_address with Controller VM ID service_vm_id reported by component component_name.`
Cause	Various
Impact	Cluster performance may be significantly degraded.

Resolution	Contact Nutanix support.
KB Article	3827
Severity	Critical

Table 192: CPS Deployment Evaluation Mode [130093] [A130093]

Name	CPS Deployment Evaluation Mode
Description	Checks if the CPS deployment is running in evaluation mode.
Alert message	`CPS deployment of cluster with cluster id cluster_id is running in evaluation mode. Please apply a license by following the CPS deployment guide.`
Cause	CPS is running in evaluation mode.
Impact	The deployment cannot be in evaluation and the user has to apply a license.
Resolution	Apply a license to the deployment.
Severity	Warning

Table 193: Metro connectivity lost [130119] [A130119]

Name	Metro Connectivity Lost
Description	Metro connectivity is lost
Alert message	`Metro connectivity from local site 'local_stargate_handler' to remote site 'remote_name' is lost due to 'reason'`
Cause	Remote site(s) not reachable
Impact	Metro connection of the remote site is unavailable
Resolution	Check remote site(s) connectivity and re-enable Metro if necessary
Severity	Warning

Table 194: Alert raised on cloud remote site {remote_name}: {alert_message} [130135] [A130135]

Name	Cloud Remote Alert
Description	Alert generated on cloud remote site.
Alert message	`Alert raised on cloud remote site remote_name: alert_message`
Cause	Various
Impact	Various
Resolution	Resolve the issue stated in the alert. If you cannot resolve the issue, please contact Nutanix support.
Severity	Warning

Table 195: Snapshot Reserve on SSD is Full. [130139] [A130139]

Name	LWS Store Full
Description	SSD reserved for the Nearsync snapshots is full.
Alert message	`The SSD space that is reserved for the snapshots is full.`
Cause	Snapshot reserve on the SSD is full or has high IO workload.
Impact	RPO cannot be met and will be temporarily transitioned to 60 minutes.
Resolution	Reduce the retention time of the snapshots.
Severity	Warning

Table 196: Two node cluster state changed [130151] [A130151]

Name	Change in state of two node cluster
Description	Change in state of two node cluster
Alert message	`Two node cluster state changed: message`
Cause	There is a change in availability of a node (or network between the nodes).
Impact	If the cluster is now in stand-alone mode, it cannot handle the remaining nodes failure.
Resolution	Two node cluster will auto-heal when node (or network between the nodes) is available again.
KB Article	8284
Severity	Critical

Table 197: Two node cluster changed state to stand-alone mode. [130152] [A130152]

Name	Two node cluster changed state to stand-alone mode
Description	One of the nodes in a two node cluster is currently unavailable, the cluster is now operating in stand-alone mode.
Alert message	`Two node cluster state change: message`
Cause	There was an unplanned event that changed the availability of a node (or network between the nodes).
Impact	The cluster is not node fault tolerant until the cluster auto-heals.
Resolution	Two node cluster will auto-heal when node (or network between the nodes) is available again.
KB Article	8285
Severity	Warning

Table 198: Two node cluster state change to stand-alone mode [130153] [A130153]

Name	Two node cluster changed now in stand-alone mode
Description	Change in state of two node cluster
Alert message	`Two node cluster state change: message`
Cause	There is a change in availability of a node (or network between the nodes). This event was planned.
Impact	The cluster is not node fault tolerant until the cluster auto-heals.
Resolution	Two node cluster will auto-heal when the planned event is completed successfully.
Severity	Warning

Table 199: State changed for two node cluster [130154] [A130154]

Name	Two node cluster changed state back to a two-node mode.
Description	Change in state of two node cluster
Alert message	`Two node cluster state changed: message`
Cause	A previously unavailable node or service is now available again.
Impact	No impact.
Resolution	No resolution necessary.
Severity	Info

Table 200: Witness is unreachable from node. [130155] [A130155]

Name	Witness VM is unreachable from a node in two node cluster
Description	Unable to reach the Witness VM over the network, the cluster may not be able to handle a node failure.
Alert message	`The node cannot issue witness APIs to witness: message`
Cause	Node is unable to issue APIs to witness VM.
Impact	The cluster may not be able to handle node failure until the witness is reachable again.
Resolution	Check the witness VM to verify it is up and verify the cluster nodes can communication with the Witness VM over the network.
Severity	Warning

Table 201: Drive Removal Stuck [130156] [A130156]

Name	Drive Removal Stuck
Description	Drive Removal Stuck

Alert message	`Removal of drive disk_id is stuck.`
Cause	Various
Impact	Drive cannot be removed from the cluster.
Resolution	Contact Nutanix support.
Severity	Warning

Table 202: Unable to retrieve the endpoint information for the Availability Zone. [130173] [A130173]

Name	Unable to get Availability Zone Endpoint
Description	Unable to retrieve the endpoint information for the Availability Zone
Alert message	`Unable to retrieve the endpoint information for the Availability Zone availability_zone_physical_name`
Cause	Prism Central, Prism Element, or Xi - Availability Zone may be unreachable
Impact	VMs may not be fully protected
Resolution	Check if the Prism Central, Prism Element, or Xi - Availability Zone are reachable. If the problem persists, contact Nutanix support.
KB Article	8274
Severity	Critical

Table 203: Unable to communicate with the Data Center Manager [130175] [A130175]

Name	Failure to communicate with DCM
Description	Unable to retrieve the endpoint information for the Availability Zone from the Data Center Manager
Alert message	`Unable to retrieve the endpoint information for the Availability Zone availability_zone_name from the Data Center Manager`
Cause	Data Center Manager may be unreachable or has provided an invalid response
Impact	VMs may not be fully protected
Resolution	Check if Data Center Manager is reachable from the Xi - Availability Zone
KB Article	6901
Severity	Critical

Table 204: Data At Rest Encryption key backup warning [130184] [A130184]

Name	Data At Rest Encryption key backup warning.
Description	New encryption keys not backed up.
Alert message	`Data At Rest Encryption key backup warning: message`

Cause	New encryption keys have been created since the last backup was downloaded.
Impact	If the encryption keys are lost, the data in the storage container cannot be accessed.
Resolution	Download the latest copy of the cluster's key backup.
Severity	Warning

Table 205: Local key manager master key rotation warning. [130185] [A130185]

Name	Local key manager master key rotation warning.
Description	Unable to rotate local key manager's master key.
Alert message	`Local key manager master key rotation warning. Last failure reason: failure_message`
Cause	One or more nodes in the cluster are unavailable.
Impact	If the cluster experiences any more node failures, the data on the cluster may be inaccessible.
Resolution	Ensure Mantle service is UP on all nodes. If the warning persists, contact Nutanix support.
Severity	Warning

Table 206: Remote Availability Zone Latency is High [130187] [A130187]

Name	Latency to Remote Availability Zone is High
Description	Network latency to the remote Availability Zone is high.
Alert message	`Network latency to the remote Availability Zone 'availability_zone_physical_name' exceeds the threshold limit.`
Cause	Network latency to the remote Availability Zone exceeds the threshold limit.
Impact	Replications will fail
Resolution	Check if the remote Availability Zone is reachable over the network.
Severity	Warning

Table 207: Cluster Join to Domain Failure [130189] [A1154]

Name	ClusterJoinToDomainFailure
Description	Failed to join AOS cluster to domain.
Alert message	`Failed to join AOS cluster to domain domain_name: reason`
Cause	Cluster could not be joined to domain because of incorrect login credentials or other causes.
Impact	Workflows such as backup service (volume shadow copy service) and Kerberos for SMB will not work.

Resolution	Please ensure that the specified domain account credentials are valid and that the virtual computer object representing the AOS cluster does not already exist.
KB Article	8080
Severity	Critical

Table 208: Background Encryption Stuck [130194] [A130194]

Name	Background Encryption Stuck
Description	Background Encryption Stuck
Alert message	`Encryption of container container_id is not progressing. Contact Nutanix support for further assistance.`
Cause	Oplog entries, extent groups, NFS entries or other components are not being encrypted.
Impact	Encryption is not progressing.
Resolution	Contact Nutanix support.
Severity	Warning

Table 209: Failed to reserve host memory for Atlas networking [130202] [A130202]

Name	Failed to reserve host memory for Atlas networking
Description	Failed to reserve host memory for Atlas networking.
Alert message	`Failed to reserve host memory for Atlas networking`
Cause	Failed to reserve host memory for Atlas networking.
Impact	If hypervisor memory is exhausted, random VMs may be powered off to reclaim memory.
Resolution	Power off VMs to reduce host memory consumption.
KB Article	8525
Severity	Critical

Table 210: Possible Service Degradation. [130338] [A130338]

Name	Service instances receiving bad degraded node scores
Description	Service instances on a few nodes of the cluster are possibly in a bad state.
Alert message	`Possible degradation of service_name instances on the Controller VM IDs service_vm_id_list.`
Cause	Various
Impact	Cluster performance may be significantly degraded.
Resolution	Follow KB9390 to troubleshoot reasons for the service receiving bad scores.

KB Article	9390
Severity	Warning

Table 211: Cross-container disk migration task is paused [130340] [A130340]

Name	Cross-container disk migration task is paused
Description	Cross-container disk migration task is paused.
Alert message	`Migration of VM disk 'vdisk_name' (VM 'vm_name') from source storage container 'src_container_name' to destination storage container 'dst_container_name' is paused`
Cause	Cross-container disk migration task is paused due to lack of free space in the destination storage container.
Impact	Cross-container disk migration cannot be completed.
Resolution	Free up additional space in the destination storage container or abort the parent task.
KB Article	10248
Severity	Warning

Table 212: Zookeeper server unresponsive. [130344] [A130344]

Name	Zookeeper server unresponsive for a long time
Description	Zookeeper server might be unresponsive on a CVM for a long time.
Alert message	`Zookeeper server on the Controller VM ID service_vm_id is unresponsive since death_timestamp_secs.`
Cause	Zookeeper server on one of the nodes has been unresponsive for a long time.
Impact	Cluster performance may be significantly degraded.
Resolution	Follow KB1587 to troubleshoot reasons for Zookeeper server being unresponsive.
KB Article	1587
Severity	Warning

Table 213: Detected iSCSI discovery or login through cluster external IP address [130346] [A130346]

Name	Invalid iSCSI portal
Description	Detected iSCSI discovery or login through cluster external IP address
Alert message	`Detected iSCSI access_type through cluster external IP address`
Cause	The iSCSI discovery or login portal is incorrectly configured
Impact	The iSCSI service won't be highly available and the performance would be affected

Resolution	Use the cluster external data services IP address for iSCSI discovery
Severity	Warning

Table 214: MEM-NVMe Tiering Cluster Constraints Check Failure [130360] [A130360]

Name	MEM-NVMe tiering disabled
Description	The cluster requirements for enabling the MEM-NVMe storage tier feature are not met. MEM-NVMe storage tiering will be disabled.
Alert message	`Cluster fails to meet the MEM-NVMe tiering constraints due to reason: reason`
Cause	Failing to meet the MEM-NVMe tiering constraints results in MEM-NVMe tiering being disabled.
Impact	All the MEM-NVMe disks will be treated as SSD-PCIe disks.
Resolution	Configure the cluster to meet the MEM-NVMe tiering constraints. All the nodes in the cluster should have at least one MEM-NVMe disk and the cluster should not contain any {tiers_not_allowed} disk.
KB Article	11650
Severity	Warning

Table 215: Pulse cannot connect to REST server endpoint [140001] [A140001]

Name	REST endpoint connection status
Description	Check if Pulse can connect to REST server endpoint https://insights.nutanix.com:443.
Alert message	`Pulse cannot connect to REST server endpoint https://insights.nutanix.com:443. Connection Status: connection_status, Pulse Enabled: enable, Error Message: message`
Cause	Pulse cannot connect to REST server endpoint https://insights.nutanix.com:443.
Impact	Data driven serviceability and customer support cannot be performed.
Resolution	Ensure that the REST server endpoint https://insights.nutanix.com:443 is reachable from Pulse.
KB Article	5490
Severity	Warning

Table 216: Pulse is disabled [140002] [A140002]

Name	Pulse Disabled
Description	Pulse is disabled on the paired Prism Central.
Alert message	`Pulse is disabled on the paired Prism Central cluster_uuid`

Cause	Pulse enablement might be missed or disabled.
Impact	Monitoring Nutanix Service Health cannot be performed.
Resolution	Enable Pulse on current cluster and all of its registered clusters.
KB Article	6377
Severity	Info

Table 217: Metering server endpoint connection unavailable. [140003] [A140003]

Name	Metering endpoint connection status
Description	Checks if the connection to the metering server endpoint is available and if the metering stats can be successfully sent to server endpoint.
Alert message	`Cannot publish metering stats to server endpoint on a metering enabled cluster.`
Cause	Unable to verify metering enablement state and/or unable to establish a connection to the metering server endpoint.
Impact	Metering and billing stats cannot be collected from this cluster violating the terms of service.
Resolution	Ensure that the metering server endpoint is resolvable and reachable from cluster and that the network connection is stable.
KB Article	10651
Severity	Critical

Table 218: Pulse is not enabled on cluster. [140004] [A140004]

Name	Pulse enablement check
Description	Checks if the pulse is disabled on cluster.
Alert message	`Pulse is not enabled on cluster`
Cause	Pulse is not enabled on cluster.
Impact	Telemetry data cannot be sent thus impacting Nutanix's ability to provide proactive support
Resolution	Ensure that pulse is enabled on the cluster to allow Nutanix to provide proactive support.
KB Article	10743
Severity	Info

Table 219: OVS Service Restart [150002] [A150002]

Name	OVS Service Restarted
Description	OVS service has unexpectedly restarted on one or more AHV hosts.

Alert message	`Node uuid_node: detected unexpected restart of ovs-vswitchd process on host host_ip from old_pid old_pid to new_pid new_pid. OVS service configuration was restored after restart. Guest VMs might have experienced temporary network disconnect.`
Cause	OVS service was restarted.
Impact	Guest VMs may temporarily lose network connectivity.
Resolution	Bridges and flow rules on hosts have been rebuilt. No user action is required.
KB Article	6151
Severity	Warning

Table 220: Flow visualization statistics collector service restart detected [150003] [A150003]

Name	Flow visualization statistics collector service monitor
Description	Check if flow visualization statistics collector service was restarted.
Alert message	`Flow visualization statistics collector service was restarted max_restart_count times in last time_interval_min minutes on host host_ip.`
Cause	Flow visualization statistics collector service was restarted 10 times in the last 15 minutes.
Impact	Flow visualization may not show real time data.
Resolution	Review KB 8911.
KB Article	8911
Severity	Warning

Table 221: Inconsistent Virtual Switch State Detected [150004] [A150004]

Name	Inconsistent virtual switch state was detected
Description	Virtual switch state is inconsistent across nodes in the cluster.
Alert message	`Virtual switch vs_name is in an inconsistent state for the host_info. Reason: error_message`
Cause	A virtual switch is not present on at least one node in the cluster.
Impact	VMs using this virtual switch may lose network connectivity when migrated to the affected node.
Resolution	(A) Review virtual switch configuration and ensure every host (storage-only hosts are optional) in the cluster is included. (B) Ensure all the hosts and their physical NICs used in the virtual switch configuration are present.
Severity	Critical
Severity	Warning

Table 222: Default Virtual Switch Error [150005] [A150005]

Name	Default virtual switch error
Description	Default virtual switch has error in the cluster.
Alert message	`Default virtual switch cannot be created on the host_info. Reason: error_message`
Cause	Default virtual switch could not be created on a node in the cluster.
Impact	Networking operations may be impacted on the affected node(s) in the cluster.
Resolution	Check Acropolis service logs and refer to KB 9233 to identify the problem and steps to restore the default virtual switch in the cluster.
KB Article	9233
Severity	Critical

Table 223: IGMP Snooping needs to be re-configured on newly migrated virtual switch. [150008] [A150008]

Name	IGMP Snooping needs to be re-configured on newly migrated virtual switch
Description	Notifies when IGMP Snooping has been configured on an OVS bridge that's migrated to a virtual switch
Alert message	`IGMP Snooping was found to be configured on bridge bridge_name, when it was migrated to virtual switch vs_name with UUID vs_uuid. This configuration won't be carried over to the virtual switch vs_name, please re-configure it.`
Cause	IGMP Snooping was previously enabled on the bridge before its migration to a virtual switch.
Impact	Existing IGMP Snooping configuration on the OVS bridge is not reflected in the migrated virtual switch configuration. Any further update of virtual switch will reconfigure IGMP Snooping based on its configuration in virtual switch.
Resolution	Please re-configure IGMP Snooping on the bridge using the migrated virtual switch.
Severity	Warning

Table 224: Inconsistent SPAN Session State Detected [150009] [A150009]

Name	Inconsistent SPAN session state was detected
Description	SPAN session state is inconsistent on one of the nodes in the cluster.
Alert message	`SPAN session session_name is in an inconsistent state for the host_info. Reason: message`
Cause	A SPAN session is in an inconsistent state on one of the nodes in the cluster.
Impact	The SPAN session mirroring traffic from the host NIC(s) to a destination VM NIC is no longer active.

Resolution	Ensure that the source ports and/or destination ports for the SPAN session are connected and in active state.
KB Article	10440
Severity	Warning

Table 225: IPFIX exporter cannot be enabled on the cluster [150010] [A150010]

Name	IPFIX exporter cannot be enabled on the cluster
Description	Cannot enable IPFIX exporter on the cluster
Alert message	`IPFIX exporter exporter_name cannot be enabled on cluster with UUID cluster_uuid. Reason: cause`
Cause	IPFIX exporter could not be enabled on the cluster.
Impact	Network connection details cannot be exported via IPFIX from this cluster.
Resolution	Check the software version(s) of Prism Element and AHV and upgrade them to the minimum required version(s). Make sure the connection between Prism Central and Prism Element is healthy.
Severity	Warning

Table 226: Analytics VM one or more component failure. [160064] [A160064]

Name	Analytics VM one or more component failure.
Description	One or more components of the File Analytics VM {avm_ip} are not functioning properly or have failed.
Alert message	`One or more components of the File Analytics VM avm_ip are not functioning properly or have failed.`
Cause	One or more File Analytics VM components have failed.
Impact	File Server analytics will not perform appropriately.
Resolution	Contact Nutanix Support.
KB Article	8239
Severity	Critical

Table 227: Analytics VM high CPU usage. [160065] [A160065]

Name	Analytics VM high CPU usage.
Description	File Analytics VM {avm_ip} has a high CPU usage of {usage_percent}%
Alert message	`File Analytics VM avm_ip has a high CPU usage of usage_percent%`
Cause	File Analytics VM is performing too many CPU intensive operations.
Impact	File Analytics VM's performance will be impacted.

Resolution	Monitor which component has a high CPU usage and reduce the processing load.
Severity	Warning

Table 228: Analytics VM high disk usage. [160066] [A160066]

Name	Analytics VM high disk usage.
Description	File Analytics VM {avm_ip} has a high disk usage of {usage_percent}%
Alert message	`File Analytics VM avm_ip has a high disk usage of usage_percent%`
Cause	Extensive disk usage on the File Analytics VM
Impact	File Analytics VM's performance will be impacted.
Resolution	Add more capacity to the File Analytics VM or delete some data.
Severity	Warning

Table 229: Analytics VM low memory available. [160067] [A160067]

Name	Analytics VM low memory available.
Description	The memory available on File Analytics VM {avm_ip} is {memory_available} bytes which is low.
Alert message	`The percent memory available on File Analytics VM avm_ip is memory_available which is low`
Cause	Excessive memory usage on the File Analytics VM
Impact	File Analytics VM's performance will be impacted.
Resolution	Add more RAM
Severity	Warning

Table 230: File Server upgrade task stuck [160075] [A160075]

Name	File Server Upgrade Task Stuck
Description	File Server Upgrade Task Stuck.
Alert message	`Upgrade of file server file_server_name is stuck due to reason.`
Cause	Check alert message for details
Impact	File server has not been upgraded.
Resolution	Refer to KB article 8438 for more details. Contact Nutanix support if issue still persists or assistance needed. Once the issue is resolved, execute 'afs infra.resume_fs_upgrade' on cvm to resume Nutanix Files upgrade.
KB Article	8438
Severity	Critical

Table 231: Files cluster HA takeover failure [160076] [A160076]

Name	File server HA Takeover Failure
Description	Files cluster HA takeover failed
Alert message	`Files cluster file_server_name HA takeover has failed multiple times on FSVM fsvm_internal_ip due to message.`
Cause	Files cluster HA takeover process failed for multiple attempts
Impact	Files cluster may be unavailable.
Resolution	Refer to KB article 8419 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8419
Severity	Critical

Table 232: File Server Alert [160080] [A160080]

Name	File Server Alert
Description	Files server alert
Alert message	`file_server_alert_msg`
Cause	Check alert summary
Impact	Check alert summary
Resolution	Check alert summary
Severity	Critical
Severity	Warning
Severity	Info

Table 233: File Server Manager Upgrade Failed. [160085] [A160085]

Name	File Server Manager Upgrade Failed.
Description	The File Server Manager upgrade task failed on the cluster.
Alert message	`reason`
Cause	Please check the alert message for details.
Impact	The File Server Manager upgrade failed on the cluster; File Server operations may not work as expected.
Resolution	Please refer to KB 8342 at the Nutanix Portal.
Severity	Warning

Table 234: Cassandra status is not normal on few nodes of file server [160092] [A160092]

Name	File Server Cassandra Status Failure

Description	Cassandra status is not normal on few nodes of file server
Alert message	*message*
Cause	See details in the alert message.
Impact	Some file server features may not be functioning
Resolution	Contact Nutanix support.
Severity	Critical

Table 235: Inconsistent state found in File Server [160093] [A160093]

Name	File Server Invalid IDF State
Description	One or more inconsistent entities have been found. This may cause upgrades to fail.
Alert message	**One or more inconsistent entities have been found. This may cause upgrades to fail.**
Cause	IDF is in an inconsistent state.
Impact	File server upgrades may fail.
Resolution	Contact Nutanix support.
Severity	Critical

Table 236: Post upgrade handler failed [160094] [A160094]

Name	File Server Master One Time Upgrade Failure
Description	Post upgrade handler failed.
Alert message	*message*.
Cause	Master One Time Upgrade failed.
Impact	File server upgrade failed. Some file server features may not be functioning.
Resolution	Contact Nutanix support.
Severity	Critical

Table 237: File Server SMB Share is unavailable [160095] [A160095]

Name	File Server Multi-Protocol Backward Compatibility Failure
Description	File Server SMB Share is unavailable.
Alert message	*message*
Cause	Multi-protocol share backward compatibility failed.
Impact	File Server SMB Share is unavailable.
Resolution	Please retry enabling multi-protocol access, if the issue persists, contact Nutanix support.

Severity	Critical

Table 238: Auto Snapshot Failed [160096] [A160096]

Name	File Server Auto Snapshot Failure
Description	File server auto snapshot failed.
Alert message	`message`
Cause	Auto snapshot was running for a long time. System is overloaded.
Impact	File server auto snapshot failed.
Resolution	Contact Nutanix support.
Severity	Critical

Table 239: Some files are possibly corrupted [160097] [A160097]

Name	File Server Corrupted Files
Description	Some files are possibly corrupted.
Alert message	`message`
Cause	Some files are possibly corrupted for file server.
Impact	Affected data may not be readable until fixed.
Resolution	Contact Nutanix support.
Severity	Critical

Table 240: Some File server share(s) are unavailable [160098] [A160098]

Name	File Server Shares Unavailable
Description	Some File server share(s) are unavailable.
Alert message	`message`
Cause	See details in the alert message.
Impact	Some shares may be unavailable.
Resolution	Contact Nutanix support.
Severity	Critical

Table 241: File server rebalance is aborted [160099] [A160099]

Name	File Server Rebalance Aborted
Description	File server rebalance is aborted
Alert message	`reason_and_resolution_msg`

Cause	See details in the alert message.
Impact	Rebalance is not performed.
Resolution	See details in the alert message. Contact Nutanix support if necessary.
KB Article	7035
Severity	Critical
Severity	Warning

Table 242: File server share is inaccessible or read-only [160102] [A160102]

Name	File Server Share Access Changed
Description	Share has become readonly or inaccessible.
Alert message	`share share_name on File Server file_server_name is message`
Cause	File analytics has set a filter to set share readonly or inaccessible.
Impact	Share may be inaccessible or read-only.
Resolution	Check File Analytics page for available actions. Contact Nutanix support if problem persists.
Severity	Critical
Severity	Warning

Table 243: File server metadata scan is aborted [160122] [A160122]

Name	File Server Metadata Scan Aborted
Description	File server metadata scan is aborted
Alert message	`message`
Cause	See details in the alert message.
Impact	Metadata scan is not completed.
Resolution	See details in the alert message.
Severity	Warning

Table 244: File server external network not functional. [160123] [A160123]

Name	File Server External IP Failure
Description	One or more of the external IP(s) may not be operational or reachable.
Alert message	`message`
Cause	Most probable cause is either network unreachability on the external interface or core Nutanix files cluster services are down/unstable. See details in the alert message.
Impact	Shares may become unavailable.

Resolution	Ensure that external network is reachable, please refer to KB10434 for more details. Contact Nutanix support if the issue still persists.
KB Article	10434
Severity	Critical

Table 245: File server persistent store volume-group unavailable [160124] [A160124]

Name	File Server Persist Store VG Unavailable
Description	One or more of persist store Volume Group(s) may not be operational.
Alert message	`message`
Cause	Most probable cause can be underneath storage issues on AOS or Files cluster virtual IP connectivity problems. See details in the alert message.
Impact	Some Persist Store Volume Group(s) may be unavailable which can lead to issues with Continuously Available shares.
Resolution	Execute Cluster Health check to ensure there are no underlying AOS issues, please refer to KB10435 for more details. Contact Nutanix support if the issue still persists.
KB Article	10435
Severity	Critical

Table 246: File server Inode usage is high [160125] [A160125]

Name	File Server Inode Usage High
Description	The number of free Inodes in the cluster is getting low. Inode contains information about data files attributes.
Alert message	`message`
Cause	There can be multiple causes including overall high storage utilization. See details in the alert message.
Impact	The system may experience unavailability.
Resolution	Execute Cluster Health check to ensure there are no underlying storage high utilization issues, please refer to KB10436 for more details. Contact Nutanix support if the issue still persists.
KB Article	10436
Severity	Critical

Table 247: File server iSCSI service is not healthy [160126] [A160126]

Name	File Server iSCSI Service Unavailable
Description	iSCSI service is not in an healthy state on the cluster and can lead to storage unavailability.

Alert message	*message*
Cause	iSCSI service is in a hung state or not in a functional state.
Impact	Shares may become unavailable.
Resolution	Execute Cluster Health check to ensure there are no underlying cluster issues, please refer to KB10437 for more details. Contact Nutanix support if the issue still persists.
KB Article	10437
Severity	Critical

Table 248: Platform is not reachable at required ports from the fileserver VM. [160134] [A160134]

Name	File Server Cvm Port Unreachable
Description	Platform is not reachable at required ports from the fileserver VM.
Alert message	*message*
Cause	Either Platform services are down OR Platform ports are blocked by firewall.
Impact	File Server services are impacted.
Resolution	Check Platform status. Ensure the File Server Port Requirements are met entirely, Refer Nutanix Files User Guide. Refer to KB article 11022. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11022
Severity	Critical

Table 249: File server slog partitioning failed. [160140] [A160140]

Name	File Server Slog Partition Failure
Description	File server slog partition failure for volume group.
Alert message	*message*
Cause	Device partition creation failed.
Impact	Share write performance is impacted.
Resolution	Refer to KB article 11400. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11400
Severity	Critical

Table 250: File server slog update task failed. [160141] [A160141]

Name	File Server Slog Update Task Failure
Description	File server slog update task failure for volume groups.

Alert message	*message*
Cause	Storage update with slog device failed.
Impact	File server is in heterogenous state.
Resolution	Refer to KB article 11412. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11412
Severity	Critical

Table 251: Cluster Connectivity Status [200000] [A200000]

Name	Cluster Connectivity Status
Description	Tests whether the cluster connectivity is fine
Alert message	*component* **data from cluster** *cluster_name* **is not up-to-date.**
Cause	The connection between Prism Central and Prism Element is down or Services on Prism Central are crashing.
Impact	Cluster data shown in the Prism Central is not up to date.
Resolution	Ensure that cluster network connectivity is up and all CVM services are up.
KB Article	3379
Severity	Warning

Table 252: IDF Source to Remote Sync Status [200001] [A200001]

Name	IDF DB to DB sync heartbeat check
Description	Checks if IDF replica is receiving requests from IDF master.
Alert message	**IDF data on this cluster is not in sync with the IDF data on cluster:***cluster_name*
Cause	Cluster network connectivity or cluster services such as insights server, insights uploader, insights receiver, Aplos, or Prism gateway could be down.
Impact	IDF data that is replicated from master cluster is not up to date.
Resolution	Ensure that cluster network connectivity is up and all cluster services are up.
KB Article	5582
Severity	Warning

Table 253: Dynamic scheduling failure. [200201] [A200201]

Name	Acropolis Dynamic Scheduler Status
Description	One or more nodes have resource contention. This imbalance can cause performance bottlenecks on the node(s) affected.
Alert message	**Dynamic scheduling failure.** *reason*

Cause	Cluster may have insufficient CPU or Controller VM resources.
Cause	Cluster may have insufficient resources to satisfy VM group affinity policies or VM host affinity policies.
Impact	Cluster performance may be significantly degraded.
Resolution	Shut down unneeded VMs to free cluster resources.
Resolution	Expand the cluster to add resources.
Severity	Warning

Table 254: Prism Central vCPU Availability Check [200301] [A200301]

Name	Prism Central vCPU Availability Check
Description	Checks if the number of vCPUs is sufficient for the number of VM entities in Prism Central.
Alert message	`The Prism Central does not have enough vCPUs for the number of VM entities it has.`
Cause	Prism Central doesn't have enough vCPUs to support the current number of VMs
Impact	Prism Central performance may be degraded.
Resolution	1) Reduce the number of VM entities in Prism Central or 2) If using 1 node small Prism Central scale up to 1 node large Prism Central or scaleout to 3 node small Prism Central. 3) If using 1 node large Prism Central scaleout it to 3 node large Prism Central.
KB Article	4381
Severity	Warning

Table 255: Insufficient Space for UVMs deployed on PC [200302] [A200302]

Name	PC Sufficient Disk Space Check
Description	Checks if the amount of storage is sufficient for the number of VM entities in Prism Central.
Alert message	`The PC does not have enough storage for the number of VM entities it has.`
Cause	Too many VMs in Prism Central for the amount of storage in Prism Central.
Impact	Prism Central may run out of disk space to store data.
Resolution	Reduce the number of VM entities in Prism Central or provide more storage for Prism Central.
KB Article	4340
Severity	Warning

Table 256: Prism Central Memory Availability Check [200303] [A200303]

Name	Prism Central Memory Availability Check
Description	Checks if the amount of memory is sufficient for the number of VM entities and the services enabled in Prism Central.
Alert message	`The Prism Central does not have enough memory for the current number of VMs.`
Cause	Prism Central doesn't have enough memory to support the current number of VMs
Impact	Services running in Prism Central may run out of memory and crash.
Resolution	1) Reduce the number of VM entities in Prism Central or 2) If using 1 node small Prism Central scale up to 1 node large Prism Central or scaleout to 3 node small Prism Central. 3) If using 1 node large Prism Central scaleout it to 3 node large Prism Central.
KB Article	4315
Severity	Warning

Table 257: Prism Central VM Limit Check [200304] [A200304]

Name	Prism Central VM Limit Check
Description	Checks if the number of VM entities is within the limit.
Alert message	`The Prism Central cannot handle these many VM entities.`
Cause	Too many VMs in Prism Central.
Impact	Prism Central performance may be degraded.
Impact	Prism Central may run out of disk space to store data.
Impact	Services running in Prism Central may run out of memory and crash.
Resolution	Reduce the number of VM entities in Prism Central.
KB Article	4296
Severity	Warning

Table 258: Policy not applicable to any Host. [200305] [A200305]

Name	No viable Hosts to place VMs in PE
Description	Policy as defined does not apply to any of the Hosts.
Alert message	`Rule rule_uuid applies to no Host in cluster cluster_uuid`
Cause	Addition of a new policy(s) or update to existing policy(s) or Attaching/ Detaching categories to Hosts.
Impact	VMs will not have any Hosts to run on.
Resolution	Make appropriate updates to affinity policies or Host categories in the affected PE.

Severity	Warning

Table 259: Prism Central cluster version check [200310]

Name	Prism Central cluster version check
Description	Checks if Prism Central is running GA version.
Cause	Non GA version detected on Prism Central.
Impact	Prism Central is running a Non GA version.
Resolution	Nutanix recommends running officially released GA versions of Prism Central software. For assistance refer to KB 6325
KB Article	6325

Table 260: Prism Central Version EOL [200311] [A200311]

Name	Prism Central version EOL check
Description	Checks if Prism Central is running an EOL version.
Alert message	*info_msg*. Upgrade Prism Central at the earliest.
Cause	EOL version running on Prism Central
Impact	Support may not be able to resolve issues since this version will no longer be maintained.
Resolution	Move to a recent released version immediately, to avoid disruption in support.You can refer to the EOL Schedule on the Nutanix support portal.
KB Article	6325
Severity	Info

Table 261: Inadequate Resource Configuration for the Prism Central VM [200327] [A200327]

Name	Prism Central Resource Resize Check
Description	Configured resource for the Prism Central VM is inadequate.
Alert message	*inadequate_resources* for the *pc_vm_type* Prism Central VM is inadequate on *non_resized_pc_vm_ips*.
Cause	The current feature capabilities of Prism Central require resource on the Prism Central VM to be increased for optimum performance.
Impact	Prism Central performance may be degraded.
Resolution	Increase the resource of the Prism Central VM. Please refer KB-8932 for guidance.
KB Article	8932
Severity	Warning

Table 262: Prism Central EXT4-fs Error Check [200329]

Name	Prism Central EXT4-fs Error Check
Description	Captures Prism Central EXT4-fs error messages
Cause	File system inconsistencies are present on the Prism Central VM node.
Impact	Inability for the Prism Central VM to boot or for the upgrade pre-checks to run.
Resolution	Look for any problems in the file system. Review KB 9622.
KB Article	9622

Table 263: Prism Central Backup Sync has failed. [200331] [A200331]

Name	Prism Central Backup Sync check
Description	Checks if Prism Central backup sync is stopped for more than 3 hours
Alert message	`Prism Central Backup Sync to these Prism Element(s) : pe_names is halted for the last : delays`
Cause	Prism Central Backup Sync to designated Backup Prism Element(s) is halted for more than 3 hours.
Impact	Prism Central may not be restorable during disaster.
Resolution	Refer to KB 10788 for further details.
KB Article	10788
Severity	Critical

Table 264: VM forcibly powered off. [200401] [A200401]

Name	VM forcibly powered off
Description	We have detected that VM {vm_name} (vm_id: {vm_id}) was running on multiple hosts. To avoid data corruption this VM was forcibly powered off on host {host_address}. {reason}
Alert message	`VM named vm_name with id vm_id has been forcibly powered off on host host_address`
Cause	Unexpected restart of Acropolis service during VM migration.
Cause	If VM does not have local disks, but using network-attached storage, it will be forcibly powered off to prevent data corruption on network-attached storage.
Impact	If the VM does not have any disks on Nutanix storage and is only using network-attached storage, data corruption on network-attached storage is possible.
Resolution	Please collect the log bundle and engage Nutanix support.
Resolution	Check the integrity of data located on network-attached storage that is in use by this VM.
Severity	Warning

Table 265: Power cycle VMs before performing upgrade or migrate operation [200404] [A200404]

Name	Power cycle VMs before performing upgrade or migrate operation
Description	Power cycle all the VMs before you perform the upgrade or migrate operation.
Alert message	`Power cycle the VMs before you perform the upgrade or migration operation`
Cause	A node with a lower CPU feature set has been added to the cluster.
Impact	You will not be able to migrate the existing VMs to newly added node(s). Operations that require VM migration, including but not limited to hypervisor upgrade, will fail.
Resolution	Ensure that all the VMs in the cluster are power cycled.
Severity	Warning

Table 266: Report Generation Failure [200501] [A200501]

Name	Report Generation Failure
Description	Report Generation failed. Database may be down or cluster may have insufficient storage.
Alert message	`Report generation failed for report report_name scheduled for timestamp. reason`
Cause	Cluster may have insufficient storage.
Impact	Some of the reports can be missed.
Resolution	(A) Free the space by deleting older reports that have been generated. (B) Manually run the report instance scanner utility.
Severity	Warning

Table 267: Send Report Through E-mail Failure [200502] [A200502]

Name	Send Report Through E-mail Failure
Description	Sending of report through E-mail failed. SMTP might be down or report might have been deleted.
Alert message	`Sending of report through E-mail failed for report report_name. reason`
Cause	(A) SMTP server might be unreachable or might not have been configured. (B) Generated report may have been deleted.
Impact	Some of the report emails can be missed.
Resolution	Check the SMTP configuration and ensure that its reachable.
Severity	Warning

Table 268: Report Quota Scan Failure [200503] [A200503]

Name	Report Quota Scan Failure
Description	Report Quota scan failed. Cassandra might be unreachable.
Alert message	`Report quota scan failed for the report report_name at timestamp. reason`
Cause	Cassandra may not be reachable for some time.
Impact	There might be some stale reports on the cluster that are taking unnecessary space on the cluster.
Resolution	Manually run the report instance scanner utility.
Severity	Warning

Table 269: VPC VMs detected in Flow Security Policy [200608] [A200608]

Name	VPC VMs detected in Flow Security Policy
Description	VMs that have NICs in one or more VPCs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for these VMs.
Alert message	`num_vms VPC VM(s) detected in Flow Security Policy rule_name (network_security_rule_uuid). VMs: vms. Flow Security Policies will not be enforced for the affected VMs.`
Cause	VMs with NICs in VPC subnets have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the affected VMs from VPC subnets for Flow Security Policies to be enforced on these VMs.
Severity	Warning

Table 270: VMs with kDirectNics detected in Flow Security Policy [200609] [A200609]

Name	VMs with kDirectNics detected in Flow Security Policy
Description	VMs that have one or more kDirectNICs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for the traffic through these kDirectNics.
Alert message	`num_vms VM(s) with kDirectNics detected in Flow Security Policy rule_name (network_security_rule_uuid). VMs: vms. Flow Security Policies will not be enforced for the traffic through these kDirectNics.`
Cause	VMs with kDirectNics have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the kDirectNics from the vms for Flow Security Policies to be enforced.
Severity	Warning

Table 271: Availability Zone Connection Failure [200801] [A200801]

Name	Availability Zone Connection Failure
Description	The remote availability zone {az_url} is unreachable.
Alert message	`The remote availability zone az_url is unreachable.`
Cause	(A) Services on remote site might be down. (B) Connection might be missing on remote site. (C) Access tokens might have expired.
Impact	Disaster recovery operations might fail due to unreachability of remote site.
Resolution	Unpair from the remote availability zone, if already paired. Pair with the remote availability zone again.
KB Article	8252
Severity	Critical

Table 272: PE-PC Connection Failure [200802] [A200802]

Name	PE-PC Connection Failure
Description	The remote {remote_vm_type} {cluster_ip} is unreachable.
Alert message	`The remote remote_vm_type cluster_ip is unreachable.`
Cause	The name servers might be missing.
Cause	Services on PC/PE might be down.
Cause	Access tokens might have expired.
Impact	Management tasks for various entities, including VM, might fail due to unreachability of remote {remote_vm_type}.
Resolution	Ensure that the appropriate name servers are configured.
Resolution	Check that no services are crashing on the PC or PE.
Resolution	Contact Nutanix Support to troubleshoot this issue further.
Severity	Warning

Table 273: Entity Sync Failure [500101] [A500101]

Name	Entity Sync failed
Description	Entity Sync failed
Alert message	`message`
Cause	An update made to the entity locally might not have synced to the connected Availability Zone(s).
Cause	One or more internal services might be down or not working as expected.
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.

Resolution	Contact Nutanix support.
Severity	Warning

Table 274: Entity Sync Failure for the Protection Policy [500102] [A500102]

Name	Entity Sync failed for Protection Policy
Description	Entity Sync failed for the Protection Policy
Alert message	`message`
Cause	An update made to the Protection Policy locally might not have synced to the connected Availability Zone(s).
Cause	One or more internal services might be down or not working as expected.
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Resolution	Contact Nutanix support.
Severity	Warning

Table 275: Entity Sync Failure for the Recovery Plan [500103] [A500103]

Name	Entity Sync failed for Recovery Plan
Description	Entity Sync failed for the Recovery Plan
Alert message	`message`
Cause	An update made to the Recovery Plan locally might not have synced to the connected Availability Zone(s).
Cause	Protection Policy configuration for the entities specified in the Recovery Plan may be different across connected Availability Zones.
Cause	One or more internal services might be down or not working as expected.
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Resolution	Ensure that the Protection Policy configuration is same across connected Availability Zones for the entities specified in the Recovery Plan.
Resolution	Contact Nutanix support.
Severity	Warning

Table 276: Entity Sync Failure for the Availability Zone [500104] [A500104]

Name	Entity Sync failed for Availability Zone
Description	Entity Sync failed for the Availability Zone

Alert message	`message`
Cause	Unable to send request to the Availability Zone.
Cause	Request timed out.
Impact	Data Protection and Recovery might be impacted.
Resolution	Ensure Availability Zone is reachable and do a 'Forced Entity Sync' to ensure that all the entities are in sync. If the issue still persists, contact Nutanix support.
Resolution	Ensure there are no sync related tasks running under the 'Tasks' page on all the connected Availability Zone(s) and then invoke 'Forced Entity Sync' from the chosen site.
KB Article	10241
Severity	Critical

Table 277: Translated Addresses Retrieval Failure in NGT [600102] [A600102]

Name	Unable to get translated addresses in NGT
Description	NGT is unable to retrieve the XAT address translations.
Alert message	`Unable to retrieve translated addresses in NGT because the Prism Element IDF service may be unreachable.`
Cause	(A) Prism Element IDF may be unreachable. (B) Translation information is not available in IDF.
Impact	NGT features may not work on the User VMs.
Resolution	Check if the Prism Element IDF service is reachable.
Severity	Warning

Table 278: Data provider collector is in crashloop [650000] [A650000]

Name	Data provider collector is in crashloop
Description	Data provider collector is restarting frequently
Alert message	`data_provider_name data collector with uuid metrics_data_provider_instance_uuid is in crashloop`
Cause	Collector heartbeat may not be reaching Prism Central
Impact	Users may notice missing stats for some of the endpoints that are configured through monitoring integration feature.
Resolution	Manually attempt to start mercury service on all Prism Central VMs. if situation persists, or any assistance is needed please contact Nutanix Support.
KB Article	9258
Severity	Warning

Table 279: Data provider collector has violated sizing limits [650001] [A650001]

Name	Data provider collector has sizing violation
Description	Data provider collector metrics usage has violated sizing limits
Alert message	`data_provider_name` data collector has reached `threshold`% of the sizing limits.
Cause	Total no. of instances or Metric collection has reached the supported limit by Data Provider Collector
Impact	Users may not be able to configure new instances through monitoring integration feature.
Resolution	Delete some existing instance to configure new instances if situation persists, or any assistance is needed please contact Nutanix Support.
KB Article	9473
Severity	Critical
Severity	Warning

Table 280: Cluster is not healthy [802004] [A802004]

Name	Failure to configure a cluster for Advanced Networking
Description	Failure to configure a cluster for Advanced Networking.
Alert message	Cluster (`cluster_uuid`) failed to be configured for Atlas Networking
Cause	Cluster cannot accept Advanced Networking Configuration.
Impact	Ability to make network-related configurations may be impacted.
Resolution	(A) Verify health of cluster. (B) Verify cluster is reachable from Prism Central. (C) Verify cluster AOS and AHV version are compatible with Advanced Networking.
Severity	Warning

Table 281: Security Planning is disabled [803001] [A803001]

Name	Security Planning is disabled
Description	Security Planning is disabled due to de-registration of Prism Element from Prism Central.
Alert message	Security Planning is disabled. Reason: `details`
Cause	The cluster that is hosting the Epoch data collector was de-registered from Prism Central.
Impact	The Security Planning feature has been automatically disabled.
Resolution	Enable Security Planning from the Prism Central UI to deploy a new instance of the Epoch data collector.
Severity	Warning

Table 282: ID Firewall lost connectivity to domain controller [803003] [A803003]

Name	ID Firewall is unable to connect to one or more domain controllers
Description	ID Firewall is unable to connect to one or more domain controllers used for scraping identities.
Alert message	`Domain controller domain_controller on domain domain is not reachable and accepting LDAP or WMI connections. Details: details.`
Cause	The domain controller may not be reachable via the network.
Cause	The domain controller may not be accepting LDAP/WMI connections due to permissions issues.
Cause	The domain controller may not be running or fully booted up.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until access to the domain controller is restored.
Resolution	Check network connectivity from Prism Central to the domain controller. Please refer article KB-10219 for more details.
Resolution	Check that the domain controller is fully up and accepting WMI and LDAP connections. Please refer article KB-10219 for more details.
Resolution	Check that the service account registered for ID Firewall is active and has permissions for both WMI and LDAP access, and that the correct password is entered in Prism Central. Please refer article KB-10219 for more details.
KB Article	10219
Severity	Warning

Table 283: ID Firewall unable to locate mapped Active Directory object [803004] [A803004]

Name	ID Firewall mapped object not found
Description	ID Firewall was unable to find a mapped object in Active Directory.
Alert message	`The mapped object_type object_dn with GUID object_guid was not found in domain.`
Cause	The mapped object in Active Directory which is referenced for ID Firewall could not be found because it was deleted from Active Directory.
Impact	Flow VDI security policies written around the mapped object will not function as expected.
Resolution	The mapping should be checked, and if the mapped object does not exist, or if it has been deleted and re-created, the mapping must also be deleted and possibly re-created. Security policies should be updated to use the new mapping or not reference the mapping at all if the object was permanently deleted.
Severity	Warning

Table 284: ID Firewall did not recover state after reconnecting to a temporarily unreachable domain controller [803005] [A803005]

Name	ID Firewall state recovery failed
Description	ID Firewall was unable to recover state after reconnecting to a temporarily unreachable domain controller.
Alert message	`ID Firewall was unable to recover state from domain controller domain_controller on domain domain. Details: details.`
Cause	Recovery failed because either too much time has passed since the domain controller was last reachable, or the event log on the domain controller has rolled over and the events necessary to recover are no longer present.
Impact	The state enforced by ID Firewall may be inconsistent, and therefore applied policies may not be properly enforced until action is taken.
Resolution	In order to ensure security policies are properly enforced, all active Nutanix VDI VM users should log out and log back into their VMs. Please refer KB-10220 for further details.
KB Article	10220
Severity	Warning

Table 285: Epoch Data Collector upgrade available [803006] [A803006]

Name	Epoch Data Collector needs upgrade
Description	Checks whether there is an upgrade available for the Epoch Data Collector.
Alert message	`message`
Cause	A new version of the Epoch Data Collector is available and not yet installed.
Impact	Some new functionality or bugfixes in the Epoch Data Collector may not be available.
Resolution	Go to Epoch Data Collector Enablement page to upgrade the Epoch Data Collector.
Severity	Info

Table 286: ID Firewall service account is invalid [803007] [A803007]

Name	Service account used for ID Firewall is invalid
Description	ID Firewall has detected that the service account configured is not valid.
Alert message	`ID Firewall service account account@domain is invalid. Details: details.`
Cause	The service account password may have been changed.
Cause	The service account may have been deleted.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until the service account is valid again.
Resolution	Check that the configured service account password is up-to-date and update it in Prism if needed.

Resolution	Check that the configured service account exists, and update it in Prism if needed.
Severity	Warning

Table 287: Object Store deployment is not supported on this cluster. [805001] [A805001]

Name	Objects Store Deployment Check
Description	Checks if Object Store is deployed on an unsupported AOS cluster.
Alert message	`The registered cluster(s) affected_clusters_list is/are on one of the versions between 5.18 and 5.18.0.4 which is unsupported for Object Store Deployment.`
Cause	One or more Prism Element cluster(s) attached to this Prism Central is on version(s) between 5.18 and 5.18.0.4 which is unsupported for Object Store Deployment.
Impact	Object Store will not work as expected.
Resolution	Upgrade the cluster to 5.18.0.5 or above.
KB Article	10000
Severity	Critical

Controller VM

Table 288: CVM or Prism Central VM RAM Usage High [3023] [A1056]

Name	CVM Memory Usage
Description	Check that CVM or Prism Central VM memory usage is not high.
Alert message	`Main memory usage in vm_type ip_address is high, available_memory_kb KB is free.`
Cause	The RAM usage on the Controller VM or Prism Central VM has been high.
Impact	Cluster performance may be significantly degraded.
Resolution	Check the memory utilization of Prism Central VM or Controller VM. If abnormal behavior is seen, please collect logs and contact Nutanix Support.
CVM Available Memory Threshold	Critical Threshold: 589824 KB
KB Article	2473
Severity	Critical

Table 289: CVM/PCVM time not synchronized with external servers. [3026] [A3026]

Name	CVM/PCVM NTP Time Synchronized

Description	Checks that the {vm_type} is synchronizing time with an NTP server.
Alert message	*alert_msg*
Cause	External NTP servers are not configured or are not reachable
Impact	Workflows involving Kerberos may fail if the time difference between the Controller VM and the NTP server is greater than 5 minutes.
Resolution	Please verify that the external NTP servers are configured and are reachable.
KB Article	4519
Severity	Warning

Table 290: CVM Rebooted [3028] [A1024]

Name	CVM Rebooted Check
Description	Checks if the CVM is not rebooted recently.
Alert message	*vm_type* *ip_address* has been rebooted on *reboot_timestamp_str*.
Cause	CVM is rebooted.
Impact	During the time the Controller VM is down, cluster compute and storage capacity are reduced.
Resolution	Check CVM status. Ignore if intentional.
KB Article	2474
Severity	Critical

Table 291: CVM Service(s) Restarting Frequently [3029] [A1032]

Name	CVM Services Status
Description	Check if services have restarted recently on the CVM.
Alert message	One or more cluster services have restarted within 15 minutes in the Controller VM *ip_address*, the total restarts being >= 10.The following are the service(s) that have restarted : *failed_components*.
Cause	Faulty behavior in one or more service(s), leading to frequent crashes,in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs more than once, please collect CVM logs and contact Nutanix support.
KB Article	2472
Severity	Critical

Table 292: {services} facing segmentation fault [3036] [A3036]

Name	CPP Services Segmentation fault Check
Description	Check if services faced a segmentation fault recently in PCVM/CVM
Alert message	`More than 5 segmentation fault errors happened in the last 15 minutes in the PCVM/CVM ip_address. Services: services`
Cause	Faulty behavior in one or more service(s).
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this occurs more than 1 time, please refer to KB 2472 for more information. If unsure, please contact Nutanix Support.
KB Article	2472
Severity	Critical

Table 293: CVM Python Service(s) Restarting Frequently [3037] [A3037]

Name	Python Services Fatal Check
Description	Check if python services have crashed/restarted recently in PCVM/CVM
Alert message	`One or more cluster services have restarted within 15 minutes in the PCVM/CVM ip_address, the total restarts being >= 10. The following are the service(s) that have restarted : failed_components.`
Cause	Faulty behavior in one or more service(s) leading to frequent crashes in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs more than once, please contact Nutanix support.
KB Article	2472
Severity	Warning

Table 294: Multiple cpu_unblock processes running [3042] [A3042]

Name	cpu_unblock Process Check
Description	Check that there are no stale cpu_unblock processes running.
Alert message	`Multiple cpu_unblock processes are running on svm_ip.`
Cause	Zookeeper restarting frequently on cluster node.
Impact	Cluster performance may be significantly degraded.
Resolution	Kill all cpu_unblock processes and restart cluster services on node.
KB Article	5581
Severity	Critical

Severity	Warning

Table 295: Prism Central VM CPU Load High [6517] [A6517]

Name	CVM CPU Load High
Description	Check that Prism Central VM CPU load is not high.
Alert message	`CPU load in or Prism Central VM ip_address is high. Load average over 5 Minutes = #load_average_5min@, number of vCPUs in CVM or Prism Central VM = num_procs.`
Cause	The CPU load on the Prism Central VM has been high.
Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix Support for diagnosis. CPU on the Prism Central VM may need to be increased.
KB Article	5227
Severity	Critical
Severity	Warning

Table 296: Prism Central VM disk usage high [101059] [A101059]

Name	Prism Central VM high disk space usage
Description	Check that disk space usage on the Prism Central VM is within limits.
Alert message	`alert_msg`
Cause	High disk usage in the Prism Central VM.
Impact	Prism Central VM may run out of storage space to store data.
Resolution	Refer to KB 5228 for further details.
Prism Central VM disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

Table 297: Core dumps are enabled on this CVM or Prism Central VM. [101069] [A101069]

Name	Saltstate coreoff check
Description	Checks if saltstate is correctly set to default coreoff state to ensure services do not produce core dumps unnecessarily.
Alert message	`Core dumps are enabled on the vm_type with IP service_vm_external_ip.`

Cause	Core dumps are enabled for services running on this CVM or Prism Central VM.
Impact	Core dumps are useful for isolated troubleshooting but leaving enabled may unnecessarily fill the node /home system partition, and this can affect node stability and availability.
Resolution	Follow KB7586 to validate and fix coreoff saltstate. Resolution may involve graceful node reboot, so ensure the cluster is otherwise healthy and resiliency status is OK prior to this action.
KB Article	7586
Severity	Critical

Table 298: Kernel crash dump found. [101073] [A101073]

Name	Kernel crash dump check
Description	Checks if Kernel crash dump logs exists on a node.
Alert message	`alert_msg`
Cause	Kernel dump logs are present on the node.
Impact	Neglecting old kernel crash dump files may unnecessarily fill the system partition, in turn affecting node stability.
Resolution	Please refer to KB 11831 to validate core dump.
KB Article	11831
Severity	Warning

Table 299: Reset Current Passwords for admin and nutanix SSH accounts on {cluster_name} [101081] [A101081]

Name	Password Strength Status
Description	High-strength password policy has been enabled, please reset the password.
Alert message	`High-strength password was recently enabled, please reset current passwords for admin and nutanix SSH accounts on cluster_name - cluster_uuid.`
Cause	High-strength password policy has been enabled.
Impact	High-strength password policy has been enabled but low-strength passwords may still be in use on the cluster.
Resolution	Please reset the current passwords. Follow the Security Guide's Hardening Instructions on Nutanix Support Portal for the password requirements.
Severity	Warning

Table 300: Host time not synchronized with any external servers. [103090] [A103090]

Name	Host time synchronized

Description	Checks to ensure that the host is configured and synchronizing time with an NTP server.
Alert message	*alert_msg*
Cause	External NTP servers are not configured or are not reachable.
Impact	Logs may have different timestamps in the host and the CVMs/PCVM. Host may not work as expected.
Resolution	Please verify that the NTP servers are configured and are reachable from the host.
KB Article	4519
Severity	Warning

Table 301: CVM TCP checksum error rate check [103104]

Name	CVM TCP checksum error rate check
Description	Check if corrupted packets with compromised checksums are reaching the CVM and not dropped by the host NICs.
Cause	Hardware failure or NIC firmware/driver compatibility issue.
Impact	Cluster performance may be significantly degraded.
Resolution	Check for any hardware failures or issues with firmware/driver/cables. Please refer to KB 9623 for more details.
KB Article	9623

Table 302: Node in Maintenance Mode [130001] [A1013]

Name	Node Maintenance Mode Status
Description	Node in Maintenance Mode.
Alert message	`Controller VM` *ip_address* `is put in maintenance mode due to` *reason*`.`
Cause	Node removal has been initiated.
Impact	The node will be removed from the cluster.
Resolution	No action is necessary.
Severity	Info

Table 303: Kerberos Clock Skew Failure [130024] [A1083]

Name	Kerberos Clock Skew Status
Description	Kerberos clock skew failure.
Alert message	`Kerberos authentication failure occurred on the Controller VM` *ip_address* `due to clock skew between the Controller VM and the Domain Controller (reason:` *reason*`).`

Cause	There is considerable clock skew between the Controller VM and the Domain Controller.
Impact	The Controller VM cannot manage the host, which may lead to node instability.
Resolution	Ensure that the time on the Controller VM is synchronized with the time of the host. It can be accomplished by configuring the NTP server properly.
KB Article	8036
Severity	Warning

Table 304: Metadata Volume Snapshot Persistent Failure [130026] [A1087]

Name	Metadata Volume Snapshot Persistent
Description	Metadata Volume Snapshot persistent failure.
Alert message	`The last failure_count metadata volume snapshots have failed. reason`
Cause	The EBS snapshot service is unavailable, possibly because AWS could not be reached.
Impact	If the cloud instance fails, backups of the Guest VMs may be unavailable.
Resolution	Please check connectivity to cloud instance and reach out to Nutanix Support for assistance.
KB Article	8413
Severity	Warning

Table 305: Metadata Volume Snapshot Timeout [130027] [A1088]

Name	Metadata Volume Snapshot Status
Description	Metadata Volume Snapshot timeout failure.
Alert message	`Metadata volume snapshot could not be created in duration_mins minutes.`
Cause	Metadata volume snapshot frequency is too high.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.
Severity	Warning

Table 306: Physical Disk Added To Slot [130034] [A1102]

Name	Physical Disk Add Check
Description	Physical Disk added to slot.
Alert message	`Disk with serial disk_serial and model disk_model was added in drive bay disk_location on Controller VM service_vm_external_ip. Action Required: action. Visit Prism hardware page for further actions.`

Cause	A drive was inserted into an empty slot.
Impact	The drive may not be available to the cluster until further configuration is completed.
Resolution	Ensure that the drive is added correctly from the Hardware page.
KB Article	9288
Severity	Critical
Severity	Warning
Severity	Info

Table 307: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	`Drive disk_id with serial disk_serial and model disk_model in drive bay disk_location on Controller VM service_vm_external_ip has failed. Node serial is node_serial_number and Node position node_position. #rf1_custom_message@`
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 308: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical drive is removed from the slot.
Alert message	`Drive with serial disk_serial and model disk_model was removed from drive bay disk_location on Controller VM service_vm_external_ip. Node serial is node_serial_number and Node position node_position. #rf1_custom_message@`
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the drive will start.
Resolution	Remove the drive from the slot and reseat it. Contact the Hardware Vendor if issue persists.
KB Article	8031
Severity	Critical

Table 309: Self-encrypting Drive Operation Failure [130051] [A1111]

Name	SED Operation Status
Description	Self-encrypting drive operation failure.
Alert message	`Self-encrypted Drive operation operation has failed for disk disk_id with serial disk_serial on node service_vm_external_ip.`
Cause	A self-encrypting drive operation could not be performed.
Impact	The desired action could not be performed.
Resolution	Retest key management server configuration to ensure connectivity and certificates are valid then retry the command.
Severity	Warning

Table 310: Stargate Temporarily Down [130054] [A1030]

Name	Stargate Status
Description	Stargate Temporarily Down.
Alert message	`Stargate on Controller VM ip_address is down for downtime seconds. #rf1_custom_message@`
Cause	Zookeeper service not running.
Cause	Cassandra service not running.
Cause	Stargate keeps timing out during restart.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If CVM node running the zeus service was powered off, power it on. Otherwise contact Nutanix support.
Resolution	Contact Nutanix support.
Resolution	Contact Nutanix support.
KB Article	3784
Severity	Critical
Severity	Warning

Table 311: Unsupported Configuration For Redundancy Factor 3 [130055] [A1092]

Name	FT2 Configuration
Description	Unsupported Configuration For Redundancy Factor 3.
Alert message	`Controller VM service_vm_id with IP address service_vm_external_ip has actual_ram_size_gbGB RAM which does not meet the configuration requirement of min_ram_required_gbGB RAM to support 'Redundancy Factor 3' feature.`

Cause	To support 'Redundancy Factor 3' feature all controller VMs in the cluster must meet the minimum requirements for RAM.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Increase RAM on the Controller VM to meet the minimum requirement. Contact Nutanix support for assistance.
Severity	Warning

Table 312: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	`Drive disk_id with serial disk_serial in drive bay disk_location on Controller VM service_vm_external_ip has failed diagnostic test. This drive is failing or has failed. #rf1_custom_message@`
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 313: Prism Central VM platform disk space usage high [200328] [A200328]

Name	Prism Central VM platform disk space usage check
Description	Check that platform disk space usage on the Prism Central VM is within limits.
Alert message	`Platform disk space usage in Prism Central VM svm_ip exceeds percentage_exceed% for disk(s): disk_paths.`
Cause	High disk usage in the Prism Central VM.
Impact	Prism Central VM may run out of storage space to store data.
Resolution	Refer to KB 5228 for further details.
Disk Usage Percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

Table 314: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed

Description	Programming a Flow rule failed.
Alert message	`Flow policy operation failed for rule_name. reason`
Cause	Microsegmentation service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.
Resolution	Check PC microsegmentation service, PE acropolis service, PC to PE connection, PE to AHV connection.
Severity	Warning

Table 315: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed
Description	Flow failure event.
Alert message	`Flow operation failed on remote_uuid. reason`
Cause	Microsegmentation service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microsegmentation service, Check PE acropolis service, AHV microsegmentation module.
Severity	Warning

Table 316: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	`Flow mode change failed on remote_uuid. reason`
Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect cvm_logs and contact Nutanix Support
Severity	Warning

Table 317: VPC VMs detected in Flow Security Policy [200608] [A200608]

Name	VPC VMs detected in Flow Security Policy
Description	VMs that have NICs in one or more VPCs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for these VMs.
Alert message	`num_vms VPC VM(s) detected in Flow Security Policy rule_name (network_security_rule_uuid). VMs: vms. Flow Security Policies will not be enforced for the affected VMs.`
Cause	VMs with NICs in VPC subnets have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the affected VMs from VPC subnets for Flow Security Policies to be enforced on these VMs.
Severity	Warning

Table 318: VMs with kDirectNics detected in Flow Security Policy [200609] [A200609]

Name	VMs with kDirectNics detected in Flow Security Policy
Description	VMs that have one or more kDirectNICs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for the traffic through these kDirectNics.
Alert message	`num_vms VM(s) with kDirectNics detected in Flow Security Policy rule_name (network_security_rule_uuid). VMs: vms. Flow Security Policies will not be enforced for the traffic through these kDirectNics.`
Cause	VMs with kDirectNics have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the kDirectNics from the vms for Flow Security Policies to be enforced.
Severity	Warning

DR

Table 319: Remote Replication is Lagging for Protection Domain Snapshot [110260] [A110260]

Name	PD Snapshot Lagging Remote Replication
Description	Remote replication of the protection domain snapshot is lagging.
Alert message	`Snapshot 'snapshot_uuid' of the protection domain 'protection_domain_name' has replication to remote 'remote_name' that is lagging by lagging_minutes minutes.`
Cause	Replication of the protection domain snapshot to the remote cluster is too slow.
Impact	Recovery point objective for the protection domain will be affected.

Resolution	(A) Ensure that the number of entities are within the supported limit. (B) Ensure that replication bandwidth to remote cluster is unaffected.
KB Article	10197
Severity	Warning
Severity	Info

Table 320: NearSync Replication of Protection Domain has not progressed. [110264] [A110264]

Name	Nearsync Replication Stuck
Description	Nearsync replication of the protection domain has not progressed.
Alert message	`NearSync replication of the protection domain protection_domain_name to the remote site remote_name for the snapshot snapshot_id has not made any progress in the last time_string.`
Cause	Network outage, saturated network or unavailability of services like Stargate, Cerebro on the remote site. Snapshots of the protection domain will not be available on the remote site.
Cause	SSD snapshot reserve on the remote site does not have enough free disk space.
Impact	Protection domain may transition out of NearSync.
Impact	Snapshots of the protection domain will not be available on the remote site.
Resolution	Check the network connection and status of the remote site.
Resolution	Ensure that the remote site has enough disk space available in the SSD snapshot reserve.
KB Article	10249
Severity	Warning

Table 321: Availability Zone configured in Protection Policy is not accessible [110401] [A110401]

Name	Availability Zone Accessibility Check
Description	Check if availability zones configured in protection policy are accessible.
Alert message	`Availability Zone(s) fault_azs configured in the Protection Policy protection_rule_name is/are not accessible.`
Cause	Availability zone configured in the protection rule is not accessible.
Impact	Availability zones of the backup site will not be reachable for taking backup.
Resolution	Check the health status of the configured availability zone.
KB Article	5685
Severity	Critical

Table 322: Protection Policy Max entities per Category Check Failed. [110402] [A110402]

Name	Protection Policy Max entities Per Category Exceeded
Description	Checks if the entity count for a category specified in Protection Policy exceeds the maximum allowed limit.
Alert message	`Maximum number of entities for a category in a Protection Policy should not exceed max_entity_count. Following categories exceeds entities limit : categories.`
Cause	Number of entities for the specified categories in the Protection Policy exceeds the limit.
Impact	Specified category will not be considered for the recovery as the Recovery Plan supports categories with limited number of entities.
Resolution	Reduce the protected entity count for the specified categories in the Protection Policy.
KB Article	6306
Severity	Warning

Table 323: Disconnected Availability Zones are affecting some entities. [110403] [A110403]

Name	Entities Affected By Disconnected Availability Zones
Description	Lists all entities which will be affected due to the disconnected Availability Zones.
Alert message	`entity_type of name name has been affected by disconnected availability zones.`
Cause	Availability zone configured in the protection rule has been disconnected.
Impact	Operations related to the affected entities will be affected.
Resolution	Check the health status of the configured availability zone.
KB Article	6283
Severity	Info

Table 324: Stale VM present in DR system [110457] [A110457]

Name	Stale VM present in the DR system
Description	A stale VM in the DR protection conflicts with the live VM
Alert message	`A stale VM stale_vm_name with ID(stale_vm_hs_id, vm_ha_uuid) against VM vm_name with ID(vm_id, vm_ha_uuid). Please unprotect the stale VM entry protection_information.`
Cause	An in-place restore workflow might have been performed on the VM by backup software resulting in newly created VM with the same Nutanix ID
Impact	Some or all DR workflows on this VM might fail
Resolution	Please unprotect the stale VM and if needed, protect the restored VM. If the issue still persists contact Nutanix Support

KB Article	10545
Severity	Critical

Table 325: Secure Tunnel To Remote Site Down [130000] [A1090]

Name	Remote Site Tunnel Status
Description	Secure Tunnel To Remote Site Down.
Alert message	`Secure tunnel to remote site remote_name is down.`
Cause	Incorrect remote site configuration or network connectivity issue.
Impact	Replication to remote site will fail.
Resolution	Check if the IP address specified in the remote site is reachable.
Severity	Critical
Severity	Warning
Severity	Info

Table 326: Metro Availability Is Disabled [130002] [A1124]

Name	Metro Connectivity
Description	Metro availability is disabled.
Alert message	`Metro availability for the protection domain 'protection_domain_name' to the remote site 'remote_name' is disabled because of 'reason'.`
Cause	Remote site unreachable.
Impact	Metro availability operation is disabled.
Resolution	Check network connectivity and health of the remote cluster. If issue persists reach out to Nutanix Support
KB Article	8228
Severity	Critical
Severity	Warning

Table 327: Duplicate Remote Cluster ID [130018] [A1038]

Name	Duplicate Remote Cluster ID Check
Description	Duplicate remote cluster ID.
Alert message	`Remote site already exists with the name 'conflicting_remote_name'.`
Cause	Two remote sites with different names or different IP addresses have same cluster ID. This can happen in two cases: (a) A remote cluster is added twice under two different names (through different IP addresses) or (b) Two clusters have the same cluster ID.
Impact	Protected data is not replicated to the remote site.

Resolution	In case (a) remove the duplicate remote site. In case (b) verify that the both clusters have the same cluster ID and contact Nutanix support.
Severity	Warning

Table 328: Entities Restored But Unprotected [130019] [A1134]

Name	Restored Entities Protected
Description	Entities are restored but unprotected.
Alert message	`Some entities became unprotected after restoring protection_domain_name. Count: num_entities. Entity names: entity_names. Reason: reason.`
Cause	Some other protected entities conflict with the entities being recovered.
Impact	Some entities became unprotected after restoring the snapshot.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 329: Entities Skipped During Restore [130020] [A1132]

Name	Entities Restored Check
Description	Entities were skipped during restore.
Alert message	`Some entities skipped during restore of protection_domain_name. Count: num_entities. Entity names: entity_names. Reason: reason.`
Cause	Existing files conflict with the files to be recovered.
Impact	Not all entities could be restored from the snapshot.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	8415
Severity	Critical

Table 330: Entity conflict in Consistency Groups [130022] [A1136]

Name	Consistency Group Configuration
Description	Entity conflict in consistency group(s)
Alert message	`Failed to create a snapshot of protection domain protection_domain_name as some files in consistency group consistency_group_name overlap other consistency groups.`
Cause	VMs with common files are not in the same Consistency Group.
Impact	Snapshot will fail.
Resolution	Put VMs with common files in the same Consistency Group.

Severity	Warning

Table 331: Metadata Volume Snapshot Persistent Failure [130026] [A1087]

Name	Metadata Volume Snapshot Persistent
Description	Metadata Volume Snapshot persistent failure.
Alert message	`The last failure_count metadata volume snapshots have failed. reason`
Cause	The EBS snapshot service is unavailable, possibly because AWS could not be reached.
Impact	If the cloud instance fails, backups of the Guest VMs may be unavailable.
Resolution	Please check connectivity to cloud instance and reach out to Nutanix Support for assistance.
KB Article	8413
Severity	Warning

Table 332: Metadata Volume Snapshot Timeout [130027] [A1088]

Name	Metadata Volume Snapshot Status
Description	Metadata Volume Snapshot timeout failure.
Alert message	`Metadata volume snapshot could not be created in duration_mins minutes.`
Cause	Metadata volume snapshot frequency is too high.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.
Severity	Warning

Table 333: Metro Availability Configuration Failed [130028] [A1123]

Name	Metro Availability
Description	Metro availability start failed.
Alert message	`Failed to operation for the protection domain 'protection_domain_name' to the remote site 'remote_name'. Reason: 'reason'.`
Cause	Check the alert message for the reason of failure
Impact	Metro availability operation could not be started.
Resolution	Resolve the issue as stated in the alert message and retry the Metro operation. If the issue persists contact Nutanix Support
KB Article	10250
Severity	Critical

Table 334: Stale NFS Mount [130029] [A1135]

Name	Metro Vstore Mount Status
Description	Stale NFS mount.
Alert message	`Unable to process NFS requests for vstore id vstore_id due to reason.`
Cause	Stale NFS file handle.
Impact	Unable to process NFS requests.
Resolution	Unmount the datastore from all hosts and remount again.
Severity	Warning

Table 335: Compliance Failure of the Protected VMs. [130037] [A1109]

Name	Protected VM CBR Capability
Description	Protected VM is not Nutanix backup and recovery compliant.
Alert message	`message`
Cause	Protected VM may have been modified to include files that are not on a Nutanix Storage Container.
Impact	Any data associated with the VM may not be backed up or replicated to the remote site.
Resolution	Remove or detach files from the VM that are not on the Nutanix Storage Containers and make the VM Nutanix backup and recovery compliant.
Severity	Warning

Table 336: Protected VM Not Found [130038] [A1010]

Name	Protected VM Not Found
Description	Protected VM not found.
Alert message	`Unable to locate the VM 'vm_name (ID vm_id)' protected by protection domain 'protection_domain_name'.`
Cause	The protected VM cannot be found and may have been deleted.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove the VM from the protection domain.
Severity	Info

Table 337: Protection Domain Activation Failed [130040] [A1043]

Name	PD Active
Description	Protection domain activation failed.

Alert message	`Unable to make protection domain 'protection_domain_name' active on remote site 'remote_name' due to 'reason'.`
Cause	Various
Impact	Protected VMs could not be started during failover to a remote site.
Resolution	Resolve the stated reason for the failure. If issue persists, contact Nutanix support.
KB Article	8468
Severity	Critical

Table 338: Protection Domain Activation or Migration Failure [130041] [A1060]

Name	PD Change Mode Status
Description	Protection domain activation or migration failed.
Alert message	`Protection domain protection_domain_name activate/deactivate failed with the error : reason`
Cause	(A) Protection domain cannot be activated or migrated. (B) Protection domain with same name might be active on remote site. (C) Remote sites might not be configured correctly, please check the remote sites on source and target clusters. (D) Protection Domain has one or more missing VMs/VGs.
Impact	Protected VMs could not be started during failover to a remote site.
Impact	Some of the VMs might have been powered off. User has to power on those VMs.
Resolution	Resolve the stated reason for the failure. If issue persists, please contact Nutanix Support.
KB Article	8021
Severity	Critical

Table 339: Registration Of One Or More VMs Failed [130043] [A1093]

Name	PD VM Registration Status
Description	Registration of one or more VMs failed.
Alert message	`Failed to register one or more VMs of the protection domain protection_domain_name.`
Cause	VMs in the snapshot have invalid VM configuration or the local hypervisor reported an error.
Impact	VMs restored from the snapshot will not be unavailable.
Resolution	Check the existence of the VM in the snapshot and check any errors reported in the hypervisor management software. If you cannot resolve the error, contact Nutanix support.
KB Article	8005
Severity	Critical

Table 340: Protection Domain Replication Expired [130045] [A1003]

Name	PD Replication Expiry Status
Description	Protection domain replication expired.
Alert message	`Protection domain protection_domain_name replication to the remote site remote_name has expired before it is started.`
Cause	Replication is taking too long to complete before the snapshots expire.
Impact	The snapshot that was scheduled to be replicated could not be replicated. Data that is expected to be protected may not be protected.
Resolution	Review replication schedules taking into account bandwidth and overall load on systems. Confirm retention time on replicated snapshots. If issue persists, please contact Nutanix Support.
Severity	Warning

Table 341: Protection Domain Replication Failure [130046] [A1015]

Name	PD Replication Status
Description	Protection Domain Replication Failure
Alert message	`Protection domain protection_domain_name replication to remote site remote_name failed. reason.`
Cause	Various
Impact	Snapshots of protected VMs were not replicated to the remote site.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 342: Skipped Replication Of Snapshot For Protection Domain [130047] [A1113]

Name	PD Replication Skipped Status
Description	Skipped replication of the snapshot.
Alert message	`Replication was skipped for protection domain protection_domain_name of the latest snapshot snapshot_name to remote site remote_name. reason There could be other snapshots that may have been skipped.`
Cause	Snapshot replication was skipped because a newer snapshot is available.
Impact	Snapshot is not present on the remote site.
Resolution	Make sure that enough bandwidth is available for replication. If the issue persists, consider reducing the frequency of replication. Perform a manual resync of the skipped snapshot, if required.
KB Article	4282
Severity	Warning

Table 343: Failed to Receive Snapshot for the Protection Domain [130048] [A1127]

Name	Protection Domain Receive Snapshot Failure
Description	Protection Domain Receive Snapshot Failure.
Alert message	`Failed to receive snapshot for the protection domain 'protection_domain_name' from the remote site remote_name. reason`
Cause	Incorrect configuration on DR site or Network Error.
Impact	The snapshot receive operation could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 344: Protection Domain Snapshot Failure [130049] [A1064]

Name	PD Snapshot Status
Description	Protection Domain Snapshot Failure.
Alert message	`Protection domain protection_domain_name snapshot 'snapshot_id' failed. reason.`
Cause	Protection domain has VMs being protected by other vstore(s) or in case of a cluster where Metro Availability is configured, Metro protection domain has more entities than supported.
Impact	A requested snapshot of guest VMs and files in the protection domain did not succeed.
Resolution	Unprotect VMs from other vstore(s) before snapshotting the concerned protection domain. In case of a cluster where Metro Availability is configured, make sure number of entities are within the supported limit.
KB Article	8038
Severity	Critical
Severity	Warning

Table 345: Remote Site Is Unhealthy [130050] [A1125]

Name	Remote Site Health
Description	Remote site is unhealthy.
Alert message	`Remote 'remote_name' is unhealthy due to reason 'reason'.`
Cause	Various
Impact	Replication to remote site may fail.
Resolution	Check if the IP address specified in the remote site is reachable.
Severity	Warning

Table 346: Snapshot Partially Crash Consistent [130052] [A1110]

Name	Snapshot Crash Consistent
Description	Snapshot partially crash consistent.
Alert message	`Failed to create an application-consistent snapshot for one or more VMs in snapshot snapshot_name of protection domain protection_domain_name. A crash-consistent snapshot has been created for these VMs instead.`
Cause	VSS or hypervisor error
Impact	Recovery may take longer and involve manual steps.
Resolution	No action is necessary if this issue is intermittent. If this error persists, contact Nutanix support.
Severity	Warning

Table 347: VM Action Error [130057] [A1033]

Name	VM Action Status
Description	VM action error.
Alert message	`Failed to action VM with name 'vm_name and internal ID (ID vm_id)' due to reason.`
Cause	A VM could not be restored because of a hypervisor error, or could not be deleted because it is still in use.
Impact	The requested VM action (restore or delete) could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10658
Severity	Warning

Table 348: VM Virtual Hardware Version Not Compatible [130058] [A1133]

Name	VM Virtual Hardware Version Compatible
Description	Virtual hardware version of the VM is not compatible.
Alert message	`Virtual hardware version of the VM vm_name is not supported by any of the nodes on the remote site remote_name. Protection domain: protection_domain_name. Reason: reason.`
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote site.
Impact	VM may not register properly on the remote site on restore/clone.
Resolution	Upgrade the hypervisor version on the remote site to support virtual hardware version for the VM.
Severity	Warning

Table 349: Missing Network Mapping for the Remote Site. [130062] [A1156]

Name	Remote Site Network Configuration
Description	Network mapping not specified for the remote site.
Alert message	`No network mapping specified for the remote site remote_name.`
Cause	Network mapping not specified for the remote site.
Impact	When the VM is restored or cloned on the remote site, networking configuration may not be recovered.
Resolution	Specify network mapping for the remote site.
Severity	Warning

Table 350: Remote Site {remote_name} Network Mapping Invalid [130063] [A1157]

Name	Remote Site Network Mapping Configuration
Description	Invalid Network Mapping Specified.
Alert message	`Invalid network mapping specified for remote site remote_name: reason.`
Cause	Either the source or destination network configuration is not present in the network mapping specified on the remote site, or the networks specified in the network mapping do not exist.
Cause	The network associated with one or more VM(s) does not have the corresponding recovery network mapping in the remote site.
Impact	When the VM is restored or cloned on the remote site, networking configuration may not be recovered.
Resolution	Delete existing network mappings and create the network mapping with the associated networks present on the source and the destination cluster for the remote site.
Resolution	Update the network mapping of the remote site with all the network(s) associated with the VM(s).
KB Article	10243
Severity	Warning

Table 351: Protection Domain Snapshot Operation Skipped [130066] [A1147]

Name	PD Snapshot Skipped Status
Description	Skipped snapshot operation.
Alert message	`Replication was skipped for protection domain protection_domain_name to remote site remote_name. reason`
Cause	Snapshot operation was skipped.
Impact	Snapshot is not present on both the local and remote sites.
Resolution	Wait for the completion of the upgrade process before initiating snapshot schedules.

Severity	Info

Table 352: Volume Group Action Error [130071] [A1163]

Name	Volume Group Action Status
Description	Volume group action error.
Alert message	`Failed to action volume group with name 'vg_name' and internal ID 'vg_uuid' because reason in protection domain 'protection_domain_name'`
Cause	A volume group could not be restored or could not be deleted because it is still in use.
Impact	The requested volume group action (restore or delete) could not be completed.
Resolution	Detach the volume group from VMs and external initiators before recovery. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10120
Severity	Critical

Table 353: VSS Snapshot Failed [130073] [A130073]

Name	VSS Snapshot Status
Description	VSS snapshot failed.
Alert message	`VSS snapshot failed for the VM(s) vm_names protected by the protection_domain_name in the snapshot snapshot_id because reason.`
Cause	Guest is not able to quiesce VM due to internal error.
Impact	Crash consistent snapshot is taken instead of application consistent.
Resolution	Review the logs in the guest VM and take necessary actions. If issue persists contact Nutanix Support.
Severity	Warning

Table 354: Nutanix Guest Tools Not Installed [130074] [A130074]

Name	NGT Configuration
Description	Nutanix Guest Tools is not Installed.
Alert message	`VSS is enabled but Nutanix Guest Tools are not installed on the VM(s) vm_names protected by protection_domain_name.`
Cause	VSS is enabled but Nutanix Guest Tools are not installed on the VM(s).
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Install Nutanix Guest Tools on the VM(s).
Severity	Warning

Table 355: Protection Domain might have symlinks. [130077] [A130077]

Name	Snapshot Symlink Check
Description	Protection domain snapshot has symlinks.
Alert message	`Protection Domain protection_domain_name might have symlinks that have been skipped in the snapshot.`
Cause	The protection domain might have symlinks that have been skipped in the snapshot.
Impact	The symlinks in the protection domain have been skipped in the snapshot.
Resolution	Delete the symlinks or protect the entities individually.
Severity	Warning

Table 356: Entity Restore Aborted [130078] [A130078]

Name	EntityRestoreAbort
Description	Entity restore aborted.
Alert message	`Restoring VMs failed for the snapshot snapshot_id protected by the protection_domain_name because reason.`
Cause	Check the alert message for the reason of failure.
Impact	Snapshot expiry is removed to prevent data loss, and you should remove it manually.
Resolution	Restart the restore process. Snapshot expiry is removed to allow subsequent restore operations. Snapshot will need to be deleted manually to free up the space if required. If this error persists, contact Nutanix support.
Severity	Warning

Table 357: Remote Site Is Incompatible For Replication [130079] [A130079]

Name	Remote Site Snapshot Replication Status
Description	Skipped replication of the snapshot.
Alert message	`Replication was skipped for protection domain protection_domain_name to remote site remote_name. reason`
Cause	Snapshot replication was skipped because the remote is not fully upgraded.
Impact	Snapshot is not present on the remote site.
Resolution	(A) Wait for the completion of the upgrade process on the remote before initiating snapshot schedules. (B) Manually replicate the snapshot from Prism UI after the remote site upgrade process completes or wait for next scheduled snapshot to replicate data.
Severity	Info

Table 358: Nutanix Guest Tools Agent Is Not Reachable On The VM [130081] [A130081]

Name	VSS VM Reachable
Description	VSS on VM(s) is not reachable.
Alert message	`VM(s) vm_names Guest Agent Service is not reachable, protected by protection_domain_name.`
Cause	The communication link to the VMs Nutanix Guest Agent service seems to be down.
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Please check and restart Nutanix Guest Agent service inside the VM.
Severity	Warning

Table 359: Protected VM(s) Not Found [130083] [A130083]

Name	Protected VMs Not Found
Description	Protected VM(s) Not Found.
Alert message	`Unable to locate VM(s) vm_names protected by protection domain 'protection_domain_name'.`
Cause	Protected VM(s) cannot be found and may have been deleted.
Impact	Any data associated with the VMs may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Verify that the VM(s) are registered on the host. If the VM(s) were unregistered or deleted, remove VM(s) from the protection domain. If the CVM on the host the VM(s) are registered on was shutdown or rebooting, retry the snapshot operation.
Severity	Warning

Table 360: Protected Volume Groups Not Found [130084] [A130084]

Name	Protected Volume Groups Not Found
Description	Protected volume groups not found.
Alert message	`Unable to locate volume group(s) vg_names protected by protection domain 'protection_domain_name'.`
Cause	Protected volume group(s) cannot be found and may have been deleted.
Impact	Any data associated with the volume group(s) may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Remove volume group(s) from the protection domain.

Severity	Warning

Table 361: VSS Software or (pre_freeze/post_thaw) Scripts Not Installed [130085] [A130085]

Name	VSS Scripts Not Installed
Description	VSS software or pre_freeze/post_thaw scripts are not installed
Alert message	`VSS is enabled but VSS software or pre_freeze/post_thaw scripts are not installed on the guest VM(s) vm_names protected by protection_domain_name.`
Cause	VSS is enabled but VSS software or pre_freeze/post_thaw scripts are not installed on the VM(s).
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Install VSS software or (pre_freeze/post_thaw) scripts on the guest VM(s).
Severity	Warning

Table 362: VStore Snapshot Status [130086] [A130086]

Name	VStore Snapshot Status
Description	vStore snapshot status.
Alert message	`Snapshot status for vstore vstore_name: reason.`
Cause	Various
Impact	The requested VStore snapshot action could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 363: Failed To Snapshot Entities [130088] [A130088]

Name	Failed To Snapshot Entities
Description	Failed to snapshot entities.
Alert message	`Failed to snapshot the entities entity_names in the protection domain protection_domain_name, snapshot snapshot_id.`
Cause	Conflicting hypervisor tasks might be running on the entities.
Impact	The entities are not captured in the snapshot.
Resolution	Retry the operation after the hypervisor tasks are complete.
KB Article	8485
Severity	Warning

Table 364: Associated entities are not protected together. [130090] [A130090]

Name	Related Entity Protection Status
Description	Protection status of a related entity.
Alert message	*error_message*.
Cause	Related entity is not protected in the same protection domain.
Impact	Related VM/Volume Group will not be snapshotted and recovered.
Resolution	Protect the related entity in the same consistency group.
Severity	Warning

Table 365: Remote Site Operation Mode ReadOnly [130092] [A1189]

Name	Remote Site Operation Mode ReadOnly
Description	Operation Mode of Remote Site changed to kReadOnly.
Alert message	Remote site *remote_name* moved to read-only state, outgoing replications to this remote will fail.
Cause	SSD failure on remote site in case of single node or two-node cluster.
Cause	Cloud virtual appliance ran out of memory in case of cloud connect.
Cause	Data stored on cloud appliance exceeded 20TB.
Impact	All outgoing replications to remote site will be aborted.
Resolution	Replace the failed SSD on the remote site.
Resolution	Delete some remote snapshots to reduce memory or data usage in case of cloud connect and wait for some time.
Resolution	Increase the cloud CVM size. Refer to the Nutanix documentation for instructions.
KB Article	8082
Severity	Warning

Table 366: VM renamed on restoration [130094] [A130094]

Name	VM Renamed On Conversion
Description	VM Renamed On Restoration
Alert message	VM '*vm_name* (ID *vm_id*)' will be registered as '*new_vm_name*' during VM restore to avoid duplicate VM names.
Cause	Directory with same name as VM name exists on Storage Container.
Impact	The restored VM will be registered under the new name.
Resolution	No action is necessary.
Severity	Info

Table 367: Failed to reconfigure Nutanix Guest Tools during VMs recovery for protection domain [130095] [A130095]

Name	Failed To Recover NGT Information
Description	Failed to enable Nutanix Guest Tools during VMs recovery for protection domain.
Alert message	`Failed to reconfigure Nutanix Guest Tools for VMs in protection domain protection_domain_name due to reason.`
Cause	Nutanix Guest Tools service might be down.
Impact	VM
Resolution	Enable and mount Nutanix Guest Tools, and restart Nutanix Guest Agent service within the VM, for VMs in protection domain. If you cannot resolve the error, contact Nutanix support.
KB Article	10079
Severity	Warning

Table 368: Failed to reconfigure Nutanix Guest Tools for the recovered VM [130096] [A130096]

Name	Failed To Recover NGT Information for VM
Description	Failed to reconfigure Nutanix Guest Tools for a VM in protection domain.
Alert message	`Failed to reconfigure Nutanix Guest Tools for VM vm_name in protection domain protection_domain_name.`
Cause	Virtual IP address of the cluster might not have been configured.
Cause	Error in generating certificate for VM.
Cause	Guest VM information could not be retrieved.
Impact	VM
Resolution	Configure the virtual IP address of the cluster (if not configured).
Resolution	Enable and Mount Nutanix Guest Tools on the failed VM, and restart the Nutanix Guest Agent service within the VM.
Resolution	Resolve any issues on the Guest VM level to make sure NGT can communicate to the VM.
KB Article	10077
Severity	Warning

Table 369: Failed to reconfigure Nutanix Guest Tools for a VM in protection domain [130097] [A130097]

Name	Failed To Mount NGT ISO On Recovery of VM
Description	Failed to mount ISO image as part of Nutanix Guest Tools reconfiguration for a VM in protection domain

Alert message	`Failed to reconfigure Nutanix Guest Tools for VM vm_name in protection domain protection_domain_name due to ISO mount failure.`
Cause	Virtual IP address of the cluster might not have been configured.
Cause	Guest VM information could not be retrieved.
Cause	NGT ISO could not be mounted.
Impact	VM
Resolution	Configure the virtual IP address of the cluster (if not configured).
Resolution	Mount Nutanix Guest Tools ISO on the failed VM, and restart Nutanix Guest Agent service within the VM.
Resolution	Make sure there is at least one free CD-ROM slot on the Guest VM. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 370: External iSCSI Attachments Not Snapshotted [130099] [A130099]

Name	External iSCSI Attachments Not Snapshotted
Description	External iSCSI Attachments Not Snapshotted.
Alert message	`error_message. Volume Group external iSCSI attachments cannot be snapshotted for these VMs.`
Cause	VMs have duplicate IQNs.
Cause	VMs have configured iSCSI target IP addresses that do not belong to the local Nutanix cluster.
Cause	IQN could not be resolved to VM ID.
Impact	Volume group external iSCSI attachments cannot be snapshotted and restored for these VMs.
Resolution	Make sure all VMs in the protection domain have different IQNs.
Resolution	Remove iSCSI target IP addresses that do not belong to the local Nutanix cluster from the VMs.
Resolution	Install Nutanix Guest Tools on all the VMs using volume group external iSCSI attachments.
Severity	Warning

Table 371: Volume Group Attachments Not Restored [130101] [A130101]

Name	Volume Group Attachments Not Restored
Description	Volume Group Attachments Not Restored.
Alert message	`error_message for the protection domain 'protection_domain_name'.`
Cause	Acropolis failed to attach volume groups to the VMs.
Impact	VMs will not be able to access volume group iSCSI disks.
Resolution	Manually attach volume groups to the VMs.

Severity	Warning

Table 372: Self service restore operation failed [130102] [A130102]

Name	Self service restore operation Failed
Description	Self service restore operation failed
Alert message	`Unable to 'operation' disk from VM 'vm_name (ID vm_id)' due to 'reason'.`
Cause	(A) Snapshot disk could not be recovered. (B) Recovered disk could not be attached to the VM. (C) Self-service restore disk could not be detached from VM.
Impact	In the case of attachment failure, the user will not be able to restore files from snapshot.
Impact	In the case of detachment failure, the guest VM will have an extraneous disk attached.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 373: Witness is not configured [130114] [A130114]

Name	Witness Not Configured
Description	Witness is not configured
Alert message	`Failed to operation for the protection domain 'protection_domain_name' to the remote site 'remote_name'. Reason: Witness is not configured or a different Witness VM is configured on the remote site`
Cause	Witness is not configured
Impact	Metro availability operation failed
Resolution	Configure Witness
KB Article	10119
Severity	Critical

Table 374: Witness is not reachable [130115] [A130115]

Name	Witness Not Reachable
Description	Witness is not reachable
Alert message	`Failed to operation since unreachable_since for the protection domain 'protection_domain_name' to the remote site 'remote_name'. Reason: Witness is not reachable or the Witness credentials are incorrect`
Cause	Witness or the network is down.
Cause	Witness credentials are incorrect.

Impact	Metro availability failures will not be automatically handled
Resolution	Check if Witness is up on the network.
Resolution	Check Witness credentials.
KB Article	5483
Severity	Critical

Table 375: Metro Availability Is Promoted [130116] [A130116]

Name	Automatic Promote Metro Availability
Description	Metro availability is promoted.
Alert message	`Metro availability for the protection domain 'protection_domain_name' to the remote site 'remote_name' is promoted because of 'reason'.`
Cause	Remote site unreachable.
Impact	Metro availability is promoted on standby cluster and storage is now inaccessible in the active side.
Resolution	(A) Check if cluster service is healthy at remote cluster. (B) Ensure VMs are running first at the promoted side. In the case of storage only outage in order to restore VM availability follow the procedure in the Prism Web Console Guide, section: <Recovering from a Storage-only Failure on the Primary Site>. After that verify the previously active cluster is healthy before re-enabling metro availability.
KB Article	8279
Severity	Critical
Severity	Warning

Table 376: Failed to update Metro Availability failure handling [130117] [A130117]

Name	Updating Metro Failure Handling Failed
Description	Error in updating failure-handling on Metro Availability protection domain.
Alert message	`Error in updating Metro Availability failure handling for the protection domain 'protection_domain_name' to the remote site 'remote_name'. Reason: 'reason'.`
Cause	Failed to update Metro Availability Failure Handling due to possible network error
Impact	Metro Availability failure-handling could not be updated.
Resolution	Resolve network issues. If the issue persists contact Nutanix Support
KB Article	8480
Severity	Critical

Table 377: Failed to update Metro Availability failure handling on the remote site [130118] [A130118]

Name	Updating Metro Failure Handling Remote Failed
Description	Error in updating failure handling on the remote Metro Availability protection domain
Alert message	`Error in updating Metro Availability failure handling on the remote site 'remote_name' for the protection domain 'protection_domain_name'. Reason: 'reason'.`
Cause	Failed to update Metro Availability Failure Handling due to possible network error
Impact	Failure handling on this Metro Availability protection domain might be affected
Resolution	Resolve any network related issues. If issue persists contact Nutanix Support
KB Article	8481
Severity	Critical
Severity	Warning

Table 378: Authentication failed in Witness [130123] [A130123]

Name	Authentication Failed in Witness
Description	Authentication failed in Witness
Alert message	`Failed to operation for the protection domain 'protection_domain_name' to the remote site 'remote_name'. Reason: Authentication failure in Witness`
Cause	Witness configured with incorrect credentials
Impact	Metro availability failures will not be automatically handled
Resolution	Configure Witness with correct credentials
KB Article	4376
Severity	Critical

Table 379: Metro Availability Is Promoted [130125] [A130125]

Name	Manual Promote Metro Availability
Description	Metro availability is promoted.
Alert message	`Metro availability for the protection domain 'protection_domain_name' to the remote site 'remote_name' is promoted.`
Cause	An administrator promoted metro availability.
Impact	Metro availability is promoted on standby site.
Resolution	Re-enable to get the benefit of Metro protection domain.
Severity	Info

Table 380: Metro Availability Is Disabled [130126] [A130126]

Name	Manual Break Metro Availability
Description	Metro availability is disabled.
Alert message	`Metro availability for the protection domain 'protection_domain_name' to the remote site 'remote_name' is disabled.`
Cause	An administrator disabled metro availability.
Impact	Metro availability operation is disabled.
Resolution	Re-enable Metro on the protection domain to get the benefit of Metro availability
Severity	Info

Table 381: VMs in the standby site of a Metro Availability protection domain are running at suboptimal performance. [130129] [A130129]

Name	Metro Protection domain VMs running at Sub-optimal performance
Description	VMs in the standby site of a Metro Availability protection domain are running at suboptimal performance.
Alert message	`Warning: num_of_vms VMs are running on the standby site 'remote_name' of protection domain 'protection_domain_name'. However, all the I/O operations are being forwarded to this site. This can significantly affect the performance of VMs. It is recommended to move VMs to this site or update the DRS rules. The affected VMs list is: vm_names`
Cause	VMs are hosted at the standby site of the Metro Availability protection domain.
Impact	Experience of decreased performance
Resolution	Host the VMs at the active site of the Metro Availability protection domain.
Severity	Warning

Table 382: Protection domain contains more than the specified VMs [130130] [A130130]

Name	Protection Domain VM Count Check
Description	Protection domain contains multiple VMs.
Alert message	`It is not recommended to protect more than one VM in a protection domain protection_domain_name if you are replicating to the backup only remote site remote_name.`
Cause	Replicating a protection domain with multiple VMs to a backup only remote site.
Impact	Selective retrieval of any VM will take more time as the snapshot will capture all the VMs in the protection domain.
Impact	Replications may be skipped because of a large snapshot size.
Resolution	Protect each VM in a separate protection domain.

Severity	Info

Table 383: Snapshot contains entities from the storage container that have deduplication enabled. [130131] [A130131]

Name	Replication Of Deduped Entity
Description	Snapshot contains entities from the storage container that have deduplication enabled.
Alert message	`Do not replicate protection domain protection_domain_name comprising entities from the storage container that have deduplication enabled to a remote_site_type remote site remote_name.`
Cause	Replicating a snapshot with entities from the storage container that have deduplication enabled to a single node backup or cloud remote site.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable dedup on the Storage Containers, or use a multi-node AOS cluster as a replication target.
Severity	Warning

Table 384: Cloud Appliance Deployment Failed [130133] [A130133]

Name	Cloud Appliance Deployment Failed
Description	Cloud Appliance Deployment Failed
Alert message	`Failed to configure cloud appliance for remote site remote_site_name. reason.`
Cause	Resource quota for buckets/instances/storage accounts exhausted for cloud account.
Cause	AWS credentials do not have permissions to perform S3/SG/EC2 related operations.
Cause	Cloud CVM deployed was not able to reach the Internet.
Impact	Cloud appliance could not be configured as remote site.
Resolution	Make sure that account has enough quota to create new buckets/instances.
Resolution	Make sure that AWS IAM user has permissions to perform S3/SG/EC2 operations.
Resolution	Make sure that the subnet in which cloud appliance is being deployed allows Internet connectivity.
Severity	Warning

Table 385: vStore is being replicated to backup only remote site. [130134] [A130134]

Name	Vstore Replication To Backup Only Remote
Description	vStore is being replicated to backup only remote site.

Alert message	`vStore protection domain protection_domain_name is being replicated to backup only remote site remote_name.`
Cause	vStore is being replicated to backup only remote site.
Impact	Low RTO can be seen as selective retrieve of VMs is not supported and snapshot of whole protection domain, having all the VMs in it, has to be retrieved.
Resolution	One VM per protection domain is the recommended configuration to achieve higher RTO.
Severity	Info

Table 386: Replication of protection domain {protection_domain_name} has not progressed. [130137] [A130137]

Name	Protection Domain Replication Stuck
Description	Replication of protection domain has not progressed.
Alert message	`Replication of protection domain protection_domain_name to remote site remote_name for the snapshot snapshot_id has not made any progress in last one hour.`
Cause	Network outage with remote.
Cause	Remote cluster does not have enough free disk space available.
Cause	Single-node cluster might be in read-only mode, due to one of the metadata SSDs failure.
Impact	Snapshots of protection domain will not be available on the remote site.
Resolution	Check network connection with remote.
Resolution	Make sure that the remote cluster has enough disk space available.
Resolution	Replace the failed SSD.
Severity	Warning

Table 387: Snapshot Reserve on SSD is Full. [130139] [A130139]

Name	LWS Store Full
Description	SSD reserved for the Nearsync snapshots is full.
Alert message	`The SSD space that is reserved for the snapshots is full.`
Cause	Snapshot reserve on the SSD is full or has high IO workload.
Impact	RPO cannot be met and will be temporarily transitioned to 60 minutes.
Resolution	Reduce the retention time of the snapshots.
Severity	Warning

Table 388: Snapshot Reserve Allocation Failure [130140] [A130140]

Name	LWS store allocation too long
Description	Failed to allocate the snapshot reserve.
Alert message	`Failed to allocate the snapshot reserve in duration_mins minutes.`
Cause	Snapshot reservation timed out.
Impact	Protection domain will not be able to meet less than 15 minutes RPO.
Resolution	Check SSD performance from Prism Web console and if issue persists contact Nutanix support
Severity	Warning

Table 389: Failed to allocate snapshot reserve on the remote site [130141] [A130141]

Name	LWS store allocation in remote too long
Description	Failed to allocate snapshot reserve on the remote site.
Alert message	`Failed to allocate snapshot reserve on the remote site 'remote_name' in duration_mins minutes.`
Cause	Snapshot reserve allocation on the remote site timed out.
Impact	Failed to meet less than 15 minute RPO for the remote site.
Resolution	Contact Nutanix support
Severity	Warning

Table 390: Protection Domain transitioning to lower frequency snapshotting. [130142] [A130142]

Name	Protection Domain transitioning to lower frequency snapshot schedule
Description	Protection Domain transitioning to lower frequency snapshot schedule.
Alert message	`Protection Domain 'protection_domain_name' has degraded to 'rpo_string' RPO. reason.`
Cause	High RPO cannot be sustained due to system constraints.
Cause	New entity has been protected in the protection domain which needs to be seeded to the remote.
Impact	All ongoing high frequency replications will be aborted and RPO for data protection will be affected.
Resolution	Protection Domain will automatically attempt to transition to a higher RPO after system constraints are resolved.
Resolution	Protection Domain will automatically attempt to transition to a higher RPO after seeding for newly added entities completes.
Severity	Warning
Severity	Info

Table 391: Protection Domain falling back to lower frequency replications to remote. [130143] [A130143]

Name	Protection Domain fallback to lower frequency replications to remote
Description	Protection Domain fallback to lower frequency replications to remote.
Alert message	`Protection Domain 'protection_domain_name' has degraded to 'rpo_string' RPO to remote 'remote_name'. reason.`
Cause	High frequency replications to remote site cannot be sustained due to system constraints.
Impact	Recovery Points available on the remote for Data Protection will be affected
Resolution	Protection Domain will automatically attempt to transition to higher frequency replications when resources are available.
Severity	Warning

Table 392: Protection Domain Transition to Higher Frequency Snapshotting. [130144] [A130144]

Name	Protection Domain transitioning to higher frequency snapshot schedule
Description	Protection domain transitioning back to the configured RPO.
Alert message	`Protection Domain 'protection_domain_name' has been upgraded to rpo_string RPO.`
Cause	High frequency snapshot schedule has been enabled.
Impact	RPO for Data Protection will be met.
Resolution	Reduce snapshot schedule frequency or delete the schedule if high frequency snapshots are not needed.
Severity	Info

Table 393: Snapshot queued for replications to remote site [130145] [A130145]

Name	Protection Domain Snapshots Queued for Replication to Remote Site
Description	Protection domain {protection_domain_name} has one or more snapshots queued for replication to remote site {remote_name}
Alert message	`Protection domain protection_domain_name has one or more snapshots queued for replication to remote site remote_name`
Cause	Various
Impact	One or more snapshots cannot be replicated to the remote site to meet the configured recovery point objective
Resolution	Maintain enough cross-cluster network bandwidth and throttle primary IO
Severity	Warning

Table 394: Protection Domain {protection_domain_name} Seeding Process Pending to the Remote Site {remote_name} [130146] [A130146]

Name	Protection Domain Seeding Process Pending to Remote Site
Description	Protection Domain {protection_domain_name} will maintain {rpo_string} RPO until seeding process completes to the remote_site {remote_name}
Alert message	`The Protection Domain protection_domain_name will maintain rpo_string RPO until seeding process completes to the remote_site remote_name. The seeding process gets completed when snapshot reserve has been allocated on the local and remote site`
Cause	Various
Impact	The Protection Domain {protection_domain_name} will maintain {rpo_string} RPO until seeding process completes to the remote_site {remote_name}
Resolution	Maintain enough network bandwidth to help the seeding transfer to succeed. Protection domain will wait for the seeding transfer and then automatically enable the configured RPO.
Severity	Warning

Table 395: VM registration failed [130147] [A130147]

Name	VM Registration Failure
Description	VM Registration Failure
Alert message	`Registration of virtual machine, vm_name, failed because reason`
Cause	This might happen for several reasons. Check the alert message for more details.
Impact	VMs cannot be registered on the cluster.
Resolution	Verify the existence of the VM in the snapshot and check whether any errors are reported in the hypervisor management software. Check the VM version and the maximum VM version supported by the host. If you cannot resolve the error, contact Nutanix support
Severity	Critical
Severity	Warning

Table 396: Snapshot Creation Delayed [130148] [A130148]

Name	Protection Domain Snapshots Delayed
Description	Snapshot creation for the protection domain {protection_domain_name} has been delayed
Alert message	`Snapshot creation for the protection domain protection_domain_name has been delayed.`
Cause	Many entities might be scheduled for the snapshot creation at the same time.
Impact	One or more snapshots cannot be created and hence cannot be replicated to the remote sites to meet the configured recovery point objective

Resolution	Stagger the schedules for different protected entities. Additionally, you can throttle other administrative operations in the cluster.
Severity	Warning

Table 397: VM Recovery Point Creation Failed. [130157] [A130157]

Name	VM Recovery Point Failed
Description	VM Recovery Point creation failed.
Alert message	`Failed to capture the Recovery Point for VM 'vm_name'.`
Cause	Nutanix Data Protection service error.
Impact	VM Recovery Point will not be available to recover the VM.
Resolution	Retry the VM Recovery Point operation. If issue persists please contact Nutanix Support.
Severity	Warning

Table 398: VM Protection Failed. [130158] [A130158]

Name	VM Protection Failed
Description	VM protection failed.
Alert message	`Failed to protect VM: 'vm_name', because of 'reason'.`
Cause	Nutanix Data Protection service error.
Impact	VM cannot be recovered in case of disaster.
Resolution	VM protection will get retriggered. In case if this situation persists, please contact nutanix support.
KB Article	8281
Severity	Warning

Table 399: Protection Rule Conflict Occurred. [130161] [A130161]

Name	Protection Rule Conflicts
Description	Unable to protect VM due to conflicting protection rules.
Alert message	`Failed to apply the protection rule on the VM 'vm_name', because of 'reason'.`
Cause	One or more protection rules protecting the same entity.
Impact	VM cannot be recovered if disaster occurs.
Resolution	Resolve all the conflicts between the protection rules or modify the VM to ensure that it is protected under one protection rule.
KB Article	8275
Severity	Warning

Table 400: Recovery Point Replication Skipped. [130162] [A130162]

Name	Recovery Point Replication Skipped
Description	Recovery Point Replication Skipped.
Alert message	`Skipped replication of Recovery point created at 'recovery_point_create_time' UTC of the VM: 'vm_name' because of 'reason'.`
Cause	Latest recovery point is in the queue for the replication.
Impact	One of the recovery point replication will be skipped.
Resolution	Trigger the replication of the snapshot manually.
Severity	Warning

Table 401: Failed to find the Target cluster for Replication [130163] [A130163]

Name	Fetch Replication Target Failure
Description	Cannot find a target cluster for replication on the Remote Availability Zone.
Alert message	`Cannot find the replication target on the Recovery Location 'availability_zone_physical_name'.`
Cause	None of the attached clusters support recovery.
Cause	Virtual IP is not set on the remote cluster.
Cause	Target cluster could not be found for replication due to some internal error.
Impact	Replication of recovery points to this remote Availability Zone will be affected.
Resolution	Ensure that at least one of the clusters, attached to the Remote Availability Zone, supports recovery.
Resolution	Configure the virtual IP address on the remote cluster.
Resolution	Retry the operation. Contact Nutanix support if the issue persists.
Severity	Warning

Table 402: VM Replication Failure [130164] [A130164]

Name	VM Replication Failure
Description	VM Replication Failure
Alert message	`Replication of VM Recovery Point created at recovery_point_create_time UTC for VM vm_name to the Recovery Location availability_zone_physical_name failed.`
Cause	(A) Remote cluster is in read only state. (B) Remote Availability Zone does not support PowerPC architecture. (C) Replication of Recovery Point with dedup entity to a Recovery Availability Zone without dedup support is not allowed. (D) Remote Availability Zone does not support Volume Groups. (E) Replication of Recovery Point with VMs having attached File Level Restore (FLR) disks to a Recovery Availability Zone without FLR support is not allowed.

Impact	VM Recovery Point was not replicated to the Recovery Location.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 403: Application-Consistent Recovery Point Failed [130165] [A130165]

Name	VSS Recovery Point Status
Description	Application-consistent Recovery Point failed.
Alert message	`Application-Consistent Recovery Point created at recovery_point_create_time UTC for the VM vm_name failed because reason.`
Cause	Failed to quiesce the VM.
Impact	Crash-consistent Recovery Point is taken instead of application-consistent.
Resolution	Review the logs in the guest VM and take necessary actions. If issue persists contact Nutanix Support.
Severity	Warning

Table 404: Recovery Point expired prior to start of replication. [130166] [A130166]

Name	VM Replication Expiry Status
Description	Recovery Point expired prior to start of replication.
Alert message	`Replication of Recovery Point created at recovery_point_create_time UTC for VM vm_name to the recovery location az_name expired prior to start of replication.`
Cause	Previous Recovery Point replication took longer than RPO.
Impact	The Recovery Point that was scheduled to be replicated could not be replicated. Data that is expected to be protected may not be protected.
Resolution	Review the RPO and retention time of the Recovery Point based on available bandwidth.
Severity	Warning

Table 405: VM Recovery Point Replication Time Exceeded the RPO [130167] [A130167]

Name	VM Recovery Point Replication Time Exceeds RPO
Description	VM Recovery Point Replication Time Exceeded the RPO Limit.
Alert message	`Replication time of the snapshot created at 'recovery_point_create_time' UTC for VM vm_name exceeded the RPO limits.`
Cause	Latest replication took more time than the RPO limit.
Impact	Replications for the next snapshots will be delayed.
Resolution	Please check your network.

Severity	Warning

Table 406: NGT on VM is not reachable [130168] [A130168]

Name	NGT is unreachable
Description	NGT on the VM is not reachable.
Alert message	`Communication link to Nutanix Guest Tools on the VM vm_name is not working.`
Cause	Communication link to Nutanix Guest Tools on the VM is not working.
Impact	Crash consistent snapshot will be taken instead of application consistent snapshot and static IP address preservation/mapping will fail.
Resolution	Restart the Nutanix Guest Tools service in the VM.
Severity	Warning

Table 407: VM Replication has not progressed. [130169] [A130169]

Name	VM Recovery Point Replication Stuck
Description	Replication of the VM has not progressed.
Alert message	`Replication of the Recovery Point created at recovery_point_create_time UTC for VM vm_name to Recovery Location az_name has not made any progress in last one hour.`
Cause	Network Connectivity issues between the Source and the recovery location.
Cause	Recovery Location might not have enough free disk space available.
Cause	Single-node cluster might be in read-only mode, due to an SSD failure.
Impact	VM Recovery Point will not be available at the Recovery Location.
Resolution	Check network connection between the Source and the Recovery Location.
Resolution	Make sure that the Recovery Location has enough disk space available.
Resolution	Replace the failed SSD.
Severity	Warning

Table 408: Nutanix Guest Tools Not Installed on the VM [130170] [A130170]

Name	VM NGT Configuration
Description	Nutanix Guest Tools not installed.
Alert message	`Nutanix VSS is enabled but Nutanix Guest Tools is not installed on the VM vm_name.`
Cause	Nutanix VSS is enabled but Nutanix Guest Tools is not installed on the VM.
Impact	Application Consistent Recovery Points will fail. Instead Crash consistent recovery point will be taken.

Resolution	Install Nutanix Guest Tools on the guest VM.
Severity	Warning

Table 409: Nutanix VSS provider or (pre_freeze/post_thaw) Scripts Not Installed. [130171] [A130171]

Name	VSS Scripts Not Installed On VM
Description	Nutanix VSS provider or pre_freeze/post_thaw Scripts Not Installed.
Alert message	`VSS is enabled, but Nutanix VSS provider or pre_freeze/post_thaw scripts are not installed on the VM vm_name.`
Cause	VSS is enabled, but Nutanix VSS provider or pre_freeze/post_thaw scripts are not installed on the VM.
Impact	Crash-consistent Recovery Point will be taken instead of application-consistent Recovery Point.
Resolution	Install Nutanix VSS provider or (pre_freeze/post_thaw) scripts on the VM.
Severity	Warning

Table 410: Degraded VM Recovery Point. [130172] [A130172]

Name	Partial VM Recovery Point
Description	Partial VM Recovery Point.
Alert message	`Recovery Point for VM vm_name failed to capture associated policies and categories because reason.`
Cause	Recovery Point for VM failed to capture associated policies and categories.
Impact	VM will be restored without associated policies and categories.
Resolution	Check the connection between PC and the PE, then retry the Recovery Point operation.
Severity	Info

Table 411: Recovery Point Replication Failed [130174] [A130174]

Name	Recovery Point Replication Failed
Description	Recovery Point replication failed.
Alert message	`Failed to replicate recovery point created at: 'recovery_point_create_time' UTC of the VM: 'vm_name' to the recovery location: availability_zone_physical_name.`
Cause	Network Connectivity issues between the Primary and the Recovery Availability Zone.
Cause	Data Protection and Replication service is not working as expected. The service could be down.

Cause	VM migration process is in progress.
Cause	If NearSync is enabled, the container name on remote cluster should match the one on source cluster hosting the VM.
Cause	Virtual IP address has not been configured on the remote cluster.
Cause	Remote clusters are unhealthy.
Cause	Recovery point was created for an entity having datasource backed virtual disks.
Cause	(A) Remote cluster is in read only state. (B) Remote Availability Zone does not support PowerPC architecture. (C) Replication of Recovery Point with dedup entity to a Recovery Availability Zone without dedup support is not allowed. (D) Remote Availability Zone does not support Volume Groups. (E) Replication of Recovery Point with VMs having attached File Level Restore (FLR) disks to a Recovery Availability Zone without FLR support is not allowed.
Cause	Replication to the remote site have been suspended for an entity by the user through CLI.
Impact	Recovery Point will not be replicated to the recovery location.
Resolution	Check network connection between the Primary and the Recovery Availability Zone.
Resolution	Contact Nutanix support.
Resolution	Retry the Recovery Point replication operation after the migration is complete.
Resolution	Create container(s) on which VM is running with the same name on the remote cluster.
Resolution	Configure the Virtual IP address and then retry the Recovery Point replication operation.
Resolution	For a manually taken Recovery Point, retry replication again. For a scheduled Recovery Point replication, ensure all the emote clusters are healthy. Then wait for the scheduled Recovery Point replication.
Resolution	Retry the Recovery Point replication operation after storage migration for the datasource backed virtual disks is complete.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Resolution	Remove the suspended replication timer associated with the entity using the CLI.
KB Article	8460
Severity	Warning

Table 412: Invalid Availability Zone details configured in Protection Policy. [130177] [A130177]

Name	Availability Zone Validation Failed
Description	Availability Zone referenced in the Protection Policy is not Valid.
Alert message	**Protection Policy *protection_rule_name* has invalid Availability Zone configured.**

Cause	Invalid Availability Zone configured in the Protection Policy.
Impact	Protection Policy will not be processed.
Resolution	Update the Availability Zone in Protection Policy.
Severity	Warning

Table 413: Entity Unprotection Failed [130178] [A130178]

Name	Entity Unprotection Failed
Description	Unable to unprotect the VM because of some internal error.
Alert message	`Failed to unprotect the VM: 'vm_name', because of 'reason'.`
Cause	Data Protection service may be down.
Impact	VM will remain protected.
Resolution	If this issue persists for a long time, contact Nutanix support for further assistance.
Severity	Warning

Table 414: Data Protection tasks are not progressing. [130179] [A130179]

Name	Cerebro tasks Stuck
Description	Replication tasks are not progressing.
Alert message	`Multiple tasks are not making progress on PE:cluster_name because of long running replication tasks.`
Cause	Replication tasks are running for long.
Impact	Tasks can stuck forever.
Resolution	Please reach out to Nutanix Support to help abort the long running tasks.
Severity	Warning

Table 415: Recovery Location is in Read Only mode [130180] [A130180]

Name	Entity Centric Remote Site Operation Mode ReadOnly
Description	Operation Mode of the Recovery Location changed to Read Only mode.
Alert message	`Availability zone: availability_zone_physical_name is in read-only state. Replications to this site will fail.`
Cause	Single node or two-node clusters: Possible SSD failure on the Recovery Location.
Impact	Replications to this Recovery Location will be aborted.
Resolution	Single node or two-node clusters: Replace the failed SSD on the recovery location.
KB Article	8324

Severity	Warning

Table 416: Snapshot has Entities from the Storage Container that have Deduplication Enabled [130181] [A130181]

Name	Replication Of Deduped Entity
Description	Snapshot contains entities from the storage container that have deduplication enabled.
Alert message	`Do not replicate the recovery point comprising entities from the storage container that have deduplication enabled to the recovery location : availability_zone_physical_name.`
Cause	A snapshot with entities from the storage container that have deduplication enabled was replicated to a single node backup or cloud recovery location.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable the deduplication on the storage containers or use a multi-node AOS cluster as a replication target.
Severity	Warning

Table 417: Nutanix Guest Tools not Upgraded [130182] [A130182]

Name	NGT needs upgrade
Description	Nutanix Guest Tools not upgraded.
Alert message	`Nutanix Guest Tools are not upgraded to version version or higher on the guest VM(s) vm_names protected by protection_domain_name.`
Cause	VSS is enabled, but Nutanix Guest Tools are not upgraded on the VM(s).
Impact	Crash-consistent snapshot will be taken instead of application consistent snapshot.
Resolution	Install the latest version of Nutanix Guest Tools available on the cluster on the guest VM(s).
Severity	Warning

Table 418: Protected VM(s) May Not Be Recoverable [130183] [A130183]

Name	Protected VMs May Not Be Recoverable
Description	Protected VM(s) May Not Be Recoverable
Alert message	`VM(s) vm_names protected by protection domain 'protection_domain_name' are running from VMWare snapshots. Recovery of the VMs from the PD snapshot 'snapshot_id' may not succeed. Manual intervention is required to boot up the VMs.`
Cause	Protected VM(s) are running from VMWare snapshots and may not be recoverable.

Impact	Recovery of the VMs from the PD snapshot may not succeed
Resolution	Ensure that the VMs are not running from VMWare snapshot at the time of scheduled Protection Domain snapshot. Refer Nutanix KB: 7023 for further details.
Severity	Warning

Table 419: VM Protection Failed [130186] [A130186]

Name	VM Protection Failed
Description	Failed to protect the VM by using the defined Protection Policy.
Alert message	`Failed to apply protection policy 'protection_rule_name' on the VM 'vm_name' because 'reason'`
Cause	VM is already protected in a legacy protection domain.
Cause	Cluster hosting the VM does not support VM in multiple schedules or replication to more than one target site.
Cause	(A) VM has volume group attachments. (B) VM is CBR incapable. (C) Source or the target cluster does not support replication within the same Availability Zone.
Impact	VM will not be protected.
Resolution	Unprotect the VM from the Protection Domain to protect it using Entity Centric Workflow (Leap).
Resolution	Upgrade cluster to version 5.19 or above.
Resolution	Resolve the stated reason for the failure.
KB Article	8472
Severity	Warning

Table 420: VM Recovery May Fail [130188] [A130188]

Name	VM Recovery May Fail
Description	VM may not start on the recovery location.
Alert message	`VM vm_name may not start on the recovery location 'availability_zone_name'`
Cause	Primary and Remote locations are on different hypervisors and VM does not have NGT information.
Cause	VM uses vNUMA. Cannot recover vNUMA VMs on Xi.
Impact	VM will not start on the recovery location.
Resolution	Install NGT on the VM (if not already installed).
Resolution	Don't replicate VMs which use vNUMA to Xi.
Severity	Warning

Table 421: Virtual IP Address not Configured on the Cluster. [130191] [A130191]

Name	Virtual IP not set on Cluster
Description	Virtual IP address is not configured on the cluster.
Alert message	`Virtual IP address is not configured on the cluster cluster_name. Protection and replication of VMs will be affected.`
Cause	Virtual IP address has not been configured on the cluster.
Impact	Protection and replication of VMs will be affected.
Resolution	Configure the virtual IP address on the cluster.
Severity	Warning

Table 422: Full Replication Started for the VM [130193] [A130193]

Name	VM Recovery Point Full Replication
Description	A full replication was triggered for the VM.
Alert message	`Replication for the VM vm_name recovery point with the create time recovery_point_create_time to recovery location az_name will be a full replication.`
Cause	A reference recovery point could not be found.
Impact	Replication will take longer to finish.
Resolution	No action is necessary.
Severity	Info

Table 423: Protected VM Not Recoverable [130195] [A130195]

Name	VM Not Recoverable
Description	Protected VM may not be recoverable.
Alert message	`VM vm_name is running from the VMware snapshots. Recovery of the VM from the Recovery Point created at recovery_point_create_time UTC may not succeed. Manual intervention is required to boot up the VM.`
Cause	Protected VM is running from the VMware snapshot and may not be recoverable.
Impact	Recovery of the VM from the Recovery Point may not succeed.
Resolution	Ensure that the VM is not running from the VMware snapshot at the time of scheduled Recovery Point.
Severity	Warning

Table 424: VM Virtual Hardware Version Incompatible [130196] [A130196]

Name	VM Virtual Hardware Version Incompatible

Description	VM Virtual Hardware Version Incompatible
Alert message	`Virtual hardware version of the VM vm_name is not supported on the availability zone availability_zone_name.`
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote Availability Zone.
Impact	VM may not register properly on the remote Availability Zone on restore/ clone.
Resolution	Upgrade the hypervisor version on the remote Availability Zone to support the virtual hardware version of the VM.
Severity	Warning

Table 425: VM Restoration Failed. [130197] [A130197]

Name	VM Restoration Failed
Description	VM Restoration Failed.
Alert message	`Failed to restore VM: 'vm_name' due to reason.`
Cause	Snapshot is not present.
Cause	VM Restoration failed due to Internal error.
Impact	VM will not be restored.
Resolution	Retry the VM Restoration operation with different snapshot.
Resolution	Retry the VM Restoration operation.
KB Article	8487
Severity	Critical

Table 426: VM Migration Failed [130199] [A130199]

Name	VM Migration Failed
Description	VM Migration Failed.
Alert message	`Failed to migrate VM: 'vm_name' as the user may have manually powered ON the VM on the source after migration process powered it off or due to an internal error.`
Cause	User may have manually powered ON the VM on the source after migration process powered it off.
Impact	VM will not be migrated.
Resolution	Retry the VM Migration operation. If this situation persists, contact Nutanix support.
KB Article	8520
Severity	Critical

Table 427: VSS Snapshot is not supported for some VMs. [130200] [A130200]

Name	VSS Snapshot Not Supported
Description	VSS snapshot is not supported for some VMs.
Alert message	`VSS snapshot is not supported for the following VMs protected by Protection Domain 'protection_domain_name'. VMs list: 'vm_names'`
Cause	Some VMs in the Protected Domain have unsupported configuration.
Impact	Hypervisor based application consistent snapshot is taken on ESX. Crash consistent snapshot is taken for other hypervisors.
Resolution	For more details, please review the alerts generated at Data Protection page and fix the unsupported configuration.
Severity	Warning

Table 428: Recovery Location is not in Good State [130203] [A130203]

Name	Recovery Location Not In Good State
Description	Recovery location is not in good state.
Alert message	`Recovery location: availability_zone_physical_name is not in good state. Replications to this site will fail.`
Cause	If Recovery Availability Zone is different from the Primary Availability Zone, check Network connectivity issues between the two Availability Zones.
Cause	Data Protection and Replication service is not working as expected. The service could be down.
Cause	The Cluster(s) attached to the affected Availability Zone may have been destroyed and recreated.
Impact	Replications to this recovery location will fail.
Resolution	Check network connection between the Primary and the Recovery Availability Zone.
Resolution	Contact Nutanix support.
Resolution	Disconnect and reconnect the Recovery Availability Zone.
Severity	Warning

Table 429: Entity Transition to a Higher Frequency Snapshot Schedule. [130204] [A130204]

Name	Entity transitioning to higher frequency snapshot schedule
Description	Entity is being transitioned to a higher frequency snapshot schedule.
Alert message	`Entity 'vm_name' has been upgraded to rpo_string RPO.`
Cause	High frequency snapshot schedule has been enabled and higher RPO can be sustained as the system is stable.
Impact	RPO for Data Protection will be affected.

Resolution	If high-frequency snapshots are not required, reduce the snapshot schedule frequency in the Protection Rule or remove the VM from the Protection Rule with high frequency schedule.
Severity	Info

Table 430: Entity is being transitioned to a lower frequency snapshot schedule. [130205] [A130205]

Name	Entity transitioning to lower frequency snapshot schedule
Description	Entity is being transitioned to a lower frequency snapshot schedule.
Alert message	`Entity vm_name has degraded to rpo_string RPO. reason.`
Cause	High RPO cannot be sustained due to system constraints.
Impact	All ongoing high frequency replications will be aborted and RPO for data protection will be affected.
Resolution	Entity will automatically attempt to transition to a higher RPO after system constraints are resolved.
Severity	Warning

Table 431: Subnet configured for Remote Site mis-matches with remote cluster subnet [130206] [A130206]

Name	Remote Subnet Mismatch With Provided Value
Description	Subnet configured for Remote Site mis-matches with remote cluster subnet.
Alert message	`Remote remote_name is configured with the subnet user_specified_subnet, which is not matching with the subnet remote_returned_subnet of the remote cluster.`
Cause	Subnet specified in the remote site is not matching with the subnet of the remote cluster.
Impact	Improper firewall and routing settings.
Resolution	Reconfigure the remote subnet parameter.
KB Article	10276
Severity	Warning

Table 432: Hosting of Virtual IP of the Network Segmentation DR Service Failed. [130207] [A130207]

Name	Failed Hosting Network Segmented VIP
Description	Hosting of virtual IP of the DR service network segmentation failed.
Alert message	`Virtual IP virtual_ip of the DR service network segmentation could not be hosted on the network interface network_interface.`

Cause	Network configuration might have been changed or has not been configured properly.
Impact	Replication of the protected entities will be affected.
Resolution	Reconfigure the DR service network configuration.
KB Article	8258
Severity	Critical

Table 433: Snapshot Replication to Remote Site is Lagging. [130210] [A130210]

Name	Recovery Point Lagging Remote Replication
Description	Replication of the snapshot to the remote cluster is lagging.
Alert message	`Replication of the recovery point for the VM 'vm_name'to the remote site 'remote_az' is lagging by 'lagging_minutes' minutes.`
Cause	Clusters are not sized properly according to the workloads that are running on them.
Cause	Recovery Point Replication requires more bandwidth than the allocated bandwidth.
Impact	Recovery point objective for the protection domain will be affected.
Resolution	Make sure that the clusters are sized properly according to the workload.
Resolution	Maintain sufficient network bandwidth for successful seeding transfer.
Severity	Warning
Severity	Info

Table 434: VM has been Recovered at an Alternate Path [130211] [A130211]

Name	VM Recovered at Alternate Path
Description	VM has been Recovered at an Alternate Path
Alert message	`The following files of the VM old_vm_name conflict with the files present on the target: conflicting_file_paths. The VM new_vm_name has been restored at an alternate path by adding the prefix 'path_prefix' after the container name to its file paths.`
Cause	Existing VM files conflict with the VM files that you are recovering.
Impact	The VM with the conflicting files has been restored at an alternate path.
Resolution	The VM containing the conflicting files has been recovered at an alternate path. Clean up the conflicting stale files.
Severity	Warning

Table 435: Cluster does not support Synchronous Replication. [130212] [A130212]

Name	Cluster does not support Synchronous Replication

Description	Cluster does not support Synchronous Replication.
Alert message	`Protection Rule: 'protection_rule_name' is processed but following issue is observed - 'reason'.`
Cause	The cluster given in the Protection Rule does not support Synchronous Replication.
Impact	Protection Rule might not be processed.
Resolution	Please use a cluster that supports Synchronous Replication.
Severity	Critical

Table 436: Failed to transition to a higher frequency snapshot schedule. [130213] [A130213]

Name	Failed to transition to higher frequency snapshot schedule
Description	Protection Domain failed to transition to a higher frequency snapshot schedule.
Alert message	`Failed to enable nearsync for protection_domain_name. reason.`
Cause	Low RPO cannot be sustained due to system constraints.
Impact	Failed to enable high frequency replications and configured RPO for data protection will be affected.
Resolution	Protection Domain will automatically attempt to transition to a high frequency snapshot schedule after system constraints are resolved.
Severity	Warning

Table 437: Nearsync Replication is lagging for Protection Domain snapshot [130214] [A130214]

Name	Nearsync Replication is Lagging
Description	Replication of the nearsync protection domain snapshot is lagging.
Alert message	`Nearsync replication for protection_domain_name to remote remote_name is lagging by replication_lagging_seconds seconds.`
Cause	Replication of the protection domain nearsync snapshot is slow.
Impact	Configured RPO to the remote might get affected.
Resolution	Ensure that the number of entities are within the supported limit.
Severity	Warning

Table 438: VM SyncRep Container not found [130335] [A130335]

Name	VM SyncRep Container not found
Description	VM SyncRep container is not found at the target cluster
Alert message	`Failed to replicate the VM: 'vm_name' synchronously since 'details'.`
Cause	Container does not exist or it is marked to be removed at the target cluster.

Impact	VM will not be synchronously replicated.
Resolution	Retry the VM SyncRep operation after creating the missing container at target cluster. If this situation persists, contact Nutanix support.
Severity	Critical

Table 439: Cluster does not support the planned failover of synchronously replicated entities. [130339] [A130339]

Name	The cluster does not support the planned failover of synchronously replicated entities
Description	Cluster does not support the planned failover of synchronously replicated entities.
Alert message	`Synchronous Replication is enabled for the entity vm_name but cluster cluster_name does not support planned failover.`
Cause	The cluster does not support the planned failover of synchronously replicated entities.
Impact	Planned failover of the entity will fail.
Resolution	Please upgrade the source and target clusters to a version later than 5.17.1. It is recommended to have source and target clusters on the same version.
Severity	Warning

Table 440: VM Protection Might Fail Post Recovery [130341] [A130341]

Name	VM Protection Might Fail Post Recovery
Description	VM Protection Might Fail Post Recovery
Alert message	`After recovery to remote availability zone remote_availability_zone_name, VM protection might fail because remote cluster remote_cluster_name does not support multisite protection.`
Cause	The remote cluster to which replication is happening does not support multisite protection.
Impact	After recovery to the specified remote site, VM protection might fail.
Resolution	Upgrade the remote Prism Element to AOS 5.19 or above, which supports multisite protection.
Severity	Warning

Table 441: Remote Site network connectivity unstable [130342] [A130342]

Name	Remote Site Connectivity Unstable
Description	Remote Site network connectivity is unstable
Alert message	`Network connectivity to remote site 'remote_name' is unstable`
Cause	Connection to remote site is unstable

Impact	VMs in Metro Protection Domain can experience intermittent I/O failures due to OS disk timeouts
Resolution	Check network connectivity between local and remote sites and manually break metro if VM I/O is stalled
KB Article	9713
Severity	Warning

Table 442: Volume Group Recovery Point Creation Failed. [130343] [A130343]

Name	Volume Group Recovery Point Failed
Description	Volume Group Recovery Point creation failed.
Alert message	`Failed to capture the Recovery Point for Volume Group 'volume_group_name'.`
Cause	Nutanix Data Protection service error occurred.
Impact	Volume Group Recovery Point will not be available to recover the Volume Group.
Resolution	Retry the Volume Group Recovery Point operation. If issue persists please contact Nutanix Support.
Severity	Warning

Table 443: Cluster does not support the live migration of synchronously replicated entities. [130345] [A130345]

Name	The cluster does not support the live migration of synchronously replicated entities
Description	Cluster does not support the live migration of synchronously replicated entities.
Alert message	`Synchronous Replication is enabled for the entity vm_name but cluster cluster_name does not support live migration.`
Cause	The cluster does not support the live migration of synchronously replicated entities.
Impact	Live migration of the entity will fail.
Resolution	Please upgrade the source and target clusters to version 5.19 or later. It is recommended to have source and target clusters on the same version.
Severity	Warning

Table 444: Consistency Group Conflict Occurred. [130349] [A130349]

Name	Consistency Group VM Conflicts
Description	Unable to protect VM due to conflicting protection policy and Consistency Group configurations.

Alert message	`Failed to protect the VM vm_name belonging to the Consistency Group consistency_group_name, because of reason.`
Cause	Multiple protection policies are applied on some of the members present in the Consistency Group.
Cause	No protection policy is applied to protect some of the members present in the Consistency Group.
Cause	Members in the same Consistency Group have different protection policies that can be applied.
Cause	VM belongs to a Consistency Group that contains VM(s) or Volume Group(s) which have been deleted.
Impact	VM cannot be recovered if disaster occurs.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that the members in a Consistency Group are protected by the same protection policy.
Resolution	Remove the stale/deleted VM(s) or Volume Group(s) from the Consistency Group.
KB Article	10571
Severity	Warning

Table 445: Consistency Group Conflict Occurred. [130350] [A130350]

Name	Consistency Group VG Conflicts
Description	Unable to protect Volume Group due to conflicting protection policy and Consistency Group configurations.
Alert message	`Failed to protect the Volume Group volume_group_config_name belonging to the Consistency Group consistency_group_name, because of reason.`
Cause	Multiple protection policies are applied on some of the members present in the Consistency Group.
Cause	No protection policy is applied to protect some of the members present in the Consistency Group.
Cause	Members in the same Consistency Group have different protection policies that can be applied.
Cause	Volume Group belongs to a Consistency Group that contains VM(s) or Volume Group(s) which have been deleted.
Impact	Volume Group cannot be recovered if disaster occurs.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.

Resolution	Ensure that the members in a Consistency Group are protected by the same protection policy.
Resolution	Remove the stale/deleted VM(s) or Volume Group(s) from the Consistency Group.
KB Article	10572
Severity	Warning

Table 446: Volume Group Protection Failed. [130351] [A130351]

Name	Volume Group Protection Failed
Description	Volume Group protection failed.
Alert message	`Failed to protect Volume Group: 'volume_group_name', because of 'reason'.`
Cause	Nutanix Data Protection service error occurred.
Impact	Volume Group cannot be recovered in case of disaster.
Resolution	Volume Group protection will get retriggered. If issue persists, please contact nutanix support.
KB Article	10560
Severity	Warning

Table 447: Volume Group Protection Might Fail Post Recovery [130352] [A130352]

Name	Volume Group Protection Might Fail Post Recovery
Description	Volume Group Protection Might Fail Post Recovery
Alert message	`After recovery to remote availability zone remote_availability_zone_name, Volume Group protection might fail because remote cluster remote_cluster_name does not support multisite protection.`
Cause	The remote cluster to which replication is happening does not support multisite protection.
Impact	After recovery to the specified remote site, Volume Group protection might fail.
Resolution	Upgrade the remote Prism Element to AOS 5.19 or above, which supports multisite protection.
KB Article	10573
Severity	Warning

Table 448: Automatic Synchronous Replication pause is triggered by Witness [130353] [A130353]

Name	Automatic Synchronous Replication Pause is triggered by Witness

Description	Automatic Synchronous Replication pause is triggered by Witness for {vm_count} VMs protected by Protection Rule '{protection_rule_name}'
Alert message	`Automatic Synchronous Replication pause has been triggered for vm_count VMs. 'reason'.`
Cause	Network connection to remote site or to Witness is lost probably because of a network partition or the remote site being unavailable.
Cause	One or more internal services are down or not working as expected.
Impact	Synchronous Replication for {vm_count} VMs will be paused.
Resolution	Ensure the remote site and Witness are reachable and resume Synchronous Replication.
Resolution	Contact Nutanix support.
Severity	Info

Table 449: Automatic Failover triggered by Witness [130354] [A130354]

Name	Automatic Failover triggered by Witness
Description	Automatic Failover triggered by Witness
Alert message	`Unplanned Failover of the Recovery Plan recovery_plan_name has been triggered by Witness. 'reason'.`
Cause	Automatic Failover triggered by Witness probably because of a network partition or source site being down.
Cause	One or more internal services are down or not working as expected.
Impact	Effected VMs are restored on remote site.
Resolution	Restore the connectivity between source and remote sites and the restored VMs can be migrated back to source site.
Resolution	Contact Nutanix support.
Severity	Info

Table 450: Volume Group Recovery Point Replication Failed. [130355] [A130355]

Name	Volume Group Recovery Point Replication Failed
Description	Volume Group Recovery Point replication failed.
Alert message	`Failed to replicate recovery point created at: recovery_point_create_time UTC of the volume group: volume_group_name to the recovery location: availability_zone_physical_name.`
Cause	Network Connectivity issues between the Primary and the Recovery Availability Zone.
Cause	Data Protection and Replication service is not working as expected. The service could be down.
Cause	Volume Group migration process is in progress.
Cause	Virtual IP address has not been configured on the remote cluster.

Cause	Remote clusters are unhealthy.
Impact	Volume Group Recovery Point will not be replicated to the recovery location.
Resolution	Check network connection between the Primary and the Recovery Availability Zone.
Resolution	Contact Nutanix support.
Resolution	Retry the Recovery Point replication operation after the migration is complete.
Resolution	Configure the Virtual IP address and then retry the Recovery Point replication operation.
Resolution	For a manually initiated Volume Group Recovery Point replication, retry again. For a scheduled Volume Group Recovery Point replication, ensure all the emote clusters are healthy. Then wait for the next scheduled Recovery Point replication.
KB Article	10754
Severity	Warning

Table 451: VM is part of multiple Recovery Plans configured with Witness [130356] [A130356]

Name	Same VM is in multiple Witness configured Recovery Plans
Description	VM is part of multiple Recovery Plans configured with Witness
Alert message	`The VM 'vm_name' is part of multiple Recovery Plans configured with Witness. reason`
Cause	Same VM is part of multiple Recovery Plans configured with Witness
Impact	The VM will be managed only by one of the Recovery Plans configured with Witness
Resolution	Ensure the VM is part of only one Witness configured Recovery Plan
KB Article	11001
Severity	Warning

Table 452: Volume Group Protection Failed [130357] [A130357]

Name	Volume Group Protection Failed On PC
Description	Volume Group protection failed.
Alert message	`Failed to protect Volume Group: 'volume_group_config_name', because of 'reason'.`
Cause	Data Protection service error occurred.
Cause	Volume Group is already protected in a legacy protection domain.
Cause	The Protection Policy applied on Volume Group has a NearSync or a Synchronous snapshot schedule applied on it.
Impact	Volume Group cannot be recovered in case of disaster.

Resolution	Volume Group protection will get retriggered. If issue persists, please contact Nutanix Support.
Resolution	Unprotect the Volume Group from the Protection Domain to protect it using Leap.
Resolution	NearSync and Synchronous protection is not supported on Volume Groups. Therefore, ensure that the Protection Policy protecting the Volume Groups does not have NearSync or Synchronous snapshot schedules
KB Article	11308
Severity	Warning

Table 453: Consistency Group Contains Stale Members [130358] [A130358]

Name	Consistency Group Contains Stale Members
Description	Consistency Group contains one or more members that have been deleted.
Alert message	`Consistency Group 'consistency_group_config_name' contains stale members.`
Cause	One or more VMs or Volume Groups belonging to the Consistency Group may have been deleted.
Impact	The VMs and Volume Groups part of the Consistency Group will not be protected.
Resolution	Remove the deleted members from the Consistency Group configuration.
KB Article	11413
Severity	Warning

Table 454: Synchronous Replication is not supported on the VM [130359] [A130359]

Name	Synchronous Replication is not supported on the VM
Description	Synchronous Replication is not supported on the VM
Alert message	`Synchronous replication is not supported on the VM vm_name. Reason: message`
Cause	VM has disks that are backed by a datasource external to the AOS cluster.
Impact	VM will not be synchronously replicated until the seeding from the external datasource is complete.
Resolution	Enable background storage migration on the disk backed by external datasource to complete the seeding.
KB Article	11344
Severity	Warning

Table 455: Partial Volume Group Recovery Point. [130361] [A130361]

Name	Partial Volume Group Recovery Point

Description	Partial Volume Group Recovery Point.
Alert message	`Recovery Point for Volume Group volume_group_name failed to capture associated policies and categories because reason.`
Cause	Recovery Point for Volume Group failed to capture associated policies and categories.
Impact	Volume Group will be restored without associated policies and categories.
Resolution	Check the connection between PC and the PE, then retry the Recovery Point operation.
KB Article	11744
Severity	Info

Table 456: Volume Group Recovery Point Replication Time Exceeded the RPO [130362] [A130362]

Name	Volume Group Recovery Point Replication Time Exceeds RPO
Description	Volume Group Recovery Point Replication Time Exceeded the RPO Limit.
Alert message	`Replication time of the snapshot created at 'recovery_point_create_time' UTC for Volume Group volume_group_name exceeded the RPO limits.`
Cause	Latest replication took more time than the RPO limit.
Impact	Replications for the next snapshots will be delayed.
Resolution	Please check your network.
KB Article	11745
Severity	Warning

Table 457: Recovery Plan Execution Failure [300401] [A300401]

Name	Recovery Plan Execution Failure
Description	Execution of the Recovery Plan failed.
Alert message	`operation_type failed for the Recovery Plan recovery_plan_name. One or more entities in this Recovery Plan might have failed to be started/ recovered.`
Cause	One or more entities were not recovered as part of the Recovery Plan execution.
Impact	One or more entities or applications in this Recovery Plan might not have started completely.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.
KB Article	8278
Severity	Critical

Table 458: Validation Warnings found during Recovery Plan Execution [300402] [A300402]

Name	Recovery Plan Execution Failure due to Validation Warnings
Description	Execution of the Recovery Plan failed.
Alert message	*operation_type* for Recovery plan *recovery_plan_name* resulted in validation warnings. You can choose to execute the Recovery Plan anyway or you can fix the issues.
Cause	The Recovery Plan validation resulted in warnings.
Impact	The Recovery Plan execution could not be started due to warnings in the validation report.
Resolution	Review the Recovery Plan validation report for detailed information.
Severity	Warning

Table 459: Subnet Creation Failure for the Recovery Plan [300405] [A300405]

Name	Subnet creation failure for Recovery Plan
Description	Subnet creation failed for the Recovery Plan.
Alert message	Subnet *network_cidr* creation in the virtual network *atlas_virtual_network_name* failed for the *operation_type* of the Recovery Plan *recovery_plan_name* executed through the recovery plan job *recovery_plan_job_name*.
Cause	Networks with conflicting CIDR might be present.
Impact	Some VMs will be recovered without vNICs.
Resolution	Resolve the issue. If you cannot resolve the issue, contact Nutanix support.
Severity	Critical

Table 460: Invalid Network Settings for the Recovery Plan [300407] [A300407]

Name	Recovery Plan Periodic Check for Invalid Network Settings
Description	Network settings have Availability Zones that are not in accordance with Availability Zone order list of Protection Policy applied to VMs in Recovery Plan.
Alert message	Invalid network settings found for the Recovery Plan *recovery_plan_name*. *reason*.
Cause	One or more Availability Zones listed in Protection Policy do not match the Availability Zones specified in the Recovery Plan network settings.
Impact	Recovery Plan will not be synced to Availability Zones defined in Protection Policy.
Resolution	Update Availability Zones in the Recovery Plan network settings to be in accordance with Availability Zone order in the Protection Policy.
Severity	Warning

Table 461: Recovery Plan Validation Failed With Errors [300408] [A300408]

Name	Validation Failed For Recovery Plan with Errors
Description	Recovery Plan validation failed with errors.
Alert message	`Validation failed with errors for the Recovery Plan recovery_plan_name.`
Cause	There might be various reasons for this issue.
Impact	Entities recovery might fail.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.
KB Article	8276
Severity	Critical

Table 462: Recovery Plan Validation Failed with Warnings [300409] [A300409]

Name	Validation Failed For Recovery Plan with Warnings
Description	Recovery Plan validation failed with warnings.
Alert message	`Validation failed with warnings for the Recovery Plan recovery_plan_name.`
Cause	There might be various reasons for this issue.
Impact	Entities recovery might fail.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.
Severity	Warning

Table 463: Recovery plans have conflicting network mappings [300410] [A300410]

Name	Recovery Plan network mapping conflict check
Description	Checks network mapping conflicts in multiple Recovery Plans.
Alert message	`Recovery plans #rp_1@ and #rp_2@ have different source networks mapped to same destination network target_network`
Cause	Multiple Recovery Plans have different source networks mapped to same destination network.
Impact	VMs isolated on source network might be able to reach each other on the destination after recovery.
Resolution	Avoid mapping multiple networks to same network in one or more Recovery Plans.
KB Article	7175
Severity	Info

Table 464: Planned Failover or Unplanned Failover operations will fail. [300411] [A300411]

Name	Delete the failed over entities from either the Primary or the Recovery Location
Description	Delete the failed over entities on either the primary or the recovery location.
Alert message	`Unplanned Failover for the Recovery Plan recovery_plan_name has completed successfully. Validate and clean up the entities from the site where it is not required, for new unplanned or planned failover operations for this recovery plan to succeed.`
Cause	Unplanned failover for the Recovery Plan completed successfully.
Impact	Planned Failover and Unplanned Failover operations are not possible without cleanup.
Resolution	Post unplanned failover entities are now active on the recovery Availability Zone. For synced entities, delete the entities on the primary Availability Zone to avoid performance issues. For entities protected by asynchronous replication, performance will not be affected and deletion can happen on either Availability Zone as required.
Severity	Warning

Table 465: Test Failover on Recovery Plan has not been executed recently [300412] [A300412]

Name	Recovery Plan Last Test Failover run status check
Description	Checks when last Test Failover was run on the Recovery Plan.
Alert message	`time days have elapsed since a Test Failover was run on Recovery Plan recovery_plan_name.`
Cause	Test Failover on the Recovery Plan has not been executed.
Impact	There might be issues in the recovery of VMs, which may be difficult to identify if the Test Failover is not run periodically.
Resolution	Run a Test failover on the Recovery Plan periodically, to ensure Recovery Plan has no errors.
KB Article	5793
Severity	Info

Table 466: Multiple Recovery Plans associated with a Category [300413] [A300413]

Name	Same Category in multiple Recovery Plans check.
Description	Checks if same category belongs to multiple Recovery Plans.
Alert message	`A specific category should be associated with only one Recovery Plan. Recovery Plans : rp_list are associated with the same entity_abbrv category category.`
Cause	Multiple Recovery Plans have Same Category
Impact	Execution of the following Recovery Plans will be affected and only one of those Recovery Plan will be able to recover VMs from specified category.

Resolution	Update the Recovery Plans such that the category is unique between them.
KB Article	5669
Severity	Warning

Table 467: Number of entities in Recovery Plan exceeds the threshold [300414] [A300414]

Name	Recovery Plan Entity Limit Check
Description	Checks if the entity count exceeds the threshold in Recovery Plan.
Alert message	`Maximum number of entities in a recovery plan should not exceed max_entity_count. Recovery Plan recovery_plan_name have entity_count entities.`
Cause	Number of entities in the Recovery Plan exceeds the limit.
Impact	Entity Recovery prone to failure.
Resolution	Reduce the number of entities in the Recovery Plan.
KB Article	5717
Severity	Warning

Table 468: VMs are part of multiple stages in Recovery Plan [300415] [A300415]

Name	VM in Multiple Stages Check
Description	Checks whether a VM is part of more than one stage in Recovery Plan.
Alert message	`A VM should be part of a single stage in a Recovery Plan. duplicate_entities`
Cause	VMs are part of multiple stages in the Recovery Plan.
Impact	VM specified in the first stage will be considered for recovery, later stages will be skipped.
Resolution	Update the Recovery Plan to ensure that the VMs are not part of multiple stages.
KB Article	5715
Severity	Info

Table 469: The same floating IP is associated with multiple VMs belonging to different Recovery Plans [300416] [A300416]

Name	Conflicting Floating IP Check
Description	Checks if VMs which are part of different Recovery Plan have same Floating IPs.
Alert message	`Same Floating IPs should not be assigned to multiple VMs and should not be part of multiple Recovery Plans. Floating IP floating_ip is assigned to VMs alert_msg.`

Cause	VMs belonging to different Recovery Plans are assigned the same Floating IP.
Impact	Floating IP assignment post VM recovery may fail.
Resolution	Update the Recovery Plans to ensure that a Floating IP address is to be assigned to only one VM.
KB Article	5797
Severity	Warning

Table 470: Some of the entities in the Recovery Plan are unprotected [300417] [A300417]

Name	Recovery Plan Entities Protection Policy Check
Description	Checks if entities configured in the Recovery Plan are protected.
Alert message	`entities in the Recovery Plan recovery_plan_name are not protected.`
Cause	Entities in the Recovery Plan are not protected.
Impact	Entities present in the Recovery Plan can't be recovered.
Resolution	Update associated entities in Recovery Plan to have Protection Policy.
KB Article	5795
Severity	Warning

Table 471: Entities not cleaned up following the Test Failover for Recovery Plan [300418] [A300418]

Name	Recovery Plan Checks for Cleanup of Entities after Test Failover
Description	Checks that entities recovered using Recovery Plan after Test Failover have been cleaned up.
Alert message	`Entities recovered following the Test Failover of Recovery Plan recovery_plan_name have not been cleaned up. Manual clean up of the entities is required.`
Cause	Entities recovered following the Test Failover of the Recovery Plan have not been cleaned up.
Impact	Subsequent recovery of the entities from Test Failover operation might get affected.
Resolution	Perform cleanup of the test entities for the Recovery Plan.
KB Article	5794
Severity	Warning

Table 472: Subnets Deletion Failure for the Recovery Plan [300419] [A300419]

Name	Dynamic Subnets deletion failure for Recovery Plan
Description	Deletion of dynamic subnets failed.

Alert message	`Deletion of the subnets networks_list failed for the operation_type of the Recovery Plan recovery_plan_name executed through the recovery_plan_job_name due to reason.`
Cause	Some of the VMs might be using the subnets.
Impact	Unused subnets will be billable.
Resolution	Resolve the issue. If you cannot resolve the issue, contact Nutanix Support.
Severity	Critical

Table 473: Floating IPs Allocation Failed [300420] [A300420]

Name	Floating IPs Allocation Failure.
Description	Floating IPs allocation failed.
Alert message	`Failed to allocate num_failed_floating_ips floating IPs, required for the dynamic allocation, for operation_type on the Recovery Plan recovery_plan_name.`
Cause	The system may not have enough floating IP addresses.
Impact	VMs will come up without floating IPs.
Resolution	Try to manually assign pre-allocated floating IPs for the vNICs.
Severity	Warning

Table 474: Floating IPs Deallocation Failed [300421] [A300421]

Name	Floating IPs Deallocation Failure.
Description	Floating IPs deallocation failed.
Alert message	`Failed to deallocate num_failed_floating_ips floating IPs, which were allocated dynamically but are not associated with any vNIC, for operation_type on the Recovery Plan recovery_plan_name.`
Cause	The floating IPs might be in use.
Impact	Floating IPs will remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.
Severity	Warning

Table 475: Recovery Plan Execution Failure due to Validation Errors [300422] [A300422]

Name	Recovery Plan Execution Failure due to Validation Errors
Description	Execution of the Recovery Plan failed.
Alert message	`The operation_type for the recovery plan recovery_plan_name could not be started due to errors in validation.`
Cause	The Recovery Plan validation failed.

Impact	The Recovery Plan execution could not be started.
Resolution	Review the Recovery Plan validation report for detailed information.
KB Article	8277
Severity	Critical

Table 476: Recovery Plan has multiple Availability Zone Orders [300424] [A300424]

Name	Recovery Plan Multiple Availability Zone Order Check
Description	Checks if the Recovery Plan has any entity protected to multiple Availability Zones.
Alert message	`Recovery Plan recovery_plan_name has entities protected to multiple Availability Zones.`
Cause	Recovery Plan contains entities protected to multiple Availability Zones.
Impact	The Recovery Plan update will not be allowed.
Resolution	Update associated entities in Recovery Plan to have single Availability Zone Order.
KB Article	6591
Severity	Warning

Table 477: Recovery Plan contains VMs with Unsupported CHDR VM configuration. [300425] [A300425]

Name	Recovery Plan CHDR VM Config Check
Description	Checks if Recovery Plan contains VMs with Unsupported CHDR VM configuration.
Alert message	`Recovery Plan recovery_plan_name contains VMs vm_list with Unsupported CHDR VM configuration.`
Cause	Recovery Plan contains VMs with unsupported configuration for Cross hypervisor disaster recovery.
Impact	The VMs will not be recovered.
Resolution	Install NGT on all the VMs in the Recovery Plan, ensure that you do not have UEFI boot, and have no SATA/PCI or delta disk attached.
KB Article	6835
Severity	Warning

Table 478: Incompatible Recovery Availability Zones for some VMs in the Recovery Plan [300426] [A300426]

Name	Recovery Plan Incompatible Recovery Availability Zone Check

Description	Checks if Recovery Plan contains VMware VMs and snapshots for these VMs are replicated to a recovery Availability Zone that doesn't support recovery of VMware VMs.
Alert message	`Incompatible Recovery Availability Zones for Recovery Plan recovery_plan_name. Recovery of VMs incompatibles_vms on Recovery Availability Zones incompatbile_target_availability_zone_names will fail.`
Cause	Target Availability Zone is running Prism Central version less than AOS 5.11 version and hence it does not support the recovery of VMware VMs.
Impact	The VM recovery will fail.
Resolution	Upgrade the Target Availability Zone to 5.11 or later version, or Remove the entity mentioned in the description of the alert from the Recovery plan.
KB Article	6872
Severity	Warning

Table 479: Floating IPs Deallocation Failed after Failback from Xi. [300427] [A300427]

Name	Unassociated Floating IPs Deallocation Failure after Failback from Xi.
Description	Failed to deallocate unassociated floating IPs after failback from Xi that were earlier dynamically created.
Alert message	`Failed to deallocate unassociated floating IPs failed_floating_ips after failback from Xi that were earlier dynamically created, for operation_type on the Recovery Plan recovery_plan_name.`
Cause	The floating IPs might be in use.
Impact	Floating IPs will continue to remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.
Severity	Warning

Table 480: Recovery Plan contains VMs with special configuration and no empty CDROM [300428] [A300428]

Name	Recovery Plan NGT Config Empty CD ROM Check
Description	Checks if VMs part of the recovery plan have an empty CDROM available to mount nutanix guest tools, so NGT reconfiguration tasks can be completed on restore.
Alert message	`Recovery Plan recovery_plan_name contains VMs vm_list with special configuration that require NGT reconfiguration on restore. Nutanix Guest tools cannot be mounted as there is no empty CD ROM slot on the VM.`
Cause	VMs on Source Availability Zone do not have a free CDROM slot.
Impact	VM will be recovered however post recovery VM reconfiguration tasks, including NGT reconfiguration, static IP reconfiguration will fail.
Resolution	Add an empty CDROM to the VM.

KB Article	7054
Severity	Warning

Table 481: Recovery Plan Execution Exceeded the Time Limit. [300429] [A300429]

Name	Recovery Plan Execution Exceeded the Maximum Expected Time Limit.
Description	Recovery Plan execution exceeded the time limit.
Alert message	`Execution of Recovery Plan 'recovery_plan_name' exceeded the estimated completion time by excess_time_taken.`
Cause	Restore or Power On operation may have taken longer time.
Impact	Recovery Plan execution is prone to failure.
Resolution	Contact Nutanix customer service.
Severity	Warning

Table 482: Entities in category(s) present in multiple fault domains. [300432] [A300432]

Name	Category Protected entities Multiple Fault Domain Check
Description	Checks if entities protected by category(s) in Recovery Plan are present in multiple fault domains.
Alert message	`alert_message`
Cause	Protection Policy is applied on category(s) with Live entities on Local AZ and Remote AZ.
Cause	Entities from category(s) are getting replicated between clusters registered with Local Availability Zone and are configured for failover operations for different Availability Zones in Recovery Plan.
Impact	Selection of entities for recovery will be affected.
Resolution	Ensure category(s) present in Recovery Plan contain entities from only single Availability Zone.
Resolution	Ensure category(s) present in Recovery Plan contain entities from only one fault domain. Refer KB 9472 for further guidance. Contact Nutanix support if the issue persists.
KB Article	9472
Severity	Warning

License

Table 483: Nutanix Objects License Expiry [14008] [A14008]

Name	ObjectLicenseExpiry
Description	Check if Nutanix Objects license is about to expire.

Alert message	`This cluster_type has expiring or expired Nutanix Objects license(s). Detailed license expiry info: license_expiry_msg license_warning_msg`
Cause	The Nutanix Objects license is set to expire or has expired.
Impact	The Nutanix Objects instance will be out of licensing compliance.
Resolution	Contact Nutanix to extend the expiration of the Object license, or apply a new Object license.
KB Article	2469
Severity	Warning

Guest VM

Table 484: Node Failure [130030] [A1137]

Name	Node Status
Description	Node failure.
Alert message	`Host hypervisor_address appears to have failed. High Availability is restarting VMs on failover_host_info.`
Cause	Host is not accessible.
Impact	VMs will restart on another host.
Resolution	Check for network connectivity or a hardware failure.
Severity	Critical
Severity	Warning

Table 485: Protected VM Not Found [130038] [A1010]

Name	Protected VM Not Found
Description	Protected VM not found.
Alert message	`Unable to locate the VM 'vm_name (ID vm_id)' protected by protection domain 'protection_domain_name'.`
Cause	The protected VM cannot be found and may have been deleted.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove the VM from the protection domain.
Severity	Info

Table 486: VM Action Error [130057] [A1033]

Name	VM Action Status
Description	VM action error.

Alert message	`Failed to action VM with name 'vm_name and internal ID (ID vm_id)' due to reason.`
Cause	A VM could not be restored because of a hypervisor error, or could not be deleted because it is still in use.
Impact	The requested VM action (restore or delete) could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10658
Severity	Warning

Table 487: Failure To Restart VMs For HA Event [130064] [A1145]

Name	VMs Restart Status
Description	Failure to restart VMs for an HA event.
Alert message	`Failed to restart one or more VMs that were running on failed host hypervisor_address.`
Cause	Not enough memory or CPU resources within the cluster.
Impact	Some VMs will remain powered off.
Resolution	Shut down unneeded VMs to free cluster resources; expand the cluster to add resources; or wait until the failed host recovers. Once resources are available, power on the VMs.
KB Article	8199
Severity	Critical
Severity	Warning

Table 488: HA Host Evacuation Failure [130065] [A1146]

Name	Host Evacuation Status
Description	Failure to evacuate host while entering maintenance mode or reserving host for HA.
Alert message	`Failed to evacuate some VMs from host hypervisor_address while reason.`
Cause	Not enough memory/CPU resources within the cluster.
Impact	The host will not enter requested state.
Resolution	Shut down unneeded VMs to free cluster resources or expand the cluster to add resources.
Severity	Warning

Table 489: HA Healing Failure [130067] [A1155]

Name	VM HA Healing Status

Description	HA healing failure.
Alert message	`Could not restore VM High Availability.`
Cause	Not enough memory/CPU resources within the cluster.
Impact	VMs are no longer protected against host failure.
Resolution	Shut down unneeded VMs to free cluster resources, expand the cluster to add resources, or disable protection against host failures. VMs cannot be powered on or migrated until the system has sufficient resources to protect against the configured number of host failures.
Severity	Warning

Table 490: VM High Availability Failure [130068] [A130068]

Name	VM HA Status
Description	High availability failure.
Alert message	`VMs are no longer protected against host failure. Reason: reason.`
Cause	Not enough memory/CPU resources within the cluster.
Impact	VMs are no longer protected against host failure.
Resolution	Shut down unneeded VMs to free cluster resources, expand the cluster to add resources, or disable protection against host failures. VMs cannot be powered on or migrated until the system has sufficient resources to protect against the configured number of host failures.
Severity	Warning

Table 491: iSCSI Configuration Failed [130100] [A130100]

Name	iSCSI Configuration Failed
Description	iSCSI Configuration Failed
Alert message	`Failed to re-configure iSCSI settings on the recovered VM 'vm_name'. reason.`
Cause	Nutanix Guest Tools failed to execute some iSCSI commands on the guest VM.
Impact	iSCSI disks may become unavailable on the guest VM.
Resolution	(A) If IQN and iSCSI target IP addresses of the VM have been updated by Nutanix Guest Agent, Discover and connect to new targets after rebooting the VM. (B) Manually configure iSCSI settings on the guest VM. (C) Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10118
Severity	Critical

Table 492: Failed to Mount Nutanix Guest Tools [130103] [A130103]

Name	NGT Mount Failure
Description	Nutanix Guest Tools mount may have failed.
Alert message	`Nutanix Guest Tools mount failed for vm_name. reason.`
Cause	CDROM slot is not free. This can be because of hypervisor-related error.
Impact	Nutanix Guest Tools will not be mounted on the VM.
Resolution	Check the logs on the CVM. If required, also check the VM logs for any error information to fix the issue.
Severity	Warning

Table 493: VSS Snapshot is not supported for the VM. [130105] [A130105]

Name	VSS Snapshot Not Supported
Description	VSS snapshot is not supported for the VM.
Alert message	`VSS snapshot is not supported for the VM 'vm_name', because reason.`
Cause	VM has unsupported configuration.
Impact	Hypervisor based application consistent snapshot is taken on ESX. Crash consistent snapshot is taken for other hypervisors.
Resolution	Please look at the alert message and fix the invalid configuration.
Severity	Warning

Table 494: VSS Snapshot Aborted [130106] [A130106]

Name	VSS Snapshot Aborted
Description	VSS snapshot is aborted by the Guest VM.
Alert message	`VSS snapshot is aborted for the VM 'vm_name (ID vm_id)', because reason.`
Cause	Guest VM aborted VSS snapshot operation.
Impact	Crash consistent snapshot is taken.
Resolution	Please look at the alert message for details of the failure and take necessary actions. If issue persists contact Nutanix Support.
Severity	Warning

Table 495: Recovered VM Disk Configuration Update Failed [130107] [A130107]

Name	Disk Configuration Update Failed
Description	Disk Configuration Update Failed
Alert message	`Failed to make some disks online on the recovered VM 'vm_name'. reason.`
Cause	Nutanix Guest Tools failed to automatically bring the disks online.

Impact	Disks may become offline on the recovered VM.
Resolution	Manually bring the disks online on the recovered VM. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	8208
Severity	Critical

Table 496: Execution of the PostThaw Script Failed [130113] [A130113]

Name	PostThaw Script Execution Failed
Description	Execution of the PostThaw Script Failed
Alert message	`Failed to execute the post_thaw script during the creation of the application consistent snapshot with uuid 'snapshot_uuid' for the VM 'vm_name'. Error : error_message.`
Cause	Guest VM failed to execute the post_thaw script during the creation of the application consistent snapshot.
Impact	Some of the applications that got stopped by the pre_freeze script may not start after application consistent snapshot is created for the VM.
Resolution	Manually run the post_thaw script on the guest VM. Fix the cause of the script failure to avoid any further execution failures. If you cannot resolve the error, contact Nutanix support.
KB Article	8116
Severity	Critical

Table 497: Failed to take the application-consistent snapshot for the VM. [130127] [A130127]

Name	Application Consistent Snapshot Skipped
Description	Failed to take the application-consistent snapshot for the VM.
Alert message	`Failed to take the application-consistent snapshot for the VM 'vm_name', because reason.`
Cause	VMware tools is not installed on the VM.
Cause	VMware tools is not running on the VM.
Cause	Removal of temporary hypervisor-based snapshot of the VM failed.
Impact	Crash consistent snapshot is taken for the VM.
Resolution	Install the VMware tools on the VM.
Resolution	Start the VMware tools on the VM.
Resolution	Manually remove temporary hypervisor-based snapshot of the VM.
Severity	Warning

Table 498: VM Registration Warning [130132] [A130132]

Name	VM Registration Warning
Description	VM Registration caused warning
Alert message	`VM with name vm_name has been recovered successfully, however the following issue was observed - reason`
Cause	Processor features may not be compatible.
Impact	VM registered successfully during the disaster recovery but with warnings.
Resolution	Ensure that the recovered VM is working as expected.
Severity	Warning

Table 499: Guest Power Operation through NGT Failed [130149] [A130149]

Name	VM Guest Power Op Failed
Description	Guest Power Operation Failed
Alert message	`Failed to perform operation 'operation'on VM 'vm_name'. reason.`
Cause	Power operation failed in guest operating system.
Impact	Desired power operation may not be completed.
Resolution	(A) Manually shutdown the VM by logging into the VM and running the appropriate shutdown command. (B) Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.
KB Article	10242
Severity	Warning

Table 500: Removal of the Temporary Hypervisor Snapshot Failed for the VM [130150] [A130150]

Name	Temporary Hypervisor Snapshot Cleanup Failed
Description	Removal of the temporary hypervisor snapshot that got created while taking the host-based application-consistent snapshot has failed for the VM.
Alert message	`Removal of the temporary hypervisor snapshot that got created while taking the Nutanix host-based application-consistent snapshot has failed for the VM vm_name.`
Cause	VM might be experiencing high I/O when host-based application-consistent snapshot cleanup is attempted.
Impact	Future application-consistent snapshots might fail.
Resolution	Wait until the merging of the snapshots is complete and then verify that there are no temporary Nutanix host-based snapshots before taking an application-consistent snapshot for the VM.
Severity	Warning

Table 501: Nutanix Guest Tools Failed to Initiate VM Reboot. [130190] [A130190]

Name	NGT VM Reboot Failure
Description	Nutanix Guest Tools failed to initiate reboot on the VM.
Alert message	`Nutanix Guest Tools failed to initiate the VM reboot after NGT install/upgrade on the VM 'vm_name' due to reason.`
Cause	Failed to initiate reboot on the VM. Please check the alert message for exact cause of the failure.
Impact	VM is not rebooted after NGT install/upgrade. VM Mobility drivers may not be updated.
Resolution	Resolve the stated reason for the failure in the alert message. If the issue persists contact Nutanix Support.
Severity	Warning

Table 502: Conflicting NGT policies [130192] [A130192]

Name	NGT policy conflicts
Description	There are conflicting NGT policies for the VM.
Alert message	`Conflicting NGT policies found for the VM 'vm_uuid'.`
Cause	More than one NGT policy is applicable for this VM. Please resolve the conflicts.
Impact	VM is not rebooted after NGT install/upgrade. Scheduled reboot of the VMs will be impacted.
Resolution	Check if more than 1 NGT policy exists for each VM and remove such policies via the Prism Central.
KB Article	10694
Severity	Warning

Table 503: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	`Failed to evacuate VMs off of node node_uuid: error`
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.
KB Article	8471
Severity	Critical

Table 504: NGT CD-ROM not Unmounted on the VM [130334] [A130334]

Name	NGT CD-ROM Not Unmounted
Description	NGT CD-ROM not unmounted on the VM.
Alert message	`NGT CD-ROM not unmounted on the VM 'vm_name' with uuid vm_id`
Cause	NGT CD-ROM not unmounted on the VM because NGT version of the VM is incompatible with the NGT version of the cluster.
Impact	NGT workflow will not be impacted, however it is advised to unmount the stale ISO.
Resolution	Unmount the NGT ISO on the VM.
Severity	Info

Table 505: Node Removal is Stuck [130337] [A130337]

Name	Node Removal Stuck
Description	Cannot proceed with the node removal as the node is not in maintenance mode.
Alert message	`Cannot proceed with the node removal due to: error`
Cause	Insufficient cluster resources are available for the VMs to migrate off the host that is marked for removal.
Cause	Exit Maintenance Mode action triggered manually on the node that is marked for removal, and hence the node removal task cannot proceed.
Cause	The node that is being removed is either the source or the destination for VMs that are being live migrated to/from another cluster, and hence the node removal task cannot proceed.
Impact	The node marked for removal cannot be removed from the cluster until this condition is resolved.
Resolution	Ensure enough resources are available on other hosts for the VMs to evacuate from the host being removed. Consider manually shutting down unnecessary VMs and reviewing the HA settings. Put the affected node back in Maintenance Mode now.
Resolution	Put the affected node back in Maintenance Mode.
Resolution	Wait for cross cluster live migration of VM(s) to finish.
KB Article	8471
Severity	Warning

Table 506: VM is in Paused state [130348] [A130348]

Name	VM is in Paused state
Description	VM found in paused state.
Alert message	`VM vm_name with ID vm_uuid has been found in paused state`

Cause	Failed migration or some hypervisor issue.
Impact	VM in paused state most likely means downtime and it's in an irrecoverable state.
Resolution	Power cycle the VM to get it out of this state.
KB Article	10834
Severity	Warning

Table 507: File Analytics VM component failure [160081] [A160081]

Name	File Analytics VM component failure
Description	One or more components of the File Analytics VM {avm_ip} are not functioning properly or have failed.
Alert message	`One or more components of the File Analytics VM avm_ip are not functioning properly or have failed.`
Cause	One or more File Analytics VM components have failed.
Impact	File Server analytics will not perform appropriately.
Resolution	Contact Nutanix Support.
Severity	Critical

Table 508: File Analytics VM high CPU usage. [160082] [A160082]

Name	File Analytics VM high CPU usage.
Description	The File Analytics VM {avm_ip} has a high CPU usage of {usage_percent}%
Alert message	`The File Analytics VM avm_ip has a high CPU usage of usage_percent%`
Cause	The File Analytics VM is performing too many CPU intensive operations.
Impact	The File Analytics VM's performance will be impacted.
Resolution	Monitor which component has a high CPU usage and reduce the processing load.
Severity	Warning

Table 509: File Analytics VM high disk usage. [160083] [A160083]

Name	File Analytics VM high disk usage.
Description	The File Analytics VM {avm_ip} has a high disk usage of {usage_percent}%
Alert message	`The File Analytics VM avm_ip has a high disk usage of usage_percent%`
Cause	Extensive disk usage on the File Analytics VM
Impact	The File Analytics VM's performance will be impacted.
Resolution	Add more capacity to the File Analytics VM or delete some data.

Severity	Warning

Table 510: File Analytics VM low memory available. [160084] [A160084]

Name	File Analytics VM low memory available.
Description	The memory available on File Analytics VM {avm_ip} is {memory_available} bytes which is low.
Alert message	`The percent memory available on File Analytics VM avm_ip is memory_available which is low`
Cause	Excessive memory usage on the File Analytics VM
Impact	The File Analytics VM's performance will be impacted.
Resolution	Add more RAM
Severity	Warning

Table 511: vNUMA VM Pinning Failure. [200402] [A200402]

Name	vNUMA VM Pinning Failure
Description	Cluster was unable to provide requested vNUMA pinning for VM.
Alert message	`VM 'vm_name' with uuid vm_id is running with a non-compliant vNUMA pinning following operation.`
Cause	Insufficient memory and/or CPU resources available.
Impact	VM performance may degrade.
Resolution	Migrate VM to a host which has adequate CPU and memory resources. Please refer to KB-10223.
KB Article	10223
Severity	Warning

Table 512: VM stored on RF1 container failed to start after host/CVM maintenance. [200405] [A200405]

Name	VM stored on RF1 container failed to start after host/CVM maintenance
Description	VM stored on RF1 container failed to start after host/CVM maintenance.
Alert message	`VM vm_name (UUID: vm_uuid) failed to start with the following error: message`
Cause	Insufficient resources available to start the VM.
Impact	VM which was running before the host/CVM maintenance will stay powered off until manual intervention.
Resolution	Free up resources and manually start the VM. Refer to KB-10804.
KB Article	10804

Severity	Critical

Hardware

Table 513: Node Failure [130030] [A1137]

Name	Node Status
Description	Node failure.
Alert message	`Host hypervisor_address appears to have failed. High Availability is restarting VMs on failover_host_info.`
Cause	Host is not accessible.
Impact	VMs will restart on another host.
Resolution	Check for network connectivity or a hardware failure.
Severity	Critical
Severity	Warning

Table 514: Non Self Encryption Drive Disk Inserted [130031] [A1122]

Name	Non SED Disk Inserted Check
Description	Non-self-encrypting drive is inserted.
Alert message	`Non encrypted disk with serial disk_serial was added in drive bay disk_location on Controller VM service_vm_external_ip. It is not usable because the rest of the cluster is protected using encrypted drives.`
Cause	A cluster with self-encrypting drives has a non-self encrypting drive installed.
Impact	The new drive was not mounted because the cluster cannot be fully secured without all drives supporting encryption.
Resolution	Ensure all installed drives are of self-encrypting type. If more help is needed call Nutanix support.
KB Article	8417
Severity	Critical

Table 515: Disk Failed/Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Drive is marked offline.
Alert message	`Disk disk_position with id disk_id on node node_position of block block_position has been marked offline due to IO errors. Serial number of the disk is disk_serial in host host_ip of block block_serial.`
Cause	A drive on the node has failed and was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.

Resolution	Replace the failed drive. Refer to the Nutanix documentation for instructions.
KB Article	8453
Severity	Critical

Table 516: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	`Drive disk_id with serial disk_serial and model disk_model in drive bay disk_location on Controller VM service_vm_external_ip has failed. Node serial is node_serial_number and Node position node_position. #rf1_custom_message@`
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 517: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical drive is removed from the slot.
Alert message	`Drive with serial disk_serial and model disk_model was removed from drive bay disk_location on Controller VM service_vm_external_ip. Node serial is node_serial_number and Node position node_position. #rf1_custom_message@`
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the drive will start.
Resolution	Remove the drive from the slot and reseat it. Contact the Hardware Vendor if issue persists.
KB Article	8031
Severity	Critical

Table 518: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	`Drive disk_id with serial disk_serial in drive bay disk_location on Controller VM service_vm_external_ip has failed diagnostic test. This drive is failing or has failed. #rf1_custom_message@`

Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 519: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	`Failed to evacuate VMs off of node node_uuid: error`
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.
KB Article	8471
Severity	Critical

Node

Table 520: Latency between CVMs is high [6005] [A6005]

Name	Inter-CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	`Latency between vm_types is higher than ping_threshold ms.`
Cause	Malfunctioning or misconfigured network infrastructure
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.
Ping Time Threshold	Critical Threshold: 15 ms
KB Article	3690
Severity	Critical

Table 521: Hardware difference is noticed [106068] [A106068]

Name	Hardware differences alert
Description	Any change in any parameter of the hardware will raise this alert
Alert message	`alert_msg`
Cause	Hardware configuration is changed

Impact	NA
Resolution	NA
Severity	Info

Table 522: SSH keys missing [111089] [A111089]

Name	Node Missing SSH Keys check
Description	Check if /home/nutanix/ssh_keys directory contents are missing.
Alert message	*alert_msg*
Cause	Some of SSH keys are missing from /home/nutanix/ssh_keys directory.
Impact	Missing SSH keys can jeopardise ongoing infrastructure workflows in a silent manner such as node un-configuration.
Resolution	Please refer to KB 10736 for more information.
KB Article	10736
Severity	Critical

Table 523: Kernel mitigation configuration doesn't match current state. [111090] [A111090]

Name	Kernel mitigations check
Description	Check if spectre/meltdown mitigation is synced up with mitigations value in /proc/cmdline.
Alert message	*alert_msg*
Cause	Spectre/meltdown configuration doesn't match current state.
Impact	Uncertain configuration.
Resolution	Reboot CVM. Please refer to KB 10938 for more information.
KB Article	10938
Severity	Critical

Table 524: Disk utilization for audit partition is more than 75%. [111092] [A111092]

Name	AHV audit disk space check
Description	Checks if the disk usage for audit partition is higher than 75%.
Alert message	`On host` *host_ip*, *alert_msg*
Cause	Disk utilization for audit partition is more than 75%.
Impact	There may not be enough disk space left for new audit records.
Resolution	Refer KB11680 to address the disk utilization over 75%.
KB Article	11680
Severity	Warning

Table 525: Certificate Creation Error [130013] [A1112]

Name	Certificate Creation Status
Description	Certificate signing request creation failed.
Alert message	`Failed to create a certificate signing request on node service_vm_external_ip in the cluster.`
Cause	Either the OpenSSL library had a failure or a non-supported subject field, subject value, or subject alternative name was provided.
Impact	Without a CSR then data at rest encryption cannot be enabled.
Resolution	Ensure that the certificate signing request fields are valid and retry the command, or contact Nutanix support if the OpenSSL configuration file is missing or corrupted.
Severity	Info

Table 526: Agent VM Restoration Failure [130091] [A130091]

Name	Agent VM Restoration
Description	Failure to Restore Agent VM
Alert message	`Failed to restore agent VM with UUID vm_uuid.`
Cause	Not enough memory/CPU resources on this node.
Impact	VMs on this host may not function correctly because of the missing agent VM.
Resolution	Shut down unneeded VMs to free cluster resources or expand the cluster to add resources.
KB Article	8010
Severity	Critical

Table 527: VM Group Snapshot and Current State Mismatch [130110] [A130110]

Name	VM group snapshot and current mismatch
Description	Check that snapshot to restore matches current VM group
Alert message	`VM group vm_group_uuid snapshot and current state mismatch: message`
Cause	VM group changed after snapshot was created.
Impact	VMs might be scheduled to run on the same host instead of a different host.
Resolution	Recreate VM snapshot to ensure up-to-date VM group information.
Severity	Critical
Severity	Warning

Table 528: Protection Domain is in a Decoupled State [130112] [A130112]

Name	Protection Domain Decoupled Status
Description	Protection domain is in a decoupled state.
Alert message	`Protection domain 'protection_domain_name' protecting the Storage Container 'container_name' is in a decoupled state. impact`
Cause	Standby site has been promoted to the primary site.
Impact	Metro availability for the protection domain on the primary is unavailable.
Resolution	Re-enable metro availability from the remote site.
Severity	Warning

Table 529: Host Network Uplink Configuration Failed. [130160] [A130160]

Name	ClusterVswitch Configuration Update Failed
Description	Failed to apply host uplink configuration on the host(s).
Alert message	`Host network uplink configuration update on Host(s):'node_uuid' failed due to the reason: cause`
Cause	Host uplink configuration failed to apply on the host(s).
Impact	Host OVS configuration may not be in sync with cluster vswitch configuration.
Resolution	Use manage_ovs to configure the host individually after correcting errors.
Severity	Warning

Table 530: Failed to configure host for Atlas networking [130201] [A130201]

Name	Failed to configure host for Atlas networking
Description	Failed to configure host for Atlas networking.
Alert message	`Failed to configure host hypervisor_address for Atlas networking: message`
Cause	Failed to configure host for Atlas networking.
Impact	VMs on Atlas networks cannot run on this host.
Resolution	Check the hypervisor version, and upgrade if necessary.
KB Article	8526
Severity	Critical

Table 531: Virtual Switch configuration deployment failed. [150006] [A150006]

Name	Virtual Switch configuration deployment failed
Description	Notifies when a validated Virtual Switch configuration is unable to be applied to host network infrastructure

Alert message	`Virtual Switch vs_name configuration deployment on Host(s):'node_uuid' failed due to the reason: cause`
Cause	Unable to communicate with one or more host(s) to apply the configuration change.
Cause	Lost connectivity while applying the Virtual Switch configuration.
Cause	OVS did not accept or apply the requested configuration change.
Cause	Unexpected error with service(s) required for Virtual Switch configuration change (Acropolis, Genesis, IDF, Prism Gateway).
Impact	Unable to apply Virtual Switch configuration on the host(s).
Resolution	Ensure all hosts are up and reachable by all CVMs.
Resolution	Ensure the external network is properly configured to retain connectivity using the applied Virtual Switch configuration.
Resolution	Review the alert message and affected host OVS log files for further details.
Resolution	Review the alert message and related service Acropolis and Genesis log files for further details.
KB Article	9406
Severity	Critical
Severity	Warning

Table 532: IPFIX exporter host update failed [150011] [A150011]

Name	IPFIX exporter host update failed
Description	IPFIX exporter update failed on the host in the cluster.
Alert message	`IPFIX exporter update failed on host hypervisor_address because of error: cause`
Cause	IPFIX exporter update failed on the host in the cluster.
Impact	IPFIX exporter on this host in the cluster may not be working as expected.
Resolution	Check the error message and follow the appropriate steps to fix the IPFIX exporter state on the host.
Severity	Warning

Table 533: OVN Connection Unhealthy [150012] [A150012]

Name	OVN Connection Unhealthy
Description	Hypervisor node is disconnected from the Advanced Networking Controller.
Alert message	`Node uuid_node is disconnected from the Advanced Networking Controller`
Cause	The Chassis' ovn-controller does not have a Healthy connection to the Advanced Networking Controller.
Impact	Ability to make network-related configurations may be impacted.
Impact	Overlay networking may be disconnected on this Hypervisor node.

Resolution	Verify that the Hypervisor node can connect to the Prism Central hosting the Advanced Networking Controller. Verify the health of the Advanced Networking Controller.
KB Article	11129
Severity	Critical

Table 534: Home partition is not expanded to 50GB. [200330] [A200330]

Name	Prism Central home partition expansion check
Description	Checks if home partition is expanded to 50GB after upgrade.
Alert message	`Home partition is not expanded to 50GB on Prism Central pc_ip after upgrade.`
Cause	During an upgrade, home partition extension has failed.
Impact	Prism Central performance may be degraded.
Resolution	Refer to KB 9525 for further details.
KB Article	9525
Severity	Warning

Table 535: SSL Certificate is expiring soon [200705] [A200705]

Name	Check SSL Cert Expiry
Description	Notify if SSL Certificate is expiring soon
Alert message	`SSL Certificate is expiring in the next 30 days. Please regenerate`
Cause	SSL Certificate is expiring soon
Impact	SSL Certificates are expired and Cannot be accessed.
Resolution	Regenerate the ssl certificate.
KB Article	11493
Severity	Critical
Severity	Warning

Network

Table 536: Latency between CVMs is high [6005] [A6005]

Name	Inter-CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	`Latency between vm_types is higher than ping_threshold ms.`
Cause	Malfunctioning or misconfigured network infrastructure

Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.
Ping Time Threshold	Critical Threshold: 15 ms
KB Article	3690
Severity	Critical

Table 537: CVM TCP checksum error rate check [103104]

Name	CVM TCP checksum error rate check
Description	Check if corrupted packets with compromised checksums are reaching the CVM and not dropped by the host NICs.
Cause	Hardware failure or NIC firmware/driver compatibility issue.
Impact	Cluster performance may be significantly degraded.
Resolution	Check for any hardware failures or issues with firmware/driver/cables. Please refer to KB 9623 for more details.
KB Article	9623

Table 538: OVA Upload Interrupted [111076] [A111076]

Name	OVA Upload Interrupted
Description	OVA upload interrupted
Alert message	`Upload interrupted for OVA ova_name`
Cause	Network might have been disconnected.
Cause	Prism UI may have been closed.
Cause	OVA file might not be accessible.
Impact	OVA will not be available for use.
Resolution	Validate stable network connection.
Resolution	Ensure Prism UI is open and logged in.
Resolution	Ensure OVA file is accessible.
Severity	Info

Table 539: Failed to configure host for Atlas networking [130201] [A130201]

Name	Failed to configure host for Atlas networking
Description	Failed to configure host for Atlas networking.
Alert message	`Failed to configure host hypervisor_address for Atlas networking: message`
Cause	Failed to configure host for Atlas networking.

Impact	VMs on Atlas networks cannot run on this host.
Resolution	Check the hypervisor version, and upgrade if necessary.
KB Article	8526
Severity	Critical

Table 540: Failed to reserve host memory for Atlas networking [130202] [A130202]

Name	Failed to reserve host memory for Atlas networking
Description	Failed to reserve host memory for Atlas networking.
Alert message	`Failed to reserve host memory for Atlas networking`
Cause	Failed to reserve host memory for Atlas networking.
Impact	If hypervisor memory is exhausted, random VMs may be powered off to reclaim memory.
Resolution	Power off VMs to reduce host memory consumption.
KB Article	8525
Severity	Critical

Table 541: IPFIX exporter host update failed [150011] [A150011]

Name	IPFIX exporter host update failed
Description	IPFIX exporter update failed on the host in the cluster.
Alert message	`IPFIX exporter update failed on host hypervisor_address because of error: cause`
Cause	IPFIX exporter update failed on the host in the cluster.
Impact	IPFIX exporter on this host in the cluster may not be working as expected.
Resolution	Check the error message and follow the appropriate steps to fix the IPFIX exporter state on the host.
Severity	Warning

Table 542: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed
Description	Programming a Flow rule failed.
Alert message	`Flow policy operation failed for rule_name. reason`
Cause	Microsegmentation service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.

Resolution	Check PC microsegmentation service, PE acropolis service, PC to PE connection, PE to AHV connection.
Severity	Warning

Table 543: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed
Description	Flow failure event.
Alert message	`Flow operation failed on remote_uuid. reason`
Cause	Microsegmentation service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microsegmentation service, Check PE acropolis service, AHV microsegmentation module.
Severity	Warning

Table 544: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	`Flow mode change failed on remote_uuid. reason`
Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect cvm_logs and contact Nutanix Support
Severity	Warning

Table 545: Maximum VPN BGP route limit reached [801001] [A801001]

Name	Max VPN eBGP route limit reached
Description	On-prem VPN gateway is advertising more than the maximum number of routes accepted by the VPN gateway over eBGP.
Alert message	`A total of num_routes routes received from the on-prem VPN gateway for connection "vpn_connection_name". Only max_num_routes routes will be processed.`

Cause	On-prem VPN gateway is advertising more than the maximum number of routes accepted by the VPN gateway over eBGP.
Impact	Connectivity between AZs via the VPN connection might be impacted.
Resolution	Consider aggregating more specific prefixes into larger prefixes to reduce the number of routes advertised.
KB Article	8469
Severity	Critical

Table 546: VPN IPSEC tunnel between VPN endpoints is down [801002] [A801002]

Name	VPN IPSEC Tunnel down
Description	The IPSEC tunnel between the VPN endpoints is down.
Alert message	`VPN IPSEC tunnel for VPN connection "vpn_connection_name" is down.`
Cause	The IPSEC tunnel between the VPN endpoints is down.
Impact	Connectivity between AZs via VPN connection will be impacted.
Resolution	Follow the on-prem VPN device vendor's troubleshooting steps. If you suspect that the problem is not with the on-prem VPN device, please contact Nutanix support.
KB Article	8466
Severity	Critical

Table 547: eBGP session between VPN endpoints is down [801003] [A801003]

Name	eBGP session between VPN gateway peers down
Description	The eBGP session between VPN endpoints is down.
Alert message	`eBGP session for VPN connection "vpn_connection_name" is down.`
Cause	The eBGP session between VPN endpoints is down.
Impact	VPN connectivity between the VPN endpoints will be impacted because routes cannot be exchanged.
Resolution	Follow the on-prem VPN device vendor's troubleshooting steps. If you suspect that the problem is not with the on-prem VPN device, please contact Nutanix support.
KB Article	8465
Severity	Critical

Table 548: Invalid routes received by VPN connection "{vpn_connection_name}". [801004] [A801004]

Name	Invalid routes received from on-prem VPN
Description	Invalid routes received from on-prem VPN gateway.

Alert message	`The following invalid routes received by VPN connection "vpn_connection_name" will be rejected: message.`
Cause	Invalid routes received from on-prem VPN gateway.
Impact	VPN connectivity may be impacted.
Resolution	Remove the invalid routes being advertised from on-prem VPN gateway over eBGP.
Severity	Warning

Table 549: Deletion of the Layer-2 subnet extension failed on the remote site. [801101] [A801101]

Name	Layer-2 subnet extension deletion failure
Description	Deletion of the Layer-2 subnet extension failed on the remote site.
Alert message	`Deletion of the Layer-2 subnet extension "#layer2_stretch_name@" failed on the remote site.`
Cause	Deletion of the remote Layer-2 subnet extension failed (possibly due to AZ connection issues).
Impact	The remote site's Prism Central will show the remote subnet as being extended to the local site, even though the subnet extension has been deleted locally.
Resolution	Delete the Layer-2 subnet extension on the remote site by logging into the remote site's Prism Central and navigating to the Subnets page.
KB Article	10310
Severity	Critical

Table 550: ANC version does not support Layer-2 subnet extension. [801102] [A801102]

Name	ANC version does not support Layer-2 subnet extension
Description	ANC version does not support Layer-2 subnet extension.
Alert message	`ANC version does not support Layer-2 subnet extension.`
Cause	ANC version does not support Layer-2 subnet extension.
Impact	ARP messages (and other unknown unicast messages) from peer AZ will be dropped, causing reachability issues between VMs on the stretched subnet across AZs.
Resolution	Upgrade the ANC to the right version that has support for Layer-2 subnet extension.
KB Article	10388
Severity	Critical

Table 551: VPN gateway version does not support Layer-2 subnet extension. [801103] [A801103]

Name	VPN gateway version does not support Layer-2 subnet extension
Description	VPN gateway version does not support Layer-2 subnet extension.
Alert message	`VPN gateway version does not support Layer-2 subnet extension "#layer2_stretch_name@".`
Cause	VPN gateway version does not support Layer-2 subnet extension.
Impact	ARP messages (and other unknown unicast messages) from peer AZ will be dropped, causing reachability issues between VMs on the stretched subnet across AZs.
Resolution	Upgrade the VPN gateway to the right version that has support for Layer-2 subnet extension.
KB Article	10389
Severity	Critical

Table 552: VPN connection associated with Layer-2 subnet extension not found. [801104] [A801104]

Name	VPN connection associated with Layer-2 subnet extension not found
Description	VPN connection associated with Layer-2 subnet extension not found.
Alert message	`VPN connection associated with Layer-2 subnet extension "#layer2_stretch_name@" not found.`
Cause	VPN connection associated with Layer-2 subnet extension was deleted.
Impact	VMs across AZs on the stretched subnet cannot communicate with each other.
Resolution	Delete the Layer-2 subnet extension and create it again when VPN connection is restored.
KB Article	10390
Severity	Critical

Table 553: Peer AZ involved in Layer-2 subnet extension is not reachable. [801105] [A801105]

Name	Peer AZ involved in Layer-2 subnet extension is not reachable
Description	Peer AZ involved in Layer-2 subnet extension is not reachable.
Alert message	`Peer AZ involved in Layer-2 subnet extension "#layer2_stretch_name@" is not reachable.`
Cause	The cause could either be VPN connection between the AZs is down, or Atlas service in the peer AZ is down (or not responsive).
Impact	VMs in the subnet extended across AZs may be unable to communicate with each other.

Resolution	Fix the underlying issue of VPN connection being down, or Atlas service in peer AZ being down.
KB Article	10391
Severity	Critical

Table 554: Subnet involved in Layer-2 subnet extension not found. [801106] [A801106]

Name	Subnet involved in Layer-2 subnet extension not found
Description	Subnet involved in Layer-2 subnet extension not found.
Alert message	`Subnet involved in Layer-2 subnet extension "#layer2_stretch_name@" not found.`
Cause	Subnet involved in Layer-2 subnet extension was deleted.
Impact	If subnet is not present, no vNICs can be created using this network UUID.
Resolution	Delete the Layer-2 subnet extension, and create it again after the subnet has been restored.
KB Article	10392
Severity	Critical

Table 555: CIDR of subnets involved in Layer-2 subnet extension don't match. [801107] [A801107]

Name	CIDR of subnets involved in Layer-2 subnet extension don't match
Description	CIDR of subnets involved in Layer-2 subnet extension don't match.
Alert message	`CIDR of subnets involved in Layer-2 subnet extension "#layer2_stretch_name@" don't match.`
Cause	CIDR of subnets involved in Layer-2 subnet extension don't match.
Impact	Some VMs may not be able to communicate with other VMs on the extended subnet.
Resolution	Modify one of the subnets to fix the CIDR mismatch.
KB Article	10393
Severity	Critical

Table 556: DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses. [801108] [A801108]

Name	DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses
Description	DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses.

Alert message	`DHCP of subnets involved in Layer-2 subnet extension "#layer2_stretch_name@" have overlapping IP addresses or include VPN interface IP addresses.`
Cause	DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses.
Impact	Some VMs may not be able to communicate with other VMs on the extended subnet.
Resolution	Modify the subnets involved in the Layer-2 subnet extension to ensure DHCP pools do not overlap, and do not contain the local or remote VPN interface IP.
KB Article	10394
Severity	Critical

Table 557: Local VPN interface IP involved in Layer-2 subnet extension is in use in peer AZ. [801109] [A801109]

Name	Local VPN interface IP involved in Layer-2 subnet extension is in use in peer AZ
Description	Local VPN interface IP involved in Layer-2 subnet extension is in use in peer AZ.
Alert message	`Local VPN interface IP involved in Layer-2 subnet extension "#layer2_stretch_name@" is in use in peer AZ.`
Cause	A vNIC in peer AZ was assigned the same IP address as VPN interface IP used in the local AZ.
Impact	Some VMs in the subnets involved in Layer-2 subnet extension will be unable to communicate with other VMs in peer AZ.
Resolution	Resolve the IP conflict by ensuring the VPN interface IP is not used for any user VMs in any of AZs involved in the Layer-2 subnet extension.
KB Article	10395
Severity	Critical

Table 558: Remote VPN interface IP involved in Layer-2 subnet extension is in use in this AZ. [801110] [A801110]

Name	Remote VPN interface IP involved in Layer-2 subnet extension is in use in this AZ
Description	Remote VPN interface IP involved in Layer-2 subnet extension is in use in this AZ.
Alert message	`Remote VPN interface IP involved in Layer-2 subnet extension "#layer2_stretch_name@" is in use in this AZ.`
Cause	A vNIC in this AZ was assigned the same IP address as VPN interface IP used in the remote AZ.
Impact	Some VMs in the subnets involved in Layer-2 subnet extension will be unable to communicate with other VMs in peer AZ.

Resolution	Resolve the IP conflict by ensuring the VPN interface IP is not used for any user VMs.
KB Article	10396
Severity	Critical

Table 559: Some IP address(es) are common across the subnets involved in the Layer-2 subnet extension. [801111] [A801111]

Name	Some IP address(es) are common across the subnets involved in the Layer-2 subnet extension
Description	Some IP address(es) are common across the subnets involved in the Layer-2 subnet extension
Alert message	`Some IP address(es) are common across subnets involved in Layer-2 subnet extension "#layer2_stretch_name@". Detailed error text: error_detail`
Cause	Some UVMs are allocated IP addresses that are in use in the peer AZ involved in Layer-2 subnet extension.
Impact	Some VMs in the subnets involved in Layer-2 subnet extension will be unable to communicate with other VMs in peer AZ.
Resolution	Resolve the IP conflict by ensuring the IP addresses allocated to UVMs are unique across the AZs involved in the Layer-2 subnet extension.
KB Article	10604
Severity	Critical

Table 560: Advanced Networking Controller is not Healthy [802001] [A802001]

Name	An Advanced Networking Controller service did not respond
Description	Atlas could not connect to an Advanced Networking Controller service.
Alert message	`The service at url did not respond`
Cause	The Advanced Networking Controller is not Healthy.
Impact	Ability to make network-related configurations may be impacted.
Resolution	Verify the health of the Advanced Networking Controller.
KB Article	9366
Severity	Critical

Table 561: Unable to resolve an Advanced Networking Controller service's DNS name [802002] [A802002]

Name	An Advanced Networking Controller service's DNS name was not resolvable
Description	Atlas was unable to resolve an Advanced Networking Controller service's DNS name.

Alert message	`The IP corresponding to url could not be resolved with the following DNS servers: dns_servers`
Cause	PC nameserver configuration is incorrect.
Impact	Ability to make network-related configurations may be impacted.
Resolution	Verify nameserver configuration on PC.
KB Article	9365
Severity	Critical

Table 562: Unable to apply Virtual Private Cloud Routing Policy. [802003] [A802003]

Name	A Virtual Private Cloud Routing Policy of type reroute is inactive
Description	Atlas was unable to apply a VPC Routing Policy of type reroute.
Alert message	`Routing Policy with priority priority and type reroute in VPC vpc_name is not applied as VM/service with reroute IP: reroute_ip is down.`
Cause	VM(s) with reroute IP(s) does not exist or is powered off.
Impact	Traffic in VPC will not hit the inactive Routing Policy.
Resolution	(A) Create VM with reroute IP in the given VPC. (B) Power on VM with reroute IP in the given VPC.
KB Article	9484
Severity	Warning

Table 563: ID Firewall lost connectivity to domain controller [803003] [A803003]

Name	ID Firewall is unable to connect to one or more domain controllers
Description	ID Firewall is unable to connect to one or more domain controllers used for scraping identities.
Alert message	`Domain controller domain_controller on domain domain is not reachable and accepting LDAP or WMI connections. Details: details.`
Cause	The domain controller may not be reachable via the network.
Cause	The domain controller may not be accepting LDAP/WMI connections due to permissions issues.
Cause	The domain controller may not be running or fully booted up.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until access to the domain controller is restored.
Resolution	Check network connectivity from Prism Central to the domain controller. Please refer article KB-10219 for more details.
Resolution	Check that the domain controller is fully up and accepting WMI and LDAP connections. Please refer article KB-10219 for more details.
Resolution	Check that the service account registered for ID Firewall is active and has permissions for both WMI and LDAP access, and that the correct password is entered in Prism Central. Please refer article KB-10219 for more details.

KB Article	10219
Severity	Warning

Table 564: ID Firewall did not recover state after reconnecting to a temporarily unreachable domain controller [803005] [A803005]

Name	ID Firewall state recovery failed
Description	ID Firewall was unable to recover state after reconnecting to a temporarily unreachable domain controller.
Alert message	`ID Firewall was unable to recover state from domain controller domain_controller on domain domain. Details: details.`
Cause	Recovery failed because either too much time has passed since the domain controller was last reachable, or the event log on the domain controller has rolled over and the events necessary to recover are no longer present.
Impact	The state enforced by ID Firewall may be inconsistent, and therefore applied policies may not be properly enforced until action is taken.
Resolution	In order to ensure security policies are properly enforced, all active Nutanix VDI VM users should log out and log back into their VMs. Please refer KB-10220 for further details.
KB Article	10220
Severity	Warning

Table 565: ID Firewall service account is invalid [803007] [A803007]

Name	Service account used for ID Firewall is invalid
Description	ID Firewall has detected that the service account configured is not valid.
Alert message	`ID Firewall service account account@domain is invalid. Details: details.`
Cause	The service account password may have been changed.
Cause	The service account may have been deleted.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until the service account is valid again.
Resolution	Check that the configured service account password is up-to-date and update it in Prism if needed.
Resolution	Check that the configured service account exists, and update it in Prism if needed.
Severity	Warning

Nutanix Files

Table 566: File Server Space Usage High [160000] [A160000]

Name	File Server Space Usage
Description	The usage of file server space is high.
Alert message	`File Server space usage for file_server_name is at usage_pct%.`
Cause	File Server Storage utilization is high.
Impact	If storage usage keeps growing, the shares will be marked as read-only.
Resolution	Expand File Server storage size. Ask users to delete unused files. Refer to KB article 8475 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8475
Severity	Warning

Table 567: File Server Space Usage Critical. [160001] [A160001]

Name	File Server Space Usage Critical
Description	File server space usage critical.
Alert message	`File Server space usage for file_server_name is at usage_pct%.`
Cause	File Server storage utilization has reached a critical value.
Impact	The shares will be marked as read-only until there is more free space on the File Server.
Resolution	Expand File Server storage size. Refer to KB article 8475 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8475
Severity	Critical

Table 568: File Server Unreachable. [160002] [A160002]

Name	File Server Unavailable Check
Description	File server is unreachable.
Alert message	`File server file_server_name is unreachable.`
Cause	FSVMs are not reachable.
Impact	No operations can be executed against the File Server.
Resolution	From Controller VM run 'ncc health_checks fileserver_checks run_all' and resolve the reported issues. If you are unable to resolve the issue reported by NCC, contact Nutanix Support for assistance.
Severity	Warning

Table 569: File Server storage is not available. [160003] [A160003]

Name	File Server Storage Status
Description	File Server storage is not available.
Alert message	`Storage for File Server file_server_name is unavailable.`
Cause	File Server storage is unavailable due to network connectivity issues between FSVM and CVM.
Impact	Read and write operations on the File Server will fail.
Resolution	From Controller VM run 'ncc health_checks run_all' and resolve the reported issues. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 570: File Server scale-out failed. [160004] [A160004]

Name	File Server Scale-out Status
Description	File Server scale-out failed.
Alert message	`File Server file_server_name scale-out failed because of reason`
Cause	See details in the alert message.
Impact	Unable to scale-out the file server.
Resolution	From Controller VM run 'ncc health_checks run_all' and resolve the reported issues. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 571: File server could not join the AD Domain [160005] [A160005]

Name	File Server Join Domain Status
Description	File server could not join the AD domain.
Alert message	`File Server file_server_name could not join the Domain because reason.`
Cause	See details in the alert message. Error could be due to one of the following reasons: Domain credentials are invalid. Correct Port for connecting to AD is not being used. File Server AD Object already exists in the domain.
Impact	Domain operations can not be performed on the File Server as it has not joined the domain
Resolution	The resolution is to identify the reason from one of the possible causes in the cause list. Please refer Nutanix Files Guide 'Pre-requisite' section for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 572: Node Failed To Join Domain [160006] [A160006]

Name	File Server Node Join Domain Status
Description	The node could not join the domain.
Alert message	`The node could not join the domain for file server file_server_name as reason`
Cause	DNS is not correctly configured or is not reachable. Domain Controllers are not reachable or the reachable Domain controller is having issues. Domain credentials are invalid. Correct Port for connecting to AD is not being used. File Server AD Object already exists in the domain.
Impact	The node could not join the domain.
Resolution	The resolution is to identify the reason from one of the possible causes in the cause list. Please refer Nutanix Files Guide 'Pre-requisite' section for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 573: File Server Time Difference High [160007] [A160007]

Name	FSVM Time Drift Status
Description	The time drift between the FSVMs is beyond the acceptable range
Alert message	`For file_server_name, time drift between the FSVMs (lower_time_ip and higher_time_ip) is more than the acceptable value of time_difference_limit_secs`
Cause	NTP is not configured correctly or the NTP service is not functioning.
Impact	The Nutanix Files cluster may become inaccessible
Resolution	Check that the NTP service is running and is reachable from the FSVMs Refer KB article 8135 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8135
Severity	Warning

Table 574: File Server Storage Cleanup Failed [160012] [A160012]

Name	File Server Storage Cleanup Failure
Description	Failed to clean up storage for the file server.
Alert message	`Storage message for the file server file_server_name is not cleaned up.`
Cause	Acropolis service may be down on the cluster.
Impact	File server storage will not be released.
Resolution	Contact Nutanix support if issue still persists or assistance needed.
Severity	Critical

Table 575: File server cannot connect with AD server [160013] [A160013]

Name	File Server AD Connectivity Failure
Description	File server cannot connect with AD server with configured information
Alert message	`File server file_server_name cannot connect with AD server`
Cause	The machine account credentials may have been changed. There is a network connectivity issue for the AD server.
Impact	Domain operations cannot be performed for the file server. Client connectivity with file-server is disabled.
Resolution	Check that the AD server is reachable. Refer to KB article 8122 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8122
Severity	Critical

Table 576: File Server performance optimization recommended [160015] [A160015]

Name	File Server Performance Optimization Recommended
Description	File server has a recommendation to optimize performance by using scale-up, scale-out or rebalance.
Alert message	`A recommendation is available to optimize the performance on one or more nodes of file server file_server_name.`
Cause	File server has one or more nodes under extensive load.
Impact	File server performance may be impacted and new client connections may be refused.
Resolution	Run 'Performance optimization' for the specified file server. Please refer Nutanix Files 'System limits' section in the release notes for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	7035
Severity	Warning
Severity	Info

Table 577: User Quota Assignment Failed [160016] [A160016]

Name	File Server Quota allocation failed for user
Description	Failed to assign the specified quota to the user
Alert message	`Failed to apply quota for user: user_name on share: share_name of File Server file_server_name`
Cause	The user has used more space on the File Server than the specified quota value.
Impact	User has no quota limit

Resolution	Notify user to reduce their usage below the quota value. Notify user to increase the quota allocation for the user.
Severity	Warning

Table 578: Share utilization reached configured limit [160017] [A160017]

Name	Share Utilization Reached Configured Limit
Description	Share is no longer writable.
Alert message	`Utilization on share_name on file server file_server_name has reached the configured limit`
Cause	Share utilization reached its configured limit
Impact	Writes to the share will fail until one of the actions is taken.
Resolution	Expand File Server storage size. Refer to KB article 8475 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8475
Severity	Critical

Table 579: File Server failed to get updated CVM IP address. [160018] [A160018]

Name	File Server CVM IP update failed
Description	File server Controller VM IP update failed.
Alert message	`The file server file_server_name could not get updated CVM IP address.`
Cause	AOS cluster is not able to contact file server.
Impact	File server is not reachable
Resolution	Ensure correct Cluster IP Address Configuration procedure is followed. Refer to KB article 2290 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	2290
Severity	Critical

Table 580: Appropriate Site Not Found in Active Directory [160019] [A160019]

Name	File Server Site Not Found
Description	Unable to determine an appropriate site in Active Directory
Alert message	`File server file_server_name's client network is not mapped to a site on the Active Directory.`
Cause	File server client network is not part of any site in Active Directory.
Impact	File server may take a long time to join the domain, depending on which domain controller is selected.

Resolution	Add the client network of file server to a local site in Active Directory for optimal performance.
Severity	Warning

Table 581: File server DNS Updates Pending [160020] [A160020]

Name	File Server DNS Updates Pending
Description	DNS updates are pending after a file server operation.
Alert message	`File server file_server_name's DNS entries need to be programmed on nameservers. message`
Cause	Failed to add or remove DNS entries
Impact	The file server cannot be used by clients by its name, but can be used with the IP addresses.
Resolution	Correct the required DNS entries under 'DNS' on the file server page. Refer KB article 3722 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3722
Severity	Warning
Severity	Info

Table 582: File Server activation failed [160021] [A160021]

Name	File Server Activation Failed
Description	File Server Activation Failed
Alert message	`File server file_server_name activation failed due to reason`
Cause	Check alert message for details
Impact	File Server is not usable.
Resolution	Try to activate file server again, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 583: Failed to set VM-to-VM anti-affinity rule [160023] [A160023]

Name	Failed To Set VM-to-VM Anti Affinity Rule
Description	Failed to set VM-to-VM anti affinity rule
Alert message	`Failed to set VM-to-VM anti-affinity rule for file_server_name. reason_and_resolution_msg`
Cause	More than one FSVM of a File Server entity is may be deployed hosted on a single physical host.

Impact	More than one FSVM of a File Server entry is hosted on a single physical host, which will cause data unavailability in case that host fails.
Resolution	Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 584: Discovery of iSCSI targets failed. [160025] [A160025]

Name	File Server iSCSI Discovery Failure
Description	Failed to discover iSCSI targets on the CVM during the discovery process.
Alert message	`Discovery of iSCSI targets failed for file server file_server_name`
Cause	The 'external_data_services_ip' or the CVM IP addresses are not reachable from the FSVM.
Impact	Share operations and File server HA will fail.
Resolution	Please refer article 8216 for more details on Changing Data Services IP Address. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8216
Severity	Critical
Severity	Warning

Table 585: File Server upgrade failed [160026] [A160026]

Name	File Server Upgrade Failed
Description	File Server Upgrade Failed.
Alert message	`Upgrade of file server file_server_name failed due to reason.`
Cause	Check alert message for details
Impact	File server has not been upgraded.
Resolution	From Controller VM run 'ncc health_checks fileserver_checks run_all' and resolve the reported issues. If you are unable to resolve the issue reported by NCC, contact Nutanix Support for assistance.
Severity	Warning

Table 586: Incompatible File Server activated [160028] [A160028]

Name	Incompatible File Server Activation
Description	Activating an incompatible File Server
Alert message	`Incompatible File server file_server_name is activated`
Cause	File server is activated which is incompatible with current AOS version
Impact	Next AOS upgrade will fail.

Resolution	Upgrade file server to the latest. Please refer Software Product Interoperability page on portal for reference. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 587: File Server in heterogeneous state. [160032] [A160032]

Name	File Server In Heterogeneous State
Description	File server in heterogeneous state. Nodes do not match in their CPU or memory configuration.
Alert message	`File server file_server_name is in heterogeneous state, as File Server VM configuration update operation failed due to reason`
Cause	Check alert message for details.
Impact	Performance of the file server is less than optimal.
Resolution	Get the largest configured CPU and memory values among VMs of reported File Server in VM page. Free up resources on hypervisor hosts and retry File Server VM configuration update with these values in File Server page. Contact Nutanix support if issue still persists or assistance needed.
Severity	Critical

Table 588: Failed to correct File Server data and meta data inconsistencies [160034] [A160034]

Name	Failed To Run File Server Metadata Fixer Successfully
Description	Failed to Run File Server Metadata Fixer tool successfully
Alert message	`Failed to run File Server file_server_name Metadata fixer task successfully. reason_and_resolution_msg`
Cause	Metadata fixer task might have timed out
Impact	Some of the top level directories would be unavailable on file server
Resolution	Run 'afs share.fix_tld_metadata' command form one of the FSVM shell to fix the metadata inconsistency. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 589: File server share deletion failed [160035] [A160035]

Name	File Server Share Deletion Failed
Description	Failed to delete share.
Alert message	`Failed to delete share share_name on File Server file_server_name. message`
Cause	Multiple causes possible. Most likely, a failure in destroying file systems from FSVMs.

Impact	Unable to clean up share storage space.
Resolution	Please retry Share delete operation. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 590: File server compatibility check skipped [160036] [A160036]

Name	Skipped File Server Compatibility Check
Description	File server compatibility check skipped
Alert message	`File server compatibility check skipped`
Cause	File server compatibility check was skipped during AOS upgrade
Impact	Some management functions of the file server may not be available.
Resolution	Upgrade file server to the latest. Please refer Software Product Interoperability page on portal for reference. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 591: Failed to add one or more file server administrator users or groups [160038] [A160038]

Name	Failed to add one or more file server admin users or groups
Description	Failed to add one or more users or groups as file server administrators
Alert message	`File server file_server_name: failure_msg Update file server administrators on the file server page`
Cause	One or more users or groups could not be resolved on Active Directory.
Impact	File server administrator users or groups may not function as administrators
Resolution	Update file server administrators on the file server page
Severity	Info

Table 592: Maximum connections limit about to reach on a file server VM [160039] [A160039]

Name	Maximum connections limit reached on a file server VM
Description	Maximum connections limit is about to reach on a file server VM
Alert message	`File server VM fsvm_name on file server file_server_name has reached conn_percentageconnections limit.`
Cause	File server has one or more nodes under extensive load.
Impact	File server performance may be impacted and new client connections may be refused.

Resolution	Run 'Performance optimization' for the specified file server. Refer to KB article 7371 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	7371
KB Article	7371
Severity	Critical
Severity	Warning

Table 593: File Server Clone failed [160040] [A160040]

Name	File Server Clone Failed
Description	File Server Clone Failed
Alert message	`File server file_server_name clone failed due to reason`
Cause	Check alert message for details
Impact	Clone File Server not available.
Resolution	Check alert message, Retry clone operation after rectifying, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 594: File Server rename failed [160041] [A160041]

Name	File Server Rename Failed
Description	File Server Rename Failed
Alert message	`File server rename from file_server_name to file_server_new_name failed due to reason`
Cause	Check alert message for details
Impact	File server could not be renamed.
Resolution	Check alert message, Retry rename operation after rectifying, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 595: File Server AntiVirus - ICAP Server Down [160042] [A160042]

Name	File Server AntiVirus - ICAP Server Down
Description	ICAP server is not responding to scan requests.
Alert message	`ICAP server scan_server_host is not responding from file server file_server_name.`
Cause	Failed to reach ICAP server from FSVM.
Impact	File server antivirus scanning is degraded.

Resolution	Check ICAP server and connectivity between FSVMs to ICAP server. Refer to KB article 8414 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8414
Severity	Warning

Table 596: File Server AntiVirus - All ICAP Servers Down [160043] [A160043]

Name	File Server AntiVirus - All ICAP Servers Down
Description	All configured ICAP servers are not responding to scan requests.
Alert message	`All configured ICAP servers are not responding from file server file_server_name.`
Cause	Failed to reach all configured ICAP server from FSVM.
Impact	File server antivirus scanning is not possible. Access is decided by scan policy.
Resolution	Check ICAP server and connectivity between FSVMs to ICAP server. Refer to KB article 8414 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8414
Severity	Critical

Table 597: File Server network change failed [160044] [A160044]

Name	File Server Network Change Failed
Description	File Server Network Change Failed
Alert message	`File server file_server_name network change failed due to reason`
Cause	Check alert message for details
Impact	File Server is not usable.
Resolution	Try to change file server network again, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 598: File Server AntiVirus - Scan Queue Full on FSVM [160045] [A160045]

Name	File Server AntiVirus - Scan Queue Full on FSVM
Description	Antivirus Scan Queue is Full on FSVM.
Alert message	`Antivirus Scan Queue is Full on FSVM nvm_name of file server file_server_name.`
Cause	ICAP Daemon is not able to keep up with in coming requests.
Impact	File server antivirus scanning is slow, clients may observe high latency.

Resolution	Check ICAP server status. Check number of ICAP servers configured. Check workload on FSVM.
Severity	Critical

Table 599: File Server AntiVirus - Scan Queue Piling Up on FSVM [160046] [A160046]

Name	File Server AntiVirus - Scan Queue Piling Up on FSVM
Description	Antivirus Scan Queue is Piling Up on FSVM.
Alert message	`Antivirus Scan Queue is Piling Up on FSVM nvm_name of file server file_server_name.`
Cause	Incoming scan request rate is higher than the scan processing rate.
Impact	File server antivirus scanning is slow, clients may observe high latency.
Resolution	Check ICAP server status. Check number of ICAP servers configured. Check workload on FSVM.
Severity	Warning

Table 600: File Server AntiVirus - Excessive Quarantined / Unquarantined Files [160047] [A160047]

Name	File Server AntiVirus - Excessive Quarantined / Unquarantined Files
Description	Excessive Quarantined / Unquarantined Files on File Server.
Alert message	`File Server file_server_name has a lot of quarantined / unquarantined files.`
Cause	Number of quarantined / unquarantined files on file server is approaching limit.
Impact	File server will stop scanning the files once the limit is reached. Default action specified in the scan policy will be applied.
Resolution	Delete unwanted quarantined / unquarantined files, check ICAP server for false positives or check system for virus activity.
Severity	Warning

Table 601: File Server AntiVirus - Quarantined / Unquarantined Files Limit Reached [160048] [A160048]

Name	File Server AntiVirus - Quarantined / Unquarantined Files Limit Reached
Description	Quarantined / Unquarantined Files Limit is Reached on File Server.
Alert message	`File Server file_server_name has reached quarantined / unquarantined files limit.`
Cause	Number of quarantined / unquarantined files on file server reached limit.
Impact	File server stopped scanning files. Default action in the scan policy is applied.

Resolution	Delete unwanted quarantined / unquarantined files, check ICAP server for false positives or check system for virus activity. Refer to KB article 8416 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8416
Severity	Critical

Table 602: File Server time is out of sync with the Active Directory [160049] [A160049]

Name	File Server time is out of sync with the Active Directory
Description	File server time is out of sync with the Active Directory domain controllers.
Alert message	`Time drift between the file server VMs and the Active Directory is at time_drift_secs seconds on file server file_server_name. If the drift becomes more than 300 seconds, the clients may get disconnected from the fileserver.`
Cause	NTP is not configured correctly or the NTP service is not functioning.
Impact	The Nutanix Files cluster may become inaccessible.
Resolution	Check that the NTP service is running and the NTP servers are reachable from the file server VMs. Refer to KB article 8404 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8404
Severity	Warning

Table 603: File Server User Management Configuration Failed [160050] [A160050]

Name	File Server User Management Configuration Failed
Description	Given user management options could not be configured for the file server.
Alert message	`The values provided for File server file_server_name's user management options are incorrect.`
Cause	Failed to configure user management with given options
Impact	The file server may have restricted or insecure access from various clients.
Resolution	Update correct entries against fields in the User Management tab. Refer to KB article 9123 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	9123
Severity	Warning

Table 604: /home partition usage on a file server VM higher than threshold [160051] [A160051]

Name	File Server Disk Usage Warning

Description	Space consumption under /home partition on file server VM is higher than the set threshold
Alert message	`/home partition usage on File server VM fsvm_name on file server file_server_name at usage_pct (greater than threshold threshold_pct).`
Cause	File server has one or more nodes under extensive load.
Impact	Deleting a share from a File server fails.
Resolution	Remove unwanted data or files from under /home for the specified file server. Refer to KB article 8478 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8478
Severity	Critical
Severity	Warning

Table 605: File server DNS records cannot be refreshed [160052] [A160052]

Name	File Server DNS - Unable to refresh DNS records
Description	File server DNS records periodic refresh failed
Alert message	`File server file_server_name DNS records cannot be updated`
Cause	DNS entries may not be created with Nutanix Files computer account due to insufficient permissions.
Impact	If scavenging is active on the DNS servers, then Nutanix Files DNS records may get removed. This can cause clients to not connect with Nutanix Files using its DNS hostname.
Resolution	Create static DNS entries for the fileserver manually using 'DNS' button. Update required permissions for the Nutanix Files computer account. Refer to KB article 8405 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8405
Severity	Critical

Table 606: File Server Share Backup diff path translation failed [160053] [A160053]

Name	File Server Backup Diff Path Translation Warning
Description	Backup diff found inodes whose path translation failed
Alert message	`Backup diff found inodes whose path translation failed file_server_name, error information: failure_msg`
Cause	File server share may have hardlinks.
Impact	Some files may not have been backed up.
Resolution	Refer to KB article 8874 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8874

Severity	Warning

Table 607: File Server Partner Server Connectivity Down [160054] [A160054]

Name	File Server Partner Server Connectivity Down
Description	Partner server is not responding to file notifications.
Alert message	`Partner server partner_server_host is not responding from file server file_server_name.`
Cause	Failed to reach partner server from File Server VM.
Impact	File server stopped notifying file operation events.
Resolution	Ensure that the partner server is functioning and that there is connectivity between File Server VMs and the partner server.
Severity	Critical

Table 608: File Server DNS resolver IP connectivity Failure [160055] [A160055]

Name	File Server DNS resolver IP connectivity Failure
Description	File server cannot resolve its domain name using configured DNS resolver IP addresses.
Alert message	`File server file_server_name cannot resolve its domain name using configured DNS resolver IP addresses. message`
Cause	None of the configured DNS resolver IP addresses are reachable. None of the configured DNS resolver IP addresses are resolving the file server domain name.
Impact	File server services may be unavailable and new clients may not be authenticated.
Resolution	Ensure that the servers specified by DNS resolver IP addresses are functional. Refer to KB article 8406 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8406
Severity	Warning

Table 609: File Server PD action to incompatible Remote Site AOS [160057] [A160057]

Name	File Server PD action warning for non compatible Remote Site AOS
Description	Remote Site AOS version is not compatible with File Server version.
Alert message	`File server file_server_name protection domain protection_domain_name enabled on non-compatible Remote site remote_site_name`
Cause	Remote Site AOS version is not compatible with File Server version.
Impact	Migrate-to or Activate-on Remote Site would fail.

Resolution	Upgrade Remote Site AOS version similar or above of primary site. Please refer Software Product Interoperability page on portal for reference. Contact Nutanix support if issue still persists or assistance needed.
Severity	Critical
Severity	Warning

Table 610: File Server NTP servers connectivity failure [160058] [A160058]

Name	File Server NTP servers connectivity failure
Description	File server cannot connect to NTP server.
Alert message	`File Server file_server_name cannot reach NTP servers specified. reason`
Cause	None of the configured NTP servers are reachable.
Impact	File server services may become inaccessible.
Resolution	Ensure that the NTP servers are functional and are reachable from the file server VMs. Refer to KB article 8135 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8135
Severity	Critical
Severity	Warning
Severity	Info

Table 611: File server services got interrupted [160059] [A160059]

Name	File Server services got interrupted
Description	One of the fileserver services has restarted.
Alert message	`File server file_server_name services got interrupted. Error Message: message`
Cause	One of the fileserver services has restarted.
Impact	Some fileserver service would be unavailable to NAS clients momentarily.
Resolution	If the service continues to be down, consider restarting minerva_ha service.
Severity	Warning

Table 612: Share usage reaching to configured limit [160061] [A160061]

Name	File Server Share Space Usage Reaching Configured Limit
Description	Share will be no longer be writable.
Alert message	`Share share_name on file server file_server_name is approaching capacity`
Cause	Share usage reaching its configured limit
Impact	Writes to the share will fail if share gets full.

Resolution	Increase storage allocation to the share or direct users to free up space by removing unneeded files. Please refer KB article 8690 for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
KB Article	8690
Severity	Warning

Table 613: Pulse cannot connect to REST server endpoint on File Server [160062] [A160062]

Name	File Server Failed to Reach Pulse REST Endpoint
Description	Pulse cannot connect to REST server endpoint on File Server.
Alert message	`Pulse cannot connect to REST server endpoint on File Server file_server_name. Connection Status: connection_status, Pulse Enabled: enabled, Error Message: message`
Cause	REST server endpoint is not reachable from Pulse.
Impact	Data driven serviceability and customer support cannot be performed for affected File Server.
Resolution	Ensure that the REST server endpoint is reachable from Pulse on File Server.
Severity	Info

Table 614: Storage pool usage reaching its limit [160063] [A160063]

Name	File Server Storage Pool Usage Reaching Configured Limit
Description	Shares on storage pool will be no longer writable.
Alert message	`share_names shares on file server file_server_name are about to run out of space. command`
Cause	Storage pool usage is reaching its limit
Impact	Writes to the shares may fail if their storage pool gets full.
Resolution	Increase storage allocation to the share or direct users to free up space by removing unneeded files. Please refer KB article 8690 for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 615: Duplicate IP address detected for a File Server VM. [160068] [A160068]

Name	File Server Duplicate IP Address Detected
Description	File server VMs have IP address conflict.
Alert message	`Duplicate IP address detected for a file server VMs for file_server_name. Error Message: message`
Cause	Another host has started using a file server VM IP address.
Impact	Some file server operations may fail due to invalid network configuration.

Resolution	Update IP address on host which has an IP address conflict with file server VM. Refer to KB article 8951 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8951
Severity	Warning

Table 616: File Server Unique Fsid failure [160069] [A160069]

Name	File Server Unique Fsid Failure
Description	Assigning Unique Fsid failed
Alert message	`Found unique fsid failure file_server_name, error information: failure_msg`
Cause	Duplicate Unique Fsid present
Impact	NFS handles may not be unique.
Resolution	Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 617: File Server VM Time Drift from NTP Servers [160072] [A160072]

Name	File Server VM Time Drift from NTP Servers
Description	Difference between FSVM clock and NTP servers is too large
Alert message	`File Server VM fsvm_name clock is different from NTP servers for time_diff_secs seconds. Acceptable time difference time_diff_limit_secs seconds.`
Cause	NTP servers are not publishing time correctly. File server VM time is not syncing with NTP servers.
Impact	File server cluster may become inaccessible
Resolution	Verify NTP servers are functional and reachable. Refer to KB article 8135 for more details. Contact Nutanix support if the issue still persists or assistance is needed.
KB Article	8135
Severity	Warning
Severity	Info

Table 618: File Server Service in Crash Loop [160074] [A160074]

Name	File server service in crash loop
Description	File server service in crash loop
Alert message	`File server service crashed_service_name has crashed crash_count times in crash_interval_secs seconds.`

Cause	File server service has crashed multiple times in short interval
Impact	Some file server features may not be functioning
Resolution	Refer to KB article 8418 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8418
Severity	Critical

Table 619: File server service got interrupted [160101] [A160101]

Name	File server service got interrupted
Description	One of the fileserver services has restarted.
Alert message	`File server file_server_name service got interrupted. Error Message: message`
Cause	One of the fileserver services has restarted.
Impact	Some fileserver service would be unavailable to NAS clients momentarily.
Resolution	If the service continues to be down, consider restarting minerva_ha service.
Severity	Warning

Table 620: File Server Active Directory operation failed during Disaster Recovery [160103] [A160103]

Name	File Server Active Directory operation failed during Disaster Recovery
Description	Active Directory operations during failover or failback operation as part of Disaster Recovery failed.
Alert message	`File server file_server_name Disaster Recovery operation failed to update Active Directory. message`
Cause	Various. See message for details.
Impact	Clients connected to previous file server may not work with the current file server after Disaster Recovery.
Resolution	Various. See message for details.
KB Article	9805
Severity	Critical

Table 621: File Server Disaster Recovery - target site not reachable [160105] [A160105]

Name	File Server Disaster Recovery - target site not reachable
Description	File server target site is not reachable for Disaster Recovery
Alert message	`File server file_server_name Disaster Recovery - Target site is not reachable. message`

Cause	Network connectivity issue between the source and the target.
Impact	Disaster Recovery operation may be impacted. Share data may not be replicated on time.
Resolution	Please check the network connectivity between the source and the target.
KB Article	9857
Severity	Critical

Table 622: File Server Disaster Recovery - not able to meet RPO [160106] [A160106]

Name	File Server Disaster Recovery - not able to meet RPO
Description	File server Disaster Recovery feature is not able meet current RPO standard.
Alert message	`File server file_server_name Disaster Recovery feature is not able meet current RPO standard. message`
Cause	An error on source or target file server can cause replication to proceed slowly.
Impact	Disaster Recovery related share replication may not be up to date as per the RPO standard.
Resolution	Please check the alert message for details.
KB Article	9872
Severity	Warning

Table 623: File Server Disaster Recovery - both source and target are active file servers [160107] [A160107]

Name	File Server Disaster Recovery - both source and target are active file servers
Description	File server Disaster Recovery feature has both source and target file servers set as active with read-write shares.
Alert message	`File server file_server_name Disaster Recovery feature has both source and target file servers set as active with read-write shares. message`
Cause	Configuration error during Disaster Recovery setup on source and/or target file servers.
Impact	Data inside shares may be modified on both source and target and may not remain in sync.
Resolution	Please make the target file server as standby with read-only shares.
KB Article	9873
Severity	Warning

Table 624: File Server connected share path was not found [160108] [A160108]

Name	File Server Connected Share Path not found

Description	Path for connected share is not available.
Alert message	`File server file_server_name connected share child_share_name path submount_path was not found.`
Cause	The submount path doesn't exist.
Impact	Share won't be accessible from submount path
Resolution	Remove submount path and reconfigure the share with new or recreated submount path
Severity	Warning

Table 625: File Server Disaster Recovery - source file server has more nodes than target file server [160110] [A160110]

Name	File Server Disaster Recovery - source file server has more nodes than target file server
Description	File server Disaster Recovery - number of nodes on source and target file servers must be the same.
Alert message	`File server file_server_name Disaster Recovery - Target site expansion required. message`
Cause	Number of file server nodes on source and target has changed due to scale-out operation.
Impact	Data inside shares may not remain in sync.
Resolution	Expand or reduce the node count on source and target file servers to make the node count identical on both file servers.
KB Article	9858
Severity	Critical

Table 626: File Server Disaster Recovery - target file server has more nodes than source file server [160111] [A160111]

Name	File Server Disaster Recovery - target file server has more nodes than source file server
Description	File server Disaster Recovery - number of nodes on source and target file servers must be the same.
Alert message	`File server file_server_name Disaster Recovery - Source site expansion required. message`
Cause	Number of file server nodes on source and target has changed due to scale-out operation.
Impact	Data inside shares may not remain in sync.
Resolution	Expand or reduce the node count on source and target file servers to make the node count identical on both file servers.
KB Article	9858
Severity	Warning

Table 627: File Server Disaster Recovery - failover operation failed [160112] [A160112]

Name	File Server Disaster Recovery - failover operation failed
Description	File server - failed to take over primary file server during Disaster Recovery Failover operation.
Alert message	`Failed to take over primary file server during Disaster Recovery Failover operation on file_server_name. message`
Cause	Various. Please see alert message for more details.
Impact	Primary file server clients may not access the shares/exports.
Resolution	Correct the error and retry the operation.
Severity	Warning

Table 628: File Server Disaster Recovery - failback operation failed [160113] [A160113]

Name	File Server Disaster Recovery - failback operation failed
Description	File server - failed to take over current file server resource during Disaster Recovery Failback operation.
Alert message	`Failed to take over current file server resources during Disaster Recovery Failback operation on file_server_name. message`
Cause	Various. Please see alert message for more details.
Impact	Current file server clients may not access the shares/exports.
Resolution	Correct the error and retry the operation.
Severity	Warning

Table 629: File Server Disaster Recovery - Source and Target Have Mismatched Protocols Configured [160114] [A160114]

Name	File Server Disaster Recovery - Source and Target Have Mismatched Protocols Configured
Description	File server - failed to establish source and target relationship for shares due to mismatched protocol configuration.
Alert message	`One or more shares are not protected for file_server_name. message`
Cause	Source and target file servers are not configured to support the same protocols.
Impact	One or more shares will not be protected.
Resolution	Configure source and target file servers with the same protocols.
Severity	Warning

Table 630: File Server Disaster Recovery - SSR snapshot schedule migration operation failed [160116] [A160116]

Name	File Server Disaster Recovery - SSR snapshot schedule migration operation failed
Description	File server - failed to migrate SSR snapshot schedule.
Alert message	`Failed to migrate SSR snapshot schedules on file_server_name. message`
Cause	The Alert message includes the cause and the details of the SSR schedule migration failure.
Impact	SSR snapshots will not be replicated to target.
Resolution	Run the SSR snapshot migration via afs cli. Please refer to KB article 11593 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11593
Severity	Warning

Table 631: File Server Domain is Offline [160120] [A160120]

Name	File Server Domain is Offline
Description	File server domain is offline.
Alert message	`File Server file_server_name domain domain_name is offline on node node_name.`
Cause	There is a network connectivity issue for the AD server.
Impact	Domain operations cannot be performed for the file server. Client connectivity with file-server is disabled.
Resolution	Check that the AD server is reachable. Refer to KB article 10074 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	10074
Severity	Critical

Table 632: File Server Domain Controller is detected slow. [160121] [A160121]

Name	File Server Domain Controller is detected slow
Description	File Server domain controller is detected slow.
Alert message	`File Server file_server_name domain controller dc_name is detected slow on node node_name.`
Cause	There could be network connectivity issue the AD server or the kerberos user ticket is getting resolved after multiple redirections to child domain having the information.
Impact	Domain operations cannot be performed for the file server. Client connectivity with file-server is disabled.

Resolution	Check that the AD server is reachable. Refer to KB article 10076 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	10076
Severity	Warning

Table 633: AIDE Baselines configuration changes on File Server. [160130] [A160130]

Name	File Server AIDE Baseline Configuration Changes Check
Description	Checks the log files of aide for changes to baseline configs if aide is enabled
Alert message	`Detected changes to baseline configs on a File Server. message`
Cause	Baseline configurations changed on a File Server.
Impact	Aide execution may fail depending on the files that were changed
Resolution	Check changed files and revert their changes if needed
KB Article	10461
Severity	Warning

Table 634: Remote syslog server forwarding failed on a File Server. [160131] [A160131]

Name	File Server Remote Syslog Server Forwarding Check
Description	Check remote syslog server log forwarding on a File Server.
Alert message	`Remote syslog server forwarding failures detected on a File Server. message`
Cause	Remote syslog server forwarding failed on a File Server.
Impact	Remote syslog server will not be receiving log output from the File Server.
Resolution	Check for forwarding failure causes with rsyslog server.
KB Article	10735
Severity	Warning

Table 635: File Server Tiering Service is not healthy [160135] [A160135]

Name	File Server Tiering Service is not healthy
Description	Tiering service is not in a healthy state and can lead to temporary unavailability of tiered files.
Alert message	`message`
Cause	Tiering service might be overloaded or is unavailable.
Impact	Delayed processing of tiering operations and can lead to temporary unavailability of tiered files.
Resolution	Check file server cluster health and refer to KB11073 for more details.

KB Article	11073
Severity	Warning

Table 636: File server Tiering Service failed to mark the bucket for cleanup [160136] [A160136]

Name	File server Tiering Service failed to mark the bucket for cleanup
Description	Tiering service failed to mark the bucket for garbage collection.
Alert message	*message*
Cause	The following causes are possible:File Server not in healthy stateProblems with Object Store Connectivity
Impact	Lifecycle policies are not accurately configured on the bucket and can retain objects longer than intended.
Resolution	Please cleanup the Object Store. Refer to KB11074 for more details. Contact Nutanix support for assistance.
KB Article	11074
Severity	Warning

Table 637: File Server/Share(s) are deleted; Object Store cleanup is required [160137] [A160137]

Name	File Server/Share(s) are deleted; Object Store cleanup is required
Description	File Server/Share(s) are deleted, please cleanup the Object Store.
Alert message	*message*
Cause	File Server/Share(s) are deleted on the cluster.
Impact	Objects can remain longer than intended. After the file server is deleted, we cannot accurately track the objects for garbage collection.
Resolution	Please cleanup the Object Store. Refer to KB11075 for more details. Contact Nutanix support for assistance.
KB Article	11075
Severity	Info

Table 638: Disk space usage for audit volume high on a file server. [160138] [A160138]

Name	File Server System Audit Volume Space Usage High
Description	Checks if FSVM system audit volume usage is within threshold to ensure uninterrupted operations.
Alert message	`System Audit Volume Space Usage High on a file server.` *message*
Cause	Increased FSVM system audit volume usage due to excessive logging, or miscellaneous files being placed inside.
Impact	If FSVM system audit volume is 100% utilized, then auditd will stop and further system auditing will not be recorded.

Resolution	Remove files other than audit.log* from /home/log/audit. If unsure, please engage Nutanix Support for assistance.
KB Article	1523
Severity	Critical
Severity	Warning

Table 639: File Server Share Conversion Alert [160139] [A160139]

Name	File Server Share Conversion Alert
Description	File server share conversion - see message for details
Alert message	`Share conversion for file_server_name: message`
Cause	See details in the alert message
Impact	See details in the alert message
Resolution	See details in the alert message
Severity	Info

Table 640: Versioning is suspended for the bucket in the file server tiering profile. [160142] [A160142]

Name	File Server Tiering - Bucket versioning is suspended
Description	Versioning is suspended for the bucket in the file server tiering profile.
Alert message	`message`
Cause	Versioning is suspended for the bucket in the file server tiering profile.
Impact	Tiering is not allowed if versioning is suspended for the bucket in the file server tiering profile.
Resolution	Enable versioning for the bucket in the file server tiering profile. Please refer to the vendor-specific documentation for bucket versioning configuration.
Severity	Warning

Other

Table 641: Detected recovery points on prism central managing clusters with high density nodes [110452] [A110452]

Name	High density nodes recovery points check
Description	Checks if recovery points are detected on prism central managing clusters with high density nodes
Alert message	`Recovery points are detected on prism central managing clusters with high density nodes`

Cause	Detected recovery points on prism central managing clusters with high density nodes
Impact	Recovery points are not currently supported on clusters with high density nodes. They may impact cluster operations
Resolution	Delete recovery points on prism central hosted by clusters with high density nodes
KB Article	7003
Severity	Info

Table 642: Detected protection policies on Prism Central managing clusters with high density nodes [110453] [A110453]

Name	High density nodes protection rules check
Description	Checks if protection policies are detected on Prism Central managing clusters with high density nodes
Alert message	`Protection policies are detected on Prism Central managing clusters with high density nodes`
Cause	Protection policies are detected on a Prism Central managing clusters with high density nodes
Impact	Data protection is currently not supported for clusters with high density nodes, this may impact cluster operations
Resolution	Delete protection policies configured on Prism Central managing clusters with high density nodes
KB Article	7003
Severity	Info

Table 643: Metro Availability Operation Failed [130124] [A130124]

Name	Metro Availability Operation Failed
Description	Metro availability operation failed
Alert message	`For the protection domain 'protection_domain_name', operation failed to the remote site 'remote_name'. Reason: 'reason'.`
Cause	Check the alert message for the reason of failure
Impact	Metro availability operation could not be started.
Resolution	Resolve the issue as stated in the alert message and retry the Metro operation. If the issue persists contact Nutanix Support
KB Article	10111
Severity	Critical

Table 644: Recovery Point Objective Cannot Be Met [130138] [A130138]

Name	Recovery Point Objective Cannot Be Met
Description	Recovery Point Objective Cannot Be Met
Alert message	`Recovery point objective cannot be met because 'reason'`
Cause	Various
Impact	Recovery plan could be affected.
Resolution	Contact Nutanix support if this condition persists
Severity	Warning

Table 645: Network Segmentation Config Update Failed [130363] [A130363]

Name	NS Config Update Failed
Description	Network Segmentation Config update during AOS upgrade failed
Alert message	*message*
Cause	Network Segmentation Configuration update during AOS upgrade failed
Impact	All Network Segmentation operations will be blocked
Resolution	Try CLI command network_segmentation fix_config_after_upgrade on any CVM. Contact Nutanix Support in case of command failure.
Severity	Warning

Table 646: Memory configuration inconsistent. [200306] [A200306]

Name	Prism Central VM same memory level check
Description	Check all Prism Central VMs have the same memory level.
Alert message	`The Prism Central VMs are not configured to have the same amount of memory.`
Cause	Memory configuration among Prism Central VMs is inconsistent.
Impact	The Prism Central VM will not perform at the level necessary to manage the cluster.
Resolution	Fix Prism Central VM memory configuration by providing same amount of memory for all the Prism Central VMs
KB Article	5377
Severity	Warning
Severity	Info

Table 647: Prism Central VM type or annotation not set. [200307] [A200307]

Name	PCVM type and annotation check

Description	Check the type and annotation of all PCVMs are set.
Alert message	`The type or annotation of Prism Central VMs is not set`
Cause	Prism Central VMs do not have annotation set.
Impact	Guardrails to avoid potentially disruptive user actions on a Prism Central will not be applied.
Resolution	Set the annotation for the Prism Central VM. Refer to KB 6027.
KB Article	6027
Severity	Warning

Table 648: Unequal metadata partition size(s) across Prism Central VMs. [200308] [A200308]

Name	Prism Central VM same disk size check
Description	Checks if metadata partition size(s) are the same among Prism Central VMs.
Alert message	`The metadata partitions across Prism Central VMs are not consistent due to mismatch in sizes: disks_info`
Cause	Metadata partition size(s) of one or more Prism Central VMs is inconsistent with other Prism Central VMs.
Impact	Manageability of the cluster may be lost. Cluster may be significantly degraded and ultimately fail.
Resolution	Contact Nutanix support to resize the Prism Central VM disks.
KB Article	6028
Severity	Warning

Table 649: High time difference between Prism Central and registered Prism Elements. [200309] [A200309]

Name	PC-PE time drift check
Description	Checks ntp sync between the PC and the registered PEs.
Alert message	`Time differs between Prism Central and registered Prism Elements.`
Cause	Time is not synchronized between Prism Central and registered Prism Elements.
Impact	Stale statistics or configuration may appear in Prism Central.
Impact	Statistics may not be available in Prism Central.
Impact	Components that make use of the statistics or configuration in Prism Central may not function as expected.
Resolution	Ensure that NTP is configured properly and that the cluster has connectivity with the NTP server. Refer to KB 6121.
KB Article	6121
Severity	Warning

Table 650: PE-PC incompatible AOS versions. [200312] [A200312]

Name	PC-PE incompatible versions
Description	Checks if the Prism Element and Prism Central AOS versions are compatible.
Alert message	`The AOS version of the registered cluster is incompatible with the Prism Central version.`
Cause	The AOS version of the registered cluster is incompatible with the Prism Central version.
Impact	Prism Central features will not work as expected for this cluster.
Resolution	Upgrade the AOS version of the registered cluster to a version compatible with this Prism Central.
KB Article	6323
Severity	Warning

Table 651: Prism Central VM active upgrade check [200313]

Name	Prism Central VM active upgrade check
Description	Checks if there are any active upgrades
Cause	An upgrade is in process for this entity
Impact	Maintenance activities will have to wait until upgrade is finished.
Resolution	Wait until the upgrade is completed before doing any other actions.Contact Nutanix support for assistance.
KB Article	6644

Table 652: Prism Central Upgrades are disabled on {cvm_ip} [200314] [A200314]

Name	Prism Central VM upgrades are disabled check
Description	Checks if upgrades are disabled on Prism Central VM
Alert message	`Prism Central Upgrades are disabled on cvm_ip`
Cause	Prism Central Upgrade has been paused manually.
Impact	Upgrades from Prism won't proceed.
Resolution	Upgrades have been disabled and need to be re-enabled.Contact Nutanix support for assistance.
KB Article	6644
Severity	Info

Table 653: Free Xi Account Expired [200701] [A200701]

Name	Free Xi Account Expired

Description	Free period using Xi expired. If this issue persists, you may experience an interruption in service.
Alert message	`Free period using Xi expired. If this issue persists, you may experience an interruption in service.`
Cause	The free period is limited by time.
Resolution	Select a payment plan and method in the billing section under the top right gear menu.
Severity	Critical

Table 654: Xi Subscription Expired [200702] [A200702]

Name	Xi Subscription Expired
Description	Xi subscription period expired. If this issue persists, you may experience an interruption in service.
Alert message	`Xi subscription period expired. If this issue persists, you may experience an interruption in service.`
Cause	The subscription period is limited by time.
Resolution	Select a new payment plan in the billing section under the top right gear menu.
KB Article	8255
Severity	Critical

Table 655: Xi Payment Missed [200703] [A200703]

Name	Xi Payment Missed
Description	Unable to process payment. If this issue persists, you may experience an interruption in service.
Alert message	`Unable to process payment. If this issue persists, you may experience an interruption in service.`
Cause	This is likely due to an expired or canceled credit card.
Resolution	Update payment method in the billing section under the top right gear menu.
KB Article	8254
Severity	Critical

Table 656: Xi Payment Missed [200704] [A200704]

Name	Xi Payment Missed
Description	Unable to process payment. If this issue persists, you may experience an interruption in service.

Alert message	`Xi tenant tenant_uuid invoice payment for plan_name plan failed at timestamp_usecs`
Cause	This is likely due to a change in address or a forgotten or misplaced invoice.
Resolution	Ensure invoices are sent to the correct contact information provided in the billing section under the top right gear menu. Confirm with your team that the invoice wasn't misplaced or forgotten. Contact support if needed.
KB Article	8253
Severity	Critical

Table 657: Recovery Plan has conflicting IP mappings with other Recovery Plan(s) [300430] [A300430]

Name	Recovery Plan IP mappings check
Description	Checks for conflicting IP mappings across multiple Recovery Plans for the paired AZs.
Alert message	`Recovery Plan rp has conflicting IP mappings with other Recovery Plans. conflicting_ips conflicting_vms`
Cause	IPs mapped in Recovery Plans are conflicting. Same IP from a subnet is specified for multiple VMs.
Cause	IPs mapped in Recovery Plans are conflicting. Different IPs are specified for the same VM in different Recovery Plans.
Impact	IP mapping will fail if a VM has been recovered with the same IP from a different Recovery Plan.
Impact	VM will come up without the specified IP after recovery.
Resolution	Ensure that the same IP from a subnet is not specified for multiple VMs.
Resolution	Ensure that different IPs are not specified for the same VM in different Recovery Plans.
KB Article	7866
Severity	Warning

Table 658: Recovery Plan has errors in Consistency Group(s) configuration. [300434] [A300434]

Name	Recovery Plan consistency group configuration check
Description	Checks if Recovery Plans have errors in the configuration of consistency groups.
Alert message	`Recovery Plan recovery_plan_name has errors in configuration of following Consistency Group(s): partial_cg_names.`
Cause	Recovery Plan {recovery_plan_name} has the following partial Consistency Groups: {partial_cg_names}.
Impact	Partial Consistency Group might be recovered on failover.

Impact	Consistency Group will not be reformed post failover of entities in the Recovery Plan.
Impact	Recovered VM(s), Volume Group(s) might be in inconsistent state.
Resolution	Ensure that all VM(s) and Volume Group(s) belonging to a consistency group are present in Recovery Plan
KB Article	10532
Severity	Warning

Table 659: Recovery Plan has networks which are modified [300435] [A300435]

Name	Recovery Plan modified network check
Description	Checks if network names specified in the Recovery Plan are modified.
Alert message	`Recovery Plan recovery_plan_name has following networks which are modified: missing_networks.`
Cause	Network name(s) specified in the Recovery Plan are modified.
Cause	Network(s) specified in the Recovery Plan are deleted.
Impact	VMs will be recovered without vNIC upon failover
Resolution	Ensure that Network name(s) specified in the Recovery Plan are correct.
Resolution	Ensure that Network(s) specified in the Recovery Plan exists.
KB Article	11044
Severity	Warning

Table 660: Recovery Plan has VM(s) with iSCSI attached Volume Group(s) that do not have a compatible NGT installed/enabled. [300436] [A300436]

Name	Recovery Plan check for NGT version installed in VM(s) with iSCSI attached Volume Group(s)
Description	Checks if VM(s) with iSCSI attached Volume Group(s) specified in Recovery Plan have a compatible NGT installed.
Alert message	`Recovery Plan recovery_plan_name has VM(s) vm_list with iSCSI attached Volume Group(s) that do not have a compatible NGT installed/enabled.`
Cause	VM(s) with iSCSI attached Volume Group(s) do not have a compatible NGT installed.
Impact	Volume Group(s) attachment post recovery of a VM(s) would fail.
Resolution	Ensure that the VM(s) with iSCSI attached Volume Group(s) have NGT version of 2.3 or higher.
KB Article	11649
Severity	Warning

Table 661: Recovery Plan has VM to Volume Group(s) attachment(s) configuration that are different from the live VM to Volume Groups(s) attachment(s). [300437] [A300437]

Name	Recovery Plan VM to Volume Group(s) attachment(s) configuration check.
Description	Checks if the VM to Volume Group(s) attachment(s) configuration which are specified in the Recovery Plan are the same as the live VM to Volume Group(s) attachment(s).
Alert message	`Recovery Plan recovery_plan_name has VM to Volume Group(s) attachment(s) attachment_list configuration specified that are not the same as the live VM to Volume Group(s) attachment(s).`
Cause	VM to Volume Group(s) attachment(s) which are specified in the Recovery Plan are not present in the live state.
Cause	VM to Volume Group(s) attachment(s) present in the live state are not specified in the Recovery Plan.
Cause	VM to Volume Group(s) attachment(s) authentication configuration present in the live state are not the same as the configuration specified in the Recovery Plan.
Impact	VM to Volume Group(s) attachment(s) might not be done post recovery.
Resolution	Ensure that VM to Volume Group(s) attachment(s) which are specified in the Recovery Plan, are also present in the live state.
Resolution	Ensure that all VM to Volume Group(s) attachment(s) present in the live state are also specified in the Recovery Plan.
Resolution	Ensure that VM to Volume Group(s) attachment(s) authentication configuration present in the live state matches with the configuration specified in the Recovery Plan.
KB Article	11629
Severity	Warning

Table 662: Nucalm Internal Service has Stopped Working [400101] [A400101]

Name	Nucalm Service Down
Description	Nucalm internal service is down
Alert message	`Discovered that the nucalm internal service is not working: 'service_names'`
Cause	Nucalm internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-7080
KB Article	7080
KB Article	8621
Severity	Critical

Table 663: Epsilon Internal Service has Stopped Working [400102] [A400102]

Name	Epsilon Service Down
Description	Epsilon internal service is down
Alert message	`Discovered that the epsilon internal service is not working: 'service_names'`
Cause	Epsilon internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-7080
KB Article	7080
KB Article	8622
Severity	Critical

Table 664: Applications Archive ready for download. [400104] [A400104]

Name	Applications Archive Ready
Description	Archive ready for download.
Alert message	`Application logs from 'start_time' to 'end_time' are now available for download and archival. Please go to Calm > Settings to download them.`
Cause	Archival process archived the run logs.
Impact	You will not be able to download it after a week.
Resolution	Make sure you download the archive.
Severity	Critical

Table 665: Dataservice IP is unreachable [400105] [A400105]

Name	Data service IP connectivity check
Description	Data service IP is not reachable
Alert message	`Data service ip 'data_service_ip' is unreachable from prism central node_name`
Cause	Invalid or empty data service ip
Impact	Nutanix Calm containers will be inaccessible
Resolution	Please provide correct data service ip. Review KB 5199
KB Article	5199
Severity	Critical

Table 666: Calm showback is unable to reach beam service [400107] [A400107]

Name	Beam is unreachable

Description	Beam service is unreachable. Calm showback calculations will be affected
Alert message	`Beam service is unreachable. Calm showback calculations will be affected`
Cause	Prism central may not have internet connectivity or beam service is unavailable.
Impact	Impact on Calm Showback data.
Resolution	Review KB-6664 for further troubleshooting.
KB Article	6664
Severity	Critical

Table 667: AHV Prism Element attached [400108] [A400108]

Name	AHV Prism Element attached
Description	A new AHV Prism Element is attached.
Alert message	`AHV Prism Element attached name: pe_name`
Cause	A new AHV Prism Element is attached.
Impact	New Calm provider setting is added without showback price list.
Impact	Prism Element unusable in Calm if not associated with Calm project.
Resolution	(A) Associate the Prism Element in Calm project. (B) Define the price list for added Calm provider setting.
Severity	Info

Table 668: AHV Prism Element detached [400109] [A400109]

Name	AHV Prism Element detached
Description	One of the existing AHV Prism Element is detached.
Alert message	`AHV Prism Element detached name: pe_name`
Cause	One of the existing AHV Prism Element is detached.
Impact	Actions other than soft delete cannot be run on Calm applications that are provisioned on this Prism Element.
Impact	Calm provider settings corresponding to this Prism Element is deleted.
Resolution	In case the Prism Element is permanently detached run soft delete on the Calm applications that are provisioned on this Prism Element.
Severity	Info

Table 669: Calm Version Mismatch [400110] [A400110]

Name	Calm Version Mismatch
Description	Calm Version Mismatch

Alert message	`All nodes are not running on same Calm version`
Cause	Failed LCM upgrade
Impact	Calm will not work as expected
Resolution	Refer to KB 7528 for resolution
KB Article	7528
Severity	Critical

Table 670: Epsilon Version Mismatch [400111] [A400111]

Name	Epsilon Version Mismatch
Description	Epsilon Version Mismatch
Alert message	`All nodes are not running on same Epsilon version`
Cause	Failed LCM upgrade
Impact	Epsilon will not work as expected
Resolution	Refer to KB 7528 for resolution
KB Article	8694
Severity	Critical

Table 671: Calm Containers are unhealthy [400112] [A400112]

Name	Calm containers health check
Description	Check for Calm Container's state
Alert message	`Nutanix Calm or Epsilon container is in unhealthy state`
Cause	Internal services of each calm containers are down, Docker plugin is not working properly
Impact	Nutanix Calm services may be inaccessible or performs incorrectly.
Resolution	Please check internal services of each calm containers.
KB Article	7343

Table 672: Calm Policy Engine Internal Service has Stopped Working [400114] [A400114]

Name	Calm Policy Engine Service Down
Description	Calm Policy Engine internal service is down
Alert message	`Discovered that the Calm policy engine internal service running on` *`policy_engine_ip`* `is not working: '`*`service_names`*`'`
Cause	Calm Policy Engine internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-9456

KB Article	9456
Severity	Critical

Table 673: Threshold set for the quota has been crossed [400116] [A400116]

Name	Quota threshold crossed
Description	Quota threshold crossed
Alert message	`Quota utilisation has crossed threshold threshold for quota_attributes at quota_level`
Cause	Number of virtual machine(s) deployed under quota exceeds the threshold
Impact	Quota consumption can go beyond the threshold value
Resolution	Either delete virtual machines or increase threshold value set for the quota
Severity	Warning

Prism Central VM

Table 674: XPilot Retry Approval Pending [111081] [A111081]

Name	XPilot Retry Approval Pending
Description	XPilot Playbook is waiting for Approval to retry
Alert message	`XPilot Playbook action_rule_name is awaiting approval to begin retry no. retry_count`
Cause	XPilot Playbook is waiting for approval to begin retry
Impact	Playbook will fail after wait period of 7 days
Resolution	Resume playbook to give approval. Refer to KB 9961 for further details
KB Article	9961
Severity	Info

Storage

Table 675: VSS Snapshot of Container Share Failed [20020] [A20020]

Name	VSS snapshot of container share failed
Description	A VSS initiated snapshot of the container share failed due to too many files in the container.
Alert message	`A backup application requires a container id container_id to be backed up, however the VSS snapshot operation on it failed as there are files_present files on the container. Only a container with a maximum of files_max files can be snapshotted through VSS.`
Cause	Too many files in the container.

Impact	Third party backup applications using VSS to backup VMs from a container share will fail.
Resolution	Reduce the number of files in each container to less than or equal to 1000.
Severity	Warning

Table 676: Fingerprinting Disabled [130021] [A1068]

Name	Fingerprinting On Write
Description	Fingerprinting Disabled
Alert message	`Disabling fingerprinting (deduplication) for future writes. Calculated metadata usage of metadata_usage bytes exceeds safety limit of metadata_limit bytes on Controller VM service_vm_id. The metadata disk size is metadata_disk_size.`
Cause	The metadata disk is used for metadata and data. Once the metadata portion on the disk exceeds the safety threshold AOS disables fingerprinting (deduplication) of future writes. Metadata sizes can increase due to a variety of reasons such as having too many snapshots, or simply because too much data is fingerprinted.
Impact	The advantages of deduplication are not available.
Resolution	Once the metadata size is brought under the safety limit (for example, by removing unneeded snapshots, or by expanding the cluster size), the fingerprinting will be auto-enabled.
Severity	Warning

Table 677: Metadata Usage High [130025] [A1101]

Name	Metadata Usage
Description	Metadata Usage High.
Alert message	`Metadata usage on Controller VM service_vm_external_ip has exceeded critical_pct%.`
Cause	The cluster either has too many snapshots or too much data is being fingerprinted.
Impact	Cluster performance may be significantly degraded.
Resolution	Reduce metadata size by removing unneeded snapshots, creating snapshots less frequently, creating snapshots of fewer VMs, or expanding the cluster size.
Severity	Warning

Table 678: Non Self Encryption Drive Disk Inserted [130031] [A1122]

Name	Non SED Disk Inserted Check
Description	Non-self-encrypting drive is inserted.

Alert message	`Non encrypted disk with serial disk_serial was added in drive bay disk_location on Controller VM service_vm_external_ip. It is not usable because the rest of the cluster is protected using encrypted drives.`
Cause	A cluster with self-encrypting drives has a non-self encrypting drive installed.
Impact	The new drive was not mounted because the cluster cannot be fully secured without all drives supporting encryption.
Resolution	Ensure all installed drives are of self-encrypting type. If more help is needed call Nutanix support.
KB Article	8417
Severity	Critical

Table 679: On-Disk Dedup Disabled [130032] [A1089]

Name	On-Disk Dedup Status
Description	On-disk deduplication is disabled.
Alert message	`Disabling further on-disk deduplication (fingerprinting). Controller VM with ram_sizeGB RAM and ssd_sizeGB SSD does not meet the minimum requirements of required_ram_sizeGB RAM and required_ssd_sizeGB SSD to support the on-disk deduplication feature.`
Cause	To support on-disk deduplication (fingerprinting) feature, the Controller VM must meet minimum requirements for RAM and SSD. If this threshold is not satisfied, on-disk deduplication is disabled.
Impact	The advantages of on-disk deduplication are not available.
Resolution	Once the Controller VMs are upgraded to meet the minimum requirements on RAM and SSD, on-disk deduplication can be enabled. Contact Nutanix support for assistance.
Severity	Warning

Table 680: Disk Failed/Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Drive is marked offline.
Alert message	`Disk disk_position with id disk_id on node node_position of block block_position has been marked offline due to IO errors. Serial number of the disk is disk_serial in host host_ip of block block_serial.`
Cause	A drive on the node has failed and was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Replace the failed drive. Refer to the Nutanix documentation for instructions.
KB Article	8453
Severity	Critical

Table 681: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	`Drive disk_id with serial disk_serial and model disk_model in drive bay disk_location on Controller VM service_vm_external_ip has failed. Node serial is node_serial_number and Node position node_position. #rf1_custom_message@`
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 682: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical drive is removed from the slot.
Alert message	`Drive with serial disk_serial and model disk_model was removed from drive bay disk_location on Controller VM service_vm_external_ip. Node serial is node_serial_number and Node position node_position. #rf1_custom_message@`
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the drive will start.
Resolution	Remove the drive from the slot and reseat it. Contact the Hardware Vendor if issue persists.
KB Article	8031
Severity	Critical

Table 683: Space Reservation Violated [130053] [A1021]

Name	Space Reservation Status
Description	Space reservation violated.
Alert message	`Space reservation configured on vdisk vdisk_name belonging to Storage Container id container_id could not be honored due to insufficient disk space resulting from a possible disk or node failure.`
Cause	A drive or a node has failed.
Impact	The space reservations on the cluster can no longer be met. Writes by guest VMs may fail if expected storage space is not available.
Resolution	Change space reservations to total less than 90% of the available storage, and replace the drive or node as soon as possible. Refer to the Nutanix documentation for instructions.

Severity	Warning

Table 684: Volume Group Action Error [130071] [A1163]

Name	Volume Group Action Status
Description	Volume group action error.
Alert message	`Failed to action volume group with name 'vg_name' and internal ID 'vg_uuid' because reason in protection domain 'protection_domain_name'`
Cause	A volume group could not be restored or could not be deleted because it is still in use.
Impact	The requested volume group action (restore or delete) could not be completed.
Resolution	Detach the volume group from VMs and external initiators before recovery. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10120
Severity	Critical

Table 685: Unable To Remount Datastore [130076] [A130076]

Name	Datastore Remount Status
Description	Unable to remount a datastore.
Alert message	`Unable to operation Datastore 'datastore_name'host_info.`
Cause	Datastore information unavailable
Impact	VM HA and DRS functionality might be impaired.
Resolution	Check container exists on ESXi, rescan datastore and mount the container on the hosts from Prism Element. If issue persists contact Nutanix Support
Severity	Warning

Table 686: Remounted Datastore [130080] [A130080]

Name	Datastore Remount Success
Description	Remounted datastore.
Alert message	`Remounted Datastore 'datastore_name' as metro availability is being enabled.`
Cause	Datastore has been remounted
Impact	VM HA and DRS functionality might be impaired.
Resolution	As the datastore has been mounted successfully rescan the datastore in vCenter in order to complete the changes necessary to support the enhanced planned failover functionality

Severity	Info

Table 687: Protected VM(s) Not Found [130083] [A130083]

Name	Protected VMs Not Found
Description	Protected VM(s) Not Found.
Alert message	`Unable to locate VM(s) vm_names protected by protection domain 'protection_domain_name'.`
Cause	Protected VM(s) cannot be found and may have been deleted.
Impact	Any data associated with the VMs may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Verify that the VM(s) are registered on the host. If the VM(s) were unregistered or deleted, remove VM(s) from the protection domain. If the CVM on the host the VM(s) are registered on was shutdown or rebooting, retry the snapshot operation.
Severity	Warning

Table 688: Protected Volume Groups Not Found [130084] [A130084]

Name	Protected Volume Groups Not Found
Description	Protected volume groups not found.
Alert message	`Unable to locate volume group(s) vg_names protected by protection domain 'protection_domain_name'.`
Cause	Protected volume group(s) cannot be found and may have been deleted.
Impact	Any data associated with the volume group(s) may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Remove volume group(s) from the protection domain.
Severity	Warning

Table 689: VStore Snapshot Status [130086] [A130086]

Name	VStore Snapshot Status
Description	vStore snapshot status.
Alert message	`Snapshot status for vstore vstore_name: reason.`
Cause	Various
Impact	The requested VStore snapshot action could not be completed.

Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 690: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	`Drive disk_id with serial disk_serial in drive bay disk_location on Controller VM service_vm_external_ip has failed diagnostic test. This drive is failing or has failed. #rf1_custom_message@`
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 691: Associated entities are not protected together. [130090] [A130090]

Name	Related Entity Protection Status
Description	Protection status of a related entity.
Alert message	`error_message.`
Cause	Related entity is not protected in the same protection domain.
Impact	Related VM/Volume Group will not be snapshotted and recovered.
Resolution	Protect the related entity in the same consistency group.
Severity	Warning

Table 692: iSCSI Configuration Failed [130100] [A130100]

Name	iSCSI Configuration Failed
Description	iSCSI Configuration Failed
Alert message	`Failed to re-configure iSCSI settings on the recovered VM 'vm_name'. reason.`
Cause	Nutanix Guest Tools failed to execute some iSCSI commands on the guest VM.
Impact	iSCSI disks may become unavailable on the guest VM.
Resolution	(A) If IQN and iSCSI target IP addresses of the VM have been updated by Nutanix Guest Agent, Discover and connect to new targets after rebooting the VM. (B) Manually configure iSCSI settings on the guest VM. (C) Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10118

Severity	Critical

Table 693: Recovered VM Disk Configuration Update Failed [130107] [A130107]

Name	Disk Configuration Update Failed
Description	Disk Configuration Update Failed
Alert message	`Failed to make some disks online on the recovered VM 'vm_name'. reason.`
Cause	Nutanix Guest Tools failed to automatically bring the disks online.
Impact	Disks may become offline on the recovered VM.
Resolution	Manually bring the disks online on the recovered VM. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	8208
Severity	Critical

Table 694: System-Defined Flash Mode Usage Limit Exceeded [130120] [A130120]

Name	System Defined Flash Mode Usage Limit
Description	Check that usage for flash-mode-enabled vDisks is within system limits.
Alert message	`System has down-migrated the data of flash-mode-enabled vDisks.`
Cause	Too many vDisks are assigned to the flash tier, or the vDisks assigned to the flash tier are too large.
Impact	The advantages of flash mode may not be available.
Resolution	Reduce the number of flash-mode-enabled vDisks or increase the capacity of the flash tier.
Severity	Warning

Table 695: NFS metadata usage high [130122] [A130122]

Name	NFS Metadata Size Overshoot
Description	NFS Metadata Usage High
Alert message	`NFS metadata usage is too high. Calculated metadata usage of metadata_usage bytes exceeds safety limit of metadata_limit bytes on Controller VM service_vm_id.`
Cause	Either a NFS datastore is used as a file store which is not supported, or too many files are present in the datastore.
Impact	Cluster performance and stability may be significantly degraded.
Resolution	If a NFS datastore is used as a file store, remove the corresponding files and wait for up to 24 hours to allow for compaction.
KB Article	6012

Severity	Warning

Table 696: Guest Power Operation through NGT Failed [130149] [A130149]

Name	VM Guest Power Op Failed
Description	Guest Power Operation Failed
Alert message	`Failed to perform operation 'operation'on VM 'vm_name'. reason.`
Cause	Power operation failed in guest operating system.
Impact	Desired power operation may not be completed.
Resolution	(A) Manually shutdown the VM by logging into the VM and running the appropriate shutdown command. (B) Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.
KB Article	10242
Severity	Warning

Table 697: Protected VM(s) May Not Be Recoverable [130183] [A130183]

Name	Protected VMs May Not Be Recoverable
Description	Protected VM(s) May Not Be Recoverable
Alert message	`VM(s) vm_names protected by protection domain 'protection_domain_name' are running from VMWare snapshots. Recovery of the VMs from the PD snapshot 'snapshot_id' may not succeed. Manual intervention is required to boot up the VMs.`
Cause	Protected VM(s) are running from VMWare snapshots and may not be recoverable.
Impact	Recovery of the VMs from the PD snapshot may not succeed
Resolution	Ensure that the VMs are not running from VMWare snapshot at the time of scheduled Protection Domain snapshot. Refer Nutanix KB: 7023 for further details.
Severity	Warning

Table 698: External Repository Access Failure [130208] [A130208]

Name	External Repository Access Failure
Description	Failed to access an NFS v3 external repository.
Alert message	`Failed to access external repository external_repository_name. Failure reason: message.`
Cause	Network connectivity issues between the AOS CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	The NFS server may not have white-listed access to the CVM subnet.

Cause	The NFS server may not support the v3 protocol version.
Cause	The repository is no longer exported by the NFS server.
Cause	The NFS server may not allow clients accessing it using non-privileged TCP ports (port number higher than 1023).
Cause	The NFSv3 external repository does not have the correct access permission.
Cause	I/O error on the NFS server while mounting the exported directory.
Impact	Mounting an NFS v3 external repository fails
Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure the CVM subnet is (all AOS CVMs have) white-listed for the export on the NFS server.
Resolution	Make sure that the NFS server support NFS v3.
Resolution	Make sure the repository is exported by the NFS server.
Resolution	Make sure the NFS server allows clients accessing it using non-privileged TCP ports.
Resolution	Make sure that the exported directory of the NFS external repository is mountable.
Resolution	Make sure the NFS server is in a good state.
Severity	Warning

Table 699: IO failures to a data source in an external repository [130209] [A130209]

Name	IO failures to data sources in external repositories
Description	IO to an external data source within an external repository failed.
Alert message	`IO to the external data source with path relative_uri in the external repository external_repository_name failed. Failure reason: message.`
Cause	I/O to the external datasource failed with I/O error from the NFS server.
Cause	Network connectivity issues between the AOS CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	Write failed because of no disk space available or hitting quota limitation in the NFS external repository.
Cause	Write request is attempted while the external repository is configured in READ_ONLY mode.
Cause	The datasource within the NFSv3 external repository does not have the correct access permission.
Cause	The datasource is not available in the NFS server.
Impact	IO to an data source within an external repository fails
Resolution	Make sure the NFS server is in a good state.
Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure that NFS server has enough space/quota for the write to this datasource.

Resolution	Change external repository access mode from READ_ONLY to READ_WRITE.
Resolution	Make sure that the datasource is available for the CVM subnet and allows requests such as read and/or write.
Resolution	Make sure the datasource is available or recreate the VM disk with the correct datasource configuration.
KB Article	8362
Severity	Warning

Table 700: External datasource storage migration suspended due to unrecoverable errors [130347] [A130347]

Name	External datasource storage migration suspended due to unrecoverable errors
Description	Storage migration from an external datasource within an external repository suspended due to unrecoverable errors.
Alert message	`Storage migration from the external datasource datasource_uuid with path relative_uri in the external repository external_repository_name suspended. Suspension reason: message.`
Cause	Insufficient free storage space.
Cause	Reads to the external datasource failed with I/O error from the NFS server.
Cause	Network connectivity issues between the CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	The datasource within the NFSv3 external repository does not have the correct access permission.
Cause	The datasource is not available in the NFS server.
Impact	Storage migration from a datasource within an external repository is suspended
Resolution	Make sure the destination container has enough free capacity.
Resolution	Make sure the NFS server is in a good state.
Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure that the datasource is available for the CVM subnet and allows requests such as read and/or write.
Resolution	Make sure the datasource is available or recreate the VM disk with the correct datasource configuration.
KB Article	10864
Severity	Warning

Table 701: Prism Central VM System Root Partition Space Usage High [200316] [A200316]

Name	Prism Central VM System Root Partition Usage
Description	Checks if Prism Central system root partition usage is within threshold to ensure uninterrupted operations.

Alert message	`Disk space usage for root partition mount_path on entity-ip_address has exceeded threshold%.`
Cause	Increased Prism Central VM system root partition usage due to excessive logging or incomplete maintenance operation.
Impact	If Prism Central VM system root partition is highly utilised then certain maintenance operations, such as upgrades, may be impacted. If Prism Central VM system root partition is 100% utilized then services may stop and impact Prism Central cluster management functions.
Resolution	Reduce Prism Central VM system root partition usage by removing any known temporary or unneeded files. Refer to KB 5228 for further details.
Prism Central VM System Root Partition Usage Threshold	Critical Threshold: 95 %
KB Article	5228
Severity	Critical

Table 702: Prism Central VM home partition disk usage high [200317] [A200317]

Name	Prism Central VM home partition disk usage
Description	Checks if Prism Central home partition usage is within threshold to ensure uninterrupted operations.
Alert message	`alert_msg`
Cause	Increased Prism Central VM home partition usage due to excessive logging or incomplete maintenance operation.
Impact	If Prism Central VM home partition is highly utilised then certain maintenance operations, such as upgrades, may be be impacted. If Prism Central VM home partition is 100% utilised then services may stop and impact cluster storage availability.
Resolution	Reduce Prism Central VM home partition usage by removing any known temporary or unneeded files. Refer to KB 5228 for further details.
Prism Central VM home partition disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

System Indicator

Table 703: Tomcat is restarting frequently. [700101] [A700101]

Name	Tomcat Frequent Restart
Description	Tomcat is restarting frequently.
Alert message	`Tomcat is restarting frequently on cvm_ip. Reason: message.`
Cause	Tomcat is restarting frequently, please check prism, prism_monitor and catalina logs for exact cause of failure.
Impact	Performance and availability impact for Prism
Resolution	Check the logs for error information or contact Nutanix support.
KB Article	8524
Severity	Critical

VM

Table 704: Non Compliance with Host Affinity policies. [200403] [A200403]

Name	VM not in compliance with the defined affinity policies.
Description	VM not in compliance with the defined affinity policies.
Alert message	`VM vm_uuid in cluster cluster not in compliance with policies policy_list_str`
Cause	A new policy was added, an existing policy was updated, or a category was attached to or detached from the VM or host.
Impact	VM might not run on hosts it is expected to run on.
Resolution	Make appropriate updates to affinity policies OR VM or Host categories.
Severity	Warning

6

SYSTEM MANAGEMENT

You can configure various system settings for Prism Central.

- You can configure an HTTP proxy, which can be used to communicate with a Nutanix service center (see Configuring an HTTP Proxy on page 508
- You can specify one or more name servers (see Configuring Name Servers on page 514).
- You can specify one or more NTP servers for setting the system clock (see Configuring NTP Servers on page 514).
- You can configure SNMP (see Configuring SNMP on page 516).
- You can configure a login banner page (see Configuring a Banner Page on page 530).
- You can change the language setting for displayed text (see Changing the Language Settings (Prism Central) on page 531).

Configuring Prism Central When a Cluster Uses Proxy Servers

Note:

- Prism Central and its managed clusters are not supported in environments deploying Network Address Translation (NAT).
- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see Prism Central Proxy for Pulse Data on page 936.

To communicate with the Nutanix service center directly, you need to configure the HTTP proxy. See Configuring an HTTP Proxy on page 508 for information on how to configure the HTTP proxy through Prism Central web console. After the HTTP proxy is configured, all the HTTP or HTTPS communication initiated by Prism Element or the Prism Central is routed through the proxy server.

The communication from Prism is routed through the proxy server until a whitelist target entry indicates otherwise. See Who Needs to Use the Whitelist Method on page 507 for information on who needs to whitelist the IP addresses. To overcome a communication failure between a Prism Element and the registered Prism Central, you need to manually whitelist the target entries from the Prism web console or nCLI. These options enable you to add Prism Central and its managed/registered clusters to a whitelist, where any HTTP proxy settings are ignored. This configuration allows network traffic between them, bypassing any proxy servers configured in the cluster. The whitelist also enables you to register new clusters with Prism Central successfully where clusters are using an HTTP proxy.

See *Step 3* of Configuring an HTTP Proxy on page 508 for information on how to configure the whitelist IP address through the Prism Central web console.

Alternatively, you can use the **`ncli http-proxy add-to-whitelist`** and **`ncli http-proxy delete-from-whitelist`** nCLI command options. See Whitelisting Prism Central and Its Managed

Clusters (nCLI) on page 513 for information on how to configure the whitelist IP address through the nCLI.

- You can add or delete one whitelist entry at a time.
- Each whitelist entry cannot exceed 253 characters.
- A maximum of 1000 whitelist entries are supported.
- When deleting an entry from a whitelist, delete the target, not the target type.
- The commands do not support the IPv4 network mask network prefix or * (asterisk) prefix notation. You can specify a subnet and netmask, such as 10.0.0.0/255.0.0.0 or 192.168.1.0/255.255.255.0. This results in adding the specified subnet to the whitelist.
- When applying a whitelist, domain names like `contoso.com` are not processed the same as `www.contoso.com` and are treated as separate, distinct entities.
- Use fully qualified domain names to whitelist.

Who Needs to Use the Whitelist Method

SSL port 9440 needs to be open in both directions between the Prism Central VM and any registered clusters or clusters to be registered. For the complete list of required ports, see Port Reference.

If you are implementing a proxy server in your cluster environment with this port open as shown in this simple graphic, you do not need to allow Prism Central and its managed/registered clusters.

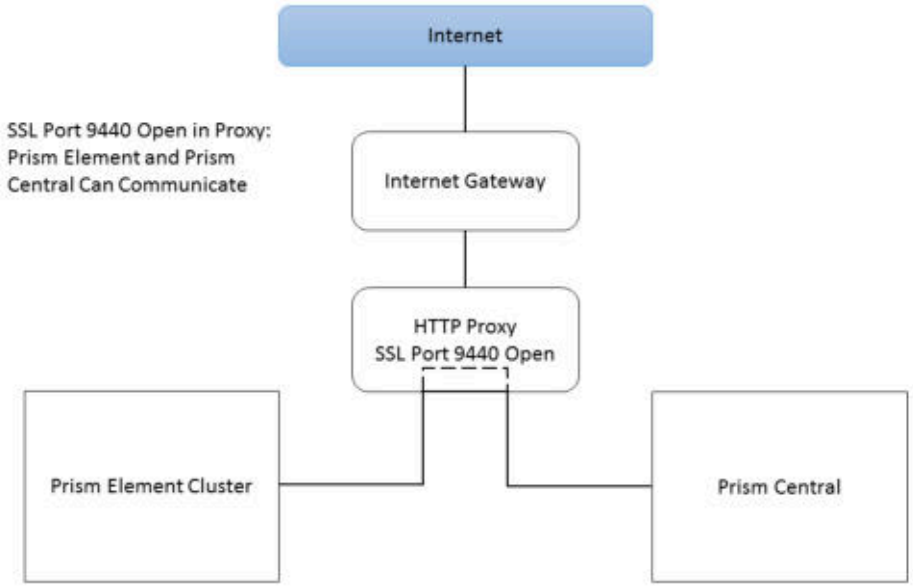

Figure 215: Prism Central, Proxy Server with SSL Port 9440 Open

If you are implementing a proxy server in your cluster environment with this port closed as shown in this simple graphic, you must allow direct communication between Prism Central and its managed/registered clusters as described in *Step 3* of Configuring an HTTP Proxy on

page 508 or alternatively in Whitelisting Prism Central and Its Managed Clusters (nCLI) on page 513.

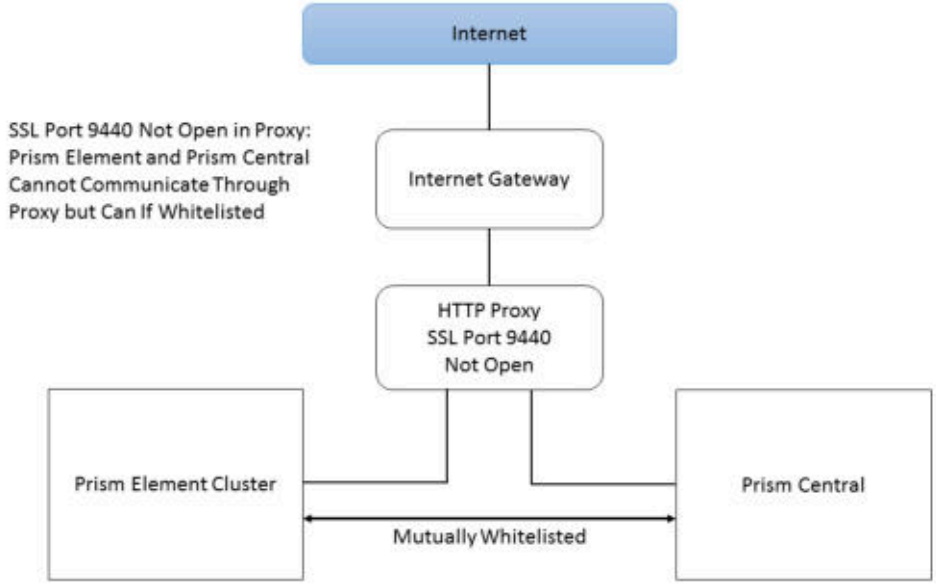

Figure 216: Prism Central, Proxy Server with SSL Port 9440 Closed and with Whitelisting

Configuring an HTTP Proxy

About this task

Note:

- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see Prism Central Proxy for Pulse Data on page 936.

About this task

If Prism Central cannot send traffic to a Nutanix Service center directly, an HTTP proxy is required. To configure an HTTP proxy, do the following:

Procedure

1. Click the gear icon in the main menu and then select **HTTP Proxy** in the **Settings** page. The HTTP Proxies dialog box appears.

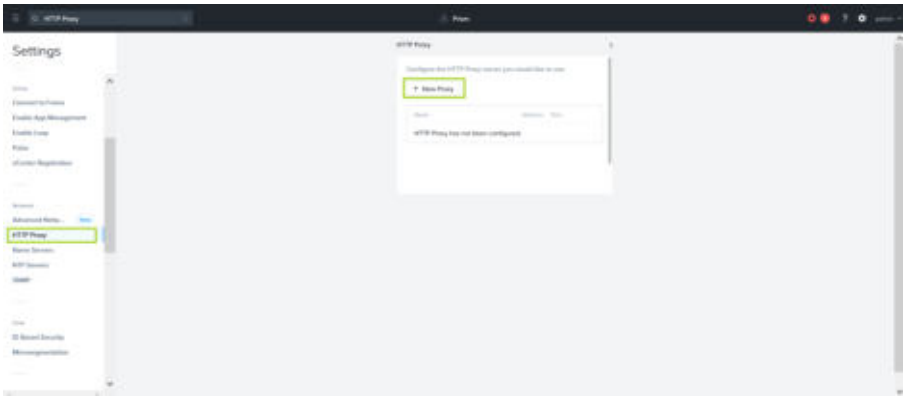

Figure 217: HTTP Proxy Window

2. To add an HTTP proxy, click the **New Proxy** button and do the following in the displayed fields:

   > Note: You can configure only one HTTP proxy at a time. If one exists currently, you must first delete it before creating a new one.

   a. **Name**: Enter a proxy server name.

   b. **Address**: Enter an IP address or host name for the proxy server.

   c. **Port**: Enter the port number to use.

   d. **Username**: Enter a user name.

   e. **Password**: Enter a password.

   f. **Protocols**: Select (check) the protocol to use, **HTTP**, **HTTPS**, or both.

      Select **HTTP** to route all HTTP requests through the proxy.

      Select **HTTPS** to route all HTTPs requests through the proxy.

      Select **HTTP** and **HTTPS** to route both HTTP and HTTPS requests through the proxy.

   g. When all the fields are correct, click the **Save** button (lower right).

      This saves the configuration and redisplays the dialog box with the new HTTP proxy entry appearing in the list.

      > Note: To return to the HTTP Proxy window without saving, click the **Cancel** button.

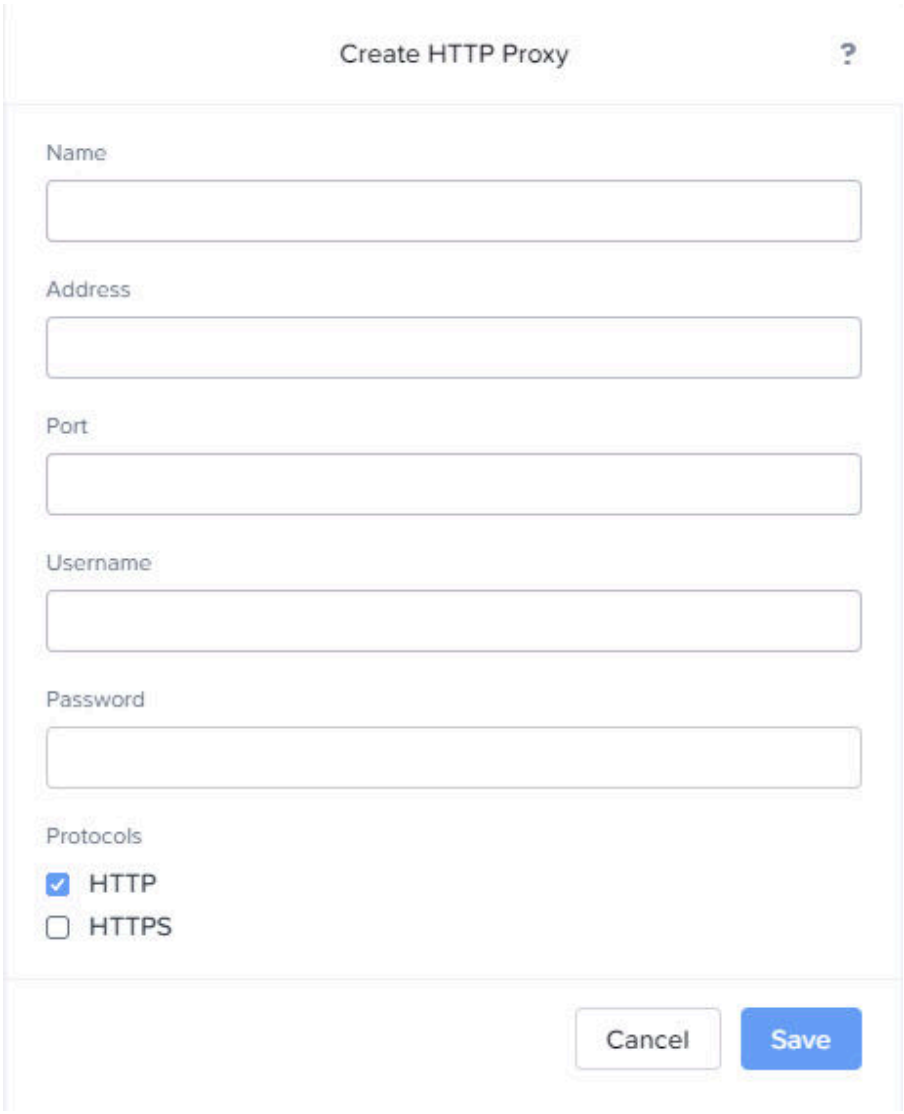

Figure 218: Create HTTP Proxy Window

3. To edit an HTTP proxy entry, click the pencil icon on the line for that entry, update one or more of displayed field entries as desired, and then click the **Save** button.

   The Update HTTP Proxy dialog box appears with the same fields as the Create HTTP Proxy dialog box plus the option (below the protocol check boxes) to add whitelist entries.

   To configure HTTP proxy whitelist entries, do the following:

   - To add a whitelist target, click the **+ Create** link. This step opens a line to enter a target address. Enter the target IP address or the network address and the subnet mask and then click the **Save** link in that field.

     Prism Central sends traffic to the whitelist IP addresses directly rather than through the HTTP or HTTPS proxy.
   - To edit a whitelist target, click the pencil icon for that target and update as needed.
   - To delete a whitelist target, click the X icon for that target.

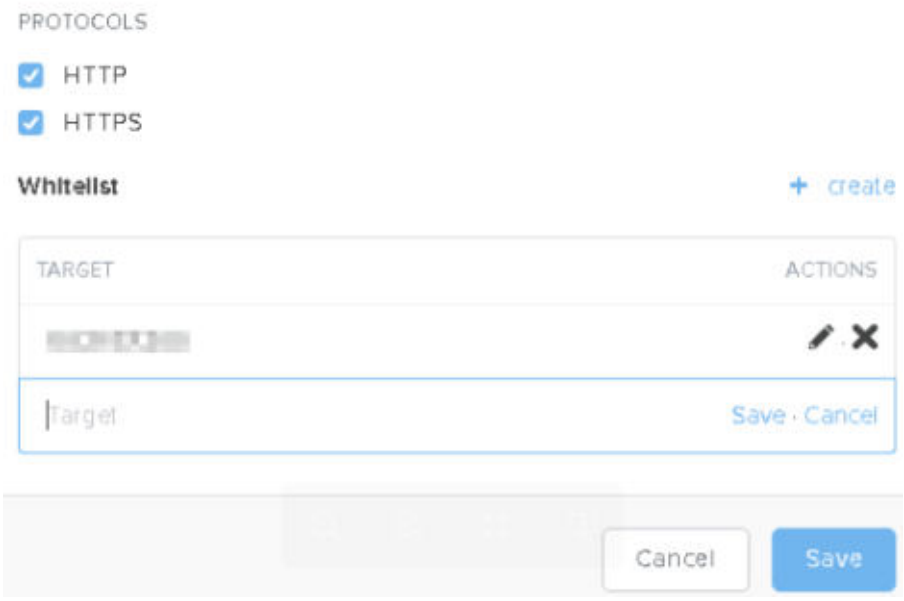

Figure 219: Whitelist Targets

Note: Where proxy whitelist configuration is needed, it is recommended that you add the cluster virtual IP and all external CVM IPs of each registered cluster to the proxy whitelist for Prism Central. The virtual IP of Prism Central and the IP of each Prism Central VM (PCVM) should be added to the proxy whitelist on the registered Prism Element clusters. For easier management, you can add the subnets instead of adding the individual IPs.

4. To delete an HTTP proxy entry, click the X icon for that entry.

   A window prompt appears to verify the action; click the **OK** button. The entry is removed from the HTTP proxy list.

Whitelisting Prism Central and Its Managed Clusters (nCLI)

About this task

In this example, bypass a proxy server used by a managed Prism Element cluster and allow network traffic between Prism Central and the cluster. Previously, if you attempted to register a cluster that implemented a proxy server, the registration failed.

Procedure

1. Open a SSH session to any Controller VM in the cluster to be managed by Prism Central.
2. In this example, add the Prism Central VM IP address to the whitelist, then ensure the Prism Central VM IP address was added to the whitelist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address target=10.4.52.40
```

```
nutanix@cvm$ ncli http-proxy get-whitelist
Target Type : IPV4_ADDRESS
Target : 10.4.52.40
```

   Note: Repeat this step for additional Prism Central VM IP addresses and the virtual IP address in the case of scaleout Prism Central.

3. Open a SSH session to the Prism Central VM managing the cluster where you just modified the HTTP whitelist.
4. Add the cluster virtual IP address to the whitelist, then ensure the IP address was added to the whitelist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address target=10.4.52.10
```

```
nutanix@cvm$ ncli http-proxy get-whitelist
Target Type : IPV4_ADDRESS
Target : 10.4.52.10
```

   Note:

   - Repeat this step for individual CVM IP addresses.
   - For large clusters or for multiple clusters in a network it can be more efficient to add a network to the http-proxy whitelist. For example,

     The following is a sample command to add a network to the http-proxy whitelist.

     ```
     nutanix@PCVM:~$ ncli http-proxy add-to-whitelist target-type=ipv4_network_mask
      target=10.4.52.0/255.255.255.0
     ```

     The following is sample output of the **get-whitelist** command, which displays the whitelisted network added above.

     ```
     nutanix@CVM:~$ ncli http-proxy get-whitelist
     ```

     ```
     Target Type : IPV4_NETWORK_MASK Target : 10.4.52.0/255.255.255.0
     ```

   In this case, Prism Central and its managed cluster can communicate, with network traffic bypassing any proxy servers configured in the cluster.

Configuring Name Servers

About this task

Name servers are computers that host a network service for providing responses to queries against a directory service, such as a DNS server. To add (or delete) a name server, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Name Servers** in the **Settings** page. The Name Servers dialog box appears.

Figure 220: Name Servers Window

2. To add a name server, enter the server IP address in the **Server** field and then click the **Add** button to the right of that field.

   The server is added to the **IP Address** list (below the **Server** field).

   > Note: Changes in name server configuration may take up to 5 minutes to take effect. Functions that rely on DNS may not work properly during this time. You can configure a maximum of three name servers.

3. To delete a name server entry, click the X icon for that server in the **Host Name or IP Address** list.

   A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

4. Click the **Close** button to close the Name Servers window.

Configuring NTP Servers

About this task

Network Time Protocol (NTP) is a protocol for clock synchronization between computers, and Prism Central must connect to an NTP server to synchronize the system clock. To add (or delete) an NTP server entry, do the following:

> Note: If Prism Central is running on Hyper-V, you must specific the IP address of the Active Directory Domain Controller server, not the hostname. Do not use DNS hostnames or external NTP servers.

Procedure

1. Click the gear icon in the main menu and then select **NTP Servers** in the **Settings** page.

   The NTP Servers dialog box appears.

Figure 221: NTP Servers Window

2. To add an NTP server entry, enter the server IP address or fully qualified host name in the **NTP Server** field and then click the **Add** button to the right of that field.

   The name or address is added to the **Server** list (below the **NTP Server** field).

3. To delete an NTP server entry, click the delete icon for that server in the **Servers** list.

   A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

4. Click the **Close** button to close the NTP Servers window.

Configuring an SMTP Server

About this task

Simple Mail Transport Protocol (SMTP) is an Internet standard protocol for electronic mail transmission across Internet Protocol (IP) networks, and Prism Central uses SMTP to send alert emails and to exchange emails with Nutanix customer support. To configure an SMTP server entry, do the following:

Procedure

1. Click the gear icon in the main menu and then select **SMTP Server** in the **Settings** page.

   The SMTP Server Settings dialog box appears.

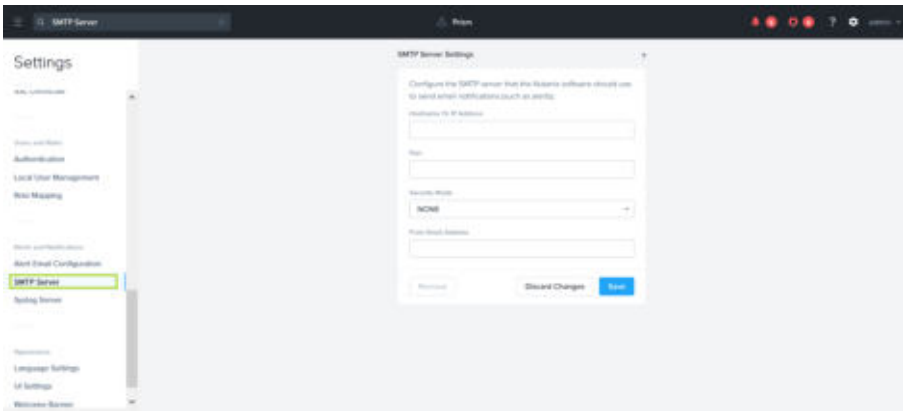

Figure 222: SMTP Server Settings Window

2. Do the following in the indicated fields:

   a. **Host Name or IP Address**: Enter the IP address or fully qualified domain name for the SMTP server.

   b. **Port**: Enter the port number to use.

      The standard SMTP ports are 25 (unencrypted), 587 (TLS), and 465 (SSL). For the complete list of required ports, see Port Reference.

   c. **Security Mode**: Enter the desired security mode from the pull-down list.

      The options are **NONE** (unencrypted), **STARTTLS** (use TLS encryption), and **SSL** (use SSL encryption).

   d. **User**: Enter a user name.

      The **User** and **Password** fields apply only when a secure option (**STARTTLS** or **SSL**) is selected. The user name might need to include the domain (*user@domain*) depending on the authentication process.

   e. **Password**: Enter the user password.

   a. **From Email Address** (optional): Enter an e-mail address that appears as the sender address.

      By default, alert and status information e-mails display "**cluster@nutanix.com**" as the sender address. You have the option to replace that address with a custom address by entering a sender address in this field.

3. When all the fields are correct, click the **Save** button.

Configuring SNMP

About this task

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. To configure SNMP on Prism Central, do the following:

Note: Prism Element (individual cluster) supports both the SNMP service (agent) and SNMP traps, but Prism Central only supports SNMP traps. Prism Central does not support the SNMP service (unlike Prism Element), so its SNMP capability is limited to just sending traps. See the *Prism Web Console Guide* for instructions on configuring SNMP for an individual cluster and for details about the Nutanix MIB.

Procedure

1. Click the gear icon in the main menu and then select **SNMP** in the **Settings** page.

   The SNMP Configuration dialog box appears.

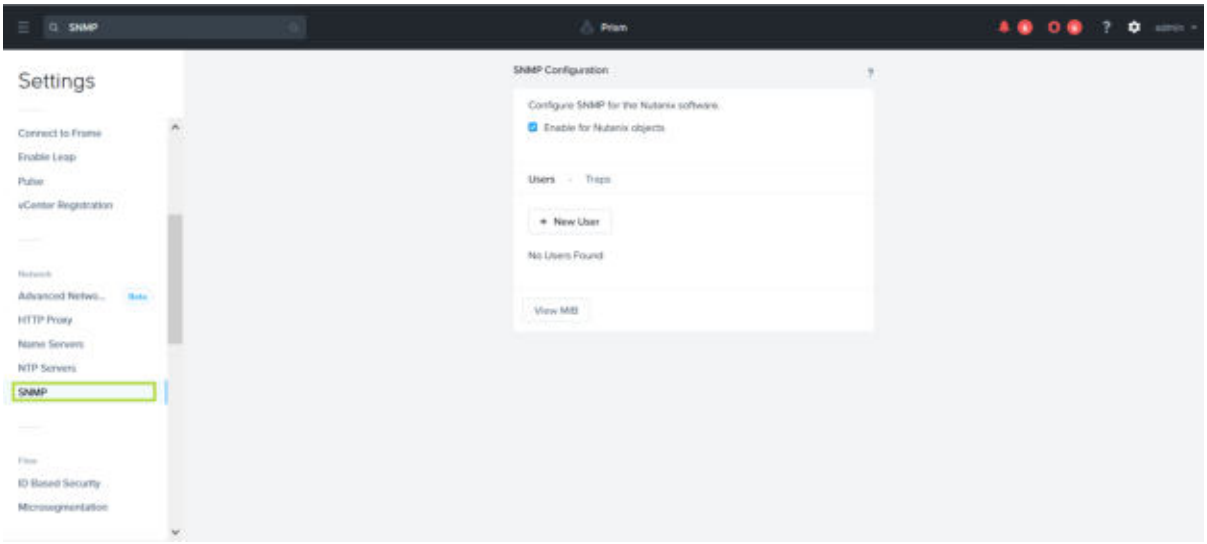

Figure 223: SNMP Configuration Window

2. To enable SNMP for Prism Central, check the **Enable SNMP** box. To disable SNMP, uncheck the box.

3. To view the Nutanix MIB (`NUTANIX-MIB.txt`), click the **View MIB** link. To download `NUTANIX-MIB.txt`, right-click and select the appropriate download action for your browser.

4. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:

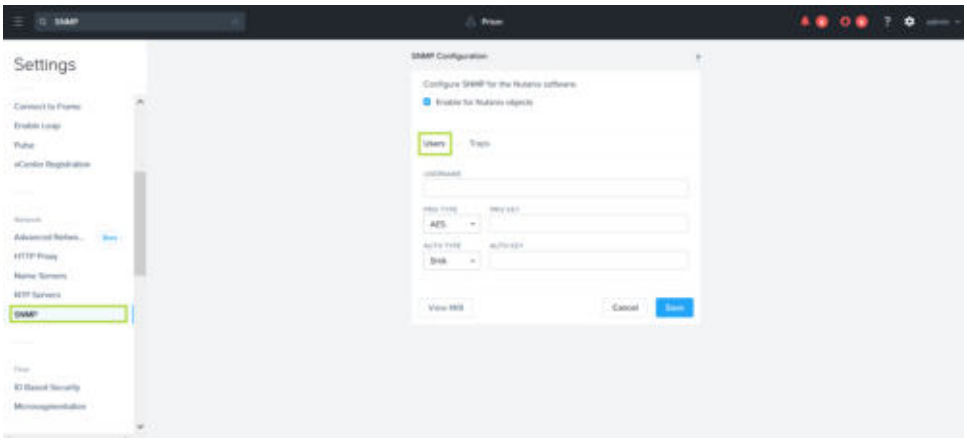

Figure 224: SNMP Configuration: Users Tab

a. **Username**: Enter a user name.

b. **Priv Type**: Select the privacy encryption type from the pull-down list.

   The only option is **AES** (Advanced Encryption Standard).

c. **Priv Key**: Enter a privacy key phrase (password) into this field.

   The key phrase is AES encrypted when the user is created.

d. **Auth Type**: Select the authentication hash function type from the pull-down list.

   The only option is **SHA** (Secure Hash Algorithm).

e. **Auth Key**: Enter an authentication key phrase (password) into this field.

   The key phrase is SHA-1 encrypted when the user is created.

f. When all the fields are correct, click the **Save** button (lower right).

   This saves the configuration and redisplays the dialog box with the new user entry appearing in the list.

5. To add an SNMP trap receiver, click the **Traps** tab and the **New Trap Receiver** button, and then do the following in the indicated fields:

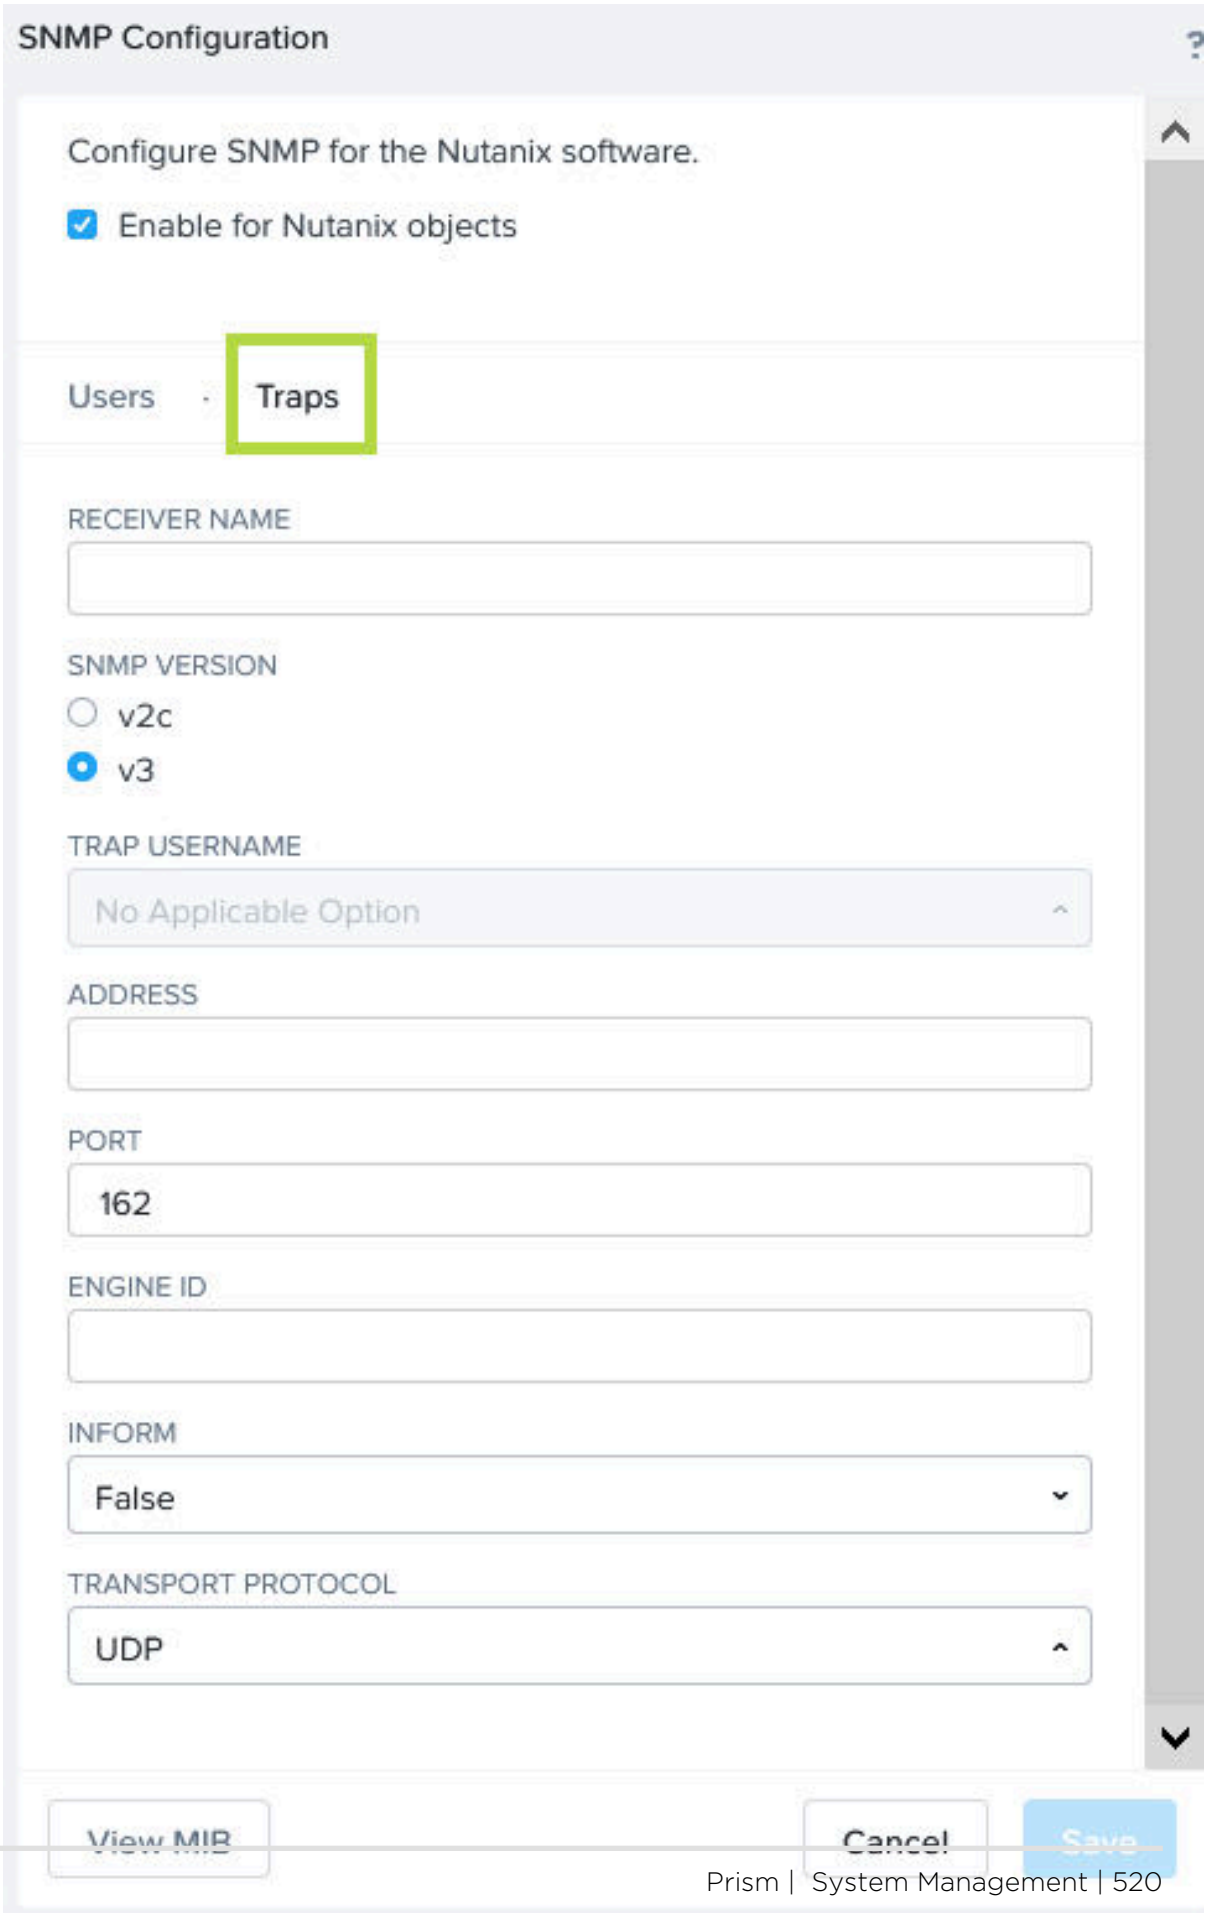

Figure 225: SNMP Configuration: Traps Tab

a. **Receiver Name**: Enter the receiver name.

b. **SNMP Version**: Select the button for the version to use, **v2c** or **v3**.

c. **Trap Username**: Select a user from the pull-down list.

   All users added previously (see step 4) appear in the pull-down list. You cannot add a trap receiver entry until at least one user has been added.

d. **Address**: Enter the target address.

   An SNMP target address specifies the destination and user that receives outgoing notifications, such as trap messages. SNMP target address names must be unique within the managed device.

e. **Port**: Enter the port number to use.

   The SNMP trap receiver uses UDP port number 162. For the complete list of required ports, see Port Reference.

f. **Engine ID** (optional): Enter an engine identifier value, which must be a hexadecimal string between 5 and 32 characters long.

   If you do not specify an engine ID, an engine ID is generated for you for use with the receiver. Every SNMP v3 agent has an engine ID that serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

g. **Inform**: Select **True** from the pull-down list to use inform requests as the SNMP notification method; select **False** to use traps as the SNMP notification method.

   SNMP notifications can be sent as traps or inform requests. Traps are one-way transmissions; they do not require an acknowledgment from the receiver. Informs expect a response. If the sender never receives a response, the inform request can be sent again. Therefore, informs are more reliable than traps. However, informs consume more resources. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform may be retried several times. The retries increase traffic and add overhead on the network. Thus, traps and inform requests provide a trade-off between reliability and resources.

h. **Transport Protocol**: Select the protocol to use from the pull-down list.

   The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.

i. When all the fields are correct, click the **Save** button (lower right).

   This saves the configuration and redisplays the dialog box with the new trap entry appearing in the list.

j. To test all configured SNMP traps, click the **Traps** tab, and then click **Test All**.

   The Nutanix cluster sends test alerts to all the SNMP trap receivers configured on the cluster.

6. To edit a user or trap receiver entry, click the appropriate tab (**Users** or **Traps**) and then click the pencil icon for that entry in the list.

   An edit window appears for that user or trap receiver entry with the same fields as the add window. (Transport entries cannot be edited.) Enter the new information in the appropriate fields and then click the **Save** button.

7. To delete an SNMP entry, click the appropriate tab (**Users** or **Traps**) and then click the X icon for that entry in the list.

   A window prompt appears to verify the delete action; click the **OK** button. The entry is removed from the list.

8. Click the **Close** button to close the SNMP Configuration window.

Configuring Syslog Monitoring

Using Prism Central, you can configure syslog monitoring to forward system logs (API Audit, Audit, Security Policy Hitlogs, and Flow Service Logs) of the registered clusters to an external syslog server. You can configure multiple remote syslog servers for Prism Central. Additionally, you can configure separate log modules to be sent to each of the rsyslog servers.

About this task

To configure syslog monitoring, do the following.

Before you begin

- You must have the IP address of the syslog server that is deployed in your environment.
- For forwarding Flow logs, the Flow feature must be enabled.

Procedure

1. Click the gear icon in the main menu and then select **Syslog Server** in the **Settings** page.
2. Click **Configure Syslog Server**.

   Note: You can configure only one syslog server per cluster.

3. In **Server Name**, enter a descriptive name for the server.
4. Enter the **IP Address** and **Port**.
5. Select the **Transport Protocol** (TCP or UDP).
6. Optionally, click the check box to enable RELP (Reliable Logging Protocol).
7. Click **Edit** against **Data Sources**.
8. Select one or more log modules from the following log types. See Syslog Modules on page 523 for details.

   - **API Audit**
   - **Audit**
   - **Security Policy Hitlogs** (policy hitlog files logs)
   - **Flow Service Logs** (Flow processes logs)

   CAUTION: Users recommend that you only configure rsyslog modules listed in Syslog Modules on page 523. Configuration of other modules using ncli commands will be over-written with any modifications done on the Prism Central web console.

9. Select the **Severity Level** from the following available options. See Syslog Modules on page 523 for details.

   Click **Save** to complete.

Syslog Modules

Refer to the following table to understand the information that is sent to the syslog server based on the selected log modules and severity levels.

Table 705: Log Modules and Severity Levels

Module Name	Severity Level	Description
API Audit	0-7	API Audit logs send information about REST API endpoints that are called and who called them from both Prism Central and Prism Element. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent.
Audit	0-7	Audit logs send information about VM, Category, and Security Policy creation, update, and delete operations. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent. See Table 707: Syslog Module: AUDIT Fields on page 524 for information on Syslog Audit fields.
Security Policy Hit Logs	All Levels	Security Policy Hit Logs send information about network flows that are acted on by a security policy. This information includes the source and destination IP address, protocol, and port. It also includes the action taken by the policy such as allow, monitor, or drop. Network statistics are also included in each message for bytes sent and received in the flow. Messages are always sent to the syslog server as level 6 - informational.
Flow Service Logs		Flow service logs contain detailed debug information used for troubleshooting the backend processes used by Flow. Note: Do not enable these unless instructed to do so by Nutanix support for troubleshooting purposes. Unlike the other data sources, the syslog level selected for Flow Service Logs has an impact on the types of data sent to the syslog server.
	0 - 5	Critical service failure messages as well as log rotation indications. Sent at the selected log level to syslog.
	6 - 7	Detailed log messages from the various components that make up the Flow product on both Prism Central and Prism Element.

The following table provides information on the syslog severity levels.

Table 706: Syslog Severity Levels

Value	Severity	Keyword	Description
0	Emergency	emerg	System is unusable.
1	Alert	alert	Should be corrected immediately.
2	Critical	crit	Critical conditions.
3	Error	err	Error Conditions.
4	Warning	warning	May indicate that an error will an action is not taken.
5	Notice	notice	Events that are unusual, but not error conditions.
6	Informational	info	Normal operational messages that require no action.
7	Debug	debug	Information useful to developers for debugging the application.

The following table provides information on the AUDIT fields of the Syslog Module.

Table 707: Syslog Module: AUDIT Fields

AUDIT Field	Description
entityType	Indicates the resource type of the entity that is associated with the audit record. Example. The entity type is `vm` for operations on virtual machines.
name	Indicates the name of the entity that is associated with the audit record. Example. Name of the virtual machine for audits or operations on virtual machines.
uuid	Indicates the unique identifier of the entity that is associated with the audit. uuid is also used to fetch the details of the entity instance.
alertUid	Indicates the name of the audit record. Example. `VmMigrateAudit` and `VmCreateAudit`.
creationTimestampUsecs	Indicates the time (in microseconds) from Epoch when the audit record was created.
defaultMsg	Provides a brief description of the current state of the operation that is captured by the audit.
opStartTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit begins.

AUDIT Field	Description
opEndTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit is completed.
operationType	Indicates the category of the audit record. Example. `Create`, `Delete`, and `Migrate`.
originatingClusterUuid	Indicates the cluster (if applicable) where the operation of is performed.
params	Additional key value pairs that capture information pertaining to the operation that is captured by the audit record.
userName	Indicates the name of the user who triggered the operation.
userUuid	Indicates the unique identifier of the user who triggered the operation.

The following table provides information on the different fields of the Security Policy Hit Logs.

Table 708: Security Policy Hit Logs Fields

Field	Description
<Timestamp>	Indicates the time of the hit log capture.
<Policy UUID>	Indicates the unique identifier of the policy that is associated with the hit log.
<Policy Name>	Indicates the name of the policy as configured by the user.
<Session Information> Create, Update, or Destroy	Indicates the session information. • Create - New session started. TCP handshake or UDP bidirectional packets. • Update - Sent every 30 seconds for active sessions. • Destroy - Session ended by TCP handshake or UDP timeout.
SRC	Indicates the source IP address.
DST	Indicates the destination IP address.
PROTO	Indicates the transport protocol.

Field	Description
ACTION	Indicates the action taken on traffic by the applied policy. • Allow - traffic allowed by policy • Drop - traffic dropped by policy • Monitor - traffic allowed by policy in monitor mode
<Packet Stats>	Indicates the packet statistics.

Modifying UI Settings

About this task

The Prism Central login page includes background animation by default, and users are logged out automatically after being idle for 15 minutes. You can change one or both of these settings.

Procedure

1. Click the gear icon and select **UI Settings** from the **Settings** > **Appearance** menu (see Settings Menu on page 16).
   The UI Settings window appears.

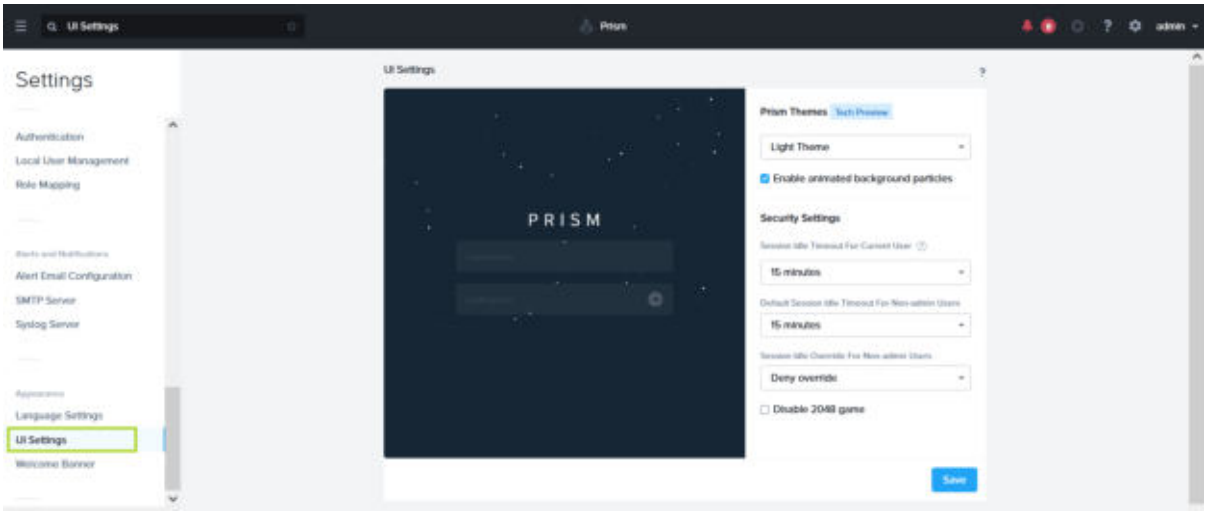

Figure 226: UI Settings Window

2. To set different variations of the Prism UI background themes, select any of the following options from the **Prism Themes** drop-down menu. The UI changes cosmetically to reflect the selection.

   Important:

   • The Prism themes feature is currently in technical preview. You may encounter visual anomalies in few settings or views where the prism themes are not applied.

For example, Licensing Settings. Nutanix recommends that you register a case on the support portal to report any anomalies.

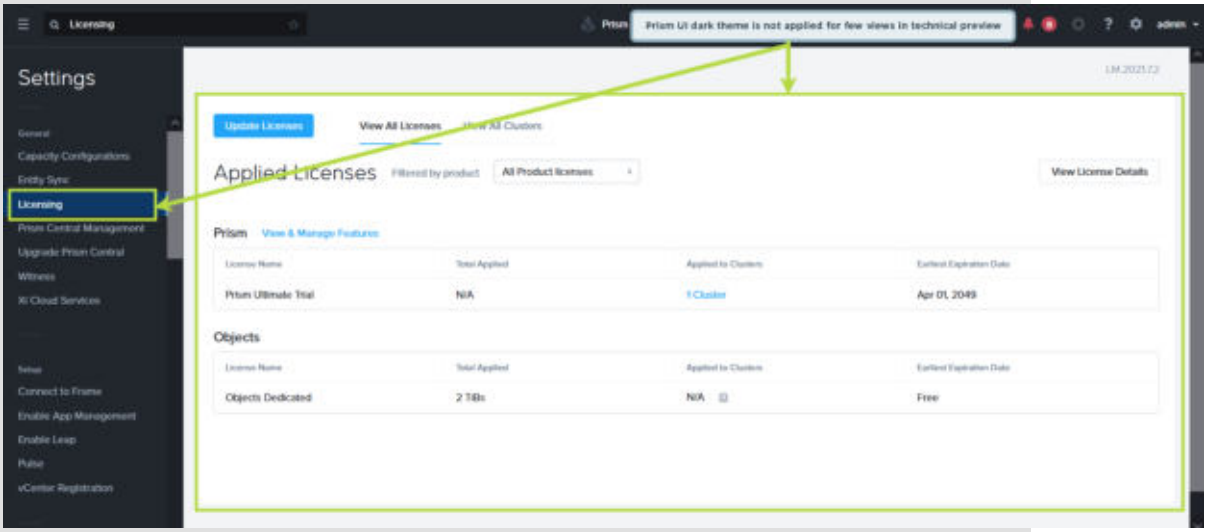

Figure 227: Example: Licensing View (Dark Theme Exception)

- You are required to Save your selection and refresh any open Prism tabs after changing a background theme.
- The Prism themes setting is not a universal feature between Prism Central and Prism Web Console. The background themes need to be configured in Prism Central and Prism Web Console respectively.

- Select **Light Theme** for light background with high contrast view. This is the default Prism background theme.
- Select **Dark Theme** for dark background with high contrast view. A pop-up appears prompting you to click **Continue** to proceed.

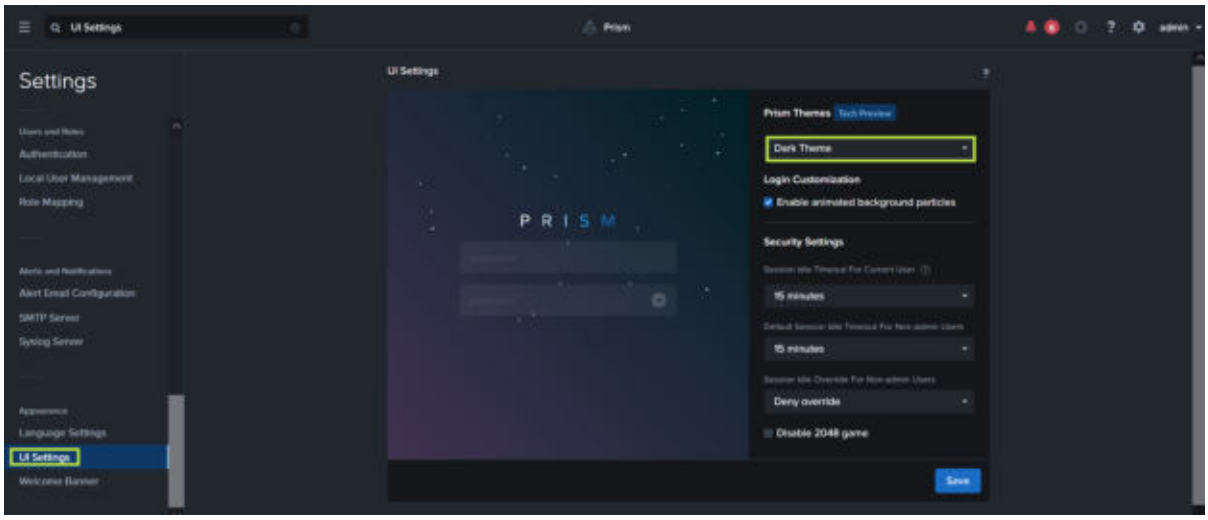

Figure 228: Prism Theme (Dark Mode)

- Select **Auto (OS defined)** to apply background themes defined in the OS. A pop-up appears prompting you to click **Continue** to proceed. For example, if you have defined

Dark mode in the OS setting, the Prism UI honors the setting and sets a dark background theme.

3. To disable the logon page background animation, uncheck the **Enable animated background particles** box (or check it to enable).

   Unchecking the **Enable animated background particles** box in the Prism UI Settings dialog box disables the creation or drawing of particles entirely. This action stops the drawing of the particles on the Prism Element logon page.

   > Note: This setting is not persistent. In other words, if the Prism service restarts, this setting is lost and must be disabled again.

   Disabling the particles allows you to conserve critical CPU resources that are used in creating and maintaining the particles.

   > Note: Disabling or enabling this setting in Prism Web Console does not propagate to Prism Central or vice versa. The setting must be disabled in Prism Web Console and Prism Central UI separately.

   You can disable the particle animation from the logon page by clicking **Freeze space-time continuum!** at the right bottom of the logon page. This action stores a setting in the local browser to stop the animation. However, this action does not stop creation or drawing of the particles itself.

   > Note: You can enable the particle animation by clicking **Engage the warp drive!** .

4. To customize the theme, background color, title text, or blurb text on the logon page, do the following:

   - Click on the top bar (displayed in the following image) in the **UI Settings** dialog box and simultaneously press the option key on the MAC system or Alt key on the Windows system. Options for customizing the theme, title text, and blurb text are displayed.
   - Select the theme from the options displayed for **Theme**. You can change the HEX codes to create your own custom gradient background color for the logon page.
   - In the **Title Text** field, enter the text to create your custom title.
   - In the **Blurb Text** field, enter the text to create your custom blurb text. This text is displayed below the password field.

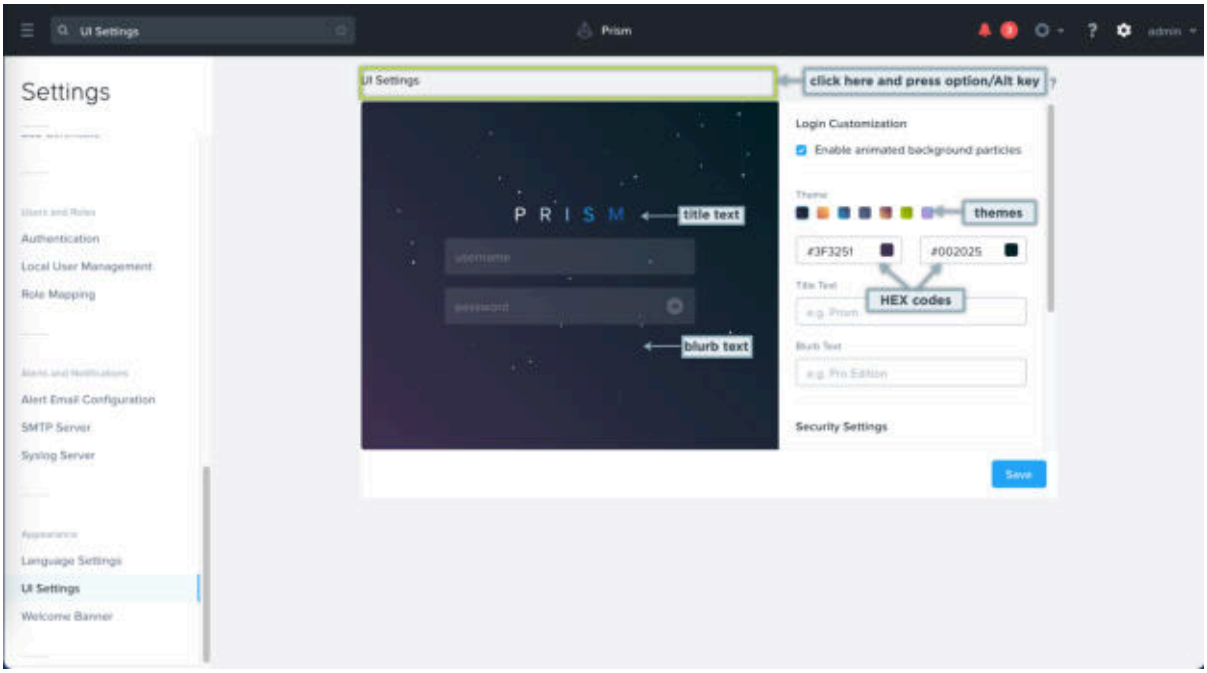

Figure 229: UI Settings Window for customizing the theme, title text, and blurb text

5. To configure session timeout, do the following under **Security Settings**:

   - Select the session timeout for the current user from the **SESSION TIMEOUT FOR CURRENT USER** drop-down list.
   - Select the default session timeout for all non-administrative users from the **DEFAULT SESSION TIMEOUT FOR NON-ADMIN USERS** drop-down list.
   - Select the appropriate option from the **SESSION TIMEOUT OVERRIDE FOR NON-ADMIN USERS** drop-down list to override the session timeout for non-administrative users.

     Note: The timeout interval for an administrator cannot be set for longer than 1 hour.

6. Clear the **Disable 2048 game** option to disable the 2048 game.
7. Click **Save** to save your changes and close the window.

Configuring a Banner Page

You have the option to create a welcome banner, which will be the first screen that appears when a user attempts to log into Prism Central. The content of the banner page is configurable, and it can include a custom message and graphics.

About this task

To configure a banner page, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Welcome Banner** in the **Settings** page.

   The Edit Welcome Banner dialog box appears.

2. Enter (paste) the desired content in HTML format in the pane on the left.

   Only "safe" HTML tags are supported. Inline event handlers, scripts, and externally-sourced graphics are not allowed.

Figure 230: Welcome Banner Window

3. Click the **Preview** button to display the banner in the pane on the right.
4. If the banner is not correct, update the HTML code as needed until the preview pane displays the desired message.
5. When the preview is correct, check the **Enable Banner** box (lower left) and then the **Save** button.

   A live banner page includes an "Accept terms and conditions" bar at the bottom. Clicking on this bar sends the user to the login page.

   You can disable (or enable) the banner at any time by unchecking (checking) the **Enable Banner** box.

Internationalization (i18n) (Prism Central)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 709: Internationalization Support for Prism Central

Supported Entities	Unsupported Entities
User-defined dashboard name	Password fields
Custom widget name	Static dashboard name
First and last name under **Update Profile**	Static widget name
User name, first name, and last name under **User Management**	
Chart name	

Localization (L10n) (Prism Central)

Nutanix localizes the user interface in Simplified Chinese and Japanese language. All the static screens are translated to the selected locale language.

You have an option to change the language settings of the cluster from English (default) to Simplified Chinese or Japanese. For more information, see Changing the Language Settings (Prism Central) on page 531.

If the Prism Element instance is launched from Prism Central, language settings of Prism Central takes precedence over Prism Element.

The dashboards (including tool tips) and menus of Prism Central are localized.

Guidelines and Limitations

- Logical entities that do not have a contextual translation available in the localized language are not translated.
- The AOS generated alerts and events are not localized to the selected locale language.
- Following strings are not localized: VM, CPU, vCPU, Language Settings, licensing details page, hardware names, storage denominations (GB, TB), About Nutanix page, EULA, service names (SNMP, SMTP), hypervisor types.

Changing the Language Settings (Prism Central)

Perform the following procedure to change the language settings in Prism Central. You can change the language setting to Simplified Chinese or Japanese.

Procedure

1. Click the gear icon in the main menu and then select **Language Settings** in the **Settings** page.

   The Language Setting window appears. The **English** language is selected by default.

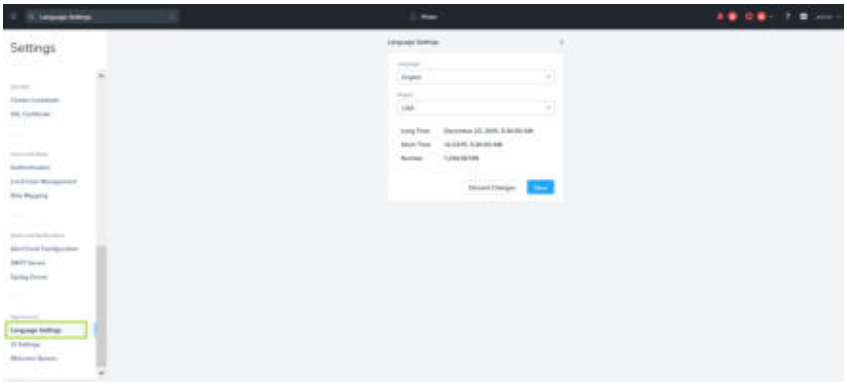

Figure 231: Language Settings Window

2. To change the setting, do the following:

   » To change the language setting of the cluster to Simplified Chinese, select **Simplified Chinese** from the drop-down menu.

   » To change the language setting of the cluster to Japanese, select **Japanese** from the drop-down menu.

   » To change the locale settings (date, time, calendar), select the appropriate region from the **Region** drop-down menu.

   By default, the locale is set to the language setting that you have set in the **Language** drop-down menu. However, you can change the **Region** to display the date, time, or calendar in some other format. This format for date, time, and calendar is applied for the entire cluster.

3. Click **Save**.

   The language and locale settings (date, time, calendar) is changed according to the selection. For example, in the below screen shot, once you click **Save** the language setting for the cluster is changed to Chinese and locale setting is changed to Russian. For more information on the entities that are supported in Simplified Chinese, see Internationalization (i18n) (Prism Central) on page 530. Also, the user interface is localized according to the

selection. For more information about localization, see Localization (L10n) (Prism Central) on page 531.

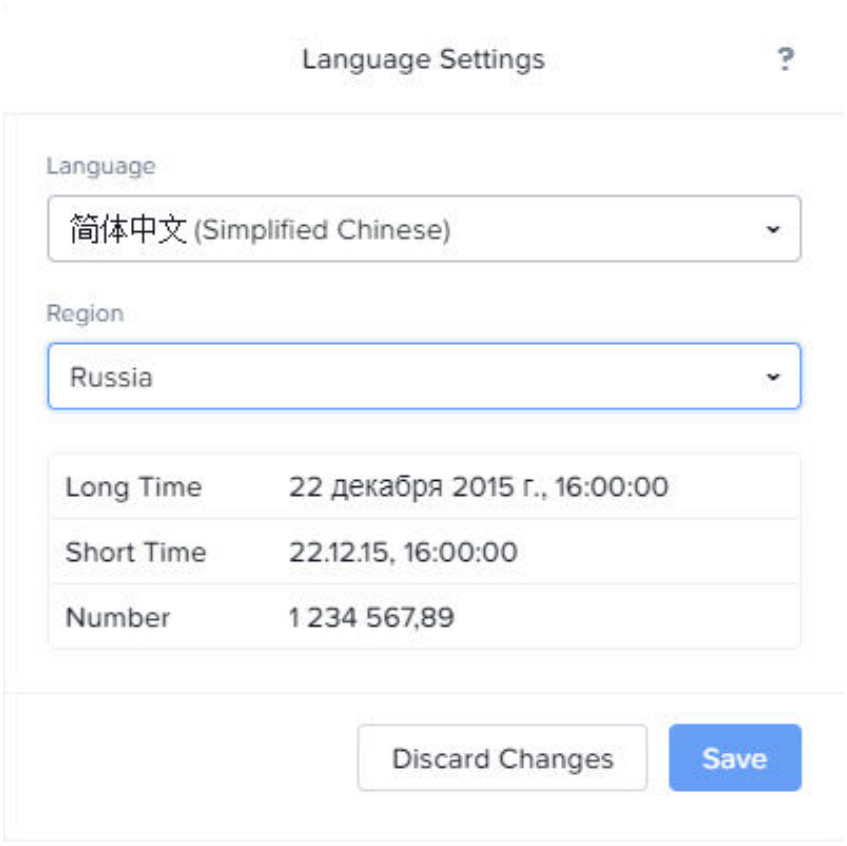

Figure 232: Localized Settings (Chinese/Russian)

7

SECURITY AND USER MANAGEMENT

Prism Central provides several mechanisms to manage security and control user access. See the following topics in the *Security Management using Prism Central (PC)* chapter of the Nutanix Security Guide.

- Set the user authentication method to local, directory service, or both (see *Configuring Authentication*)
- Add, edit, or delete local user accounts (see *Managing Local User Accounts*)
  - Update your account (see *Updating My Account*)
- Assign roles to users (see *Controlling User Access (RBAC)*)
- Install SSL certificates (see *Installing an SSL Certificate*)
- Control SSH access to Prism Central (see *Controlling Remote (SSH) Access*)
- Create and apply security-related policies (see "Security Policies" in the Flow Microsegmentation Guide).

8

COMPUTE AND STORAGE (CLUSTER) ADMINISTRATION

You can view hardware and infrastructure information about registered (individual) clusters through the Entities menu (see Hardware Entities on page 197 and Compute and Storage Entities on page 86). In addition, you can configure certain components on a registered cluster directly through Prism Central.

- To create or manage a VM on an AHV or ESXi cluster, see
  - AHV: Creating a VM (AHV) on page 536, Creating a VM (Self Service) on page 546, Managing a VM (AHV) on page 551, and Managing a VM (Self Service) on page 562
  - ESXi: Creating a VM (ESXi) on page 576 and Managing a VM (ESXi) on page 581
- To upload disk or ISO images that can by applied when creating a VM, see Image Management on page 661.
- To add images or VMs to a catalog that can be applied by a self-service portal administrator to project users, see Catalog Management on page 657.
- To configure affinity polices for an AHV cluster, see Affinity Policies Defined in Prism Central on page 599
- To perform other administrative tasks for a cluster, see Performing Other Administrative Tasks on page 706.
- To register (or unregister) an ESXi cluster with vCenter, see vCenter Server Integration on page 707.

VM Management

You can create and manage VMs directly from Prism Central when the hypervisor is either ESXi or AHV.

- ESXi
  - To create a VM, see Creating a VM (ESXi) on page 576.
  - To manage or modify a VM, see Managing a VM (ESXi) on page 581.
- AHV
  - To create a VM, see Creating a VM (AHV) on page 536.
  - To manage or modify a VM, see Managing a VM (AHV) on page 551.

- Prism Self Service (AHV)
  - To create a VM, see Creating a VM (Self Service) on page 546.
  - To manage or modify a VM, see Managing a VM (AHV) on page 551.

Creating a VM (AHV)

You can create virtual machines (VMs) in Acropolis managed clusters through Prism Central.

About this task

Figure 233: Creating a VM (AHV)

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275780958001 on Brightcove.

Note: For instructions about how to install Nutanix VirtIO on Windows VMs, see *Windows VM Provisioning* in the *AHV Administration Guide*.

To create a VM, do the following:

Note: If you are logged in as the self-service administrator or a project member, see Creating a VM (Self Service) on page 546.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 86) and click the **Create VM** button.
   The **Create VM** wizard appears.

2. In the Configuration step, do the following in the indicated fields:
   a. **Name**: Enter a name for the VM.
   b. **Description** (optional): Enter a description for the VM.
   c. **Cluster**: Select the target cluster from the drop-down list on which you intend to create the VM.
   d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.
   e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
   f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
   g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

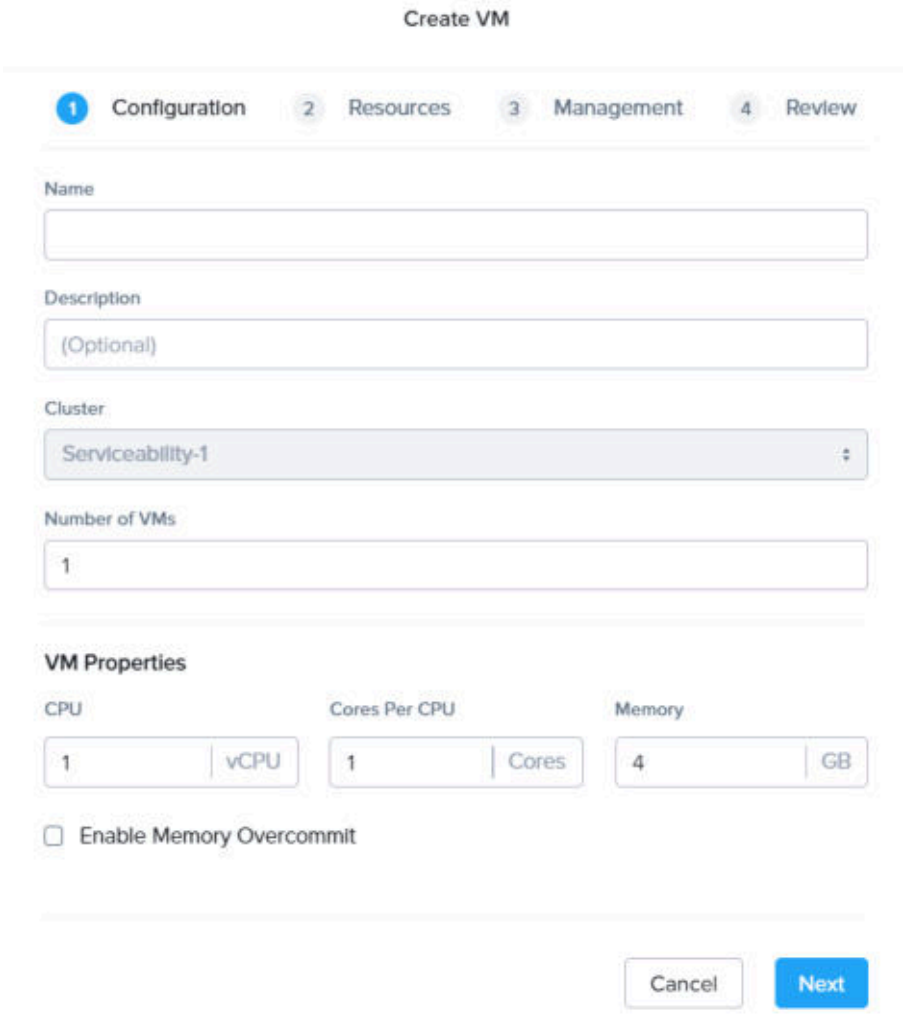

Figure 234: Create VM Window (Configuration)

3. In the **Resources** step, do the following.

   Disks: To attach a disk to the VM, click the **Attach Disk** button. The Add Disks dialog box appears. Do the following in the indicated fields:

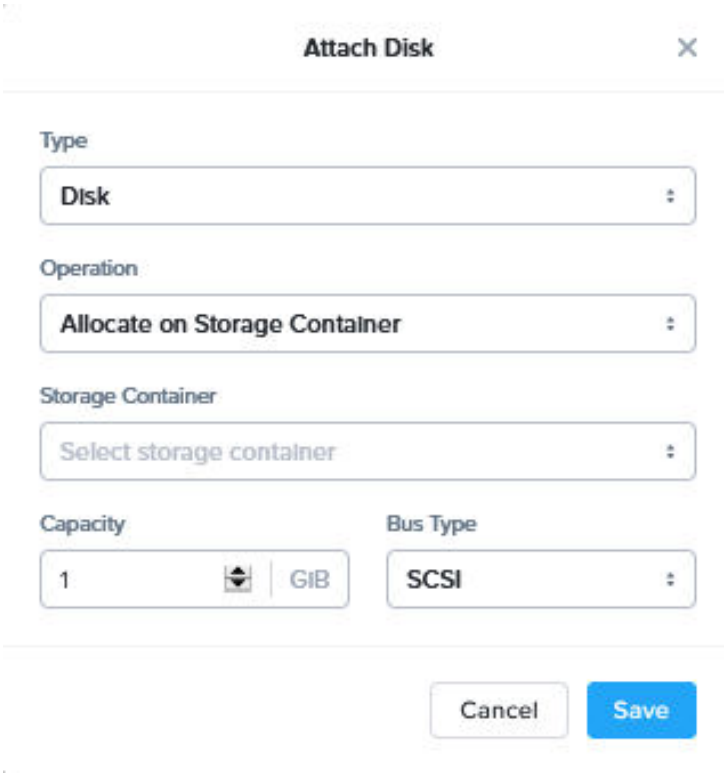

   Figure 235: Add Disk Dialog Box

   a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the drop-down list.

   b. **Operation**: Specify the device contents from the drop-down list.

      - [CD-ROM only] Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
      - [Disk only] Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
      - Select **Clone from Image** to copy an image that you have imported by using image service feature onto the disk.

   c. **Bus Type**: Select the bus type from the pull-down list.

      The options displayed in the pull-down list varies based on the storage **Type** selected in Step a.

      - For type **Disk** the choices are **SCSI**, **SATA**, **PCI**, or **SATA**.
      - For type **CD-ROM** the choices are **IDE**, or **SATA**.

   d. **Clone from Image**: Select the image that you have created by using the image service feature.

This field appears only when **Clone from Image Service** is selected. It specifies the image to copy.

> Note: If the image you created does not appear in the list, see KB-4892.

> Note: The image transfer can trigger image bandwidth throttling if a bandwidth throttling policy is associated with the image. For more information, see Bandwidth Throttling Policies on page 685.

e. **Storage Container**: Select the storage container to use from the drop-down list.

   This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.

f. **Capacity**: Enter the disk size in GiB.

g. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the Create VM dialog box.

h. Repeat this step to attach additional devices to the VM.

4. In the Resources step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The Attach to Subnet dialog box appears.

   Do the following in the indicated fields:

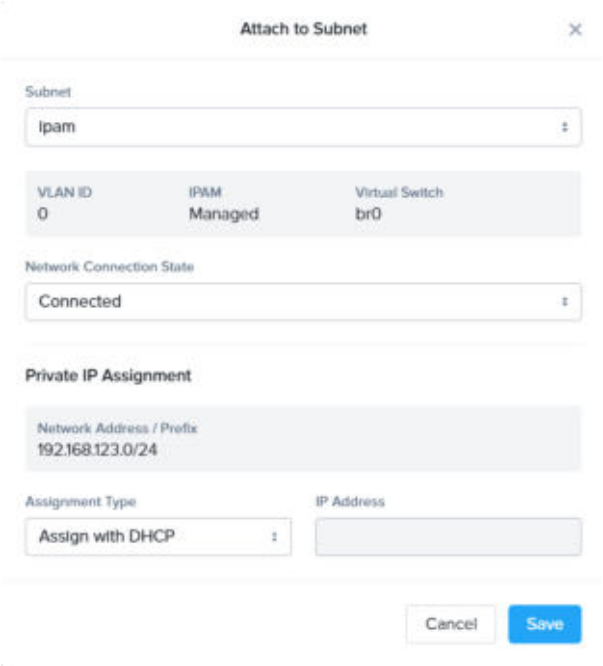

Figure 236: Attach to Subnet Dialog Box

   a. **Subnet Name**: Select the target subnet from the drop-down list.

      The list includes all defined networks (see *Network Configuration for VM Interfaces* in *Prism Web Console Guide*).

   b. **VLAN ID**: This is a read-only field that displays the VLAN ID.

   c. **IPAM**: This is a read-only field that informs you if the subnet is IPAM managed or not.

   d. **Virtual Switch**: This is a read-only field that displays the name of the virtual switch associated with the subnet.

   e. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are `Connected` or `Disconnected`.

   f. **Network Address/Prefix**: This is a read-only field that displays the network IP address and prefix.

   g. **Assignment Type**: Select the IP assignment type from the drop-down list. The options are `Assign with DHCP`, `Assign Static IP`, and `No Private IP`.

   h. **IP Address**: Enter an IP address for the VLAN.

      This field appears only if the NIC is placed in a managed network. Entering an IP address in this field is optional when the network configuration provides an IP pool. If the field is left blank, the NIC is assigned an IP address from the pool.

   i. When all the field entries are correct, click the **Save** button to create a network interface for the VM and return to the Create VM dialog box.

j. Repeat this step to create additional network interfaces for the VM.

5. In the Resources step, do the following.

   Boot Configuration: Select one of the following firmware to boot the VM.

   » **Legacy BIOS**: Select legacy BIOS to boot the VM with legacy BIOS firmware.

   » **UEFI**: Select UEFI to boot the VM with UEFI firmware. UEFI firmware supports larger hard drives, faster boot time, and provides more security features. For more information about UEFI firmware, see *UEFI Support for VM* section in the *AHV Administration Guide*.

6. In the Resources step, do the following.

   GPUs: (For GPU-enabled AHV clusters only) To add a GPU, click the **Add GPU** button. The Add GPU dialog box appears. Do the following in the indicated fields:

   a. **GPU Mode**: Click the radio button for the desired mode, either **vGPU** or **Passthrough**.

   b. If you selected **vGPU**, do the following:

   - **NVIDIA Virtual GPU License**: Select a license type from the drop-down list. This sets (filters the list of) available profiles. Click the help link (circled question mark) for information about the license types.

     Note:

     You can add multiple vGPUs to the VM only if you select the license for NVIDIA Virtual GPU software version 10.1 (440.53) or later.

     Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.

   - **vGPU Profile**: Click the radio button for the desired profile. Click the help links (Virtual Slice column and end of line) for more information about the profiles.

     Note:

     Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.

   - After you add the first vGPU, to add multiple vGPUs, see Adding Multiple vGPUs to the Same VM on page 570.

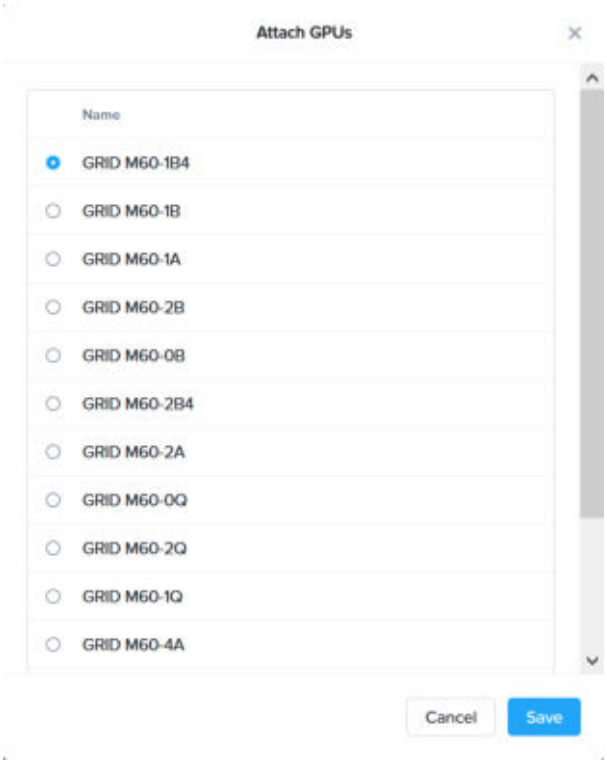

Figure 237: Add GPU Dialog Box: vGPU

c. If you selected **Passthrough**, click the radio button for the desired type.

Figure 238: Add GPU Dialog Box: Passthrough

d. Click the **Add** button.

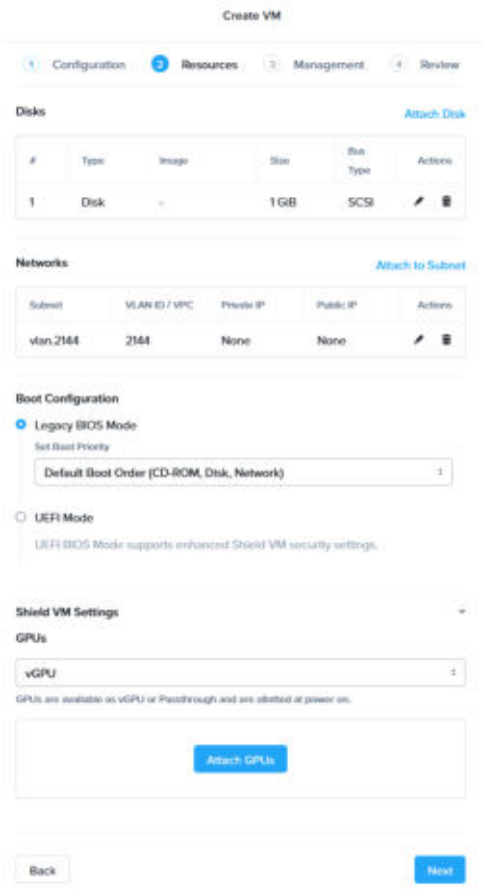

Figure 239: Create VM Window (Resources)

7. In the **Management** step, do the following.

   a. **Categories**: Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

   b. **Timezone**: Select the local timezone to use from the drop-down list.

   c. **Use this VM as an agent VM**: Select this option to make this VM as an agent VM.

      You can use this option for the VMs that must be powered on before the rest of the VMs (for example, to provide network functions before rest of VMs are powered on the

host) and must be powered off and migrated after rest of the VMs (for example, during maintenance mode operations).

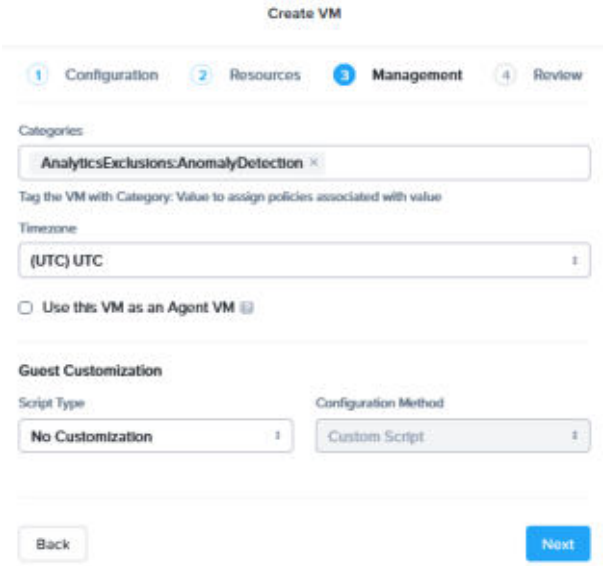

Figure 240: Create VM Window (Management)

8. In the **Management** step, do the following.

   Guest Customization: To customize the VM, select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs).

   Script Type: Select the VM script customization type from the drop-down list. The options are *No Customization*, *Sysprep(Windows)*, or *Cloud-init (Linux)*.

   Configuration Method: Select the VM script customization method from the drop-down list. This field is activated when the **Script Type** is either *Sysprep(Windows)* or *Cloud-init (Linux)*. The options are *Custom Script* or *Guided Script*.

   Field required for configuring Cloud-init and Sysprep, such as options for specifying a configuration script and option to upload script appears.

9. To specify a user data file (Linux VMs) or answer file (Windows VMs) the script for unattended provisioning, do one of the following.

   » If the file is available on your local computer, click the **Upload Script** button, choose and upload the file.

   » Create or paste the contents of the file in the text box below the **Upload Script** button.

   Note: The script type supports the following file formats.

   - Sysprep: XML
   - Cloud-init (Linux): YAML, JSON, or Shell.

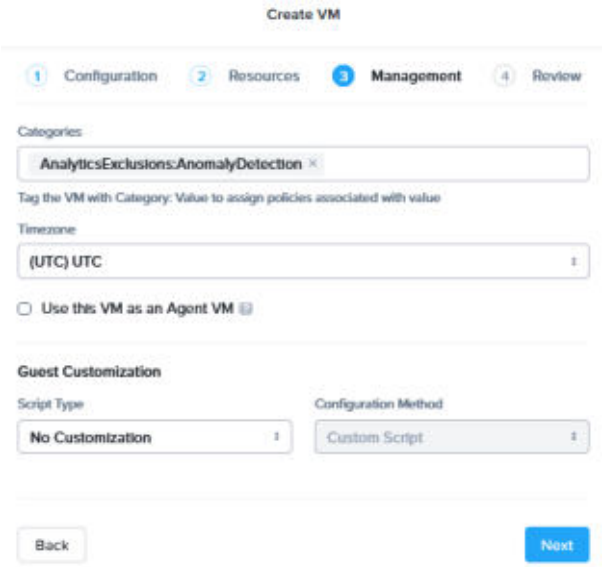

Figure 241: Create VM Window (Management)

10. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the Create VM dialog box.

    The new VM appears in the VMs entity page list.

Creating a VM (Self Service)

About this task

When you are logged in as a self-service administrator or a project member with permission to create a VM (see Prism Self Service Overview on page 713), the procedure to create a VM is different because the VM will be based on a source file stored in the Prism Central catalog (see Catalog Management on page 657). To create a VM in this case, do the following:

Note: If you are not logged in as the self-service administrator or a project member, see Creating a VM (AHV) on page 536.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 86) and click the **Create VM** button.

   The Create VM page appears.

2. In the **Select Source** tab, select the source image for the VM.

   » Click **VM Template** if you want to create a VM from a template in the catalog, and then click **Next**.

   A VM template includes all the configuration information required to create a VM. You can use a template to quickly deploy a VM without having to specify all configuration information.

   » Click **From Disk Images** if you want to create a VM from a mounted disk image, and then click **Next**.

   Disk images can be CD-ROM images such as installer ISO images or images of hard drives that contain pre-installed applications and data. Disk images enable you to share data with other VMs, but you need to specify configuration information to create the VM that will use the data.

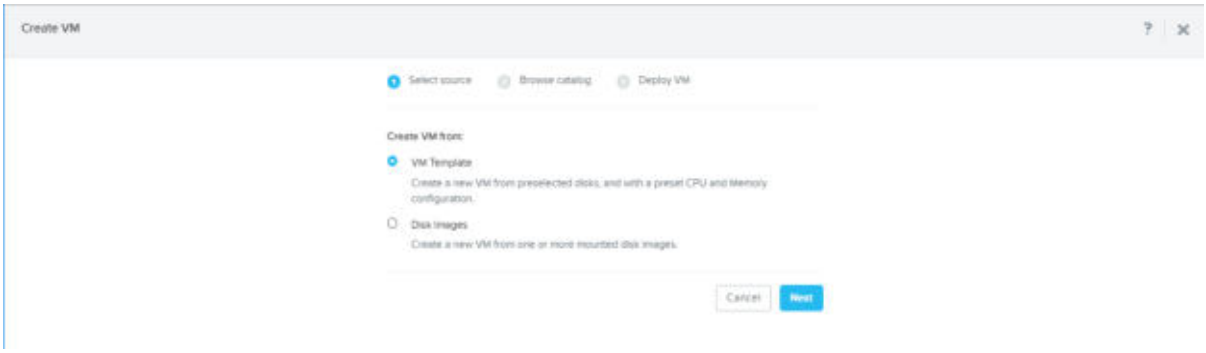

Figure 242: Create VM page: Select Source tab

3. In the **Browse Catalog** tab, do one of the following:

   - For **VM Template** select (click on) the target VM template, and then click **Next**. The available VM templates from the catalog are listed. If the target VM image does not appear in the list, you can search for it by name in the search field.

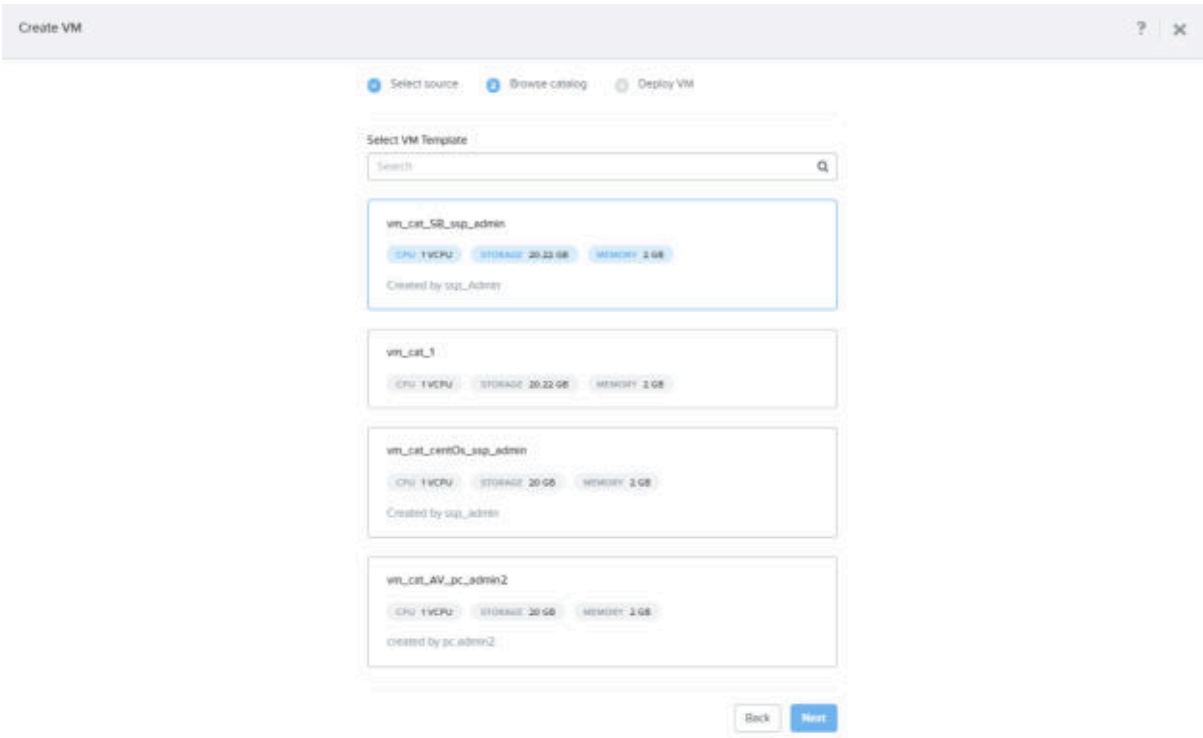

Figure 243: Create VM page: Browse Catalog tab (VM templates)

   - For **Disk Images** select (click the radio button for) one or more of the disk images, and then click **Next**. The available disk images from the catalog are listed. If the target disk image does not appear in the list, you can search for it by name in the search field.

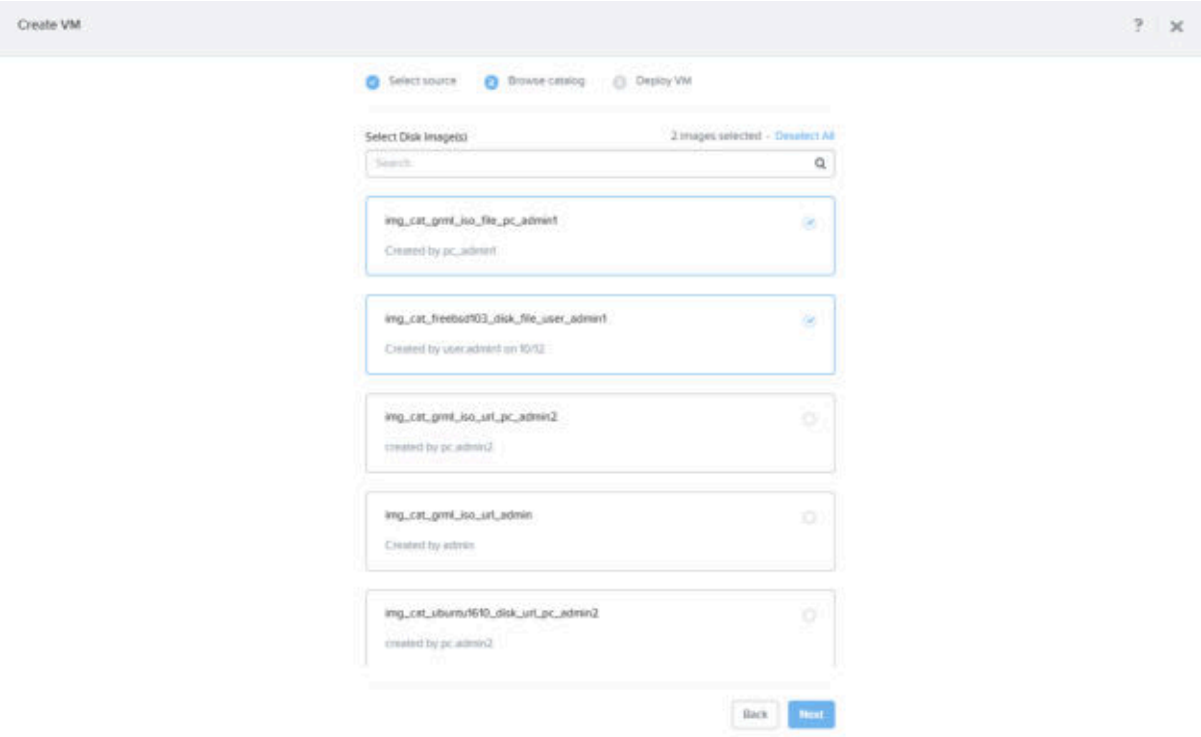

Figure 244: Create VM page: Browse Catalog tab (disk images)

4. In the **Deploy VM** tab, do the following in the indicated fields:

   a. **Name**: Enter a name for the VM.

   b. **Target Project**: Select the project associated with this VM from the pull-down list.

   c. **Disks**: Displays the disks and CD-ROMs attached to this VM.

      - For **VM Template** the list is fixed. Click the radio button for the device to boot from.
      - For **Disk Images** there are **+ New Disk** and **+ New CDROM** links above the list.
        - Click **+ New Disk** to add a disk. This displays a new line at the bottom of the list. Specify the disk size and then click **Save** at the end of line to add the disk. (The name and type field values are entered automatically.)
        - Click **+ New CDROM** to add a CD-ROM. This displays a new line at the bottom of the list. Click **Save** at the end of line to add the CD-ROM. (All field values are entered automatically.)

      Repeat this step until you have added all the desired disks and CD-ROMs. When the list is complete, click the radio button for the device to boot from.

   d. **Selected VM Template** (VM template only): Displays information about the selected VM template image. (There is nothing to do in this field.)

   e. **Guest Customization** (VM template only): If customizations were implemented when the VM template was added to the catalog, they appear here (see Adding a Catalog Item on page 658).

   f. **Network**: Select a network for the VM.

      If there is a default network specified for your project (see Creating a Project on page 718), that network is selected by default. However, you can add networks or change the default selection.

   g. **Categories** (optional): To assign the VM a category value, enter the category name and then click the plus sign (+) to the right of the field.

      This field acts like a search field; it provides a list of matching categories as you enter a string. Select the desired category value when you see it in the list. You can repeat this step for as many categories as desired.

   h. **Advanced Settings**: Check this box to specify the number of vCPUs and amount of memory allocated to this VM. Otherwise, the default values apply (1 vCPU and 2 GB of memory).

   i. When all values are correct, click the **Save** button.

      This closes the page and deploys the VM. The new VM appears in the VMs entity page list.

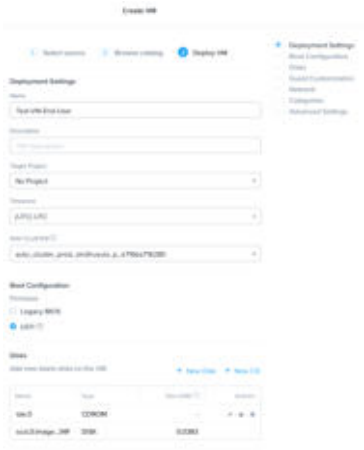

Figure 245: Create VM page: Deploy VM tab (disk images)

Figure 246: Create VM page: Deploy VM tab (network)

Managing a VM (AHV)

About this task

After creating a VM (see Creating a VM (AHV) on page 536) you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, update the network configuration, disable efficiency measurement, disable anomaly detection, power on (or off) the VM, enable flash mode for a VM and disks, assign the VM to a protection policy, migrate the VM, create VM recovery point, add the VM to a recovery plan, run a playbook, manage categories, install, manage, and upgrade Nutanix Guest Tools (NGT), configure VM host affinity, add the VM image to the catalog, manage the VM ownership, configure QoS settings and export VM as OVA.

Note: You can perform only those operations for which you have permissions from the admin.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 86) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
- Go to the details page of a selected VM (see VM Details View on page 96) and select the desired action.

> Note: The available actions appear in bold; other actions are not available. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Make the desired changes and then click the **Save** button in the **Review** step.

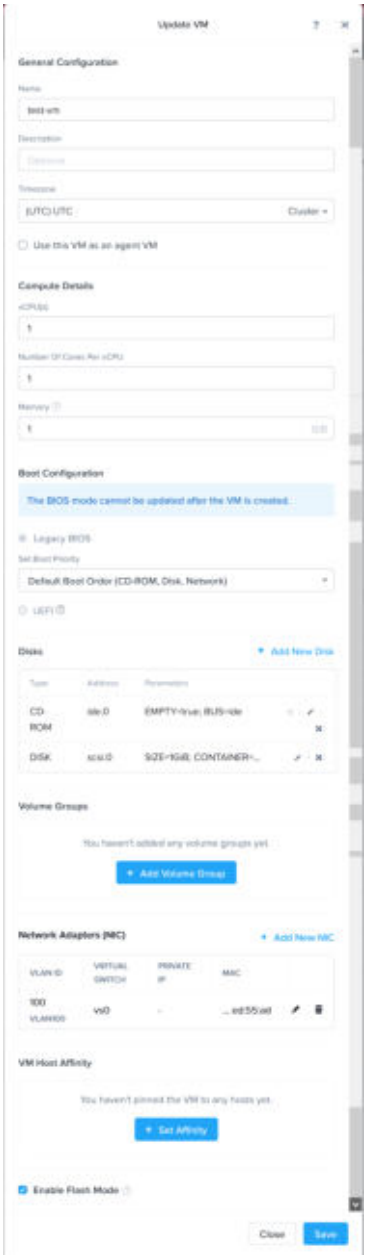

Figure 247: Update VM Window (Configuration)

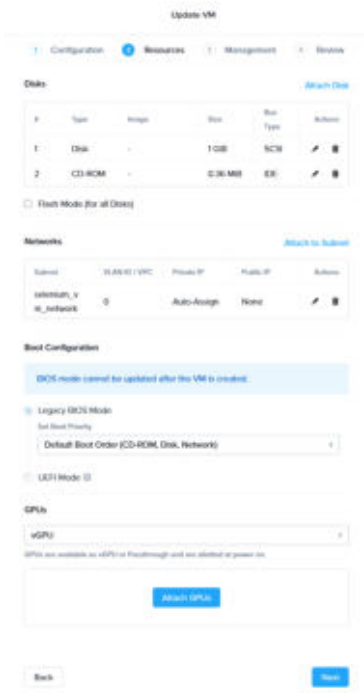

Figure 248: Update VM Window (Resources)

- **Disks**: You can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See

Creating a VM (AHV) on page 536 before you modify the disk or create a new disk for a VM. You can enable or disable the flash mode settings for the VM and VM disks.

a. To enable flash mode on the VM, click the **Enable Flash Mode** check box.

   » After you enable this feature on the VM, the status is updated in the VM table view. To view the status of individual virtual disks (disks that are flashed to the SSD), click the update disk icon in the **Disks** pane in the Update VM window.

   » You can disable the flash mode feature for individual virtual disks. To update the flash mode for individual virtual disks, click the update disk icon in the **Disks** pane and deselect the **Enable Flash Mode** check box.

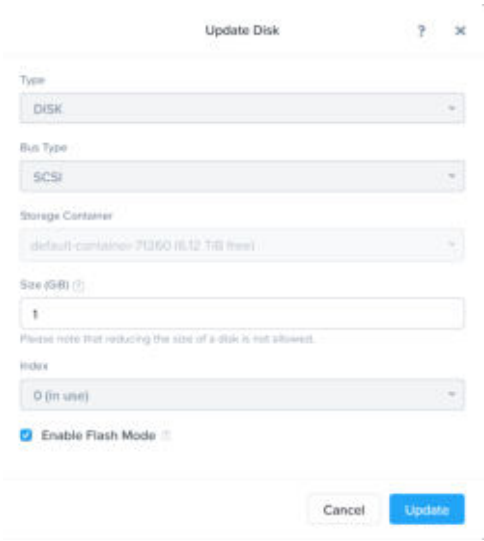

Figure 249: Update Disk Window (VM Disk Flash Mode)

Note: The flash mode setting is not supported for **CD-ROM** storage device.

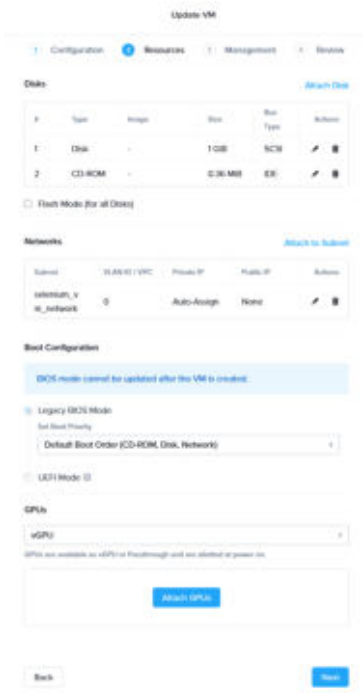

Figure 250: Update VM Window (Resources)

- **Networks:** you can attach new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (AHV) on page 536 before you modify NIC network or create a new NIC for a VM.
- **Boot Configuration:** You cannot modify the BIOS mode after the VM is created. However, if you have enabled the Secure Boot option under UEFI Mode while creating the VM, you can modify this setting during the VM update.
- **GPUs:** (For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the specified vGPU license. However, you must shut down the VM before you perform these operations.

  - Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide.*
  - Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
  - After you add the first vGPU, to add multiple vGPUs, see *Adding Multiple vGPUs to the Same VM* in the *AHV Admin Guide*.

- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

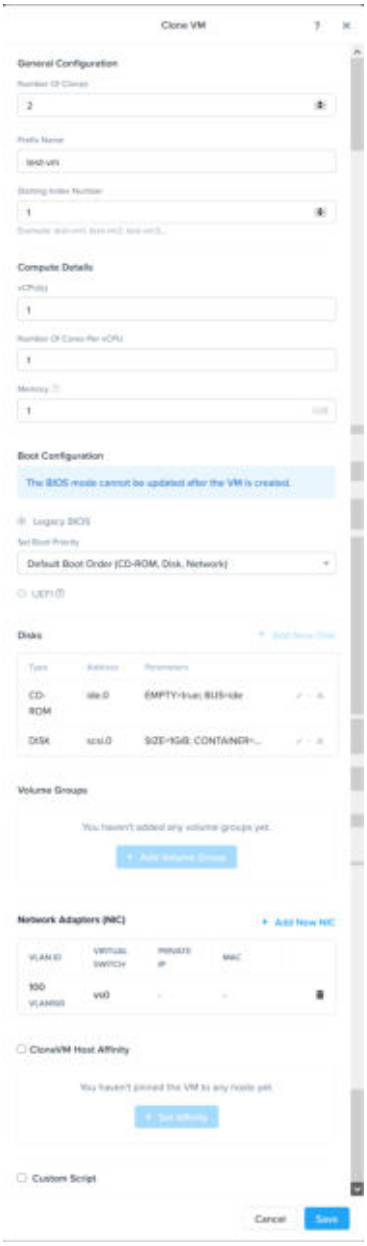

Figure 251: Clone VM Window

- To create a VM template, select **Create VM Template**.

  This displays the Create template from VM dialog box.

  > Note: Before selecting a VM for creating a VM template, ensure that the VM is powered off.

  Enter a name and description (optional) for the VM template. You can choose to customize the operating system (OS) of the VMs that will be deployed using this template. Select Sysprep for customizing the Windows OS, and Cloud-init for customizing the Linux OS. For configuring each of these script types, you can choose to either upload a custom script or opt for a guided setup. The guided setup lets you enter the authentication information such as username, password, Locale, hostname, domain join, and license key.

  You can also specify if you want to allow the template users to override the template settings at the time of VM deployment.

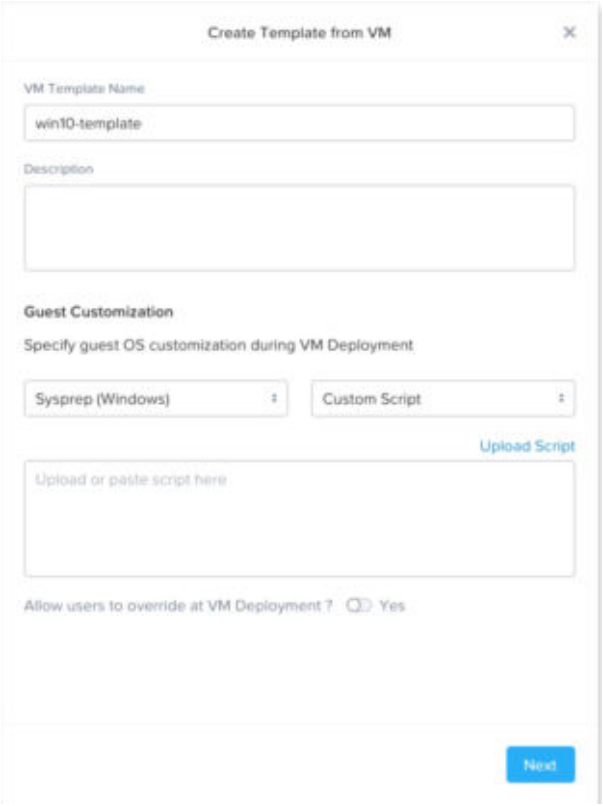

Figure 252: Create VM Template

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM

table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

> Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

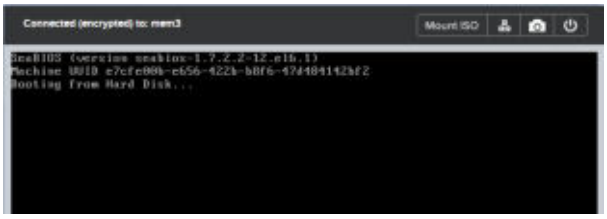

Figure 253: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To assign the VM to a protection policy, select **Protect**. This opens a page to specify the protection policy to which this VM should be assigned (see Policies Management). To remove the VM from a protection policy, select **Unprotect**.

  > Note: You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Leap and configuring the Availability Zone.

- To create VM recovery point, select **Create Recovery Point**.

  This displays the Create VM Recovery Point dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

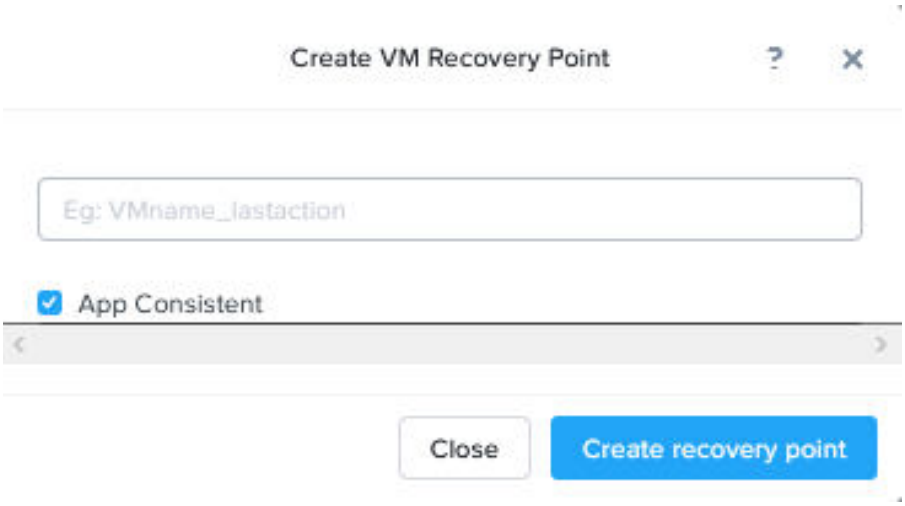

Figure 254: Create VM Recovery Point Window

- To migrate the VM to another host, select **Migrate**.

  This displays the Migrate VM dialog box. Select the target host from the pull-down list (or select the **System will automatically select a host** option to let the system choose the host) and then click the **Migrate** button to start the migration.

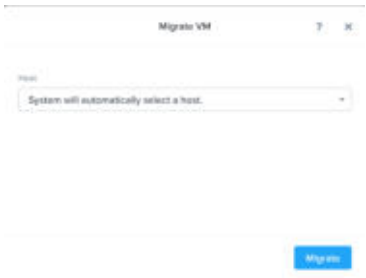

Figure 255: Migrate VM Window

> Note: Nutanix recommends to live migrate VMs when they are under light load. If they are migrated while heavily utilized, migration may fail because of limited bandwidth.

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the "Explicitly Adding a VM to a Recovery Plan" section in the Leap Administration Guide).
- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 875).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 737).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 591).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 593).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > Note: If you clone a VM, by default NGT is not enabled on the cloned VM. You must again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the *Prism Web Console Guide*.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

```
ncli> ngt mount vm-id=virtual_machine_id
```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

```
ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-
```

c1601e759987

Note:

- Self-service restore feature is not enabled by default on a VM. You must manually enable the self-service restore feature.
- If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You must unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.

- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 594).
- To establish VM host affinity, select **Configure VM Host Affinity**.

  A window appears with the available hosts. Select (click the icon for) one or more of the hosts and then click the **Save** button. This creates an affinity between the VM and the selected hosts. If possible, it is recommended that you create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure. See Affinity Policies Defined in Prism Central on page 599 for more information about the VM affinity policies.

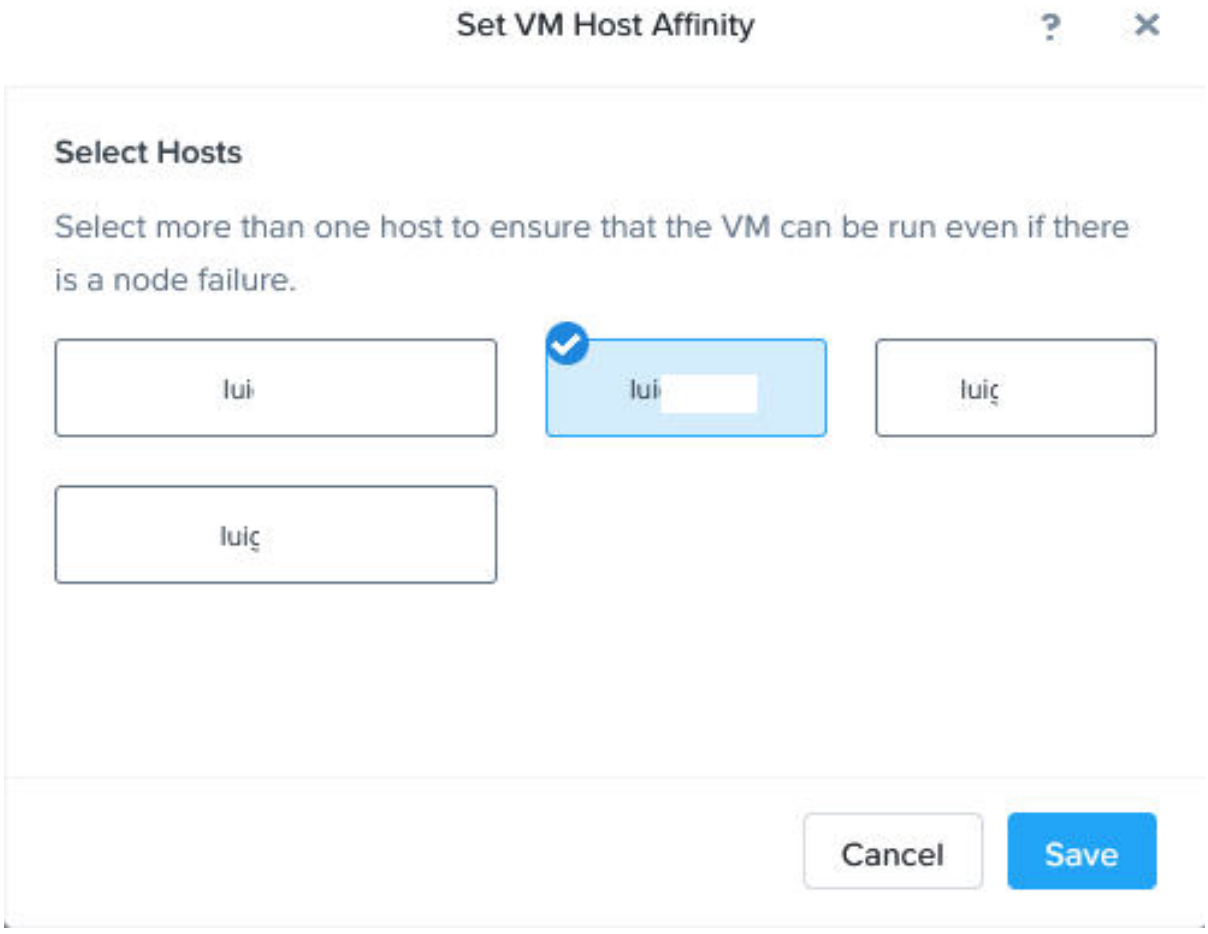

Figure 256: Set VM Host Affinity Window

- To add a VM to the catalog, select **Add to Catalog**. This displays the **Add VM to Catalog** page (see Adding a Catalog Item on page 658).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

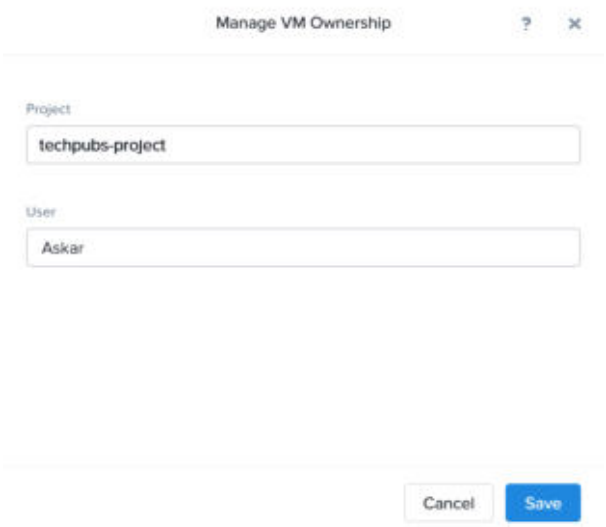

Figure 257: VM Ownership Window

- To configure quality of service (QoS) settings, select **Set QoS Attributes** (see Setting QoS for an Individual VM on page 597).
- To export as OVA, select **Export as OVA** (see OVA Management on page 688 and Uploading an OVA on page 690).

Managing a VM (Self Service)

About this task

After creating a VM (Creating a VM (Self Service) on page 546), you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, power on (or off) the VM, disable efficiency measurement, disable anomaly detection, set flash mode settings for a VM and VM disks, assign the VM to a protection policy, migrate the VM, add the VM to a recovery plan, run a playbook, manage categories, install and manage Nutanix Guest Tools (NGT), configure VM host affinity, add the VM image to the catalog, manage the VM ownership, configure QoS settings, or export VM as OVA.

> Note: You can perform only those operations for which you have permission.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 86) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.

- Go to the details page of a selected VM (see VM Details View on page 96) and select the desired action.

  Note: The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The Edit VM dialog box appears, which includes the same fields as the Create VM dialog box. Make the desired changes and then click the **Save** button.

  (For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the specified vGPU license. However, you need to power off the VM before you perform these operations.

  - Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.
  - Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
  - After you add the first vGPU, to add multiple vGPUs, see *Adding Multiple vGPUs to the Same VM* in the *AHV Admin Guide*.

Figure 258: Edit VM Window

- **Disks:** You can add new disks to the VM using the **+ New Disk** option. You can also add new CDs to the VM using the **+ New CD** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM (Self Service) on page 546 before you modify the disk or create a new disk for a VM.

- **Network:** You can attach new network to the VM using the **Connect to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (Self Service) on page 546 before you modify NIC network or create a new NIC for a VM.
- **Configuration:**You can update the VM category value, enter the category name and then click the plus sign (+) to the right of the field. See Creating a VM (Self Service) on page 546 before you modify or assign new categories for a VM.
- **Categories:**You can update the .
- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot enabled.

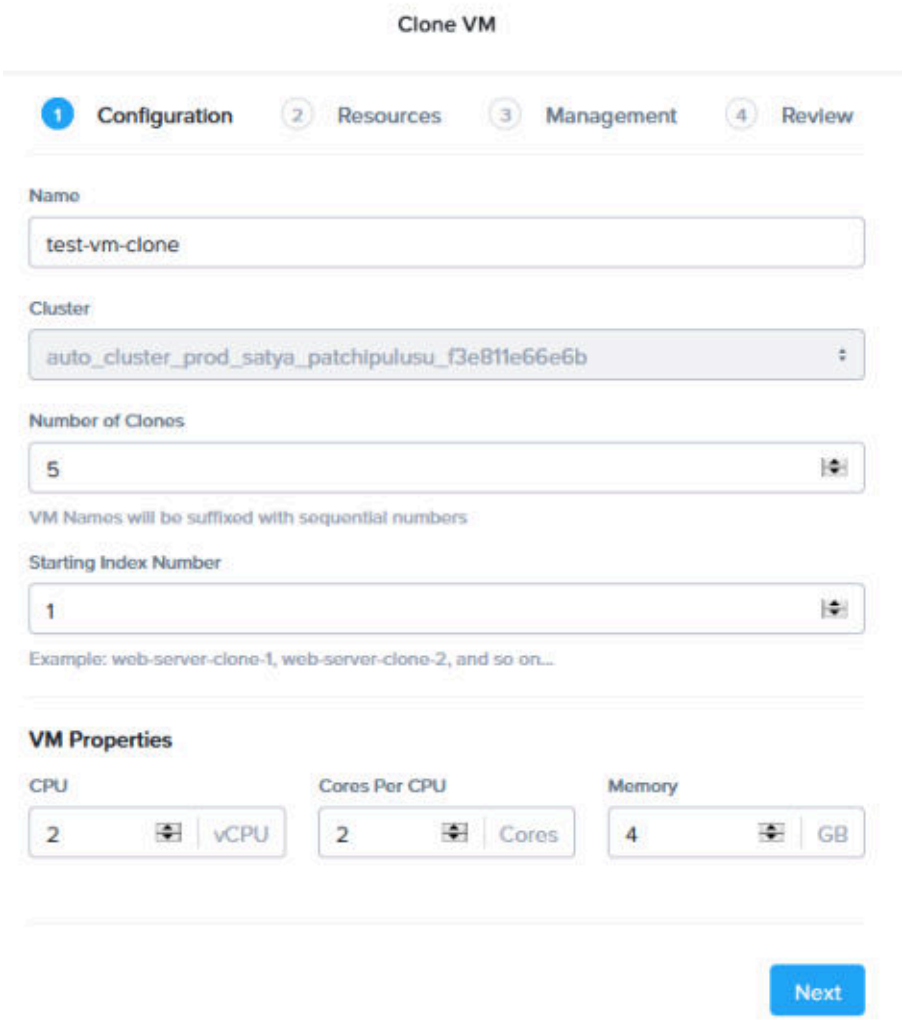

Figure 259: Clone VM Window (Configuration)

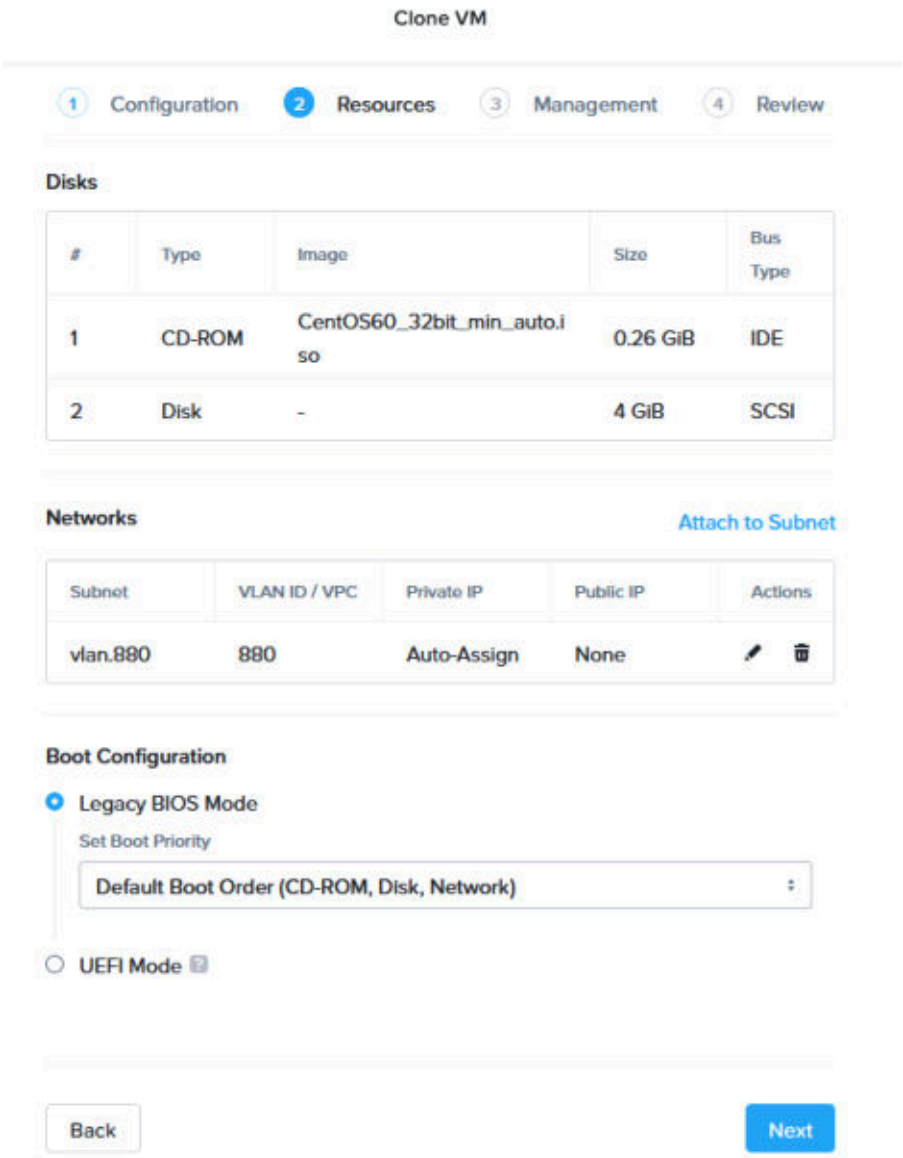

Figure 260: Clone VM Window (Resources)

Figure 261: Clone VM Window (Management)

Figure 262: Clone VM Window (Review)

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

  > Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

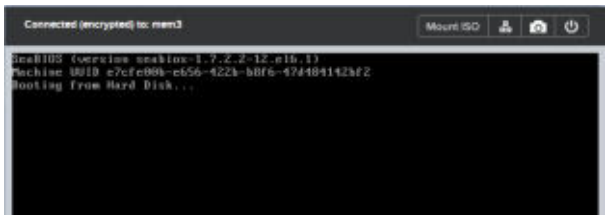

Figure 263: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To create VM recovery point, select **Create Recovery Point**.

  This displays the Create VM Recovery Point dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

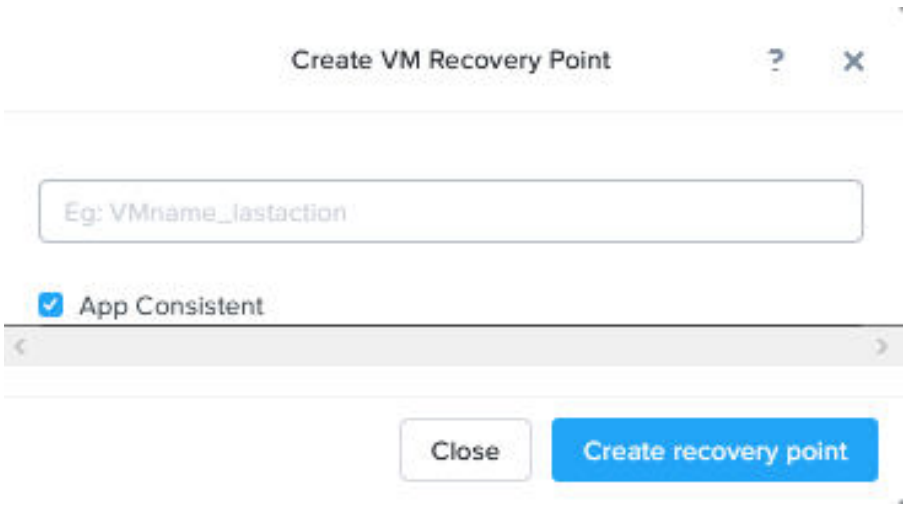

Figure 264: Create VM Recovery Point Window

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the "Explicitly Adding a VM to a Recovery Plan" section in the Leap Administration Guide).

- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 875).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 737).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 591).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 593).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > Note: If you clone a VM, by default NGT is not enabled on the cloned VM. You need to again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the *Prism Web Console Guide*.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

  ```
  ncli> ngt mount vm-id=virtual_machine_id
  ```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

  ```
  ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-
  c1601e759987
  ```

  > Note:
  >
  > - Self-service restore feature is not enabled by default on a VM. You need to manually enable the self-service restore feature.
  > - If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.
- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 594).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

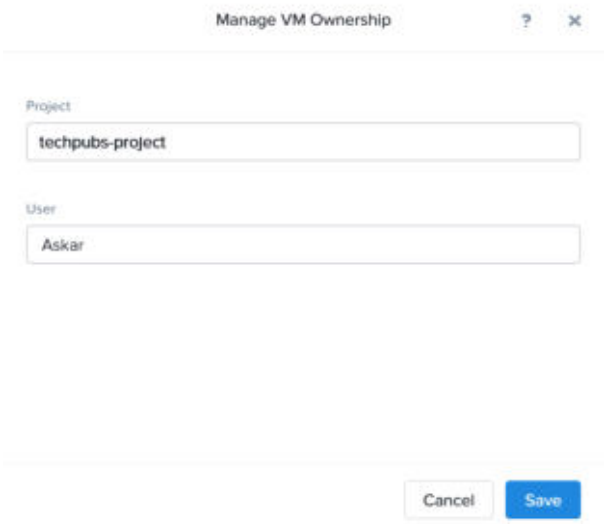

Figure 265: VM Ownership Window

- To export as OVA, select **Export as OVA** (see OVA Management on page 688 and Uploading an OVA on page 690).

Adding Multiple vGPUs to the Same VM

About this task

You can add multiple vGPUs of the same vGPU type to:

- A new VM when you create it.
- An existing VM when you update it.

> Important:
>
> Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.

After you add the first vGPU, do the following on the Create VM or Update VM dialog box (the main dialog box) to add more vGPUs:

Procedure

1. Click **Add GPU**.

2. In the Add GPU dialog box, click **Add**.

   The License field is grayed out because you cannot select a different license when you add a vGPU for the same VM.

   The **VGPU Profile** is also auto-selected because you can only select the additional vGPU of the same vGPU type as indicated by the message at the top of the dialog box.

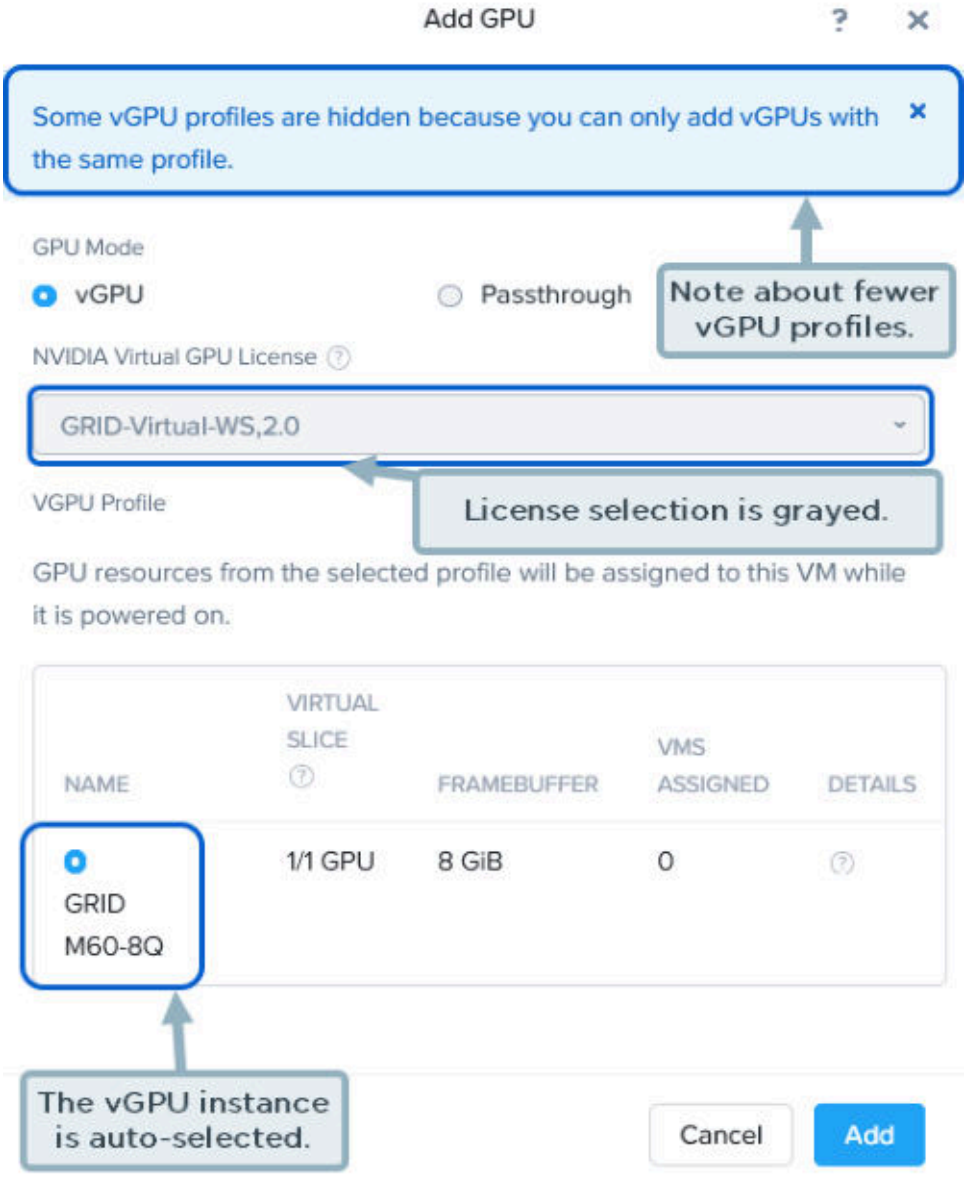

Figure 266: Add GPU for multiple vGPUs

3. In the main dialog box, you see the newly added vGPU.

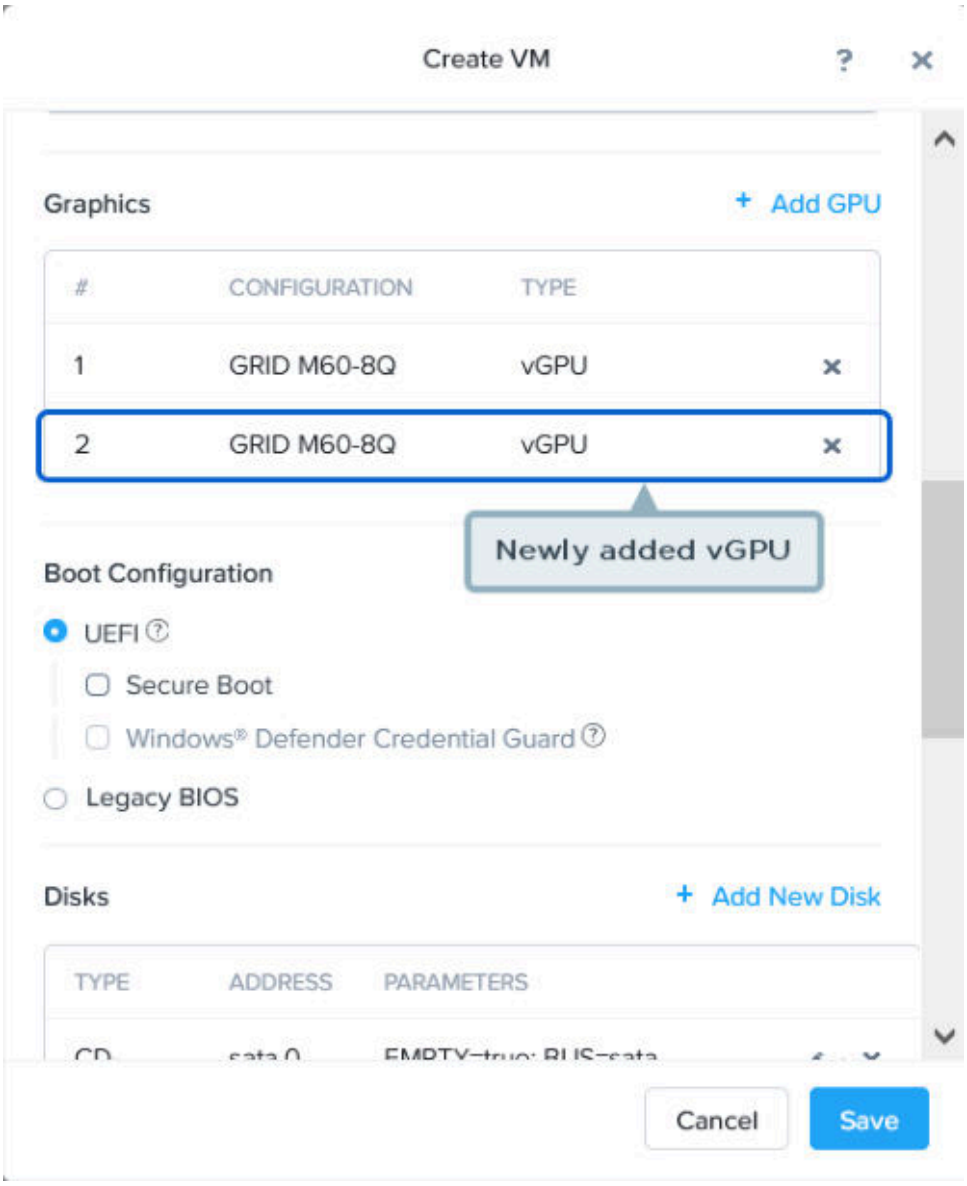

Figure 267: New vGPUs Added

4. Repeat the steps for each vGPU addition you want to make.

Live Migration of vGPU-enabled VMs

You can live-migrate a vGPU-enabled VM to another host within the cluster. You can also migrate the vGPU-enabled VM to a host outside the cluster, that is, a host in another cluster.

Note: Live migration of VMs with vGPUs is supported:

- For vGPUs created with minimum NVIDIA Virtual GPU software version 10.1 (440.53).
- Minimum AOS version that supports live migration within the same cluster is 5.18.1.
- Minimum AOS version that supports live migration across cluster is 6.1.

Ensure the following:

1. The VM is not powered off.
2. The host affinity is not set for the VM.

   If the host affinity of the VM is set to only one host, you cannot migrate the VM. However, based on the GPU resources required, if the host affinity is set to multiple hosts with the same or similar GPU resources, you can migrate the VM among the hosts with which the affinity is set.
3. You have another host in the same cluster that supports the resources necessary to support the vGPU-enabled VM that you want to live migrate.

For limitations applicable to live migration of vGPU-enabled VMs, see *Limitations for Live Migration* in the *NVIDIA GRID Virtual GPU Support on AHV* section of the *AHV Administration Guide*.

Migrating Within the Cluster

You can live-migrate a vGPU-enabled VM to another host within the cluster.

About this task

The **Migrate Within Cluster** option in the Migrate VM dialog box allows you to migrate vGPU-enabled VMs to other hosts within the clusters.

For limitations applicable to live migration of vGPU-enabled VMs, see *Limitations for Live Migration* in the *NVIDIA GRID Virtual GPU Support on AHV* section of the *AHV Administration Guide*.

To live migrate a vGPU-enabled VM to another host in the same cluster, do the following.

Procedure

1. Click **Compute & Storage** > **VMs** > **List**.
2. Select the VM you want to migrate live. Click **Actions** > **Migrate**

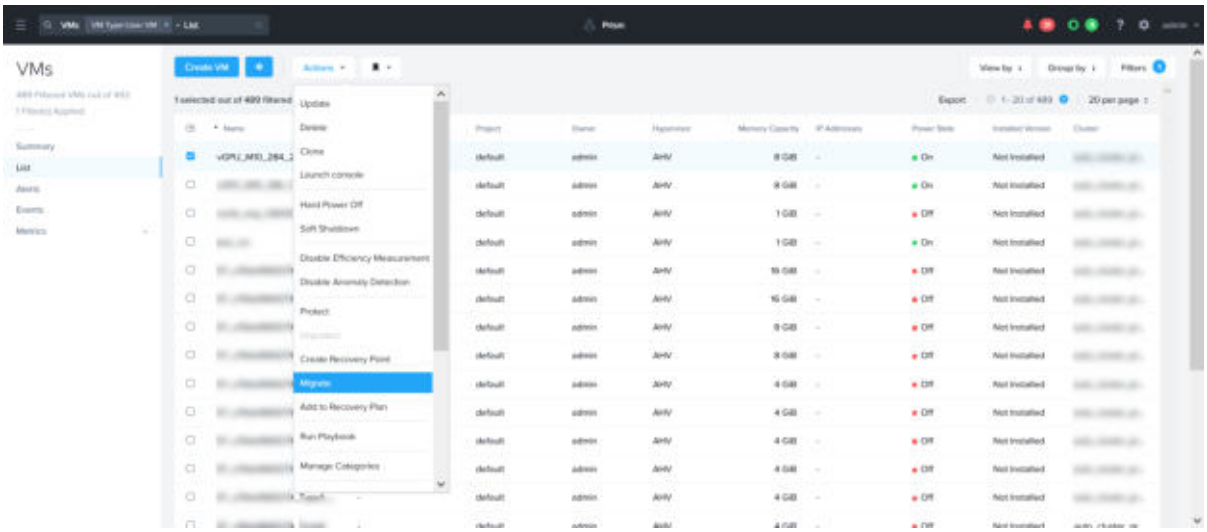

Figure 268: VM Selection

3. In the Migrate VM dialog box, the **Migrate Within Cluster** option is default selection.

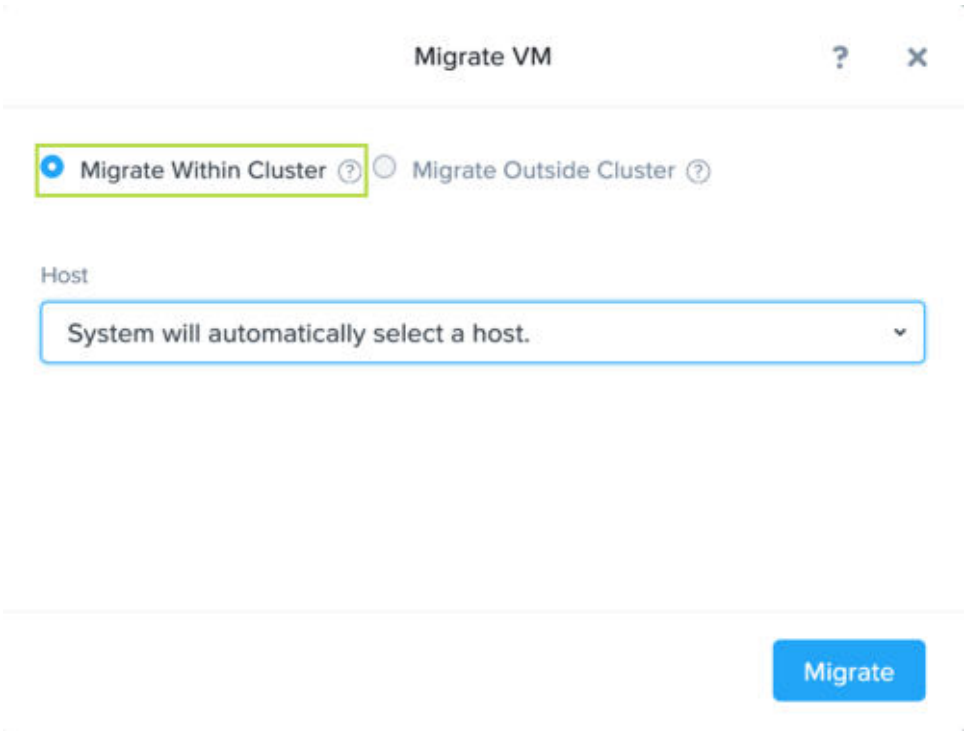

Figure 269: Migrate VM dialog box

4. In the **Host** drop-down list, do one of the following:

   - Retain the *System will automatically select a host* default option if you want to migrate the VM to a host selected by the system.

     The system selects a host based on the GPU resources available with the host as appropriate for the VM to be migrated live.

   - Select the host listed in the drop-down list that you want to migrate the VM to.

5. Click **Migrate**.

   Prism submits the task and displays the following message:

   ```
   Successfully submitted migrate operation.
   Task details
   ```

   Task details is a link to the Tasks page. Click the **Task details** link to monitor the migration task in the Tasks page.

   The host name of the VM in the List view changes to the host name to which you migrated the VM.

Migrating Outside the Cluster

You can migrate the vGPU-enabled VM to a host outside the cluster, that is, a host in another cluster.

About this task

The **Migrate Outside Cluster** option in the Migrate VM dialog box allows you to migrate vGPU-enabled VMs across clusters.

Note: Live migration to a host in another cluster is supported only if the VM is protected by protection policy with Synchronous replication schedule.

For limitations applicable to live migration of vGPU-enabled VMs, see *Limitations for Live Migration* in the *NVIDIA GRID Virtual GPU Support on AHV* section of the *AHV Administration Guide*.

To migrate a vGPU-enabled VM to a host in another cluster, do the following.

Procedure

1. Click **Compute & Storage** > **VMs** > **List**.
2. Select the VM you want to migrate live. Click **Actions** > **Migrate**

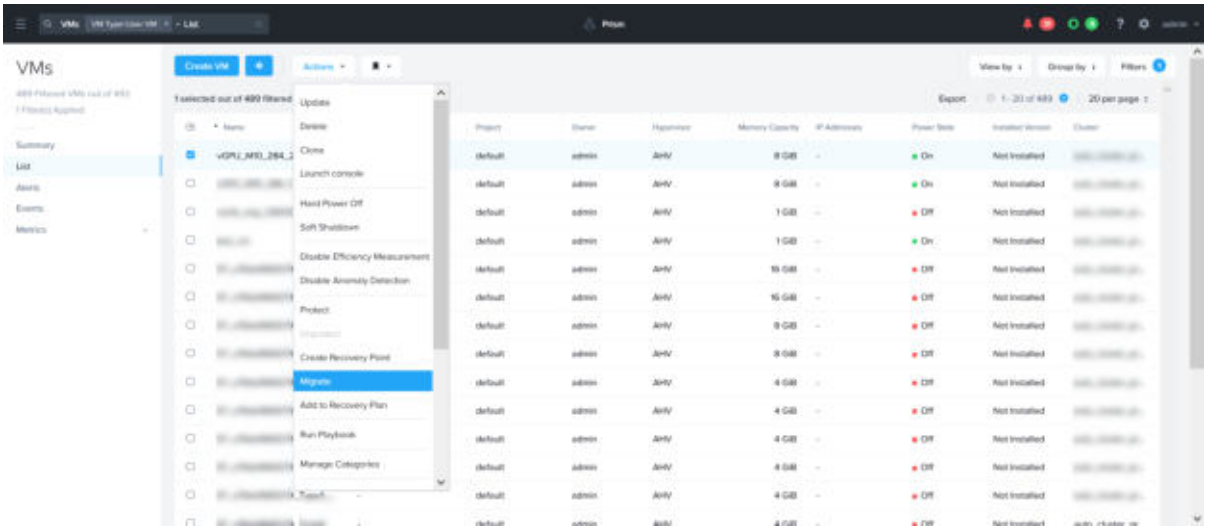

Figure 270: VM Selection

3. In the Migrate VM dialog box, select the **Migrate Outside Cluster** option.

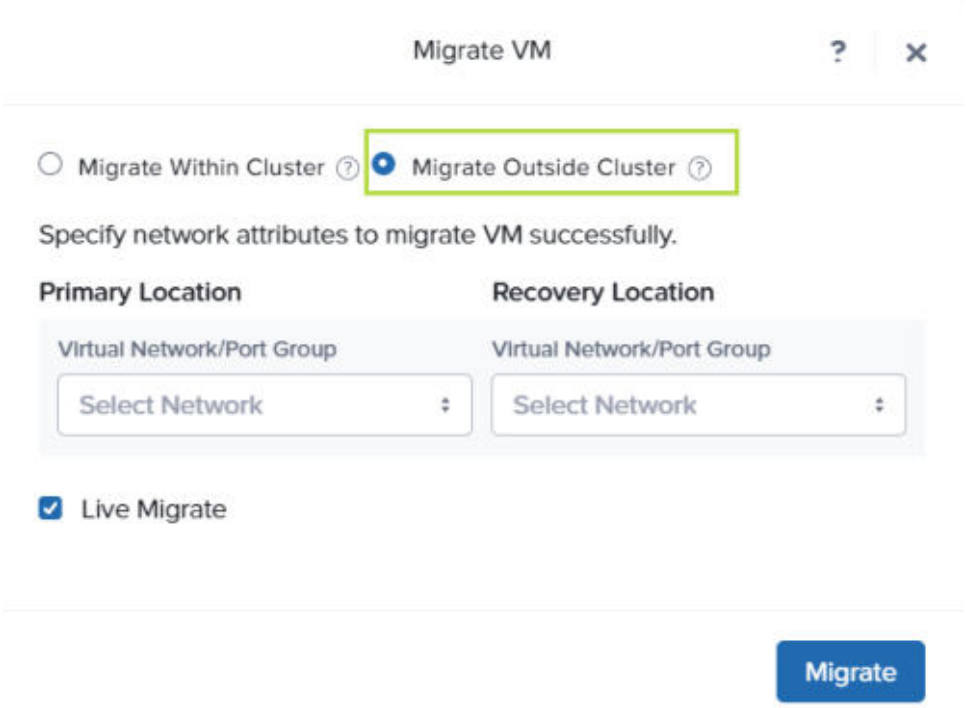

Figure 271: Migrate VM dialog box

4. In the **Select network attributes to migrate VM successfully** section, select the parameters for the following:

   Note: If the VM is already configured in a recovery plan, then the network attribute fields are already populated and in-actionable. You cannot modify or add any values in these fields.

   - **Primary Location**: Virtual Network/Port Group—Select the network in the Primary AZ to be used for the live migration.
   - **Recovery Location**: Virtual Network/Port Group—Select the network in the Recovery AZ to be used for live the migration.

5. Check the **Live Migrate** box.

6. Click **Migrate**.

   Prism submits the task and displays the following message:

   ```
   Successfully submitted migrate operation.
   Task details
   ```

   Task details is a link to the Tasks page. Click the **Task details** link to monitor the migration task in the Tasks page.

   The host name of the VM in the List view changes to the host name to which you migrated the VM.

Creating a VM (ESXi)

In ESXi clusters, you can create a new virtual machine (VM) through the Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration on page 707 before proceeding.
- Register the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 709.

About this task

To create a VM, do the following:

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 86) and click the **Create VM** button.
   The **Create VM** wizard appears.
2. In the Configuration step, do the following in the indicated fields:
   a. **Name**: Enter a name for the VM.
   b. **Description** (optional): Enter a description for the VM.
   c. **Cluster**: Select the target cluster from the pull-down list on which you intend to create the VM.
   d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.
   e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
   f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
   g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

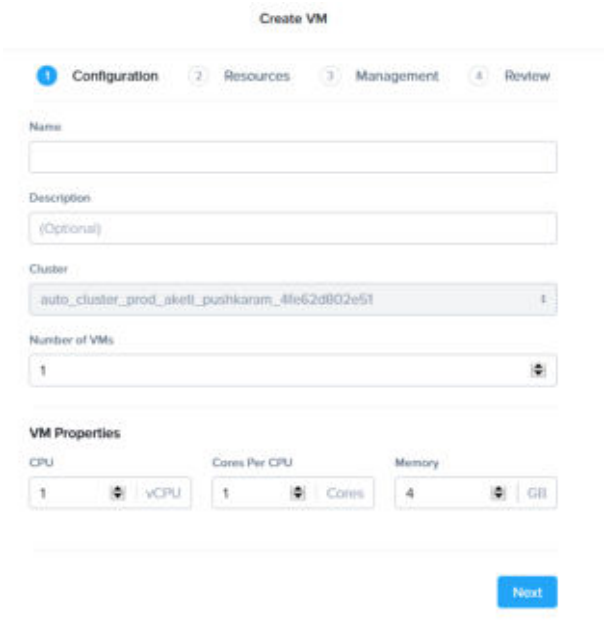

Figure 272: Create VM Window (Configuration)

3. In the **Resources** step, do the following.

   Disks: To attach a disk to the VM, click the **Attach Disk** button. The Add Disks dialog box appears. Do the following in the indicated fields:

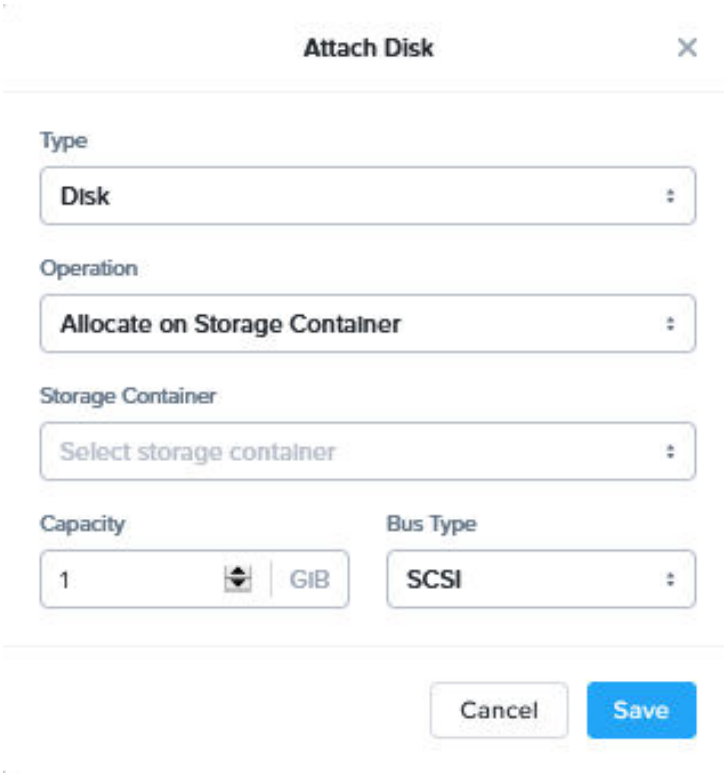

Figure 273: Add Disk Dialog Box

   a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the pull-down list.

   b. **Operation**: Specify the device contents from the pull-down list.

      - Select **Clone from NDSF file** to copy any file from the cluster that can be used as an image onto the disk.
      - [CD-ROM only] Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
      - [Disk only] Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
      - Select **Clone from Image** to copy an image that you have imported by using image service feature onto the disk.

   c. **Bus Type**: Select the bus type from the pull-down list. The choices are **IDE**, **SCSI**, **PCI**, or **SATA**.

   d. **Path**: Enter the path to the desired system image.

   e. **Clone from Image**: Select the image that you have created by using the image service feature.

This field appears only when **Clone from Image Service** is selected. It specifies the image to copy.

> Note: If the image you created does not appear in the list, see KB-4892.

f. **Storage Container**: Select the storage container to use from the pull-down list.

   This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.

g. **Capacity**: Enter the disk size in GiB.

h. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the Create VM dialog box.

i. Repeat this step to attach additional devices to the VM.

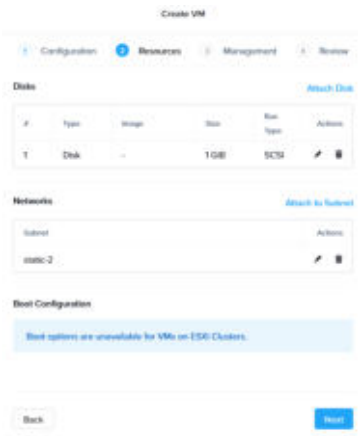

Figure 274: Create VM Window (Resources)

4. In the Resources step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The Attach to Subnet dialog box appears.

   Do the following in the indicated fields:

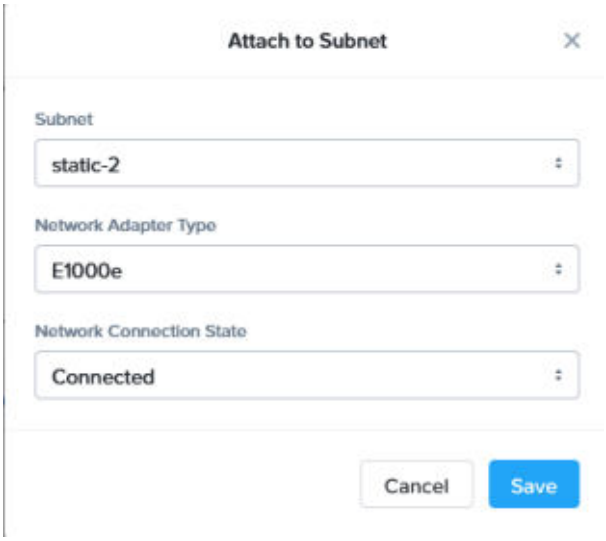

Figure 275: Attach to Subnet Dialog Box

   a. **Subnet**: Select the target virtual LAN from the pull-down list.

      The list includes all defined networks (see Configuring Network Connections on page 725).

   b. **Network Adapter Type**: Select the network adapter type from the pull-down list.

      See vCenter Server Integration on page 707 for the list of supported adapter types.

   c. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are `Connected` or `Disconnected`.

   d. When all the field entries are correct, click the **Add** button to create a network interface for the VM and return to the Create VM dialog box.

   e. Repeat this step to create more network interfaces for the VM.

5. In the **Management** step, do the following.

   a. **Categories**: Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

   b. **Guest OS**: Type and select the guest operating system.

      The guest operating system that you select affects the supported devices and number of virtual CPUs available for the virtual machine. The Create VM wizard does not install the

guest operating system. See the list of supported operating systems in vCenter Server Integration on page 707 topic.

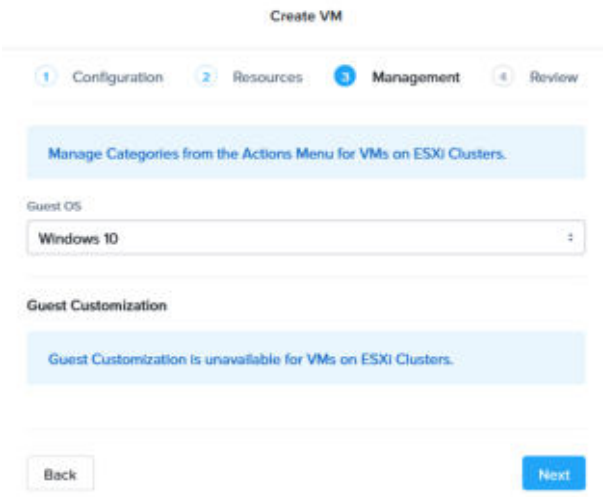

Figure 276: Create VM Window (Management)

6. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the Create VM dialog box.

   The new VM appears in the VMs entity page list.

Managing a VM (ESXi)

You can manage virtual machines (VMs) in an ESXi cluster through Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration on page 707 before proceeding.
- Ensure that you have registered the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 709.

About this task

After creating a VM (see Creating a VM (ESXi) on page 576), you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, power on (or off) the VM, enable flash mode for a VM, assign the VM to a protection policy, create VM recovery point, add the VM to a recovery plan, run a playbook, manage categories, install and manage Nutanix Guest Tools (NGT), manage the VM ownership, or configure QoS settings.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 86) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
- Go to the details page of a selected VM (see VM Details View on page 96) and select the desired action.

Note: The available actions appear in bold; other actions are unavailable. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Make the desired changes and then click the **Save** button in the **Review** step.

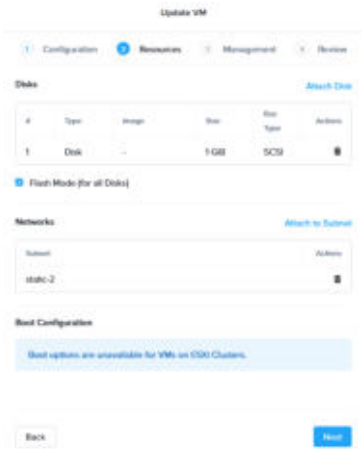

Figure 277: Update VM Window (Resources)

- **Disks:** You can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM (ESXi) on page 576 before you create a new disk for a VM. You can enable or disable the flash mode settings for the VM. To enable flash mode on the VM, click the **Enable Flash Mode** check box. After you enable this feature on the VM, the status is updated in the VM table view.
- **Networks:** You can attach new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (ESXi) on page 576 before you modify NIC network or create a new NIC for a VM.
- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

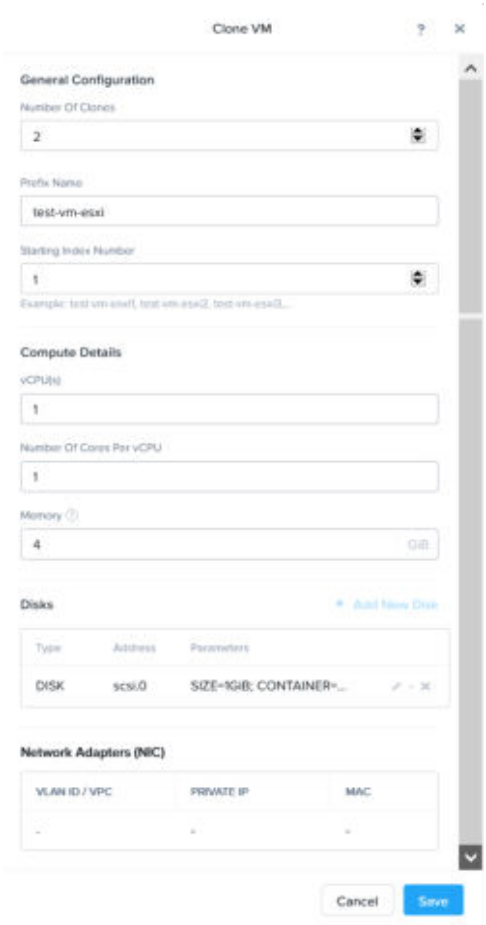

Figure 278: Clone VM Window

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM

table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

> Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

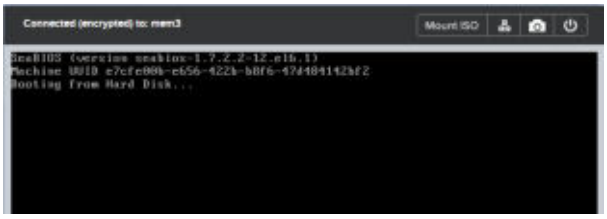

Figure 279: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To assign the VM to a protection policy, select **Protect**. This opens a page to specify the protection policy to which this VM should be assigned (see Policies Management). To remove the VM from a protection policy, select **Unprotect**.

  > Note: You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Leap and configuring the Availability Zone.

- To migrate the VM to another host, select **Migrate**.

  This displays the Migrate VM dialog box. Select the target host from the pull-down list (or select the **System will automatically select a host** option to let the system choose the host) and then click the **Migrate** button to start the migration.

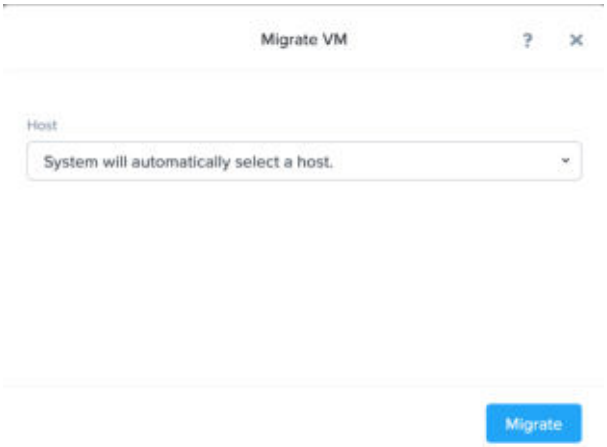

Figure 280: Migrate VM Window

> Note: Nutanix recommends to live migrate VMs when they are under light load. If they are migrated while heavily utilized, migration may fail because of limited bandwidth.

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the "Explicitly Adding a VM to a Recovery Plan" section in the Leap Administration Guide).
- To create VM recovery point, select **Create Recovery Point**.

  This displays the Create VM Recovery Point dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

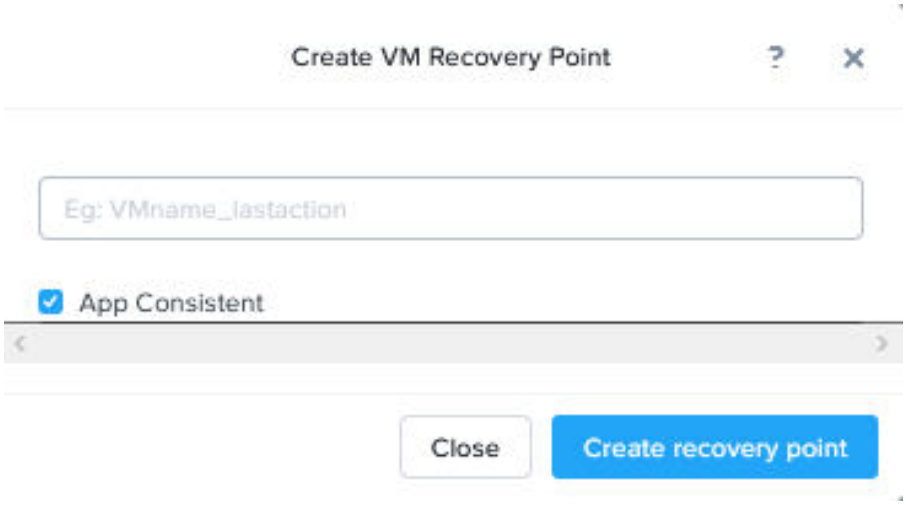

Figure 281: Create VM Recovery Point Window

- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 875).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 737).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 591).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 593).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > Note: If you clone a VM, by default NGT is not enabled on the cloned VM. You need to again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the *Prism Web Console Guide*.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

```
ncli> ngt mount vm-id=virtual_machine_id
```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

```
ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-
```

```
c1601e759987
```

> Note:
>
> - Self-service restore feature is not enabled by default on a VM. You need to manually enable the self-service restore feature.
> - If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.

- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 594).
- To establish VM host affinity, select **Configure VM Host Affinity**.

  A window appears with the available hosts. Select (click the icon for) one or more of the hosts and then click the **Save** button. This creates an affinity between the VM and the selected hosts. If possible, it is recommended that you create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure. See Affinity Policies Defined in Prism Central on page 599 for more information about the VM affinity policies.

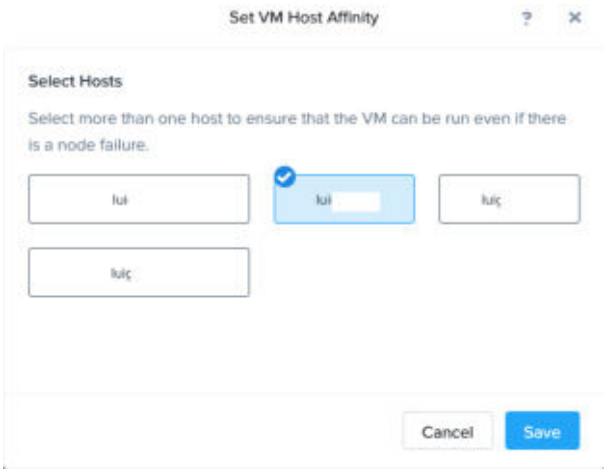

Figure 282: Set VM Host Affinity Window

- To add a VM to the catalog, select **Add to Catalog**. This displays the **Add VM to Catalog** page (see Adding a Catalog Item on page 658).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

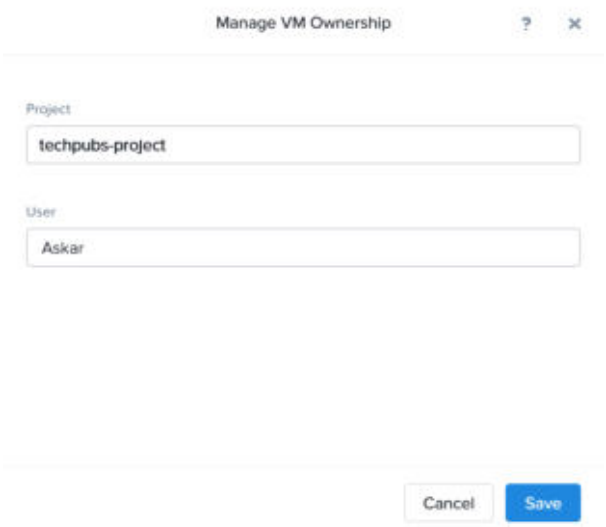

Figure 283: VM Ownership Window

- To configure quality of service (QoS) settings, select **Set QoS Attributes** (see Setting QoS for an Individual VM on page 597).

Exporting a VM as an OVA

Before you begin

If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the destination hypervisor.

About this task

As a Super Admin, you can export a VM as an OVA by selecting the **Export** as OVA action from the **Actions** dropdown list. You can export a VM as an OVA with a specified disk format for the disks configured on the VM.

> Note:
>
> - The Super Admin role was formerly called as User Admin.
> - If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the target destination hypervisor.

To export a VM as an OVA, do the following in the VMs dashboard:

Procedure

1. Select the VM in List tab.

   You can click the VM and open the VM details page.

2. In the **Actions** list, select **Export as OVA**.

   Alternatively, on the details page of the VM, select **More** and click **Export as OVA**.

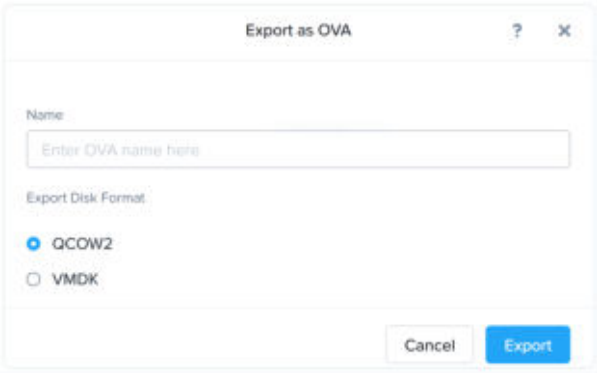

Figure 284: Export as OVA dialog box

3. In the Export as OVA window, do the following in the indicated fields:

   a. Enter a name for the OVA, in the **Name** field.

   b. Select the disk format that you want to format the disks in the VM to, in the **Export Disk Format** section.

      The default format is QCOW2. If you want to use the OVA to deploy a VM with disks formatted as VMDK, then select VMDK as the disk format.

4. Click **Export**.

   You can check the progress of the *Export as OVA* task in Tasks.

   The exported OVA is available on the OVA dashboard.

NGT Management in Prism Central

The NGT management in Prism Central feature lets you to select multiple VMs in the Prism Central from the VM Entity browser and then install, upgrade, or manage NGT applications on these VMs. You can use this feature for large-scale deployments. If you use this feature, you do not have to manually login to the VM and do the NGT installation as the installation of NGT is fully automated. Following are other important characteristics of this feature.

- Ability to enable VSS, SSR applications as part of the install workflow.
- Ability to select multiple VMs in the Prism Central from the VM Entity browser and upgrade NGT on these VMs.
- Ability to defer restarting of the VMs to a specified later time and date after NGT install or upgrade.
- Ability to define policies that allow you to defer restarting of the VMs to a specified later time and date.

For more information on the Nutanix Guest Tools, see the *Prism Web Console Guide*.

NGT Management in Prism Central Requirements

Following are the general requirements, specific requirements for Linux and Windows.

General Requirements

- Supported on the AHV and ESXi VMs.
- For AHV, both PC and PE should at least be on 5.9 version.
- For ESXi, both PC and PE should at least be on 5.11 version.
- VMs and virtual IP address of the cluster should be able to communicate with each other.
- The antivirus scanners must exclude the C:\Program Files\Nutanix directory for successful NGT installation.
- If you clone the VM from the VM where NGT was installed by using NGT Management in Prism Central feature, see Cloning from VMs With NGT on page 593 topic for how to enable NGT on the cloned VM.

The following TCP ports must be open between the Controller VMs and user VMs with NGT installed:

Port	Description	Inbound/ Outbound	Source	Destination
CVM Ports				
9440	To run Prism Element, Prism Central, and APIs.	Inbound and Outbound	Prism Central and Prism Element	Prism Central and Prism Element
2100	For Prism Central cluster creation.	Inbound	Prism Element	Prism Central
2073	For communication between NGT and other services in the cluster.	Inbound and Outbound	NGT leader node IP address and any other service in the cluster	NGT leader node IP address and any other service in the cluster
2074	Requests from user VMs come to this port in the CVM.	Inbound and Outbound	User VM IP address and CVM virtual IP address	User VM IP address and CVM virtual IP address
User VM Ports				
5986	Applies to Windows VMs only. The firewall rule must allow inbound connection over SSL on this port.	Inbound	CVM IP address	Target Windows user VM IP address

Port	Description	Inbound/ Outbound	Source	Destination
22 or any SSH port	Applies to Linux VMs only.	Inbound	NGT leader node IP address	Target Linux user VM IP address
23578	This is an internal port and required for inter-process communication (IPC) inside the user VM.	NA	NA	NA

Note: All ports are TCP.

For the complete list of required ports, see Port Reference.

Requirements for the Linux VMs

- The lvm2 package version should be 2.02.109 or later version.
- CVM should have SSH access to the VM.
- Ensure that you have at least 5 free SSH session configured in MaxSessions by navigating to `/etc/ssh/sshd_config`.
- The user should have no password sudo access on the VM.
- Ensure that you have disabled requiretty for a particular user.

```
sudo sed -i 's/Defaults    requiretty/ #Defaults    requiretty/g' /etc/sudoers
```

- For Linux operating systems, NGT is installed at `/usr/local` and NGT ISO is mounted on `/mnt/nutanix/ngt`. Hence, ensure that the user has write permissions to these locations.

Requirements for the Windows VMs

- PowerShell version on the Windows VMs should at least be 3.0.
- Local VM accounts with administrative privileges are required.
- PowerShell path must be available in the environment for NGT installation. Otherwise, NGT install fails.
- Windows Remote Manager Service (winrm) should be running.
- The winrm service with ssl configuration should be configured. Run the following commands from the Windows PowerShell as an administrator to configure winrm with SSL.

  For example, you can use the following commands to configure the winrm service.

```
> $certificate = New-SelfSignedCertificate -DnsName $env:computername -CertStoreLocation
 cert:\LocalMachine\My
```

```
> winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$env:computername`";CertificateThumbprint=`"$($certificate.ThumbPrint)`"}"
```

```
> cmd /c 'winrm set winrm/config/service/auth @{Basic="true"}'
```

```
> netsh advfirewall firewall add rule name=\"WinRM-HTTPS\" dir=in localport=5986
 protocol=TCP action=allow
```

Installing NGT on Multiple VMs

Perform the following procedure to install NGT on multiple VMs at the same time by using the NGT management in Prism Central feature.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 86).
   The NGT status column displays whether the NGT is installed on the VM, if NGT is installed whether it is the latest version, or whether an upgrade is available or not.
2. Select the VMs on which you want to install NGT and from the **Actions** drop-down menu, and then click **Install NGT**.
   The Install NGT windows displays the number of VMs on which you can install NGT. For example, if you have selected a VM on which NGT is already installed, this VM will be skipped during this operation.

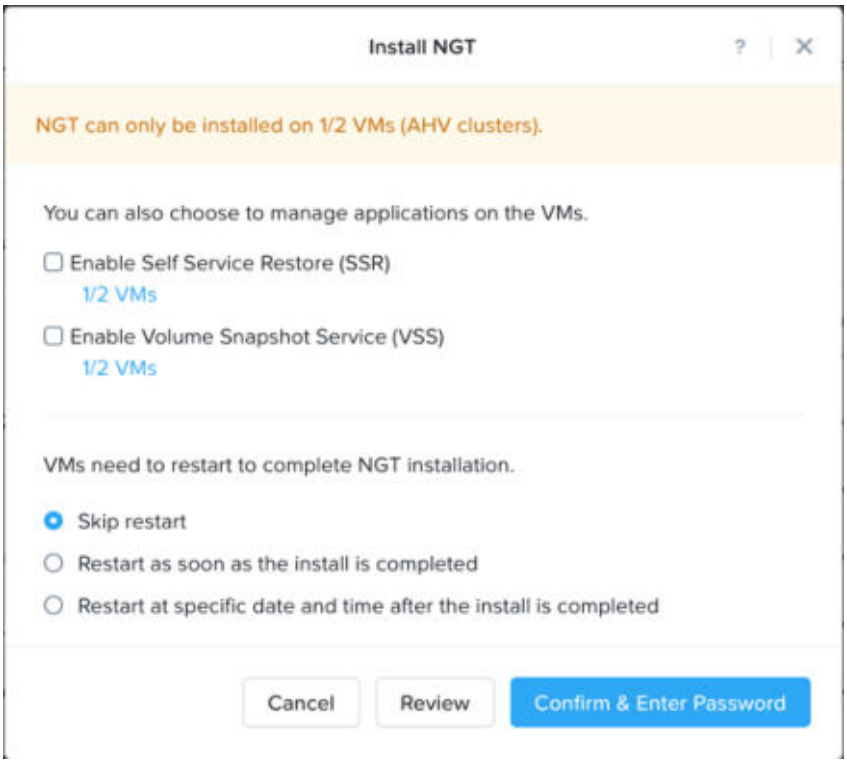

Figure 285: Installing NGT

3. To enable SSR and VSS applications, select the **Enable Self Service Restore (SSR)** and **Enable Volume Snapshot Service (VSS)** checkboxes.
   If you do not want to enable these features, you can do it later by following Managing NGT Applications on page 593.

   The VMs may have to be restarted to complete NGT installation because virtIO drivers are packaged with NGT ISO. When you install NGT on the VM, the virtIO drivers are also installed and hence VM restart is required. NGT operations does not require a restart. If you choose

to restart the VM at later point of time, all the NGT functionality will be available except the virtIO driver. VirtIO driver functionality comes into effect only after restart.

Select the restart schedule of the VMs.

- **Skip restart**: If you do not want the VMs to be restarted post NGT installation or if you want the category policy to be applied, select this option. For more information on the policies, see Creating a New Restart Policy on page 602.
- **Restart as soon as the install is completed**: The VMs will be restarted after the installation is completed.
- **Restart at specific date and time after the install is completed**: The VMs will be restarted at the specific date and time that you have configured in the **Date** and **Time** fields.

4. (Optional) Click **Review** to view your configurations.
5. Click **Confirm and Enter Password** to proceed with the installation.
   If all the VMs have same user names and password, you can enter it in the administrator **Username** and **Password** field. If the user names and passwords are different, you have to upload a JSON in the specified format.
6. Click **Done** to finish the installation process.
   Once the installation is completed, the **NGT Status** column displays NGT status as latest. This is an Async operation and takes some time for the status to be reflected in the Prism Central.

   - If you have not conform to the NGT prerequisites or if there is any issue with the installation process, the NGT installation fails and an alert message is displayed. Perform the installation process again to install NGT on the VMs.
   - You can also filter the list by clicking the **Filters** button. This displays a filter pane on the right of your screen. Click on the appropriate filter to list out the tasks that corresponds to that filter. The filters and their description are as follows.
     - Installed: Filters the VMs on which only NGT is installed.
     - SSR: Filters the VMs on which only SSR application is installed.
     - VSS: Filters the VMs on which only VSS application is installed.
     - VSS and SSR: Filters the VMs on which both VSS and SSR application is installed.
     - Not Installed: Filters the VMs on which NGT is not installed.

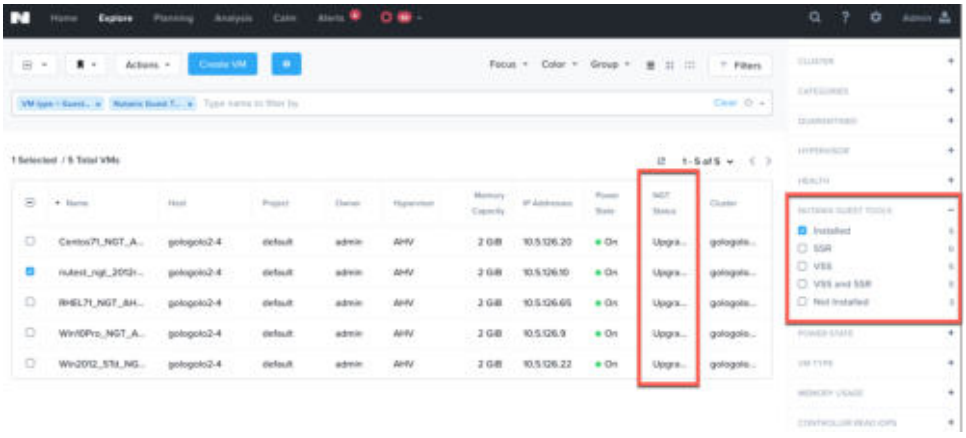

Figure 286: NGT Status and Filter Option

7. (Optional) If you do not want to install NGT and just want to mount NGT on the VM, click **Skip and Mount**.

   The VMs are registered with the NGT service. NGT is enabled and mounted on the selected virtual machines. A CD with volume label NUTANIX_TOOLS gets attached to the VMs. You can manually install the NGT by logging into the VM.

Cloning from VMs With NGT

For a VM that is cloned from an image or a VM that has NGT installed, perform the following procedure to enable NGT on these VMs. Since NGT is already present on the VM, you do not have to install NGT and only refresh of certifications is required.

Procedure

1. Select the cloned VMs and follow the NGT install process again to install NGT on the cloned VMs.
2. During the installation of the NGT on the cloned VMs, you can only mount NGT on the cloned VM by using **Skip and Mount** operation.

   For more information on selecting VMs and mounting NGT, see Installing NGT on Multiple VMs on page 591. NGT will be enabled and mounted on the cloned VMs.

   The NGT ISO will be mounted on the VM.
3. After the ISO is mounted, restart the VM.

   The NGT status is refreshed and updated after the VM starts.

Managing NGT Applications

If you have not enabled the VSS and SSR applications during the install workflow or if you want to change the application state, you can do it by using the following procedure. You can also disable these applications (if already enabled) by using the same workflow.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 86).

2. Select the VM or VMs on which you want to enable SSR and VSS applications by clicking **Manage NGT Applications**.

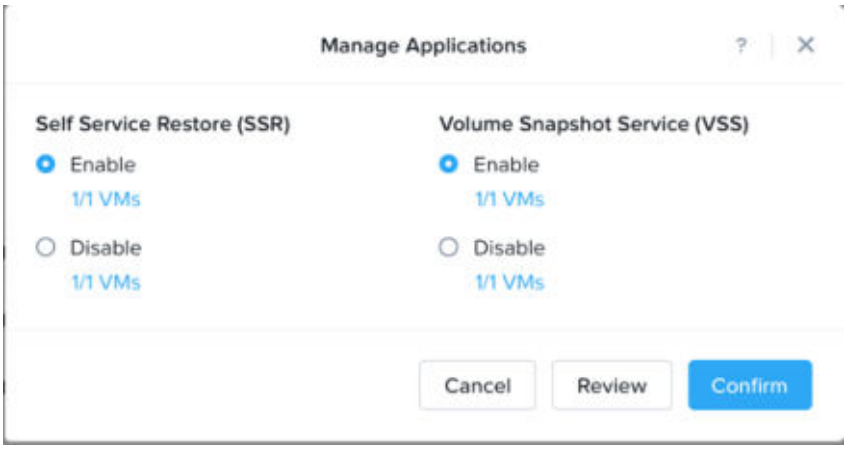

Figure 287: Manage NGT Applications

   The Manage Applications page displays whether you can enable these applications on the selected VMs or not.

3. Select **Enable or Disable** for the application, and click **Confirm**.

Upgrading NGT

Perform the following procedure to upgrade NGT on the VMs.

Before you begin

You can only upgrade from VMs that have NGT version is 1.2.0.3 and later (excluding 1.2.1, 1.2.2). If you have older versions of NGT, first upgrade NGT by using the procedure mentioned in the *Prism Web Console* guide.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 86).
   The NGT status column displays whether an upgrade is available for the selected VMs or not.

2. Select the VMs for which you want to upgrade NGT and from the **Actions** drop-down menu click **Upgrade NGT**.

   The Upgrade VM window displays the number of VMs on which you can upgrade NGT. For example, if you have selected a VM on which NGT is already installed and is latest, this VM is skipped during this operation.

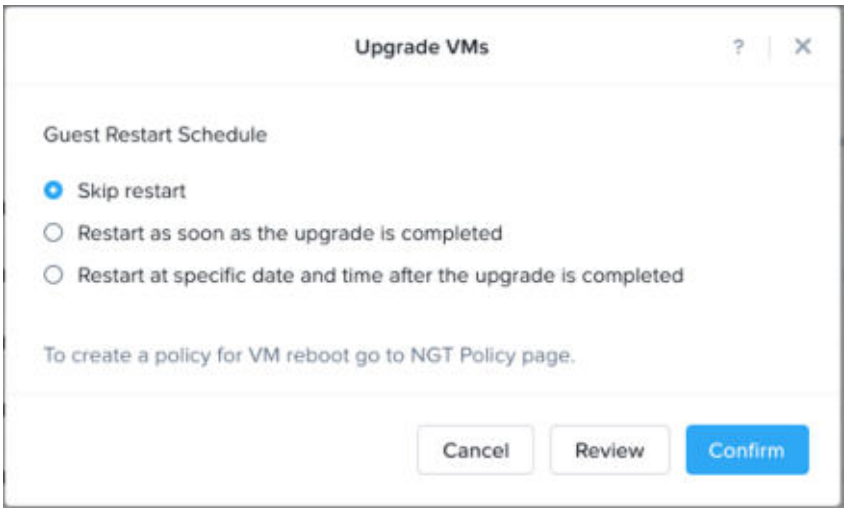

Figure 288: Upgrade NGT

3. Select the restart schedule of the VMs.

   - **Skip restart**: If you do not want the VMs to be restarted post NGT installation or if you want the category policy to be applied, select this option. For more information on the policies, see Creating a New Restart Policy on page 602.
   - **Restart as soon as the install is completed**: The VMs will be restarted after the upgrade is completed.
   - **Restart at specific date and time after the install is completed**: The VMs will be restarted at the specific date and time that you have configured in the **Date** and **Time** fields.

     Note: If you do not restart the VMs post upgrade and if the virtIO version has been upgraded, the upgraded functionality of the virtIO drivers will not be available.

4. (optional) Click **Review** to view your configurations.
5. Click **Confirm** to start the upgrade process.

   Note: If you have not conform to the prerequisites or if there is any issue with the upgrade process, the NGT upgrade fails and an alert message is displayed. Perform the upgrade process again to upgrade NGT on the VMs.

NGT Uninstallation

To uninstall NGT, you have to login into the VM and then perform the uninstallation process.

For more information on NGT uninstallation, see the *Disabling and Removing Nutanix Guest Tools* topic for the *Prism Web Console Guide*.

Storage Quality of Service (QoS)

Storage QoS provides administrators granular control to manage the performance of virtual machines and ensure that the system delivers consistent performance for all workloads. You

can use a controllable knob to limit the IOPS that the storage layer would serve for individual virtual machines. IOPS is the number of requests the storage layer can serve in a second. You can set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources.

Note:

- Storage QoS is applicable only if you have an AOS Pro or above license. See the *Nutanix Licensing Guide* for more information on AOS licenses.
- Setting storage QoS is not allowed while the AOS upgrade is in progress.

Terminology

Throttled IOPS

The maximum IOPS that the storage layer tries to admit over a sustained period. Applications that are running on the VM are not permitted to exceed this level. This prevents VMs from affecting a system beyond the set limits. However, if the system is in a state of flux, your application might not be throttled. To prevent application crashing due to misconfiguration, the throttled IOPS setting must not be less than one hundred. Otherwise, an error is raised.

Setting maximum IOPS on a VM may affect the cluster latency for that VM and may result in higher latency being displayed in the cluster latency charts. Performance of any other VM is not affected by this setting. You can view individual VM latency graph to ensure that the other VMs are not affected. You can also create a custom focus view to display the throttled IOPS for all the VMs.

Throughput

Throughput is the amount of data that an application sees from the storage layer underneath per second. The maximum throughput is calculated by multiplying the number of IOPS with the block size.

Calculation of the IOPS and Throughput

The default block size used for IOPS calculation is 32 KiB. If the block size is less than or equal to 32 KiB, then the system is rate limited by the number of IOPS and the bandwidth is calculated by multiplying the number of IOPS with the block size of the application. Otherwise, if the block size is greater than 32 KiB, the system is rate limited by the throughput and the number of IOPS is calculated by dividing the throughput by block size of the application.

For example, if you have set a throttle limit of 800 IOPS on a VM and the application is admitting IOPS with 8 KiB block size, then the application would see 800 IOPS and throughput of 6.25 MBps. However, if application is admitting IOPS with 64 KiB block size, then the application would see 400 IOPS and a throughput of 25 MBps. For more information about the block size, IOPS, and throughput relationship when throttle limit it set to 800 IOPS, see *Relationship between Block Size, IOPS, and Throughput* table.

Table 710: Relationship between Block Size, IOPS, and Throughput

Block size	IOPS	Throughput
8 KiB	800	6.25 MBps
16 KiB	800	12.50 MBps
32 KiB	800	25 MBps

Block size	IOPS	Throughput
64 KiB	400	25 MBps
128 KiB	200	25 MBps

Limitations

Storage QoS has the following limitations:

- Storage QoS is not supported for an in-place restore, out-of-place restore, and snapshot operation on VM.
- While creating a clone of a VM, any QoS attributes throttle limit set on the original VM might not be applied to the new cloned VM. QoS throttle limit has to be set on the new cloned VM separately.
- For linked clones, when you are sharing the vDisks across multiple VMs, shared disks are not a part of QoS policy and shared disks would not be throttled. Each VM might exceed the maximum limit depending on the shared vDisk usage. For example, during bootstorm in a VDI setup.
- If a VM has volume groups attached, then QoS is not applicable.
- Storage QoS is not supported for Metro Availability, synchronous replication, and Nutanix Files virtual machines.

Setting QoS for an Individual VM

About this task

Storage QoS allows you to set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources. Perform the following procedure to set the QoS for an individual VM.

Figure 289: Setting QoS for an Individual VM

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275779955001 on Brightcove.

Procedure

1. Go to the **List** tab of the VM dashboard (see VMs Summary View on page 86) and select the target VM.
2. Do one of the following:
   - » From the **Actions** drop-down list, select **Set QoS Attributes**.
   - » Click the VM and then select **Set QoS Attributes** from the **More** drop-down list.
3. Under **Input Type**, select one of the following:
   - » IOPS (For more information about IOPS, see *Throttled IOPS* in *Storage Quality of Service (QoS)*.)
   - » Throughput (For more information about throughput, see *Throughput* in *Storage Quality of Service (QoS)*.)

4. If you selected **IOPS**, enter a value for the throttled IOPS in the **Throttled IOPS**.
5. If you selected **Throughput**, enter a value for the throttled throughput in the **Throttled Throughput**.

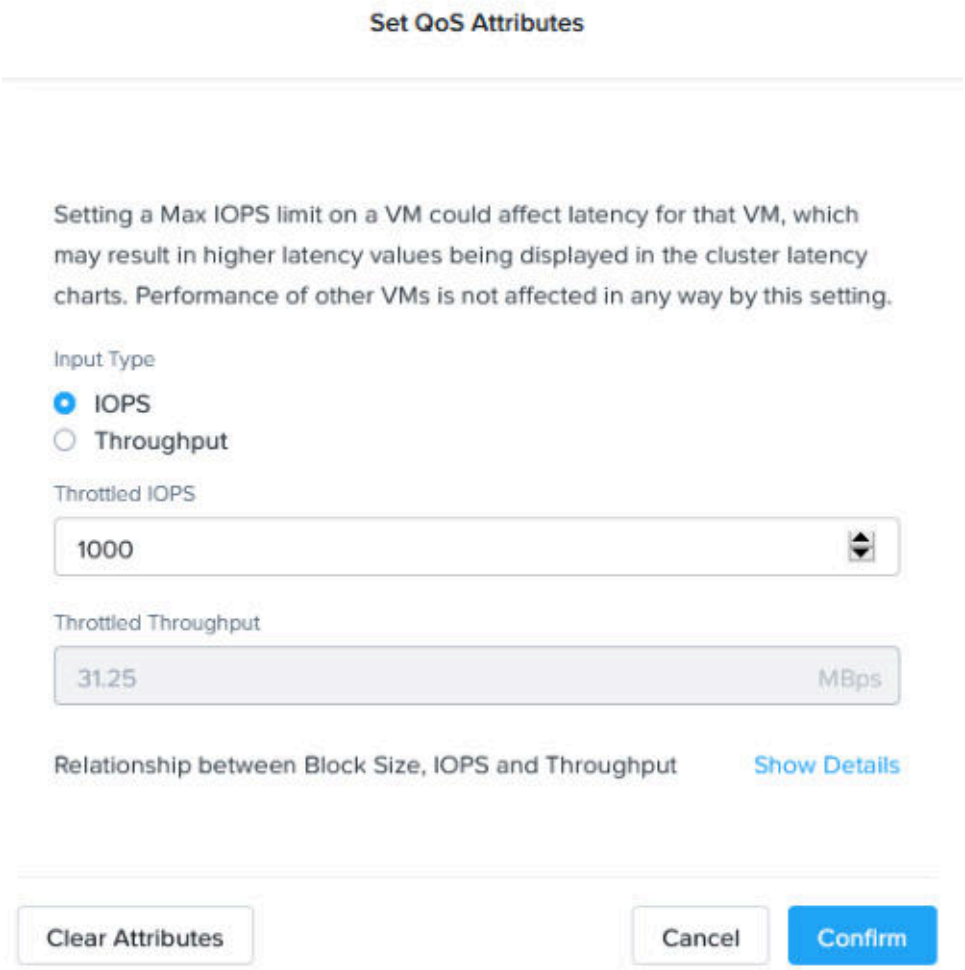

Figure 290: Setting QoS Attributes

6. Click **Show Details** to populate a table that describes the relationship between block size, IOPS, and throughput.
7. Click **Clear Attributes** to delete all the values.
8. To save details, click **Confirm**.
   This step stores the number of IOPS or throughput that would be served by a storage layer. To edit the storage QoS setting, follow the previous process and enter a new value for IOPS or throughput.

VM Policy Management

In Prism Central, you can create and manage the following types of policies for the VMs:

- Affinity Policies (see Affinity Policies Defined in Prism Central on page 599)
- NGT Policies (see NGT Policies on page 602

Affinity Policies Defined in Prism Central

In Prism Central, you can define category-based VM-Host affinity policies, where a set of VMs can be affined to run only on a particular set of hosts. Category-based affinity policy enables you to easily manage affinities for a large number of VMs. In case of any changes to the affined hosts, you only need to update the category of the host, and it updates the affinity policy for all the affected VMs.

This policy checks and enforces where a VM can be hosted when you start or migrate the VM. If there are no resources available on any of the affined hosts, the VM does not get started.

> Note:
>
> If you create a VM-Host affinity policy for a VM that is configured for asynchronous replication, you must create similar categories and corresponding policies on the remote site as well. If you define similar categories and policies on the remote site, affinity policies will be applied when the VMs are migrated to the remote site.

Limitations of Affinity Policies

Affinity policies created in Prism Central have the following limitations:

- Only a super admin can create, modify, or delete affinity policies.
- The minimum supported versions for VM-Host affinity policies are version 6.1 for Prism Element and version 2022.1 for Prism Central.
- Host category attach or detach takes around five minutes to get reflected in the applicable affinity policies.
- You cannot apply VM-Host affinity policy on a VM that is enabled for synchronous replication. Also, you cannot enable synchronous replication on a VM that is associated with a VM-Host affinity policy.

Affinity Policy Configuration Workflow

About this task

To set up an affinity policy, do the following:

Procedure

1. Create categories for the following entities:
   a. VMs
   b. Hosts
   For information about creating a category, see Creating a Category.
2. Apply the VM categories to the VMs and host categories to the hosts.
   For information about associating categories with VMs, see Associating VMs with Categories. For information about associating categories with hosts, see Associating hosts with Categories.
3. Create the affinity policy. See Creating an Affinity Policy.

Associating VMs with Categories

About this task

To associate categories with VMs, do the following:

Procedure

1. In Prism Central, in the **Entities** menu, go to **Compute & Storage** > **VMs**.
2. Select the VMs that you want to associate with a category.
3. In the **Actions** menu that is displayed, click **Manage Categories**.
4. In the **Search for a category** field, type the name of the category or value that you want to add, and then select the category and value from the list that is displayed.
5. Click the add button to add as many search fields and repeat this step for each category that you want to add. To remove a category, click the remove button beside the field.
6. Click **Save**.

Associating Hosts with Categories

About this task

To associate categories with hosts, do the following:

Procedure

1. In Prism Central, in the Entities menu, go to **Hardware** > **Hosts**.
2. Select the hosts that you want to associate with a category.
3. In the **Actions** menu that is displayed, click **Manage Categories**.
4. In the **Search for a category** field, type the name of the category or value that you want to add, and then select the category and value from the list that is displayed.
5. Click the add button to add as many search fields and repeat this step for each category that you want to add. To remove a category, click the remove button beside the field.
6. Click **Save**.

Creating an Affinity Policy

About this task

To create an affinity policy, do the following:

Before you begin

Complete the following tasks:

- Configure the categories that you need to associate with the VMs. Associate this category with all the relevant VMs. Alternatively, you can associate the category with the VMs after creation of the affinity policy.

- Configure the categories that you need to associate with the hosts. Associate this category with all the relevant hosts. Alternatively, you can associate the category with the hosts after creation of the affinity policy.

  Note: If you have configured any legacy affinity policy (non-category-based affinity policy) associated with the VMs, you must first remove those legacy affinity policies to allow the creation of category-based affinity policies associated with the same VMs.

Procedure

1. In the VMs Summary View, go to **Policies** > **Affinity Policies**, and then click **Create**.
2. In the **Create Affinity Policy** page, enter a name and description (optional) for the policy.

Figure 291: Create Affinity Policy

3. Click inside the **VM Categories** search field and select the category you want to associate with the VMs.
4. Click inside the **Host Categories** search field and select the category you associated with the hosts.
5. Click **Create**.

Updating an Affinity Policy

You can update an affinity policy. Updates to an affinity policy can result in a policy violation. Prism Central attempts to correct the violation by executing a series of actions.

About this task

To update an affinity policy, do the following:

Procedure

1. In the entity menu, go to **Policies** > **Affinity Policies**, and then click the affinity policy that you want to update.
2. From the **Actions** menu that is displayed, click **Update**.

3. The **Update Affinity Policy** page that is displayed includes the same options and settings as the **Create Affinity Policy** page.

Figure 292: Update Affinity Policy

4. Update the settings that you want, and then click **Save**.

NGT Policies

If a restart is required post installation or upgrade, you can define a policy on when the restart should happen. You can use these policies when you have different set of VMs which you would like to be restarted at different times, for example if they are in different time zones or if their down times are expected to be different. Policy workflow are defined on the categories. You have to attach the category to the VM and then you can create a policy on the VM.

Creating a New Restart Policy

Perform the following procedure to create a new restart policy.

Procedure

1. Go to the **VMs** > **Policies** dashboard (see NGT Policies View on page 187) and click the **New VM Reboot Policy** button.

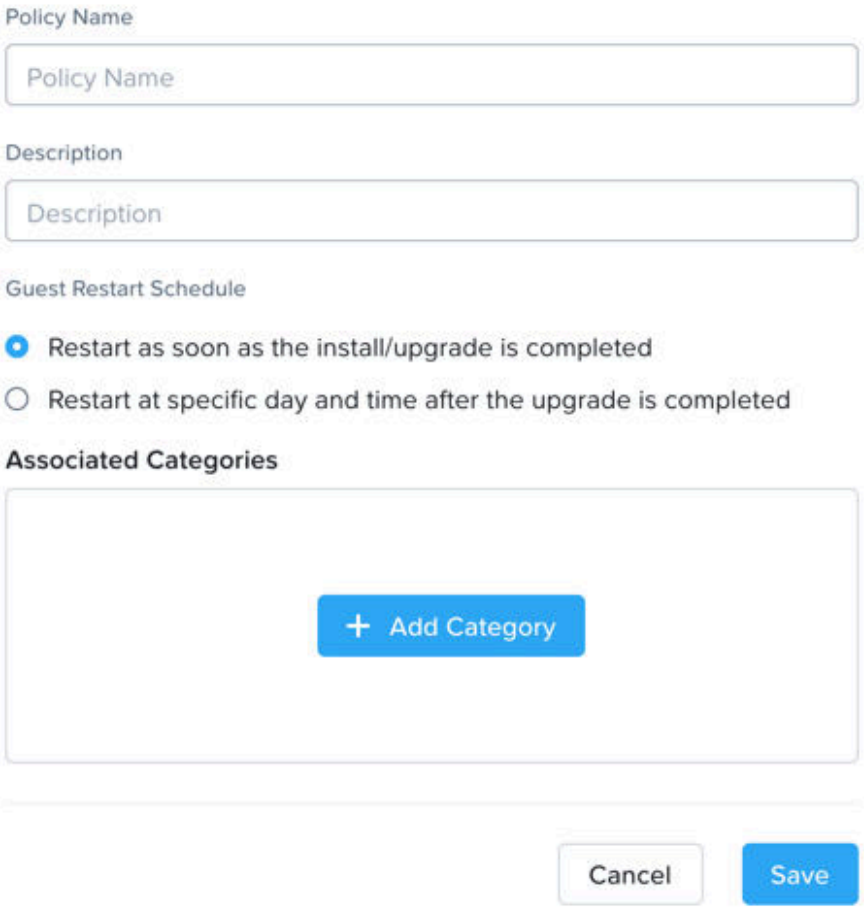

Figure 293: NGT Policy

2. In the **Policy Name** text box, type a policy name.
3. In the **Policy Description** text box, type a description of the policy.
4. Configure the guest restart schedule.

   - **Restart as soon as the install/upgrade is completed**: Select this option if you want to restart the VMs of the selected category right after the install process is completed.
   - **Restart at specific day and time after the upgrade is completed**: Select this option and choose the date and time on which you want to the restart to happen.

     Note: The VMs will be restarted depending on the Prism Element timezone settings.

5. Click **Add Category** and select the category that you want to apply this policy.

   Note: You can select multiple categories and the application of policy is dependent on the following rules.

   - 1. If you choose to restart the VM by using the **Restart as soon as the install/upgrade is completed** option either by using NGT install/upgrade workflow or by using this procedure, it will always take precedence over other workflows. Also, if you select either of these options, the policy will be created as a System policy and you will not be able to modify this policy. Once these policies are applied, it is removed from the user interface.
     2. If you schedule a restart during the NGT install or upgrade workflow, it will take precedence over other policies, for example, the policy that you have created by using this procedure.
     3. If there are multiple VM restart policies that are created at the same time by using this procedure, an alert message will be displayed.

6. Click **Save**.
   The policy is created and saved in the NGT Policies pane.

Modifying a Restart Policy

You can update and delete the restart policy that you have created by following this procedure.

Procedure

1. Go to the NGT policies dashboard (see NGT Policies View on page 187).
2. Check the box for the policy to update and then select **Update** from the **Actions** drop-down menu.
   a. Make the necessary changes to the policy and click **Save**.
3. To delete a policy, select the policy and from the **Actions** drop-down menu click **Delete**.
   a. Click Delete in the confirmation dialog box to delete the policy.

VM Template Management

In Prism Central, you can create VM templates to manage the golden image of a VM. A VM template can be considered as a master copy of a virtual machine. It captures the virtual machine configuration and the contents of the VM including the guest operating system and the applications installed on the VM. You can use this template to deploy multiple VMs across clusters.

Note: You can create or manage a VM template only as an admin user.

Limitations of VM Template Feature

The current implementation of VM template feature has the following limitations:

- <Version related limitations>

- Templates do not copy the following attributes from the source VMs:
  - Host affinity attributes
  - HA priority attributes
  - NGT installation

  The above-listed attributes must be reconfigured at the deployed VMs.

Creating a VM Template

About this task

To create a VM template, do the following:

Procedure

1. Select the VM in List tab under VM dashboard.

   Note: Before selecting a VM for creating a template, ensure that the VM is powered off.

2. In the **Actions** list, select **Create VM Template**.

   Alternatively, on the details page of the VM, select **More** and click **Create VM Template**.

3. This displays the Create template from VM dialog box. In the **Create Template from VM** window, do the following in the indicated fields:

   a. In the **Name** field, enter a name for the template.

   b. In the **Description** field, enter a description for the template. Description is an optional field.

   c. In the **Guest Customization** field, select your options for guest operating system (OS) customization.

   You can choose to customize the OS of the VMs that will be deployed using this template. In the **Script Type** field, select Sysprep (Windows) for customizing the Windows OS, and Cloud-init (Linux) for customizing the Linux OS. For configuring each of these script types, you can choose to either upload a custom script or opt for a guided setup in the **Configuration Method** field. The guided setup lets you enter the authentication information such as username, password, locale, hostname, domain join, and license key.

   You can also specify if you want to allow the template users to override the guest customization settings of the template while deploying the VM. If you select a script type and then allow the users to override the settings, the users can change the settings only for the configuration method. For example, the users can change the authentication information at the time of deploying a VM from template, or they can change from a guided setup to a custom script.

   If you select no customization at the time of creating the template but allow the users to override the guest customization settings, it gives the maximum customization control to the users. They can customize the script type and the configuration method.

   Note: If you opt for a guest customization script, ensure that the script is in a valid format. The guest customization scripts are not validated, hence even though the VM deployment

might succeed, the guest customization script may not work as expected and it will be apparent only by observing the VM after it gets deployed.

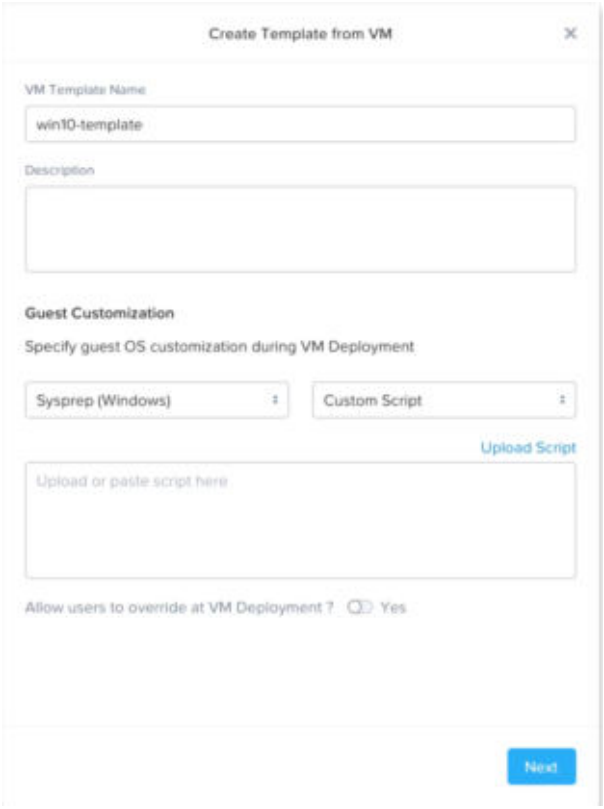

Figure 294: Create VM Template

4. Click **Next**.
   On the next page, you can review the configuration details, resource details, network details, and management details. Except the name, description and guest customization options that you chose on the previous page, you cannot modify any other setting while creating the template.

5. Click **Save** to save the inputs and create a template.
   The new template appears in the templates entity page list.

Deploying VM from a Template

About this task

After creating a VM template, you can use the template to deploy any number of VMs across clusters.

To deploy a VM from a template, do the following:

Procedure

1. You can use any of the following methods to deploy VMs:

   - Select the target template in the templates dashboard (see VM Template Summary View on page 109) and click the **Deploy VMs** button. By default, the active version of the template is used for deployment.
   - Go to the details page of a selected template (see VM Template Details View on page 111) and click the **Deploy VMs** button. By default, the active version of the template is used for deployment.
   - Go to the details page of a selected template (see VM Template Details View on page 111) and then go to **Versions** view. Select the desired version that you want to use for VM deployment, and click the **Deploy VMs** button. Here, you can choose any active or non-active version for the deployment.

2. The Deploy VM from Template page appears. By default, you see the **Quick Deploy** page. You can click**Advanced Deploy** button to access the Advanced Deploy page. In **Advanced Deploy** option, you can view and modify some VM properties and network settings.

   - **Quick Deploy**: To deploy VMs using the quick deploy method, provide inputs in the indicated fields:

     - **Name**: Enter a name for the VM.
     - **Cluster**: Select the cluster where you want to deploy the VM.
     - **Number of VMs**: Enter the number of VMs that you want to deploy.
     - **Starting Index Number**: Enter the starting index number for the VMs when you are deploying multiple VMs simultaneously. These index numbers are used in the VM names. For example, if you are deploying two VMs and specify the starting index number as 5, the VMs are named as *vm_name*-5 and *vm_name*-6.
     - **Guest Customization**: The template can have any one of the following options for guest OS customization: **No Customization**, **Sysprep (Windows)**, or **Cloud-init (Linux)**. For **Sysprep (Windows)**, or **Cloud-init (Linux)**, you can choose to either upload a custom script or opt for a guided setup. The fields are enabled for modification only if the template allows you to override its guest customization settings while deploying

the VM. If you are not allowed to override the guest customization settings, the settings that are already provided in the template are used for VM deployment.

- Click **Next** to verify the configuration details of the VMs to be deployed.
- Click **Deploy** to deploy the VMs.

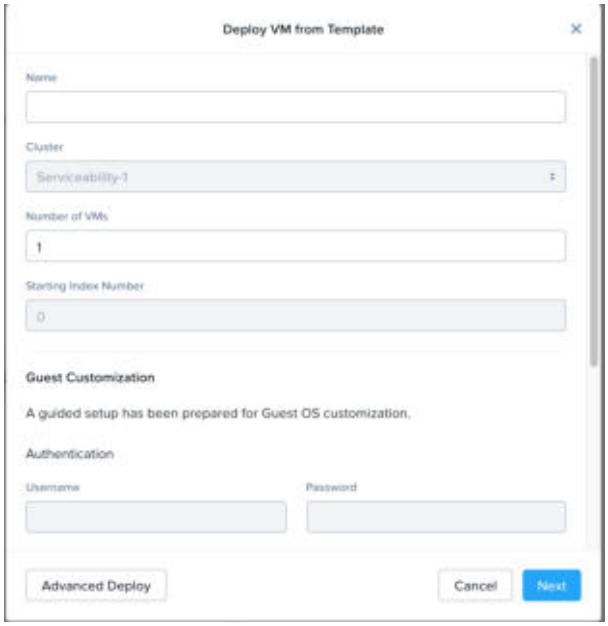

Figure 295: Deploy VM from Template using Quick Deploy Method

- **Advanced Deploy**: To deploy VMs using the advanced deploy method, provide inputs in the indicated fields and review the information displayed under various tabs:
  - **Configuration**: Provide inputs for name and description (optional) of the VM, cluster where you want to deploy the VM, Number of VMs to be deployed, and starting index

number (only if deploying multiple VMs). In this tab, you can also view and modify the VM properties such as CPU, core per CPU, and memory.

- **Resources**: Review the configuration settings for the VM resources such as disks, networks, and boot configuration. Here, you can modify the network settings but cannot modify any other settings.
- **Management**: If the template allows you to modify the settings for guest OS customization, provide inputs for the same.
- **Review**: Review the information displayed on this tab. Click **Deploy** to deploy the VMs.

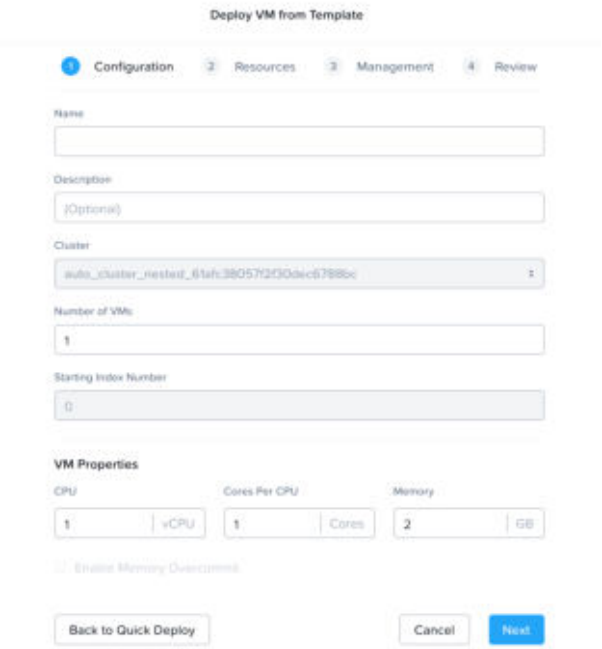

Figure 296: Deploy VM from Template using Advanced Deploy Method

Managing a VM Template

About this task

After creating a template (see Creating a VM Template on page 605), you can update the guest OS of the source VM of the template, complete guest OS update, cancel guest OS update, update the configuration of the template to create a version, or delete the template.

You can perform these tasks by using any of the following methods:

- Select the target template in the templates dashboard (see VM Template Summary View on page 109) and choose the required action from the **Actions** menu.
- Go to the details page of a selected template (see VM Template Details View on page 111) and select the desired action.

> Note: The available actions appear in bold; other actions are not available. The available actions depend on the current state of the template and your permissions.

Procedure

- To update the guest OS of the template, select **Update Guest OS**. Click **Proceed**.

  This step deploys a temporary VM and gives you access to that VM. You can also access the new VM from VM dashboard. You must start the VM, log on to the VM, and update the guest OS of that VM.

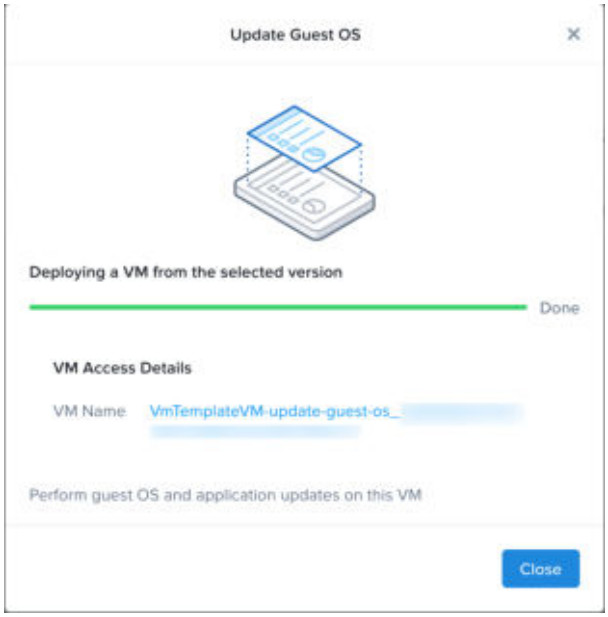

Figure 297: Update Guest OS

Once you have updated the guest OS of the temporary VM, you must complete the guest OS update from Prism Central UI. Alternatively, you can choose to cancel the guest OS update, if you choose to discontinue the guest OS update process.

The temporary VM automatically gets deleted after completion or cancellation of the guest OS upgrade process.

- To complete the process of guest OS update that you had earlier initiated, select **Complete Guest OS Update**. You must select this option only after successful update of the guest OS of the temporary VM.
- To cancel the process of guest OS update that you had earlier initiated, select **Cancel Guest OS Update**.
- To modify the template configuration, select **Update Configuration**. You cannot modify the configuration of the initial version of the template, but you can update the configuration

settings to create a new version of the template. In Select a version to Upgrade window, select the version of the template that you want to modify and create a version on top of it.

The Update Template Configuration window displays the following sections:

- **Configuration**: You can view the name of the base version that you want to update, change notes for that version, cluster name, VM properties (CPU, cores per CPU), and memory). In this section, you can modify only VM properties.
- **Resources**: You can view the information about disks, networks, and boot configuration. In this section, you can modify only network resources.
- **Management**: In this section, you can modify the guest customization settings.
- **Review**: In this section, you can review and modify the configuration settings that you are allowed to modify. You must provide a name and change notes for the new version. You can also choose to set this new version as active version. An active version is the version of the template that by default gets deployed when you click the **Deploy VMs** button.

Click **Save** to save the settings and create a new version of the template.

Figure 298: Update Template Configuration

- To delete a template, select **Delete Template**. A window prompt appears; click the **OK** button to delete the template.

Storage Management

Storage in a Nutanix cluster is organized into several components that allow you to manage capacity and performance. Nutanix clusters include storage pool, storage container, volume group, and virtual disk components. Prism Element allows you to manage all aspects of storage for a cluster, but you can also manage some storage components directly from Prism Central.

- The storage containers dashboard displays information about storage containers across the registered clusters (see Storage Containers Summary View on page 122).

- The volume groups dashboard displays information about volume groups across the registered clusters (see Volume Groups Summary View on page 139). Using Prism Central you can do the following:

  Note: Some actions are unavailable for clusters running a pre-6.0 AOS version.

  - Create a storage container (see Creating a Storage Container (AHV) on page 622, Creating a Storage Container (ESXi) on page 626, or Creating a Storage Container (Hyper-V) on page 629)
  - Update or modify a storage container (see Modifying a Storage Container on page 633)
  - Delete a storage container (see Deleting a Storage Container on page 637)
  - Create a volume group (see Creating a Volume Group on page 637)
  - Update or delete a volume group (see Modifying a Volume Group on page 641)
  - Add, update, or remove connections to external clients for a volume group (see Managing Volume Group Connections on page 642)
  - Add, update, or delete virtual disks for a volume group (see Managing Volume Group Virtual Disks on page 646)
- The storage policy dashboard displays information about the storage policies that are configured to manage the govern the storage behaviour of VMs or entities (see Storage Policies Summary View on page 133 and Storage Policy Management on page 647).
- For information about storage management tasks that are performed through Prism Element, see the *Storage Management* chapter in the *Prism Web Console Guide*.

Storage Components

Storage in a Nutanix cluster is organized into the following components.

Storage Tiers

Each type of storage hardware (SSD-PCIe (NVMe), SSD (SATA SSD), and HDD) is placed in a storage tier. You can determine the tier breakdown for disks in a storage pool through the web console (see "Storage Table View" section in the Prism Web Console guide.

Storage Pools

Storage pools are groups of physical disks from one or more tiers. Storage pools provide physical separation because a storage device can only be assigned to a single storage pool at a time. Nutanix recommends creating a single storage pool for each cluster. This configuration allows the cluster to dynamically optimize capacity and performance. Isolating disks into separate storage pools provides physical separation, but can create an imbalance of these resources if the disks are not actively used. When you expand your cluster by adding new nodes, the new disks can also be added to the existing storage pool. This scale-out architecture allows you to build a cluster that grows with your needs.

When you create a cluster, a default predefined storage pool is available. This pool includes the total capacity of all the disks on all the hosts in the cluster.

Storage Containers

A storage container is a subset of available storage within a storage pool (see "Creating a Storage Container" section). Storage containers are created within a storage pool to hold virtual disks (vDisks) used by virtual machines. By default, storage is thinly provisioned, which

means that the physical storage is allocated to the storage container as needed when data is written, rather than allocating the predefined capacity when the storage container is created. Storage efficiency features such as compression, deduplication, and erasure coding are enabled at the container level.

Volume Groups

A volume group is a collection of logically related virtual disks (or volumes). A volume group is attached to VM either directly or using iSCSI. You can add vDisks to a volume group, attach them to one or more consumers, include them in disaster recovery policies, and perform other management tasks. You can also detach a volume group from one VM and attach it to another, possibly at a remote location to which the volume group is replicated.

You manage a volume group as a single unit. When a volume group is attached to a VM, the VM can access all of the vDisks in the volume group. You can add, remove, and resize the vDisks in a volume group at any time.

Each volume group is identified by a UUID, a name, and an iSCSI target name. Each disk in the volume group also has a UUID and a SCSI index that specifies ordering within the volume group. A volume group can be configured for either exclusive or shared access.

You can backup, protect, restore, and migrate volume groups. You can include volume groups in protection domains configured for asynchronous data replication (Async DR), either exclusively or with VMs. However, volume groups cannot be included in a protection domain configured for metro availability, in a protected vStore, or in a consistency group for which application consistent snapshots are enabled.

vDisks

A vDisk is created within a storage container or volume group to provide storage to the virtual machines. A vDisk shows up as a SCSI device when it is mapped to a VM.

Containers for VMware and Hyper-V (Datastores/SMB Shares)

In vSphere, a datastore is a logical container for files necessary for VM operations. Nutanix provides the choice by supporting both iSCSI and NFS protocols when mounting a storage volume as a datastore within vSphere. NFS has many performance and scalability advantages over iSCSI, and it is the recommended datastore type.

In Hyper-V environments, storage containers are mounted as an SMB share.

> Note: Using a Nutanix storage container as a general-purpose NFS or SMB share is not recommended. For NFS and SMB file service, use Nutanix Files.

NFS Datastores. The Distributed Storage Fabric (DSF) reduces unnecessary network chatter by localizing the data path of guest VM traffic to its host. This boosts performance by eliminating unnecessary hops between remote storage devices that is common with the pairing of iSCSI and VMFS. To enable vMotion and related vSphere features (when using ESX as the hypervisor), each host in the cluster must mount an NFS volume using the same datastore name. The Nutanix web console and nCLI both have a function to create an NFS datastore on multiple hosts in a Nutanix cluster.

To correctly map the local ESX datastore to the Nutanix container:

- Map the NFS share with 192.168.5.2 (internal IP address) and not the Controller VM IP address or cluster virtual IP address.
- The name of the datastore should be same as the name of the container.

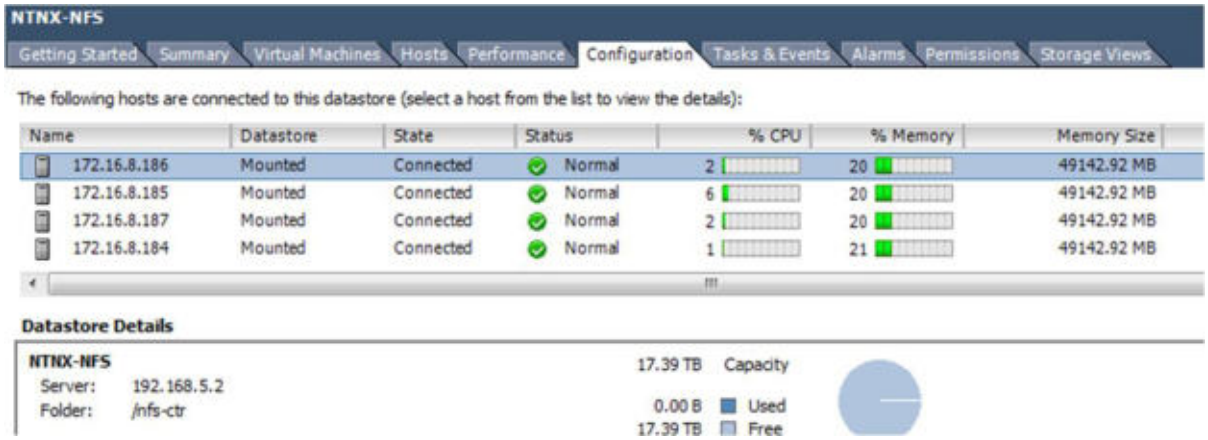

Figure 299: vSphere Configuration of NFS Datastore

SMB Library Share. The Nutanix SMB share implementation is the Hyper-V equivalent of an NFS Datastore with feature and performance parity with a vSphere configuration. The registration of a Nutanix storage container as an SMB Library share can be accomplished through a single powershell script, or through the Virtual Machine Manager GUI.

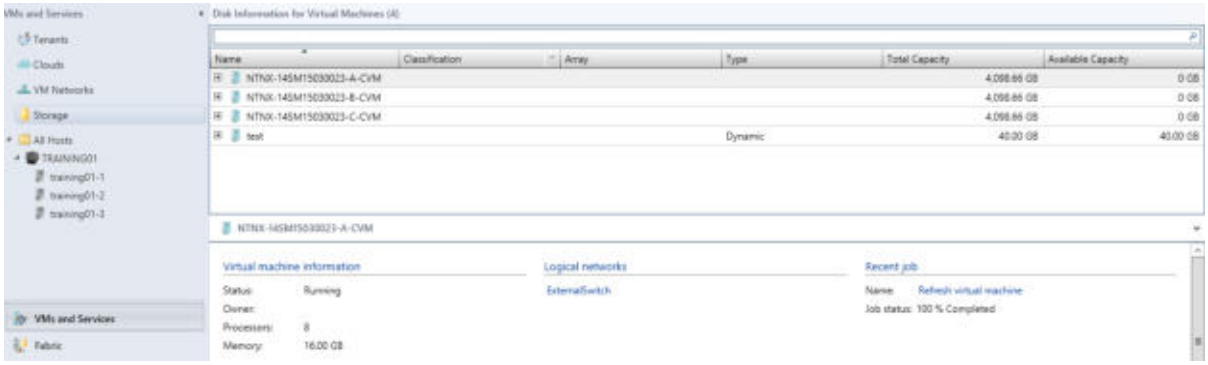

Figure 300: Hyper-V Configuration of an SMB Share

Compression

You can enable compression on a storage container. Compression can save physical storage space and improve I/O bandwidth and memory usage—which may have a positive impact on overall system performance.

> Note: If the metadata usage is high, compression is automatically disabled. If compression is automatically disabled, an alert is generated.

The following types of compression are available.

Post-process compression

Data is compressed after it is written. The delay time between write and compression is configurable, and Nutanix recommends a delay of 60 minutes.

Inline compression

Data is compressed as it is written. When you create a new storage container, inline compression is enabled by default for all license tiers. It is set to a delay of 0, compressing data immediately as it is written.

Compression Ratios

You can view compression ratios and usage savings in the Prism Web Console.

- Cluster

  In the Storage dashboard, under **Capacity Optimization**, click the **After** bar, and hover your mouse over **Compression**.

- Storage container

  In the Storage dashboard **Table** view, on the **Storage Container** tab, click the storage container for which you want to view the compression ratio. You can see the compression ratio for the selected storage container under **Storage Container Details**.

Deduplication

Deduplication reduces space usage by consolidating duplicate data blocks on Nutanix storage. You can enable either cache deduplication only or both cache and capacity deduplication on a storage container.

Cache Deduplication

You can enable cache deduplication of read cache to optimize performance. Cache deduplication means deduplication performed on the data in memory.

> Note: Cache deduplication is not enabled by default.

Capacity Deduplication

You can enable capacity deduplication of persistent data to reduce storage usage. Capcity deduplication means deduplication performed on the data in hard disk storage (HDD). Capacity deduplication is not enabled by default and is available if you have purchased a Pro or higher license. Note that you can enable capacity deduplication only if cache deduplication is enabled.

Enabling Deduplication

Deduplication is enabled at the storage container level (the **Cache** property for cache deduplication and the **Capacity** property for capacity deduplication). These storage container properties can be set in the web console or nCLI.

> Note: Deduplication enabled on storage containers having protected VMs lowers the replication speed.

In addition, Controller VMs in clusters with deduplication enabled need to be configured with more RAM:

- Cache deduplication: 24 GB RAM
- Capacity deduplication: 32 GB RAM

For more information about enabling deduplication, see Creating a Storage Container.

Deduplication Best Practices

This table shows when deduplication is recommended and when it is not.

Enable deduplication	Do not enable deduplication
• Full clones • Physical-to-virtual (P2V) migration • Persistent desktops	• Linked clones or Nutanix VAAI clones: Duplicate data is managed efficiently by DSF so deduplication has no additional benefit • Server workloads: Redundant data is minimal so may not see significant benefit from deduplication

Erasure Coding

Erasure coding increases the usable capacity on a cluster. Instead of replicating data, erasure coding uses a parity information to rebuild data in the event of a disk failure. The capacity savings of erasure coding is in addition to deduplication and compression savings.

If you have configured redundancy factor 2, two data copies are maintained. For example, consider a 6-node cluster with 4 data blocks (a b c d). In this example, we start with 4 data blocks (a b c d) configured with redundancy factor 2. In the following image, the white text represents the data blocks and the green text represents the copies.

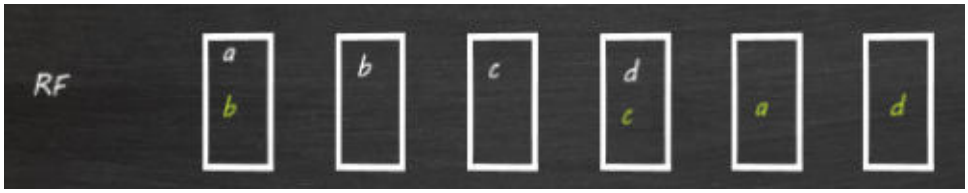

Figure 301: Data copies before Erasure Coding

When the data becomes cold, the erasure code engine computes parity “P” for the data by performing an exclusive OR operation.

Figure 302: Computing Parity

Once parity is computed, the data block copies are removed and replaced with the parity information. Redundancy through parity results in data reduction because the total data on the system is now a+b+c+d+P instead of 2 × (a+b+c+d).

> Note: Each block in the stripe is placed on a separate node to protect from a single node failure.

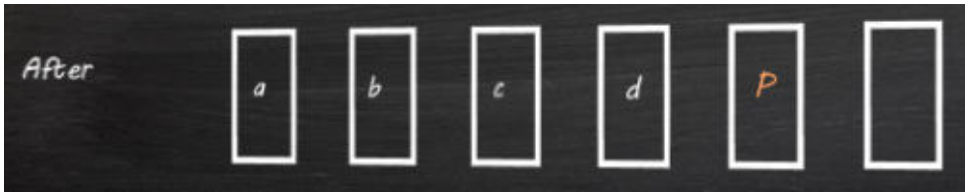

Figure 303: After Computation of Parity

If the node containing data block c fails, block c is rebuilt using the rest of the erasure coded stripe (a b d and P) as displayed in the following image.

Block c is then placed on a node that does not have any other members of this erasure coded stripe.

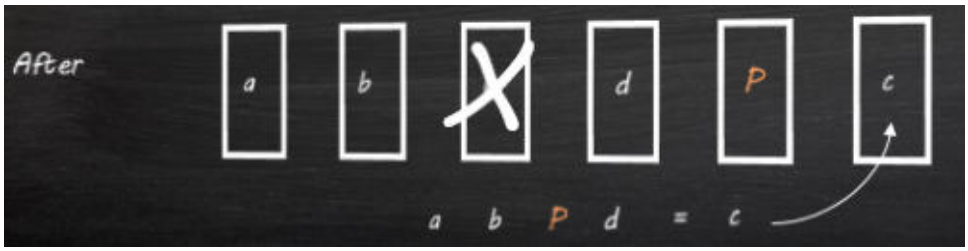

Figure 304: Post Node Failure

Note: When the cluster is configured for the redundancy factor 3, two parity blocks are maintained so that the erasure coded data has the same resiliency as the replicated data. An erasure coded stripe with two parity blocks can survive the failure of two nodes.

Example of Savings from Erasure Coding

The space savings from the erasure coding depends on the cluster size, redundancy setting, and percentage of cold data. You can view the erasure coding usage savings from the storage container summary.

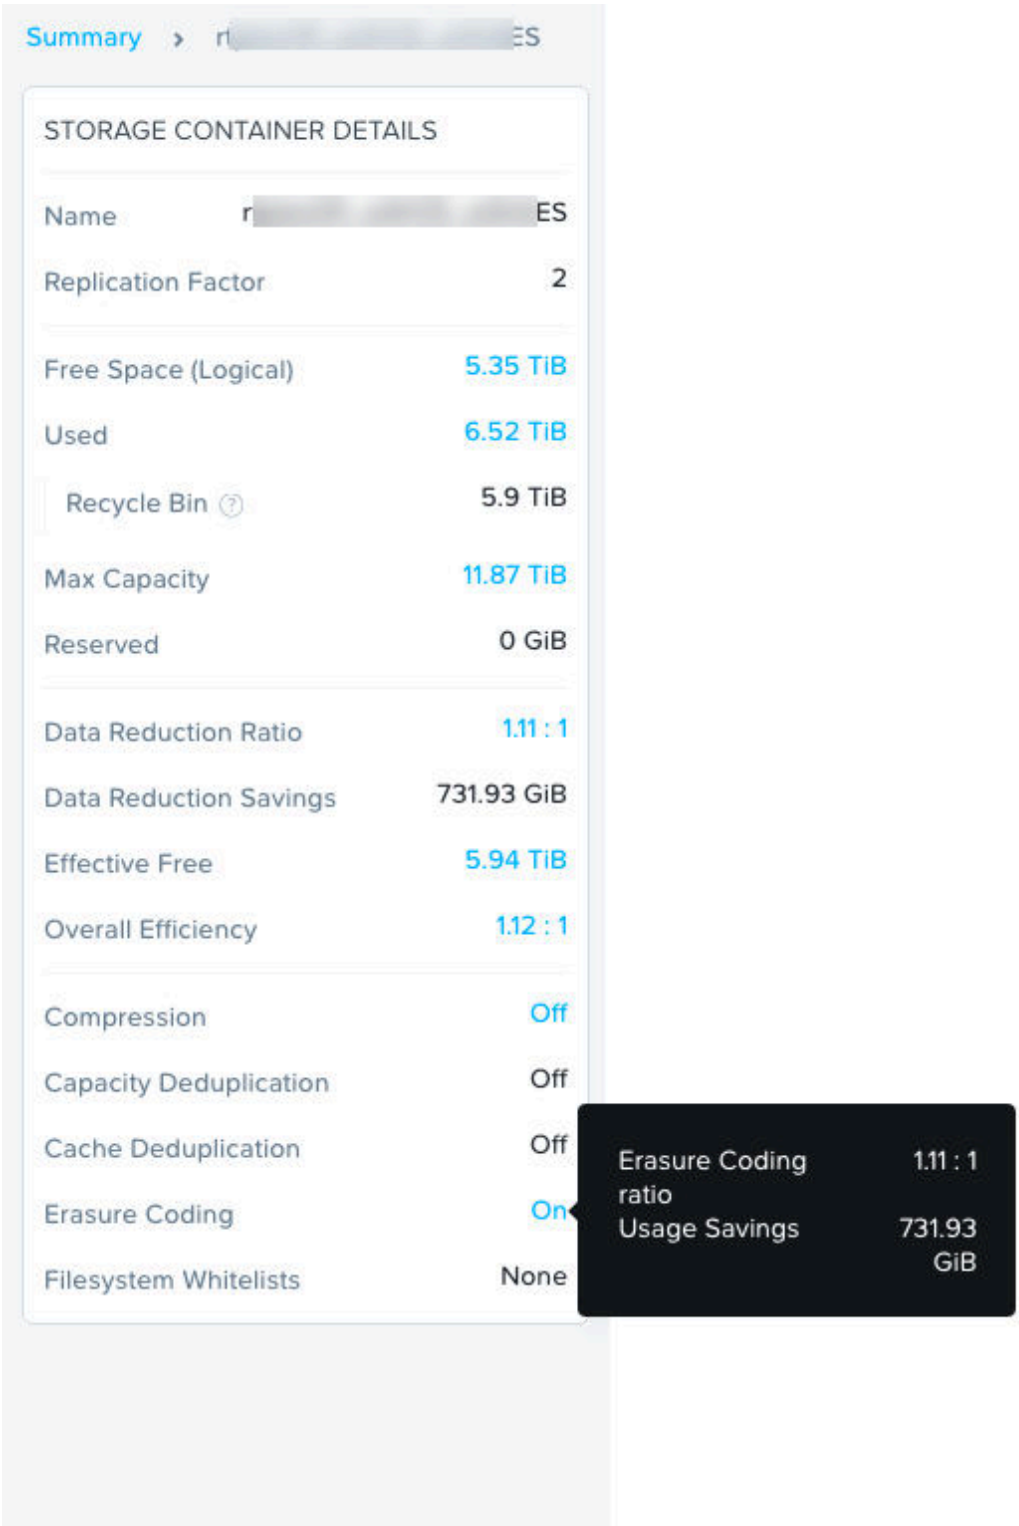

Figure 305: Storage Container Summary: Usage Savings Screen

In a 6-node cluster configured with redundancy factor 2, erasure coding uses a stripe size of 5. 4 nodes for data and 1 node for parity. The sixth node in the cluster ensures that if a node fails, a node is available for rebuild. With a stripe of 4 data to 1 parity, the overhead is 25%. Without erasure coding, the overhead is 100%.

Erasure coding stripe size adapts to the size of the cluster starting with the minimum 4 nodes with a maximum of 5 node stripe width. The following image displays the various configurations of cluster size, possible stripe widths, and approximate savings that might occur when erasure coding is enabled.

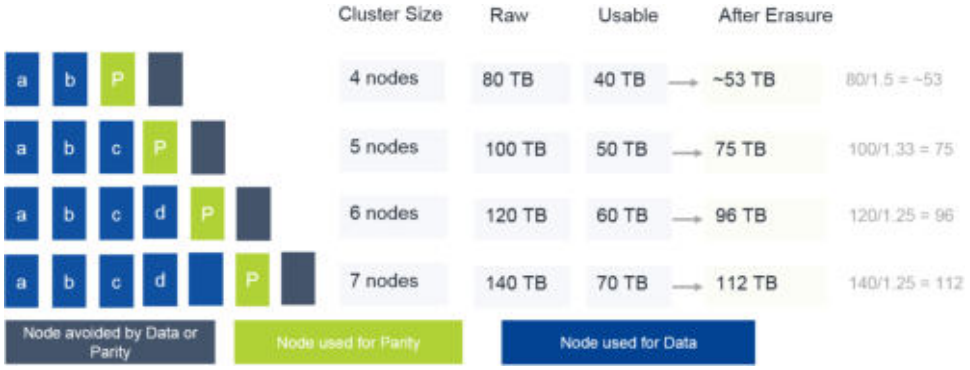

Figure 306: Example of Space Saving from Erasure Coding on 20 TiB Nodes

Erasure Coding Best Practices and Requirements

The following are the recommended best practices and requirements for implementing Erasure Coding:

- A cluster must have at least four nodes/blocks/racks to enable erasure coding. The cluster can have all four flash nodes or a combination of flash and hybrid nodes, or all hybrid nodes. If erasure coding is enabled on a storage container, a minimum of four blocks for RF2 or six blocks for RF3 is required to maintain block awareness.
- The recommended minimum configuration for multiple node removal operations is:

Table 711: Minimum Recommended Configuration for Erasure Coding

Desired Awareness Type	FT Level	Min. Units	Simultaneous Failure Tolerance
Node	1	4 nodes	1 node
Node	2	6 nodes	2 nodes
Block	1	4 blocks	1 block
Block	2	6 blocks	2 blocks
Rack	1	4 racks	1 rack
Rack	2	6 racks	2 racks

Note: We recommend that you always maintain a cluster size that is at least one node greater than the combined strip size (data + parity) to allow space to rebuild the strips if a node fails.

- AOS dynamically calculates the erasure coding strip sizes depending on the number of nodes, blocks, and racks. The maximum supported and recommended strip sizes are (4,1) or (4,2) depending on the nodes, blocks, and racks. We recommend that you should not

change the strip size. Greater strip sizes increases the space savings, however, it increases the cost of rebuild.

- Erasure coding effectiveness (data reduction savings) might reduce on workloads that have many overwrites outside of the erasure coding window. The default value for erasure coding window is seven days for write cold.
- Read performance is affected during rebuild and the amount depends on cluster strip size and read load on the system.
- Erasure coding is an asynchronous process, and hence the time taken to calculate and display space savings depends on the type and coldness of data. A minimum of two full curator scans are required to calculate the data savings.
- We recommend that you have RF+1 storage heavy or storage only nodes for all-flash clusters with storage heavy nodes. For example, if you have a four-node RF2 cluster, then you must add a minimum of three storage heavy nodes for optimum performance.

Inline Erasure Coding

Inline erasure coding encodes and creates erasure coding strips inline by erasure coding data without waiting for the data to become write cold as compared to the regular erasure coding functionality. In inline erasure coding the erasure coded strips are created on the same vDisk data, hence it is recommended to enable inline erasure coding for workloads that do not require data locality.

> Note:
>
> - Inline erasure coding can be enabled only for clusters running AOS version 5.18 or higher.
> - Erasure coding must be enabled on the container before you enable inline erasure coding.
> - If you want to use Inline erasure coding, contact Nutanix support.

Enabling Inline Erasure Coding

Inline erasure coding can be enabled only using nCLI. Inline erasure coding is added as a storage container parameter in Zeus. Note that erasure coding must be enabled on the container to enable inline erasure coding.

About this task

To manually enable erasure coding, run the following nCLI commands.

Procedure

- Enabling erasure coding on the storage container:

  ```
  ncli> container edit name=container_name erasure-code=on
  ```

  Replace `container_name` with the name of the storage container on which you want to enable erasure coding.

- Enabling inline erasure coding on the storage container.

  Run the following nCLI commands:

  ```
  ncli> container create name=container_name sp-id=storage_pool_id erasure-code=on inline-ec-enabled=true
  ```

  Replace *container_name* and *storage_pool_id* with the storage container name and storage pool ID on which you want to enable erasure coding.

  ```
  ncli> container edit inline-ec-enabled=true id=container_id
  ```

  Replace *container_id* with the ID of the storage container.

Capacity Reservation Best Practices

By default, each storage container has access to all of the unused storage in the storage pool. If a storage pool has multiple storage containers, one storage container may take all the remaining storage space and leave others with no available space. To make sure that there is space for a storage container, you can enable capacity reservation.

Capacity reservation allows you to guarantee that a storage container has a minimum amount reserved that is unavailable to other storage containers.

Following are best practices for capacity reservation.

- Reserve capacity for a storage container only if the storage pool has multiple storage containers. Unless there is a specific reason to have multiple storage containers, Nutanix recommends having a single storage pool with a single storage container.
- In total, reserve no more than 90% of the space in the storage pool.
- When setting an advertised capacity for a storage container, be aware that some extra space should be allocated beyond the projected size of any VMs placed in the container to allow room for data that has not yet been garbage collected, which can be substantial depending on the workload (10% or more of the storage capacity in some cases).

Storage Container Management

A storage pool is created during cluster configuration. There is one storage pool per cluster. Storage efficiency features including compression, deduplication, erasure coding, and replication factor are enabled at the storage container level. To access the storage containers dashboard, select **Compute & Storage** > **Storage Containers** from the entities menu (see Entities Menu on page 13).

The storage containers dashboard displays information about storage containers across the registered clusters (see Storage Containers Summary View on page 122). Using Prism Central you can do the following:

-
  - Create a storage container (see Creating a Storage Container (AHV) on page 622, Creating a Storage Container (ESXi) on page 626, or Creating a Storage Container (Hyper-V) on page 629)
  - Update or modify a storage container (see Modifying a Storage Container on page 633)
  - Delete a storage container (see Deleting a Storage Container on page 637)
- For information about storage management tasks that are performed through Prism Element, see the *Storage Management* chapter in the *Prism Web Console Guide*.

  Note:

- A storage pool and three storage container are created automatically when the cluster is created.
- A storage container is not created if you have not configured the Controller VMs with enough memory. Controller VM memory allocation requirements differ depending on the models and features that are being used. For more information, see *Controller VM Memory Configurations* in *Acropolis Advanced Administration Guide*.
- The NutanixManagementShare, SelfServiceContainer, and a default-container are created by default.
  - The NutanixManagementShare storage container is a built-in storage container for Nutanix clusters for use with the Nutanix Files and Self-Service Portal (SSP) features. This storage container is used by Nutanix Files and SSP for file storage, feature upgrades, and other feature operations. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.
  - SelfServiceContainer storage container is also a built-in storage container for Nutanix clusters for use with Image service features.

  To ensure proper operation of these features, do not delete these storage containers. Nutanix also recommends that you do not delete this storage container even if you are not using these features.

Known Issues and Limitations

RF1 containers are not supported in Prism Central. You can create RF1 containers only from Prism Element. However, you cannot update or delete Prism Element created RF1 containers from Prism Central.

Creating a Storage Container (AHV)

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers on page 514 for information on how to configure the NTP server) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor. The storage container is accessible transparently for AHV.

To create a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage containers dashboard.

4. Click the **Create Storage Container** button.
   The **Create Storage Container** page appears.

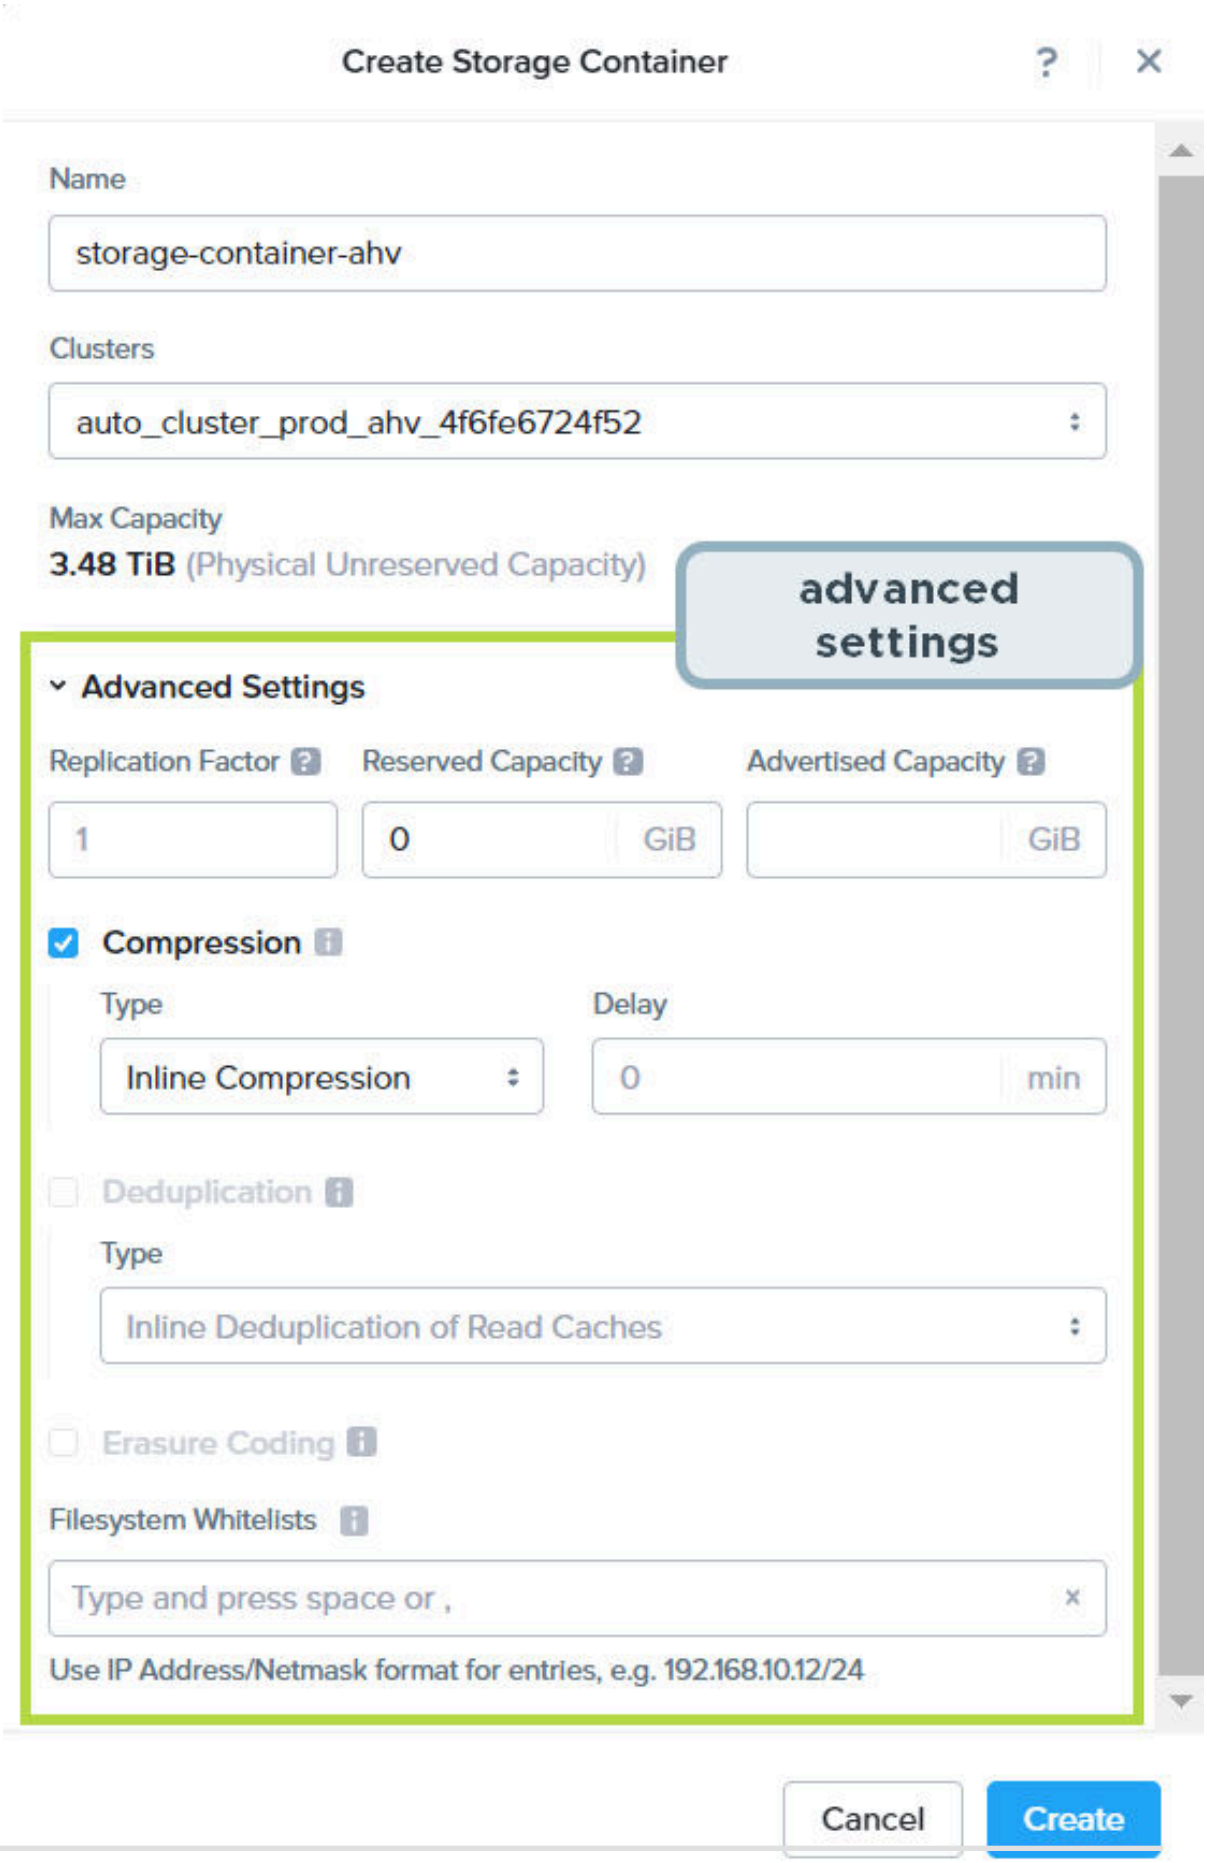

Figure 307: Create Storage Container (AHV)

5. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

      Note: This entity has the following naming restrictions.

      - Container Name Length: Maximum length is 75 characters.
      - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
      - Case Sensitivity: Container names are case sensitive.

   b. **Clusters**: Select a cluster from the drop-down list.

   c. **Max Capacity** displays the amount of free space available in the selected cluster.

6. To configure additional parameters, click the **Advanced Settings**.

   Do the following in the indicated fields:

   a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

      Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

      Note: To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See the *Increasing the Cluster Fault Tolerance Level* in the *Prism Web Console Guide*.

   b. **Reserved Capacity**: To reserve storage space for this storage container, enter the capacity (in GiB).

      You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 621 for more information.

   c. **Advertised Capacity**: To reserve the maximum storage space for this storage container, enter the capacity (in GiB).

      This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity.

   d. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 indicates that the data is compressed immediately as it is written. To configure the delay time between the write and compression, under **Type**, select **Post Process Compression** from the drop-down menu. For post-process compression, where data is compressed after it is written, Nutanix recommends settings a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

      All data in the storage container is compressed when you select **Compression**. See Compression on page 614 for information about using compression.

   e. **Deduplication**: Select the **CACHE** check box to perform inline deduplication of read caches to optimize performance. If you enable this option, the Controller VMs must be configured to have at least 24 GB of RAM. This feature is primarily recommended

for full-clone, persistent desktops, and physical to virtual migration use cases. Turning deduplication on for VAAI clone or linked clone environments is not recommended.

Select the **CAPACITY** check box to perform post-process deduplication of persistent data. This option is recommended primarily for full clone, persistent desktops, and physical to virtual migration use cases that need storage capacity savings (not just performance savings from deduplication). It is further recommended that the Controller VMs have at least 32GB of RAM and 300GB SSDs for the metadata disk to use this option.

> Important: Deduplication is only supported on clusters with a minimum of three nodes.

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 616.

> Note: Erasure coding is supported on clusters with a minimum of 4 nodes when using RF2 and a minimum of 6 nodes when using RF3.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form *ip_address/ netmask*).

A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

> Note: Setting a storage container level whitelist overrides any global whitelist for this storage container.

7. When all the field entries are correct, click the **Create** button.

Creating a Storage Container (ESXi)

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers on page 514 for information on how to configure the NTP server) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor. For ESXi, the storage container is accessible as an NFS datastore. This requires access to the vSphere APIs. Ensure that you have appropriate license of vSphere to access the APIs.

To create a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage containers dashboard.

4. Click the **Create Storage Container** button.
   The **Create Storage Container** page appears.

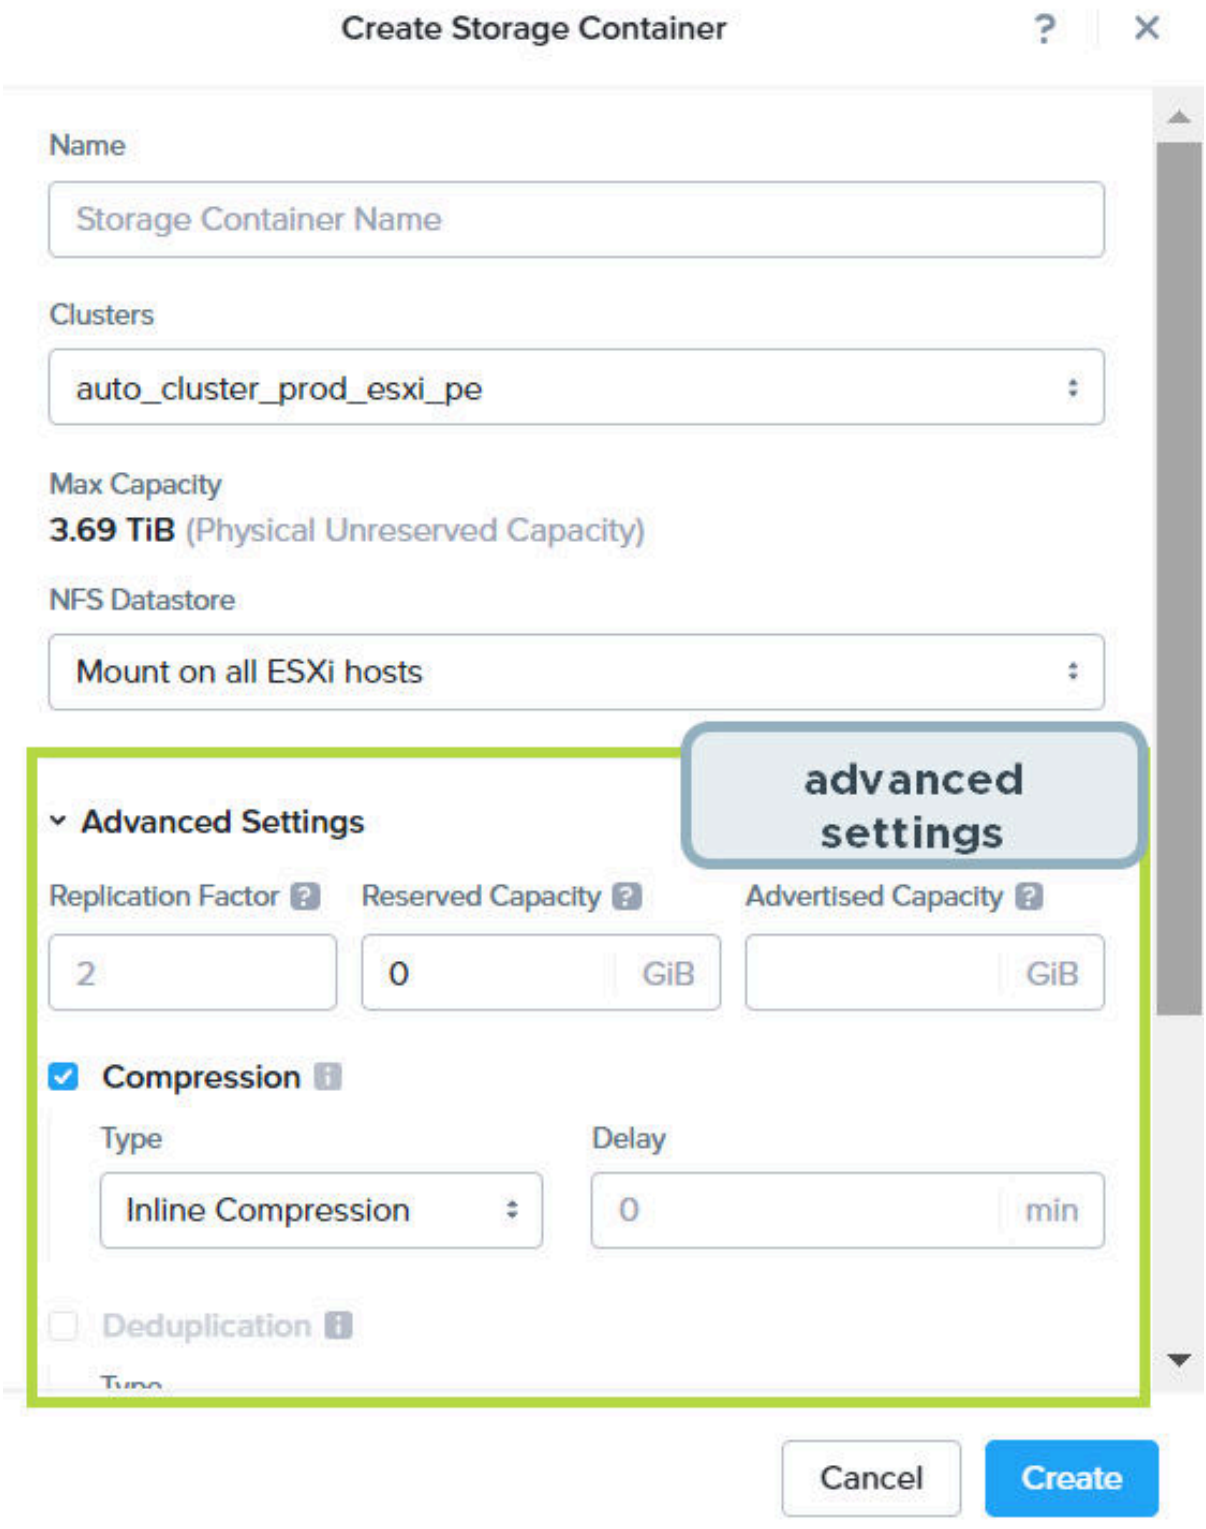

Figure 308: Create Storage Container (AHV)

5. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

   > Note: This entity has the following naming restrictions.
   >
   > - Container Name Length: Maximum length is 42 characters.
   > - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   > - Case Sensitivity: Container names are case sensitive.

   b. **Clusters**: Select a cluster from the drop-down list.

   c. **Max Capacity** displays the amount of free space available in the selected cluster.

   d. **NFS Datastore**: Select the **Mount on all ESXi hosts** button to mount the storage container on all hosts. Select the **Mount on the following ESXi hosts** button to mount the storage container on a subset of hosts, which displays a list of host IP addresses below this field. Check the boxes of the hosts to be included.

Option	Description
Make default on all Hyper-V hosts	Makes this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.
Make default on particular Hyper-V hosts	Provides you with an option to select the hosts on which you want to make this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.

6. To configure additional parameters, click the **Advanced Settings**.

   Do the following in the indicated fields:

   a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

   Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > Note: To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See the *Increasing the Cluster Fault Tolerance Level* in the *Prism Web Console Guide*.

   b. **Reserved Capacity**: To reserve storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer

available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 621 for more information.

c. **Advertised Capacity**: Sets a maximum storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

d. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is configurable. For post-process compression, where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

   All data in the storage container is compressed when you select **Compression**. See Compression on page 614 for information about using compression.

e. **Deduplication**: Select the **CACHE** check box to perform inline deduplication of read caches to optimize performance. If you enable this option, the Controller VMs must be configured to have at least 24 GB of RAM. This feature is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. Turning deduplication on for VAAI clone or linked clone environments is not recommended.

   Select the **CAPACITY** check box to perform post-process deduplication of persistent data. This option is recommended primarily for full clone, persistent desktops, and physical to virtual migration use cases that need storage capacity savings (not just performance savings from deduplication). It is further recommended that the Controller VMs have at least 32GB of RAM and 300GB SSDs for the metadata disk to use this option.

   > Important: Deduplication is only supported on clusters with a minimum of three nodes.

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 616.

   > Note: Erasure coding is supported on clusters with a minimum of 4 nodes when using RF2 and a minimum of 6 nodes when using RF3.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form `ip_address/ netmask`).

   A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

   > Note: Setting a storage container level whitelist overrides any global whitelist for this storage container.

7. When all the field entries are correct, click the **Create** button.

Creating a Storage Container (Hyper-V)

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers on page 514 for information on how to configure the NTP server) and that the time

on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor. The storage container is accessible as an SMB share for Hyper-V.

To create a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage containers dashboard.

4. Click the **Create Storage Container** button.
   The **Create Storage Container** page appears.

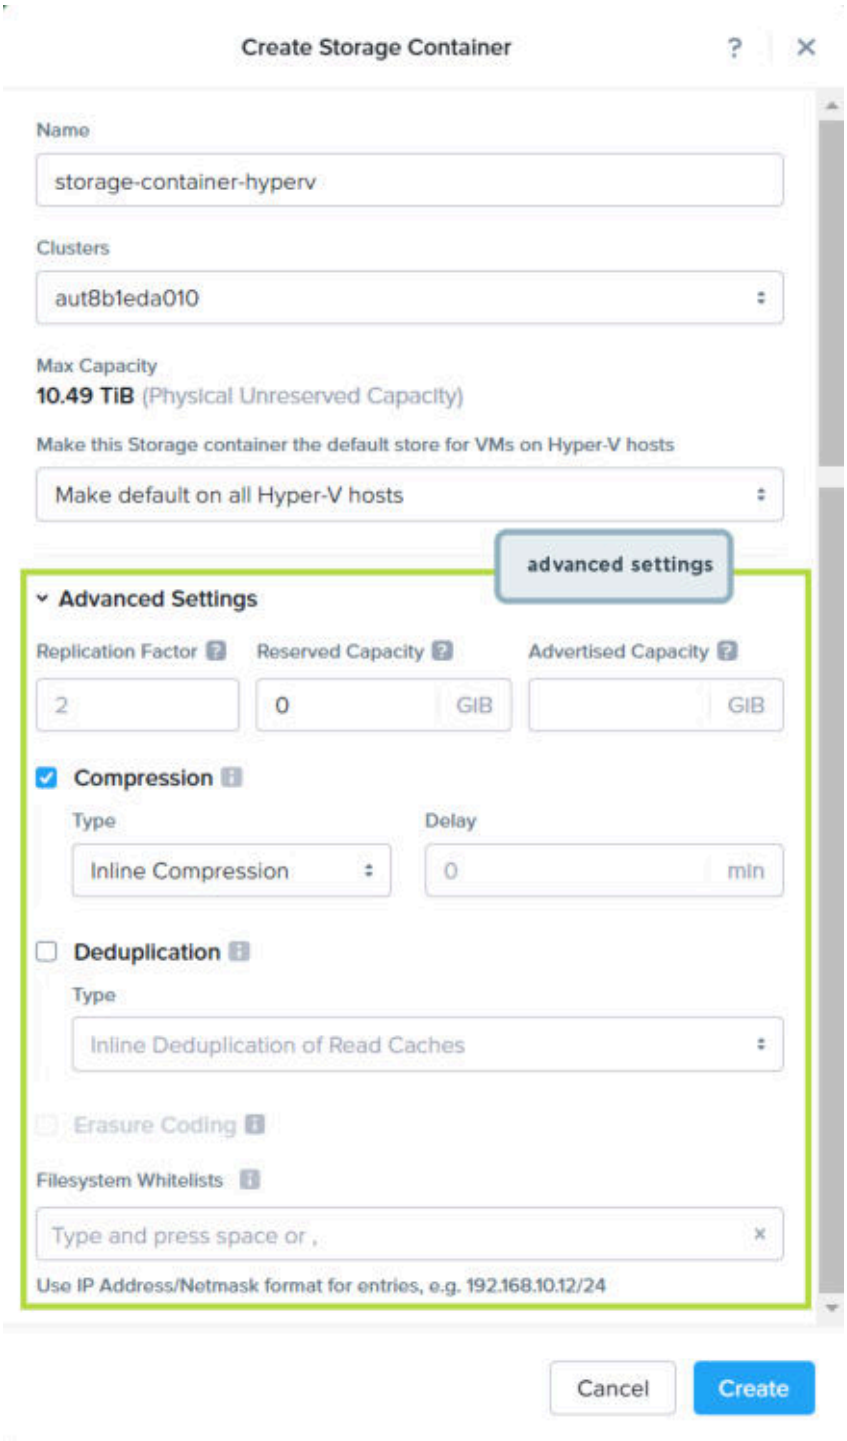

Figure 309: Create Storage Container (AHV)

5. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

   > Note: This entity has the following naming restrictions.
   >
   > - Container Name Length: Maximum length is 32 characters.
   > - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), decimal digits (0-9), dots (.), hyphens (-), and underscores(_).
   > - Case Sensitivity: Container names are case insensitive. For example, if there is a container with name ABCD, then creating another container with name abcd is restricted.

   b. **Storage Pool**: Select a storage pool from the drop-down list.

   The following field, **Max Capacity**, displays the amount of free space available in the selected storage pool.

   c. Set this storage container as default store for VMs on Hyper-V hosts.

   Depending on your selection the Create Virtual Machine Wizard of Hyper-V automatically populates the storage location with the relevant storage container.

Option	Description
Make default on all Hyper-V hosts	Makes this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.
Make default on particular Hyper-V hosts	Provides you with an option to select the hosts on which you want to make this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.

6. To configure additional parameters, click the **Advanced Settings**.

   Do the following in the indicated fields:

   a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

   Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > Note: To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See the Capacity Reservation Best Practices on page 621 for more information.

   b. **Reserved Capacity**: To reserve storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer

available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 621 for more information.

c. **Advertised Capacity**: Sets a maximum storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

d. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is configurable. For post-process compression, where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

   All data in the storage container is compressed when you select **Compression**. See Compression on page 614 for information about using compression.

e. **Deduplication**: Select the **CACHE** check box to perform inline deduplication of read caches to optimize performance. If you enable this option, the Controller VMs must be configured to have at least 24 GB of RAM. This feature is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. Turning deduplication on for VAAI clone or linked clone environments is not recommended.

   Select the **CAPACITY** check box to perform post-process deduplication of persistent data. This option is recommended primarily for full clone, persistent desktops, and physical to virtual migration use cases that need storage capacity savings (not just performance savings from deduplication). It is further recommended that the Controller VMs have at least 32GB of RAM and 300GB SSDs for the metadata disk to use this option.

   > Important: Deduplication is only supported on clusters with a minimum of three nodes.

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 616.

   > Note: Erasure coding is supported on clusters with a minimum of 4 nodes when using RF2 and a minimum of 6 nodes when using RF3.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form `ip_address/ netmask`).

   A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

   > Note: Setting a storage container level whitelist overrides any global whitelist for this storage container.

7. When all the field entries are correct, click the **Create** button.

Modifying a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be modified as conditions change.

About this task

Storage Containers can be modified to change how the data in that storage container is handled, for example to apply compression.

> Note: You cannot rename a storage container in an AHV cluster when modifying container details through the Update Storage Container dialog box. You cannot rename a storage container if it contains vdisks.

To modify a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage container dashboard.
4. To update the storage container, select the target storage container and then click the **Update** link.
   The Update Storage Container window appears, which includes the same fields as the Create Storage Container window. Through this window you can specify NFS datastore mounts, reserve storage capacity, enable (or disable) compression, enable or disable erasure

coding, select or deselect fingerprinting on writes which enables deduplication on reads, and configure filesystem whitelist.

For example,

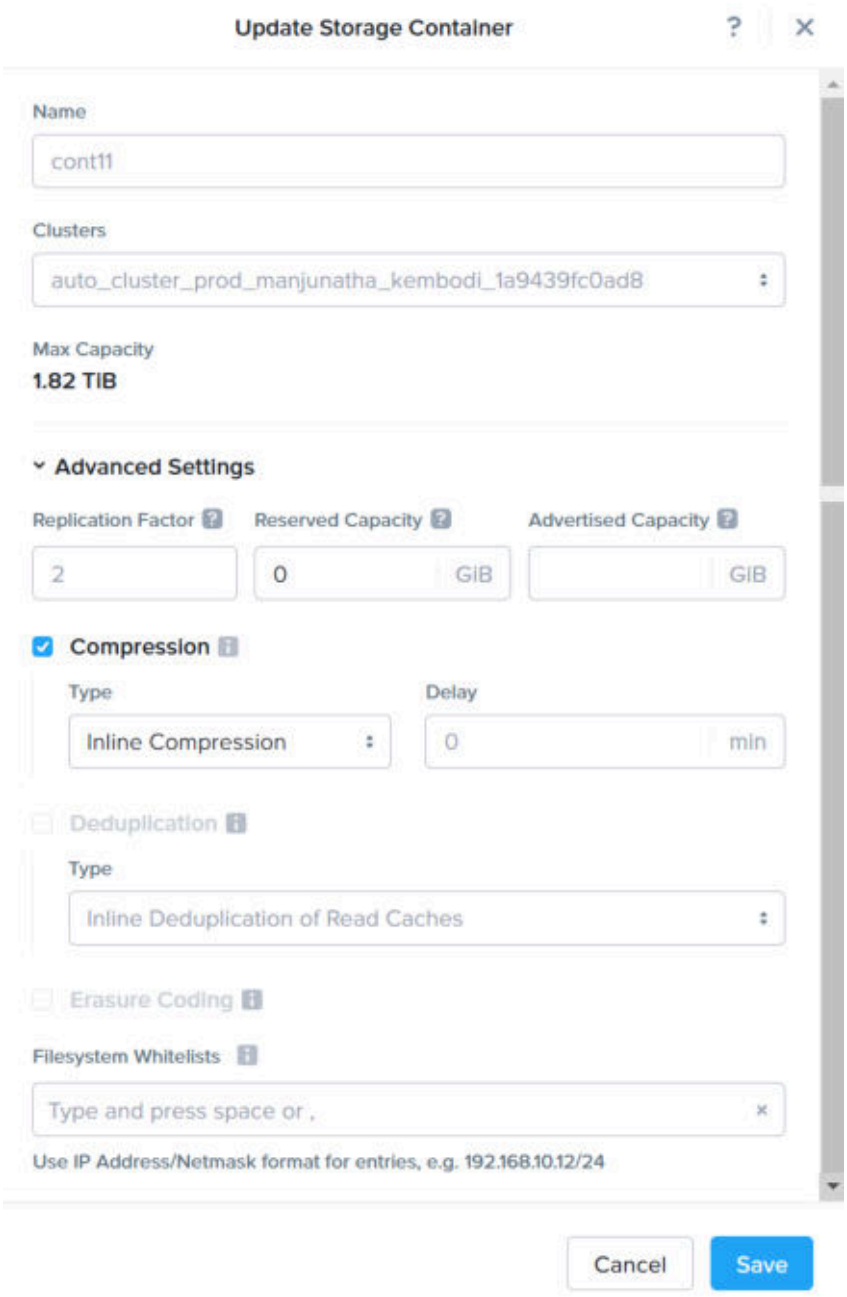

Figure 310: Update Storage Container (AHV)

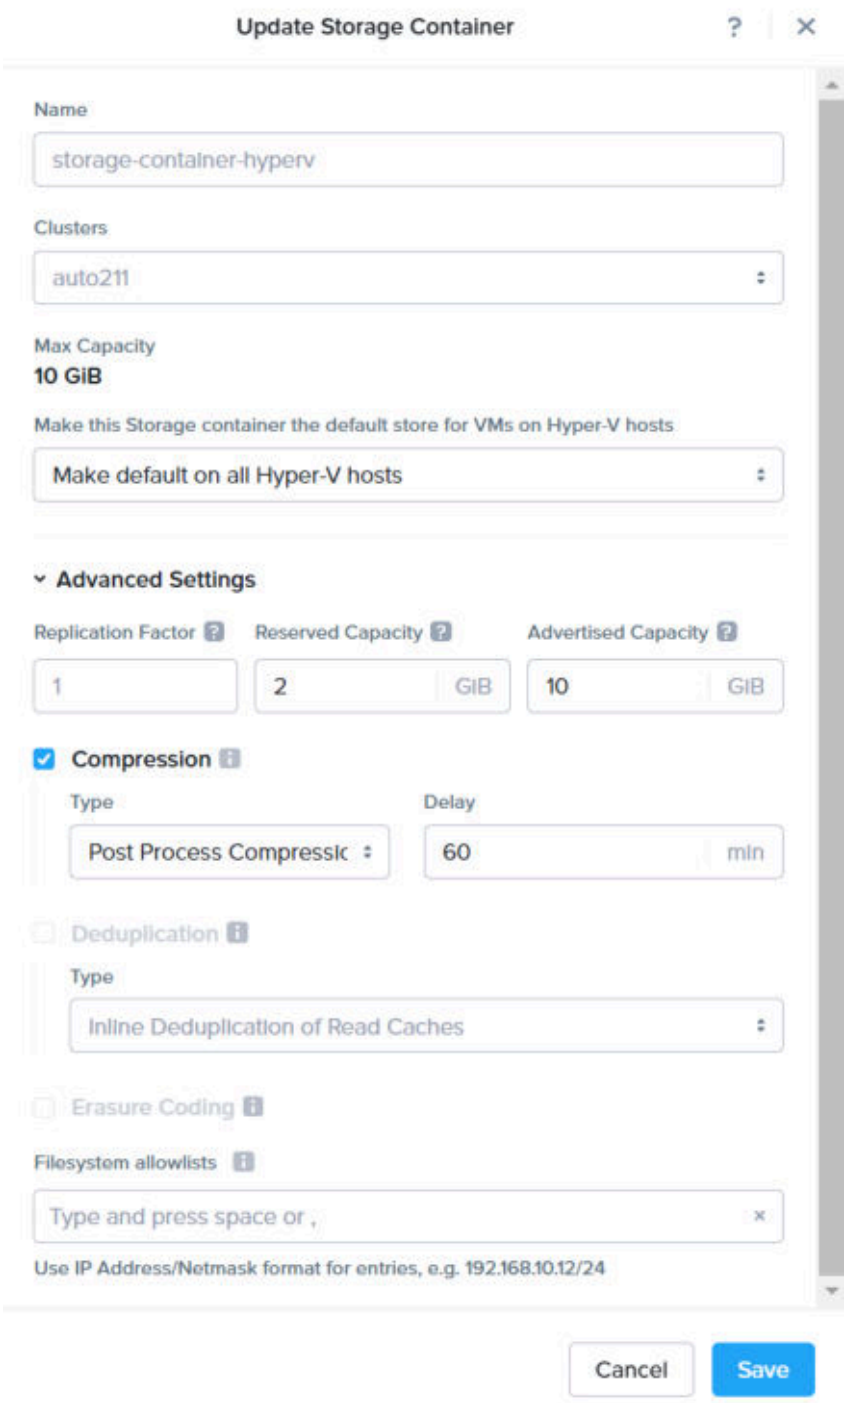

Figure 311: Update Storage Container (Hyper-V)

Note:

- If the compression policy is changed from compressed to uncompressed (or vice versa), the existing compressed (uncompressed) data in the storage container

will be uncompressed (compressed) as a background process when the next data scan detects the data that needs this change.

- The Prism Central does not provide an option to change the container replication factor. That can be done only through the nCLI (see *Increasing the Cluster Fault Tolerance Level* section in the *Prism Web Console Guide*).

Deleting a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be deleted.

About this task

To delete a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage container dashboard.
4. To delete the storage container, select the target storage container.
5. Click the **Delete** link.
   The Deleting Storage Container window appears which requests you to confirm deletion of the selected storage container.
6. Click the **Delete** button.

Creating a Volume Group

About this task

To create a volume group through Prism Central, do the following:

Figure 312: Creating a Volume Group

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275775660001 on Brightcove.

Procedure

1. Go to the **List** tab in the volume groups dashboard (see Volume Groups Summary View on page 139) and click the **Create Volume Group** button.
   The **Create Volume Group** page appears.

2. In the **Configuration** screen, do the following in the indicated fields:

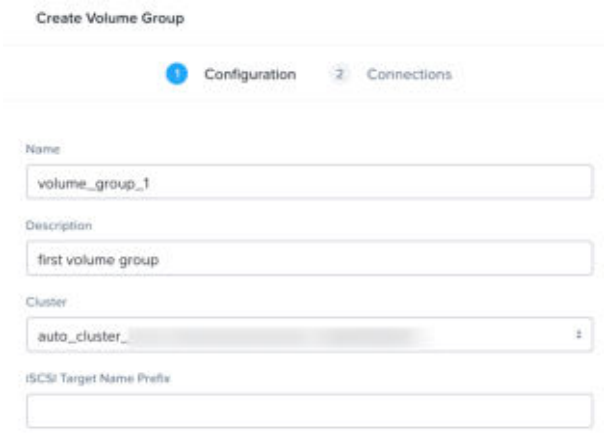

Figure 313: Create Volume Group: Configuration

a. **Name**: Enter a name for the volume group.

b. **Description**: Enter an optional description for the volume group.

c. **Cluster**: Select the cluster in which to create the volume group from the pull-down list.

   The list includes all clusters registered with this Prism Central.

d. **iSCSI Target Name** (read only): The iSCSI target name appears in the field. (You cannot change the name.)

e. **Virtual Disks**: Click the **Add Disk** button. In the **Add Virtual Disk** window, select the target volume group from the pull-down list, enter the disk size (in GiBs), and then click the **Add** button.

   The **Configuration** screen reappears with a table of virtual disks. The table includes index (assigned number), storage container (name), size, and action fields with a line for each disk in the volume group.

   - To add another disk to the volume group, click the **Add Disk** link above the table. This action redisplays the **Add Virtual Disk** window. Repeat this step for the new disk.
   - To change the size of a disk, click the pencil mark icon for that disk. An **Edit Virtual Disk** window appears in which you can change the size (but not the storage container).
   - To delete a disk, click the garbage can icon for that disk.

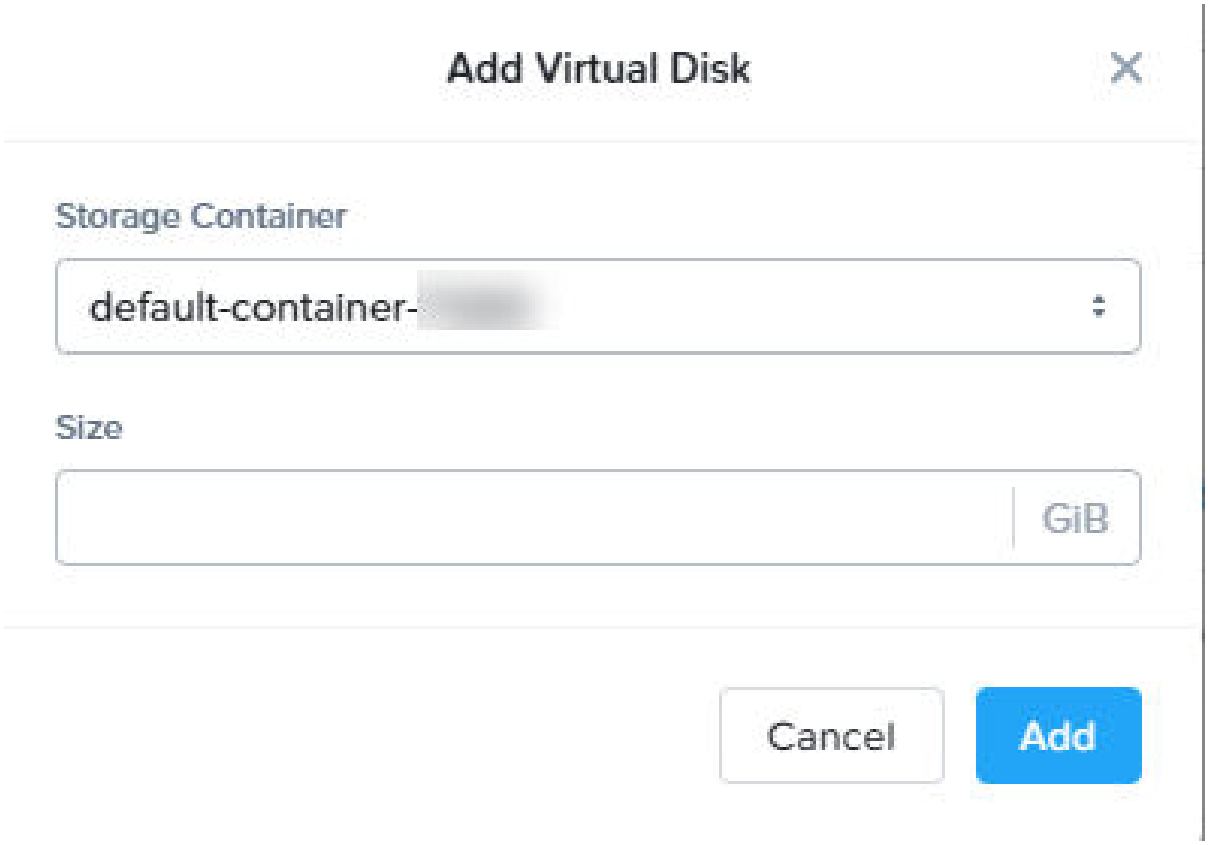

Figure 314: Add Virtual Disk Window

f. **Advanced Settings** (optional): Check the **Flash Mode** box to ensure no down migration of data occurs from the flash tier.

   Clicking **Advanced Settings** switches (display/hide) the **Flash Mode** box.

g. Click the **Next** button.

3. In the **Connections** screen, do the following in the indicated fields:

   » **Virtual Machines** - Allows you to configure direct attachment to VMs on the same cluster.

   » **External Clients** - Allows you to configure access to external clients or clients not residing on the same cluster.

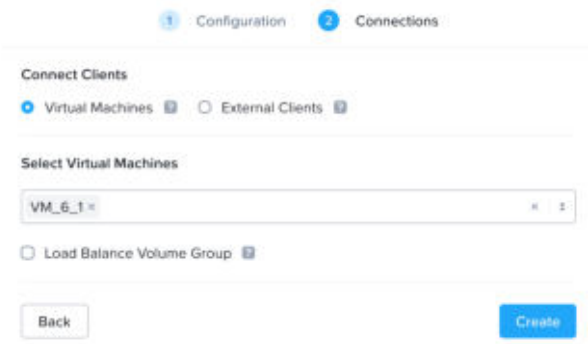

Figure 315: Create Volume Group: Connections - Virtual Machines

Configure the **Virtual Machines** settings as indicated.

- **Select Virtual Machines**: Select one or more VMs using the drop-down menu.
- (Optional) **Load Balance Volume Group**: Check to enable load balancing between the selected VMs. Virtual disks in the volume group are evenly distributed between the CVMs instead of being hosted locally.

  Note: You cannot change the load balance configuration once you attach the VMs.

Configure the **External Clients** configuration as indicated.

a. **CHAP Authentication**: Check the box to enable the Challenge Handshake Authentication Protocol (CHAP) and then enter the password in the indicated field.

b. **Attached Client**: Configure the clients as desired. The configured clients appear in the table.

   - To attach (unattach) a client, check (uncheck) the box for that client in the table.
   - To add a client, click the **Add Client** link. In the **Add External Client** window, enter the client IP address or IQN designator name in the indicated field, check the **CHAP**

Authentication box and password to enable CHAP authentication, and then click the **Add** button. The added client appears in the table.

- To update a client, click the pencil icon for that client in the table. The **Edit External Client** window appears. Check (uncheck) the box to enable (disable) CHAP authentication and then click the **Save** button. (You cannot change the client address.)

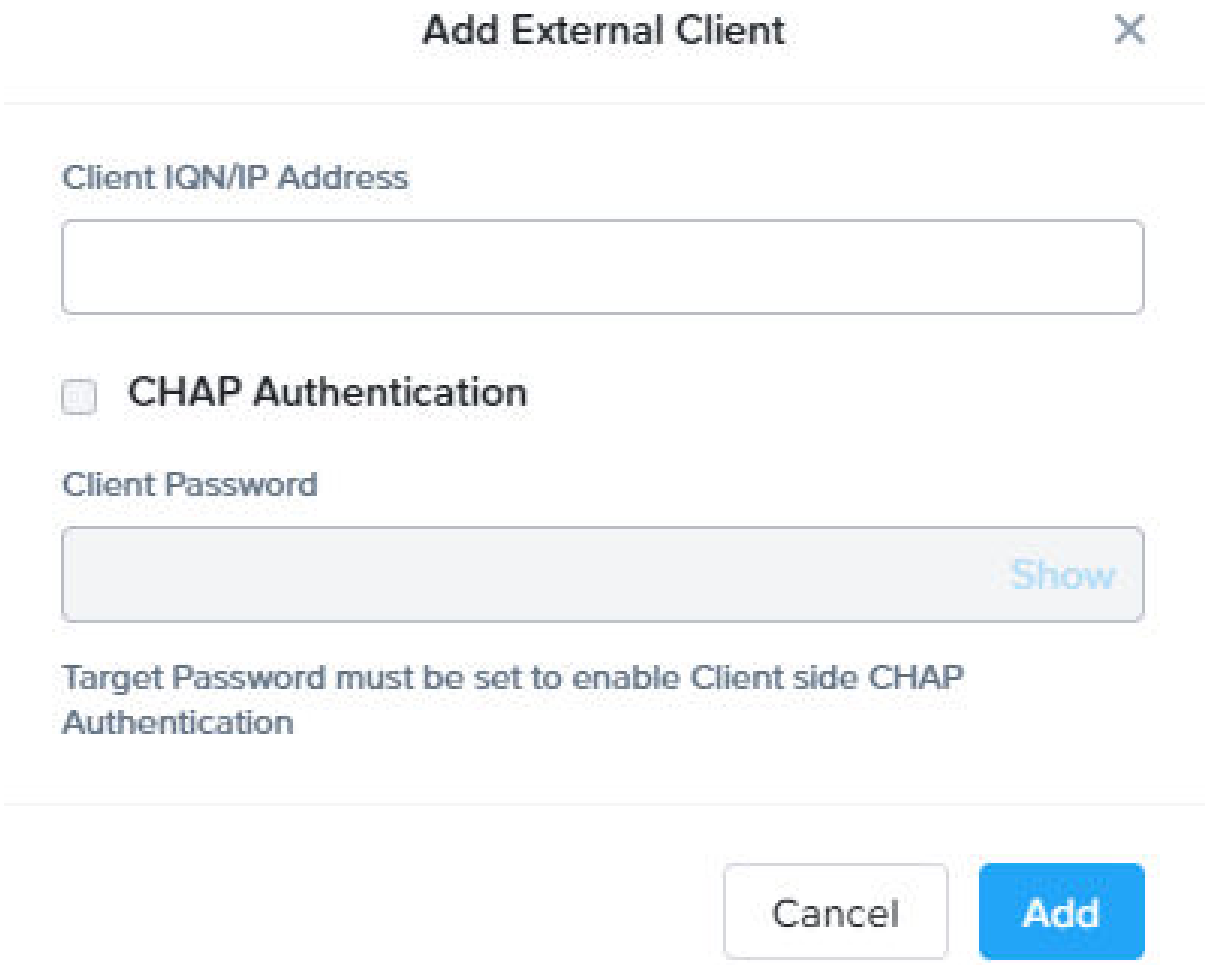

Figure 316: Add External Client Window

c. Click the **Create** button.

The volume group is created, and the new volume group appears in the list of volume groups.

Modifying a Volume Group

About this task

You can modify a volume group (update or delete) at any time after creating it.

Procedure

- To update a volume group, do the following:

  a. Do one of the following:

     » Go to the **Summary** tab in the details page of the target volume group (see Volume Group Details View on page 143) and click the **Update** button.

     » Select the target volume group in the **List** tab of the volume groups dashboard and then click the **Update** button (see Volume Groups Summary View on page 139).

  b. In the **Update Volume Group** page, update the target parameters and then click the **Save** button.

     The fields are the same as when creating the volume group (see Creating a Volume Group on page 637).

- To delete a volume group, either select the volume group in the **List** tab of the dashboard or go to the **Summary** tab in the details page and then click the **Delete** button. (In the dashboard **Delete** is an option in the **Action** drop-down menu.)

  If there are active connections, you must first remove the connections before deleting the volume group.

Managing Volume Group Connections

About this task

Connections to external clients are configured when the volume group is created (see Creating a Volume Group on page 637), but you configure authentication and add or remove connections to a volume group at any time.

Procedure

- To manage connections as a whole, do the following:

  a. Do one of the following:

     » Go to the **Summary** tab in the details page of the target volume group (see Volume Group Details View on page 143) and click the **Manage Connections** button.

     » Select the target volume group in the **List** tab of the volume groups dashboard and then select **Manage Connections** from the **Actions** pull-down menu (see Volume Groups Summary View on page 139).

     The **Manage Connections** page appears.

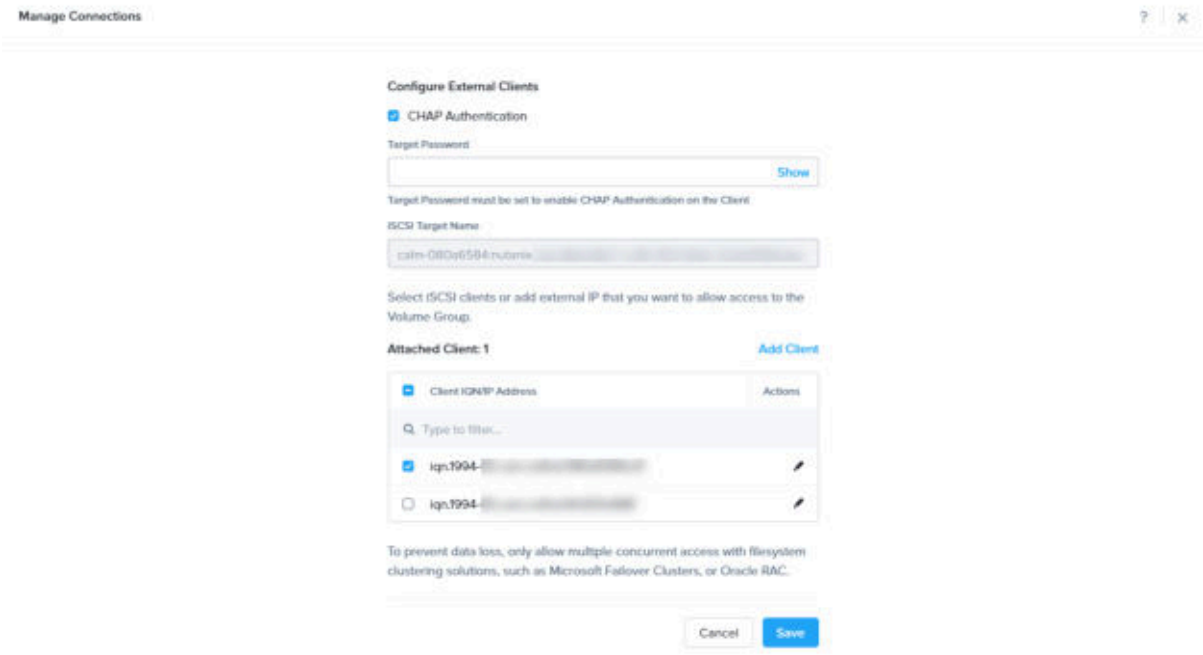

Figure 317: Manage Connections Page

  b. Check the **CHAP Authentication** box to enable the Challenge Handshake Authentication Protocol (CHAP) and then enter the password in the indicated field.

  c. The iSCSI target name appears in the **iSCSI Target Name** field. This is a read-only field because you cannot change the name.

  d. Configure the clients as desired in the **Attached Clients** section. The configured clients appear in the table.

     - To attach (unattach) a client, check (uncheck) the box for that client in the table.
     - To add a client, click the **Add Client** link. In the **Add External Client** box, enter the client IP address or IQN designator name in the indicated field, check the **CHAP Authentication** box and enter the password to enable CHAP authentication, and then click the **Add** button. The added client appears in the list.
     - To change the CHAP authentication for a client, click the pencil icon for that client in the table. The **Edit External Client** box appears. Check (uncheck) the box to enable

(disable) CHAP authentication and then click the **Save** button. (You cannot change the client address.)

e. Click the **Save** button.

- To add an external client connection to a volume group, do the following:

  a. Go to the **Connections** tab in the details page of the target volume group (see Volume Group Details View on page 143) and then click the **Add Connection** button.

  b. In the **Add External Client** window, do the following in the indicated fields:

  - **Add External Client**: Enter the client IP address or IQN designator name
  - **CHAP Authentication**: Check the box to enable CHAP authentication and enter the client password in the indicated field.
  - Click the **Add** button.

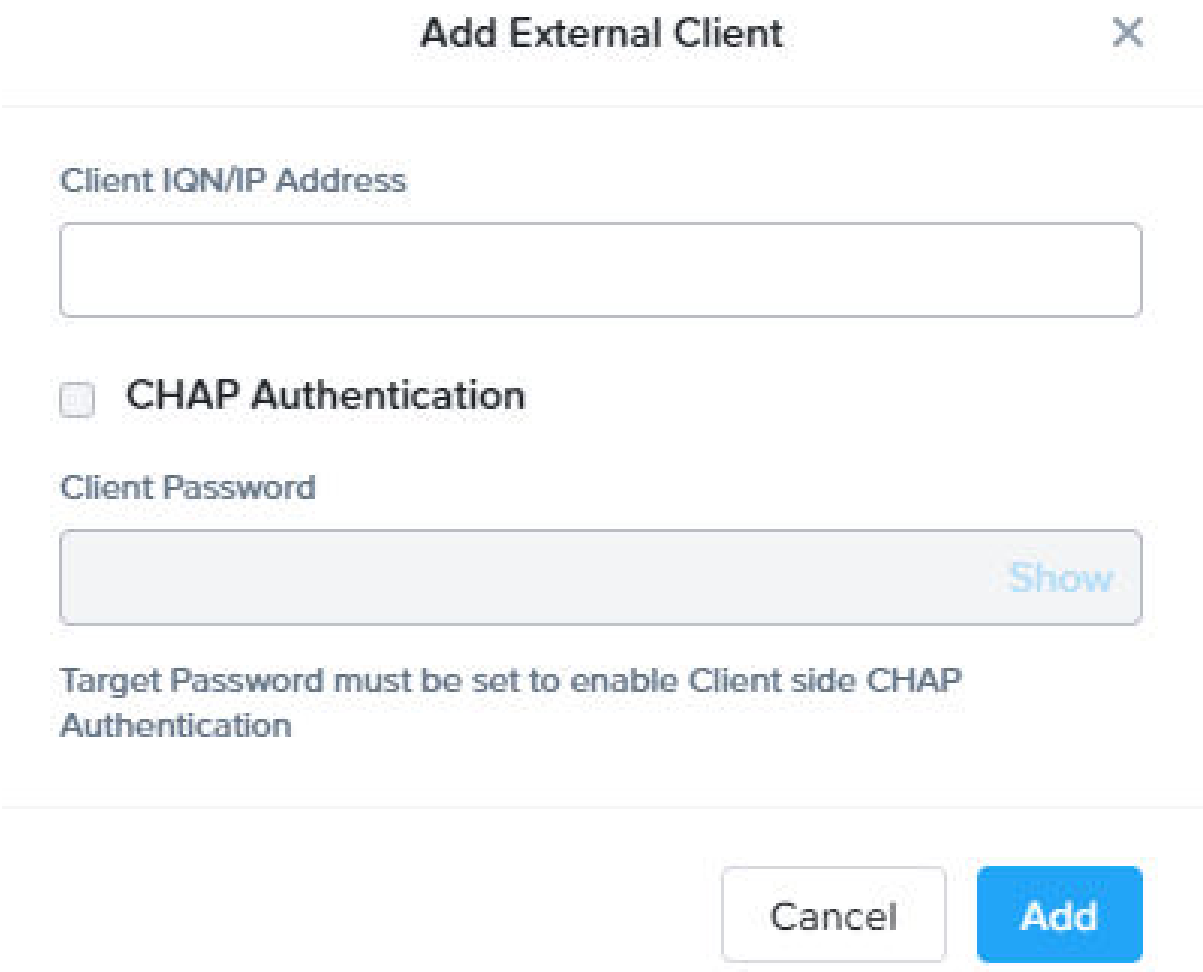

Figure 318: Add External Client Window

The connection is added, and the new client appears in the list of connections for the volume group.

- To update a client connection, do the following:

  Note: A separate CHAP authentication setting for an individual client is not allowed when CHAP authentication is not enabled for the volume group, and you cannot change the client address.

  a. Go to the **Connections** tab in the details page of the target volume group, select the target connection, and then click the **Update** button.

  b. In the **Edit External Client** window, uncheck the **CHAP Authentication** box to disable or check to enable. If you check the box, enter the client password in the indicated field.

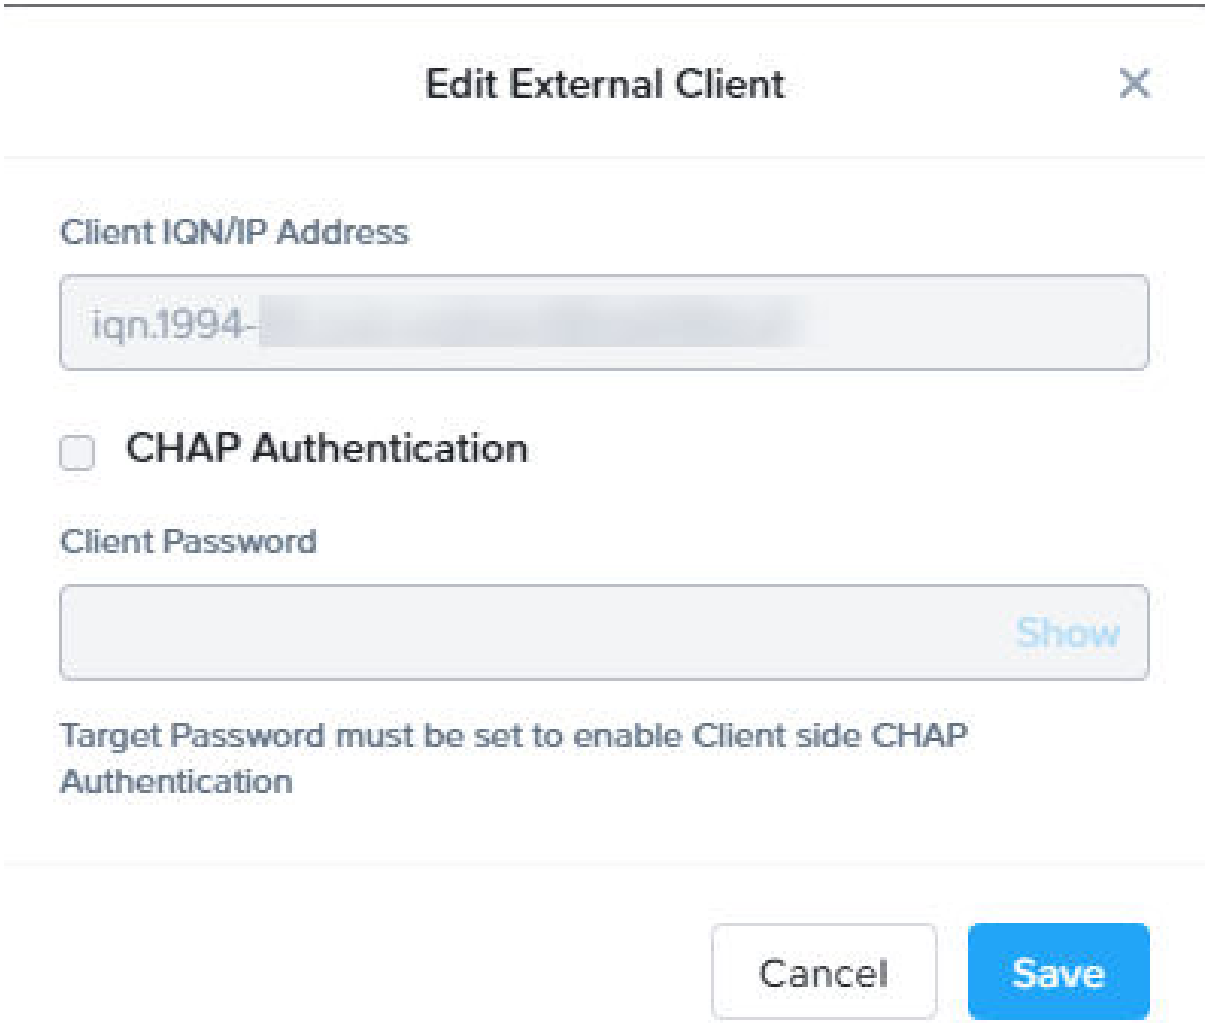

Figure 319: Edit External Client Window

- To remove a client connection, select the **Connections** tab in the details page of the target volume group, select the target connection, and then click the **Remove** button.
  The connection is removed, and the entry is deleted from the list of connections.

Managing Volume Group Virtual Disks

About this task

Virtual disks are added when the volume group is created (see Creating a Volume Group on page 637), but you add, update, or delete disks for a volume group at any time.

Procedure

- To add a virtual disk to a volume group, do the following:
  a. Go to the **Virtual Disks** tab in the details page of the target volume group (see Volume Group Details View on page 143) and then click the **Add Virtual Disk** button.
  b. In the **Add Virtual Disk** window, do the following in the indicated fields:
     - **Storage Container**: Select the target storage container from the pull-down list.
     - **Size**: Enter the desired size (in GiBs) for the virtual disk.
     - Click the **Add** button.

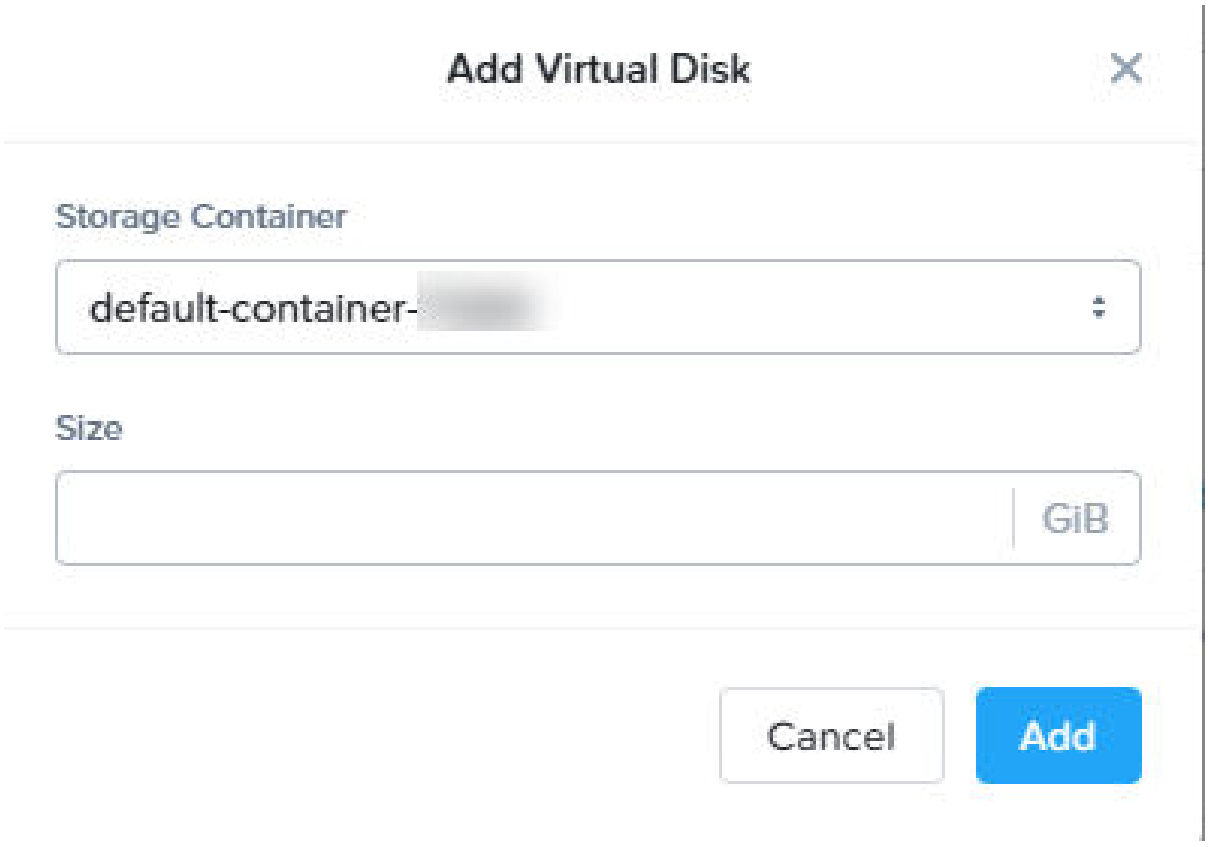

Figure 320: Add Virtual Disk Window

The disk is created, and the new disk appears in the list of virtual disks for the volume group.

- To update a virtual disk, do the following:

  Note: You can only change the disk size, not the storage container.

  a. Go to the **Virtual Disks** tab in the details page of the target volume group, select the target disk, and then click the **Update** button.

  b. In the **Edit Virtual Disk** window, update the disk size (in GiBs) as desired and then click the **Save** button.

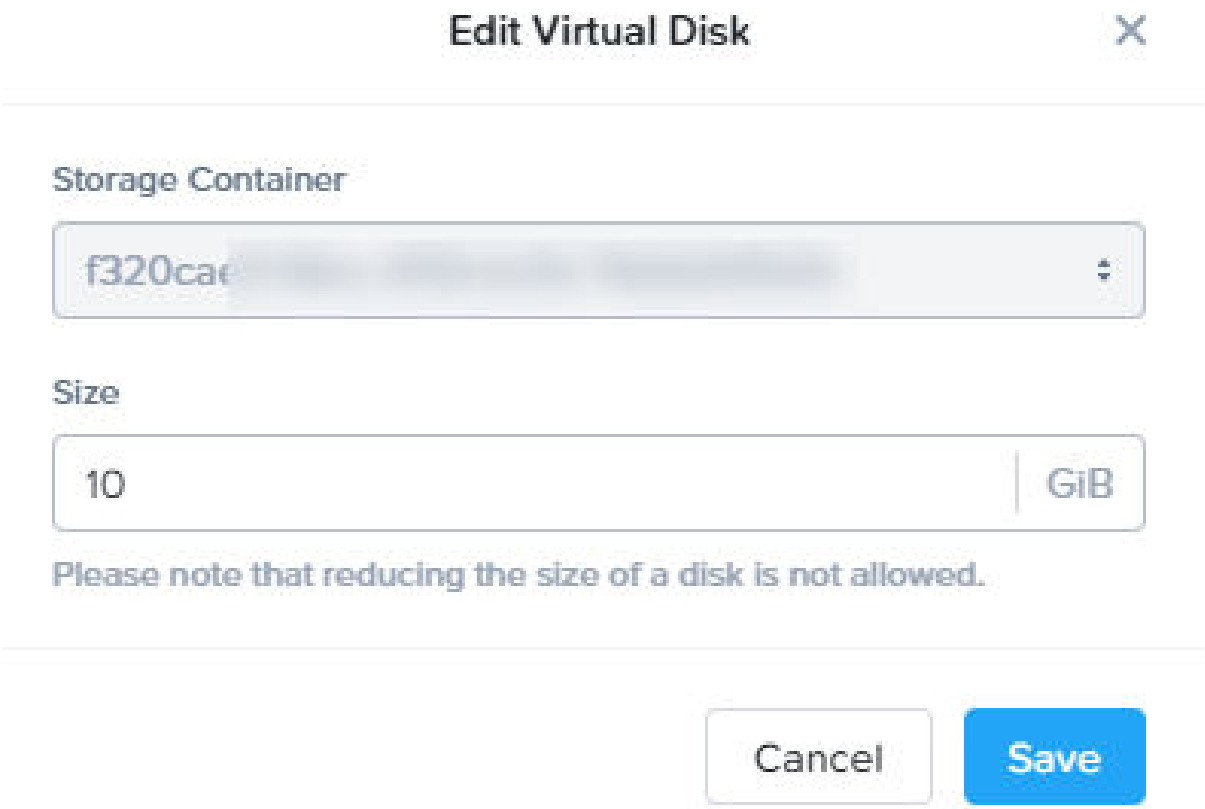

Figure 321: Edit Virtual Disk Window

- To delete a virtual disk, select the **Virtual Disks** tab in the details page of the target volume group, select the target disk, and then click the **Delete** button.

Storage Policy Management

Storage Policies in Prism Central let you manage the storage attributes like encryption, compression and QoS of entities like VMs. A single storage policy can manage the attributes of several entities that are associated with various Categories.

You can create a storage policy in Prism Central. A storage policy uses Categories to associate with entities such as VMs. Using a storage policy, you can manage parameters of entities (VMs) such as encryption, type of or lack of data compression, and IOPS or Throughput throttling values to be applied to the entities.

To create and manage storage policies, open the Storage Policies dashboard (see Storage Policies Summary View on page 133).

Requirements

Prism Central supports storage policies only on AHV. Prism Central supports storage policies only for VMs. It does not support storage policies for Volume Groups.

To successfully complete the inbuilt pre-checks and create a storage policy, ensure that:

1. You provide a unique name for every storage policy. In other words, no two storage policies can have the same name. If you use the name of an existing policy, Prism Central displays an error.
2. You provide a *defined* (non-default) value for at least one storage policy attribute or property. For example, for **Compression** the default value is **Inherit from Cluster**, so you may instead select **On** with **Inline** as the value. You could select or provide a defined value for any one ore more properties of the storage policy.

   See Creating or Updating a Storage Policy on page 650 for list of properties that you can configure during creation or updating of a storage policy.
3. The throttled IOPS value is greater than 99 or equal to -1 in the storage policies.

> Note: The storage policy engine completes the pre-checks in the same order as stated above. For example, if you repeat an existing storage policy name for a new storage policy, the policy engine fails the storage policy creation, generates the failure message stating `Duplicate name` and stops other checks. The engine triggers the storage policy creation only after all the pre-checks are validated.

Limitations

- Prism Central does not support storage policies for entities running on ESXi and Hyper-V.
- Prism Central does not support storage policies for Leap-based disaster recovery with on-prem recovery AZ (on-prem recovery cluster managed by a separate PC or Leap with multiple Prism Central deployments, see *Leap Administration Guide* for more information about on-prem AZ).
- Nutanix supports storage policies for Leap with single Prism Central. However, the storage properties defined for the container are applied to the snapshots replicated to the replication cluster. In other words, the snapshot data replicated to the remote site is transformed based on the settings defined for the remote site containers. The storage policies are applied only after the VM failover.
- Storage policies use categories to associate with entities like VMs. A storage policy cannot associate directly with an entity like a VM. Add an appropriate category to the VM and then associate the category with the storage policy.

  > Note: Storage Policies

  When you associate a category with a storage policy, all the VMs that are associated with that category are associated with the storage policy. The storage policy, thus, applies to all the VMs associated with that category.
- Prism Central does not support storage policies for volume groups (VGs).
- You can associate multiple categories with one storage policy. However, you can associate one category with only one storage policy. After associating a category with one storage policy, you cannot associate the same category with another storage policy. When you associate the same category with another storage policy, Prism Central adds the category to the **Associated Categories** list but displays an alert (an exclamation mark) and disables the **Save** button until you remove the category. The following message pops up when you hover on or click the alert (exclamation mark):

  ```
  This category has an existing storage policy (<name of storage policy>)
  ```

```
associated with it.
```

- The storage policy engine takes 30 minutes or more to apply a newly created storage policy to the associated entities like VMs. The time taken depends on the number of VMs associated with the storage policy.
- You can create a maximum of 100 storage policies irrespective of the type of Prism Central deployed (whether XLarge, Large or Small PC with Scale-out or Single Instance).
- You cannot create a storage policy for a snapshot. A storage policy applies to a live vDisk and the entire chain of snapshots of that vDisk.
- The snapshots of the vDisk that are in the storage are not encrypted if the live vDisk is deleted before AOS applies the storage policy to the vDisk.
- Storage policies support entity-centric encryption only if you did not already enable encryption on the underlying clusters in the deployment.
- After you enable encryption using storage policies, you cannot disable it.

  You can, however, delete the storage policy that has encryption enabled. In such cases, any new data written into the entities is not encrypted. The data that was written prior to the deletion of the storage policy remains encrypted.
- If two VMs share data and one or both the VMs have a storage policy with encryption enabled, then the shared data is encrypted.
- If two VMs share data and one of the VMs have a storage policy with compression enabled, even then the shared data may not be compressed.
- If you associate a VM with multiple storage policies associated, the policy engine sorts the storage policy list using a method based on the policy name and applies the highest ranking policy from the sorted list to the VM. If you add another storage policy later and this newly added policy becomes the highest ranking policy in the newly sorted list, the engine continues to apply the previously applied storage policy. To apply the newly added storage policy to the VM, remove the association of the VM with the previously applied storage policy.
- Storage policy inheritance is supported for operations on VMs performed using API v3. This support is not available for operations performed using other API versions like API v2.

For information about Storage Policy summary and detailed views, see Storage Policies Summary View on page 133.

Storage Policy Based Encryption

Storage policies allow you to make Encryption-at-Rest choices at the VM level.

Deployments can continue to have Nutanix Encryption-at-Rest capabilities scoped for the entire cluster. Storage policy provides the additional option to control the encryption scopes decisions at the VM. The Key Manager System (KMS) choice is configured as a cluster-wide setting, leveraging either Nutanix Built-In Native Key manager or leveraging a third party Key Manager. As part of the Key manager configuration, the scope of encryption-at-rest is defined either as Cluster-wide or specific entity-scoping. Entity-scoping means that the encryption-at-rest scope is defined at the container or the data policy level.

When Storage policy is used, AOS data for the VM is encrypted on the disk for all the vDisks attached to the VM. The main features include:

- Once encryption is enabled in a policy, it cannot be disabled. If the VM moves out of the policy or the policy is deleted, then the new writes to the VM are unencrypted and existing data remains encrypted.

- The system automatically generates a unique key for each storage container.
- Multi-tenant requirements having a unique key per tenant can be fulfilled by ensuring that the set of VMs belonging to each tenant is mapped to its own container. Nutanix recommends having only 50 tenants or less per cluster.
- Encryption can be enabled for new and existing VMs.
- Storage policies do not support Volume Group. For any VM with a Policy with encryption enabled, the data stored in the VG is not encrypted.
- Unencrypted VM needs to remain unencrypted during the cloning process.
- Shared data between the encrypted VMs and unencrypted VMs is encrypted.

For information about containers, see the *Storage Management* section in the *Prism Web Console Guide*.

For information about configuring key manager, see the *Security Guide*.

Creating or Updating a Storage Policy

You can create or update a storage policy in Prism Central only.

About this task

The Create Storage Policy has two tabs - **Configuration** and **Association**. The **Configuration** tab lets you configure the Data Security (Encryption), data Reduction (Compression) and QoS (IOPS or Throughput throttling) aspects of the storage policy. The **Association** tab lets you associate the storage policy with categories (and thence, entities).

The Create Storage Policy and the Update Storage Policy pages have the same storage policy attribute or property configuration fields.

Do the following on the Storage Policies dashboard:

Procedure

1. Click **Create Storage Policy**.
   On the Create Storage Policy page, the **Configuration** tab opens.

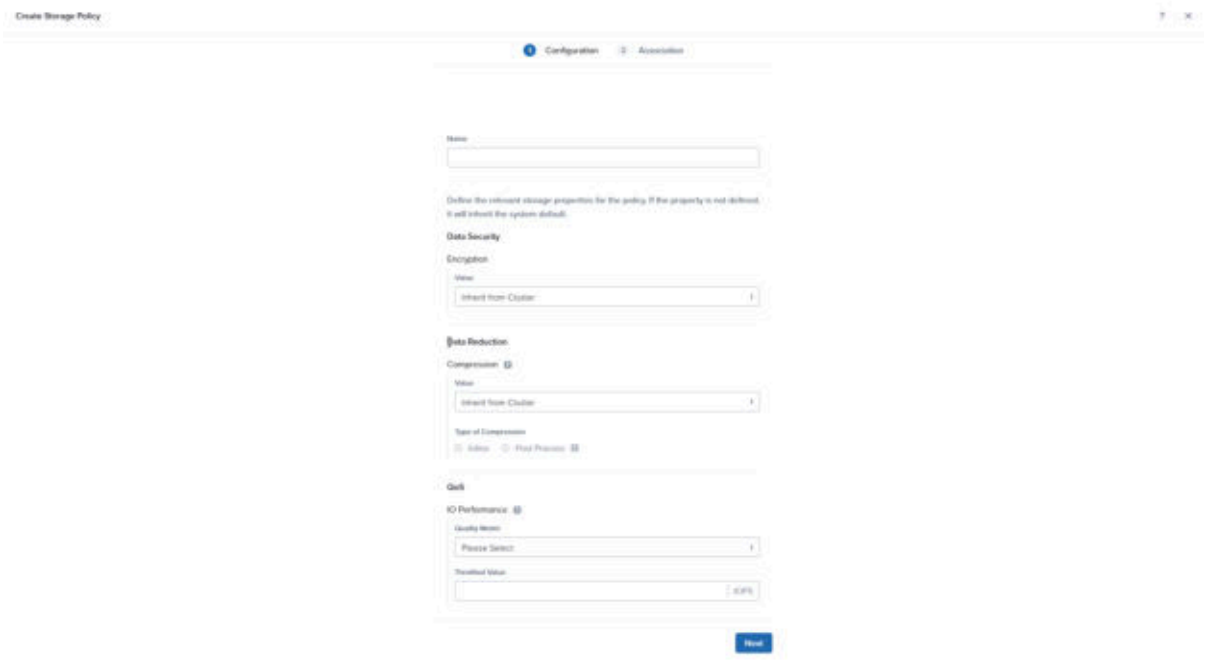

Figure 322: Create Storage Policy - Configuration Tab

Configure the following properties in the **Configuration** tab. Select a defined (non-default) value for a least one of the properties.

Table 712: Configuration Tab Fields

Field	Description and Values
Name	Enter a unique name for the storage policy. Note: Do not repeat a name previously used for another storage policy. The built-in pre-check fails if you enter a non-unique name.
Data Security	
Encryption Value	Select one of the following values from the drop-down list. • **Enabled**—to enable encryption through the storage policy. Note: After you enable encryption using storage policies, you cannot disable it. • **Inherit from Cluster** (default)—the category or entity inherits the encryption configuration from the cluster. This is the default value for Encryption.
Data Reduction	

Field	Description and Values
Compression Value	Select one of the following values from the drop-down list. • **On**—to enable compression. If you select **On** then you must select the type of compression that you want to apply. You can select **Inline** compression (default) or **Post Process** compression. If you select **Post Process**, the data is compressed with some delay (up to 3600 seconds) that is set by the system default. • **Off**—to disable compression. • **Inherit from Cluster** (default)—the category or entity inherits the compression configuration from the cluster. This is the default value for Encryption.
IO Performance	
Quality Metric	Select one of the following values from the dropdown list. • **Throughput** • **IOPS**
Throttled Value	Provide a numeric value. Note: For **IOPS** the value must be greater than 99 or equal to -1 (minus 1).

2. Click **Next**.
   If you did not select or enter an appropriate value and the pre-checks failed, the relevant error is displayed below the field that failed the pre-check. Select or enter an appropriate value and click **Next** again.

3. On the **Association** tab, in the **Add Category** field, click the category or categories you want to associate with the storage policy.
   Enter the first 2-3 characters of the category name you want to associate with in the **Add Category** field. This opens a drop-down list of categories that you can select from. Click