NUTANIX
YOUR ENTERPRISE CLOUD

Nutanix Files

Nutanix Tech Note

Version 5.2 • October 2019 • TN-2041

Copyright

Copyright 2019 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 is a trademark of Nutanix, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Contents

1. Executive Summary.....5

2. Audience and Purpose.....6

3. Nutanix Enterprise Cloud Overview..... 7
3.1. Nutanix Acropolis Architecture.....8

4. Nutanix Files Architecture.....9
4.1. Nutanix Files: File Server Virtual Machine..... 9
4.2. Exports and Shares..... 13
4.3. Load Balancing and Scaling.....20
4.4. High Availability.....22
4.5. Active Directory and SMB Operations.....26
4.6. Active Directory, LDAP, and NFS Operations.....32
4.7. Multiprotocol..... 35
4.8. Quotas.....39
4.9. Selective File Blocking.....41
4.10. File Analytics.....42
4.11. Hypervisor-Specific Support.....47

5. Backup and Disaster Recovery.....48
5.1. Self-Service Restore.....48
5.2. Protection Domains and Consistency Groups.....49
5.3. Cluster Migration, Failure, and Restoration.....49
5.4. Cloning..... 50

6. Third-Party Integration..... 51
6.1. Antivirus.....51
6.2. File Operations Monitoring.....54
6.3. Intelligent Backup.....55

7. Conclusion.....59

Appendix.....60
About Nutanix.....60

List of Figures.....61

List of Tables.....64

1. Executive Summary

Nutanix Files is a software-defined, scale-out file storage solution that provides a repository for unstructured data, such as home directories, user profiles, departmental shares, application logs, backups, and archives. Flexible and responsive to workload requirements, Files is a fully integrated, core component of the Nutanix Enterprise Cloud.

You can deploy Nutanix Files on an existing cluster or a standalone cluster. Unlike standalone NAS appliances, Files consolidates VM and file storage, eliminating the need to create an infrastructure silo. Administrators can manage Files with Nutanix Prism, just like VM services, unifying and simplifying management. Integration with Active Directory enables support for quotas and access-based enumeration, as well as self-service restores with the Windows Previous Versions feature. Nutanix Files also supports file server cloning, which lets you back up Files off-site and run antivirus scans and machine learning without affecting production.

Nutanix Files can run on a dedicated cluster or be collocated on a cluster running user VMs. Nutanix supports Files with both ESXi and AHV. Files includes native high availability and uses the Acropolis Distributed Storage Fabric (DSF) for intracluster data resiliency and intercluster asynchronous disaster recovery. The DSF also provides data efficiency techniques such as erasure coding (EC-X) and compression.

2. Audience and Purpose

This tech note is part of the Nutanix Solutions Library. We wrote it for architects and systems engineers who want to use Nutanix Files as a simple way to deliver user and group file management. This document describes how to implement and operate Files in your datacenter.

We cover the following subject areas:

- Overview of the Nutanix architecture with Files.
- Load balancing of standard and distributed shares (SMB) and exports (NFS).
- High availability.
- Backup and recovery.
- Quotas and permission management.
- Antivirus.

Table 1: Document Version History

Version Number	Published	Notes
1.0	December 2016	Original publication.
1.1	May 2017	Updated Backup and Disaster Recovery section.
1.2	September 2017	Updated for version 2.2 features.
2.0	February 2018	Updated for version 3.0.
2.1	April 2018	Solution overview update.
2.2	August 2018	SMB share and NFS export updates.
3.0	October 2018	Updated for version 3.1 and updated product naming.
4.0	January 2019	Updated for version 3.2 and AOS 5.9 and later.
5.0	April 2019	Updated for version 3.5 and AOS 5.10 and later.
5.1	August 2019	Updated for version 3.5.1.
5.2	October 2019	Updated for version 3.6.

3. Nutanix Enterprise Cloud Overview

Nutanix delivers a web-scale, hyperconverged infrastructure solution purpose-built for virtualization and cloud environments. This solution brings the scale, resilience, and economic benefits of web-scale architecture to the enterprise through the Nutanix Enterprise Cloud Platform, which combines three product families—Nutanix Acropolis, Nutanix Prism, and Nutanix Calm.

Attributes of this Enterprise Cloud OS include:

- Optimized for storage and compute resources.
- Machine learning to plan for and adapt to changing conditions automatically.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.
- Simplified one-click upgrade.
- Native file services for user and application data.
- Native backup and disaster recovery solutions.
- Powerful and feature-rich virtualization.
- Flexible software-defined networking for visualization, automation, and security.
- Cloud automation and life cycle management.

Nutanix Acropolis provides data services and can be broken down into three foundational components: the Distributed Storage Fabric (DSF), the App Mobility Fabric (AMF), and AHV. Prism furnishes one-click infrastructure management for virtual environments running on Acropolis. Acropolis is hypervisor agnostic, supporting two third-party hypervisors—ESXi and Hyper-V—in addition to the native Nutanix hypervisor, AHV.

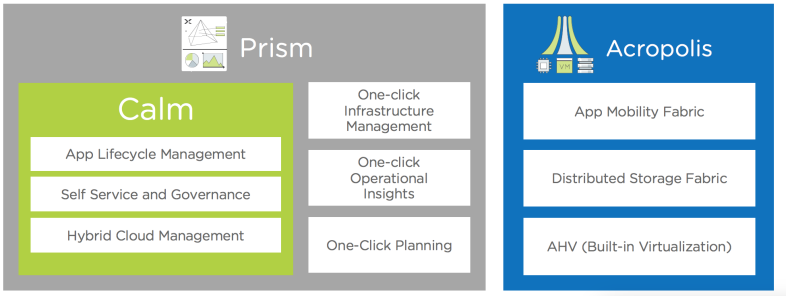

Figure 1: Nutanix Enterprise Cloud

3.1. Nutanix Acropolis Architecture

Acropolis does not rely on traditional SAN or NAS storage or expensive storage network interconnects. It combines highly dense storage and server compute (CPU and RAM) into a single platform building block. Each building block delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution requires no SAN constructs, such as LUNs, RAID groups, or expensive storage switches. All storage management is VM-centric, and I/O is optimized at the VM virtual disk level. The software solution runs on nodes from a variety of manufacturers that are either all-flash for optimal performance, or a hybrid combination of SSD and HDD that provides a combination of performance and additional capacity. The DSF automatically tiers data across the cluster to different classes of storage devices using intelligent data placement algorithms. For best performance, algorithms make sure the most frequently used data is available in memory or in flash on the node local to the VM.

To learn more about the Nutanix Enterprise Cloud, please visit the Nutanix Bible and Nutanix.com.

4. Nutanix Files Architecture

Nutanix Files is a scale-out approach that provides Server Message Block (SMB) and Network File System (NFS) file services to clients. Nutanix Files server instances are composed of a set of VMs (called FSVMs). Files requires at least three FSVMs running on three nodes to satisfy a quorum for high availability.

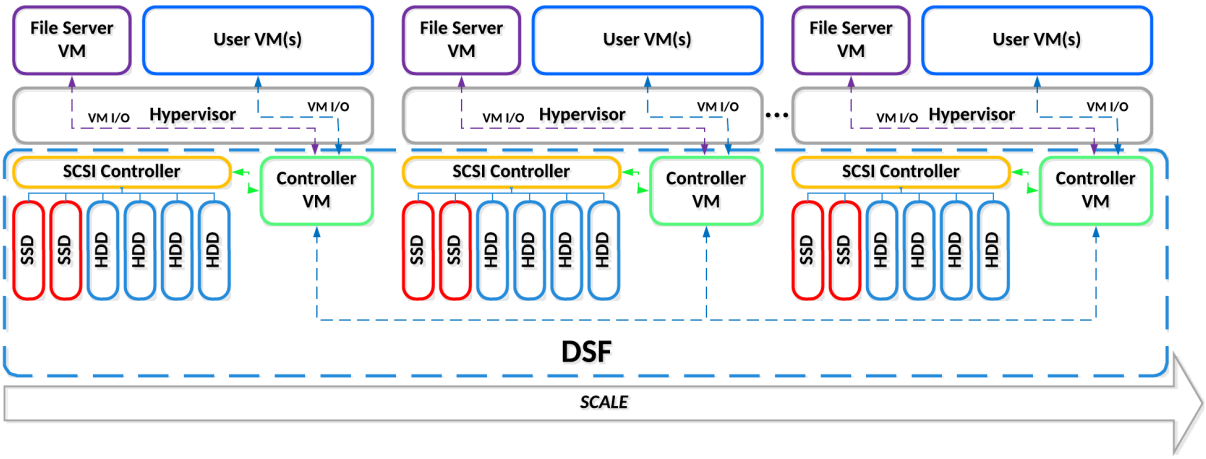

Figure 2: Nutanix Files Server Instances Run as VMs for Isolation from the DSF

4.1. Nutanix Files: File Server Virtual Machine

The File Server VM (FSVM) is based on CentOS and incorporates all the security and hardening that goes into the Nutanix Controller VM (CVM). All the FSVMs have the same configuration, starting with four vCPUs and 12 GiB of RAM. You can add more vCPUs, RAM, and FSVMs to the cluster. For each file server the number of FSVMs must be less than or equal to the number of nodes in the Nutanix cluster; however, you can create multiple file server deployments if needed. Each Nutanix Files cluster can support up to 16 FSVMs. Nutanix Files 3.1 introduced single FSVM deployments intended for one- and two-node Nutanix clusters. You can also have single FSVM deployments for larger clusters with 5.10.1 or later.

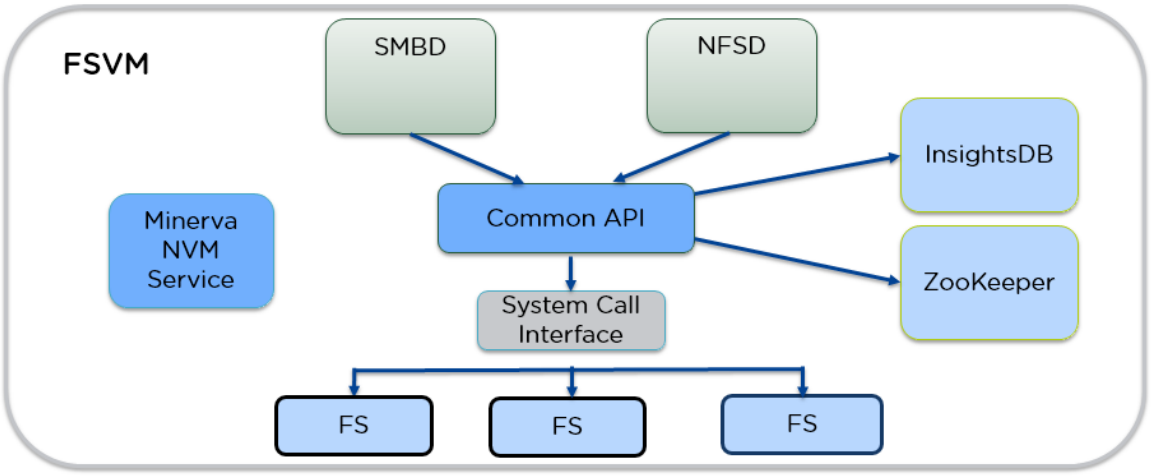

Figure 3: Data Path Architecture of Nutanix Files

Nutanix Files can support SMB and NFS from the same FSVM, and with Files 3.5 you can enable simultaneous SMB and NFS access to the same share and data, commonly referred to as multiprotocol support. Both SMB and NFS share a common library, allowing a modular approach. InsightsDB is a NoSQL database that maintains the statistics and alerts for Files. Zookeeper is a centralized service that maintains configuration information, such as domain, share or export, and IP information. The Minerva NVM Service talks to the local CVM and sends heartbeats to share health information and to help with failover.

Each FSVM stores file server data on multiple file systems that store share-specific data. The individual file system provides the snapshot capability used to provide Windows Previous Versions (WPV) support to clients. By using separate file systems for each share or export, Nutanix Files can scale to support billions of files in one cluster.

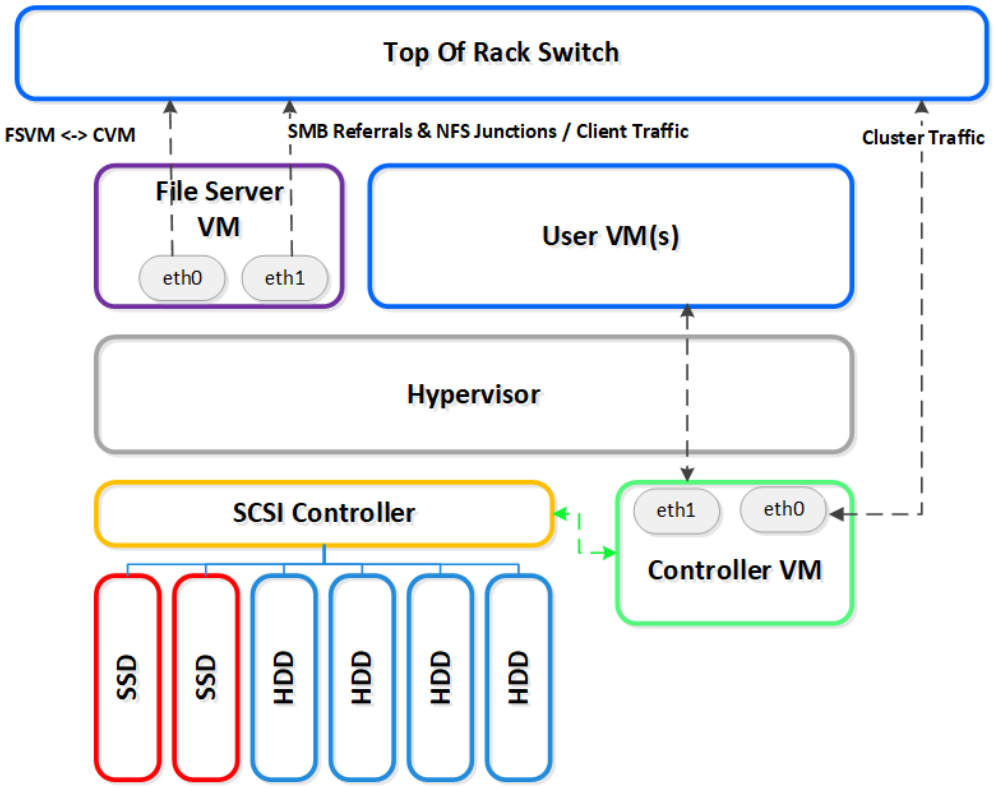

Figure 4: FSVM Internal Communication on One Node

The above diagram shows one FSVM running on a node, but you can put multiple FSVMs on a node for multitenancy.

Networking

The FSVM has two network interfaces: the storage interface and the client interface. The FSVM service that talks to the CVMs uses the storage interface, which also provides access to Nutanix Volumes iSCSI vDisks in volume groups. The storage interface helps manage deployment, failover, and maintenance, and enables control over one-click upgrades. Integration with the CVM lets the FSVM determine if a storage fault has occurred and if so, whether you must take action. The FSVM service sends a heartbeat to its local CVM service each second indicating its state.

Tip: We recommend placing the FSVM storage interface on the same network VLAN as the Nutanix CVM management interface (eth0 of the CVM). Using the same VLAN helps ensure maximum performance for the iSCSI sessions.

The client interface allows clients to connect to SMB shares and NFS exports hosted on an FSVM. A client can connect to any FSVM client network interface to access their file data. If a different FSVM provides the data, the client connection automatically redirects to the correct

FSVM interface. If an FSVM fails, the client network address for the failed FSVM moves to another to preserve data access.

Storage

As shown in the following figure, each FSVM uses three separate vDisks: a 12 GiB boot disk that contains the boot image, a 45 GiB disk (/home/nutanix) that contains logs and software state, and a 45 GiB disk for Cassandra data.

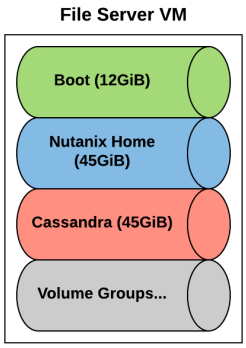

Figure 5: FSVM vDisks and Volume Groups

A Nutanix Files cluster is a single namespace that includes a collection of file systems used to support the SMB and NFS shares. These file systems use Nutanix volume groups as storage. Volume groups enable both high availability and scale-out for the file systems. A volume group is a collection of logically related vDisks (or volumes) attached to the guest via iSCSI. When an FSVM is down for maintenance or if a fault occurs, one of the surviving FSVMs takes over volume group ownership and continues servicing requests.

With the 2.x release of Nutanix Files, volume groups contain six virtual disks: two disks used for metadata and four disks used for data. The four data disks represent 40 TB of available storage for a given share using that volume group.

With the 3.x release of Nutanix Files, volume groups contain 16 virtual disks. Files 3.x versions prior to version 3.2 use 2 disks for metadata and 4 disks for data, as in the 2.x release, with the remaining disks reserved for expansion in later versions. From the 3.2 release of Files onward, all disks in the volume group are available for use by the assigned shares, representing a maximum standard share size of 140 TB.

Figure 6: FSVM Volume Group with 16 vDisks

The initial file system pool created on the volume group uses six of the drives as previously noted and supports 40 TB of data.

Figure 7: Initial 40 TB File System Pool

Once this file system pool reaches 80 percent space utilization, Nutanix Files automatically expands it to 80 TB by incorporating another four data disks from the volume group.

Figure 8: 80 TB File System Pool Following First Expansion

This expansion process can occur two more times, increasing the file system capacity to 120 TB with the second expansion and 140 TB with the final expansion. Each expansion is triggered at 80 percent space utilization of the current pool.

Figure 9: 120 TB File System Pool After Second Expansion

Figure 10: 140 TB File System Pool After Final Expansion

4.2. Exports and Shares

There are two types of SMB shares and NFS exports.

SMB shares:

1. Distributed (previously called home).
2. Standard (previously called general).

NFS exports:

1. Distributed (previously called sharded).
2. Standard (previously called nonsharded).

A standard share is an SMB share or NFS export hosted by a single FSVM. A distributed share spreads the workload by distributing the hosting of top-level directories across FSVMs, which also simplifies administration.

Distributed shares and exports don't allow files in the root.

Distributed Shares

Distributed shares and exports distribute data by dividing the top-level directories across all the FSVMs that make up the file server. Nutanix Files maintains the directory mapping for each responsible FSVM using InsightsDB. FSVMs use DFS referrals for SMB and junctions for NFSv4 to make sure that clients can connect to the right top-level directories.

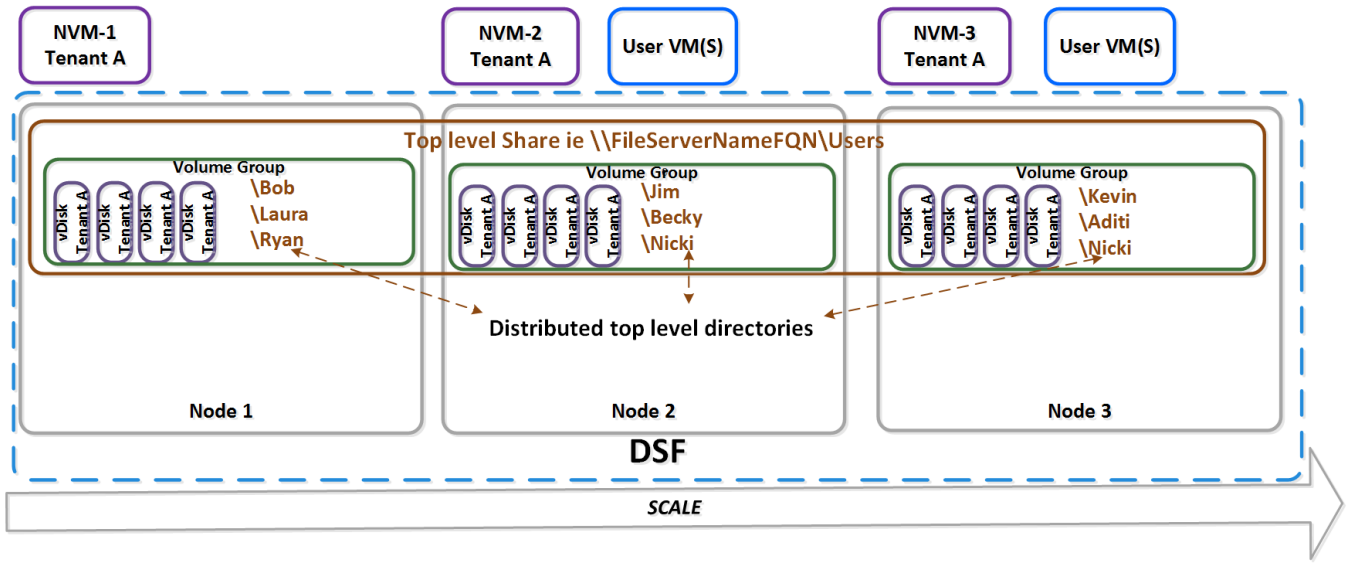

Figure 11: Distributed Directory Shares

Distributed share directories work well for home shares and exports because Nutanix Files automatically spreads the workload over multiple FSVMs per user (see the preceding figure, Distributed Directory Shares). If a user creates a share called "\\FileServer1\Users" that contains top-level directories \Bob, \Becky, and \Kevin, \Bob may be on FSVM-1, \Becky on FSVM-2, \Kevin on FSVM-3, and so on. The FSVMs use a string hashing algorithm based on the directory names to distribute the top-level directories.

This distribution can accommodate a very large number of users or directories in a single share or export. The scaling limits of more traditional designs can force administrators to create multiple shares or exports in which, for example, one set of users whose last names begin with A through M run off one controller, and users whose names begin with N through Z run off another controller. This design limitation leads to management overhead headaches and unnecessary Active Directory complexity. For these reasons, Nutanix Files expects to have one

SMB distributed share for the entire cluster. If you need to have more than one home directory share, you can create additional shares as needed.

The top-level directories act as a reparse point—essentially a shortcut. Consequently, administrators must create directories at the root of the share for optimal load balancing. Because it appears as a shortcut, we don't allow user files in the root level of the share; we recommend setting permissions at the share or export root before deploying user folders. This step allows newly created top-level directories to inherit permissions, so you don't have to adjust them after the fact using the Nutanix Files Microsoft Management Console (MMC) plug-in.

Standard shares and exports do not distribute top-level directories. A single file server always owns the files and subfolders for standard shares and exports. The following diagram illustrates two standard shares (for example, accounting and IT) on the same file server.

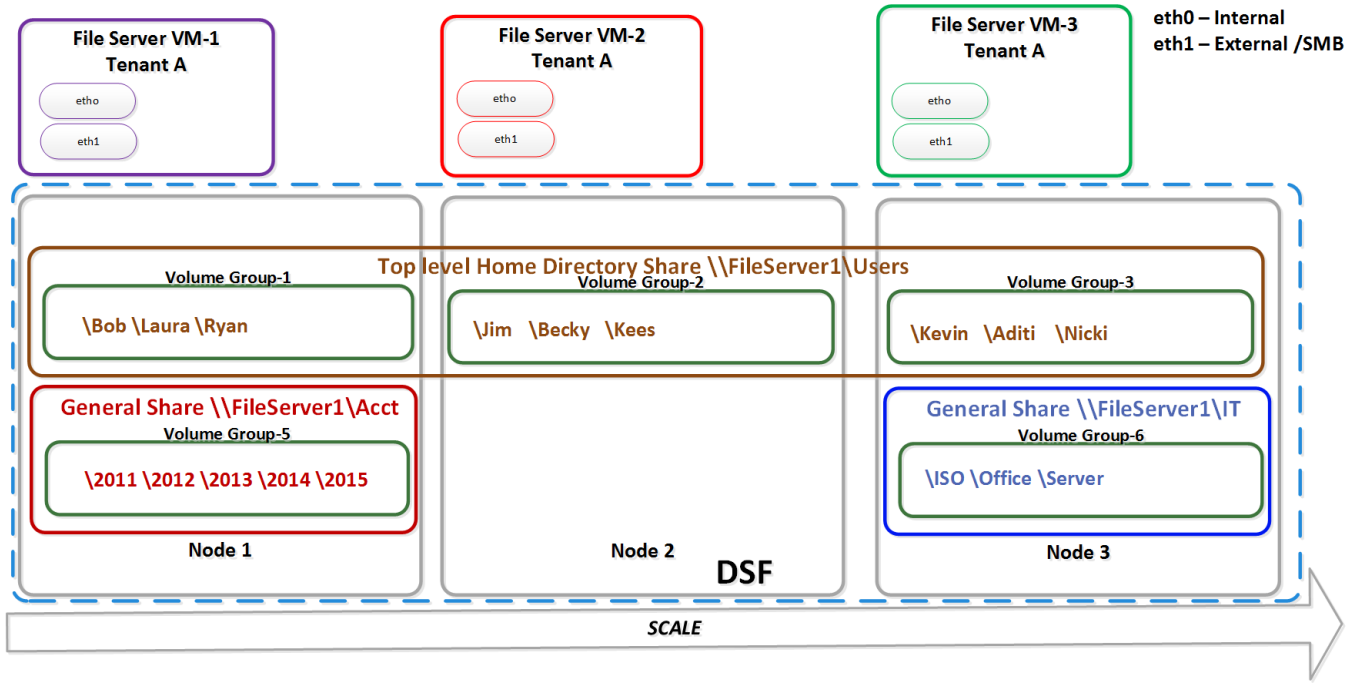

Figure 12: Two Standard Shares on the Same File Server

Unlike distributed shares and exports, standard shares and exports can store files in the root of the directory.

Managing NFS Distributed Share Directories

Distributed share directories with NFSv4 introduce some unique behaviors. To balance performance across FSVMs, each top-level directory you create becomes an automatically generated export. Nutanix Files mounts the export on demand when the directory is accessed. Since these are exports, rather than standard directories, it takes a few steps to remove a top-

level directory. Below we provide an example to demonstrate how distributed top-level directories behave and then walk through the process of deleting a top-level directory.

If you create a distributed share and mount it as **/projects**, you can see that mount on Linux using the **df** command.

```
# df /projects
Filesystem 1K-blocks Used Available Use% Mounted on
1.1.1.10:/projects 1073741824 839455744 234286080 79% /projects
```

Next, create some project directories and access them.

```
mkdir /projects/project1
mkdir /projects/project2
ls /projects/project1
ls /projects/project2
```

Accessing the directory, using **ls** in this case, mounts the automatically created top-level directory export. If you run **df** again you see two additional mount points.

```
# df | grep project
1.1.1.10:/projects 1073741824 839455744 234286080 79% /projects
1.1.1.11:/projects/project1 1073741824 839455744 234286080 79% /projects/project1
1.1.1.12:/projects/project2 1073741824 839455744 234286080 79% /projects/project2
```

These additional mount points allow a different FSVM to serve each export.

This behavior introduces additional steps when deleting a top-level directory. You can delete the project2 directory from any NFS client with the export mounted if you log on as a user with the appropriate permissions. There are three steps to deleting a distributed share export.

- Delete the contents of the share:

```
rm -rf /projects/project2/*
```

- Unmount the project2 share:

```
umount /projects/project2
```

- Delete the project2 directory:

```
rmdir /projects/project2
```

At this point you've deleted the top-level directory and it becomes inaccessible to other clients mounting the export. Processes that access the export after you delete it receive a Stale File Handle error.

Distributed shares with NFSv3 clients behave differently than those with NFSv4 clients. NFSv3 connections do not automatically mount the top-level directories of distributed shares. Because

the top-level directories are not submounted, you can rename and delete them without the additional steps required for NFSv4.

Nested Shares

Nutanix Files 3.2 introduced support for nested shares. A nested share allows you to create a folder within an existing standard or distributed share and turn that folder into a directly accessible share. Both SMB and NFS protocols support nested shares. Nested shares inherit some attributes from the parent share, but you can modify other attributes.

Nested share inherited attributes:

- Protocol type
- Maximum size
- Self-service restore
- Quota policy

Nested share modifiable attributes:

- SMB access-based enumeration (ABE)
- NFS authentication and access
- NFS advanced settings

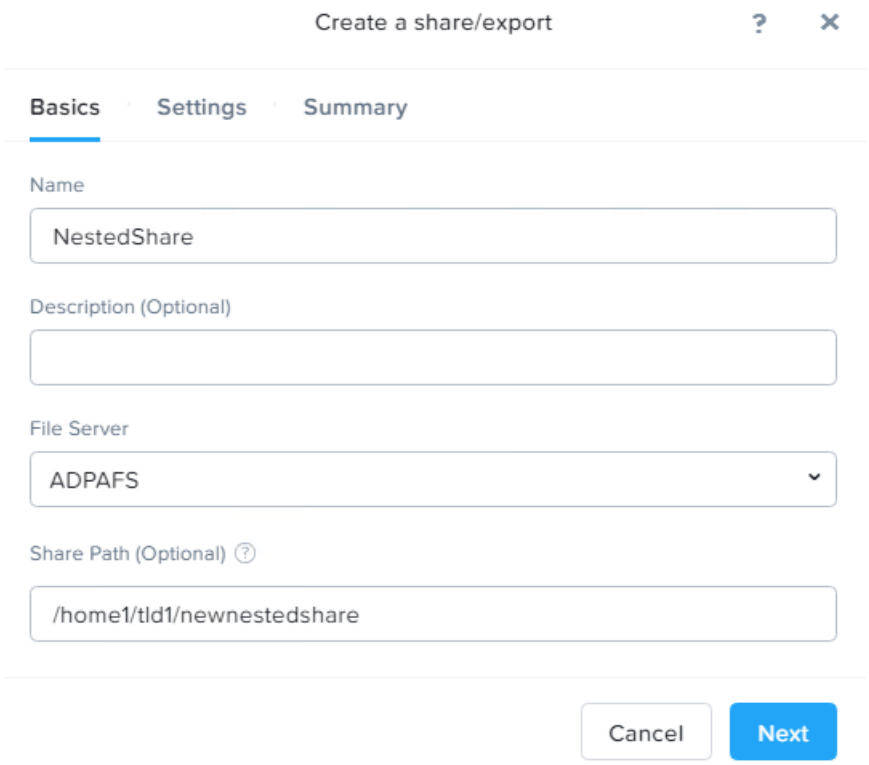

Figure 13: Create Share Menu with Nested Share Path

vDisk and Volume Group Allocation

Distributed shares and exports begin with 5 volume groups for each FSVM (for example, 15 for a three-node cluster). Nutanix Files distributes the volume groups to different FSVMs in the file server cluster. When you have a large number of users, multiple volume groups improve load balancing across the FSVMs.

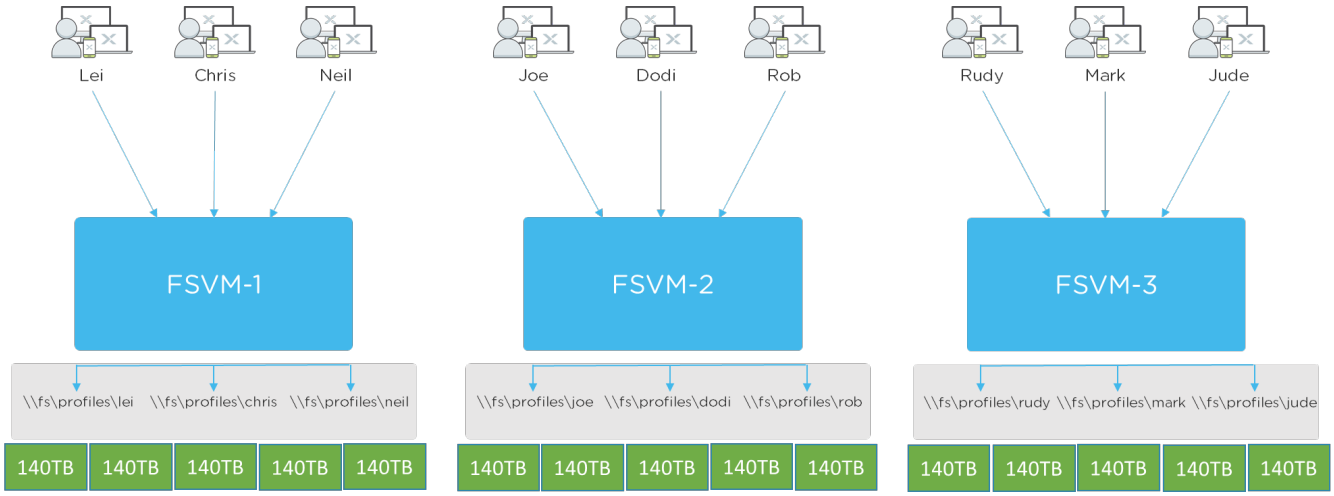

Figure 14: Distributed Share and Top-Level Directories

A single FSVM and volume group serves each standard share or export.

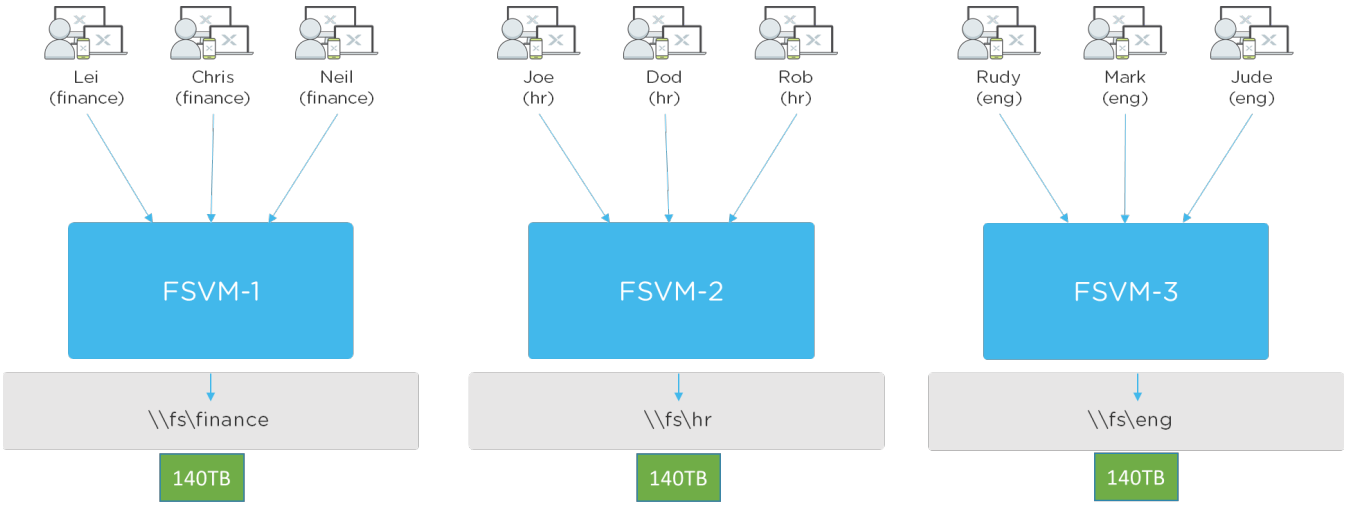

Figure 15: Three Standard Shares

Once you reach the limit of 10 volume groups per FSVM, new shares and exports use existing volume groups of the same type. For example, if you have deployed a distributed share (15 volume groups) and 15 standard shares (creating 15 additional volume groups) on a three-node physical cluster, each FSVM hosts 10 volume groups: 5 for the distributed share and 5 for the standard shares. In this situation, because each FSVM is serving the maximum of 10 volume groups, the next share created uses an existing volume group.

Every file server maps one-to-one to a container. This mapping allows you to manage compression and erasure coding individually for each file server deployed. Inline compression and erasure coding are turned on by default to save capacity. Deduplication is not recommended.

Tip: Use distributed shares for most of your use cases to ensure that user connections and data ownership are distributed across the FSVMs in the cluster.

From Files 3.6 onward, you can control compression on a share-by-share basis. When you create a share, compression is enabled by default.

Directory Layout

Nutanix Files can store millions of files within a single share and billions of files across a multinode cluster with multiple shares. To achieve good response times for high file and directory count environments, be sure to give some thought to directory design. Placing millions of files or directories in a single directory leads to slow file enumeration, which may impact some applications.

For performance-sensitive applications, Nutanix recommends limiting directory width—the number of files or directories in the root of a folder—to 100,000 objects. Increasing FSVM memory to cache metadata can help improve performance for environments with high object counts, where directory enumeration is common.

Nutanix also recommends limiting the number of top-level directories in distributed shares. The recommended number of top-level directories depends on the memory assigned to the FSVMs in the cluster. Your directory count should not exceed 3,000 times the FSVM memory of one node. For example, for a three-node cluster with 12 GB of memory assigned to each FSVM, the maximum top-level directory count is 3,000 x 12 = 36,000.

4.3. Load Balancing and Scaling

Load balancing occurs on two levels. First, a client can connect to any one of the FSVMs and users can add FSVMs as needed. Second, on the storage side, Nutanix Files can redistribute volume groups to different FSVMs for better load balancing across nodes. The following situations prompt load balancing:

1. When removing an FSVM from the cluster, Files automatically load balances all its volume groups across the remaining FSVMs.
2. During normal operation, the distribution of top-level directories becomes poorly balanced because of changing client usage patterns or suboptimal initial placement.
3. When increased user demand necessitates adding a new FSVM, its volume groups are initially empty and may require rebalancing.

Nutanix Files addresses the second and third situations by maintaining usage statistics and patterns to detect per-FSVM load (in terms of CPU and memory utilization) and per–volume

group load (in terms of user connections and latency of operations). Nutanix Files uses these statistics to make a load balancing recommendation, but the administrator must accept the recommendation before Files carries out the action. Nutanix calls this feature one-click performance optimization.

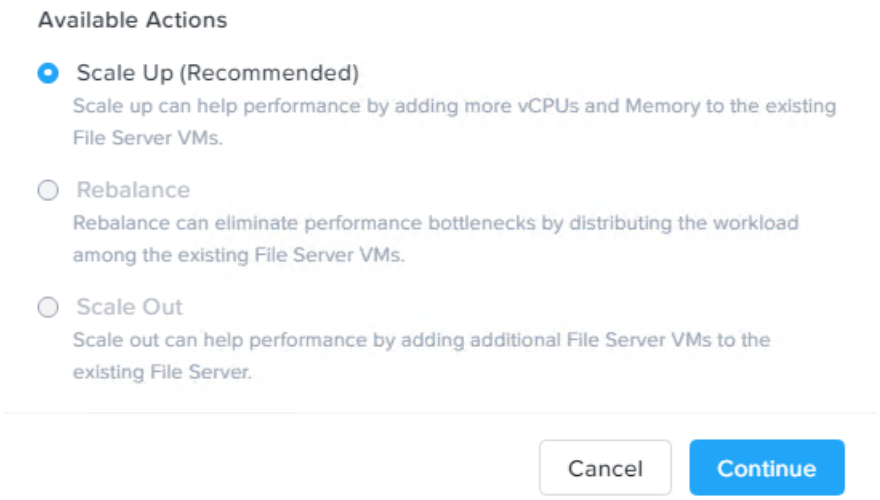

Figure 16: Performance Optimization Recommendation

Load balancing through volume group redistribution may not always improve performance. For example, if clients target a low-level share directory that cannot be further distributed among FSVMs, performance does not improve. In such cases, Nutanix Files supports scaling up the FSVMs by adding vCPUs and memory. Scaling up is seamless to end users.

In the current implementation, there is a brief outage during volume group migration and FSVM scale-out. The file share or export requires a client reconnect after migration and scaling out. When you move a volume group, the time to live (TTL) slowly decreases to limit the disconnection from the client. Today most clients try to reconnect for 50–60 seconds, limiting the overall impact.

Note: With Files 3.6, we recommend a scale-up operation for performance optimization if SMB connection limits reach 95 percent utilization over a two-hour time window. Scale-out and rebalance recommendations are disabled with Files 3.5 and later versions. Nutanix plans to enable scale-out and rebalance recommendations in a later release.

Load balancing also occurs for each virtual disk in a volume group. Nutanix Volumes requires the administrator to configure an iSCSI data service IP. The data service IP is a highly available virtual IP address used as an iSCSI discovery portal. Each virtual disk in a volume group represents its own iSCSI target that any controller VM (CVM) in the cluster can own. Nutanix Volumes uses iSCSI redirection to place and automatically load balance these sessions as

needed. The load balancing feature for Nutanix Volumes is called Acropolis Dynamic Scheduling (ADS).

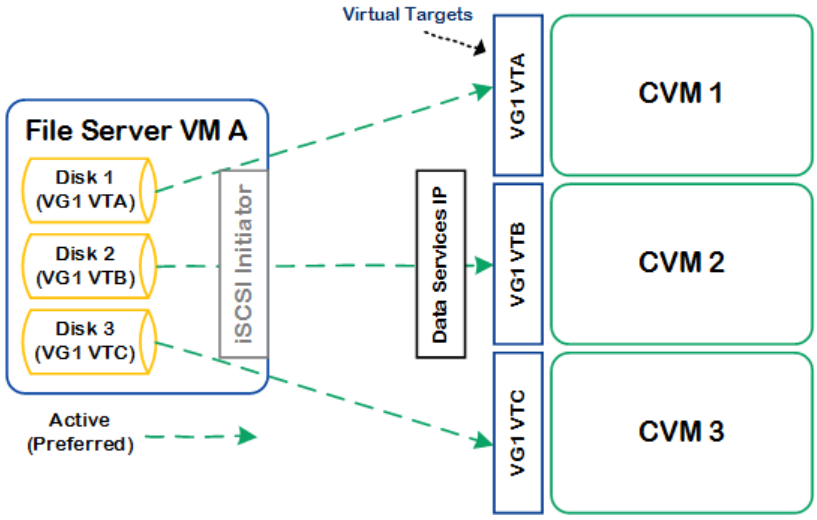

Figure 17: Load Balancing Volume Groups with Nutanix Volumes

For more detailed information, refer to the Nutanix Volumes best practices guide.

4.4. High Availability

Nutanix designed Files to recover from a range of service disruptions, including when a local CVM or FSVM restarts or fails.

CVM Failure or Upgrade

If a CVM goes offline because of failure or planned maintenance, any active sessions against that CVM are disconnected, triggering the iSCSI client to log on again. The new logon occurs through the external data services IP, which redirects the session to a healthy CVM. The following figure shows this general process.

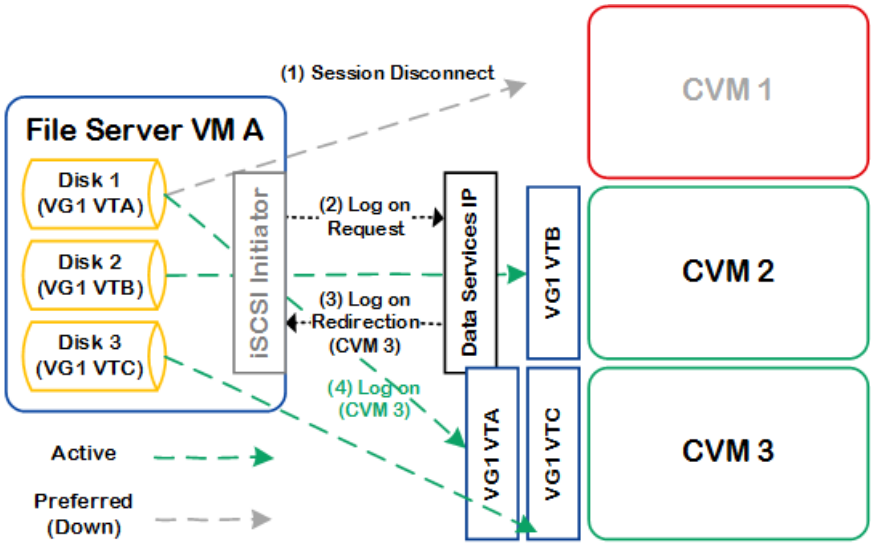

Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups

When the failed CVM returns to operation, the iSCSI session fails back. In the case of a failback, the FSVM is logged off and redirected to the appropriate CVM.

Node Failure

When a physical node fails completely, Nutanix Files uses leadership elections and the local Minerva CVM service to recover. The FSVM sends heartbeats to its local Minerva CVM service once per second, indicating its state and that it's alive. The Minerva CVM service keeps track of this information and can take action during a failover.

When an FSVM goes down, the Minerva CVM service unlocks the files from the downed FSVM and releases the external address from eth1. The downed FSVM's resources then appear on a running FSVM. The internal Zookeeper instances store this information so that they can send it to other FSVMs if necessary.

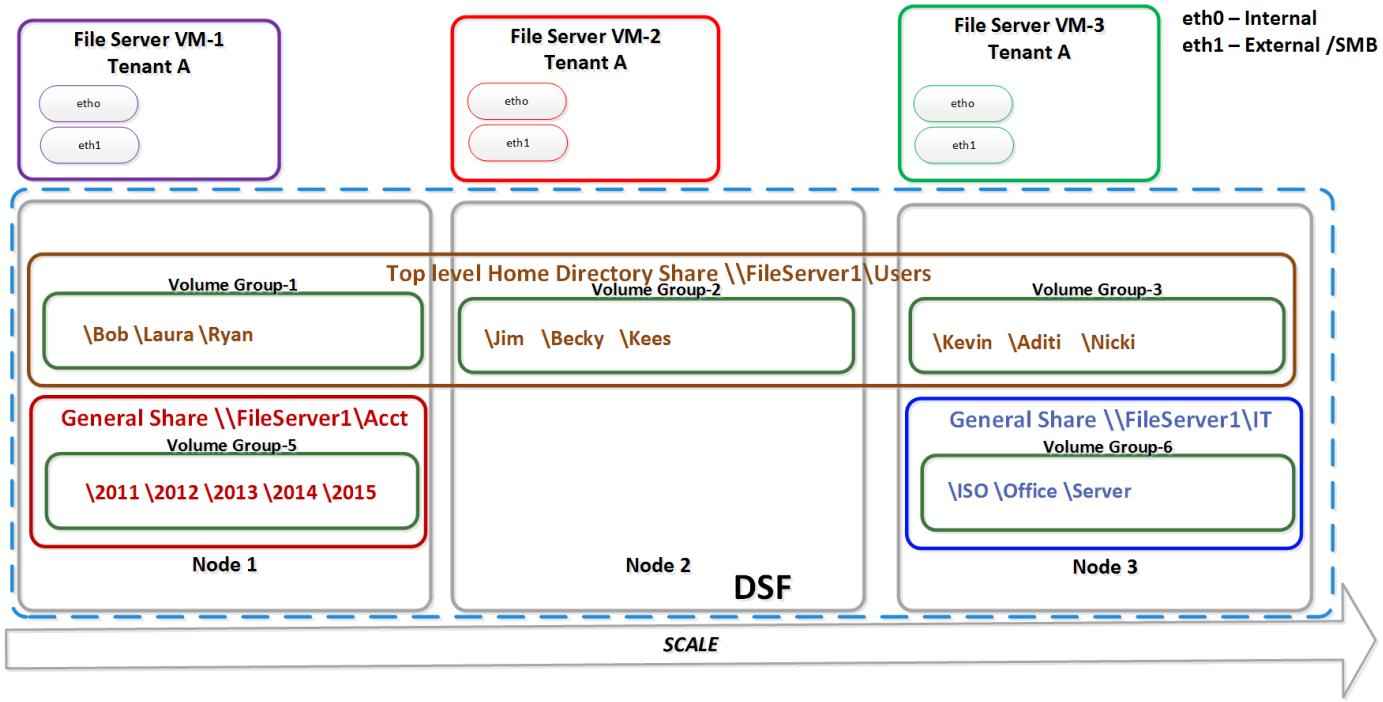

Figure 19: Each FSVM Controls Its Own Volume Groups

When an FSVM is unavailable, the remaining FSVMs volunteer for ownership of the shares and exports that were associated with the failed FSVM. The FSVM that takes ownership of the volume group informs the CVM that the volume group reservation has changed. If the FSVM that attempts to take control of the volume group is already the leader for a different volume group it has volunteered for, it relinquishes leadership for the new volume group immediately. This arrangement ensures distribution of volume groups, even if multiple FSVMs fail.

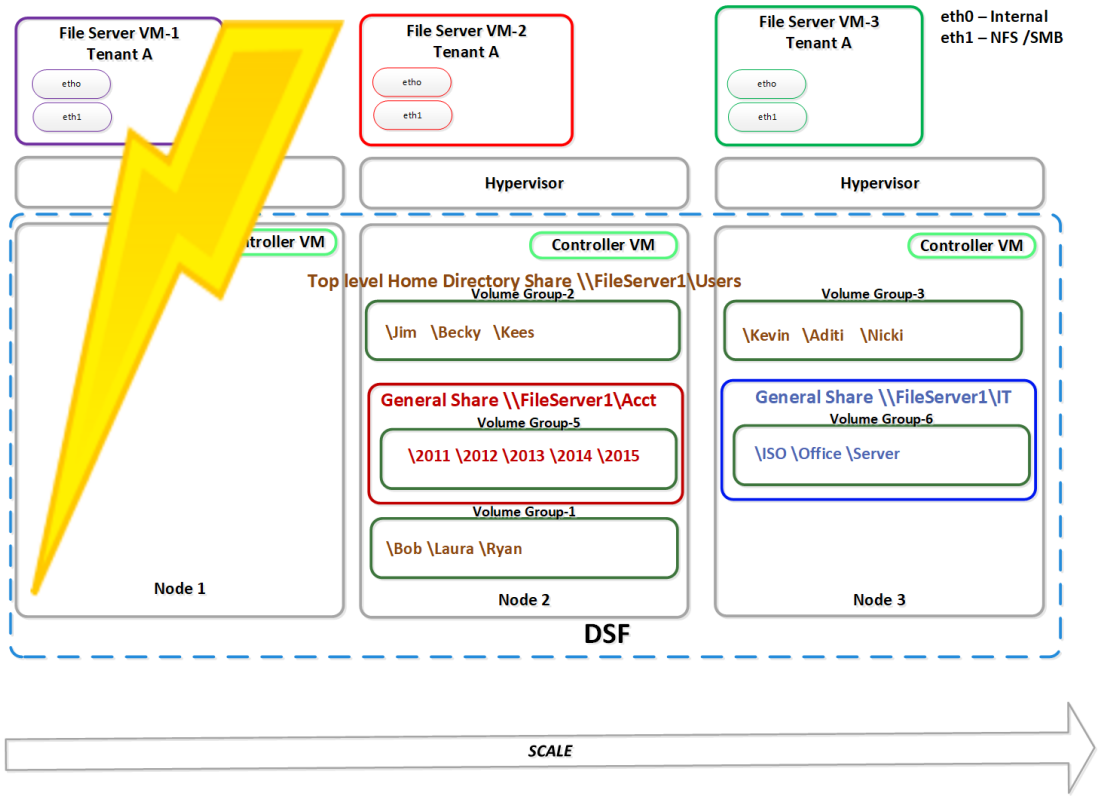

Figure 20: FSVM-1 Failure

The Nutanix Files Zookeeper instance tracks the original FSVM's ownership using the storage IP address (eth0), which does not float from node to node. Because FSVM-1's client IP address from eth1 is now on FSVM-2, client connections persist. The volume group and its shares and exports are reregistered and locked to FSVM-2 until FSVM-1 can recover and a grace period has elapsed.

When FSVM-1 comes back up and finds its shares and exports locked, it assumes that an HA event has occurred. After the grace period expires, FSVM-1 regains control of the volume group through the Minerva CVM service.

To summarize, the process Nutanix Files goes through to reinstate control is:

1. Stop SMB and NFS services.
2. Disconnect the volume group.
3. Release the IP address and share and export locks.
4. Register the volume group with FSVM-1.
5. Present new shares and exports to FSVM-1 with eth1.

A node failure drops the client connection, but the client can recover based on common timeout values.

4.5. Active Directory and SMB Operations

Nutanix Files SMB works with Active Directory. To deploy a cluster, you must have domain privileges; these privileges are necessary to create the machine account and the DNS entries for DFS referrals. The file server does not store these credentials.

Files 3.6 and later versions have simplified the permission requirements to add domains when creating a file server. The delegated domain user permissions required are as follows:

- Create computer objects
- Read servicePrincipalName
- Write servicePrincipalName

Nutanix Files can also create the required DNS entries automatically during file server creation. Automated DNS entry creation requires Microsoft Windows DNS and a user account with DNS admin permissions.

The maximum number of client connections depends primarily on the amount of RAM in the FSVM. See the following table for Nutanix configuration recommendations as of the Files 3.6 release.

Table 2: Supported Active Client Connections

RAM per FSVM	Supported Client Connections per FSVM	Supported Client Connections for Four-FSVM File Server Cluster
12 GB	500	2,000
16 GB	1,000	4,000
24 GB	1,500	6,000
32 GB	2,000	8,000
40 GB	2,750	11,000
60 GB	3,250	13,000
96 GB	4,000	16,000

You can continue deploying additional FSVMs if you have free nodes; you can also deploy multiple file servers.

Nutanix Files uses DFS referrals to direct clients to the FSVM that owns the targeted share or top-level directory. The following diagram shows what happens behind the scenes when a client sends a file access request.

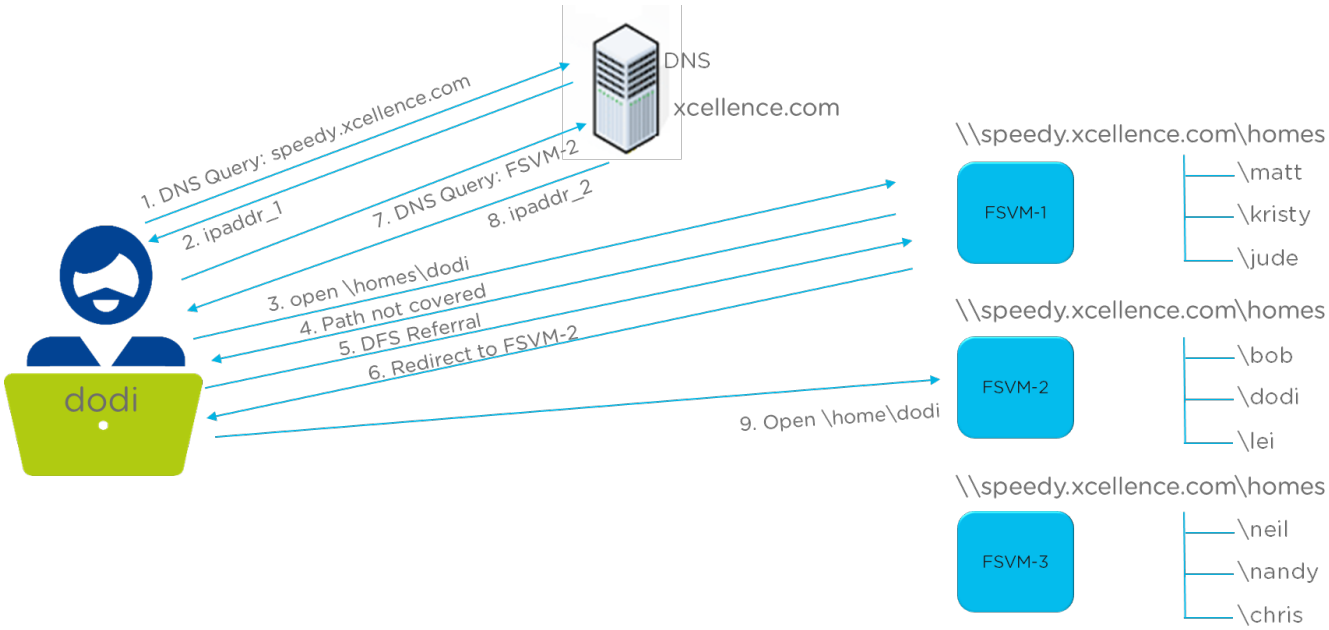

Figure 21: DNS Request for SMB

1. When the user Dodi accesses their files, they click on a shortcut that triggers a DNS request. The DNS request is first sent for the file server name.
2. Using DNS round robin, DNS replies with an FSVM IP address. In this example, the IP address for FSVM-1 returned first.
3. The client sends a create or open request to FSVM-1.
4. The \homes\dodi folder doesn't exist on this file server, so a STATUS_PATH_NOT_COVERED is returned.
5. The client then requests a DFS referral for the folder.
6. FSVM-1 looks up the correct mapping in the file server's Zookeeper and refers the client to FSVM-2.
7. A DNS request goes out to resolve FSVM-2.
8. The DNS request returns the IP address of FSVM-2.
9. The client gets access to the correct folder.

Managing Shares

You can manage Nutanix Files shares much the same way that you manage traditional file servers. For standard shares and top-level directories in distributed shares, you can assign permissions with native tools such as Windows Explorer.

The distributed share has some special requirements because of the way we use DFS referrals. DFS referrals have a single namespace even though the data contained in the share may be spread out over many FSVMs. Typically, to delete a user's home directory, you select the top-level directory and delete the entire directory subtree. However, when using a distributed export or share you cannot simply delete the top-level directory, because it is a separate share created internally by Nutanix Files. Since you cannot directly delete a top-level directory, removing it requires additional steps.

There are three options for renaming or deleting home share folders:

- Find out which FSVM is hosting the folder and delete the folder directly.
- For SMB shares, you can use Nutanix PowerShell scripts (found on the support portal) to handle mass operations.
- Use the Nutanix Files MMC snap-in to manage top-level directories. Any file server administrator can perform the MMC operations; there is no need to assign privileges manually. Establish a connection to the Files namespace, and perform the following SMB share management tasks:
  # Create, delete, rename, and change permissions for top-level directories.
  # Change permissions for shares.

Tip: Use the Nutanix Files MMC when modifying permissions at the root of a distributed share. The Files MMC ensures that permissions are propagated to the top-level directories based on inheritance.

You can manage open file locks and user sessions with the Windows-native Shared Folders MMC. Simply launch the shared folders MMC and point it to the name of the Nutanix Files instance under the **Another computer** option.

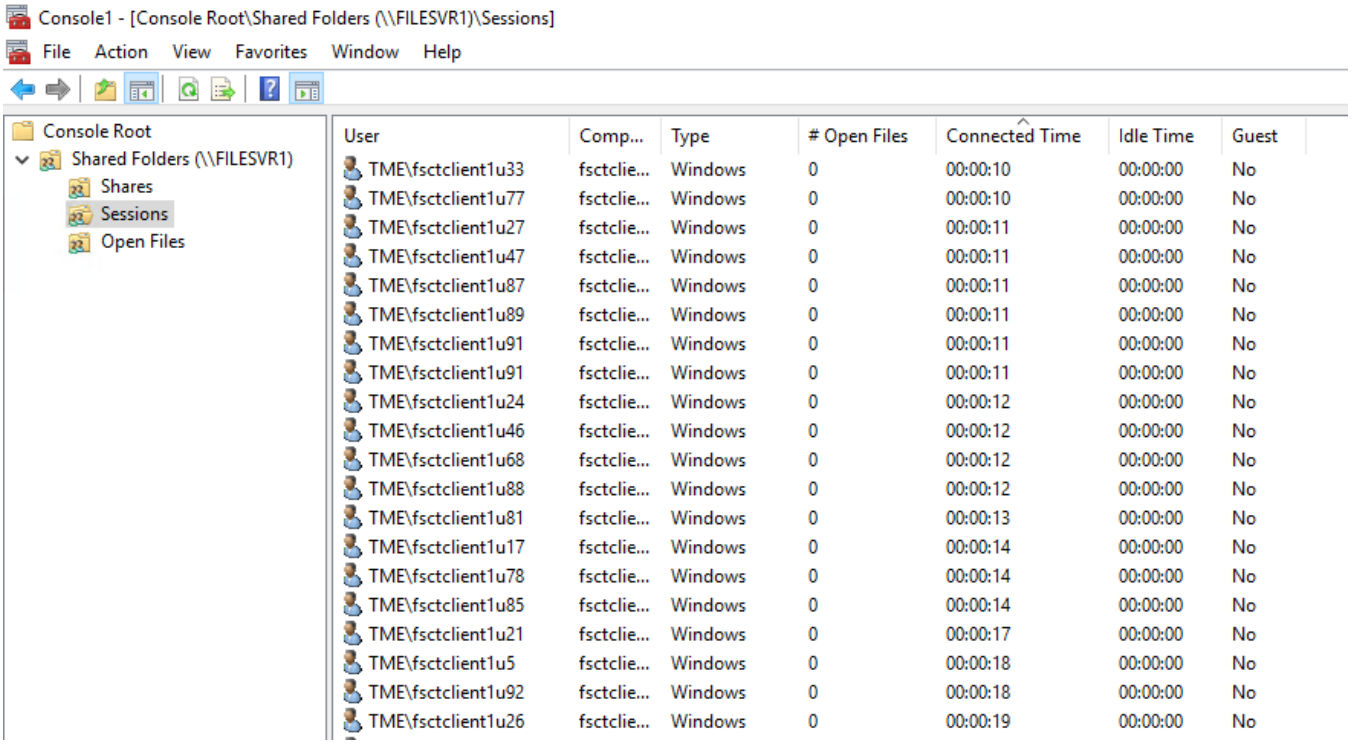

Figure 22: Shared Folders MMC

Access-Based Enumeration

Access-based enumeration (ABE) is a Windows feature (SMB protocol) that filters the list of available files and folders on the file server to include only those the requesting user can access. This filter helps both users and administrators, saving time for the user and worry for the administrator concerned about users accessing files not meant for them.

ABE does not control security permissions, and running ABE has associated overhead. Every time a user requests a browse operation, ABE must filter out objects for which the user doesn't have permission. Even if the user has permission to access all contents of the share, ABE still runs, causing additional CPU cycles and increased latency.

Tip: Home directories, where a distributed share is used, are a great example of where you shouldn't enable ABE. Most users get their home mapping when they log on and always have access to their own contents, therefore we don't recommend enabling ABE for home directory use cases.

SMB Signing

SMB signing is a feature of the SMB protocol that enables a digital signature against each network packet. Digital signing helps ensure the packet's origin and authenticity and prevent tampering, such as "man in the middle" attacks against data in-flight.

Nutanix Files supports SMB signing and honors the signing request as configured from the SMB client. Nutanix Files doesn't require a configuration setting; just enable SMB signing from the client and Nutanix Files negotiates the appropriate setting.

The SMB protocol uses AES-CMAC to compute signatures for SMB signing. This computation process can impact performance against the SMB share where you enabled signing. With Nutanix Files 3.2 and later versions, the signature computation uses the Intel processor AES-NI instruction set. The Intel processor hardware acceleration helps reduce the overhead associated with SMB signing.

Note: SMB signing lowers the maximum performance possible for SMB shares.

SMB Encryption

Nutanix Files 3.6 introduces support for in-flight encryption of SMB data between clients and shares. SMB encryption is disabled by default with Nutanix Files. You can enable encryption as needed on a share-by-share basis.

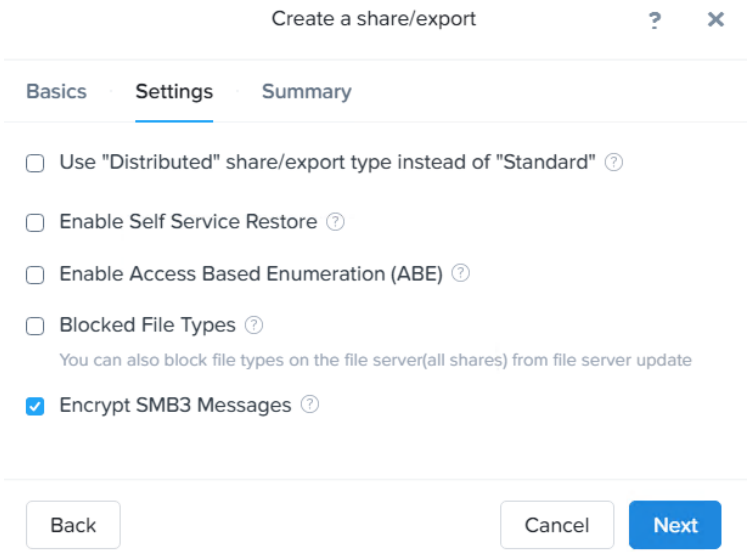

Figure 23: Enable SMB Encryption

Clients must support encryption in order to access shares with encryption enabled. Nutanix Files offloads encryption processing using the Intel processor AES-NI instruction set to limit performance impact.

> **Note:** SMB encryption lowers the maximum performance possible for SMB shares.

Durable SMB File Handles

Durable handles let SMB clients survive a temporary client-side connection loss after opening a file, allowing transparent client reconnection within a timeout. Starting with the 3.6 release, Nutanix Files supports durable handles for SMB clients with no configuration required. SMB 2.x and SMB 3.0 clients can use durable handles to reconnect transparently if there's a client-side network interruption. The system does not maintain durable handles through file server events such as FSVM failures or upgrades.

> **Note:** Durable SMB file handles are not equivalent to SMB 3.0 transparent failover (continuously available shares). Nutanix Files does not currently support SMB 3.0 transparent failover.

DFS Namespace

DFS Namespaces (DFS-Ns) are commonly used to logically organize shared folders on different servers or in different geographic locations into a single namespace. These shared folders appear to the user as a unified hierarchical directory they can navigate using any Windows client. The server names and their locations are completely hidden from the user, enabling large scale-out architectures.

The most common architecture is an Active Directory–integrated DFS-N, where the namespace is hosted on two or more domain controllers (namespace servers) and the file data is stored on member servers. Nutanix Files supports the use of DFS-N where Nutanix Files is used as a member server and Nutanix Files shares are used as folder targets in the namespace. When using Files 3.5.1 or later versions, you can use either distributed shares or standard shares.

DFS-N is also commonly used in active-active replication scenarios. DFS-N allows multiple file servers hosting the same data to support a common folder. DFS-N helps provide site affinity, based on Active Directory, for users to connect to the file server acting as a folder target. You need a replication engine to maintain the data when active-active scenarios are present. Nutanix supports the use of Peer Software as the replication engine for active-active scenarios. To learn more about the Peer Software and Nutanix integration, see Active-Active File Services for the Nutanix Enterprise Cloud.

4.6. Active Directory, LDAP, and NFS Operations

NFSv4

Nutanix Files supports NFSv4, which is more stateful than NFSv3, and includes advanced features such as strongly mandated security and DFS-like referrals. Moreover, most recent distributions of Linux, Solaris, and AIX use NFSv4 as the default client protocol.

Nutanix Files supports Active Directory, LDAP, and unmanaged access to NFSv4 exports. To make the transition from NFSv3 easier, Files does not require administrators to configure Active Directory or LDAP. You can use AUTH_SYS or AUTH_NONE authentication. AUTH_SYS authenticates at the client, just like NFSv3.

There are three different levels of Kerberos authentication when you enable Active Directory or LDAP support. Each of the following options uses Kerberos version 5:

1. krb5, DES symmetric key encryption, and an MD5 one-way hash for Nutanix Files credentials.
2. krb5i, in addition to krb5, uses MD5-based MAC on every request and response.
3. krb5p, on top of krb5 and krb5i, makes the connection between client and server private by applying DES encryption.

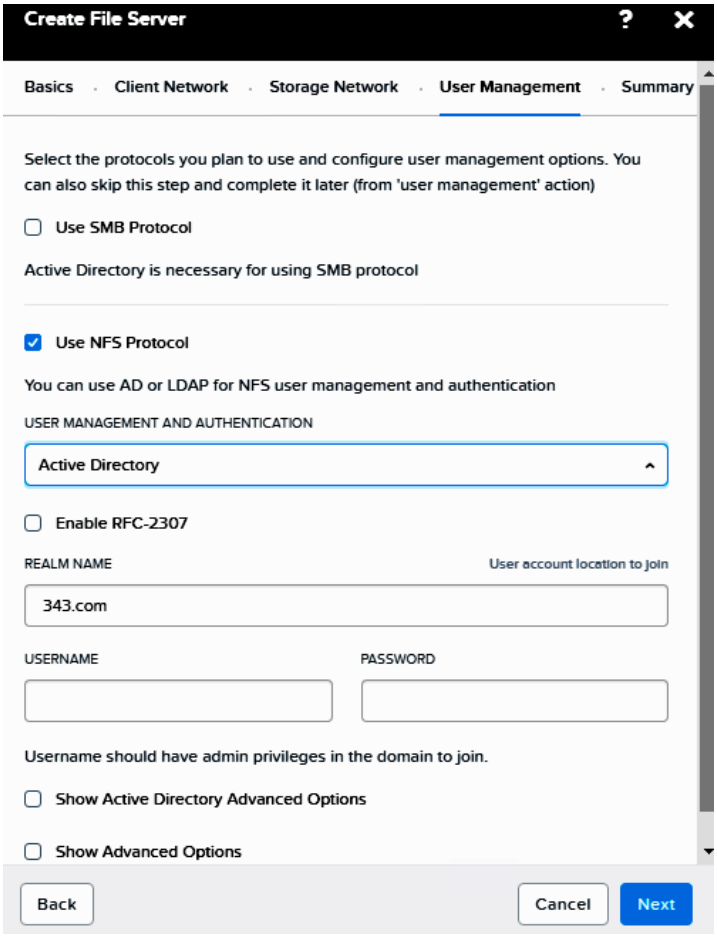

Figure 24: Nutanix Files User Management with NFS

When you deploy Nutanix Files for NFS, you select AD, LDAP, or leave it unmanaged. The following diagram shows what happens behind the scenes when a client sends a file access request using NFS.

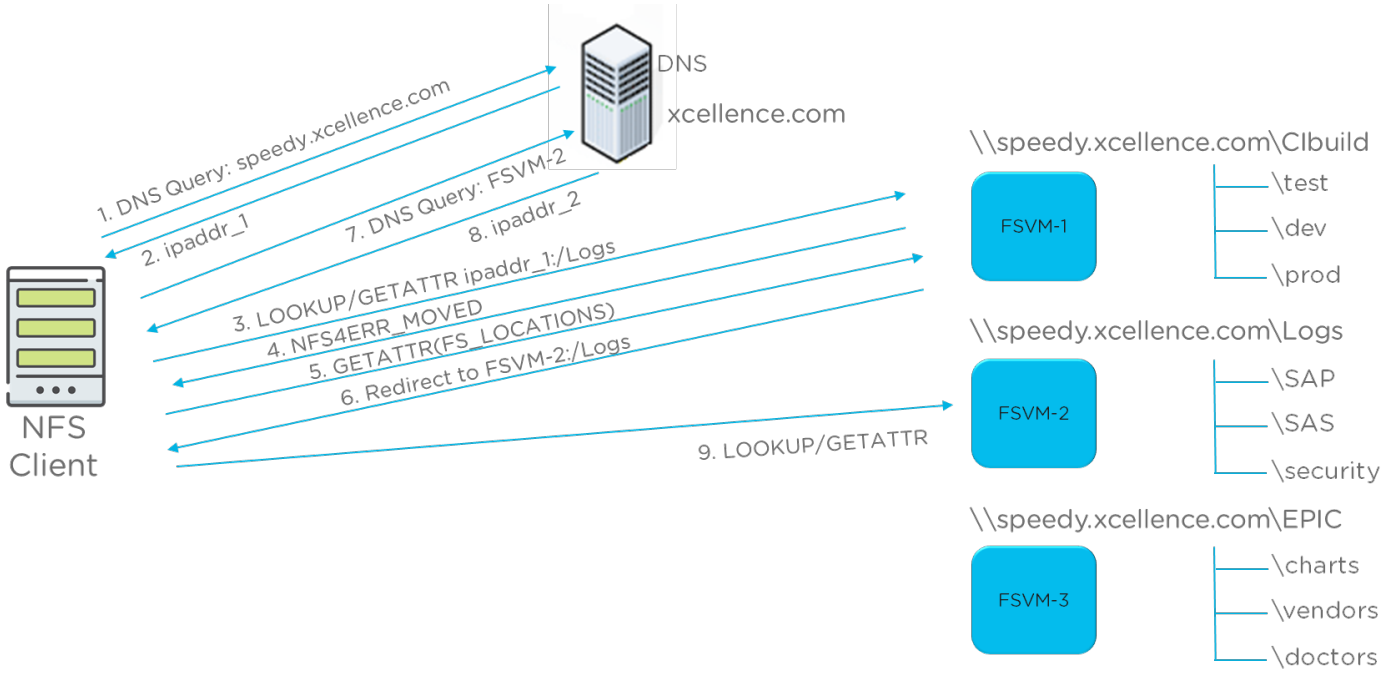

Figure 25: DNS Request for NFSv4

1. When the server wants to access files, the client first sends a DNS request for the file server name.
2. Using DNS round robin, a DNS reply returns with an FSVM address. In this example, the IP address for FSVM-1 returned first.
3. The client sends a create/open request to FSVM-1.
4. The \Logs mount doesn't exist on this file server, so it returns NFS4ERR_MOVED.
5. The client then requests a GETATTR(FS_LOCATIONS).
6. FSVM-1 looks up the correct mapping in the file server's Zookeeper and refers the client to FSVM-2.
7. A DNS request goes out to resolve FSVM-2.
8. The DNS request returns the IP address of FSVM-2.
9. The client gets access to the correct mount point.

NFSv3

With the 3.5 release, Nutanix Files supports NFSv3. NFSv3 is enabled by default for each file server but you can manually disable it. All NFS exports have NFSv3 enabled or disabled based on this file server setting. Files 3.5 supports LDAP and unmanaged exports with NFSv3 but does not support Active Directory and Kerberos for NFSv3.

Note: You must mount exports with the TCP protocol if clients use NFSv3.

Nutanix Files with NFSv3 includes support for both distributed and standard exports. Unlike NFSv4, NFSv3 does not support DFS-like referrals. Clients are not redirected to the FSVMs that host a given distributed top-level directory or standard export; they connect to the first FSVM resolved using DNS. An internal remote procedure call (RPC) manages any file access requests to exports and top-level directories not hosted by the client-connected FSVM. This RPC sends or receives data between the client-connected FSVM and the FSVM that owns the required exports and volume groups.

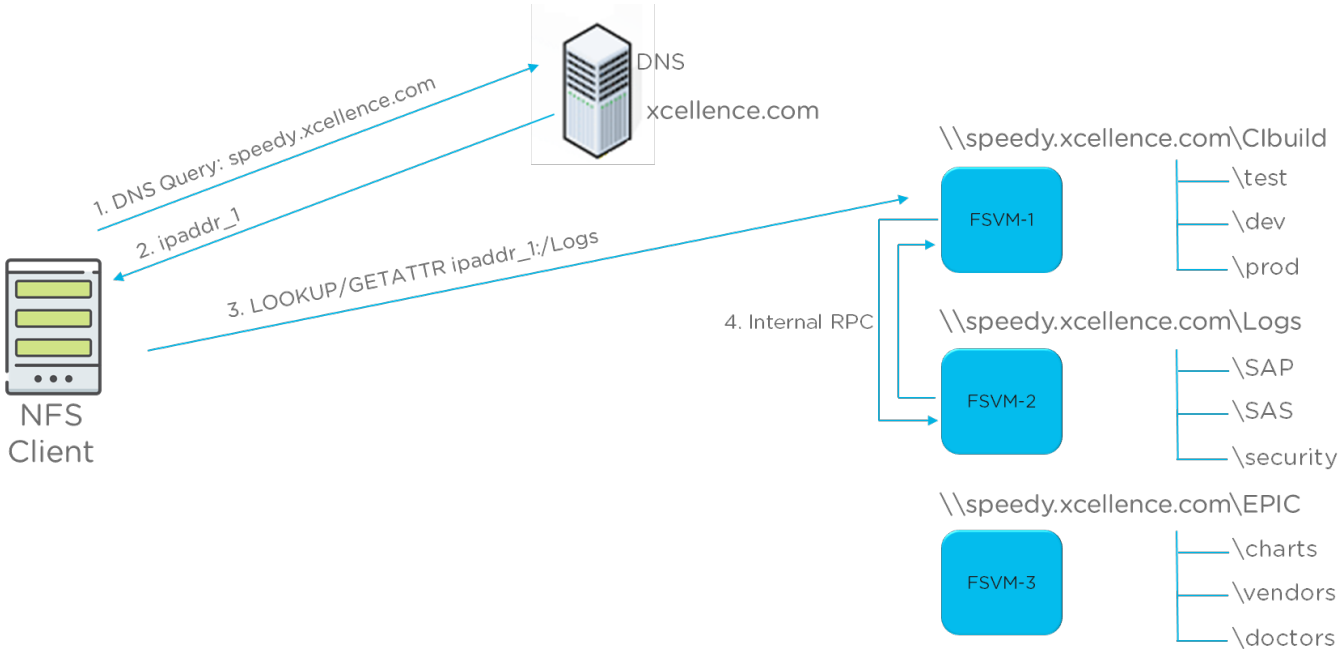

Figure 26: DNS Request for NFS v3

4.7. Multiprotocol

From the Nutanix Files 3.5.1 release onward, you can create file shares that are accessible from both SMB and NFS clients, referred to as multiprotocol-enabled shares. Multiprotocol shares allow simultaneous read access to the underlying file data from either protocol. Write access can also occur from either protocol, but not simultaneously to the same file. Authentication support for multiprotocol shares includes all protocols and authentication options available with Nutanix Files (such as Active Directory for SMB and LDAP, AUTH_SYS, AUTH_NONE, and Active Directory support for NFS).

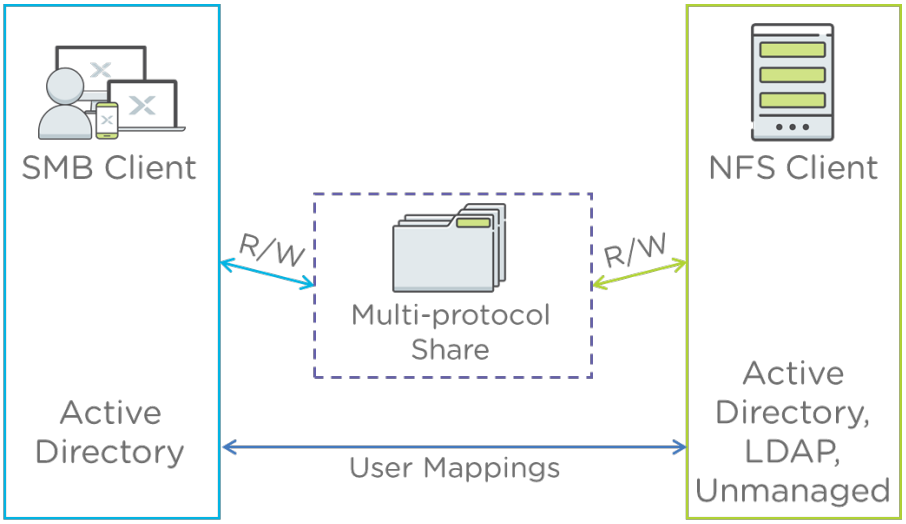

Figure 27: Multiprotocol Shares

Multiprotocol introduces the concept of a native protocol (either SMB or NFS) for a given file share. You specify the native protocol and non-native protocol when you create the share.

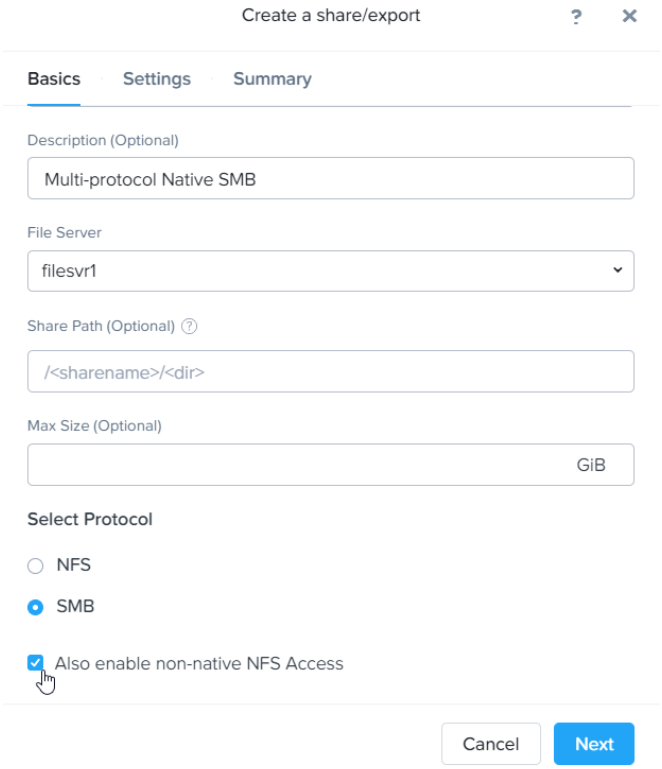

Figure 28: Enable Non-Native Protocol Access

You manage all access control for a multiprotocol-enabled share using the native protocol. For SMB, the native protocol is Windows access control lists (ACLs), and for NFS, it's Unix mode bits. When non-native protocol access occurs, Nutanix Files maps user access to the permission applied with the native protocol. A mapping must exist between the user accounts using the non-native protocol and the user accounts with permission applied through the native protocol.

If you use Active Directory for both SMB and NFS with Kerberos, you don't need any explicit mappings because both have the same users and groups. You can use Identity Management for Unix (RFC 2307) to map NFS users to your Active Directory users with attributes. For LDAP or unmanaged accounts, or for Active Directory without Kerberos, you must create a mapping between the users and groups in Active Directory. You can use Prism or the NCLI to manage Nutanix Files user mappings.

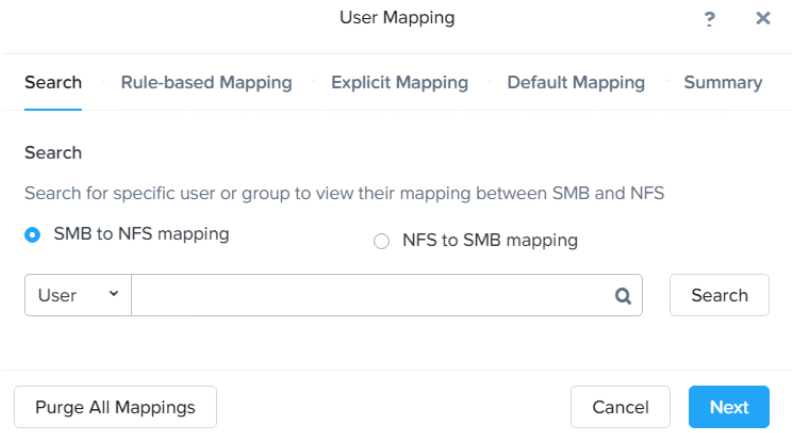

Figure 29: Multiprotocol User Mapping

You can configure a default index to map all non-native users and groups to a specific native user or group. You can also use a rule-based mapping for AD and LDAP users—specifically, a template where the SMB name matches the NFS username. Additionally, you can configure explicit mapping, which overrides rule-based mapping. Explicit mapping consists of two mapping subcategories: one-to-one mapping lists and wildcard mapping. You can use the one-to-one mapping list to manually enter or upload a CSV file that maps users or groups across protocols. Use wildcards for many-to-one mapping. You can also deny share access for a specific user or group. For more details on how to create user mappings, see the user mapping section of the Nutanix Files guide. The following table shows the user mapping requirements depending on the chosen directory service and authentication type.

Table 3: User Mapping Requirements

SMB Directory Service	NFS Directory Service	Export Authentication Type	Supported	User Mapping Required
AD	AD	Kerberos5*	Yes	No
AD	AD	System	Yes	Yes (Name to ID mapping)
AD	AD	None	Yes (Primary NFS only)	Yes (Name to ID mapping)
AD	AD + RFC2307	Kerberos5*	Yes	No

SMB Directory Service	NFS Directory Service	Export Authentication Type	Supported	User Mapping Required
AD	AD + RFC2307	System	Yes	Yes (Name to ID mapping)
AD	AD + RFC2307	None	Yes (Primary NFS only)	Yes (Name to ID mapping)
AD	LDAP	System	Yes	Yes (Name to Name mapping)
AD	LDAP	None	Yes (Primary NFS only)	Yes (Name to Name mapping)
AD	Unmanaged	System	Yes	Yes (Name to ID mapping)
AD	Unmanaged	None	Yes (Primary NFS only)	Yes (Name to ID mapping)

4.8. Quotas

Administrators can configure the default, user, and group quota space for any share. The default level is the quota limit for every user unless the administrator specifies otherwise. A user-level quota policy sets a specific amount of storage for a single user. For example, if an administrator allocates 1 GB, the user cannot take more than 1 GB. A group-level quota policy extends a user policy to include all users for an entire Active Directory group, where each user can use the assigned quota value. For example, if the administrator sets a group's quota to 10 GB, then each member of that group can use 10 GB. You can assign quotas for SMB, NFS, or multiprotocol-enabled shares. In the case of multiprotocol shares, administrators apply quotas using the native protocol. User or group policies are enforced for the non-native protocol based on the user mapping.

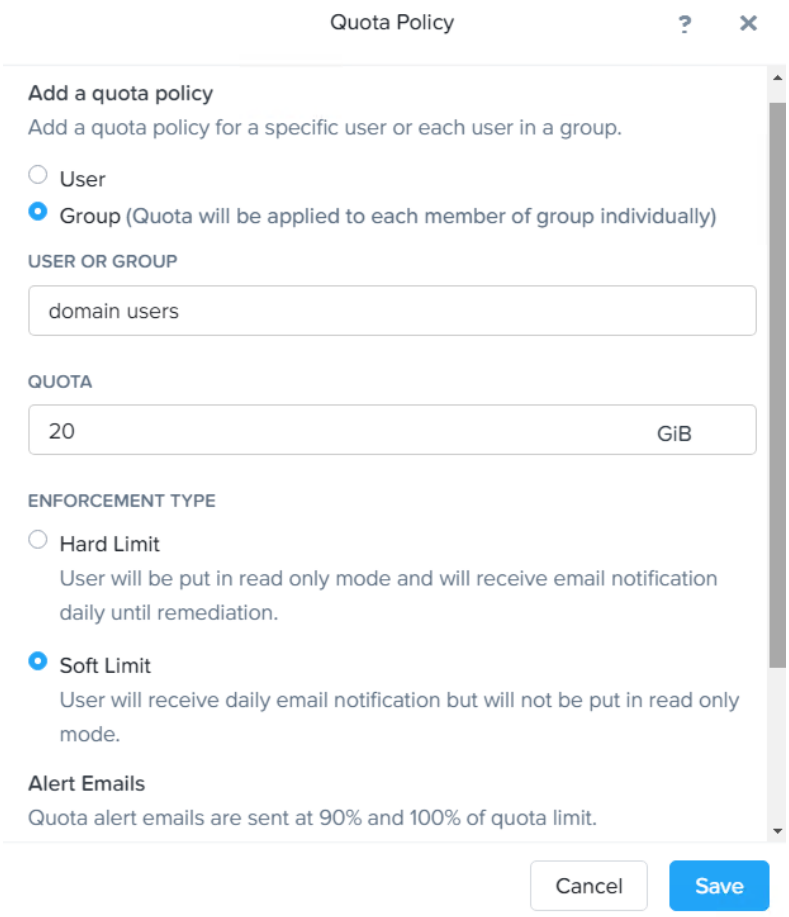

Figure 30: Setting a Share Quota

Notifications

Administrators can configure Prism to send email alerts to the user and to other recipients using the same engine that sends cluster alerts. Designated users receive email notifications when the quota is near maximum consumption—a warning email at 90 percent and an alert email at 100 percent. You can also add departmental share owners to the email notification list so they're aware that they may need to take action.

The following table shows the order of precedence when dealing with quotas for a share.

Table 4: Order of Precedence for Quotas

Order of Precedence	Policy
1	User policy
2	Group policy (the group policy with the highest quota wins)
3	Default user policy

Enforcement

The administrator can also configure enforcement types for each quota-level category. Enforcement types determine if a user or group can continue to use the share after they've consumed their quota. A hard enforcement type prevents the user from writing on the share when they reach their quota limit. If a user or application is writing to the share after it has reached the quota, the operation fails. A soft enforcement type allows a user to write even if they exceed the quota limit. Under either enforcement type, users over their quota receive an email notification every 24 hours until the issue is resolved.

4.9. Selective File Blocking

With Nutanix Files 3.6 you can define a list of file types, based on file extension, to block from being stored on SMB, NFS, and multiprotocol enabled shares. You can define a list of files to block at the server level and at the share level.

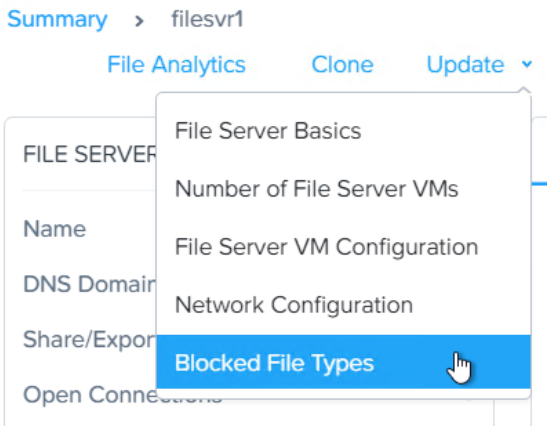

Figure 31: Blocked File Types for the File Server

A comma-separated list of file extensions defines the blocked file types. File types blocked at the server level apply to all shares. When you define a list of file types at the share level, the share-level setting overrides the server-level setting.

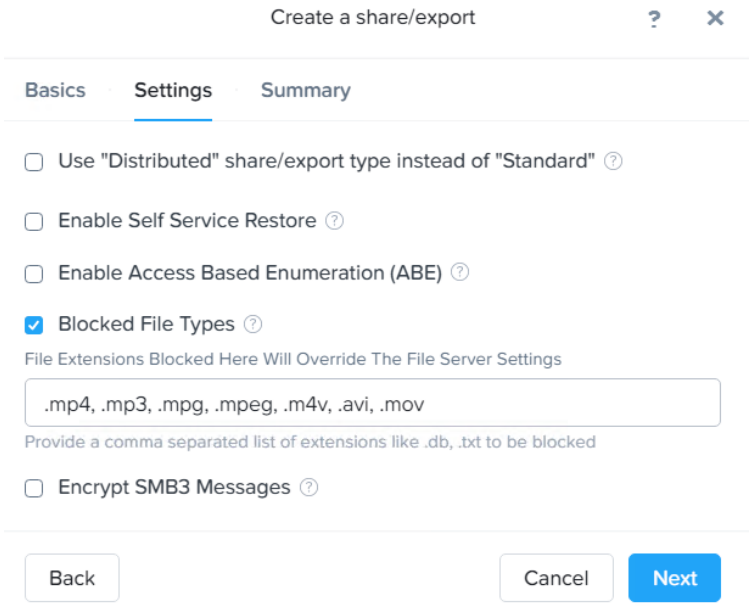

Figure 32: Blocked File Types for a Share

When you attempt to create a file with or rename a file to a blocked extension, an access denied message appears. You can read, edit, or delete existing files, created prior to a blocked file type policy. You can also use File Analytics to discover if any unwanted file types exist on a file server before applying the policy.

4.10. File Analytics

Nutanix Files offers a rich auditing API to which applications can subscribe and receive real-time notifications of file-related events, including file creation, deletion, read, write, and permissions changes. A common use of these APIs is forwarding these events to a syslog server for retention and audit trails. While logging audit trails is useful, simplifying insights into this data is a must. We developed Nutanix File Analytics to consume this native auditing API while delivering additional insights to the underlying data and user activity.

You can deploy File Analytics using Prism: you simply download the Analytics bundle and perform a one-click deployment. Once deployed, launch the File Analytics page and enable analytics against the desired file server instances running on your cluster. Enabling analytics requires a user account with administrative permissions to do an initial scan of the file server.

Intelligent Insights

File Analytics scans and then continuously captures all file activity for registered file server instances. This logging helps to form a repository of information so administrators can review what operations have occurred against specific data and by specific users. Logged events are analyzed to provide an initial dashboard of information:

- Capacity trend, which shows what's being consumed and how it's changed over time.
- Data age, which is the calculation of the last time a file was modified and the percentage of data at varying age ranges.
- Anomaly alerts, which show all file operations that exceed a given anomaly threshold, like the deletion of many files.
- Permission denials, which are the number of permission denial events for specific users over the selected time range.
- File distribution by size, which shows the number of files within a given size range.
- File distribution by type, which details storage consumption by file type, such as log files, image files, and video files.
- The top five active users based on total operations over the selected time period.
- The top five accessed files based on total operations over the selected time period.
- File operations, which detail the most frequent types of operations, like file create, read, or write, over a selected time period, including the trend over that time.

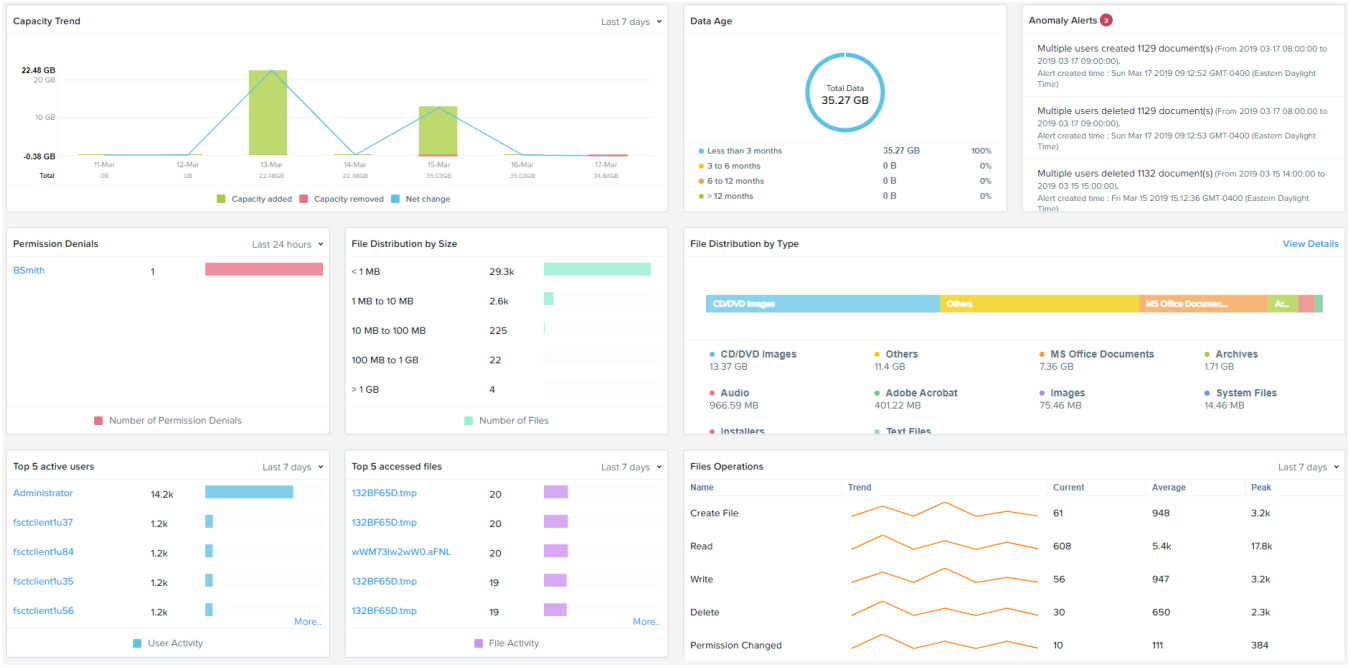

Figure 33: File Analytics Dashboard

The dashboard provides a quick and easy health check so you better understand capacity trends, data age, and file activity that could indicate malicious activity.

Audit Trails

Audit trails allow you to search for a specific file or a specific end user to find all file or user activity for a given timeframe. You can search for a given file or user based on wildcards. The frequency and types of operations, by users and against files, including the time they occurred, are displayed over your specified time period. You can further filter the audit trail based on operation type, such as open, read, write, delete, and other events.

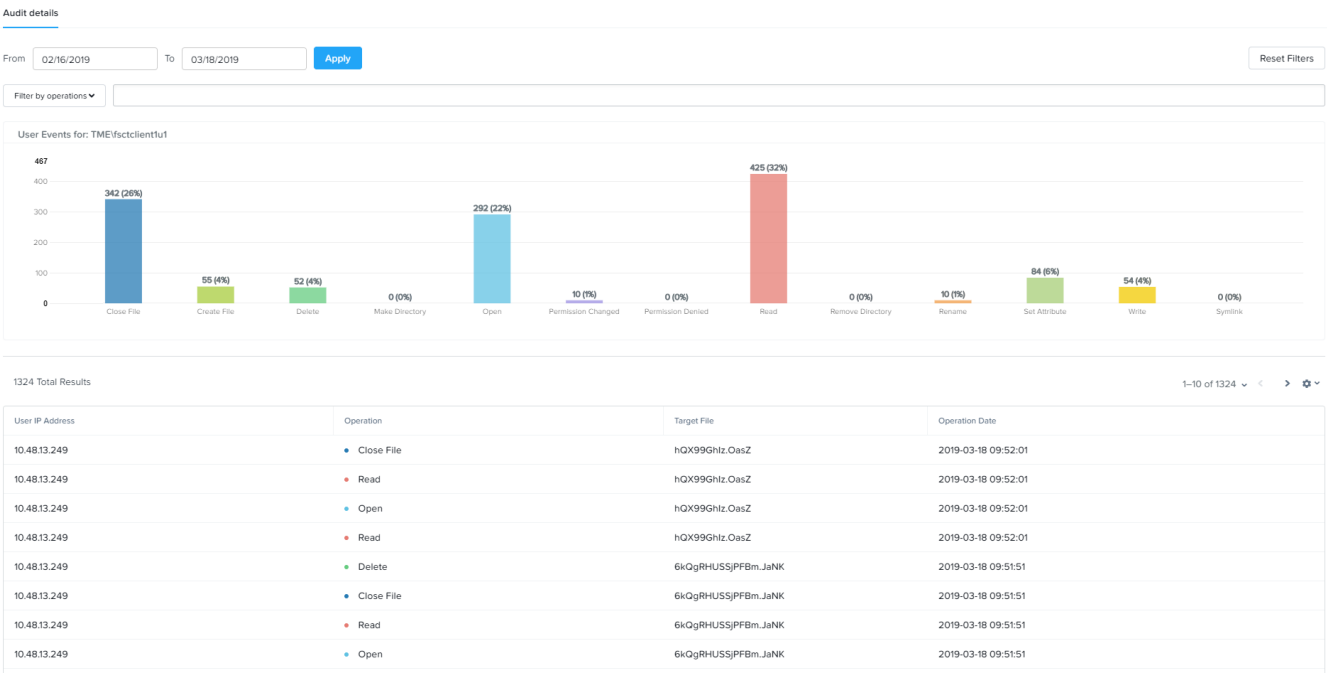

Figure 34: File Analytics Audit Trails

Being able to search user and file activity helps you monitor access activity against sensitive data. You can also download the queried activity to a JSON or CSV formatted files for further reporting activities.

Anomaly Detection

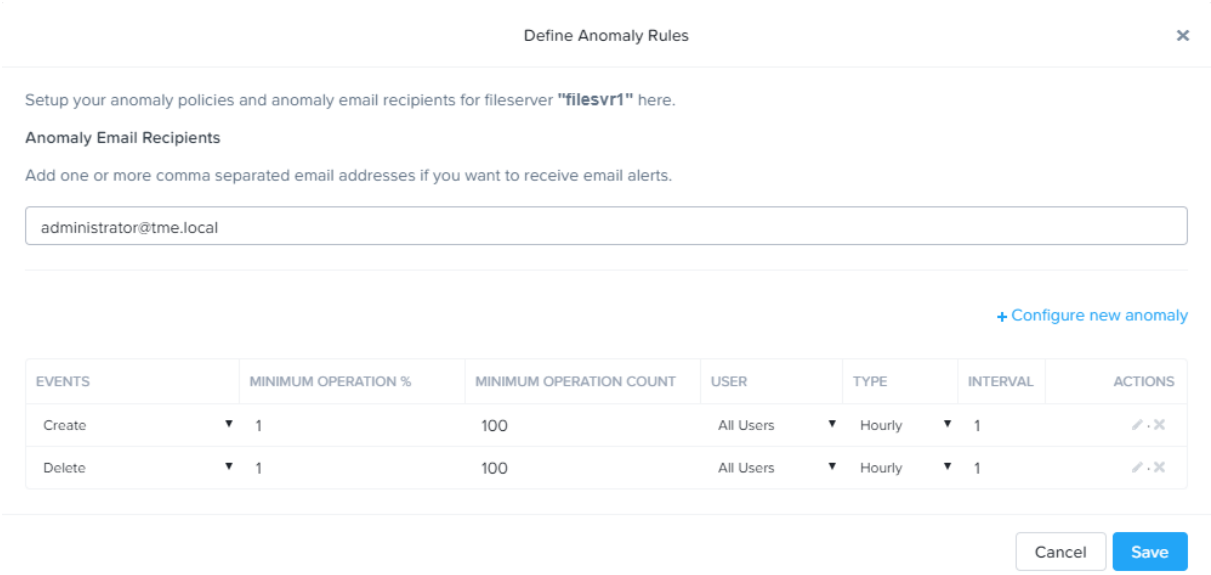

Figure 35: File Analytics Anomaly Detection

You can specify different events, such as permission changes, permission denials, file deletes, and file creates; define an operation percentage and operation count as a part of the alert; specify the interval for when these thresholds apply to the operations in question; and specify email recipients for the anomaly events.

> **Tip:** Minimum operation percentage is based on the number of files. For example, if there are 1,000 files and the minimum operation percentage is 5, it means the defined event (such as create or delete) impacted 50 files within the specified interval.

As anomaly events occur, the email addresses defined in the anomaly rule receive detailed information. You can also monitor each anomaly alert and track it over time with the usage anomalies dashboard. The dashboard shows you the users who caused the anomalies and the folders the anomalies impacted and details the types of anomalies and trends.

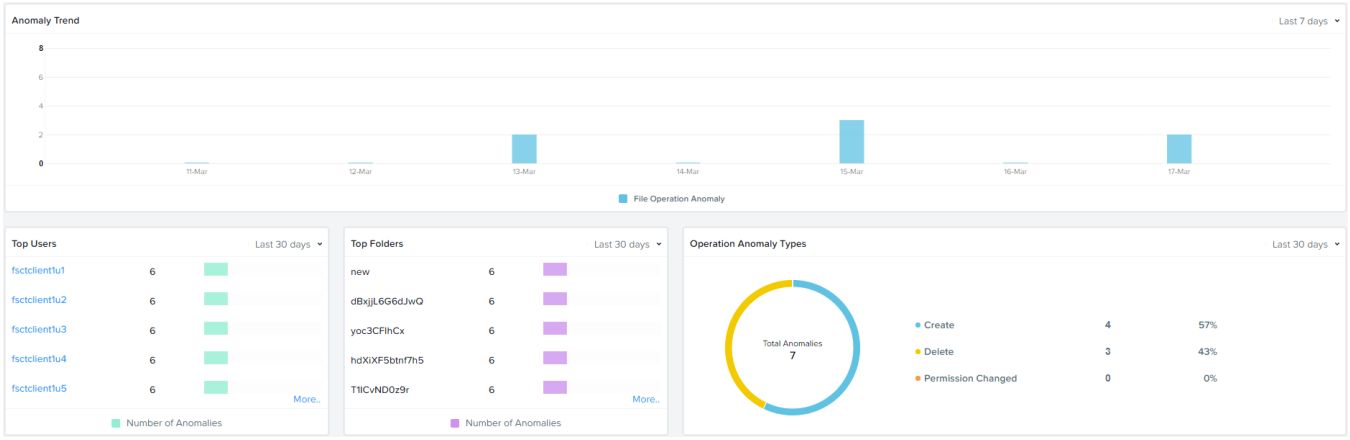

Figure 36: File Analytics Usage Anomalies Dashboard

Strange access patterns that trigger anomalies can indicate malicious user, virus, or ransomware activity. Seeing this activity and being alerted in real time can help you stop these operations and better understand user access patterns across your organization.

4.11. Hypervisor-Specific Support

Nutanix supports ESXi and AHV for Files. For ESXi support, you need vCenter credentials to deploy Nutanix Files and to create DRS rules to make sure the FSVMs are on different nodes. You must register vCenter with the Nutanix cluster instances where you deployed Files. The deployment process generates the DRS rules for AHV automatically.

Data-at-Rest Encryption

Nutanix Files supports data-at-rest encryption (DARE) with self-encrypting drives (SEDs) or software-defined DARE using AOS. Official support for AOS software encryption with Nutanix Files begins with AOS 5.10.2 in combination with Nutanix Files 3.5.

For AHV, you enable software encryption at the cluster level. For ESXi, you enable software encryption at either the cluster level or the storage container level. To enable storage container–level software encryption with Nutanix Files you must use NCLI:

```
<ncli> storage-container edit enable-software-encryption=true name=<Files_Container_Name>
 force=true
```

5. Backup and Disaster Recovery

Following modern data protection methodologies, Nutanix provides administrators and users quick restore access using Self-Service Restore (SSR) and site recovery with Nutanix-based snapshots.

5.1. Self-Service Restore

Administrators can enable SSR at any time for SMB or NFS shares. Windows Previous Version in each folder exposes SSR for SMB shares. SSR for NFS shares is exposed as a hidden snapshot directory in each folder.

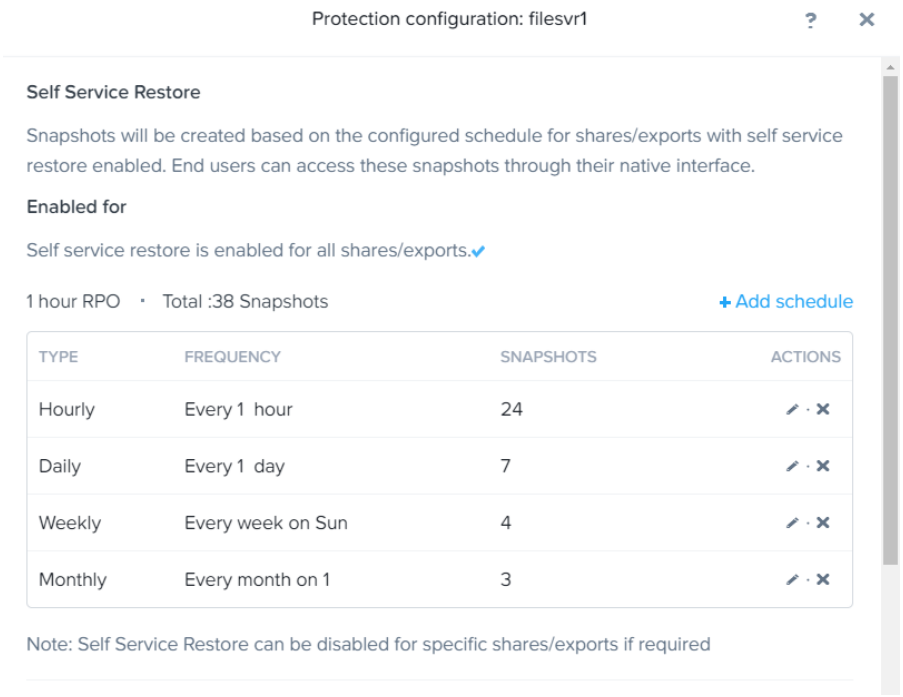

Figure 37: Protection Using SSR

SSR allows you to view snapshots of shares while the share is in use. The share snapshots are read-only, point-in-time copies (that is, copies created at a specified time).

You can view and restore removed or overwritten files, which allows you to choose a share snapshot from the same file at different times during the file's history. Administrators can configure a snapshot schedule at the file server level that applies to all shares in the file server.

Nutanix Files supports 24-hour, daily, weekly, and monthly snapshots on a fixed schedule. The default snapshot policy includes:

- Hourly, 24 per day.
- Once daily for seven days.
- Four weekly.
- Three monthly.

> **Note:** You can configure a maximum of 50 total snapshots in the SSR schedule.

Schedule frequency can vary to suit your requirements, including shorter intervals for same-day protection against accidental deletions. You can enable SSR during or after share creation. After share creation, the administrator can change the current settings using the share update workflow feature. SSR supports both standard and distributed SMB shares.

5.2. Protection Domains and Consistency Groups

Nutanix provides integrated, automated disaster recovery between Nutanix clusters. A Nutanix Files cluster can be protected with Prism and uses the same asynchronous replication with protection domains and consistency groups as any other Nutanix cluster. A protection domain is a defined group of entities (VMs and volume groups) that you back up locally on a cluster and that may replicate to one or more remote sites. A consistency group is a subset of the entities in a protection domain. Consistency groups are configured to snapshot a group of VMs or volume groups in a crash-consistent manner.

> **Note:** Nutanix Files supports both asynchronous (RPO of one hour or longer) and NearSync (down to one minute RPO) schedules. NearSync support requires a minimum of AOS 5.11.1 and Files 3.6.

When you create a file server, Prism automatically sets up a corresponding protection domain, which it annotates with the Nutanix Files cluster name. Prism also creates multiple consistency groups in a protection domain, including a group that includes all FSVMs.

Once you've protected Nutanix Files, all future operations on it (such as adding or removing FSVMs or adding or deleting volume groups) automatically update the corresponding consistency group in the protection domain.

5.3. Cluster Migration, Failure, and Restoration

In the event of a Nutanix Files cluster failure, restore Files in a remote cluster by initiating the **Activate** workflow, which restores from the last good snapshot. If you're moving your file services because you need to shut the cluster down, as with a planned outage, the **Migrate** workflow

shuts down all the FSVMs and takes a final snapshot for replication. You run the Migrate workflow from the Nutanix cluster that owns the active protection domain. You initiate both the Activate and Migrate workflows from the **Data Protection** menu in Prism.

After you run either the Activate or Migrate workflow, you also need to activate the file server instance, which you can do from the **File Server** menu in Prism. Activating the file server may require you to configure network VLANs on the replica site before Nutanix Files becomes operational again. These VLANs can be in subnets that are the same as or different from the source site. The recovery process is like creating the file server, but in this context you can change networks and IP addresses if necessary.

5.4. Cloning

Because Nutanix Files cloning does not affect the original Files cluster, it offers improved support for a variety of use cases, including:

- Backups at the primary and secondary sites.
- DR at the secondary site.
- File server recovery from a specific point in time.
- File server creation at the primary or remote site for testing or development.
- File server clone copies.

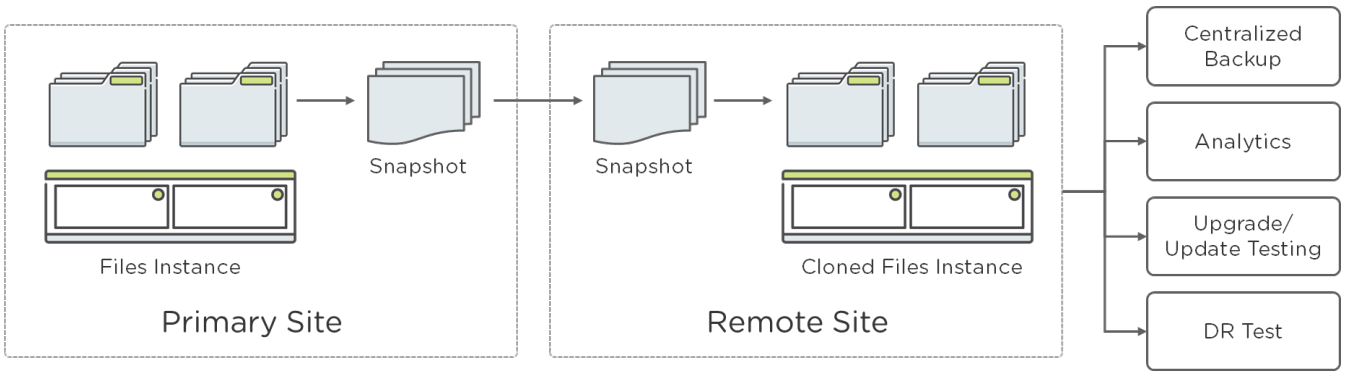

Figure 38: Nutanix Files Cloning Use Cases

Files uses Nutanix native snapshots to clone entire file servers. The clone is a thin copy that consumes minimal storage space. File server clones reside on the same container as the original and maintain the original security permissions. During the clone process you can specify new IP addresses and give the cloned file server a new name.

6. Third-Party Integration

6.1. Antivirus

To protect users from malware and viruses, you need to address both the client and the file server. Nutanix currently supports third-party vendors that use Internet Content Adaptation Protocol (ICAP) servers. ICAP, which is supported by a wide range of security vendors and products, is a standard protocol that allows file and web servers to be integrated with security products. Nutanix chose this method to give customers wide latitude in selecting the antivirus solution that works best for their specific environment.

Following is the workflow for an ICAP-supported antivirus solution:

1. An SMB client submits a request to open or close a file.
2. The file server determines if the file needs to be scanned, based on the metadata and virus scan policies. If a scan is needed, the file server sends the file to the ICAP server and issues a scan request.
3. The ICAP server scans the file and reports the scan results back to the file server.
4. The file server takes an action based on the scan results:
   a. If the file is infected, the file server quarantines it and returns an “access denied” message to the SMB client.
   b. If the file is clean, it returns the file handle to the SMB client.

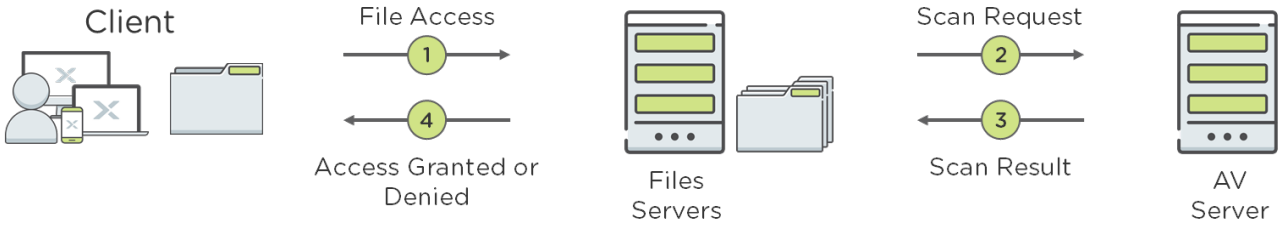

Figure 39: ICAP Workflow

The ICAP service runs on each Nutanix Files file server and can interact with more than one ICAP server in parallel to support horizontal scale-out of the antivirus server. The scale-out nature of Files and one-click optimization greatly mitigates any antivirus scanning performance overhead. If the scanning affects Nutanix Files file server VM performance, one-click optimization recommends increasing the virtual CPU resources or scaling out the file server VMs. This feature also allows both the ICAP server and Files to scale out, ensuring fast responses from the customer’s antivirus vendor.

Tip: We recommend configuring two or more ICAP servers for production.

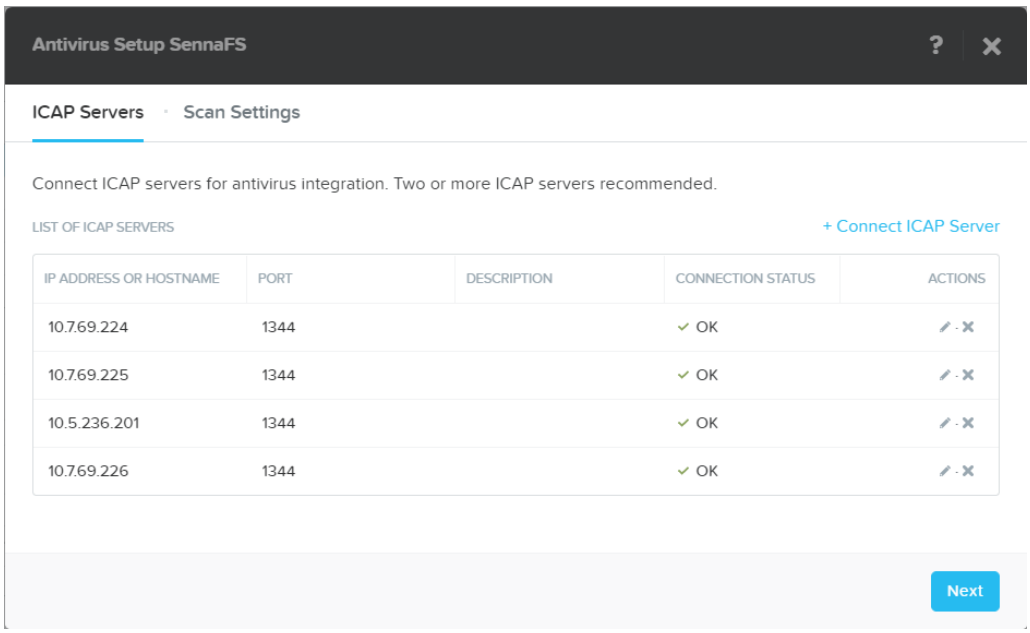

Figure 40: Configure Multiple ICAP Servers

Nutanix Files sets scanning defaults across the entire file server. You can enable scan on write and scan on read. Scan on write begins when the file is closed, and scan on read occurs when the file is opened. You can also exclude certain file types and files over a certain size. Share scan policies can override any defaults set for the file server.

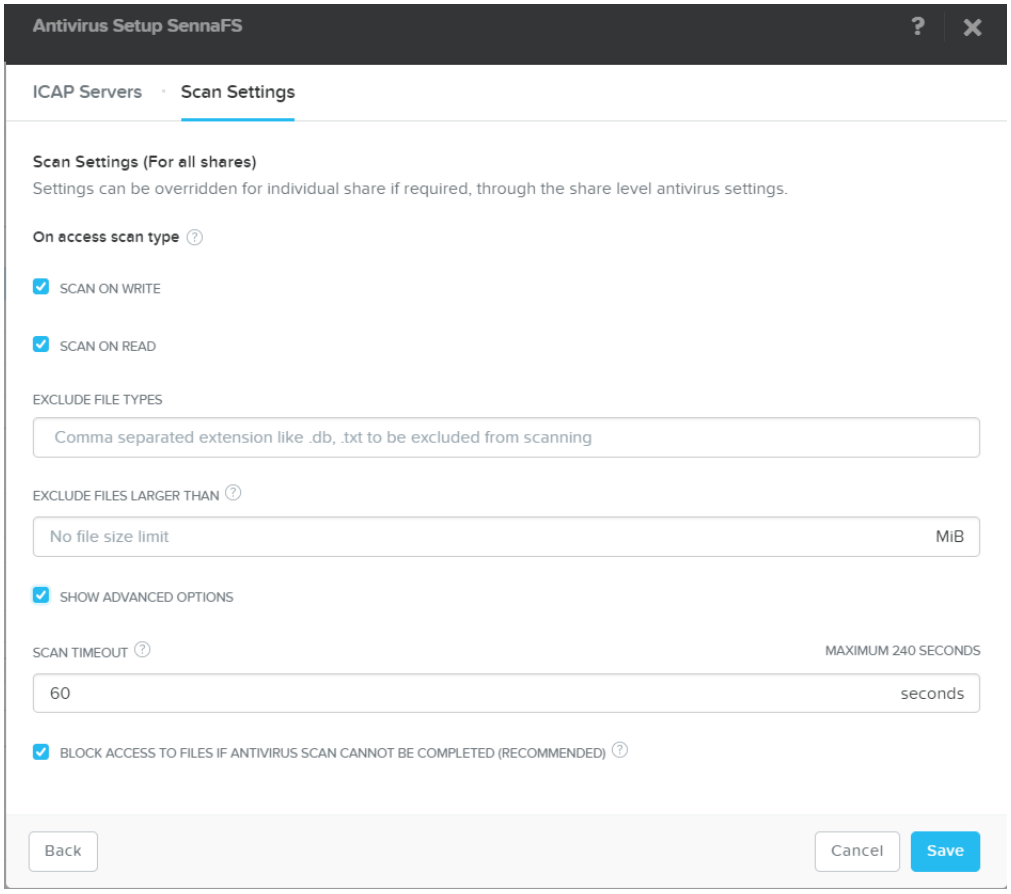

Figure 41: Default Scan Settings

For each ICAP server, Nutanix Files will spin up no more than 10 parallel connections per FSVM and randomly distribute the file scanning among all the ICAP servers. With heavier workloads, which may involve many scan requests and use all connections, the scan servers can be given more processing power to scan more files. As soon as the current scan finishes, the next file is picked up from the queue, which keeps the number of active connections at 10.

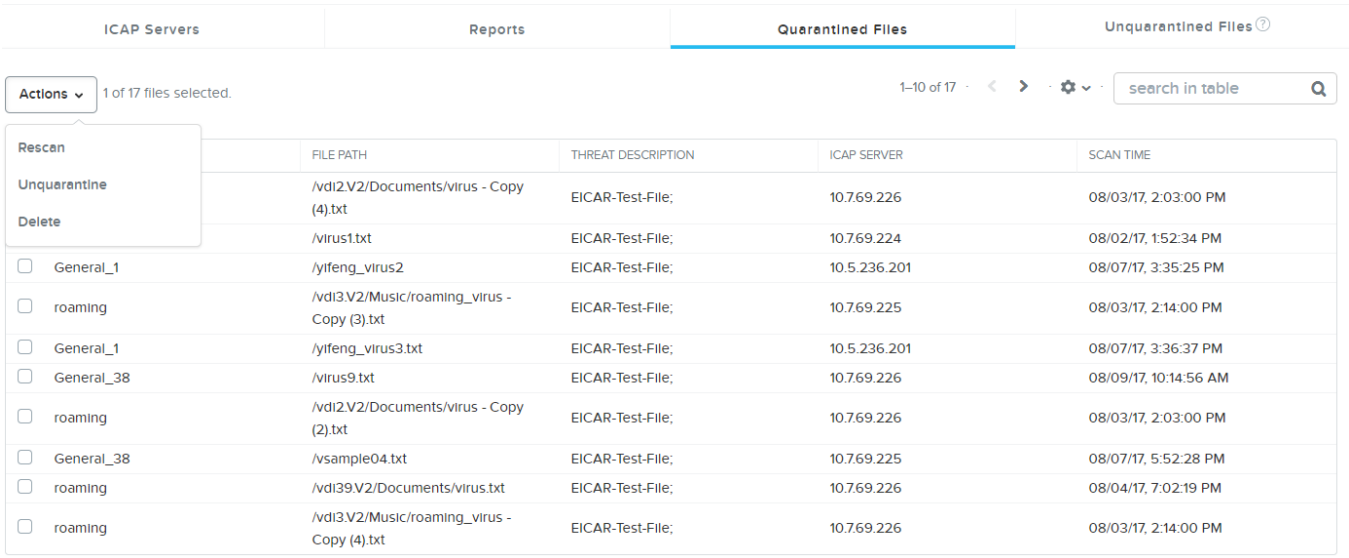

Figure 42: Quarantined Files

Once Nutanix Files quarantines a file, the administrator can rescan, unquarantine, or delete the file. You can search quarantined files if you need to restore a file quickly.

If your antivirus vendor doesn't support ICAP, you can scan the shares by installing an antivirus agent on a Windows machine, then mounting all the shares from the file server. This approach allows you to schedule scans for periods of low usage. At the desktop or client level, you can set your antivirus solution to scan on write or scan only when files are modified. You can configure high-security environments to scan inline for both reads and writes.

See the Nutanix Files release notes for the latest list of qualified ICAP server vendors. The current list of supported vendors includes Bitdefender, McAfee, Symantec, Kaspersky, and Sophos.

6.2. File Operations Monitoring

File operations monitoring is a native Nutanix Files API used to forward all SMB and NFS operations executed by clients to a registered target repository.

File operations monitoring can be broken down into two major areas for third-party vendors:

1. File activity
2. Audit

Partner software can register with the Nutanix Files instance to capture file activity events. Partner software can also make REST calls to Files to create a policy defining which file notification types you want to receive and the shares you want to receive them from. The partner

software uses a web client and communicates with the Nutanix Files file server using HTTPS requests. The HTTPS communication relies on SSL authentication. The HTTP server runs with either its unique self-signed SSL certificate or a Transport Layer Security (TLS) connection with the partner, then exchanges the keys between the two parties and sends messages to the partner server over this secure channel.

Third-party software can use different protocol communication methods such as syslog, Google, or Kafka protobuf. One common way to use the Nutanix file operations monitoring API is to forward events to a syslog server for any future auditing needs. Several vendors have integrated with this operations monitoring API:

- Peer Software

  Peer Software has integrated with Nutanix, including the file operations monitoring API, to support active-active deployments with real-time replication and file locking between heterogenous environments. To learn more about the Peer Software and Nutanix integration, please see the document Active-Active File Services for the Nutanix Enterprise Cloud.

- Netwrix

  Netwrix has integrated with the Nutanix file operations monitoring API. Netwrix provides comprehensive visibility into changes and data access across Nutanix Files instances. Netwrix offers a free add-on for Nutanix Files. See www.netwrix.com for more details.

6.3. Intelligent Backup

Traditional NAS backup approaches like Network Data Management Protocol (NDMP) have core limitations. Problems like complexity, performance (including the need for periodic full backups), and scale led Nutanix to develop a more modern backup solution: a changed file tracking (CFT) API that third-party backup vendors use. This API speeds backup times because it does not perform a metadata scan across your file server, which could contain millions of files and directories. The API enables scale as backups can occur across multiple FSVMs in parallel. The API also reassures customers aren't locked into any one solution as backups taken using the Nutanix Files API can be restored to other vendors.

CFT uses internal snapshots to track differences between backup windows. The backup software vendor requests the creation of these snapshots, and Files returns a list of the changed files in the form of a share. The backup software proceeds to mount the share onto the backup server to be read for backup.

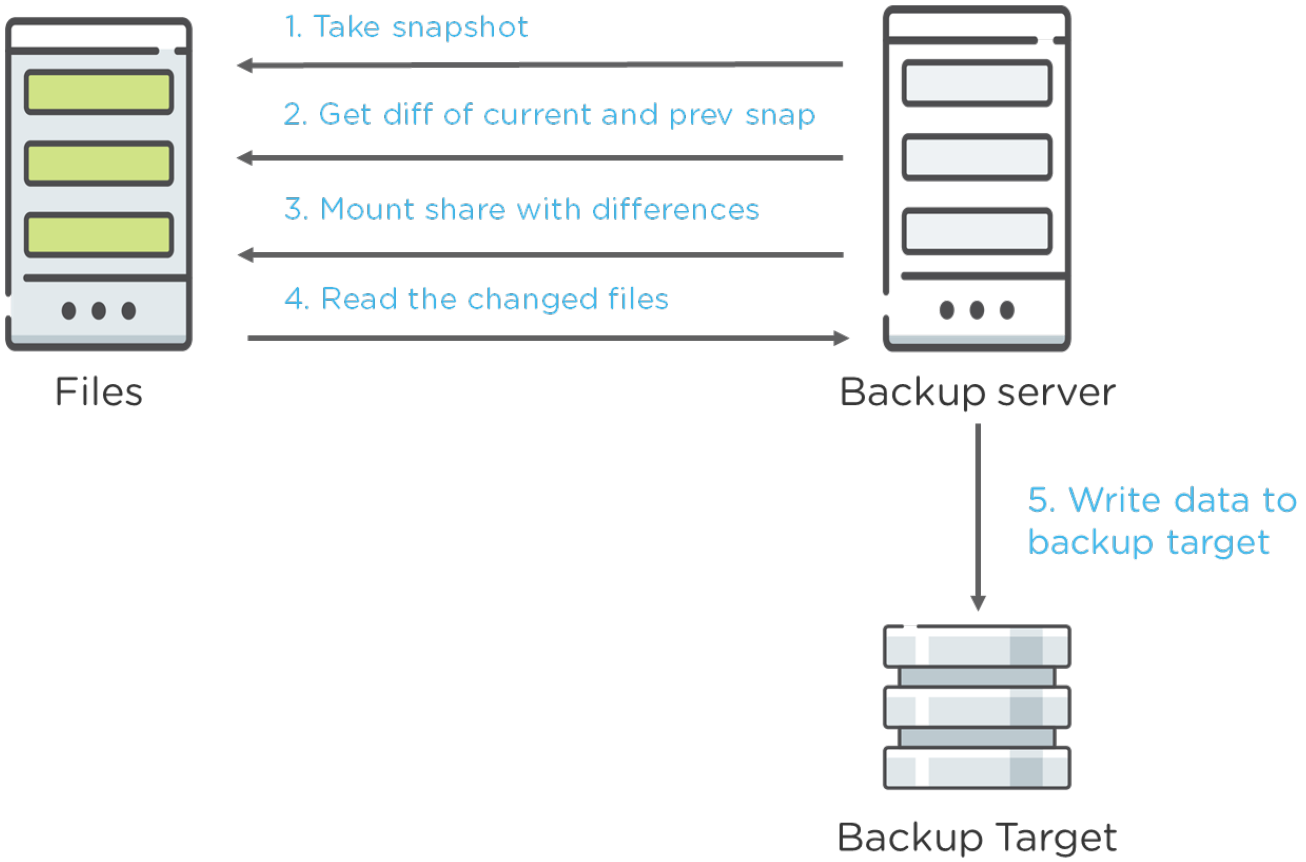

Figure 43: CFT Backup Process

Backup software can specify multiple shares and their respective snapshot information. Nutanix Files returns the list of URLs that map to the number of client streams that can start in parallel. Because shares are distributed evenly across the FSVMs based on load, the backup software can take advantage of all the FSVMs to drive throughput.

HYCU Backup and Recovery is integrated with the Nutanix Files CFT API. See www.hycu.com for more details.

Commvault version 11 service pack 15 is also integrated with the Nutanix Files CFT API as of the 3.5 release. Please see the Commvault blog post on the topic for more details.

There are two ways to back up Files shares with software that doesn't support CFT. One option is to run the backup application on a Windows machine and map the UNC path of the share as a drive that needs to be backed up.

There are also vendors (such as Commvault, Rubrik, and Veritas) that provide support for backing up file shares without mounting to a guest VM. These applications can read directly from

the UNC path. Note for this scenario that Nutanix does not qualify mounting SMB shares from Linux machines or NFS shares from Windows machines.

Tip: Do not attempt to back up the FSVMs. FSVMs are stateless and data is stored in Nutanix volume groups. Back up Nutanix Files from the share level. You can protect FSVMs using Nutanix protection domains.

Because the system spreads different standard shares across the cluster, try to back up multiple shares at the same time with multiple subclients. The distributed share allows you to configure multiple data readers to drive throughput.

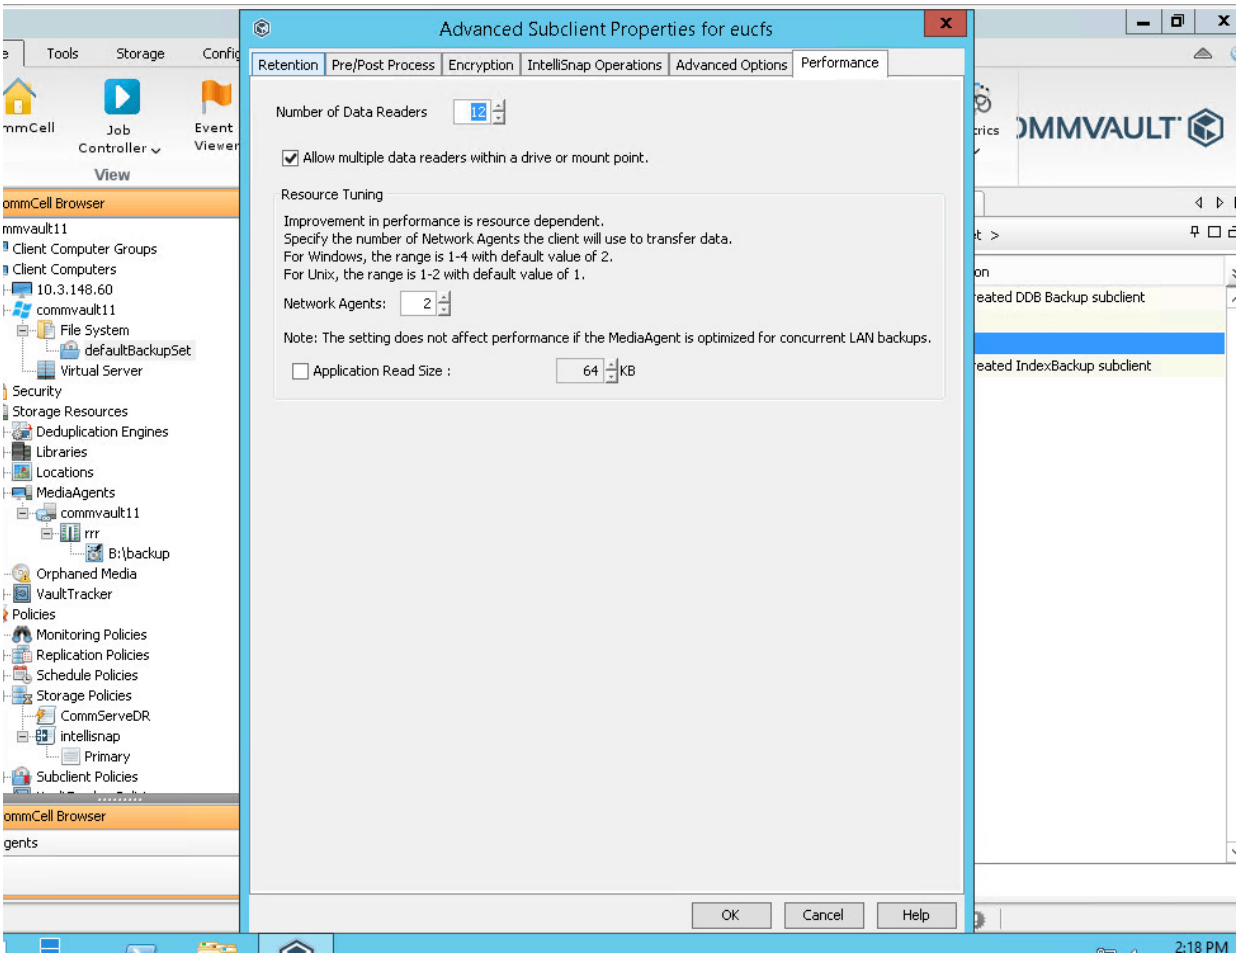

Figure 44: Adding More Readers to Drive Throughput on a Home Share

As an example of backup using Commvault, we tested 400 users spread out on three FSVMs, placing the data on a home share. We found that adding more readers for the backup job could increase performance. The bottleneck was the media agent, which was the backup destination

for the files. The media agent was virtualized and configured with only eight vCPUs. Adding more vCPUs to the media agent achieved a faster backup time.

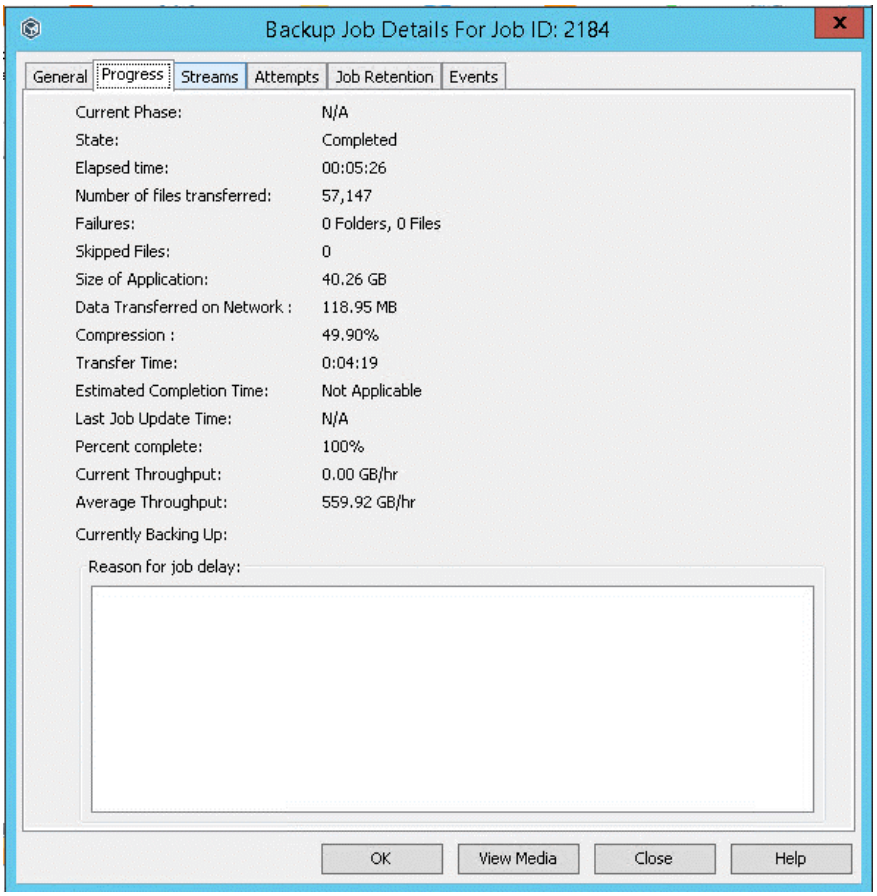

Figure 45: Backup Job Details

As the preceding figure shows, we achieved almost 600 GB average throughput per hour with a fairly small setup.

7. Conclusion

Nutanix Files delivers on-demand performance and automated provisioning to provide highly scalable file management. It reduces the administration and configuration time needed to deploy and maintain your environment, providing a public cloud experience within your private cloud.

The Acropolis Distributed Storage Fabric (DSF) and Prism make hosted file services highly resilient and easy to use—for example, you can configure the native data efficiency features of the DSF, such as compression and erasure coding (EC-X), for each individual file server. Prism lets you administer network and resource management, Active Directory, fault tolerance, and Nutanix Files share and export management all in one place, vastly improving operational efficiency.

You can deploy Files on a new or existing Nutanix environment, which means you have the flexibility to use your existing investment to take advantage of Files services.

Nutanix Files streamlines design, implementation, and maintenance, which provides an unrivaled user experience, an essential element of success for end-user computing projects that store data in shared environments.

Appendix

About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix Enterprise Cloud OS leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization, and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow us on Twitter @nutanix.

List of Figures

Figure 1: Nutanix Enterprise Cloud..... 7
Figure 2: Nutanix Files Server Instances Run as VMs for Isolation from the DSF.....9
Figure 3: Data Path Architecture of Nutanix Files..... 10
Figure 4: FSVM Internal Communication on One Node.....11
Figure 5: FSVM vDisks and Volume Groups.....12
Figure 6: FSVM Volume Group with 16 vDisks.....12
Figure 7: Initial 40 TB File System Pool.....13
Figure 8: 80 TB File System Pool Following First Expansion.....13
Figure 9: 120 TB File System Pool After Second Expansion.....13
Figure 10: 140 TB File System Pool After Final Expansion..... 13
Figure 11: Distributed Directory Shares.....14
Figure 12: Two Standard Shares on the Same File Server.....15
Figure 13: Create Share Menu with Nested Share Path.....18
Figure 14: Distributed Share and Top-Level Directories.....19
Figure 15: Three Standard Shares.....19
Figure 16: Performance Optimization Recommendation.....21
Figure 17: Load Balancing Volume Groups with Nutanix Volumes.....22
Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups.....23
Figure 19: Each FSVM Controls Its Own Volume Groups.....24
Figure 20: FSVM-1 Failure..... 25
Figure 21: DNS Request for SMB.....27

Figure 22: Shared Folders MMC..... 29
Figure 23: Enable SMB Encryption..... 30
Figure 24: Nutanix Files User Management with NFS..... 33
Figure 25: DNS Request for NFSv4..... 34
Figure 26: DNS Request for NFS v3..... 35
Figure 27: Multiprotocol Shares..... 36
Figure 28: Enable Non-Native Protocol Access..... 37
Figure 29: Multiprotocol User Mapping..... 38
Figure 30: Setting a Share Quota..... 40
Figure 31: Blocked File Types for the File Server..... 41
Figure 32: Blocked File Types for a Share..... 42
Figure 33: File Analytics Dashboard..... 44
Figure 34: File Analytics Audit Trails..... 45
Figure 35: File Analytics Anomaly Detection..... 46
Figure 36: File Analytics Usage Anomalies Dashboard..... 47
Figure 37: Protection Using SSR..... 48
Figure 38: Nutanix Files Cloning Use Cases..... 50
Figure 39: ICAP Workflow..... 51
Figure 40: Configure Multiple ICAP Servers..... 52
Figure 41: Default Scan Settings..... 53
Figure 42: Quarantined Files..... 54
Figure 43: CFT Backup Process..... 56
Figure 44: Adding More Readers to Drive Throughput on a Home Share..... 57

Figure 45: Backup Job Details.....58

List of Tables

Table 1: Document Version History..... 6

Table 2: Supported Active Client Connections..... 26

Table 3: User Mapping Requirements..... 38

Table 4: Order of Precedence for Quotas..... 41

ESPECIFICAÇÕES TÉCNICAS – CARACTERÍSTICAS DOS ITENS	
1. Licenciamento de software de gerenciamento centralizado	**Referência**
1.1. Ser fornecida licença de software para gerenciamento centralizado avançado de 1 (um) ou mais cluster, facilitando a tarefa de administração diária dos clusters localizados localmente ou distantes geograficamente. A ferramenta deverá apresentar as seguintes informações consolidadas de todos os clusters registrados:	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.1. Saúde dos Sistema clusters;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.2. Máquinas Virtuais;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.3. Armazenamento;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.4. Situação do Hardware;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.5. Painel de Análise de Performance;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.6. Painel de Alertas e Eventos.	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.2. O licenciamento será por entidade gerenciada;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.2.1. Ter uma ferramenta de planejamento de capacidade disponível, de forma a permitir a análise dos recursos e indicar máquinas virtuais subdimensionadas, superdimensionadas e inativas, para que seja possível identificação e remediação/otimização através da própria interface de gerenciamento. A ferramenta de planejamento de capacidade deve permitir simulações de provisionamento de novas aplicações com recomendações de otimização e eventuais capacidades ou equipamentos a serem adicionados ao cluster para que seja possível suportar estas novas aplicações;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-resource-planning-pc-c.html#nconcept_bcx_ldr_mx
1.2.2. Oferecer funcionalidade de planejamento de capacidade para crescimento baseado na carga de trabalho projetada;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-resource-planning-pc-c.html#nconcept_bcx_ldr_mx
1.2.3. Permitir a customização dos painéis de informação;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-dashboard-widgets-pc-r.html#nref_gdf_wzw_c5
1.2.4. Possuir funcionalidade de busca (*search*) que suporte busca contextualizada;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-search-options-pc-r.html
1.2.5. Possuir um algoritmo que determina o desvio padrão de recursos e situações (anomalias) e que são desta forma gatilhos dinâmicos para geração de alertas;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-behavioral-learning-pc-c.html#nconcept_icv_zgz_1bb
1.2.6. Possibilitar a criação e customização de relatório de consumo sobre os recursos da infraestrutura;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:ssp-report-management-ssp-pc-c.html#nconcept_yym_c5f_f1b
1.2.7. Oferecer a configuração de agendamento e envio dos relatórios via email.	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:ssp-report-management-ssp-pc-c.html#nconcept_yym_c5f_f1b
2. Licenciamento de software de microssegmentação de rede	
2.1.1. Ser fornecida licença de rede definida por software para servid or de hiperconvergência, por entidade gerenciada;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.2. Ser baseada em, ou compatível com Open Virtual Switch (OVS);	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.3. Suportar microssegmentação para prover controle granular e governança de todo o tráfego de entrada e saída de uma máquina virtual (VM) ou grupos de máquinas virtuais (VMs);	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.4. Permitir a associação de políticas de rede a VMs e aplicativos em vez de segmentos de rede específicos (por exemplo VLANs) ou identificadores (endereços IP ou MAC);	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.5. Prover visualização de todo tráfego e relacionamentos com a descoberta automática dos fluxos entre as máquinas virtuais;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.6. Prover uma estrutura de segurança orientada por políticas que inspeciona o tráfego dentro do data center, da seguinte maneira:	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.6.1. As políticas de segurança inspecionam o tráfego originado e terminado dentro de um datacenter, ajudando a eliminar a necessidade de firewalls adicionais no datacenter;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.6.2. Utilizar uma abordagem centrada na carga de trabalho em vez de uma abordagem centrada na rede, permitindo examinar o tráfego de, e para as VMs, independentemente de como as configurações de rede mudam e onde residem no data center.	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow

2.1.7. Prover uma abordagem agnóstica a estrutura de rede, centrada na carga de trabalho, permitindo que a equipe de virtualização implemente essas políticas de segurança sem depender de equipes de segurança de rede;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.8. As políticas de segurança deverão ser aplicadas às categorias (um agrupamento lógico de VMs) e não às próprias VMs, não importando quantas VMs são inicializadas em uma determinada categoria. O tráfego associado às VMs em uma categoria deverá ser protegido sem intervenção administrativa, em qualquer escala;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.9. A interface de gerenciamento deve oferecer uma abordagem baseada em visualização para configurar políticas e monitorar o tráfego ao qual uma determinada política se aplica:	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.9.1. Política de Segurança de Aplicação: quando for necessário proteger um aplicativo especificando origens e destinos de tráfego permitidos;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.9.2. Política de Isolamento do Ambiente: quando for necessário bloquear todo o tráfego, independentemente da direção, entre dois grupos de VMs identificados por sua categoria. VMs dentro de um grupo podem se comunicar umas com as outras;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.9.3. Política de Quarentena: quando for necessário isolar uma VM comprometida ou infectada e, opcionalmente, desejar submetê-la à perícia.	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.10. Garantir que seja apenas permitido o tráfego entre camadas de aplicativos ou outros limites lógicos, garantindo a proteção contra ameaças avançadas para que não sejam propagadas no ambiente virtual;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.11. Permitir a atualização automática durante todo o ciclo de vida da VM, eliminando a carga do gerenciamento de mudanças de políticas;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.12. Permitir categorizar as Máquinas Virtuais de forma a permitir a criação políticas de segurança com no mínimo as seguintes funções:	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.12.1. Isolar o tráfego de dados entre Máquinas Virtuais de Diferentes categorias;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.12.2. Isolar o tráfego de dados de Máquinas Virtuais específicas para modo de quarentena, tanto forense quanto restrita, de forma a prover uma rápida reação ao time de infraestrutura em caso de Máquinas Virtuais contaminadas ou pertencentes a usuários que foram desligados ou sob procedimento de custódia de dados;	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.12.3. Mapear o tráfego de entrada, entre as camadas e de saída de aplicações, permitindo ao administrador determinar quais servidores tem acesso de entrada na aplicação, o tipo de protocolo e o número da porta que o fluxo de dados pode ocorrer, permitir ou restringir também o fluxo de dados entre as camadas, máquinas virtuais, pertencentes à aplicação, através da especificação do protocolo e o número da porta, realizar também o mesmo procedimento para conexões de saída das camadas da aplicação, também através da especificação de protocolo e número de porta.	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
2.1.13. Possuir integração com software de terceiros que permita o redirecionamento do tráfego das VMs para ferramentas terceiras, como por exemplo, mas não limitado a softwares de detecção e prevenção de intrusos (IDS/IPS), monitoração de performance de aplicações (APM), balanceadores de carga.	https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
3. Licença de software para estrutura de diretórios distribuído	
3.1.1. Possibilizar a entrega de serviço de arquivos compartilhados dos tipos SMB e NFS;	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:Files-v3_6
3.1.2. Permitir integração com o *Nutanix Operating System* para criar serviços hospedados;	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:Files-v3_6
3.1.3. Suportar funcionalidades de balanceamento de carga, replicação nativa e cotas para:	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-server-load-balance-c.html
3.1.3.1. Diretório de arquivos do usuário, onde apenas o dono do perfil terá acesso (*Home directory*);	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-server-create-share-export-c.html#nconcept_wyw_2zp_scb
3.1.3.2. Diretórios departamentais que poderão ser acessados por qualquer usuário com as devidas permissões (*Access Based Enumeration – ABE*);	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-server-enable-abe-wc-t.html
3.1.4. Suportar tierização, compressão, desduplicação e erasure-coding;	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-share-compression-c.html
3.1.5. Suportar o protocolo - ICAP (*Internet Content Adaptation Protocol*);	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-server-anti-virus-c.html#nconcept_on4_xnh_n1b
3.1.6. Suportar recuperação de desastres com replicação assíncrona;	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-server-async-dr-c.html
3.1.7. Possuir integração com soluções de proteção de dados de terceiros;	https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:fil-file-server-anti-virus-c.html#nconcept_on4_xnh_n1b
3.1.8. Permitir a auditoria da estrutura compartilhada, possibilitando a identificação dos usuários de maior uso;	https://portal.nutanix.com/page/documents/details/?targetId=File-Analytics-v2_1:File-Analytics-v2_1
3.1.9. Ter a visibilidade dos tipos de arquivos e tamanho;	https://portal.nutanix.com/page/documents/details/?targetId=File-Analytics-v2_1:File-Analytics-v2_1
3.1.10. Suportar detecção de anomalia;	https://portal.nutanix.com/page/documents/details/?targetId=File-Analytics-v2_1:File-Analytics-v2_1
3.1.11. Suportar análise de tendência de capacidade;	https://portal.nutanix.com/page/documents/details/?targetId=File-Analytics-v2_1:File-Analytics-v2_1

3.1.12. Permitir retenções de cópias temporais com possibilidade de recuperação granular de arquivos pelos usuários sem intervenção do administrador do cluster.	https://portal.nutanix.com/page/documents/details/?targetId=File-Analytics-v2_1:File-Analytics-v2_1
4. Licença de software para estrutura de objetos	
4.1.1. Prover solução de armazenamento de objetos definida por software com interface HTML5 de alta disponibilidade e resiliência;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.2. Suportar interface de API REST compatível com protocolo S3;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.3. Prover *namespace* único global;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.4. Para maior segurança, deve contar com secret key/access key, políticas de acesso read/write baseado em usuário;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.1. Deve possibilitar configuração Data-at-Rest Encryption;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.2. Suportar certificação FIPS;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.3. Suportar ações de API comumente utilizadas, como: GET, PUT, HEAD, LIST, DELETE;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.4. Ser compatível com Nutanix Prism Central;	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
4.1.5. Possibilidade de versionamento, gerenciamento de ciclo de vida.	https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
5. Expansão de suporte e garantia - Modelo 1	
5.1.1. A garantia e suporte técnico deverão ser estendidas por mais 3 (três) anos, pelo fabricante, na modalidade 24 (horas) x 7 (dias por semana) x 365 (dias por anos) e troca de peças no próximo dia útil;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
5.1.2. A renovação deverá abranger o *hardware* e *software* existentes em cada servidor;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
5.1.3. Os chamados técnicos serão categorizados em níveis de severidade, conforme a seguir:	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.3.1. Nível 1: Equipamento total ou parcialmente inoperante, com degradação de desempenho ou ocorrência de mau funcionamento;	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.3.2. Nível 2: Equipamento operante com ocorrência de alarmes; consultas gerais sobre administração, configuração, otimização, troubleshooting ou utilização.	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.4. Os chamados técnicos deverão ser atendidos nos seguintes prazos:	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.4.1. Nível 1: Até 2h (duas horas) para o início do atendimento;	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.4.2. Nível 2: Até 24h (vinte e quatro horas) para o início do atendimento.	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.5. O atendimento de um chamado técnico será considerado solucionado conforme os seguintes critérios:	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.5.1. Nível 1: O equipamento estiver operante, sem degradação de desempenho ou ocorrência de mau funcionamento;	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.5.2. Nível 2: O equipamento estiver operante sem ocorrência de alarmes e as consultas realizadas tenham sido respondidas.	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
5.1.6. Cada servidor possui as seguintes características:	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
5.1.6.1. Dois processadores Intel(R) Xeon(R) v3 ou v4 com 10 (dez) cores cada;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
5.1.6.2. Memória RAM de 256GB (duzentos e cinquenta e seis);	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-
5.1.6.3. Até 1 (um) disco SSD de 960GB padrão SATA de 6.0 Gb/s e hot-swap;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications

5.1.1.1. Até 2 (dois) discos HDD de 8TB cada padrão SATA de 6.0 Gb/s e hot-swap;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
5.1.1.2. 2 (duas) interfaces de rede de 10Gbps SFP+;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
5.1.1.3. Licenciamento Pro ou Ultimate.	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications , https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=Licensing-Guide:lic-licensing-prism-element-c.html
6. Expansão de suporte e garantia - Modelo 2	
6.1.1. A garantia e suporte técnico deverão ser estendidas por mais 3 (três) anos, pelo fabricante, na modalidade 24 (horas) x 7 (dias por semana) x 365 (dias por anos) e troca de peças no próximo dia útil;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.2. A renovação deverá abranger o *hardware* e *software* existentes em cada servidor;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.3. Os chamados técnicos serão categorizados em níveis de severidade, conforme a seguir:	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.3.1. Nível 1: Equipamento total ou parcialmente inoperante, com degradação de desempenho ou ocorrência de mau funcionamento;	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.3.2. Nível 2: Equipamento operante com ocorrência de alarmes; consultas gerais sobre administração, configuração, otimização, troubleshooting ou utilização.	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.4. Os chamados técnicos deverão ser atendidos nos seguintes prazos:	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.4.1. Nível 1: Até 2h (duas horas) para o início do atendimento;	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.4.2. Nível 2: Até 24h (vinte e quatro horas) para o início do atendimento.	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.5. O atendimento de um chamado técnico será considerado solucionado conforme os seguintes critérios:	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.5.1. Nível 1: O equipamento estiver operante, sem degradação de desempenho ou ocorrência de mau funcionamento;	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.5.2. Nível 2: O equipamento estiver operante sem ocorrência de alarmes e as consultas realizadas tenham sido respondidas.	Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
6.1.6. Cada servidor possui as seguintes características:	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.6.1. Dois processadores Intel(R) Xeon(R) v3 ou v4 com 12 (doze) cores cada;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.6.2. Memória RAM de 512GB (duzentos e cinquenta e seis);	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.6.3. Até 2 (dois) disco SSD de 1200GB padrão SATA de 6.0 Gb/s e hot-swap;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.6.4. Até 4 (dois) discos HDD de 8TB cada padrão SATA de 6.0 Gb/s e hot-swap;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
6.1.6.5. 2 (duas) interfaces de rede de 10Gbps SFP+;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications

6.1.6.6. Licenciamento Pro ou Ultimate.		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications , https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=Licensing-Guide:lic-licensing-prism-element-c.html
7. Expansão de suporte e garantia - Modelo 3		
7.1.1. A garantia e suporte técnico deverão ser estendidas por mais 3 (três) anos, pelo fabricante, na modalidade 24 (horas) x 7 (dias por semana) x 365 (dias por anos) e troca de peças no próximo dia útil;		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.2. A renovação deverá abranger o *hardware* e *software* existentes em cada servidor;		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.3. Os chamados técnicos serão categorizados em níveis de severidade, conforme a seguir:		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.3.1. Nível 1: Equipamento total ou parcialmente inoperante, com degradação de desempenho ou ocorrência de mau funcionamento;		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.3.2. Nível 2: Equipamento operante com ocorrência de alarmes; consultas gerais sobre administração, configuração, otimização, troubleshooting ou utilização.		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.4. Os chamados técnicos deverão ser atendidos nos seguintes prazos:		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.4.1. Nível 1: Até 2h (duas horas) para o início do atendimento;		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.4.2. Nível 2: Até 24h (vinte e quatro horas) para o início do atendimento.		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.5. O atendimento de um chamado técnico será considerado solucionado Conforme os seguintes critérios:		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.5.1. Nível 1: O equipamento estiver operante, sem degradação de desempenho ou ocorrência de mau funcionamento;		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.5.2. Nível 2: O equipamento estiver operante sem ocorrência de alarmes e as consultas realizadas tenham sido respondidas.		Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
7.1.6. Cada servidor possui as seguintes características:		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.1. Dois processadores Intel(R) Xeon(R) v3 ou v4 com 8 (oito) cores cada;		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.2. Memória RAM de 64GB (sessenta e quatro);		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.3. Até 1 (um) disco SSD de 800GB padrão SATA de 6.0 Gb/s e hot-swap;		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.4. Até 5 (cinco) discos HDD de 8TB (oito) cada padrão SATA de 6.0 Gb/s e hot- swap;		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.5. 2 (duas) interfaces de rede de 10Gbps SFP+;		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.6. Licenciamento Pro ou Ultimate.		Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications , https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=Licensing-Guide:lic-licensing-prism-element-c.html

Files 2.1

File Analytics Guide

March 23, 2020

Contents

1. File Analytics 4
   - Deployment Requirements 5
     - Network Requirements 5
   - Limitations 5

2. Administration 7
   - Deploying File Analytics 7
     - Enabling Files Analytics 9
     - Disabling Files Analytics 13
   - Launching File Analytics 13
   - Analytics VM Management 13
     - Removing Analytics VMs 14
   - Updating Credentials 14
   - Upgrades 14
     - Upgrade File Analytics at a Dark Site 14

3. Dashboard 16
   - Updating File Categories 20
   - Health 21

4. Anomalies 22
   - Configuring Anomaly Detection 23
   - Configuring an SMTP Server 24

5. Audit Trails 27
   - View Audit Trails 27
   - Audit Trails - Users 27
   - Audit Trails - Folders 29
   - Audit Trails - Files 30
   - Audit Trails - Client IP 32

6. Analytics Options 36
   - Updating Data Retention 36
   - Scanning the File System 36
   - Blacklisting 37
   - Data Protection 37
     - Configuring Disaster Recovery 38
     - Activating Disaster Recovery 38
     - Deploying Analytics on a Remote Site 39

Copyright....42
License....42
Conventions....42
Default Cluster Credentials....42
Version....43

1

FILE ANALYTICS

File Analytics provides data and statistics on the operations and contents of a file server.

Once deployed, Files adds an Analytics VM to the Files cluster. A single Analytics VM supports all file servers in the cluster; however, you must enable Analytics separately for each file server. Data on the Analytics VM is protected, and is kept in a separate volume group.

Once you deploy Analytics, a new File Analytics link appears on the file server actions bar. You can access Analytics through this link for any file server where it is enabled.

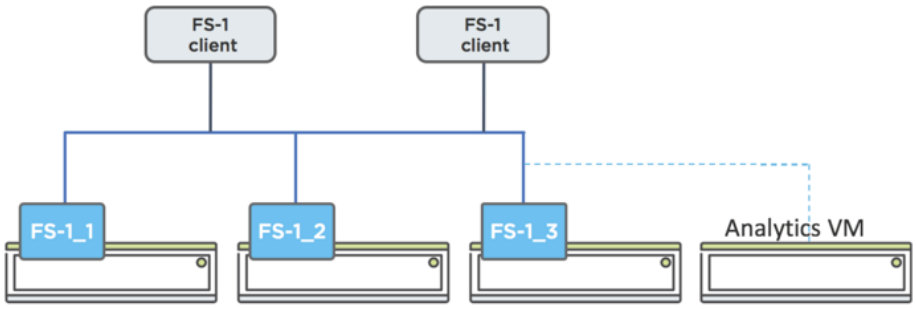

Figure 1: File Analytics VM

Display Features

The Analytics web console consists of the following display features:

- **Main menu bar**: The main menu bar appears at the top of every page of the Analytics web console. The main menu bar includes the following display features:
-
  - **Dashboard tab**: View widgets that present data on file trends, distribution, and operations.
  - **Audit Trails tab** : Search for a specific user or file and view various widgets to audit activity.
  - **Anomalies tab**: Create anomaly policies and view anomaly trends.
  - **Status icon**
  - **File server drop-down**
  - **Settings drop-down**
  - **Health icon**
  - **Admin dropdown**

Deployment Requirements

Meet the following requirements prior to deploying File Analytics.

Ensure that you have performed the following tasks and your Files deployment meets the following specifications.

- Assign the file server administrator role to an Active Directory (AD) user, see *Managing Roles* in the *Nutanix Files Guide*.
- Log on as the Prism **admin** user to deploy the File Analytics server. An Active Directory (AD) user or an AD user mapped to a Prism admin role cannot deploy Analytics.
- Configure a VLAN with one dedicated IP address for Analytics, or you can use an IP address from an existing Files external network. This IP address must have connectivity to AD, the control VM (CVM), and Files.

  Note: Do not install Analytics on the Files internal network.

- (optional) Assign the file server administrator role to an LDAP user, see *Managing Roles* in the *Nutanix Files Guide*.
- Ensure that all software components meet the supported configurations and system limits, see the *File Analytics Release Notes*.

Network Requirements

Open the required network connections before deploying Analytics.

External hosts use the following ports to connect to File Analytics.

Note: In addition to meeting the following network requirements, ensure to meet Nutanix Files port requirements.

Table 1: Incoming Connections

Port	Source	Purpose
7	Prism	ICMP ECHO PING request
3000	Client	UI Access
443	Control VM (CVM)	Gateway API access
29092	File server	Kafka Broker access

Limitations

File Analytics has the following limitations.

Note: Depending on data set size, file count, and workload type, enabling Analytics can affect performance. Latency is more common with heavy file-metadata operations (directory and file creation, deletion, permission changes, and so on). To minimize the impact on performance, ensure that the host has enough CPU and memory resources to handle the File Analytics VM (AVM), file servers, and guest VMs (if any).

- Only Prism admin users can deploy Analytics. Active Directory (AD) users and AD users mapped to Prism admin roles cannot deploy File Analytics.

- Analytics analyzes data from 1 month up to 1 year, based on configuration. Analytics automatically deletes data beyond the defined configuration.
- You cannot deploy or decommission Analytics when a file server has high-availability (HA) mode enabled.
- If file server DNS or IP changes, Analytics does not automatically reconfigure.
- Analytics does not collect metadata for files on Kerberos authenticated NFS v4.0 shares.
- If Analytics is running on a one-node file server, you cannot upgrade using the Life Cycle Manager (LCM)
- Analytics does not support hard links.
- You cannot enable Analytics on a file server clone.
- Analytics creates an unprotected Prism and an unprotected file server user for integration purposes. Do not delete these users.

2

ADMINISTRATION

Overview of administrative processes for File Analytics.

As an admin, you have the privileges to perform administrative tasks for Analytics. To add a file server admin user, see *Managing Roles* in the *Nutanix Files Guide*. The topics in this chapter describe the basics for administering your File Analytics environment. For advanced administrative options, refer to the Analytics Options on page 36 chapter.

Deploying File Analytics

Follow this procedure to deploy the File Analytics server.

Before you begin

Ensure that your environment meets all requirements prior to deployment, see Deployment Requirements on page 5.

Procedure

Deploying the File Analytics server.

1. Go to **Support Portal** > **Downloads** > **Files**.
2. Download the File Analytics QCOW2 and JSON files.
3. Log on to Prism with user name and password of the Prism administrator.

   Note: An Active Directory (AD) user or an AD user mapped to a Prism admin role cannot deploy Analytics.

4. In Prism, go to the **File Server** view and click the **Deploy File Analytics** action link.

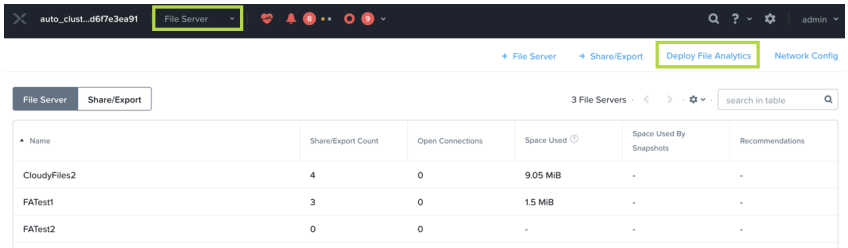

Figure 2: File Analytics

5. In the **Deploy File Analytics** window, click **Deploy**.

   » In the available versions section, click the **Install** button for the target version (continue to step 7.).

   » Install by uploading installation binary files (continue to next step).

6. Upload installation files.

   a. In the **Upload installation binary** section, click on the **upload the File Analytics binary** link to upload the File Analytics JSON and QCOW files.

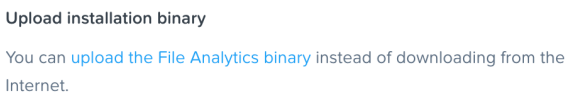

Figure 3: Upload Binary Link

   b. Under **File Analytics Metadata File (.Json)**, click **Choose File** to choose the downloaded JSON file.

   c. Under **File Analytics Instalation Binary** (.Qcow2) click **Choose File** to choose the downloaded QCOW file.

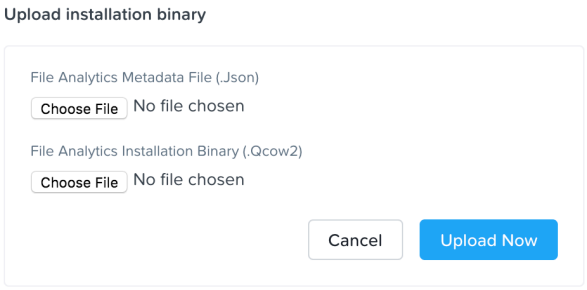

Figure 4: Upload Binary Files

   d. Click **Upload Now** after choosing the files.

7. Do the following in the indicated fields:

   a. **Name**: Enter a name for the File Analytics VM (AVM).

   b. **Storage Container:** select a storage container from the dropdown.
   The dropdown only displays file server storage containers.

   c. **Network List**: Select a VLAN.

> Note: If the selected network is **unmanaged**, enter more network details in the **Subnet Mask**, **Default Gateway IP**, and **IP Address** fields as indicated.

> Note: The AVM must use the client-side network.

8. (optional) Check the **Show Advanced Settings** box and do the following in the indicated fields.

   > Note: Refer to *System Limits* in the *File Analytics Release Notes* for information on recommended sizing. To see the number of files on your file server, refer to the **Usage** tab in the **File Server** view.

   a. **Memory (GiB)**: Enter the memory capacity for the Analytics VM.

   b. **VCPUs**: Enter the number of vCPUs.

   c. **Cores**: Enter the number of cores.

   d. **DNS Resolver IP**: Enter a coma-separated list of DNS resolver IPs.

   e. **NTP Servers**: Enter a coma-separated list of NTP servers.

9. Click **Deploy**.

10. In the main menu drop-down, select the **Tasks** view to monitor the deployment progress.

Results

Once deployment is complete, Analytics creates an AVM, CVM, and a new Files user to make REST API calls. Do not delete the CVM, AVM, or the REST API user.

Enabling Files Analytics

Steps for enabling File Analytics after deployment or disablement.

About this task

> Attention: Nutanix recommends enabling Analytics during off-peak hours.

Follow these steps to enable File Analytics after disabling the application.

> Note: Analytics saves all previous configurations.

Procedure

1. In the **File Server** view in **Prism**, select the target file server.

2. (skip to step 3 if you are re-enabling a file server) click **Manage roles** to add a file server admin user, see *Managing Roles* in the *Nutanix Files Guide*.

3. In the **File Server** view, select the target file server and click **File Analytics** in the tabs bar.

4. (skip to step 5 if you are not re-enabling a disabled instance of Analytics) to re-enable Files Analytics, click **Enable Files Analytics** in the message bar.

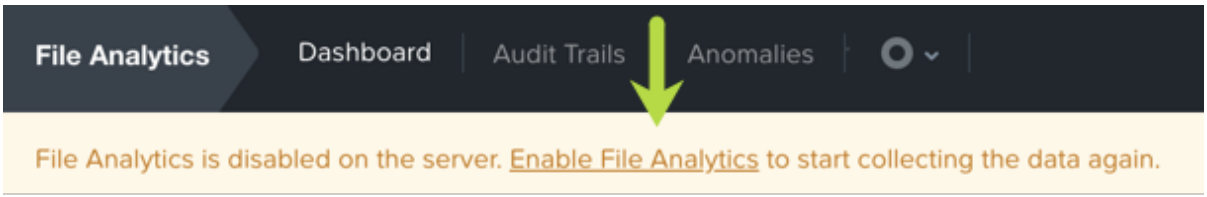

Figure 5: Enabling File Analytics Link

The **Enable File Analytics** dialog-box appears. Skip the remaining steps.

5. In the **Data Retention** field, select a data retention period.

6. In the **Enable File Analytics** dialog-box, enter the credentials as indicated:

   a. (For SMB users only) In the SMB Authentication section, do the following in the indicated fields to provide SMB authentication details:

      - **Username**: Enter the AD username for the file server administrator, see **File Analytics Prerequisites**.
      - **Password**: Enter the AD user password for the file server administrator.

   b. (For NFS users only) In the NFS Authentication section, do the following in the indicated fields to provide NFS authentication details:

      - **LDAP Server URI**: Enter the URI of the LDAP server.
      - **Base DN**: Enter the base DN for the LDAP server.
      - **Password**: Enter the LDAP user password for the file server administrator.

Enable File Analytics

File Analytics Server	Running
File Server	FATest1

Data Retention

Last 1 year

SMB Authentication

File Analytics server crawls through the files and collects data. It needs a admin user credentials for the task. Enter the Active Directory user credentials for the SMB protocol with admin user privileges

Active Directory Realm Name

child2.afs.minerva.com

Username

Password

NFS Authentication

Enter the LDAP server details

LDAP Server URI

ldap://10.4.44.37:389

Base DN

dc=devtest,dc=local

Bind DN (Optional)

cn=ldapadmin,dc=devtest,

Bind Password

Enable

7. Click **Enable** .

Results

After enablement, Analytics performs a one-time file system scan to pull metadata information. The duration of the scan varies depending on the protocol of the share. There is no system downtime during the scan.

Example

Scanning 3–4 million NFS files or 1 million SMB files takes about 1 hour.

Disabling Files Analytics

About this task

Follow the steps as indicated to disable Files Analytics.

Procedure

1. In Files Analytics click the **gear icon** > **Disable Files Analytics** .
2. In the dialog-box, click **Disable**.
   Disabling Analytics disables data collection. The following message banner appears.

```
File Analytics is disabled on the server. Enable File Analytics to start collecting the data
 again.
```

Launching File Analytics

About this task

Do the following to launch Analytics.

Procedure

1. From the Prism views drop-down, select the **File Server** view.
2. Select the target file server from the entity tab.
3. Click the **File Analytics** action button below the entity table.

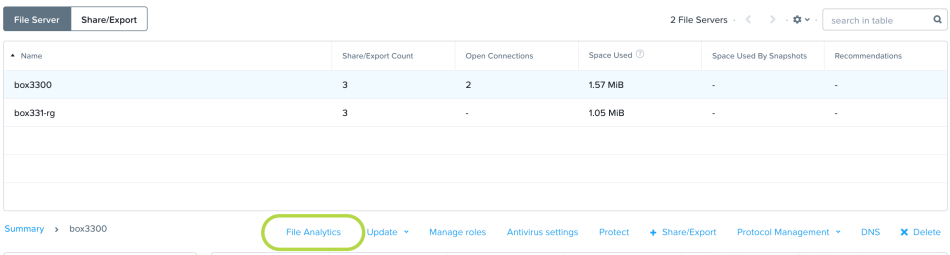

Figure 6: Launch File Analytics

Analytics VM Management

To update an Analytics VM (AVM), refer to the sizing guidelines in the File Analytics release notes and follow the steps in the *VM Management* topic of the *Prism Web Console Guide*.

Removing Analytics VMs

Remove a File Analytics VM (AVM) by disabling it and deleting it from the cluster in Prism.

About this task

Procedure

1. Disable Analytics on all file servers in the cluster, see Disabling Files Analytics on page 13.
2. Delete the AVM from the cluster in Prism Element.

   > Note: You must delete the AVM using the Prism UI.

   Refer to *Managing a VM (AHV)* in *The Prism Web Console Guide*. When you delete an Analytics VM (AVM), you also delete your File Analytics configuration. Delete the AVM only using the Prism UI.

Updating Credentials

About this task

Follow the steps as indicated to update authentication credentials for LDAP or Active Directory.

Procedure

1. Click **gear icon** > **AD/LDAP Configuration**.
2. Type the new password in the **Password** field.
3. Click **Save**.

Upgrades

Perform File Analytics upgrades using the Life Cycle Manager feature in Prism Element.

Before you upgrade Analytics, ensure that you are running a compatible version of AOS and Files. Refer to File Analytics release notes for compatibility details. You can upgrade both AOS and Files through Prism Element, see *AOS Upgrade* in the *Prism Web Console Guide*.

To upgrade Analytics, perform inventory and updates using the Life-Cycle Management feature, see the *Life Cycle Manager Guide* for instructions on performing inventory and updates.

> Note: The Analytics UI is not accessible during upgrades.

Upgrade File Analytics at a Dark Site

Upgrade File Analytics at a dark site using the Life-Cycle Manager (LCM).

About this task

Before you begin

You need a local web server reachable by your Nutanix clusters to host the LCM repository.

Procedure

1. From a device that has public Internet access, go to **Nutanix Portal** > **Downloads** > **Tools & Firmware**.
   a. Download the tar file `lcm_dark_site_version.tar.gz`.
   b. Transfer `lcm_dark_site_version.tar.gz` to your local web server and untar into the `release` directory.
2. From a device that has public Internet access, go to the Nutanix portal and select **Downloads** > **File Analytics**.
   a. Download the following files.
      - `file_analytics_dark_site_version.tar.gz`
      - `nutanix_compatibility.tgz`
      - `nutanix_compatibility.tgz.sign`
   b. Transfer `file_analytics_dark_site_version.tar.gz` to your local web server and untar into the `release` directory.
   c. Transfer the `nutanix_compatibility.tgz` and `nutanix_compatibility.tgz.sign` files to your local web server (overwrite existing files as needed).
3. Log on to Prism Element.
4. Click **Home** > **LCM** > > **Settings**.
   a. In the **Fetch updates from** field, enter the path to the directory where you extracted the tar file on your local server. Use the format `http://webserver_IP_address/release`.
   b. Click **Save**.
   You return to the Life Cycle Manager.
   c. In the LCM sidebar, click **Inventory** > **Perform Inventory**.
   d. Update the LCM framework before trying to update any other component.
   The LCM sidebar shows the LCM framework with the same version as the file you downloaded.

3

DASHBOARD

The Dashboard tab displays data on the operational trends of file server.

The main **Dashboard** tab is the opening screen that appears after launching File Analytics from Prism. The dashboard displays widgets that present data on file trends, distribution, and operations.

Table 2: Dashboard Panes

Tile Name	Description	Values
Capacity Trend	Displays capacity trends for the file server including capacity added, capacity removed, and net changes. Clicking a bar in the graph provides more granular details about the files.	Seven days, the last 30 days, or the last 1 year.
Capacity Trend Details view	Clicking on a Capacity Trend interval displays the Capacity Trends Details view for that interval.	Overall capacity data for the interval: net capacity change (GB), capacity added (GB), capacity removed (MB), total capacity (GB). Breakdown by extension: extension name, net capacity change (MB), capacity added (MB), capacity removed (B).
Data Age	Displays the percentage of data by age.	Less than 3 months, 3–6 months, 6–12 months, and > 12 months.
Anomaly Alerts	Displays alerts for configured anomalies, see Configuring Anomaly Detection on page 23.	
Permission Denials	Displays users who have had excessive permission denials and the number of denials.	[user id], [number of permission denials]
File Distribution by Size	Displays the number of files by file size. Provides trend details for top 5 files.	Less than 1 MB, 1–10 MB, 10–100 MB, 100 MB to 1 GB, greater than 1 GB).

Tile Name	Description	Values
File Distribution by Type	Displays the space taken up by various applications and file types. The file type is determined by the file extension. See the *File Types* table for more details.	MB or GB
File Distribution by Type Details view	Displays a trend graph of the top 5 file types. Displays details of file distribution including file type, current space used, current number of file, and change in space for the last 7 or 30 days.	Daily size trend for top 5 files (GB), file type (see File Type table), current space used (GB), current number of files (numeric), change in last 7 or 30 days (GB).
Top 5 active users	Lists the users who have accessed the most files and number of operations the user performed for the specified period. When there are more than 5 active users, the **more** link provides details on the top 50 users.	24 hours, 7 days, 1 month, or 1 year.
Top 5 accessed files	Lists the 5 files that are accessed most frequently. When more than 5 files are accessed, the **more** link provides details on the top 50 files.	Twenty-four hours, 7 days, 1 month, or 1 year.
Files Operations	Displays the number of times operations have been performed during current and peak intervals and average operations over the selected period. Operations include: create file, read, write, and permission changed. For example, for a 7-day period, the intervals are one day. The current column shows operations for the current day, and the peak column shows the number of operations for the day with the most operations. Average shows the average number of operations per day.	Twenty-four hours, 7 days, 1 month, or 1 year.

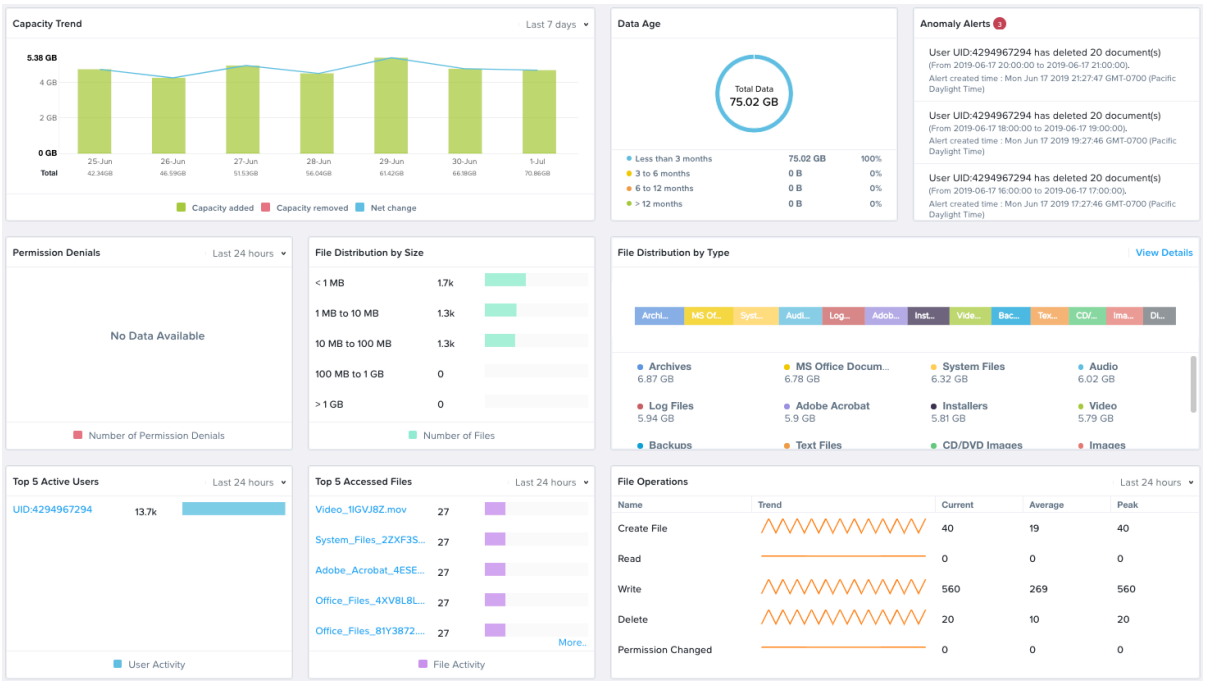

Figure 7: Analytics Dashboard

Clicking an event period in the **Capacity Trend** widget displays the Capacity Trend Details view for that period. The view includes three tabs **Share/Export**, **Folder**, and **Category**. Each tab includes columns detailing entity details: **Name**. Net Capacity Change, Capacity Added, and Capacity Removed.

Table 3: Capacity Trend Details

Category	Supported File Type
Name	Name of share/export, folder, or category.
Net Capacity Change	The total difference between capacity at the beginning and the end of the specified period.
Share Name (for folders only)	The name of the share or export that the folder belongs to.
Capacity Added	Total added capacity for the specified period.
Capacity Removed	Total removed capacity for the specified period.

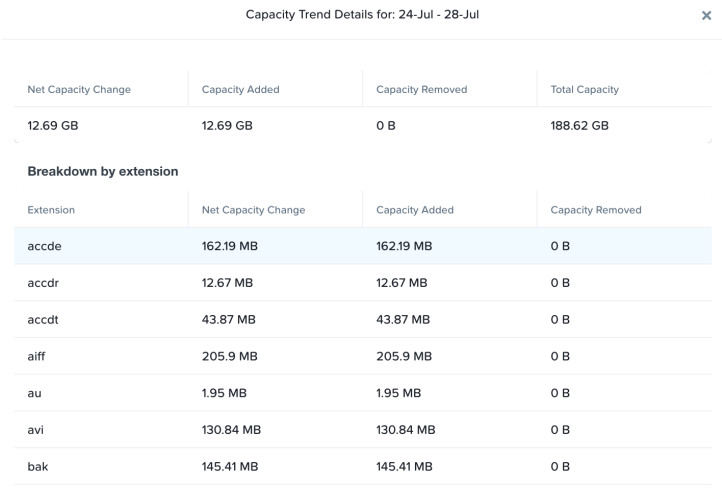

Figure 8: Capacity Trend Details View

Clicking **View Details** for the File Distribution by Type widget displays granular details of file distribution, see the **File Types** table below for details.

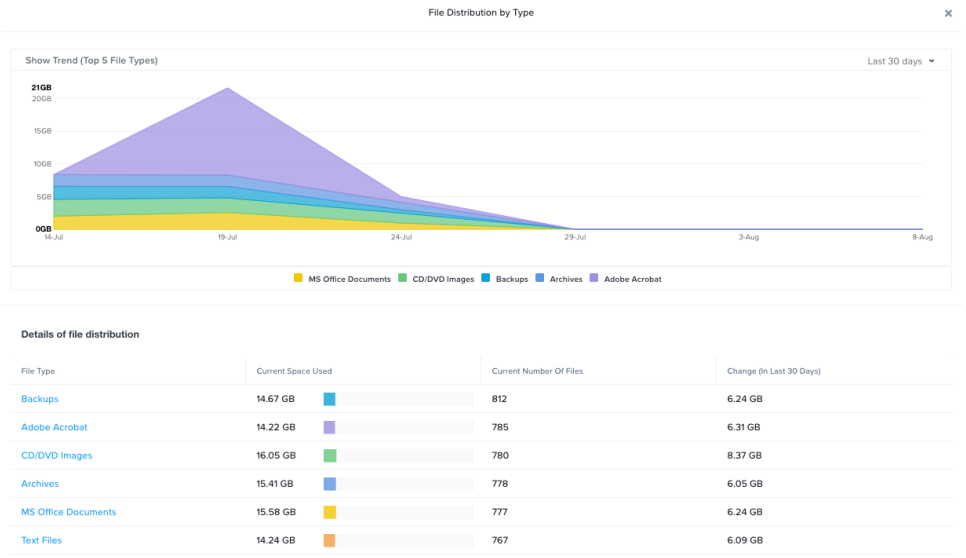

Figure 9: File Distribution by Type

Table 4: File Types

Category	Supported File Type
File Type	Name of file type
Current Space Used	Space capacity occupied by the file type
Current Number of Files	Number of files for the file type
Change (In Last 30 Days)	The increase in capacity over a 30 day period of time for the specified file type .

Table 5: File Types

Category	Supported File Type
Archives	.cab, .gz, .rar, .tar, .z, .zip
Audio	.aiff, .au, .mp3, .mp4, .wav, .wma
Backups	.bak, .bkf, .bkp
CD/DVD Images	.img, .iso, .nrg
Desktop Publishing	.qxd
Email Archives	.pst
Hard Drive images	.tib, .gho, .ghs
Images	.bmp, .gif, .jpg, .jpeg, .pdf .png, .psd, .tif, .tiff,
Installers	.msi, .rpm
Log Files	.log
Lotus Notes	.box, .ncf, .nsf, .ns2, .ns3, .ns4, .ntf
MS Office Documents	.accdb, .accde, .accdt, .accdr, .doc, .docx, .docm, .dot, .d
System Files	.bin, .dll, .exe
Text Files	.csv, .pdf, .txt
Video	.avi, mpg, .mpeg, .mov, .m4v
Disk Image	.hlog, .nvram, .vmdk, .vmx, .vmxf, .vmtm, .vmem, .vmsn, .vms

Updating File Categories

Analytics uses the file category configuration to classify file extensions.

About this task

The capacity widget in the dashboard uses the category configuration to calculate capacity details.

Procedure

1. Click **gear icon** > **Update File Category**.
2. To create a new category, click **+ New Category**. (Otherwise, move on to step 3).
   a. In the **Category** column, name the category.
   b. In the **Extensions** column, specify file extensions for the category.
3. To delete an existing category, click the **x icon** next to the category. (Otherwise, move on to step 4)
4. To modify an existing category, click the **pencil icon** next to the category and modify the specified file extensions.
5. Click **save**.

Health

The Health dashboard displays dynamically updated health information about each File Analytics component.

The Health dashboard includes the following details:

- **Data Summary** Data summary of all file servers with Analytics enabled.
- **Host Memory** Percent of used memory on the Analytics VM (AVM).
- **Host CPU Usage** Percent of CPU used by the AVM.
- **Storage Summary** Amount of storage space used on the Analytics data disk or AVM disk.
- **Overall Health** Overall health of Analytics components.
- **Data Server Summary** Data server usage by component.

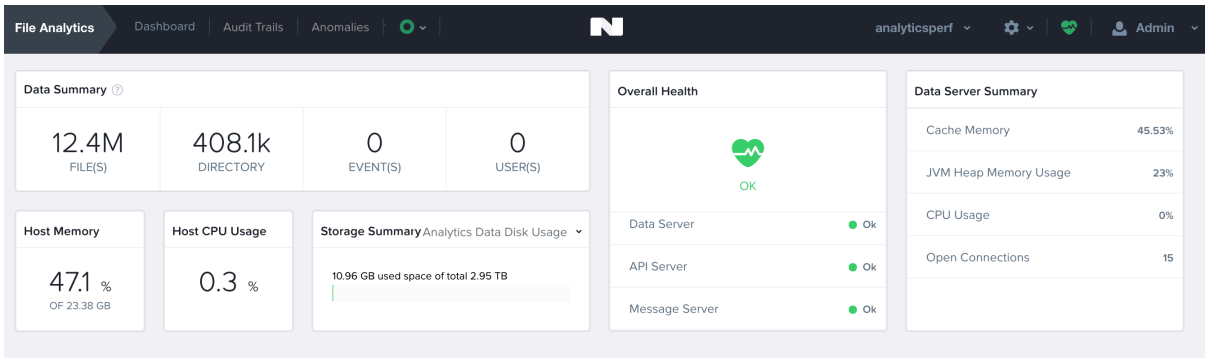

Figure 10: Health Page

4

ANOMALIES

Data panes in the **Anomalies** tab display data and trends for configured anomalies.

The Anomalies tab provides options for creating anomaly policies and dashboards for viewing anomaly trends.

> Note: Configure an SMTP server to send anomaly alerts, see Configuring an SMTP Server on page 24

You can configure anomalies for the following tasks:

- Creating files and directories
- Deleting files and directories
- Permission changes
- Permission denials

You can define anomaly rules for the following conditions:

- Users exceed operation count threshold
- Users exceed the operation percentage threshold

Table 6: Anomalies Data Pane Descriptions

Pane Name	Description	Values
Anomaly Trend	Displays the number of anomalies per day or per month.	Last 7 days, Last 30 days, Last 1 year
Top Users	Displays the users with the most anomalies and the number of anomalies per user.	Last 7 days, Last 30 days, Last 1 year
Top Folders	Displays the folders with the most anomalies and the number of anomalies per folder.	Last 7 days, Last 30 days, Last 1 year
Operation Anomaly Types	Displays the percentage of occurrences per anomaly type.	Last 7 days, Last 30 days, Last 1 year

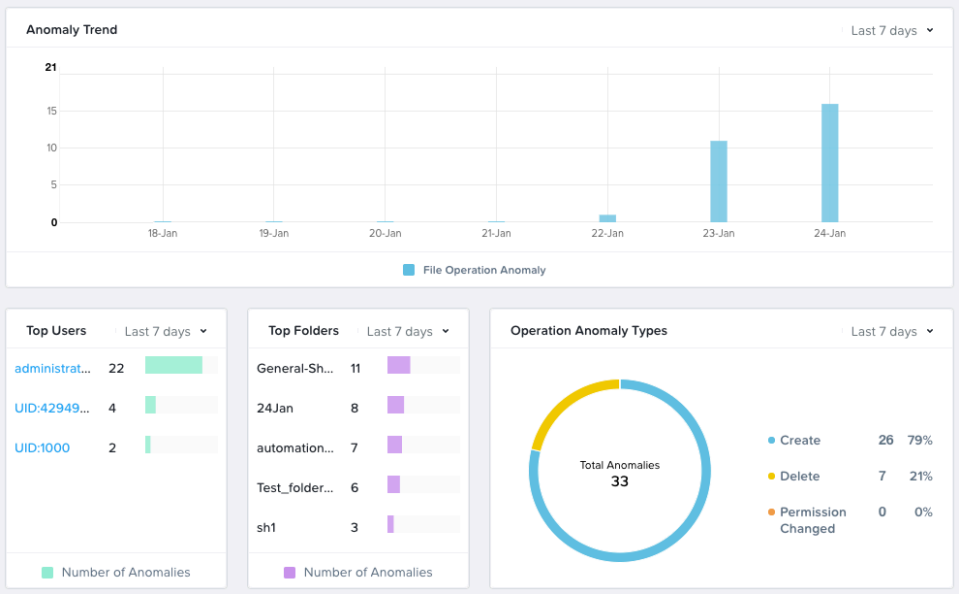

Figure 11: Anomalies Dashboard

Configuring Anomaly Detection

Steps for configuring anomaly rules.

About this task

Configure an SMTP server for Analytics to send anomaly alerts, see Configuring an SMTP Server on page 24. To create an anomaly rule, do the following.

Procedure

1. In the File Analytics web console, click the **gear icon** > **Define Anomaly Rules.** .
2. In the **Anomaly Email Recipients** field, enter a comma-separated list of email recipients for all anomaly alerts and data.

> Note: Analytics sends anomaly alerts and data to recipients whenever Analytics detects an anomaly.

3. To configure a new anomaly, do the following in the indicated fields:

   a. **Events**: Select a rule for the anomaly from one of the following:

      - Permission changed
      - Permission denied
      - Delete
      - Create

      The event defines the scenario type for the anomaly.

   b. **Minimum Operations %**: Enter a percentage value for the minimum threshold.
      Analytics calculates the minimum operations percentage based on the number of files. For example, if there are 100 files, and you set the minimum operations percentage to 5, five operations within the scan interval would trigger an anomaly alert.

   c. **Minimum Operation Count**: Enter a value for a minimum operation threshold.
      Analytics triggers an anomaly alert after meeting the threshold.

   d. **User**: Choose if the anomaly rule is applicable for **All Users** or an **Individual** user.

   e. **Type:** the type determines the interval.
      The interval determines how far back Analytics monitors the anomaly.

   f. **Interval** : Enter a value for the detection interval.

   g. (optional) **Actions**: Click the **pencil icon** to update an anomaly rule. Click the **x** icon to delete an existing rule.

Figure 12: Anomaly Configuration Fields

4. Click **Save**.

Configuring an SMTP Server

File Analytics uses a simple mail transport protocol (SMTP) server to send anomaly alerts.

About this task

To configure an SMTP server, do the following:

Procedure

1. In the File Analytics web console, click the **gear icon** > **SMTP Configuration**.

2. In the SMTP Configuration window, enter the indicated details in the following fields:

   a. **Hostname Or IP Address**: Enter a fully qualified domain name or IP address for the SMTP server.

   b. **Port**: Enter the port to use.
   The standard SMTP ports are 25 (encrypted), 587 (TLS), and 465 (SSL).

   c. **Security Mode**: Enter the desired security mode from the dropdown list.
   The options are:

      - **NONE** (unencrypted)
      - **STARTTLS** (TTL encryption)
      - **SSL** (SSL encryption)

   d. (If security mode is "NONE" go to step f.)

   e. **User Name** enter a user name for logging into the SMTP server. Depending on the authentication method, the user name may require a domain.

   f. **Password** enter password.

   g. **From Email Address:** enter the email address from which Analytics will send the anomaly alerts.

   h. **Recipient Email Address:** enter a recipient email address to test the SMTP configuration.

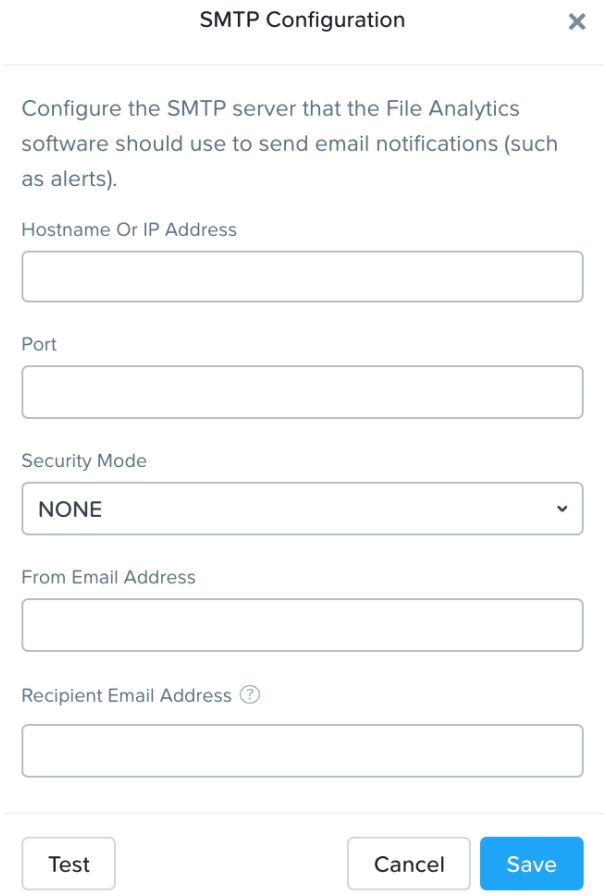

Figure 13: SMTP Configuration

3. Click **Save**.

5

AUDIT TRAILS

Use audit trails to look up operation data for a specific user, file, folder, or client.

The **Audit Trails** tab includes **Files**, **Folders**, **Users**, and **Client IP** options for specifying the audit type. Use the search bar for specifying the specific entity for the audit (user, folder, file, or client IP).

The results table presents details for entities that match the search criteria. Clicking the entity name (or client IP number) takes you to the Audit Trails dashboard for the target entity.

View Audit Trails

Audit a user, file, client, or folder.

About this task

Procedure

1. Click the **Audit Trails** tab.
2. Select the **Files**, **Folders**, **Users**, or **Client IP** option.
3. Enter the audit trails target into the search bar.
4. Click **Search**.
5. To display audit results in the Audit Trails window, click the entity name (or client IP number).

Audit Trails - Users

Details for client IP Audit Trails.

Audit Trails Search - Users

When you search by user in the **Audit Trails** tab, search results display the following information in a table.

- User Name
- Domain
- Last Operation
- Last Operation On
- Share Name
- Operation Date
- Action

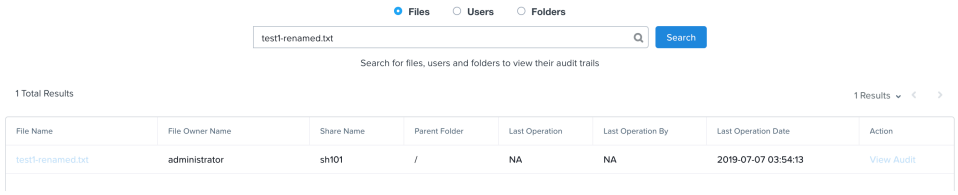

Figure 14: Users Search Results

Audit Details Page - Users

Clicking **View Audit** displays the **Audit Details** page, which shows the following audit information for the selected user.

- A **User Events** graph displays various operations the user performed during the selected period and the percentage of time each operation has occurred per total operations during the specified period.
    - The **Filter by operations** dropdown contains operation filters, which you can use to filter the audit by operation type. Operations include:
        - Create File
        - Delete
        - Make Directory
        - Permission Changed
        - Permission Denied
        - Read
        - Remove Directory
        - Rename
        - Set Attribute
        - Write
        - Symlink
    - The **filter bar**, above the User Events graph, displays the filters in use.
    - Use the **From** and **To** fields to filter by date.
- The **Results** table displays operation-specific information. See more details below.
- The **Reset Filters** button removes all filters.

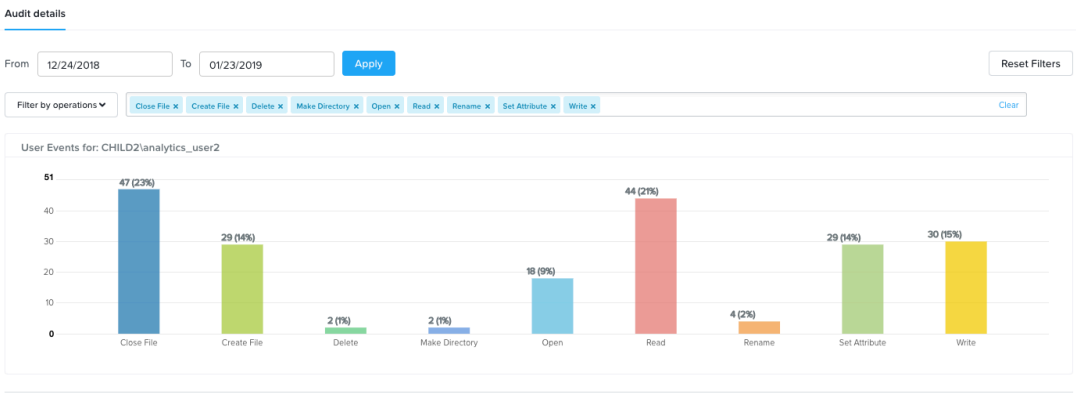

Figure 15: User Audit Details - Events

The **Results** table provides granular details of the audit results. The following data is displayed for every event.

- User Name
- User IP Address
- Operation
- Operation Date
- Target File

Click the **gear icon** for options to download the data as an xls, csv, or JSON file.

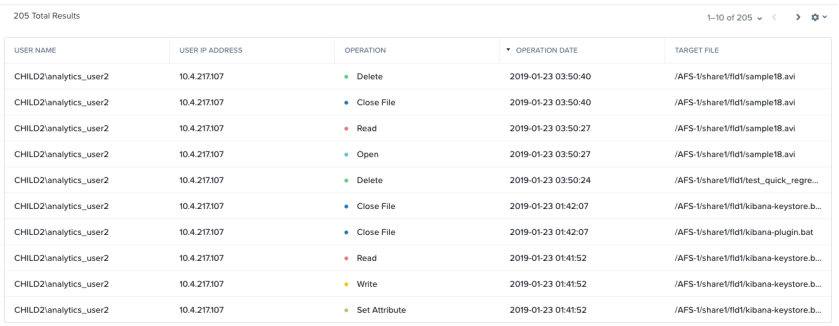

Figure 16: Results Table

Audit Trails - Folders

Dashboard details for folder audits.

The following information displays when you search by file in the Audit Trails tab.

- Folder Name
- Folder Owner Name
- Share Name
- Parent Folder
- Last Operation

- Last Operation By
- Last Operation Date
- Action

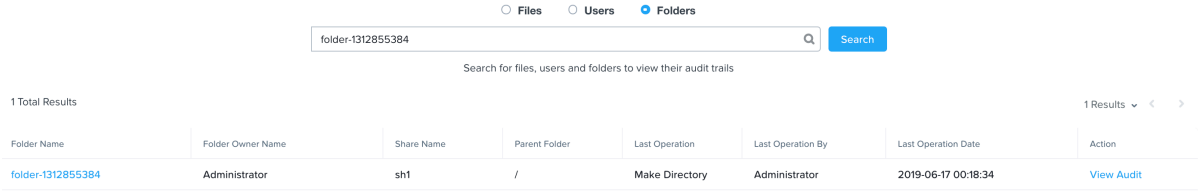

Figure 17: Folders Search Results

The **Audit Details** page shows the following audit information for the selected file.

- A **Folder Events** graph displays various operations performed on the file during the selected period, and the percentage of time each operation has occurred per total operations during that period.
    - The **Filter by operations** dropdown contains operation filters, which you can use to filter the audit by operation type. Operations include:
        - Select All
        - Make Directory
        - Permission Changed
        - Permission Denied
        - Remove Directory
        - Rename
        - Set Attribute
    - A **filter bar**, above the File Events graph displays the filters in use.
    - Use the **From** and **to** fields to filter by date.
- The **Results** table displays operation-specific details.
- The **Reset Filters** button removes all filters.

The **Results** table provides granular details of the audit results. Analytics displays the following data for every event.

- User Name
- Client IP
- Operation
- Operation Date

Click the gear icon for options to download the data as a CSV file.

Audit Trails - Files

Dashboards details for file audit.

Audit Trails for Files

When you search by file in the Audit Trails tab, the following information displays:

- File Name
- File Owner Name
- Share Name
- Parent Folder
- Last Operation
- Last Operation By
- Last Operation Date
- Action

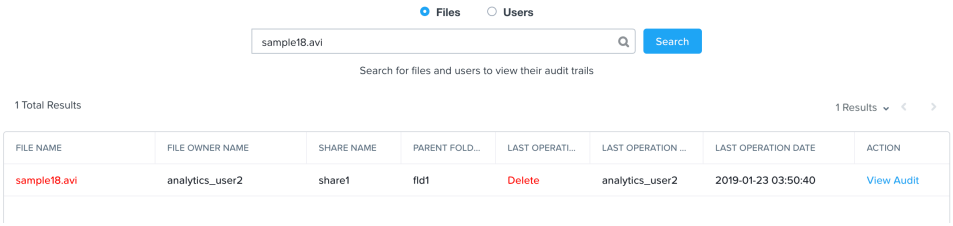

Figure 18: Files Search Results

The **Audit Details** page shows the following audit information for the selected file.

- A **File Events** graph displays various operations performed on the file during the selected period, and the percentage of time each operation has occurred per total operations during that period.
    - The **Filter by operations** dropdown contains operation filters, which you can use to filter the audit by operation type. Operation types include:
        - Close File
        - Create File
        - Delete
        - Make Directory
        - Open
        - Read
        - Rename
        - Set Attribute
        - Write
        - Symlink
    - A **filter bar**, above the File Events graph displays the filters in use.
    - Use the **From** and **to** fields to filter by date.

- The **Results** table displays operation-specific details.
    - The **Reset Filters** button removes all filters.

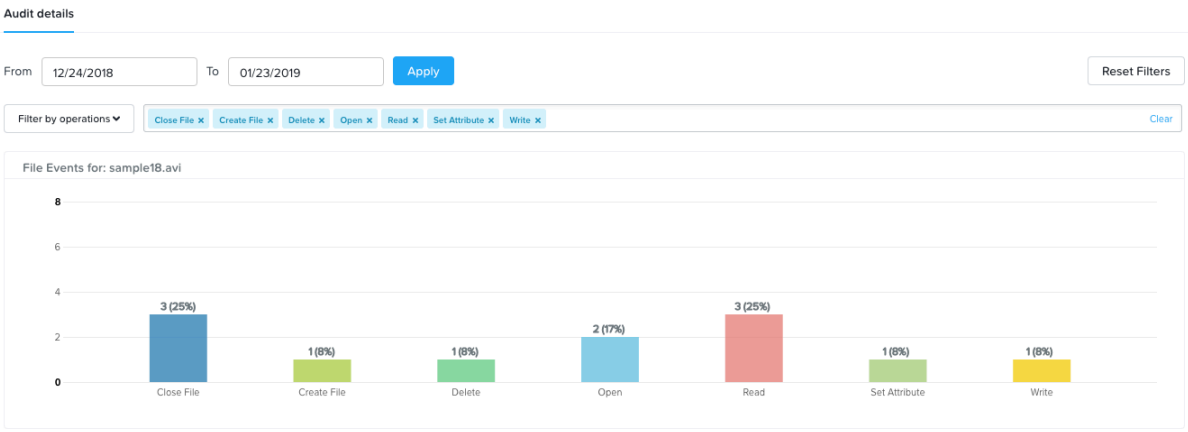

Figure 19: Files Audit Details - Events

The **Results** table provides granular details of the audit results. Analytics displays the following data for every event.

- Username
- Client IP
- Operation
- Operation Date

Click the gear icon for options to download the data as a CSV file.

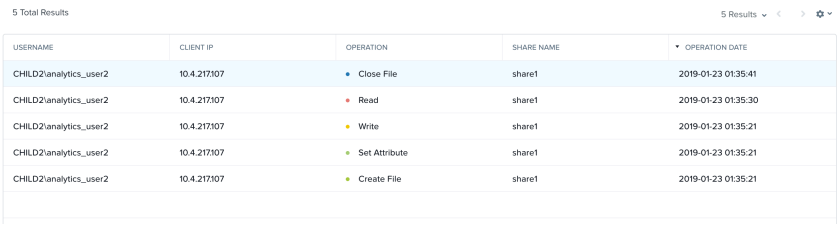

Figure 20: Results Table

Audit Trails - Client IP

Dashboard details for client IP Audit Trails.

Audit Trails Search - Client IP

When you search by client IP in the **Audit Trails** tab, search results display the following information in a table.

- Client IP
- User Name

- Domain
- Last Operation
- Last Operation On
- Share Name
- Operation Date
- Action

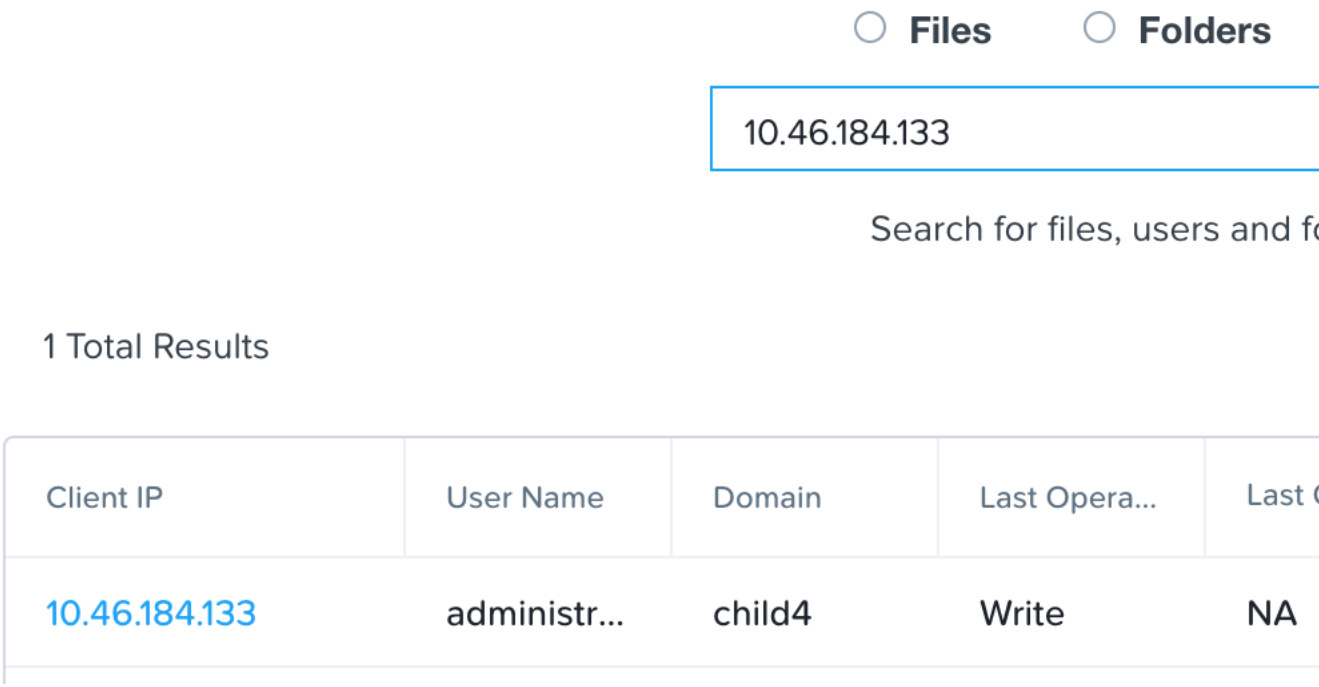

Figure 21: IP Search Results

The **Audit Details** page shows the following audit information for the selected client.

- A **User Events** graph displays various operations performed on the client during the selected period, and the percentage of time each operation has occurred per total operations during that period.
    - The **Filter by operations** dropdown contains operation filters, which you can use to filter the audit by operation type. Operation types include:
        - Create File
        - Delete
        - Make Directory
        - Permission Changed
        - Permission Denied
        - Read
        - Removed Directory
        - Rename
        - Set Attribute
        - Write
        - Symlink
        - Permission Denied (File Blocking)
    - A **filter bar**, above the File Events graph displays the filters in use.
    - Use the **From** and **to** fields to filter by date.
- The **Results** table displays operation-specific details.
    - The **Reset Filters** button removes all filters.

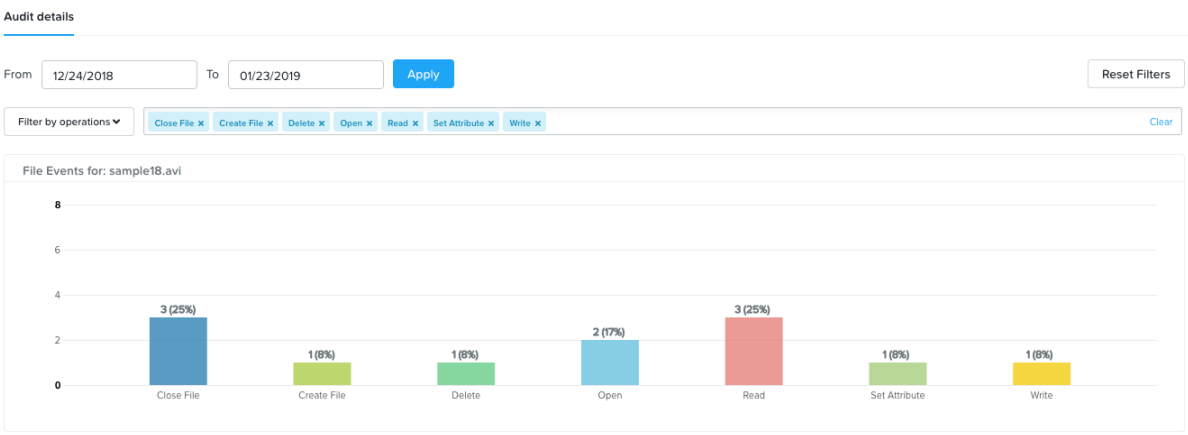

Figure 22: Files Audit Details - Events

The **Results** table provides granular details of the audit results. Analytics displays the following data for every event.

- User Name

- Operation
- Target File
- Operation Date

Click the gear icon for options to download the data as a CSV file.

6

ANALYTICS OPTIONS

You can get more insight into the usage and contents of files on your system by configuring and updating Analytics features and settings. Some options include scanning the files on your file server on demand, updating data retention, and configuring data protection.

Updating Data Retention

The data retention period determines how long File Analytics retains event data.

About this task

Follow the steps as indicated to configure data retention.

Procedure

1. In File Analytics, click **gear icon** > **Update Data Retention**.
2. In the **Data Retention Period** drop-down, select the period for data retention.
3. Click **Update**.

Scanning the File System

Once enabled, File Analytics scans the metadata of all files and shares on the system. You can perform an on-demand scan of shares in your file system.

About this task

To scan shares, perform the following task:

Procedure

1. In File Analytics, click the **gear icon**.

2. In the drop-down list, click **Scan File System**.

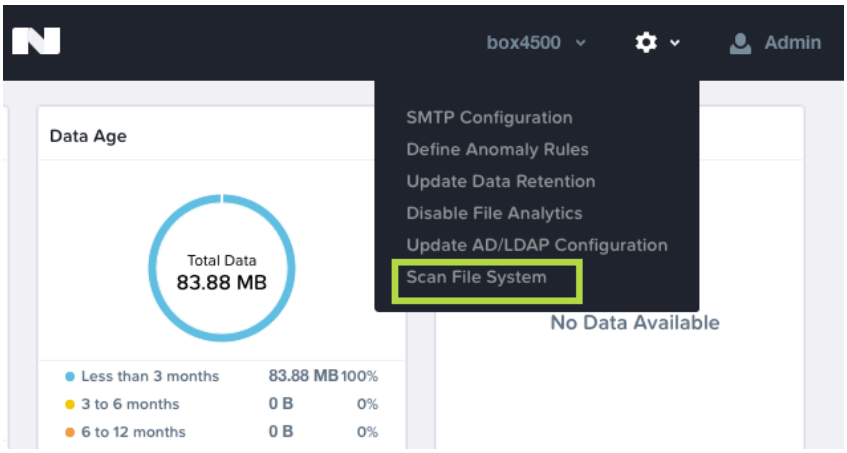

Figure 23: Scan File System Option

3. In the list of shares, select the target shares for the scan.

Figure 24: Select Scan Targets

4. Click **Scan**.
   The status of the share is **In Progress**. Once the scan is complete, the status changes to **Completed**.

Blacklisting

Blacklist users, file extensions, and client IPs.

About this task

Use the blacklisting feature to block audit events from being performed on specified file extensions or by specified users and clients.

Procedure

1. Click the **gear icon** > **Define Blacklisting Rules**.
2. Click the pencil icon in the user, file extension, or client IP row.
3. Add a comma separated list of entities that you want blocked.
4. Click save in the updated row.

Data Protection

Configure File Analytics disaster recovery (DR) using Prism Element.

Analytics only supports async disaster recovery. Analytics does not support NearSync and metro availability.

Create an async protection domain, configure a protection domain schedule, and configure remote site mapping. The remote site must have symmetric configurations to the primary site. The remote site must also deploy File Analytics to restore a File Analytics VM (AVM).

The *Data Protection* section in the *Prism Web Console Guide* provides more detail on the disaster recovery process.

Configuring Disaster Recovery

To set up disaster recovery for File Analytics, create an async protection domain, configure a protection domain schedule, and configure remote site mapping.

About this task

By default, the File Analytics volume group resides on the same container that hosts vDisks for Nutanix Files.

Procedure

1. If you have not done so already, configure a remote site for the local cluster.
   See the *Configuring a Remote Site (Physical Cluster)* topic in the *Prism Web Console Guide* for this procedure.
2. Create an async DR protection domain for the File Analytics volume group as the entity. The volume group name is `File_Analytics_VG`.
   See *Configuring a Protection Domain (Async DR)* in the **Prism Web Console Guide**.
3. In the **Schedule** tab, click the **New Schedule** button to add a schedule.
   Add a schedule, as Analytics does not provide a default schedule. See *Creating a Protection Domain Schedule (Files) Nutanix Files Guide.*
4. Configure local and remote container mapping.
   See the *Configuring Disaster Recovery (Files)* section in the *Nutanix Files Guide* for steps to configure mapping between local and remote containers.
5. Create a protection domain schedule.
   See *Creating a Protection Domain Schedule (Files)* in the *Nutanix Files Guide*.

Activating Disaster Recovery

Recover a File Analytics VM (AVM) after a planned or unplanned migration to the remote site.

About this task

Perform the following tasks on the remote site.

Procedure

1. Fail over to the protection domain for disaster recovery activation.
   See the *Failing Over a Protection Domain* topic in the *Prism Web Console Guide*.
2. Fail back the protection domain to the primary site.
   See the *Failing Back a Protection Domain* topic in the *Prism Web Console Guide*.

Deploying Analytics on a Remote Site

Deploy a File Analytics VM (AVM) after a planned or unplanned (disaster) migration to the remote site.

About this task

To perform disaster recovery, deploy and enable Analytics on the remote site. Restore the data using a snapshot of the volume group from the primary AVM.

Before you begin

Ensure that the Nutanix Files and AOS versions match the versions on the remote and primary sites.

About this task

Run the following commands from the command prompt inside the AVM.

Procedure

1. Deploy a new File Analytics instance on the remote site, see Deploying File Analytics on page 7.

   > CAUTION: Do not enable File Analytics.

   The remote site requires an iSCSI data service IP address to configure the AVM on the remote site.

2. On the remote site, create a volume group by restoring the snapshot of the `File_Analytics_VG`.

   See *Restoring an Entity from a Protection Domain* in the *Prism Web Console Guide*.

3. In the **Storage Table** view, go to the **Volumes** tab.

   a. Copy the target IQN prefix from the **Volume Group Details** column.

4. To configure the AVM on the remote, follow these steps:

   a. Log on to the AVM with SSH using the following credentials: username `nutanix` and password `nutanix/4u`.

   b. To discover all storage devices accessed by the AVM, run the following commands.

   ```
   nutanix@avm$ sudo blkid
   ```

   c. Copy the `cvm.config` file to the temporary files directory.

   ```
   nutanix@avm$ sudo cd /mnt/containers/config/common_config/
   ```

```
nutnix@avm$ sudo cp cvm.config /tmp
```

d. Log off from all iSCSI targets.

```
nutnix@avm$ sudo /sbin/iscsiadm -m node -u
```

e. Remove the disconnected target records from the discoverydb mode of the AVM.

```
nutanix@avm$ sudo /sbin/iscsiadm -m node –o delete
```

f. Discover all storage devices accessed by the AVM.

```
nutanix@avm$ sudo blkid
```

```
/dev/sr0: UUID="2019-06-11-12-18-52-00" LABEL="cidata" TYPE="iso9660"
/dev/sda1: LABEL="_master-x86_64-2" UUID="b1fb6e26-a782-4cf7-b5de-32941cc92722"
 TYPE="ext4"
```

The output does not show the **/dev/sdb** device.

g. Get the File Analytics Linux client iSCSI initiator name.

```
nutanix@avm$ sudo cat /etc/iscsi/initiatorname.iscsi
```

The output displays the initiator name.

```
InitiatorName=iqn.1991-05.com.redhat:8ef967b5b8f
```

h. Copy the iSCSI initiator name.

i. To configure the cloned volume group in Prism, whitelist the AVM client using the iSCSI initiator or IP address of the AVM client.

For this procedure, see the *Enabling Nutanix Volumes* topic in the *Prism Web Console Guide*.

j. Run the command to let the Analytics initiator discover the cluster and its volume groups.

```
nutanix@avm$ sudo /sbin/iscsiadm --mode discovery --type sendtargets --portal
 data_services_IP_address:3260
```

Clicking the Nutanix cluster name in Prism displays cluster details including the data service IP address.

The output displays the iSCSI target from step 2.

k. Connect to the volume target by specifying IQN prefix.

```
nutanix@avm$ sudo /sbin/iscsiadm --mode node --targetname iqn_name --portal
 data_services_IP_address:3260,1 --login
```

l. Restart the AVM to restart the iSCSI host adapters, which allows the discovery of the attached volume group.

```
nutanix@avm$ sudo reboot
```

m. Log on to the AVM with SSH using the following credentials: username **nutanix**, and password **nutanix/4u**.

n. Discover all storage devices accessed by the AVM.

```
nutanix@avm$ sudo blkid
```

The AVM discovers the attached iSCSI volume group and assigns to the **/dev/sdb** device.

```
/dev/sr0: UUID="2019-06-11-12-18-52-00" LABEL="cidata" TYPE="iso9660"
```

```
/dev/sda1: LABEL="_master-x86_64-2" UUID="b1fb6e26-a782-4cf7-b5de-32941cc92722"
 TYPE="ext4"
/dev/sdb: UUID="30749ab7-58e7-437e-9a09-5f6d9619e85b" TYPE="ext4"
```

o. Create a backup of the **cvm.config** file.

```
nutanix@avm$ sudo cd /mnt/containers/config/common_config/
nutanix@avm$ sudo mv cvm.config cvm_bck.config
```

p. Copy the **cvm.config** file from the **/tmp** directory to **/common_config/** on the AVM.

```
nutanix@avm$ sudo cd /tmp
nutanix@avm$ sudo mv cvm.config /mnt/containers/config/common_config/
```

COPYRIGHT

Copyright 2020 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.

Default Cluster Credentials

Interface	Target	Username	Password
Nutanix web console	Nutanix Controller VM	admin	Nutanix/4u
vSphere Web Client	ESXi host	root	nutanix/4u
vSphere client	ESXi host	root	nutanix/4u

Interface	Target	Username	Password
SSH client or console	ESXi host	root	nutanix/4u
SSH client or console	AHV host	root	nutanix/4u
SSH client or console	Hyper-V host	Administrator	nutanix/4u
SSH client	Nutanix Controller VM	nutanix	nutanix/4u
SSH client	Nutanix Controller VM	admin	Nutanix/4u
SSH client or console	Acropolis OpenStack Services VM (Nutanix OVM)	root	admin

Version

Last modified: March 23, 2020 (2020-03-23T07:56:48+05:30)

Objects 2.1

Objects User Guide

May 11, 2020

Contents

1. Nutanix Objects Overview....4
   - Usage of Objects....4
   - Terminology Reference....4
   - Objects Workflow....5

2. Enabling Objects....6
   - Objects License Management....6
     - Open Source Software Usage....7

3. Deployment and Network Prerequisites....8
   - Deployment Prerequisites....8
   - Network Configurations....14
   - Limitations....15

4. Object Store Service Deployment....16
   - Object Store Instance Naming Conventions....16
   - Creating or Deploying an Object Store (Prism Central)....17
   - Viewing Object Store Deployments....21
     - Access Objects Endpoints....22
   - Deleting an Object Store....23
   - Object Store Storage Expansion....24
     - Expanding Storage for an Object Store....24
   - Scaling Out an Object Store....27
   - Setting up SSL Certificate for an Object Store....30
   - Object Store Deployment at a Dark Site (Offline Deployment)....32

5. Types of Objects Users....35

6. Directory Configuration and Access Key Generation....36
   - Configuring Directories....36
   - Generating Access Key for API Users....39
   - Viewing API Users....42
     - Managing API Keys....43
     - Deleting API Users....43

7. Bucket Creation, Operations and Bucket Policy Configuration....45
   - Bucket Naming Conventions....45
   - Creating and Configuring a Bucket....45
   - Bucket Policy Configuration....47
     - Object Versioning....47
     - Lifecycle Policies....47
     - WORM Bucket....48
     - Configuring a Bucket for Static Website Hosting....50

Cross-Origin Resource Sharing (CORS) Overview 52
Viewing Buckets 54
Buckets Filter Options 54
Updating a Bucket 55
Sharing a Bucket 56
Bucket Access Policies 57
Listing the Shared Buckets 60
Viewing Bucket Users 61
Deleting a Bucket 62

8. Objects Replicator 64

9. Monitoring and Alerts 65
Viewing Performance of Object Stores 65
Viewing Performance of Buckets 66
Viewing Object Store Usage 67
Assigning Quota Policy to a User 68
Viewing Buckets Usage 69
Viewing Alerts 70
Nutanix Objects Specific Alerts 71

10. CRUD Operations by Using S3 APIs 75
Authentication 75
Supported and Unsupported APIs 75
Supported APIs 75
Unsupported APIs 78
Objects Tagging APIs Overview 79

11. Error Responses 81
REST Error Responses 81
List of Error Codes 81

12. Integration with Backup Applications 84

13. Objects LCM Upgrades 85
Microservices Platform (MSP) 86
Upgrading Objects Manager 86
Upgrading Objects Service 87

14. Troubleshooting Objects 89

1

NUTANIX OBJECTS OVERVIEW

Nutanix Objects™ (Objects) is a software-defined Object Store Service being introduced by Nutanix. This service is designed with an Amazon Web Services Simple Storage Service (AWS S3) compatible REST API interface to handle petabytes of unstructured and machine-generated data. Objects addresses storage-related use cases for backup, storage, long-term retention, and enables you to develop cloud-native applications by using standard S3 APIs. With Objects, the current storage offering of the Nutanix solution expands. You do not have to introduce extra storage solutions when you use Objects as it is deployed and managed as part of the Nutanix Enterprise Cloud OS.

Buckets are logical constructs in which you can store objects, which consist of data and metadata that describes the data. You can manage these buckets and objects by using Prism Central or 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*.

Usage of Objects

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

- Backup – You can integrate Objects with the back up applications such as Commvault, Commtrade 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 storage and long-term data retention. You can use built-in object versioning for deeper storage protection and search your data without the problem of tape systems. You can also maintain previous versions of the object and avoid loss from accidental overwrites or deletes. You can meet regulatory compliance for healthcare, financial, and government sectors by enabling Write-Once-Read-Many (WORM) buckets so your data is not overwritten or deleted.
- DevOps – You can use a single global namespace to PUT and GET objects with over-the-network HTTP and HTTPS commands. Objects enables you to integrate REST API calls within your programs or scripts without tracking complex directory structures. DevOps and IT ops can use an S3-compatible interface for cross-geo, cross-team collaboration, and agile development.

Terminology Reference

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

Table 1: Terminology Reference

Terminology	Description
Bucket	An organizational unit exposed to the users and contains the objects. A deployment may have multiple buckets. For example, departmental and compartmental.
Object	The actual unit (blob) of storage and the item interfaced by using the API (GET or PUT).
S3	The term used to describe the original object service introduced by Amazon Web Services (AWS). This term is now used synonymously for an object service. S3 is also used to define the object API which can be leveraged throughout the projects.
Internal Access Network	A VLAN required internally for deploying Object Store Services on Prism Element.
Client Access Network	A VLAN required for accessing the Object Store endpoints externally.
Microservices Platform (MSP)	A platform based on Kubernetes where all the Objects microservices run.

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 6.
2. Deploy object store on a desired cluster. Refer to Object Store Service Deployment on page 16.
3. Generate the access keys. Refer to Generating Access Key for API Users on page 39.
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 30.
5. Access the object store endpoints from the third-party clients. Refer to Access Objects Endpoints on page 22.
6. Create buckets and objects suitable for your use case. Refer to Creating and Configuring a Bucket on page 45.
7. Expand the object store if the storage is getting full. Refer to Expanding Storage for an Object Store on page 24.

2

ENABLING OBJECTS

For using Objects services, you need to enable Objects in Prism Central, and add license in Prism Element.

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.

   You can now use the Objects services.

Objects License Management

Nutanix provides licenses that you can apply manually or automatically 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.

Note:

- You can now manage Objects license from the Prism Central, not from the Prism Element.
- Up to 2 TB Objects storage is free for all users.

Types of Licenses

Nutanix provides following two types of Objects licenses.

- Objects (For AOS): Add this license if you want to deploy Objects on the Prism Central with AOS licenses.
- Objects (Dedicated): Add this license if you want to deploy Objects on the Prism Central without AOS licenses.

Adding Licenses to Prism Central

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

You can view the Objects usage information in the Prism Central. For more information, refer to 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.

3

DEPLOYMENT AND NETWORK PREREQUISITES

Ensure that you have met the deployment requirements and configured the network before proceeding for the object store deployment.

Deployment Prerequisites

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

General Requirements

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

- Hypervisor is AHV.
- Prism Element and Prism Central version 5.11.2 or later running in your environment.
- Recommended browser: Google Chrome
- Minimum of three nodes in a cluster running AHV with 12 vCPUs and 36 GiB of memory in each node (recommended). You can also have a single node cluster running AHV with a minimum of 12 vCPUs and 36 GiB of memory.

  Note:

  - It is not recommended to have one node or two node (worker nodes) deployments for production clusters. One or two node deployments are for test or POC use only.
  - Objects use no more than 12 vCPUs for each AHV node.

- Ensure that no AHV 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 `testobjectstore.mycompany.com`.

- Firewall must be enabled in Prism Element.
- Recommended that proxy should be able to reach the guest VM.
- Recommended to enable Pulse (only for site with Internet access).

- Recommended to have a high speed and stable Internet connection (only for site with Internet access).

  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 86.
- 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 1.0.5 or later versions and Objects 2.0 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.
- 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 Configurations on page 14 section.
- 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 32.

URL and Port Requirements

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

The following URLs and ports 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*.

- Open the following ports:

> Note: Port 80 should be opened between the Prism Central and web server, and Prism Element and web server for dark site only.

Table 2: Port Requirements

Protocol	Port Number	Source	Destination	Directionality	Description
TCP	7301	Prism Central hosting Objects	Prism Central hosting Objects	Bi-directional	This port is opened on the Objects Manager running in the Prism Central to receive the manageability requests. This port is the proxy to other ports.
TCP	7302	Prism Central hosting Objects	Prism Central hosting Objects	Bi-directional	This port is opened on the Objects Manager running in the Prism Central to receive the manageability requests. This port is the proxy to Object Manager Leader in scale-out Prism Central.
TCP	81	Prism Central hosting Objects	Load balancer VM (Internal and external network)	Uni-directional	This port is opened on the load balancer that directs traffic from the Objects Manager to the UI Gateway.

Protocol	Port Number	Source	Destination	Directionality	Description
TCP	80	External S3 clients	Load balancer VM (External network)	Uni-directional	This port is the HTTP endpoint to access Objects. This port is opened on the load balancer that directs traffic to the Object Controller pod.
TCP	443	External S3 clients	Load balancer VM (External network)	Uni-directional	This port is the HTTPS endpoint to access Objects. This port is opened on the load balancer that directs traffic to the Object Controller pod.
TCP	2009	Worker VMs	Prism Element	Uni-directional	This port is opened for the data traffic from Object Controller pod to Stargate. The data services IP address needs to be permitted for this port.
TCP	22, 53, 443, 2379, 5000	Prism Central hosting Objects	Worker VMs	Uni-directional	These ports are opened for the management traffic from Prism Central to Worker VMs.

Protocol	Port Number	Source	Destination	Directionality	Description
TCP	3205, 3260	Worker VMs	Prism Element	Uni-directional	These ports are opened for the management traffic from Worker VMs to Prism Element. The data services IP address needs to be permitted for these ports.
TCP	9440	Worker VMs	Prism Central hosting Objects	Uni-directional	This port is opened for the V3 API calls from the Worker VMs to Prism Central.
		Worker VMs	Prism Element	Uni-directional	This port is opened for the V2 API calls from Worker VMs to Prism Element.
TCP/UDP	Any port	Worker VMs	Worker VMs	Bi-directional	This port is opened for the data traffic from Worker VMs to Worker VMs. Within the network, all ports should be reachable.
TCP	31002	Prism Central hosting Objects	Worker VMs	Uni-directional	This port is opened on the Worker VMs for the service registry registration.

Protocol	Port Number	Source	Destination	Directionality	Description
TCP	5553, 81	Prism Central hosting Objects, Worker VMs	Load balancer VMs (Internal and external network)	Uni-directional	This port is opened for the traffic from the Worker VMs and the service manager running on Prism Central to IAM (Identity and Access Manager) service through load balancers.
TCP, UDP	53	Worker VMs	Name Servers	Uni-directional	This port is opened on the Name Servers for the domain name resolution.
TCP, UDP	53	Prism Central hosting Objects	Name Servers	Uni-directional	This port is opened on all the Name Servers to receive the TCP packets from the Prism Central and used to determine the health of the Name Servers.
UDP	123	Worker VMs	NTP Servers	Uni-directional	This port is opened on the NTP Servers to sync time on the worker VMs. This is not required if NTP Server configured on the Prism Central is local.

Network Configurations

You can configure and manage VLANs through Prism Element, and use these VLANs for deploying Object Store through Prism Central. These configurations are required for both online and offline deployments.

Internal Access Network is a VLAN (AHV IPAM enabled) for internal communication between the different components of the object store services. **Client Access Network** is a VLAN (AHV IPAM enabled) that is required for accessing the Object Store endpoints externally. Once you have successfully added these VLANs in Prism Element, you can use these VLANs for deploying object store through Prism Central.

For more information about network configurations, refer to *Network Configuration for VM Interfaces* in *Prism Web Console Guide*.

The following section describes each of these VLANs in more detail:

- **Internal Access Network**: Object Store uses this VLAN internally to deploy its services. This VLAN must have a minimum of 18 IP addresses in usable IP address range.

  Make note of the following requirements about the VLAN configured on the AHV cluster that will run the Object Store.

  - VLAN used should be AHV IPAM enabled. For more information, see *Network Configuration for VM Interfaces* in the Prism Web Console Guide.
  - At Least 12 IPs in the DHCP pool from this VLAN is required.
  - Two static IPs for every deployment outside the DHCP pool is required, but in the same VLAN. Later while deploying object store from Prism Central, you can use two of these static IP addresses for the Internal Access Network configuration based on the object store storage and resource requirements.

    Note: Internal Access Networks cannot overlap with `10.100.0.0/16` and `10.200.0.0/16` networks.

- **Client Access Network**: Object Store uses this VLAN (AHV IPAM enabled) for setting up the Object Store endpoints that are accessible by the users.

  Make note of the following requirements about the VLAN configured on the AHV cluster that will run the Object Store.

  - VLAN used should be AHV IPAM enabled. For more information, see *Network Configuration for VM Interfaces* in the Prism Web Console Guide.
  - At least four static IP addresses in the usable IP address range. Later while deploying object store from Prism Central, you can use four of these IP addresses for the client access network configuration based on the object store storage and resource requirements.

    Note: You can have two VLANs, each for Internal Access Network and Client Access Network, but it is not required. You can have the Internal Access Network and the Client Access Network on the same VLAN. However, it is recommended to have the Internal Access Network and the Client Access Network on different VLANs for production deployments.

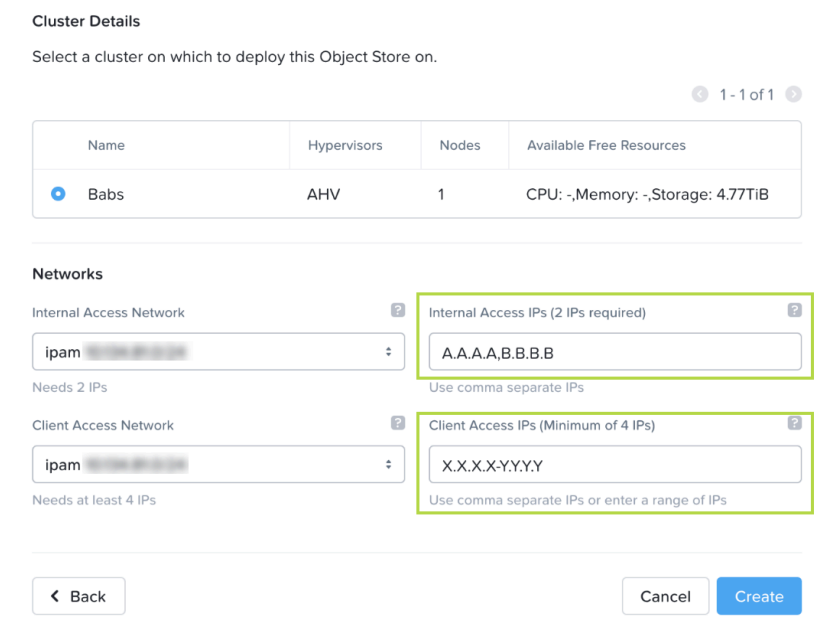

Figure 1: Network Configuration - Internal Access IPs and Client Access IPs

Note that all the IPs may not get used during the deployment. The number of IPs used depends on the size of your deployment. The unused IPs 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 Client Access IPs for your deployment. For more information, see Viewing Object Store Deployments on page 21.

Limitations

The following section lists the limitations for Objects.

System Limitations

- Maximum 240 TB HDD tier for each node is supported.
- Data Services IP address change is not supported.
- Controller VM, Microservices Platform (MSP) and Prism Central IP addresses change are not supported.
- Prism Central and Prism Element deregistration and reregistration is not supported.

4

OBJECT STORE SERVICE DEPLOYMENT

Objects is a highly available and distributed Object Store Service used for redundant and scalable data storage, which stores 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 17. You can also perform offline deployment (site without Internet access) of object store. Refer to Object Store Deployment at a Dark Site (Offline Deployment) on page 32.

You can deploy multiple Object Store Services 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 object store service:

1. Configure the directory and generate the access key. Refer to Directory Configuration and Access Key Generation on page 36.
2. Create buckets within the object store. Refer to Creating and Configuring a Bucket on page 45
3. Upload objects and meta-data to the buckets by using the S3 APIs. Refer to Supported APIs on page 75.

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.

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.
- Can contain alphanumeric or hyphen characters.
- Begin with a letter, and end with a letter or number.
- Cannot contain special characters.
- 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 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.

Before you begin

Make sure you satisfy the deployment prerequisites before starting the deployment. Refer to Deployment Prerequisites on page 8.

About this task

Note:

- Objects 2.0 new deployments are only supported in AOS 5.11.2 and later versions.
- You cannot deploy a single Object Store Service across multiple Prism Elements.
- Deployment can take a minimum of 30 minutes.

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 Prerequisites window, click **Continue** if you fulfil the prerequisites.
4. On the **Name** page, type an object store name and domain for reference, and then click **Next**.
   For guidelines on choosing a compliant name, refer to Object Store Naming Conventions.

5. On the Configure page, you can configure the following properties.

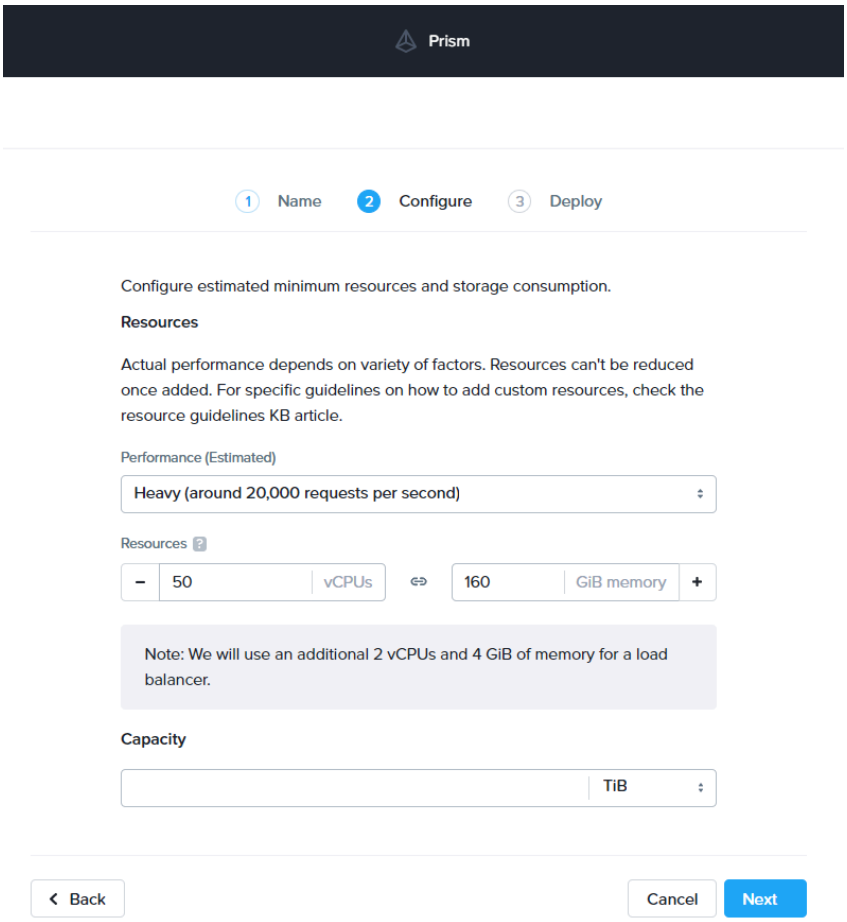

Figure 2: Configure an Object Store

a. **Resources**: Define the resources (vCPUs and memory) based on the desired performance.

> Note:
>
> - 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 adds 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.

The resources field change based on any of the following selected performance. You can also add the resources independent of the preceding performance requirements.

- When you select the **Light (around 1000 requests per sec)** performance, the estimated performance is light (around 1000 requests per sec) across all the deployed Objects worker nodes; total vCPUs are 34 and memory is 104 GiB. For light performance implementation, 6 static IP addresses and 7 DHCP IP addresses are

required (1 IP address for each of 3 worker VMs and 2 load balancer VMs, and 2 master Virtual IP addresses).

- When you select the **Medium (around 10,000 requests per sec)** performance, the estimated performance is medium (around 10,000 requests per sec) across all the deployed Objects worker nodes; total vCPUs are 46 and memory is 140 GiB. For medium performance implementation, 6 static IP addresses and 9 DHCP IP addresses are required (1 IP address for each of 4 worker VMs and 3 load balancer VMs, and 2 master Virtual IP addresses).
- When you select the **Heavy (around 20,000 requests per sec)** performance, the estimated performance is heavy (around 20,000 requests per sec) across all the deployed Objects worker nodes; total vCPUs are 58 and memory is 176 GiB. For heavy performance implementation, 6 static IP addresses and 11 DHCP IP addresses are required (1 IP address for each of 5 worker VMs and 4 load balancer VMs, and 2 master Virtual IP addresses).
- When you select the **Custom** performance, you can manually enter the total vCPUs and the total memory. You can enter in the multiple of 10 vCPUs (minimum 10 vCPUs and 32 GiB memory). For 10 and 20 vCPUs, additional 2 vCPUs and 4 GiB of memory is used for the load balancer. For 30 vCPUs, additional 4 vCPUs and 8 GiB of memory is used for the load balancer. For 40 vCPUs, additional 6 vCPUs and 12 GiB of memory is used for the load balancer. For 50 vCPUs and above, additional 8 vCPUs and 16 GiB of memory is used for the load balancer.

  Note: Objects would not always be able to deliver the selected performance. The actual performance varies depending upon various factors.

b. **Capacity**: Enter the desired capacity.
   Ensure that the available capacity of the physical cluster is sufficient.

6. Click **Next**.

7. On the Deploy page, do the following:

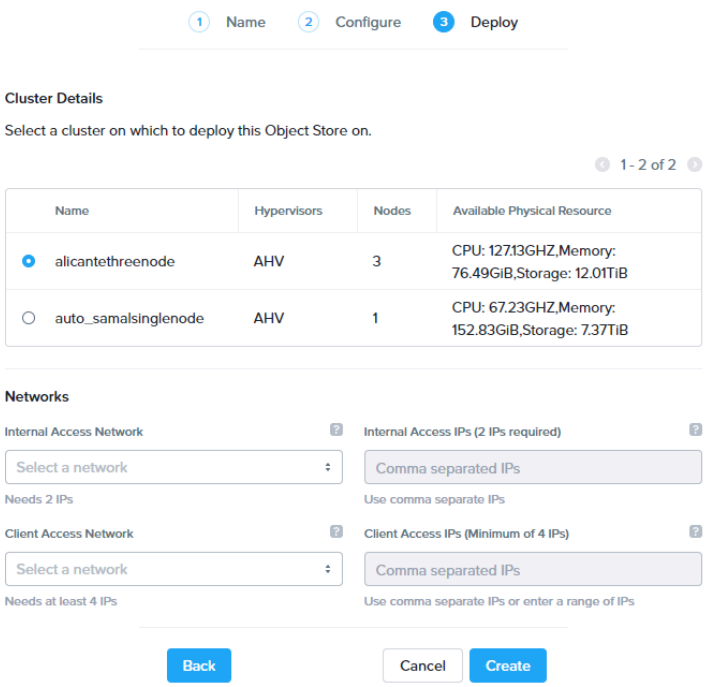

Figure 3: Deploy an Object Store

a. Select a cluster on which you want to deploy the object store.

> Note: You can hover over the help icon to get more information about the respective fields.

b. In **Internal Access Network**, select internal access networks.
   You need at least 12 free IP addresses in the DHCP pool.

c. In **Internal Access IPs**, enter two internal access IP addresses separated by a comma.
   These IP addresses are within the internal access network.

d. In **Client Access Network**, select client access networks.
   This VLAN should have at least four IP addresses in the usable IP address range. This network can be the same as the Internal Access Network.

e. In the **Client Access IPs**, enter a minimum of four client access IP addresses separated by a comma or as an IP address range.
   These IP addresses are within the client access network and used to access the object store.
   For more information about network configurations, refer to Network Configurations on page 14.

8. Click **Deploy**.

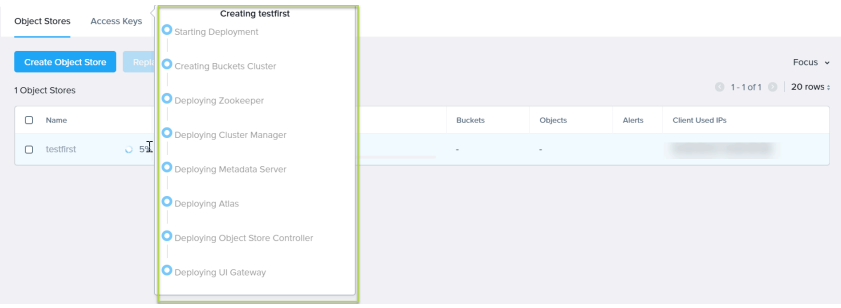

Figure 4: Deployment Steps

This procedure successfully deploys the object store by completing all the deployment steps. You can see the deployment progress in percentage and each step in the grid by hovering over the loading icon.

> Note: You must use an encrypted cluster for encrypting the bucket.

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 36.
- Access the endpoint provided as part of Client Access Network (S3 Endpoints) by using HTTP or HTTPS. Refer to Access Objects Endpoints on page 22.
- Create and configure buckets. Refer to Creating and Configuring a Bucket on page 45.
- 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 24
- Share buckets. Refer to Sharing a Bucket on page 56.

Viewing Object Store Deployments

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

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**.

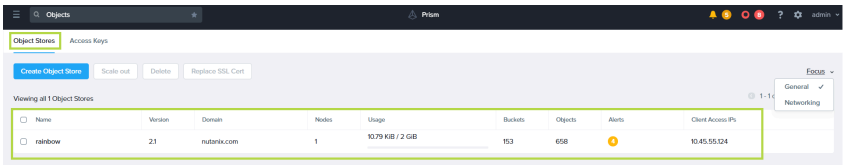

Figure 5: Object Stores List View

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 **Focus** to view the **General** 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**: 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 70.
- **Client Access 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 22.

Networking view:

- **Name**: Displays the name of the object store.
- **Cluster**: Displays the cluster in which object store is deployed.
- **Client Access Network**: Displays the VLAN required for accessing Object Store endpoints externally.
- **Client Configured IPs**: Displays the client configured IP address.
- **Internal Access Network**: Displays the VLAN required internally for deploying Object Store Services on Prism Element.
- **Internal Configured IPs**: Displays the internal configured IP addresses.

For information about Network Configurations, refer to Network Configurations on page 14.

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 for accessing buckets and objects.

- 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, whitelist the Objects FQDN in the DNS server with the wild card whitelists. 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.

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.
- Multicluster 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 can only try replacing the SSL certificate again. You cannot delete that object store deployment.

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 Store table, select the object store which you want to delete, and then click **Delete**.

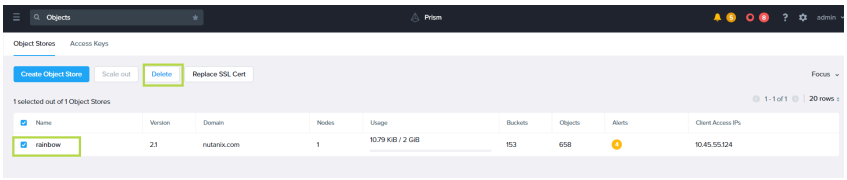

Figure 6: Delete Object Store Window

3. On the dialog box, click **Confirm**.
   A message appears to confirm the object store deletion.

Object Store Storage Expansion

If your existing object store storage is less than 85% full or you want to use different cluster for object store storage, you can expand the storage of your object store cluster by adding new nodes to the same existing cluster or by adding secondary 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 add some other new clusters or existing clusters (such as Cluster2 and Cluster3) if they have unused storage to the object store for saving your data. In this example, Cluster1 is the primary cluster on which the object store was initially deployed and hosts the worker VMs for object store. Cluster2 and Cluster3 are secondary clusters that are later added for capacity expansion.

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 for expanding the object store cluster.
- Ensure that the new clusters are registered to the same Prism Central cluster where 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 should be present in the secondary clusters.
- Firewall and IP table should be running in the cluster.
- Minimum 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.

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

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.

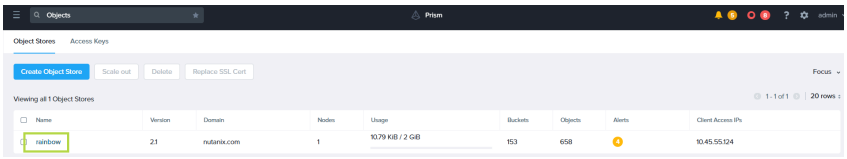

Figure 7: Object Stores List View

The object store opens on a new window.

4. Click **Clusters**.

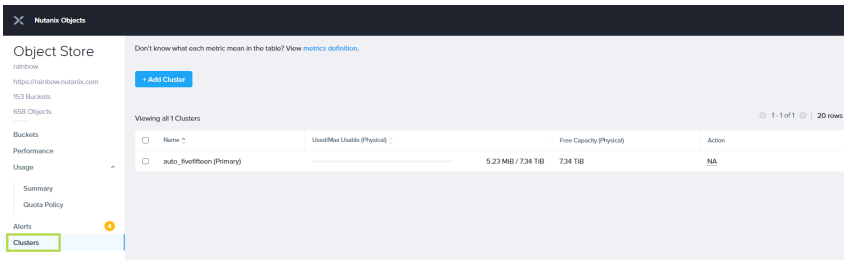

Figure 8: Object Store Clusters List

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 its 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.

   A list of clusters registered to the Prism Central will be displayed. Hypervisor type, total capacity and free 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 certain amount of total capacity on the selected cluster. If you exceed the limit, an alert is generated.

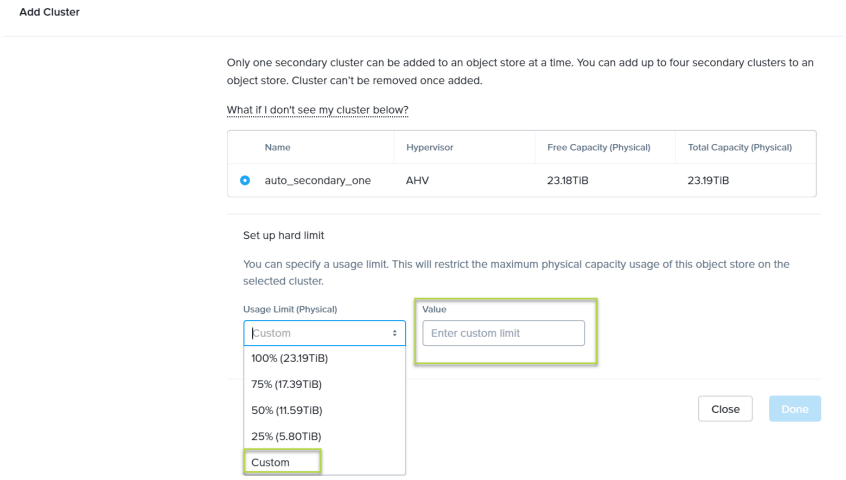

Figure 9: Setting Up Hard Limit

You can also update the limit once the secondary cluster is added by clicking **Update Limit**. You cannot update limit for the primary cluster.

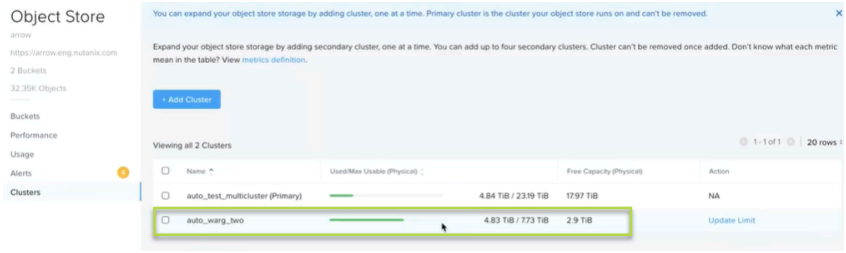

Figure 10: Setting Up Hard Limit

A new cluster is added to the object store. If any secondary cluster addition fails, you can remove that cluster. You can see the storage usage of the cluster in a bar. This bar changes color from green to yellow to red to reflect the usage status of storage. You can also see the usage of these clusters in the **Usage** tab. For more information, refer to Viewing Object Store Usage on page 67.

> 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 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.

Before you begin

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

About this task

- 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 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.

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, and then click **Scale Out**.

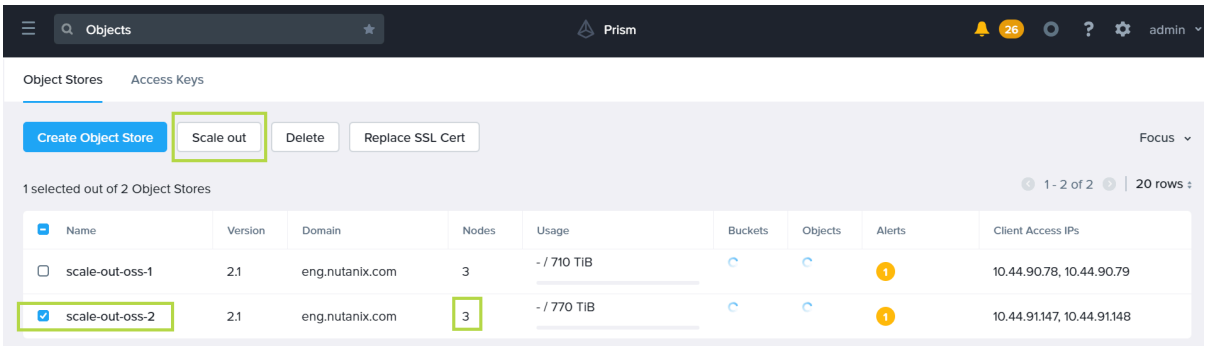

Figure 11: Object Store Selection for Scale Out

The Scale Out window appears.

4. (Optional) Click **Update Storage** check box to add more storage to the current usage, and enter the required storage in **Storage Consumption**.

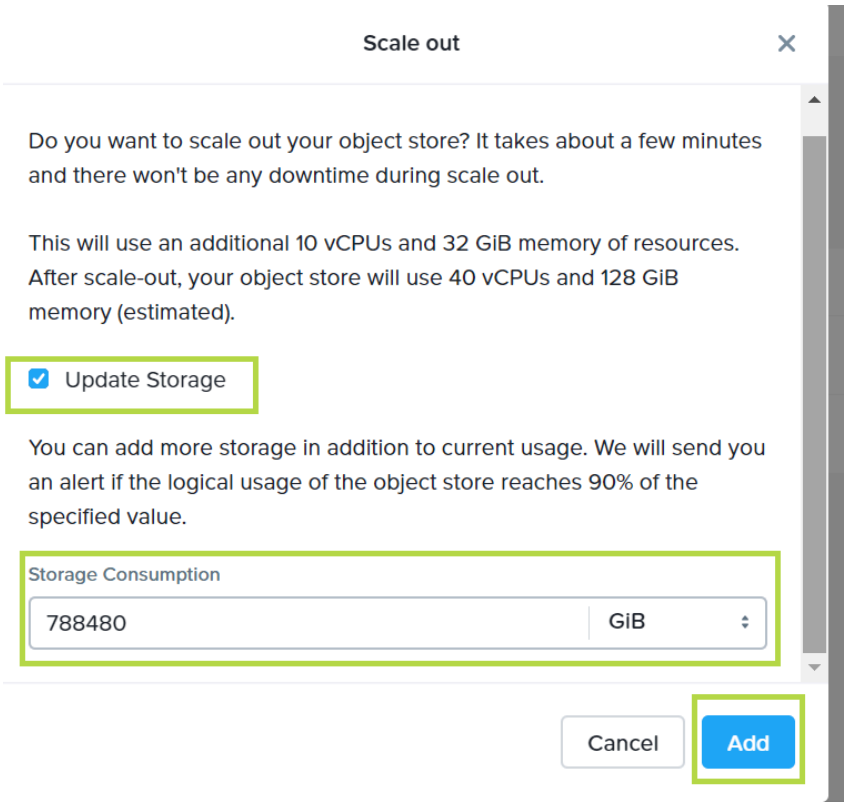

Figure 12: Scale Out Window

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 on page 70.

5. Click **Add**.

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

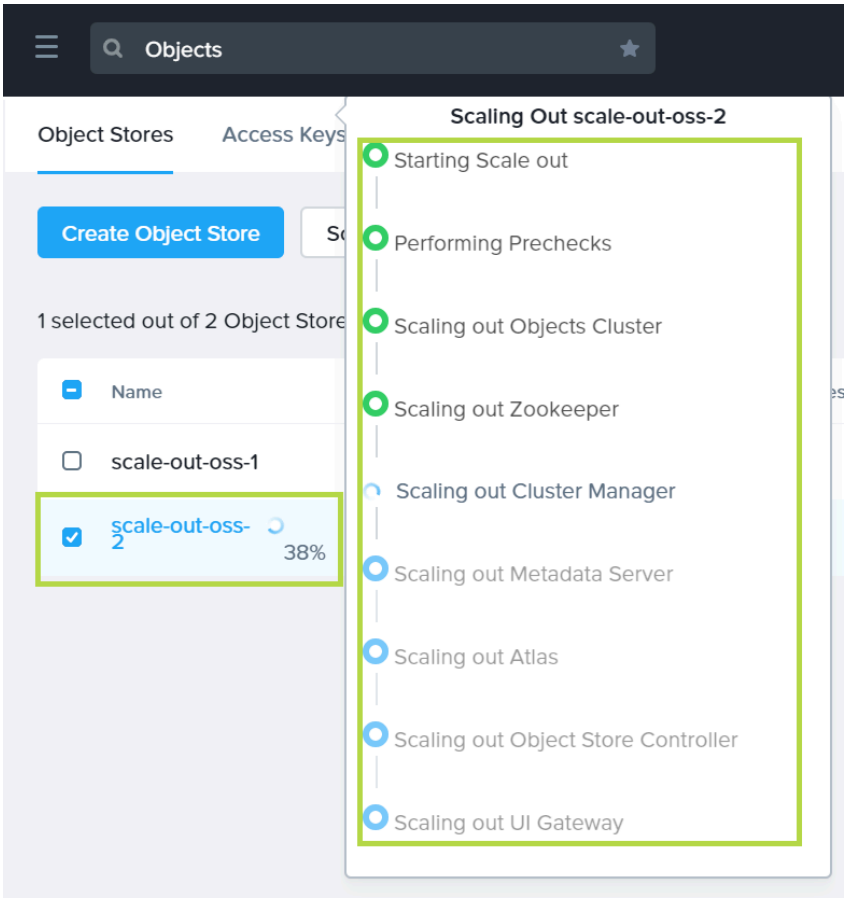

Figure 13: Object Store Scale Out Progress

Once the object store scale out is completed, new node is added to the object store.

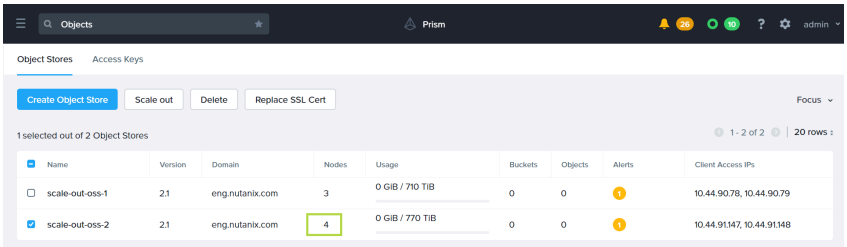

Figure 14: Scaled Out Object Store

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 web server to securely connect to the web server 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 2048 bit and in PEM format and must follow PKCS#1 standard.
- The provided public certificate must be signed by the provided CA.
- The public certificate must have the FQDN of the Object Store Service in either CN or SAN. For example, if the object store name is `objects-2019` and domain is `companyname.com`, then the FQDN should be objects-2019.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 **Replace SSL Cert**.

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

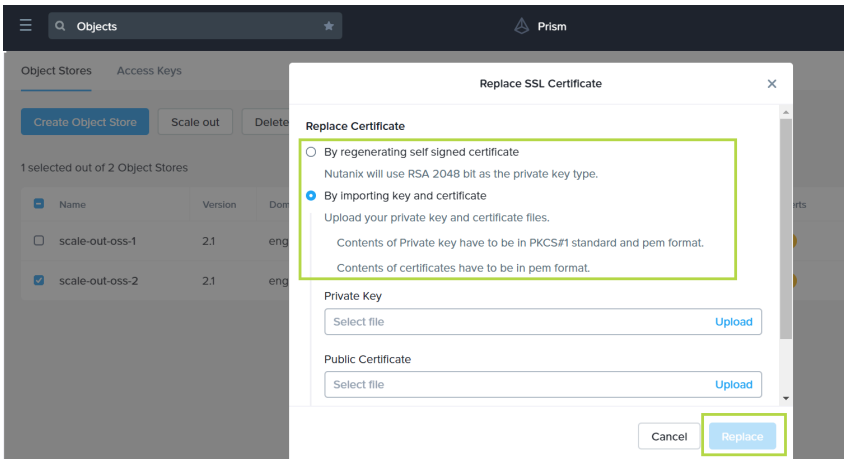

Figure 15: Replace SSL Certificate Window

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

5. Click **Replace**.
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 on page 8.

About this task

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** > **Tools & Firmware**. 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.
   KB 5901 describes several examples of setting up a web server.

   Note:
   - Set the following mime types in the LCM web server when you are using Windows Web Server:
     - .json
     - .xz
     - .qcow2
     - .sign
     - .gz
   - In step 13 of the KB 5901, use the latest LCM Dark Site Bundle.
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**. Choose the required version and download `objects-x.x.tar.gz` tar file.
   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. Enable Objects to start the msp-controller on Prism Central. Refer to Enabling Objects on page 6.

   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.
8. (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 86.

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

   ```
   admin@pcvm$ mspctl airgap --enable --lcm-server=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 airgap --enable --lcm-server=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 airgap --status
   ```

10. Deploy an object store through Prism Central. Refer to Creating or Deploying an Object Store (Prism Central) on page 17

5

TYPES OF OBJECTS USERS

Objects has two types of users.

- Prism Central users: 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 bucket of any API user with any other users.
- API users: API users cannot access Objects by using the Prism Central web console. They access bucket 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.

6

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:

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**.

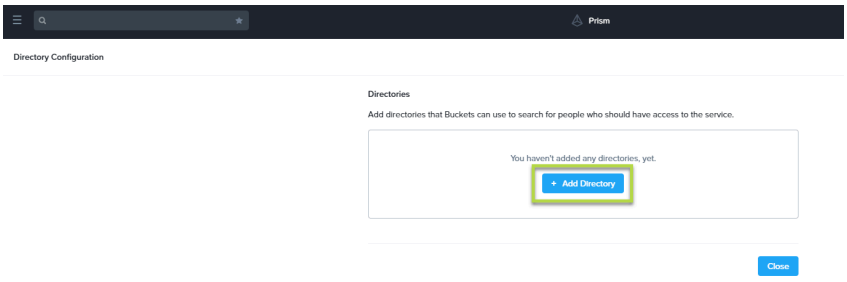

Figure 16: Directory Configuration

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:

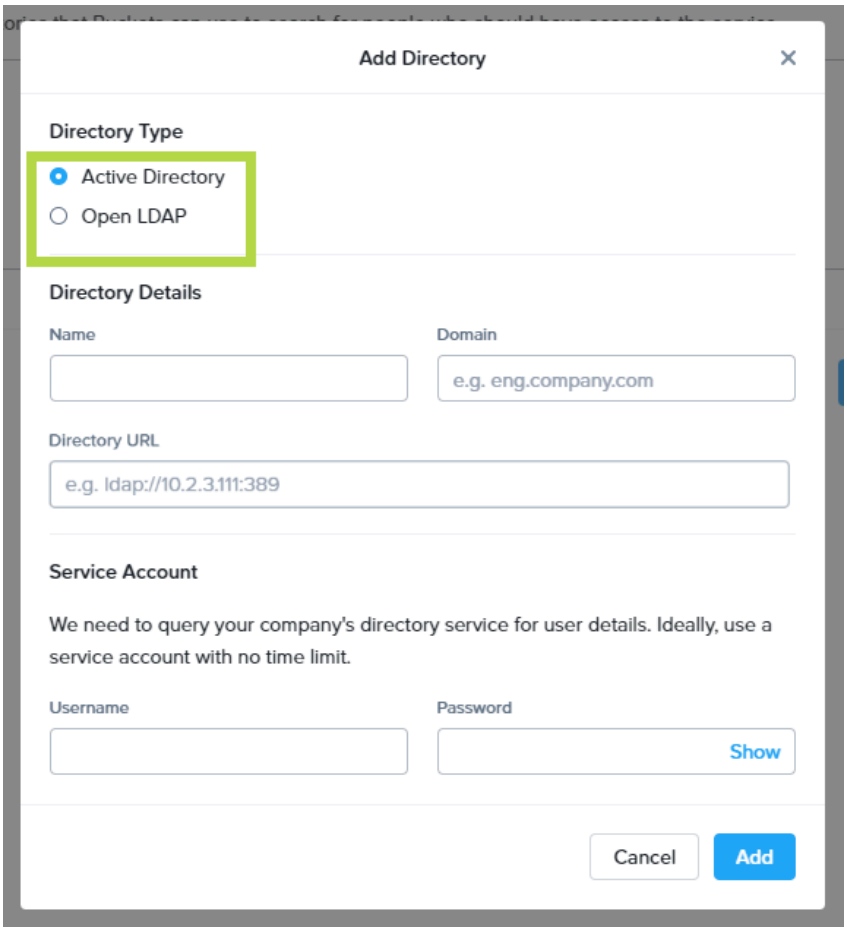

Figure 17: Add Directory Window

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 39.

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.

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**.

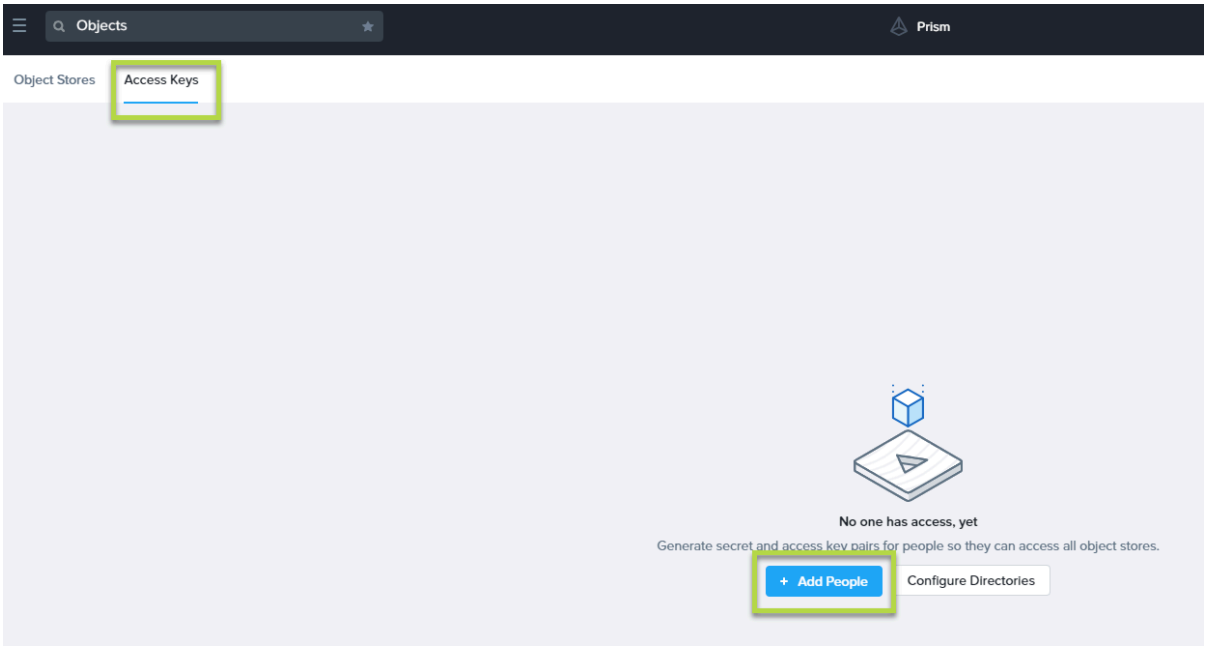

Figure 18: Adding People for Access Key Generation

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

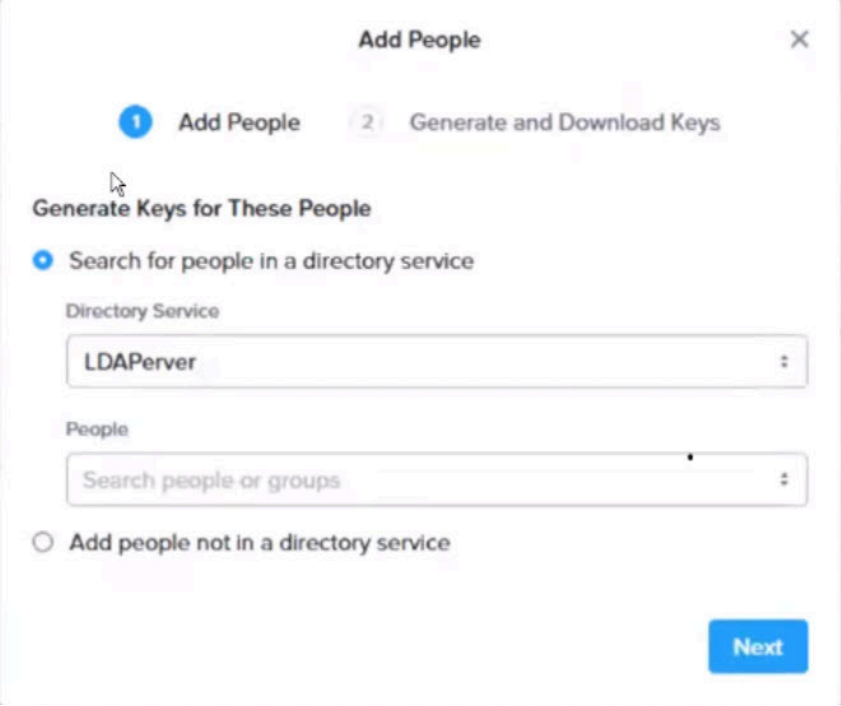

Figure 19: Add People Window: Search in Directory

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 36.

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.

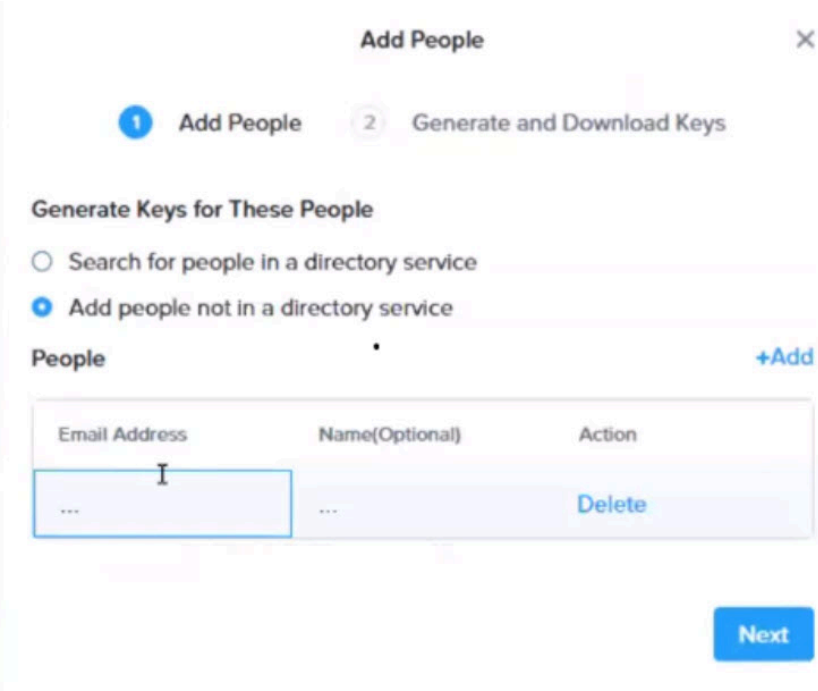

Figure 20: Add People Window: Enter Email Address

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.

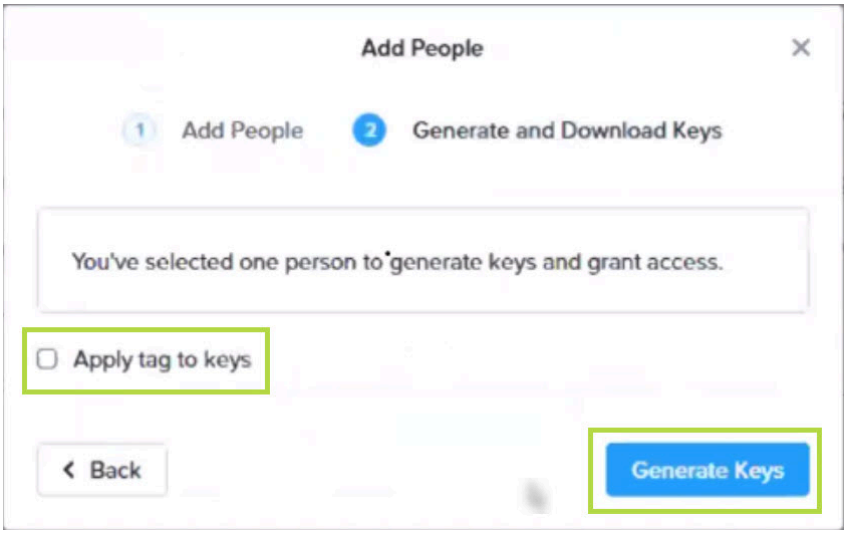

Figure 21: Generate and Download Keys Page

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.

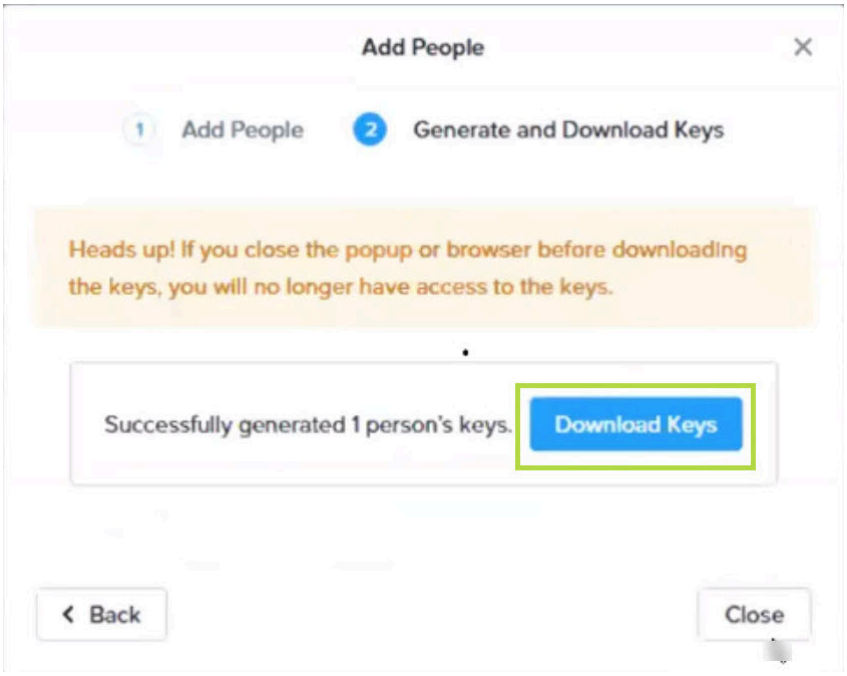

Figure 22: Generate and Download Keys Page

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.

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.

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.

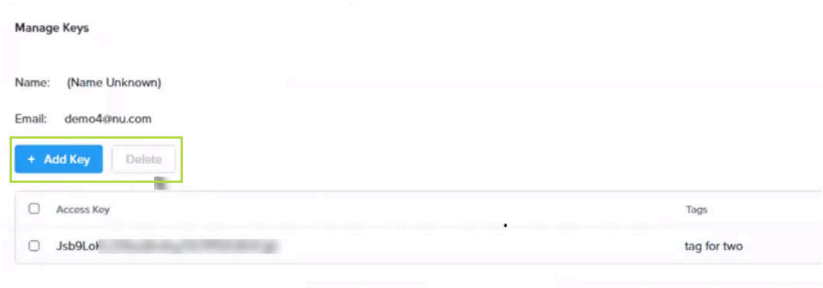

Figure 23: Manage Keys Page

   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.

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.

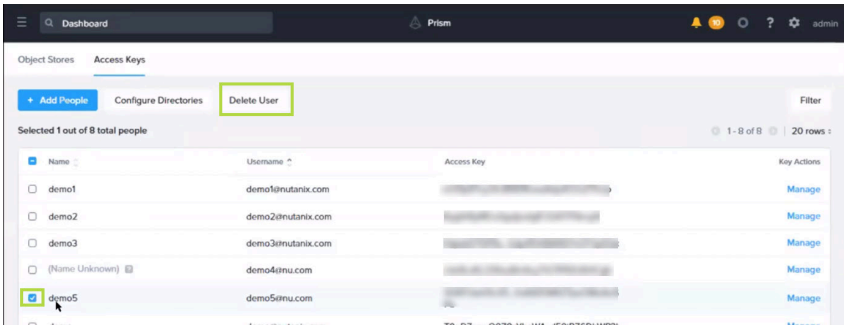

Figure 24: Access Keys Page - Deleting an API User

The user is deleted.

7

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 also configure a bucket while creating it except applying the WORM policy.

You can enable versioning, apply lifecycle policies and WORM policies to a bucket. You can also filter, update, share, and delete a bucket, and view the existing buckets list.

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 a 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.

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**.

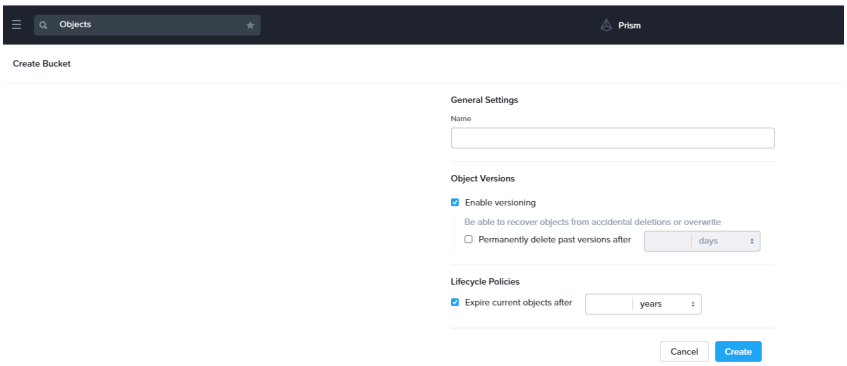

Figure 25: Create Bucket Window

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 47.
   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.
7. Click **Save**.
   The bucket is created successfully.

What to do next

After creating a bucket, you can create objects through S3 APIs.

For more information, refer to Supported S3 APIs.

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. You can also configure a bucket for static website hosting and CORS.

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 a Bucket on page 45.

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:

- 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.

Lifecycle Policies

Lifecycle policy enables you to create a set of rules that define actions that Nutanix Objects applies to a group of objects. With these policies, you can manage the objects cost effectively throughout their lifecycle.

Note: By default, the lifecycle policies are disabled for a new bucket.

You can apply these policies while creating or editing a bucket. You cannot change the policy for different objects of the same bucket.

With lifecycle policies, you can configure the following:

- Allow object to automatically delete itself after a specified number of days or months or years from the date of object creation.
- Delete past versions permanently after a specified number of days or months or years from the date of object creation.
- Set a rule for a bucket so that all the objects in the bucket (or folder) adopt that rule.

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 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 any lifecycle policy to a WORM bucket with versioning enabled, the lifecycle policy is not applicable until the WORM retention period has elapsed.

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.

When you increase the retention period of a bucket, the new retention period does not apply to the older objects. For example, if you have first applied WORM policy on a bucket on October 2018, and then you increase the retention period on January 2019. Then the new retention period is not applicable on the objects created before January 2019.

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 past.

Note: You cannot enable WORM policy while creating a bucket.

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 new version of 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 the WORM policy to any bucket while creating a bucket. 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 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.

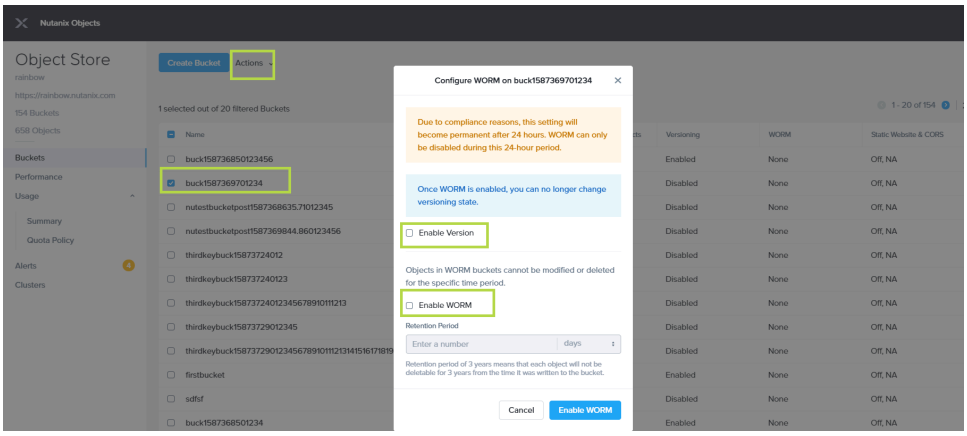

Figure 26: Configure WORM Window

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** checkbox 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.

About this task

> Note: 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.

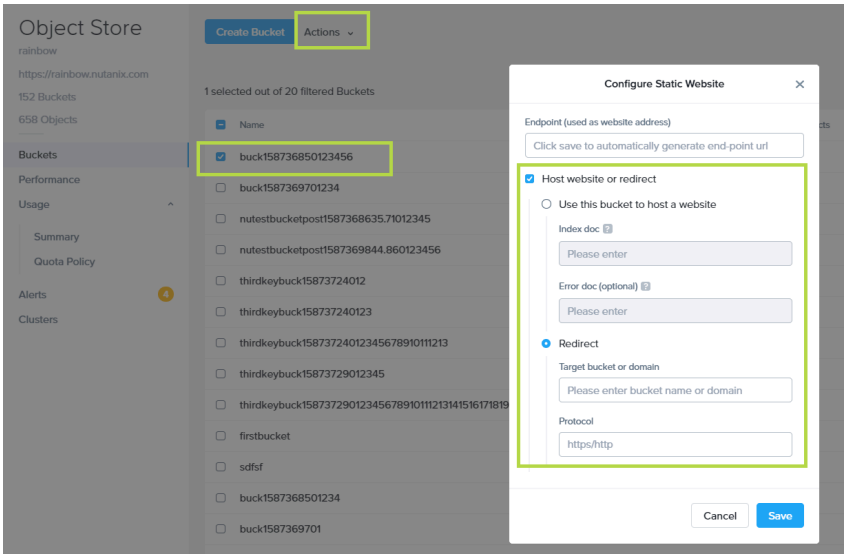

Figure 27: Configure Static Website Window

5. By default, 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, `test.png`) in a bucket (for example, test-bucket) in domain1 that contains some security-related information, and you are not allowed to make access to `test.png` from a website (for example, `www.example.com`) on domain2. Then you can configure a CORS policy for that test-bucket to allow `www.example.com` to access the resources of test-bucket.

> Note: 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 for 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.test.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.test.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.
- 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.

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**.

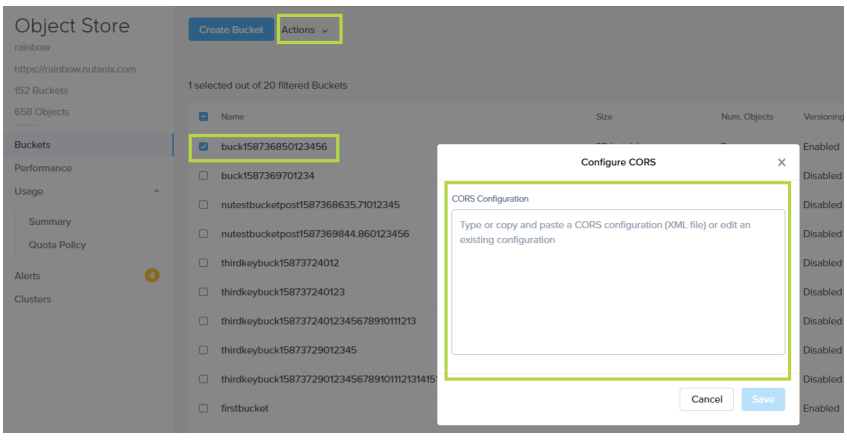

Figure 28: Configure CORS Window

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.

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.

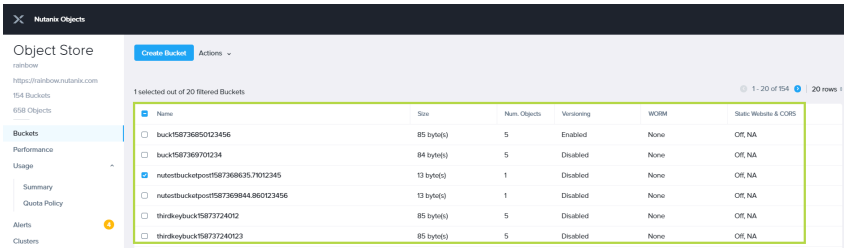

Figure 29: Buckets List View

A new window appears with the list of buckets in a tabular view. If you are in 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.
- **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.
- **Static Website & CORS**: Displays if static website or CORS is configured for a bucket.

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.

Updating a Bucket

You can update the bucket settings after creating the buckets and adding the objects to the buckets.

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.

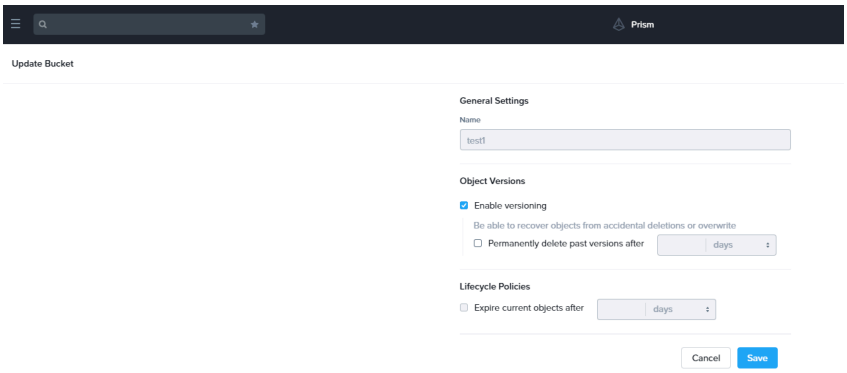

Figure 30: Update Bucket Window

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.

About this task

> Note: You can share only one bucket at a time.

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.

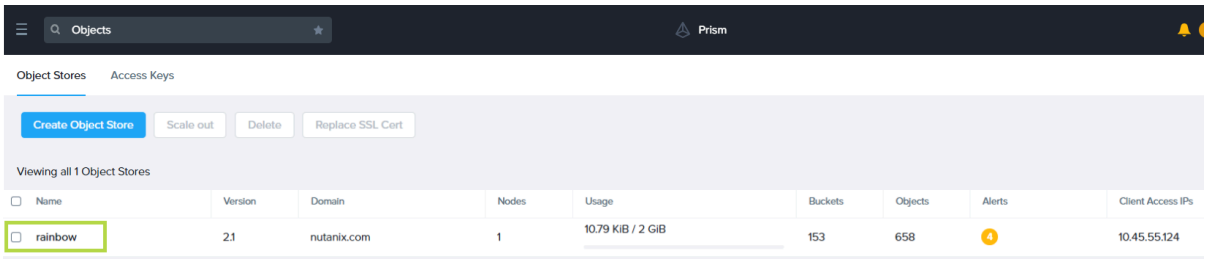

Figure 31: Object Store List

   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.

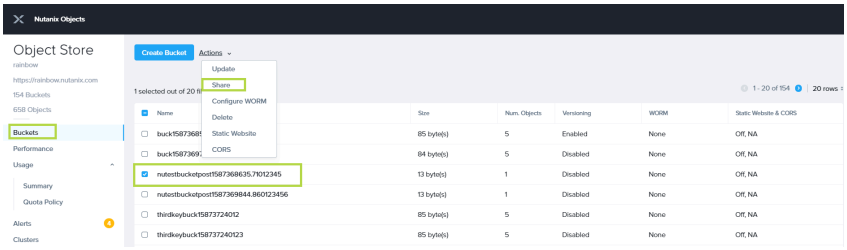

Figure 32: Sharing the Bucket

> 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.

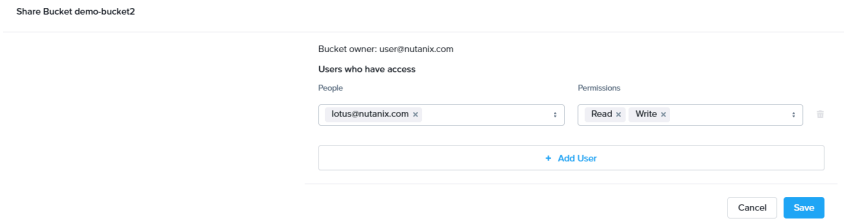

Figure 33: Share Bucket Window

To generate access keys, refer to Generating Access Key for API Users on page 39.
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 586b47ca6cf7b311f6a5200a30fc8a7065cd933f.dita.

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 3: 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.

Access Role	Description
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 4: 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 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 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 5: 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
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

S3 API	Read	Write	Read and Write
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
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.

About this task

To view the buckets usage.

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.

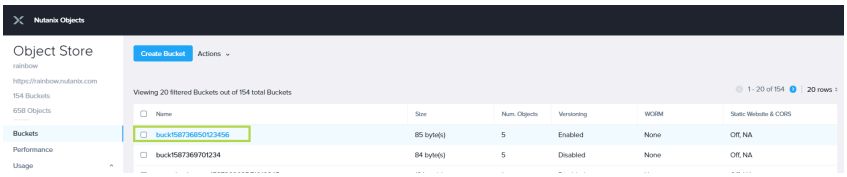

Figure 34: Buckets List

4. Click **User Access**.

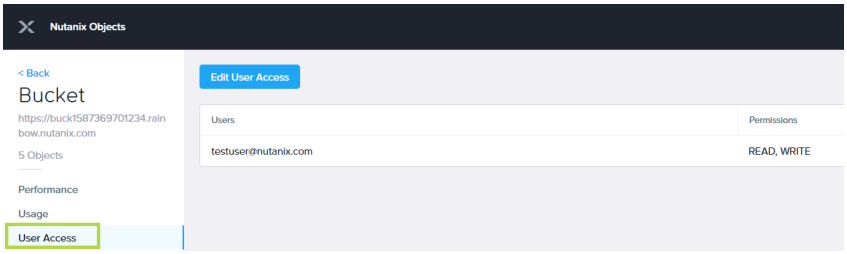

Figure 35: List of Users

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 56 topic.

Deleting a Bucket

If you no longer require a bucket, you can delete the bucket.

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**.

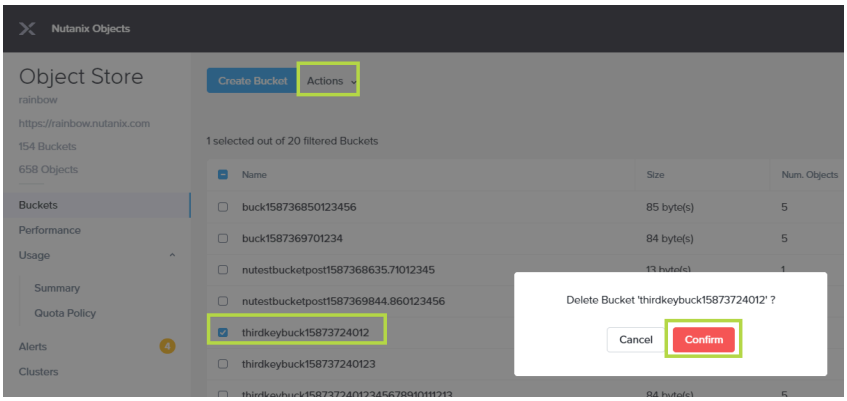

Figure 36: Delete Bucket Window

5. On the dialog box, click **Confirm**.
   A message appears to confirm the buckets deletion.

8

OBJECTS REPLICATOR

Objects Replicator is a script-based replication solution to sync objects across two object stores.

> Note: The Objects Replicator feature is a Limited Availability (LA) release.

To access the Objects Replicator tool, contact the *Accounts* team or reach out to *objects@nutanix.com*

9

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.

> 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**.

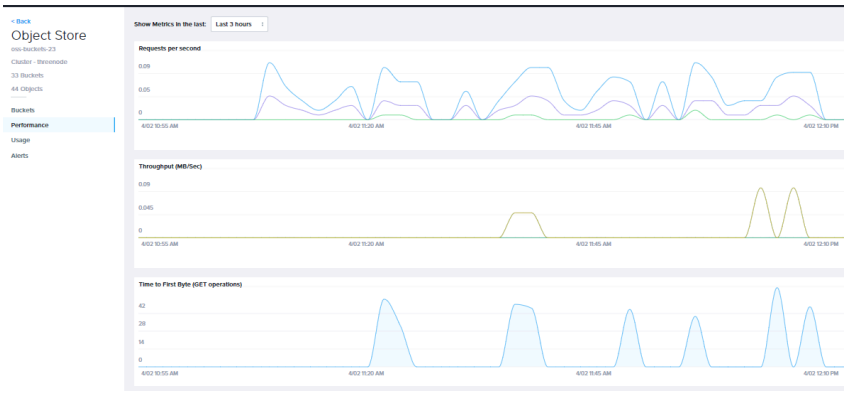

Figure 37: Object Store Performance Graph

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 **Get** graph displays the total output in each second.
- The **Put** graph displays the total input 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 and Puts 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**.

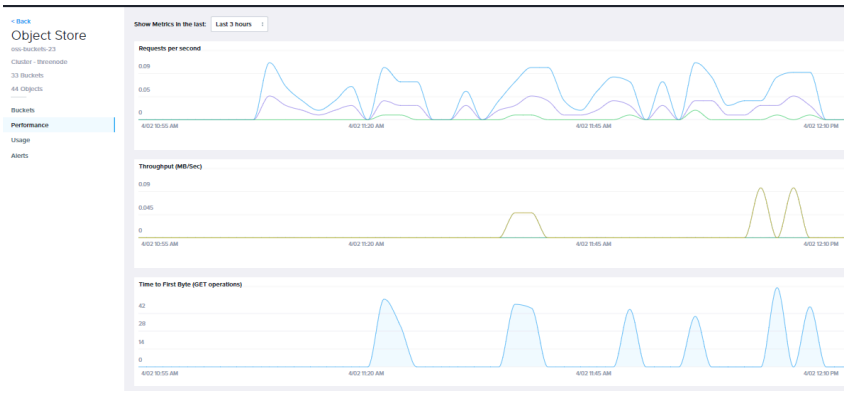

Figure 38: Bucket Performance Graph

The graph displays the performance of the bucket.

a. **Requests Per Second**: Displays the following graphs:

   - The **Total** graph displays the total I/O requests in each second in the object store or bucket.
   - The **Get** graph displays the total output in each second.
   - The **Put** graph displays the total input 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 and Puts 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 storage usage in a graphical form. 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**.

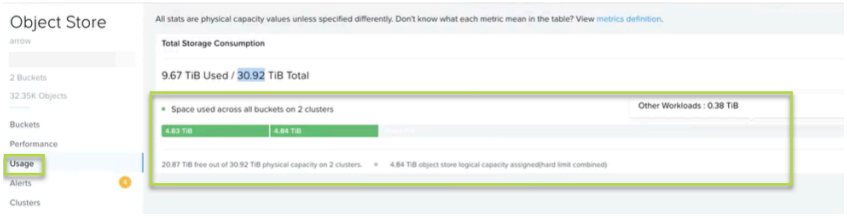

Figure 39: Object Store Usage

The storage usage section displays the following information:

- Total used and total free storage across all clusters.
- Space used across all buckets on the existing clusters and other workloads in a bar.
- Total and free physical capacity on all the clusters.
- Total and free logical capacity of the object store.

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 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.

About this task

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**.

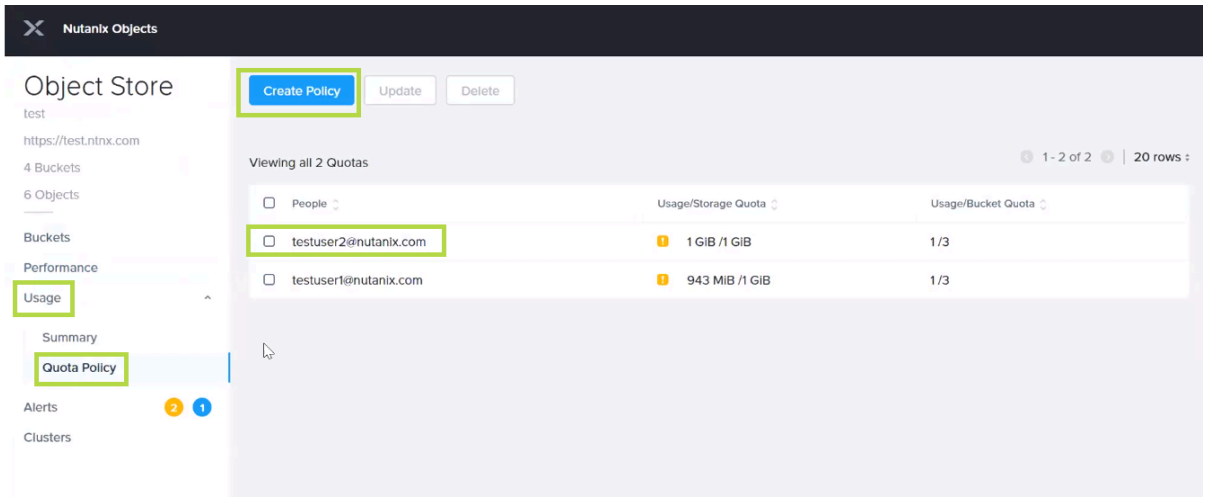

Figure 40: Create Quota Policy

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 **Storage Limit** check box and enter the soft limit for storage usage, and then click **Save**. A list of users with the configured storage quota and the current usage are displayed.

   If you are about to reach the quota or have exceeded the quota, 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 70 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.

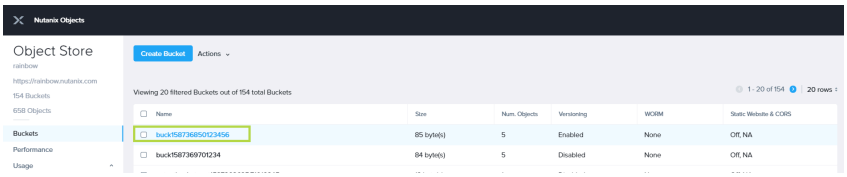

Figure 41: Buckets List

Note: Only finalized objects of a multipart upload are shown in the buckets table.

4. Click **Usage**.

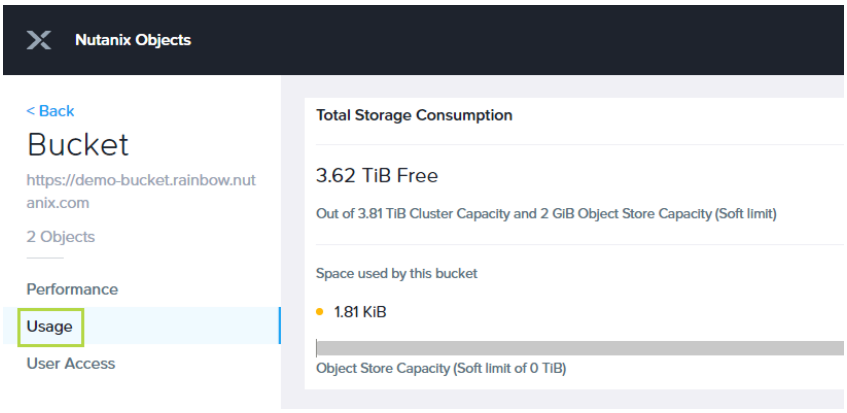

Figure 42: Bucket Usage

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.

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 43: Alerts List

The alerts table displays the following information:

Table 6: 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 7: 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 8: 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 9: 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.

10

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).

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

Supported APIs

The following table lists the supported S3 API methods.

Note:

- User-defined meta-data header names are case insensitive.
- User-defined object meta-data is supported for `PUT Object` and `PUT Object - Copy` APIs and limited to 2 KB in size.

Table 10: Supported S3 APIs and Parameters

Supported S3 APIs	Request Parameters/Request Headers/ Request Elements/Request Body
PUT Bucket	Common Headers on page 78 + CreateBucketConfiguration, LocationConstraint (By default, LocationConstraint “us-east-1” is only supported.)
PUT Bucket Lifecycle	Common Headers on page 78 + 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 78 + JSON string containing the policy contents containing the policy statements
PUT Bucket versioning	Status, VersioningConfiguration
Complete Multipart Upload	Common Headers on page 78 + CompleteMultipartUpload, Part, PartNumber, ETag
PUT Object - Copy	Common Headers on page 78 + 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-
DELETE Objects	Single query string parameter "delete" + Common Headers + Content-MD5, Content-Length, Delete, Quiet, Object, Key, VersionId
GET Object	Common Headers on page 78 + Range, If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match
HEAD Object	Common Headers on page 78 + Range, If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match
List Multipart Uploads	Common Headers on page 78 + delimiter,max-uploads, key-marker, prefix, upload-id-marker
List Object Versions	Common Headers on page 78 + delimiter, key-marker, max-keys, prefix, version-id-marker
GET Bucket (List Objects) Version 1	Common Headers on page 78 + delimiter, marker, max-keys, prefix

Supported S3 APIs	Request Parameters/Request Headers/ Request Elements/Request Body
GET Bucket (List Objects) Version 2	Common Headers on page 78 + delimiter, max-keys, prefix, list-type, continuation-token, fetch-owner, start-after
List Parts	Common Headers on page 78 + uploadId, max-parts, part-number-marker
PUT Object	Common Headers on page 78 + Content-Length, Content-MD5, Expect, x-amz-tagging, x-amz-meta-
INITIATE Multipart Upload	Common Headers on page 78 + x-amz-meta-, x-amz-tagging
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 78 + ObjectLockEnabled Rule
DELETE Bucket	Common Headers on page 78 + Bucket
Delete Multiple Objects	Common Headers on page 78 + Bucket, Delete, Quiet, Objects, Key, VersionId
POST Object	Common Headers on page 78 + post url + all the required parameters for post url
GET Object ACL	Common Headers on page 78 + Bucket, Key, VersionId

For more information about objects tagging API, refer to Objects Tagging APIs Overview on page 79.

Pre-signed URLs can be generated for Objects. For more information about generating the pre-signed URLs, refer to *Authenticating Requests: Using Query Parameters (AWS Signature Version 4)* section in the *Amazon Simple Storage Service API Reference Guide*.

The following table lists the supported APIs that use only Common Headers on page 78:

Supported APIs
GET Bucket lifecycle configuration
PUT Bucket lifecycle configuration
GET Bucket Object Lock Configuration
GET Bucket Location
GET Bucket Policy
GET Bucket versioning
HEAD Bucket
LIST Bucket

Supported APIs
DELETE Bucket Lifecycle
DELETE Bucket Policy
ABORT Multipart Upload
DELETE Object
GET Bucket acl

Common Headers

You can use the following headers while making requests:

Table 11: 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 12: Unsupported S3 APIs

Unsupported S3 APIs
GET Bucket accelerate configuration
GET Bucket acl
GET Bucket analytics configuration
GET Bucket cors
GET Bucket encryption
GET Bucket inventory configuration
GET Bucket logging
GET Bucket metrics configuration

Unsupported S3 APIs
GET Bucket notification configuration
GET Bucket replication
GET Bucket requestPayment
GET Bucket website
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 cors
PUT Bucket encryption
PUT Bucket inventory configuration
PUT Bucket logging
PUT Bucket metrics configuration
PUT Bucket notification configuration
PUT Bucket replication
PUT Bucket website
DELETE Bucket analytics configuration
DELETE Bucket cors
DELETE Bucket encryption
DELETE Bucket inventory configuration
DELETE Bucket metrics configuration
DELETE Bucket replication
DELETE Bucket website
DELETE Bucket encryption
Copy
GET 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 value pair that you can define.

You can add, retrieve, and remove tags for your objects by using the tagging APIs. Retrieving objects metadata also returns the number of tags associated with the object (if any). Tagging is also supported with a few more Object APIs.

Objects Tag Limitations

The following are the limitations of Objects tag.

- You can add up to 10 tags for each object.
- 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.

Supported Tagging API Operations

Currently, Objects only supports APIs for tagging objects. Tag-based bucket lifecycle policies 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.

  There are two scenarios 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.

Also, the following Object APIs support tagging.

- GET Object (returns tag count, if any)
- PUT Object
- POST Object
- PUT Object-Copy
- Initiate Multipart Upload

11

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, the 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 13: REST Error Responses

Name
Code
Error
Message
RequestId
Resource

For more information on Error Responses, refer to *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 14: 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

Error Code	HTTP Status Code
IncompleteBody	400 Bad Request
InlineDataTooLarge	400 Bad Request
InternalError	500 Internal Server Error
InvalidAccessKeyId	403 Forbidden
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

Error Code	HTTP Status Code
Redirect	307 Moved Temporarily
RequestIsNotMultiPartContent	400 Bad Request
RequestTimeout	400 Bad Request
RequestTimeTooSkewed	403 Forbidden
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 *List of Error Codes* section in the *Amazon Simple Storage Service API Reference Guide*.

12

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, Commtrade HYCU, Veeam and Veritas.

For more information about Commvault Integration, refer to *Commvault with Nutanix* guide on the Nutanix Support Portal.

13

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 86.

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 87.

> 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 6.
3. In LCM, perform inventory and update Objects to an available GA version.

Microservices Platform (MSP)

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. 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

You can find your MSP version by running the following command on Prism Central.

```
admin@pcvm$ mspctl controller_version
```

Upgrading MSP Controller

You can upgrade the MSP Controller from the LCM page in the Prism Central.

Note: For Objects deployment, it is recommended to upgrade to the latest MSP controller version.

To upgrade the MSP controller, do the following.

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 6.
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.

14

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/.

Prism 5.15

Prism Central Guide

April 7, 2020

Contents

Welcome vii

Help Organization viii

1. Prism Central Overview 9
- Main Menu 9
- Entities Menu 12
- Settings Menu 14
- Searching for Information 18
- Understanding Displayed Statistics 25

2. Prism Central Administration 27
- Installing or Upgrading 27
- Licensing 27
- Prism Central Management 27
  - Logging Into Prism Central 28
  - Logging Out of Prism Central 31
  - Managing Prism Central 31
  - Checking Prism Central Version 34
  - Register (Unregister) Cluster with Prism Central 34
  - Expanding (Scale Out) Prism Central 42
  - Shutdown (Start Up) Prism Central VM 46

3. Main Dashboard 47
- Creating a New Dashboard 58
- Modifying a Dashboard 59
- Adding Dashboard Widgets 61

4. Entity Exploring 66
- Virtual Infrastructure Entities 71
  - VMs Summary View 72
  - Storage Containers Summary View 87
  - Catalog Items View 97
  - Images Summary View 98
  - Categories Summary View 102
  - Recoverable Entities View 104
  - Subnets View 105
- Policies Entities 106
  - Security Policies Summary View 106
  - Protection Policies View 109
  - Recovery Plans View 110
  - NGT Policies View 111
  - Image Placement Policies Summary View 112
- Hardware Entities 114

Clusters Summary View....114
Hosts Summary View....131
Disks Summary View....141
GPUs Summary View....150
Activity Entities....153
Audits Summary View....153
Tasks View....157
Operations Entities....159
Reports View....159
Administration Entities....161
LCM View....161
Projects Summary View....161
Roles Summary View....167
Users Summary View....170
Availability Zones View....174
Services Entities....175

5. Resource Planning....176
Capacity Runway View....176
Scenarios View....177
Creating a Scenario....179
Modifying a Scenario....195
Updating Capacity Configurations....196
Behavioral Learning Tools....197

6. Performance Monitoring....202
Analysis Dashboard....202
Creating an Entity Chart....204
Creating a Metric Chart....206
Chart Metrics....207

7. Services Enablement....212
Enabling Calm....212
Enabling Karbon....213
Enabling Objects....213

8. Alert and Event Monitoring....214
Alerts Summary View....214
Alert Details....217
Events Summary View....220
Event Details....221
Configuring Alert Emails....223
Configuring Alert Policies....226
Modifying System Alert Policies....227
Adding Custom Alert Policies....229
Stopping Alerts During a Maintenance Window....244
Prism Central Logs....247
Alerts/Health checks....250
Cluster....250
Controller VM....259
DR....262
Guest VM....266

Node....267
Network....268
Other....269
Storage....275
System Indicator....276

9. System Management....277
Configuring an HTTP Proxy....277
Configuring Name Servers....280
Configuring NTP Servers....281
Configuring an SMTP Server....282
Configuring SNMP....283
Configuring Syslog Monitoring....290
Modifying UI Settings....290
Configuring a Banner Page....291
Internationalization (i18n) (Prism Central)....292
Localization (L10n) (Prism Central)....293
Changing the Language Settings (Prism Central)....293

10. Security and User Management....295

11. Virtual Infrastructure (Cluster) Administration....296
VM Management....296
Creating a VM (AHV)....297
Creating a VM (Self Service)....306
Managing a VM (AHV and Self Service)....310
Creating a VM (ESXi)....318
Managing a VM (ESXi)....321
NGT Management in Prism Central....325
Storage Quality of Service (QoS)....331
Image Management....334
Requirements....335
Limitations....335
Adding an Image....336
Modifying an Image....346
Importing Images to Prism Central....347
Catalog Management....347
Adding a Catalog Item....347
Deleting a Catalog Item....350
Category Management....351
Creating a Category....351
Modifying a Category....352
Assigning a Category....353
Affinity Policies for AHV....354
Configuring VM-VM Anti-Affinity Policy....355
Network Connections....356
Configuring Network Connections....357
Modifying Network Connections....361
Performing Other Administrative Tasks....362
vCenter Server Integration....363
Registering vCenter Server (Prism Central)....364
Managing vCenter Server Registration Changes (Prism Central)....367

12. Prism Self Service Administration....368
Prism Self Service Overview....368
Configuring Prism Self Service....369
Migrating Prism Self Service....372
Project Management....373
Creating a Project....373
Modifying a Project....376

13. Policies Management....377
NGT Policies....377
Creating a New Restart Policy....377
Modifying a Restart Policy....379
Image Placement Policies....379
Configuration Workflow....380
Associating a Registered Nutanix Cluster with Categories....380
Associating an Image with Categories....380
Configuring an Image Placement Policy....381
Updating an Image Placement Policy....382
How Prism Central Handles Conflicting Policies....382
Sample Scenarios and Configurations....383

14. Reports Management....385
Creating a New Report....386
Managing a Report....387
Generating a Report....387
Editing a Report....388
Cloning a Report....389
Deleting a Report....389
Viewing Report Instances....390
Example Report - Environment Summary....391
Views in the Reports....393
Adding Views to a Report....393
Scheduling a Report....409
Downloading Reports....410
Configuring Report Settings....411
Role Based Access Control....413

15. Codeless Task Automation (X-Play)....415
Overview....415
Action Gallery....415
Plays....416
Playbook Actions....416
Prerequisites for VM Actions in Playbooks....421
Creating Playbooks....423
Creating Playbooks Using Alerts as Triggers....423
Creating Playbooks Using Manual Triggers....425
Running a Playbook (Manual Trigger)....426
Configuring Manual Parameters....426
Parameters for Triggers and Actions....428
Integrating Slack....429
ServiceNow Integration with Prism Central (X-Play Support)....430

ServiceNow Integration Overview....430
Enabling Event Management Plug-in on ServiceNow (Production Instance).... 430
Enabling Event Management Plug-in on ServiceNow (Developer Instance)....430
Configuring Playbooks for ServiceNow Alert Management....431

16. Customer Support Services.... 433
Pulse Health Monitoring....433
Configuring Pulse.... 435
Prism Central Proxy for Pulse Data.... 438
Creating a Support Case....439
Viewing Case Status....444
Accessing the Nutanix Support Portal.... 445
Accessing the REST API Explorer.... 448

17. Help Resources....450
Accessing Online Help.... 450
Accessing the Nutanix Next Community....452
Glossary.... 453

Copyright....459
License.... 459
Conventions....459
Version....459

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 viii.
- For help with viewing this documentation, see Accessing Online Help on page 450.
- To access other Nutanix documents, see Nutanix support portal.

HELP ORGANIZATION

This documentation is organized as follows:

- Prism Central Overview on page 9 provides an overview of Prism Central including menu options and navigation tools.
- Prism Central Administration on page 27 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 47 describes the main (home) dashboard and how to customize it.
- Entity Exploring on page 66 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 176 describes how to review and analyze current and potential resource needs in a Nutanix cluster.
- Performance Monitoring on page 202 describes how to monitor and analyze performance in a cluster.
- Services Enablement on page 212 describes how to enable certain applications (Calm, Karbon, and Objects) through Prism Central.
- Alert and Event Monitoring on page 214 describes how to monitor alerts and events across the registered clusters and how to configure alert policies and notification.
- System Management on page 277 describes how to configure various system settings (NTP, SMTP, and so on) for Prism Central.
- Security and User Management on page 295 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 296 describes how to administer individual clusters through Prism Central such as how to create and manage VMs.
- Prism Self Service Administration on page 368 describes how to implement self-service features through Prism Central.
- Policies Management on page 377 describes how to create and apply policies to manage and control access and operations in a cluster.
- Reports Management on page 385 describes how to create and run reports about infrastructure resources.
- Codeless Task Automation (X-Play) on page 415 describes how to automate routine administrative tasks through Prism Central by using the X-play feature.
- Customer Support Services on page 433 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 450 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 47).
- 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 71).
- Dashboards to manage security policies, protection policies, recovery plans, Nutanix Guest Tools (NGT) policies, and image placement (see Policies Entities on page 106).
- 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 114).
- Activity monitors for alerts, events, audits, and tasks (see Activity Entities on page 153).
- Tools to analyze system activity, plan for resource needs, create usage reports, and automate routine administrative tasks (see Operations Entities on page 159).
- Dashboards to manage inventory (LCM), projects, roles, users, and availability zones (see Administration Entities on page 161).
- Service connection pages to enable Calm, Karbon, and Buckets (see Services Entities on page 175).
- Settings menu from which you can configure Prism Central functions (see Settings Menu on page 14).

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 12). The Main dashboard is the first screen that appears after logging into Prism Central (see Main Dashboard on page 47).

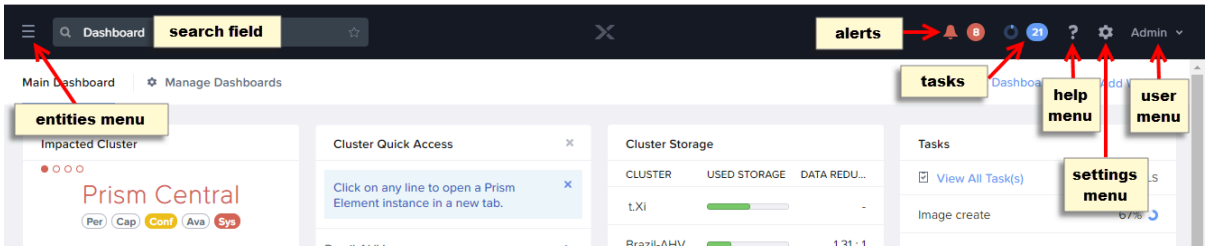

Figure 1: Prism Central Main Menu

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 18).
- 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 214).
- 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 157).

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 18).
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help on page 450).
Online Documentation	Opens the online help at the introduction page (see Accessing Online Help on page 450).
Support Portal	Opens a new browser tab (or window) at the Nutanix Support portal logon page (see Accessing the Nutanix Support Portal on page 445).

Name	Description
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 452). 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 439).

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 14).

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.
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 448).
About Nutanix	Opens the **About Nutanix** window that displays Nutanix operating system (AOS) and other version information (see Checking Prism Central Version on page 34).
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.

Name	Description
Sign Out	Logs you out of Prism Central (see Logging Out of Prism Central on page 31).
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 9) 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**, **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.
- When the entities menu is displayed, the collapse menu button changes to an X button. Click the X button to hide the menu.

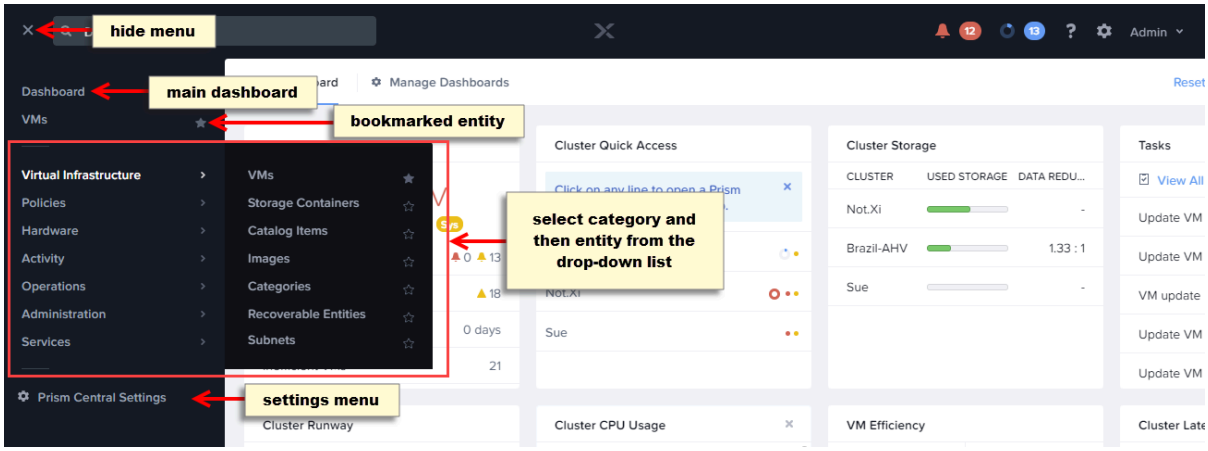

Figure 2: Entities Menu

Table 3: Entities Menu Options

Category	Name	Description
Dashboard		Displays the main dashboard (see Main Dashboard on page 47).
Virtual Infrastructure	VMs	Displays the VMs dashboard (see VMs Summary View on page 72).

Category	Name	Description
	Storage Containers	Displays the Storage Containers dashboard (see Storage Containers Summary View on page 87).
	Catalog Items	Displays the Catalog dashboard (see Catalog Items View on page 97).
	Images	Displays the Images dashboard (see Images Summary View on page 98).
	Categories	Displays the Categories dashboard (see Categories Summary View on page 102).
	Recoverable Entities	Displays the Recoverable Entities dashboard (see Recoverable Entities View on page 104).
	Subnets	Displays the Subnets dashboard (see Subnets View on page 105).
Policies	Security Policies	Displays the Security Policies dashboard (see Security Policies Summary View on page 106).
	Protection Policies	Displays the Protection Policies dashboard (see Protection Policies View on page 109).
	Recovery Plans	Displays the Recovery Plans dashboard (see Recovery Plans View on page 110).
	NGT Policies	Displays the NGT Policies dashboard (see NGT Policies View on page 111).
	Image Placement Policies	Displays the Image Placement Policies dashboard (see Image Placement Policies on page 379).
Hardware	Clusters	Displays the Clusters dashboard (see Clusters Summary View on page 114).
	Hosts	Displays the Hosts dashboard (see Hosts Summary View on page 131).
	Disks	Displays the Disks dashboard (see Disks Summary View on page 141).
	GPUs	Displays the GPUs dashboard (see GPUs Summary View on page 150).
Activity	Alerts	Displays the Alerts dashboard (see Alerts Summary View on page 214)
	Events	Displays the Events dashboard (see Events Summary View on page 220).
	Audits	Displays the Audits dashboard (see Audits Summary View on page 153).
	Tasks	Displays the Tasks dashboard (see Tasks View on page 157).
Operations	Analysis	Displays the Analysis dashboard (see Analysis Dashboard on page 202).

Category	Name	Description
	Planning	Displays the capacity Planning dashboard (see Resource Planning on page 176).
	Reports	Displays the Reports dashboard (see Reports View on page 159).
	Playbooks	Displays the list of playbooks (see Codeless Task Automation (X-Play) on page 415).
	Action Gallery	Displays the list of available actions (see Codeless Task Automation (X-Play) on page 415).
	Plays	Display the list of execution instances of playbooks (see Codeless Task Automation (X-Play) on page 415).
Administration	LCM	Displays the life cycle manage (LCM) dashboard (see LCM View on page 161).
	Projects	Displays the Projects dashboard (see Projects Summary View on page 161).
	Roles	Displays the Roles dashboard (see Roles Summary View on page 167).
	Users	Displays the Users dashboard (see Users Summary View on page 170).
	Availability Zones	Displays the Availability Zones dashboard (see Availability Zones View on page 174).
Services	Calm	Displays the Calm feature page (see Enabling Calm on page 212).
	Karbon	Displays the Karbon feature page (see Enabling Karbon on page 213).
	Objects	Displays the Objects feature page (see Enabling Objects on page 213).
Prism Central Settings		Displays the settings menu (see Settings Menu on page 14).

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 9).
- Click the collapse menu ("hamburger") button and then select **Prism Central Settings** from the entities menu (see Entities Menu on page 12).

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.

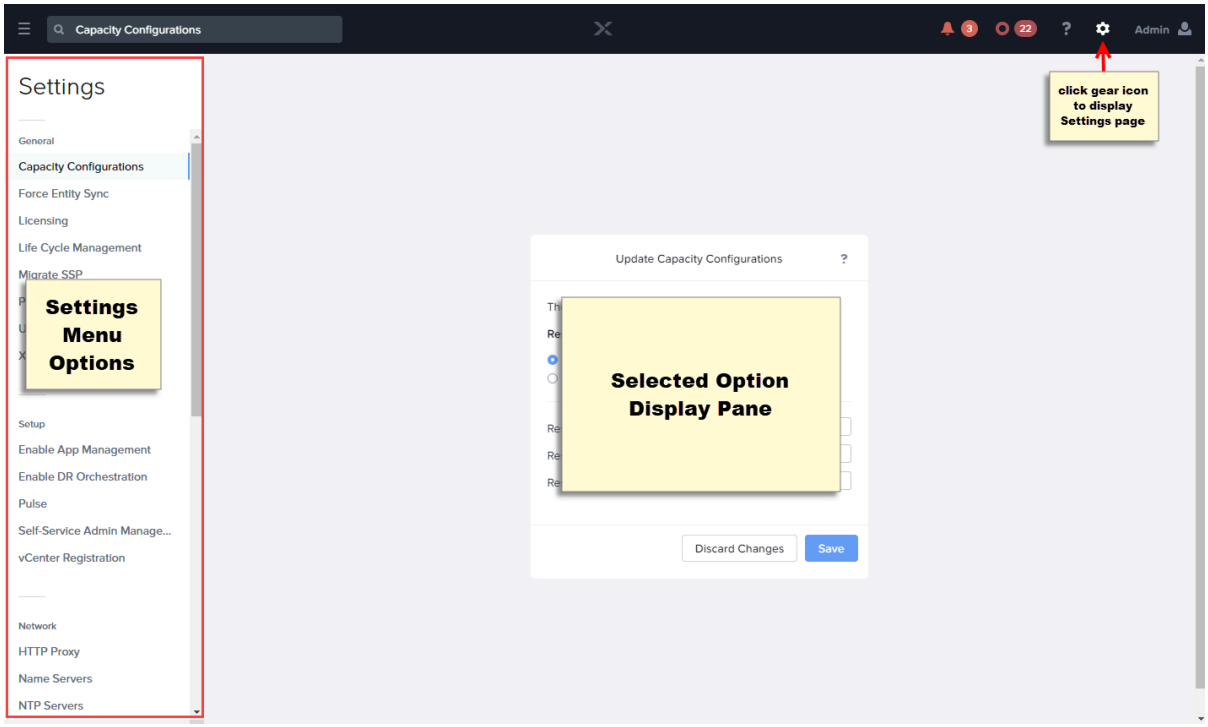

Figure 3: Settings Page

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 196). 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 27).
	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 372).
	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 31).

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 27).
	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	Enable App Management	Opens the **Enable App Management** page to enable the Nutanix Calm features (see Enabling Calm on page 212).
	Connect to Frame	Opens the **Connect to Frame** page to connect to Xi Frame. See the Xi Frame documentation for more information.
	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 435).
	Self-Service Admin Management	Opens the Self-Service Admin Management page to modify the self-service configuration (see Configuring Prism Self Service on page 369). 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 364).
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 277).
	Name Servers	Opens the **Name Servers** window to configure name servers (see Configuring Name Servers on page 280).
	NTP Servers	Opens the **NTP Servers** window to specify which NTP servers to access (see Configuring NTP Servers on page 281).
	SNMP	Opens the **SNMP Configuration** window to configure SNMP (see Configuring SNMP on page 283).
Flow	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. For details, see “” () .

Category	Name	Description
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 “” () .
Email and Alerts	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).
	Alert Policies	Opens the **Alert Policies** window, which allows you to specify what events should generate an alert and how frequently the system should check for each event type (see Configuring Alert Policies on page 226).
	SMTP Server	Opens the **SMTP Server Settings** window to configure an SMTP server (see Configuring an SMTP Server on page 282).
	Syslog Server	Opens the **Syslog Servers** window to configure a syslog server (see Configuring Syslog Monitoring on page 290).
Appearance	Language Settings	Opens the **Language Settings** window to localize the Prism Central screen text (see Changing the Language Settings (Prism Central) on page 293).
	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 290).
	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 291).

Searching for Information

The main menu includes a search field on the left (see Main Menu on page 9) that allows you to find information about selected entities in various ways. An entity is an object type such as a VM, cluster, security policy, project, or report. The search field is context-sensitive, which means it is populated automatically based on where you are in Prism Central. The first screen after logon is the main dashboard, so "Dashboard" appears in the search field initially.

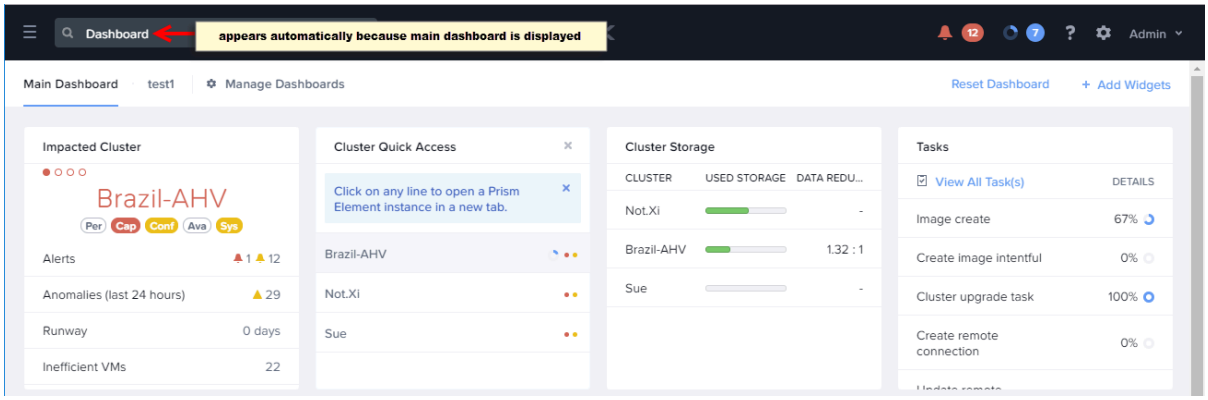

Figure 4: Search Field (initial screen)

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.

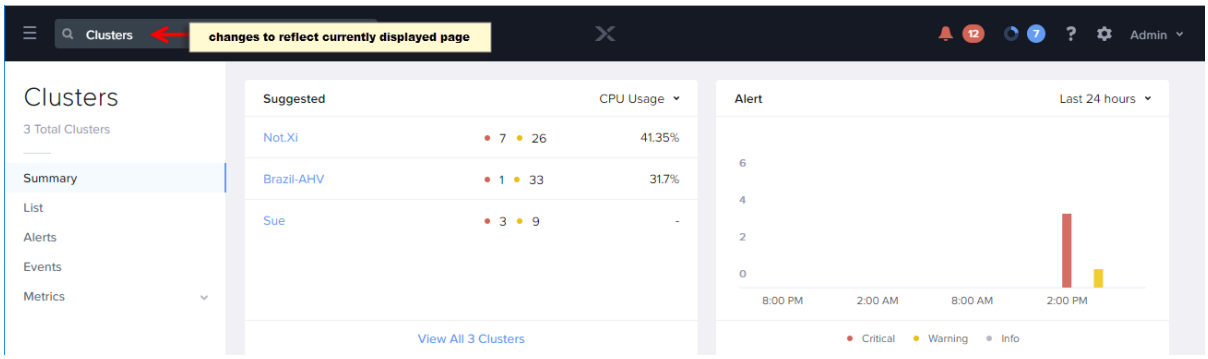

Figure 5: Search Field (from a different screen)

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:

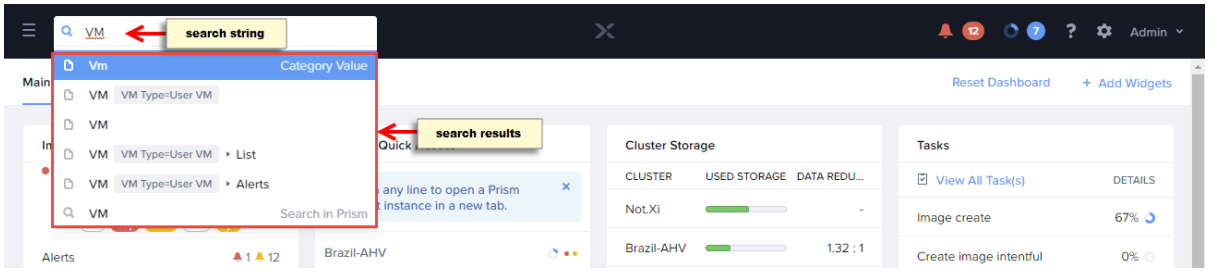

Figure 6: Search Results (for an entity type)

- "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.

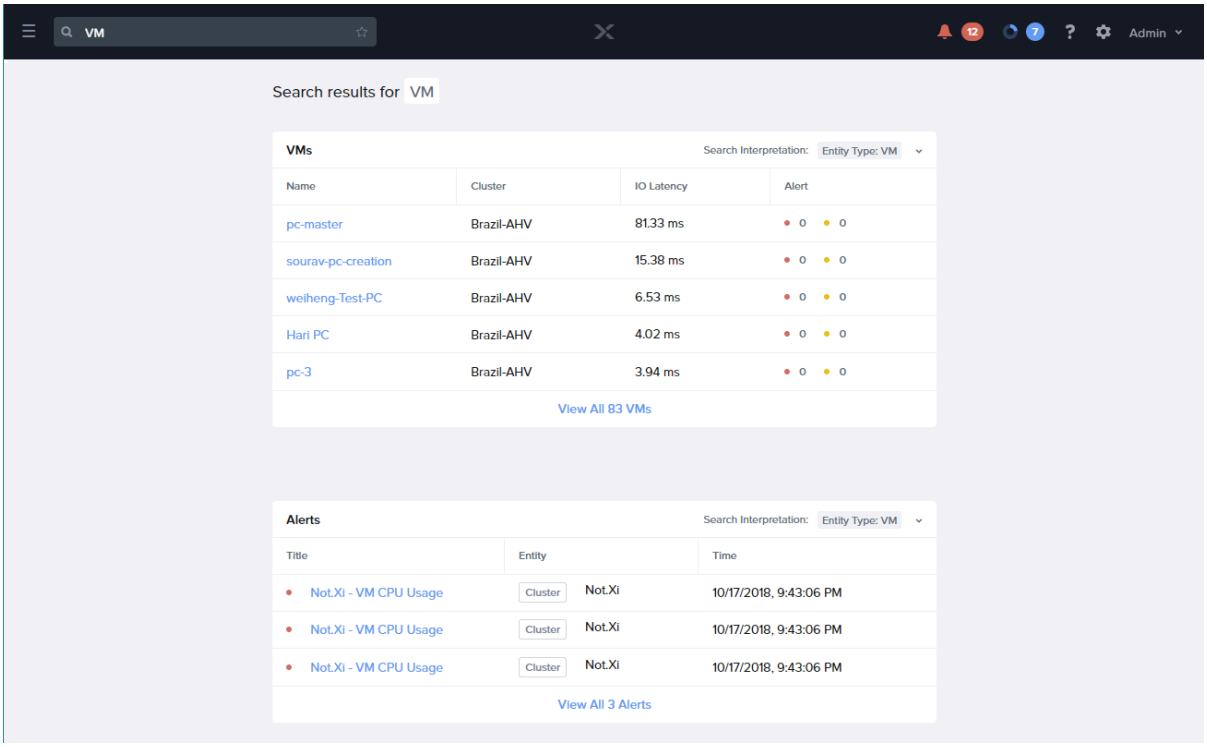

Figure 7: "Search in Prism" Results

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.

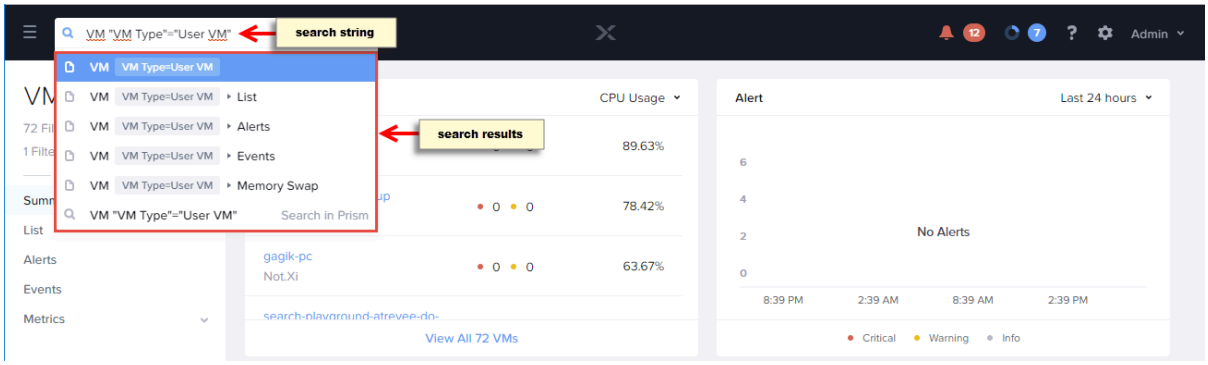

Figure 8: Search Results (for the current location)

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.

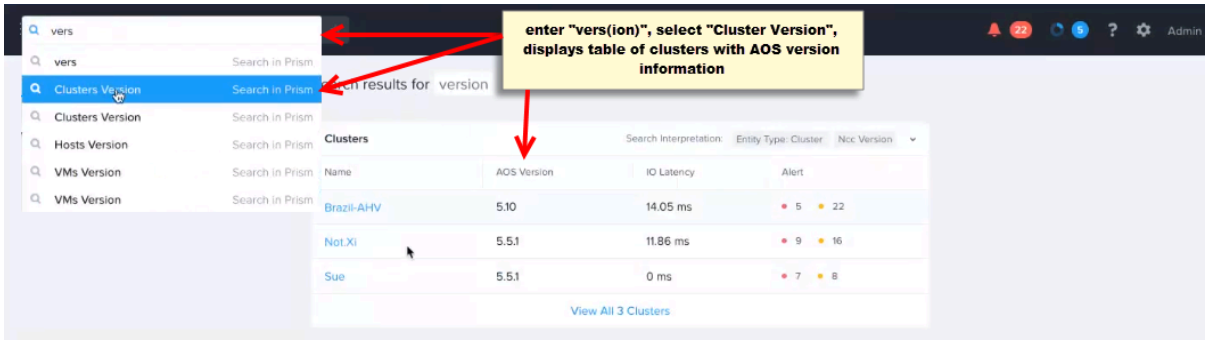

Figure 9: Search Results (for generic string)

You can bookmark a filtered page (see Entity Exploring on page 66) 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 12).

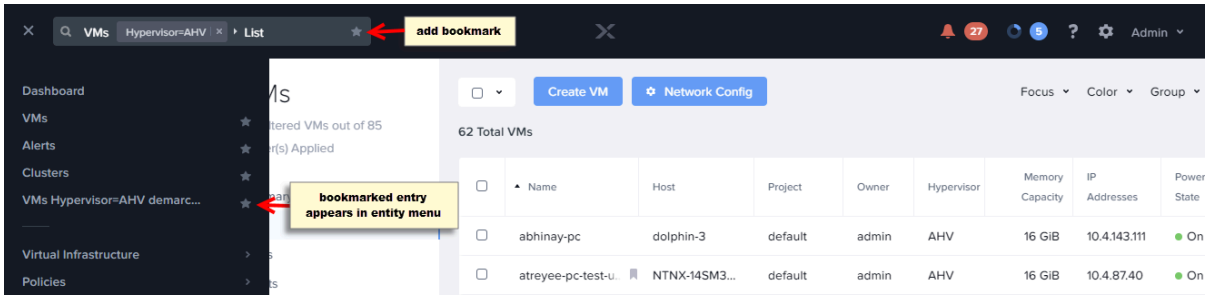

Figure 10: Bookmark Entry

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.

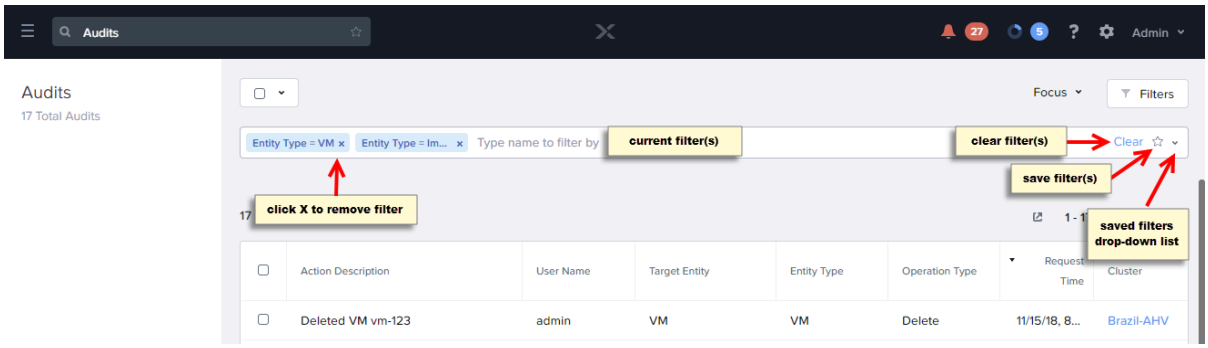

Figure 11: Local Filters Field

Search Guidelines

If you need help in navigating Prism Central or applying a search filter, 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.

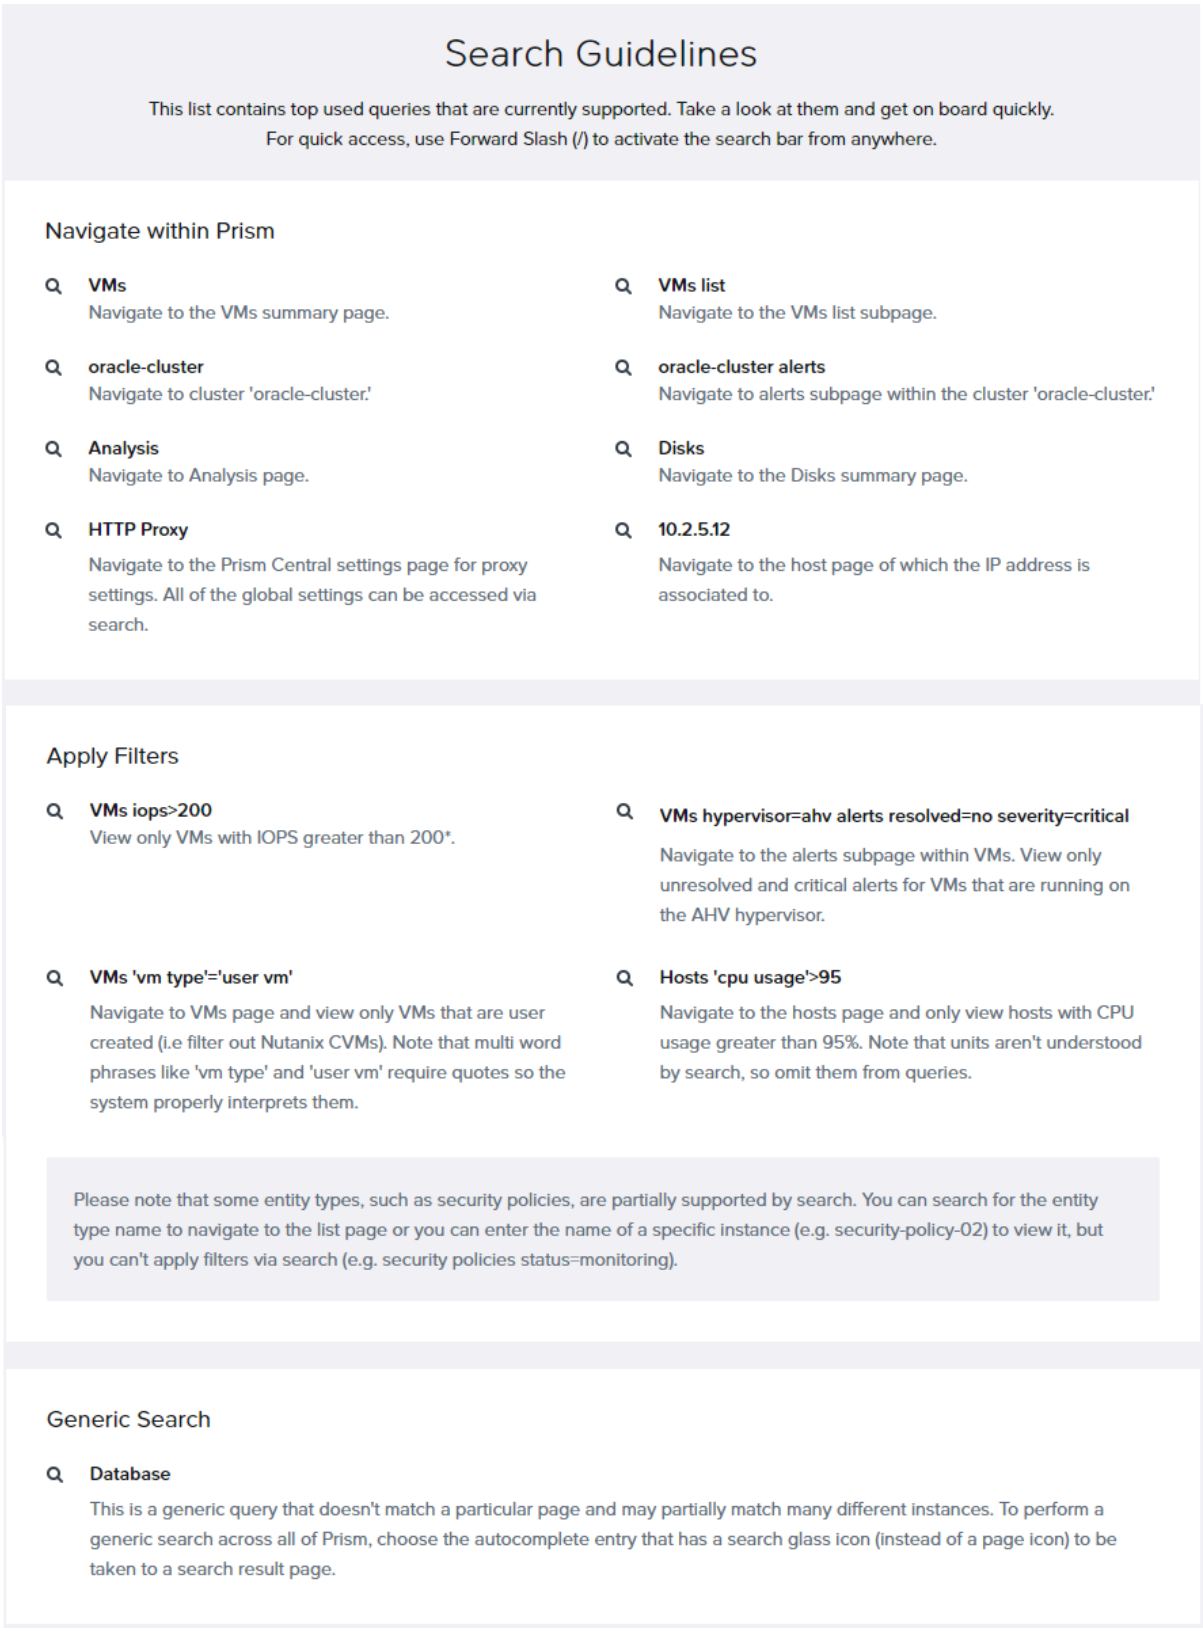

Figure 12: Search Guidelines Page

Query Rules

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. The following operators are supported in expressions:

  ```
  =  <  >  <=  >=
  ```

- <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).

Here are examples of various query types.

- Entity queries:

  ```
  <cluster_name>
  <cluster_ip_address>
  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”
```

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.
  - 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 5: Search Keywords

Object	Entity	Keywords
Entities	(n/a)	vm, cluster, node, storage 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
	Storage Container	storage container name, rf, on disk deduplication, perf-tier deduplication, compression, erasure coding

Object	Entity	Keywords
	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, resolved, unresolved, acknowledged, unacknowledged Severity levels: critical, warning, info Categories (impact types): capacity, performance, configuration, availability, system indicator
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

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 6: 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 27)
- Managing licenses (see Licensing on page 27)
- 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 27)

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 if you choose a non-Nutanix environment.

> Note: Most of the following links are to topics in the *Acropolis Upgrade Guide*.

- To install or upgrade a Prism Central instance, see “” () .
- To scale out an existing single VM Prism Central to multiple VMs, see Expanding (Scale Out) Prism Central on page 42.
- To upgrade NCC through Prism Central, see “” () .
- To upgrade one or more managed clusters through Prism Central, see “” () .
- If you use a proxy server, see “” () .

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the Nutanix Licensing 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 27), you can

- Log in and out through any supported browser (see Logging Into Prism Central on page 28 and Logging Out of Prism Central on page 31).
- Monitor the status of this Prism Central instance (see Managing Prism Central on page 31).
- Check version information (see Checking Prism Central Version on page 34).
- Modify login page and automatic log off settings (see Modifying UI Settings on page 290).

- Register clusters with this Prism Central instance (see Register (Unregister) Cluster with Prism Central on page 34).
- Expand Prism Central from a single VM to a multi-VM instance (see Expanding (Scale Out) Prism Central on page 42).
- Shutdown or Start Up a Prism Central VM (see Shutdown (Start Up) Prism Central VM on page 46)

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.

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.

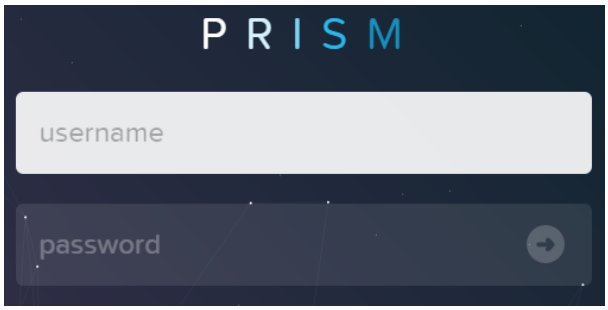

Figure 13: Login Screen

> 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.)

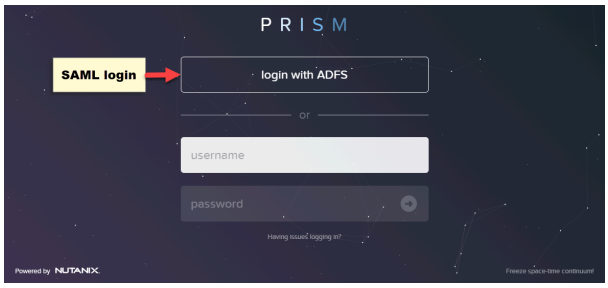

Figure 14: Login Screen (SAML)

> 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 290).

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.

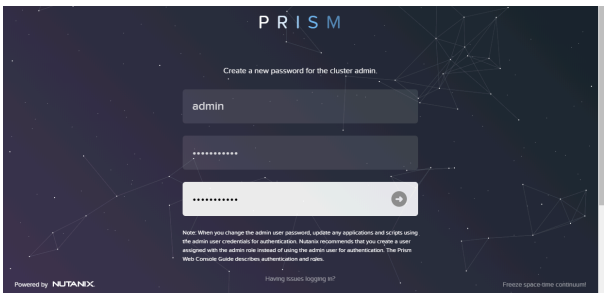

Figure 15: Login Screen (first admin login)

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:

- 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.

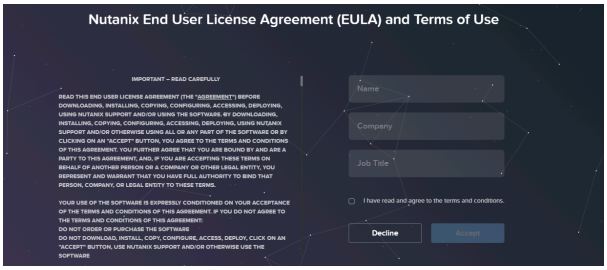

Figure 16: EULA Screen

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 14). This displays the Manage Prism Central page. This page contains the following:

> Note: This page is not available if Prism Central is hosted in a non-Nutanix environment, that is not hosted in a Nutanix cluster.

- A "Prism Central Capacity" pane on the 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 **View your clusters** link to display the clusters summary page.
  - Number of VMs that comprise this Prism Central instance; click the **See VM details** link to display the details page for the Prism Central VM.
- A "Prism Central VMs" pane in the middle that displays network address information (VLAN tag, 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 42).

- An "Alerts" pane on the right that displays a list of Prism Central-related alerts broken into Critical, Warning, and Info sections. Click an alert to see details.

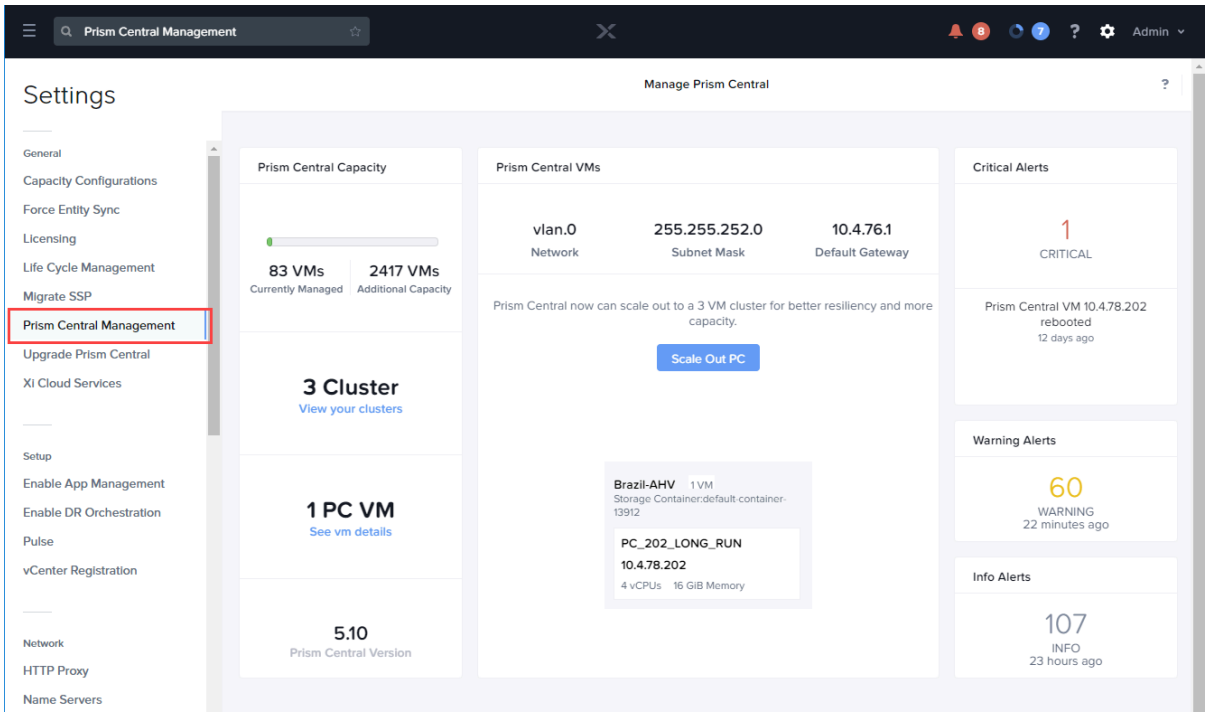

Figure 17: Manage Prism Central Page (1-VM instance)

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. 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 42 are satisfied before adding a Prism Central VM.

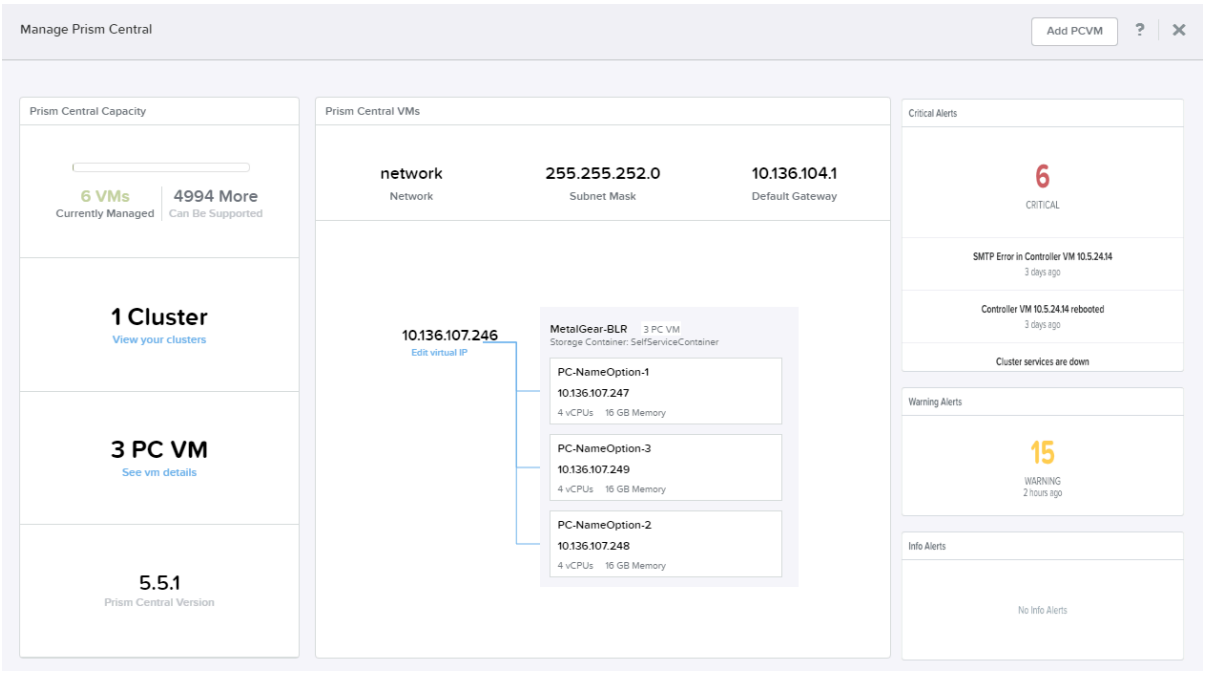

Figure 18: Manage Prism Central Page (3-VM instance)

Checking Prism Central Version

Procedure

1. To check the Prism Central version, 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.

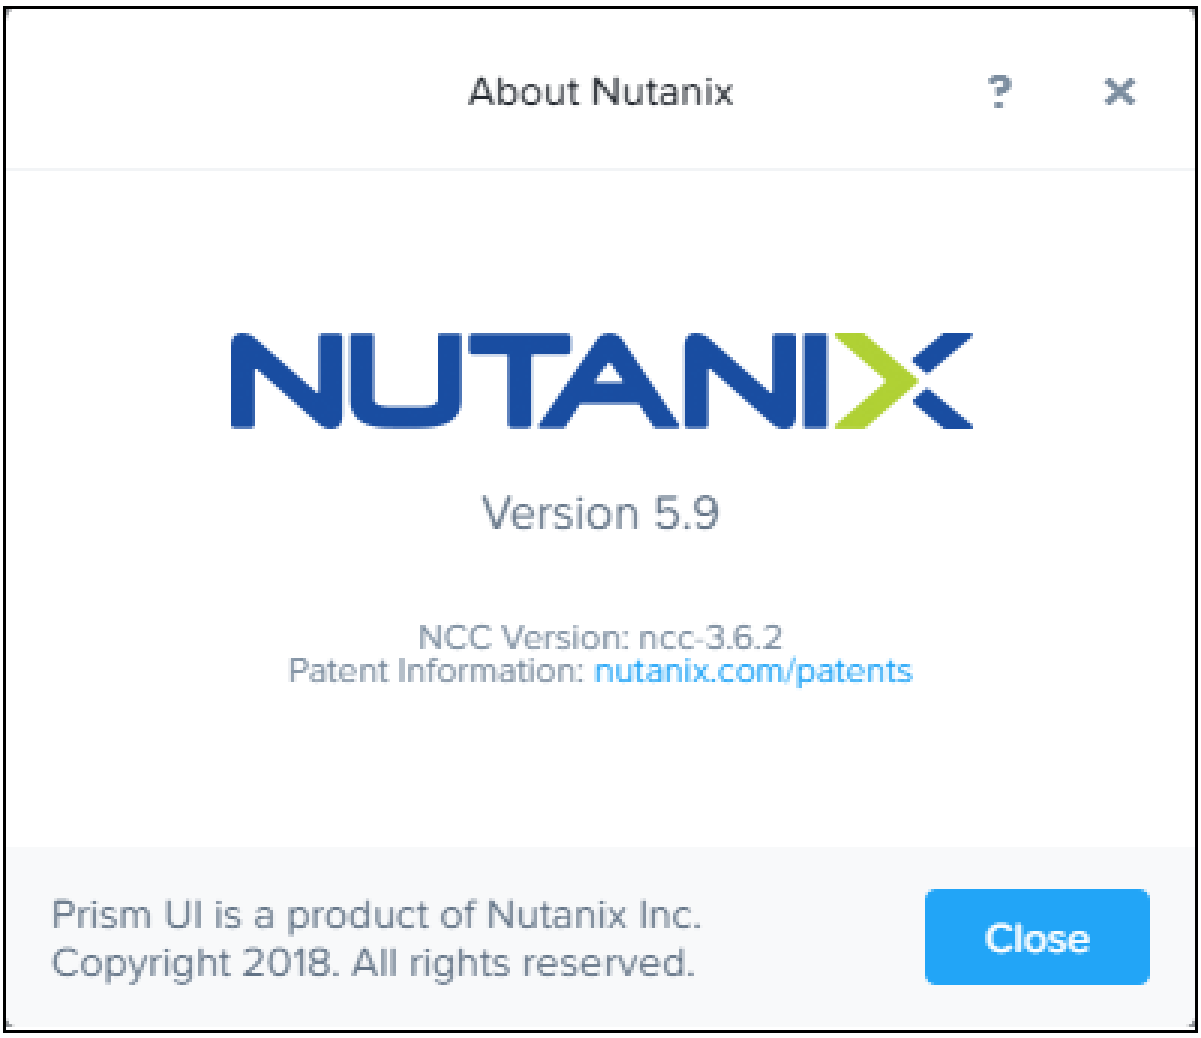

Figure 19: About Nutanix Window

2. Click the **OK** button to close the window.

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 (see Logging Into Prism Central on page 28).

- 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.
- 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. Run NCC as described in Run NCC 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 20: Home Dashboard

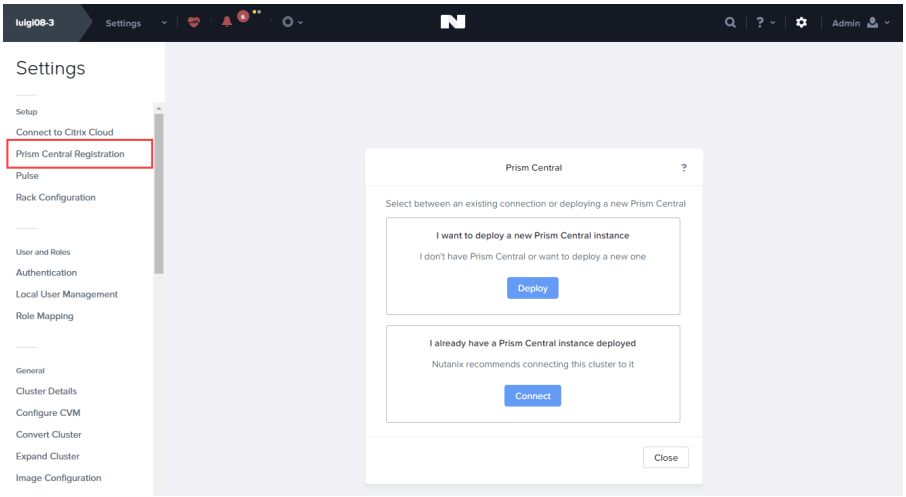

Figure 21: Settings Menu

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 “” () for instructions on how to deploy a new Prism Central instance.

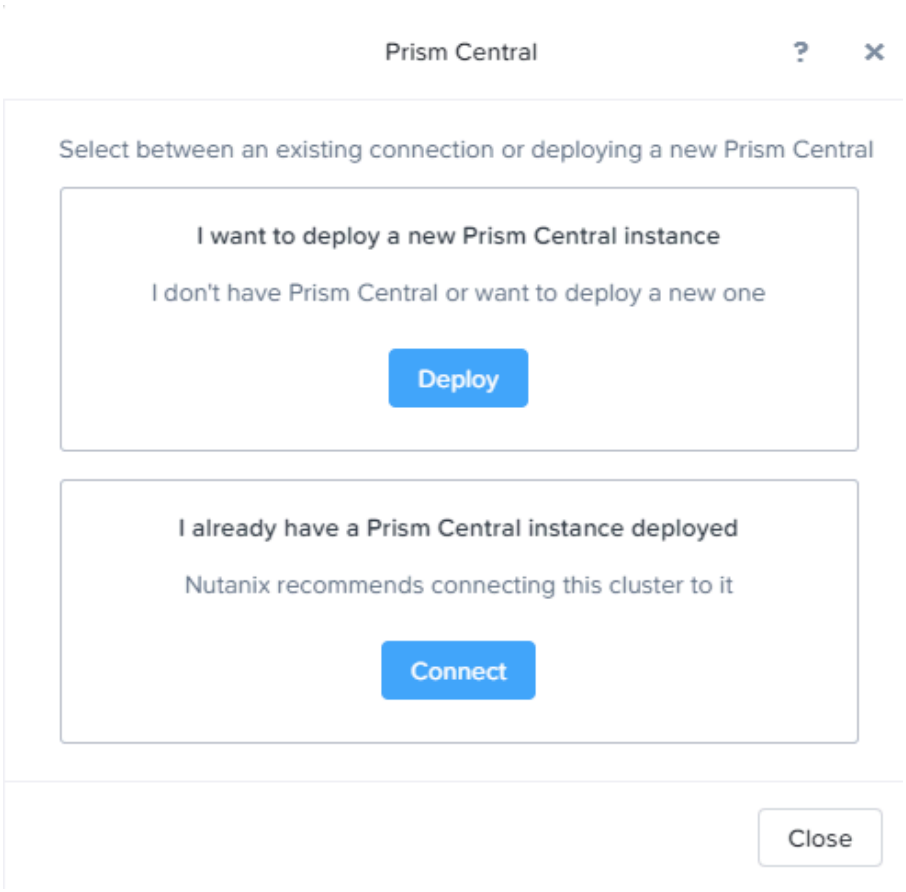

Figure 22: Prism Central Dialog Box (deploy/connect screen)

5. A screen appears outlining what services are available through Prism Element and Prism Central. After reviewing the message, click the **Next** button.

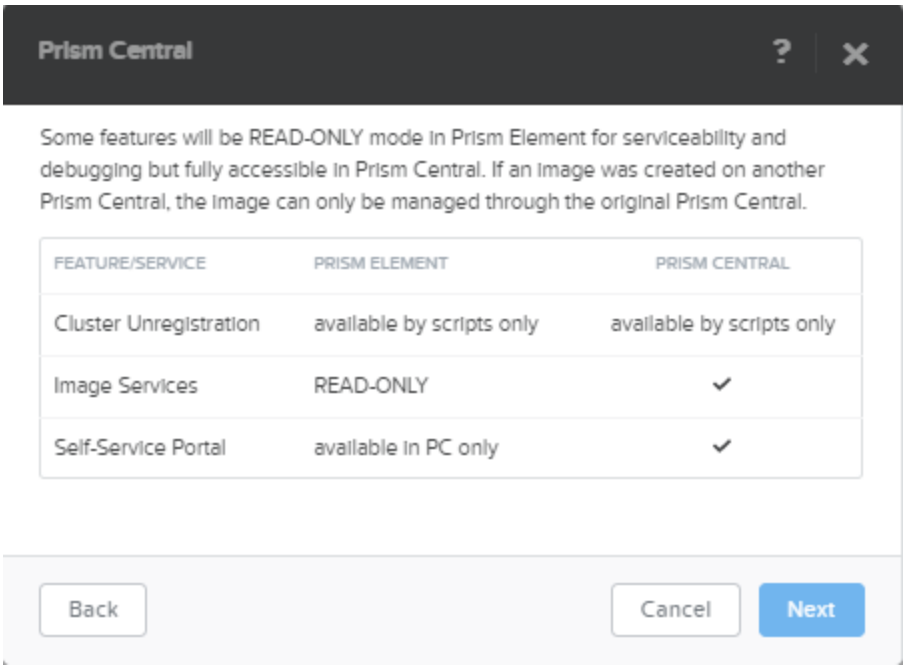

Figure 23: Prism Central Dialog Box (services screen)

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. **Username**: Enter `admin` as the Prism Central user name.
   c. **Password**: Enter the Prism Central `admin` user password.

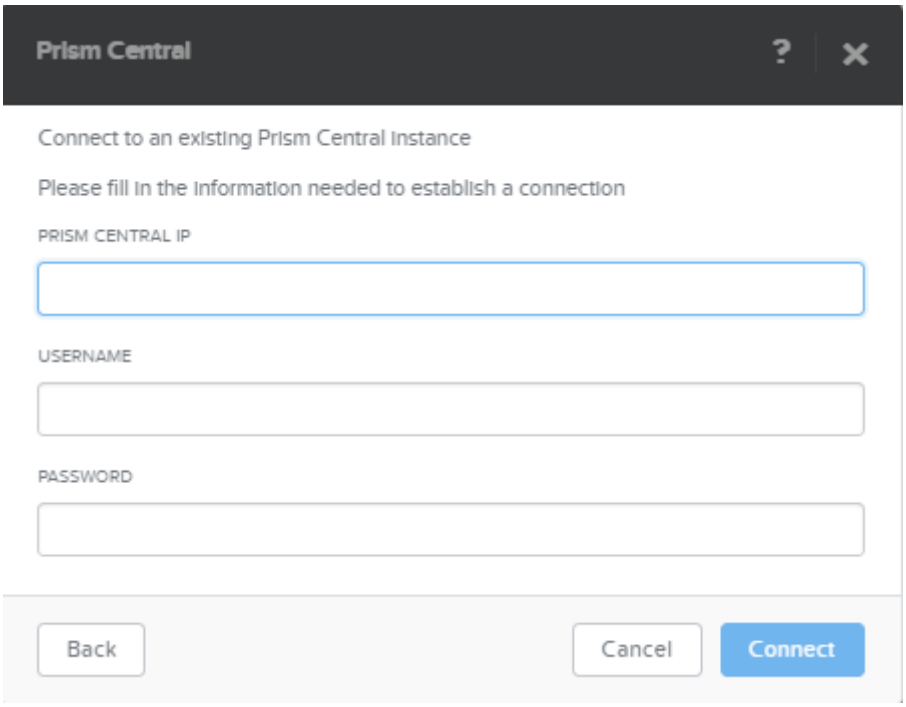

Figure 24: Prism Central Dialog Box (registration screen)

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: If the Prism Central IP address changes after registering, you must unregister and then re-register the cluster(s) with the new IP address. However, changing the Prism Central administrator password after registering any clusters does not require re-registration.

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.

```
<acropolis> exit
nutanix@NTNX-[redacted]-A-CVM:10.5.25.30:~$
nutanix@NTNX-[redacted]-A-CVM:10.5.25.30:~$ cluster status
2017-11-19 21:31:20 INFO zookeeper_session.py:110 cluster is attempting to connect to Zookeeper
2017-11-19 21:31:20 INFO cluster:2420 Executing action status on SVMs 10.5.25.30
The state of the cluster: start
Lockdown mode: Disabled

        CVM: 10.5.25.30 Up, ZeusLeader
                                Zeus   UP       [1966, 1995, 1996, 1997, 2043, 2061]
                           Scavenger   UP       [2708, 2736, 2737, 2738]
                      SSLTerminator    UP       [2104, 2136, 2137, 2138]
                     SecureFileSync    UP       [2108, 2162, 2163, 2164]
                             Medusa    UP       [2278, 2312, 2313, 2318, 2421]
                  DynamicRingChanger   UP       [2520, 2660, 2661, 2665]
                             Pithos    UP       [2524, 2586, 2587, 2615]
                             Mantle    UP       [2528, 2590, 2591, 2620]
                               Hera    UP       [2533, 2608, 2609, 2610]
                           Stargate    UP       [2823, 2893, 2894, 3250, 3258]
                         InsightsDB    UP       [2826, 2876, 2877, 3027]
                InsightsDataTransfer   UP       [2834, 2916, 2917, 3001, 3004, 3005, 3006]
                              Ergon    UP       [2842, 2963, 2964, 2965]
                             Cerebro   UP       [2901, 2984, 2985, 3448]
                             Chronos   UP       [2920, 2999, 3000, 3074]
                             Curator   UP       [2942, 3047, 3048, 3506]
                               Prism   UP       [3024, 3094, 3095, 3201, 3707, 3776]
                                 CIM   UP       [3037, 3154, 3155, 3214]
                        AlertManager   UP       [3046, 3283, 3286, 3422]
                            Arithmos   UP       [3083, 3215, 3216, 3488]
                             Catalog   UP       [3106, 3189, 3190, 3191]
                           Acropolis   UP       [3173, 3294, 3295, 3296]
                               Uhura   UP       [3192, 3389, 3390, 3396]
                                Snmp   UP       [3198, 3391, 3392, 3395]
                    SysStatCollector   UP       [3208, 3445, 3446, 3447]
                              Tunnel   UP       [3231, 3582, 3583]
                               Janus   UP       [3304, 3485, 3486]
                     NutanixGuestTools  UP       [3517, 3617, 3618, 3683]
                          MinervaCVM   UP       [4029, 4064, 4065, 4066, 4213]
```

Figure 25: Cluster Status Command

3. Unregister the cluster from Prism Central.

   ```
   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.

   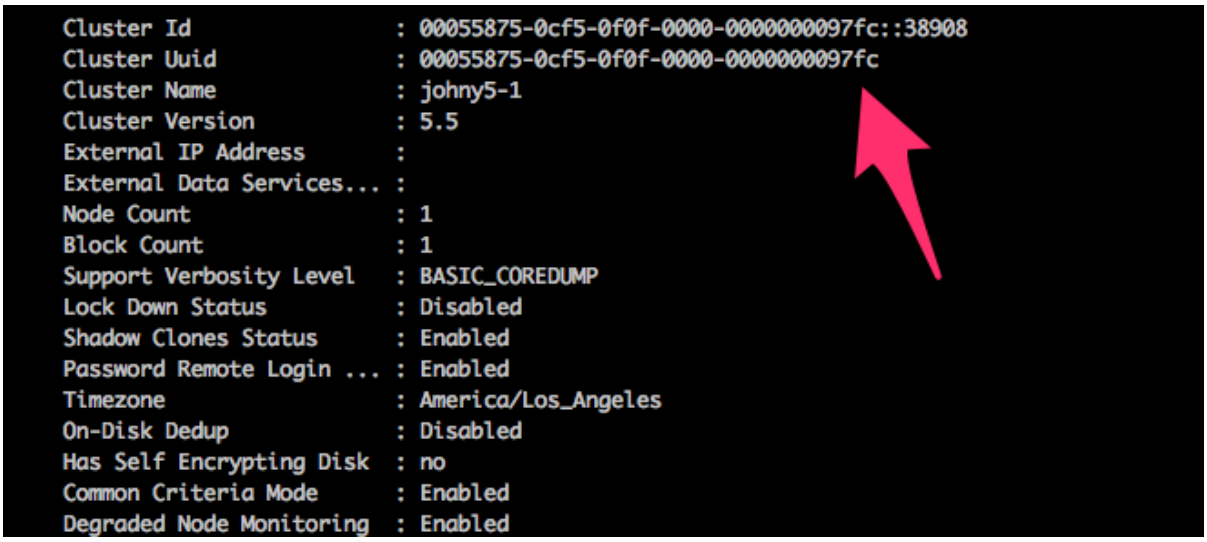

   Figure 26: Cluster Info Output

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 5. 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 5), 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 6c.

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 storage container used for deployment is mounted on all 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.

Procedure

1. Click the gear icon and then select **Prism Central Management** from the Settings menu (see Settings Menu on page 14).
   The Manage Prism Central page appears. This page provides information about this Prism Central instance.

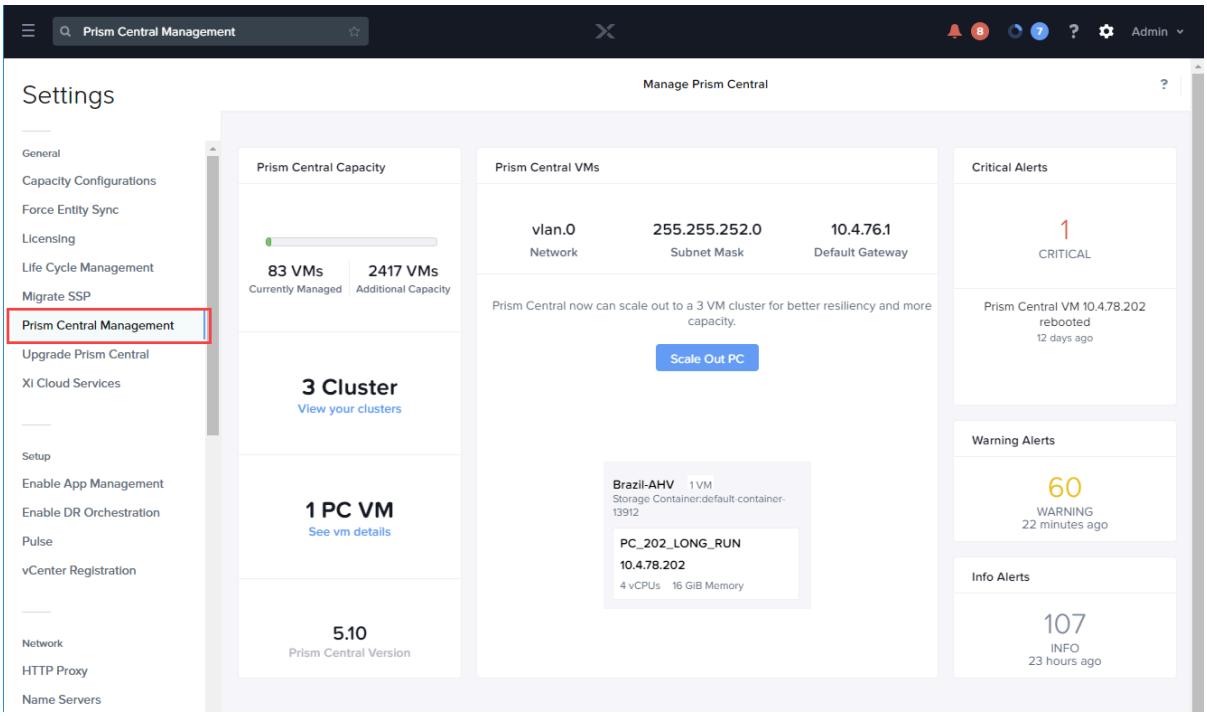

Figure 27: Manage Prism Central Page

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.

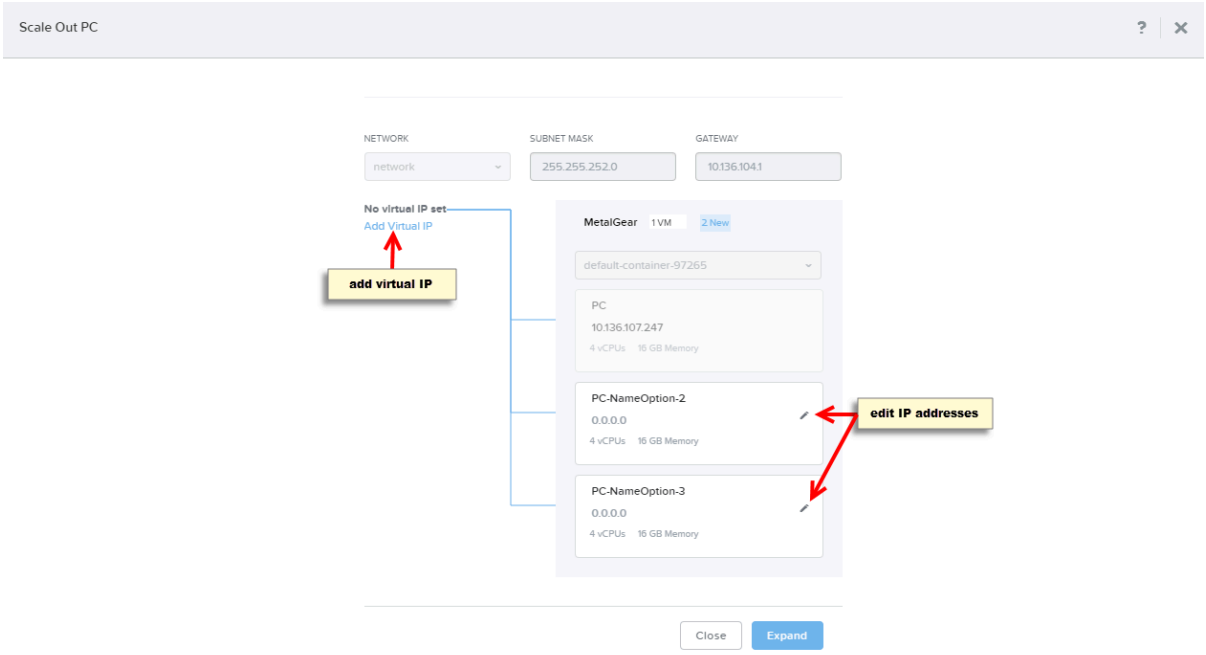

Figure 28: Scale Out PC Page

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 29: 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 157).

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.

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 on page 12).

Main Screen Layout

In addition to the main menu options (see Main Menu on page 9), 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 58).
- **Reset Dashboard** button. Click this button to reset the main dashboard to its original set of widgets.
- **Add Widget** button. Click this button to add a widget to the displayed dashboard (see Adding Dashboard Widgets on page 61).

> 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.

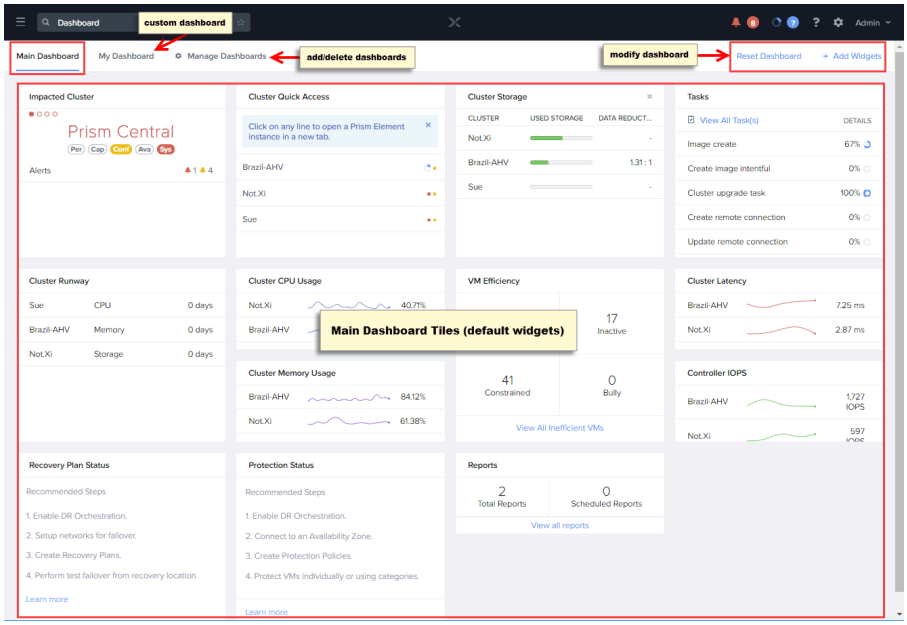

Figure 30: Home Dashboard

Main Dashboard Display

The main dashboard displays the following nine information tiles (widgets) by default. You can customize this view by adding or deleting widgets as desired (see Modifying a Dashboard on page 59).

- **Impacted Cluster**. This widget identifies registered clusters that have active (unresolved) alerts and might need attention. When multiple clusters are impacted, radio buttons appear at the top left; click a button to display the information for that cluster. The impacted cluster name appears below the buttons. (The message "you have no alerts" appears when none of the clusters are impacted currently.) This widget includes the following fields:

  - **Alerts**. Displays active alert indicators. Bell icons appear to the right with small buttons above for each alert category. The icons and buttons are color coded, red for critical alerts and yellow for warnings. Click a bell icon or category button to display information about those alerts (see Alerts Summary View on page 214). The category buttons are **Per** (performance), **Cap** (capacity), **Conf** (configuration), **Ava** (availability), or **Sys** (system indicator).

  - **Anomalies (last 24 hours)**. Displays the number of anomalies recorded in the cluster during the past 24 hours. Click the number to display a list of those anomalies in the Events page (see Events Summary View on page 220). Anomalies are detected through the VM behavioral learning engine (see Behavioral Learning Tools on page 197).

  - **Runway**. Indicates whether there is a potential storage, CPU, or memory resource constraint in the cluster. The estimated "runway" (time remaining before resource usage reaches capacity based on current trends) is displayed. A dash (-) indicates there is no runway information. For more detailed runway information, click the number to display the **Capacity** tab for that cluster (see Cluster Details View on page 121).

  - **Inefficient VMs**. Displays the number of inefficient VMs in the cluster currently. Click the number to display a list of those VMs (see VMs Summary View on page 72).

"Inefficient" VMs are detected through the VM behavioral learning engine (see Behavioral Learning Tools on page 197).

Figure 31: Impacted Cluster Widget

- **Cluster Quick Access**. This widget allows you to open 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.

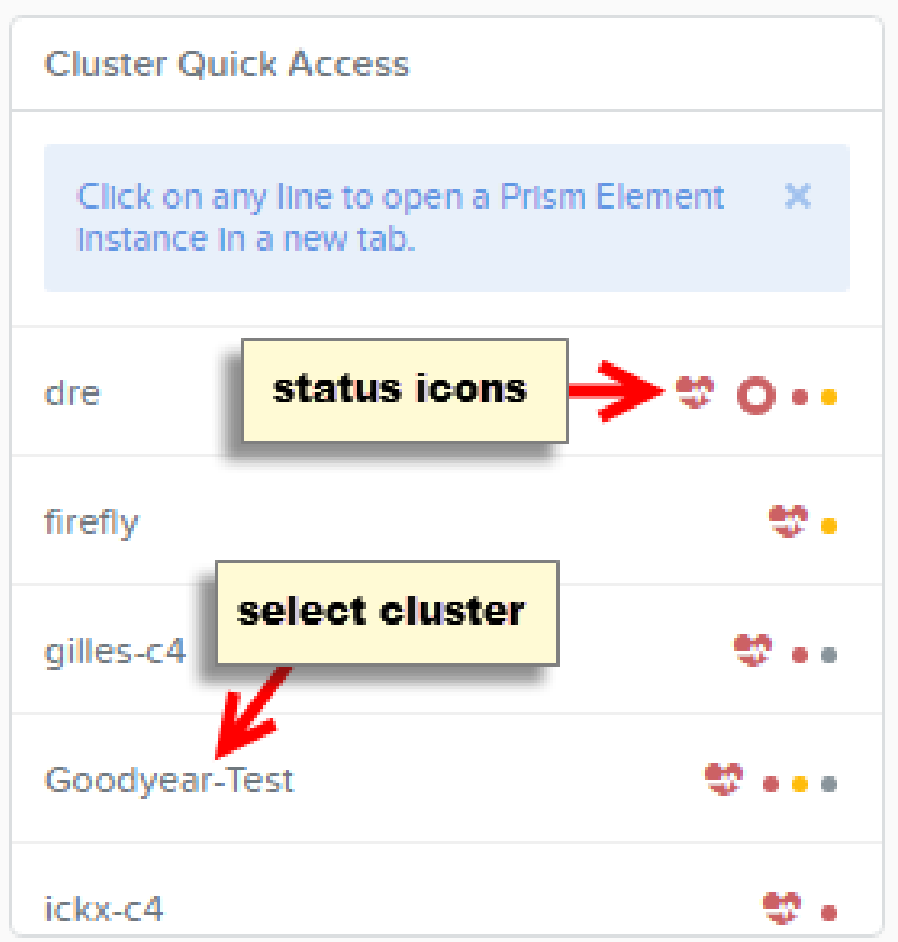

Figure 32: Cluster Quick Access Widget

- **Cluster Storage**. This widget displays information about storage use for the highest usage clusters. It includes a used storage column with a bar that visually indicates how much capacity is used currently and a data reduction column that displays the data savings ratio 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.) When the used storage exceeds a threshold, the bar color is red to highlight the situation, as illustrated

in the following figure. Clicking on a cluster name displays the summary page for that cluster (see Cluster Details View on page 121).

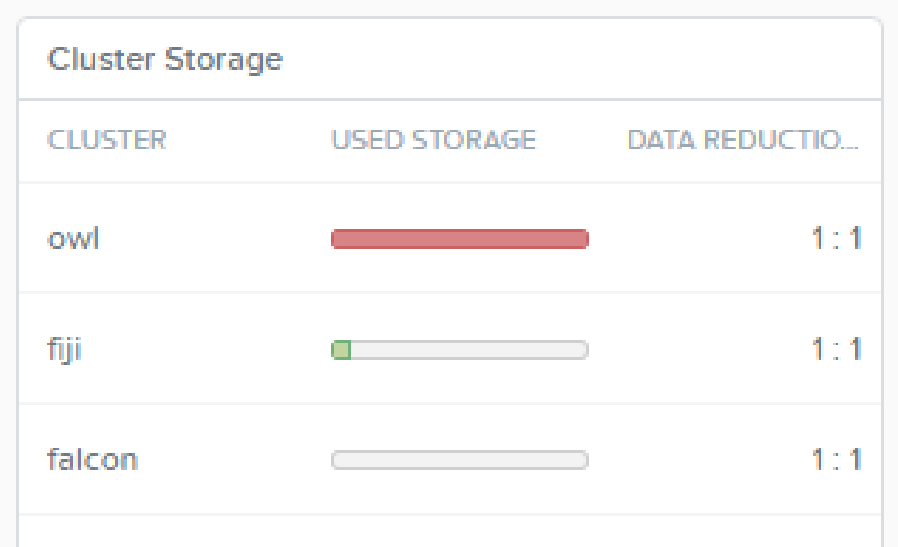

Figure 33: Cluster Storage Widget

- **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 (see Tasks View on page 157).

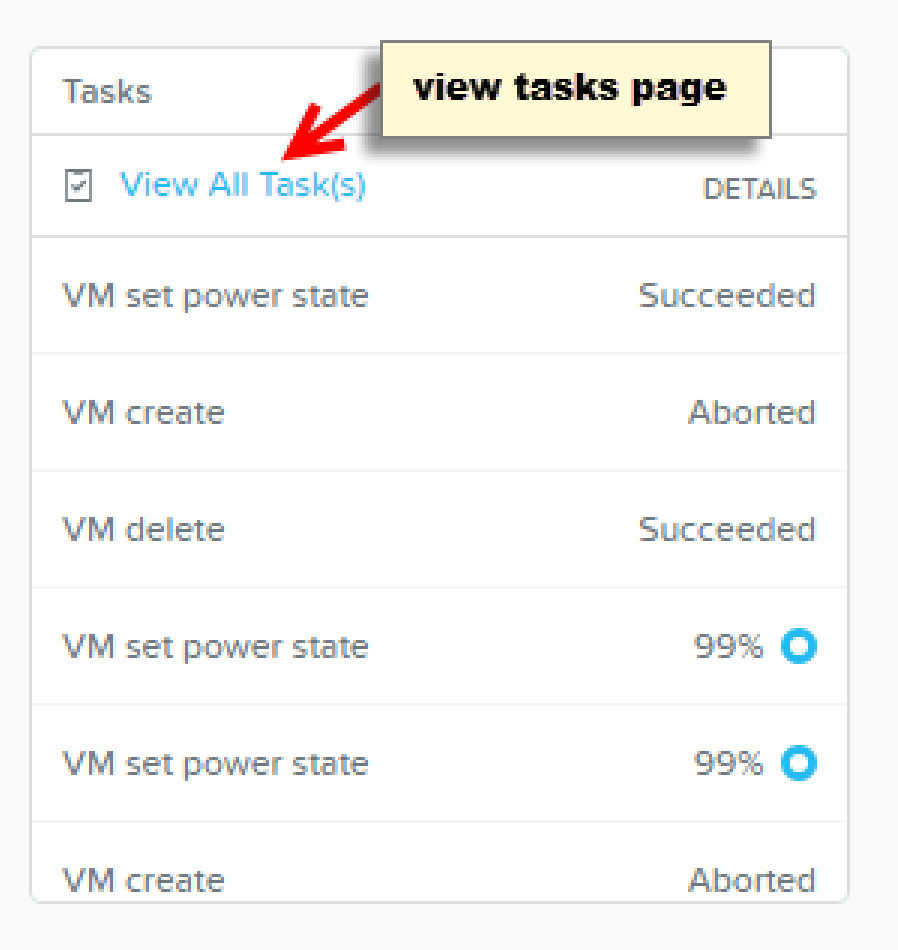

Figure 34: Tasks Widget

- **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 (see Cluster Details View on page 121).

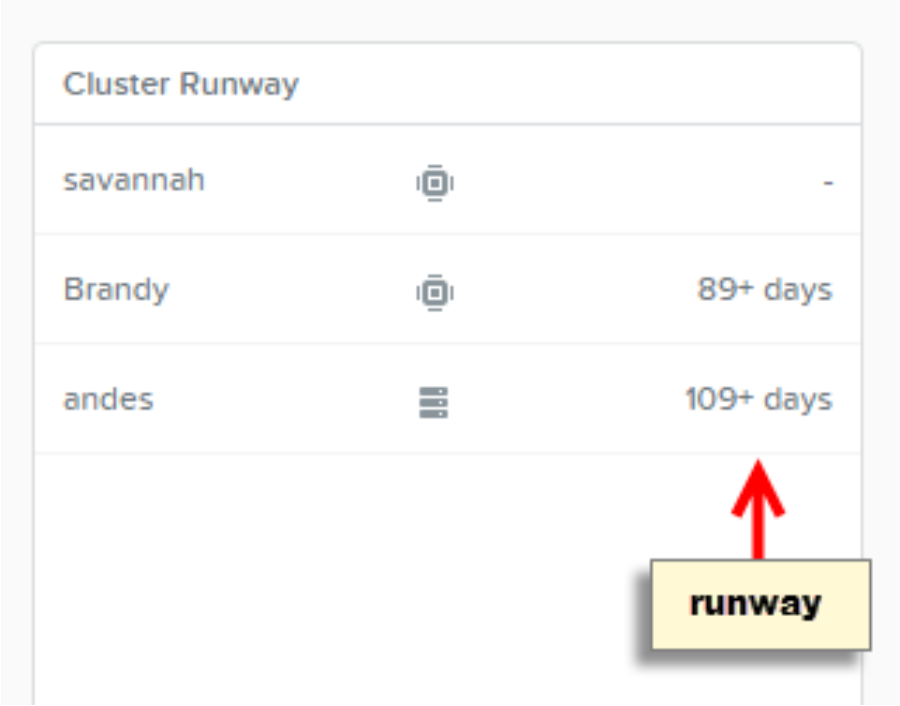

Figure 35: Cluster Runway Widget

- **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 (see Cluster Details View on page 121).

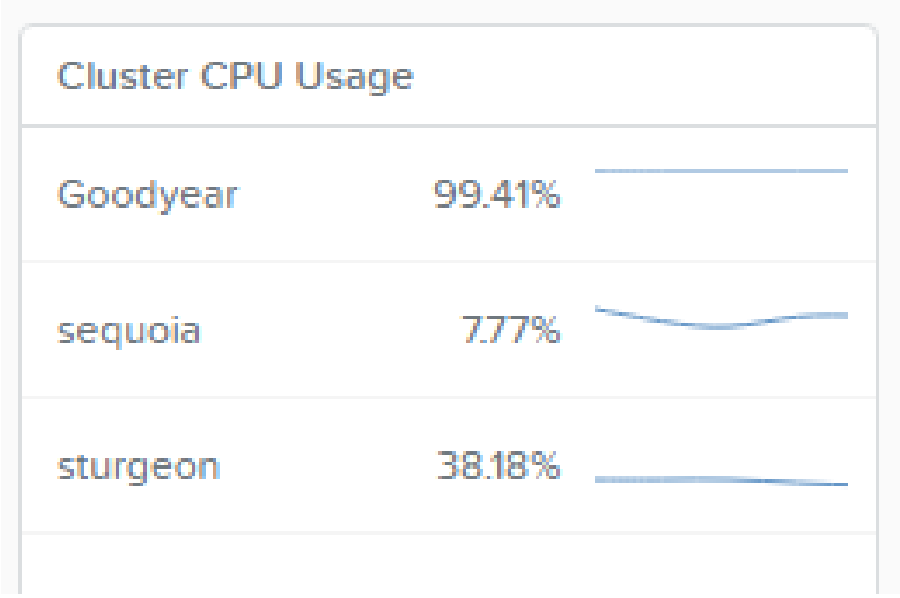

Figure 36: Cluster CPU Usage Widget

- **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 (see Cluster Details View on page 121).

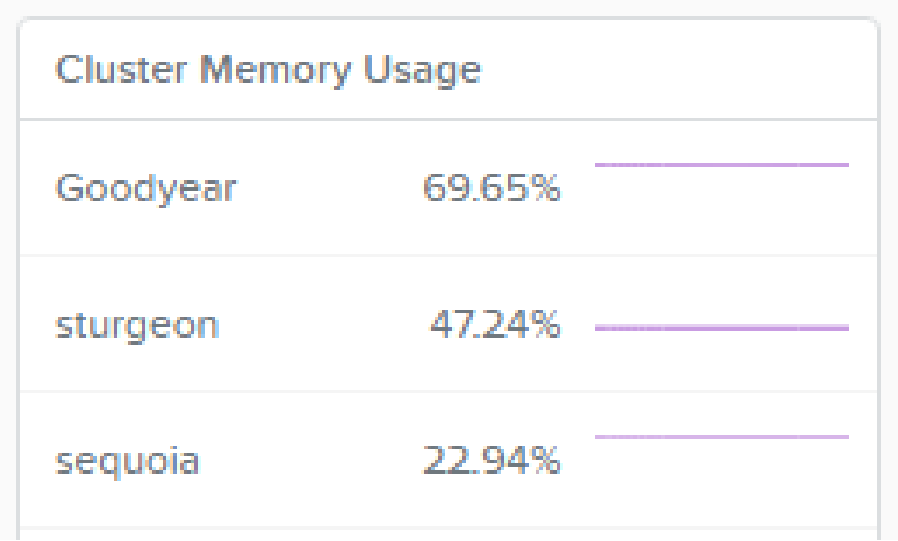

Figure 37: Cluster Memory Usage Widget

- **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 VMs Summary View on page 72). See Behavioral Learning Tools on page 197 for more information about VM efficiency.

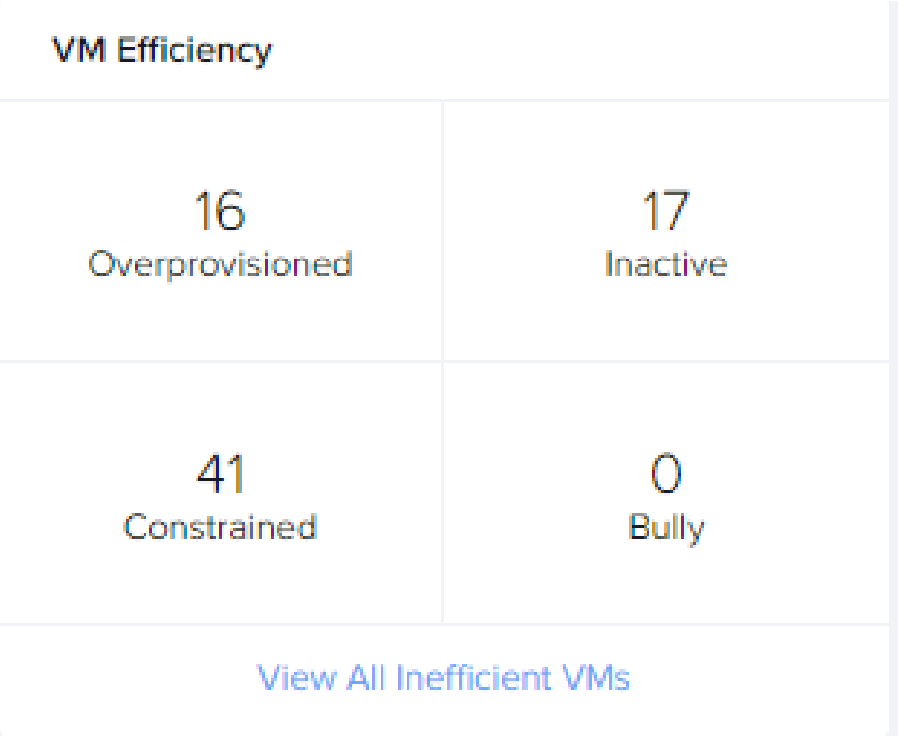

Figure 38: VM Efficiency Widget

- **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 (see Cluster Details View on page 121).

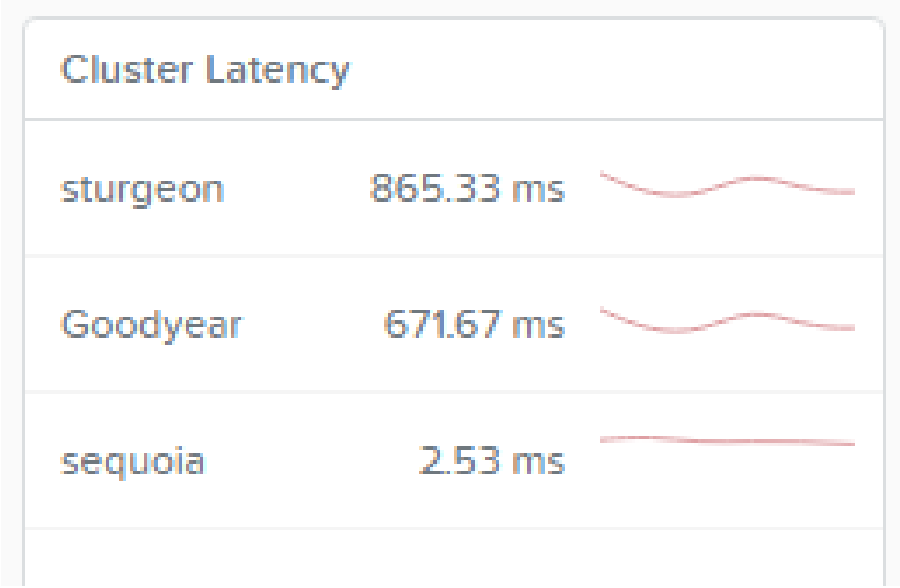

Figure 39: Cluster Latency Widget

- **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 Cluster Details View on page 121). The IOPS number comes from the controller when the hypervisor is AHV or Hyper-V and from the hypervisor when the hypervisor is ESXi.

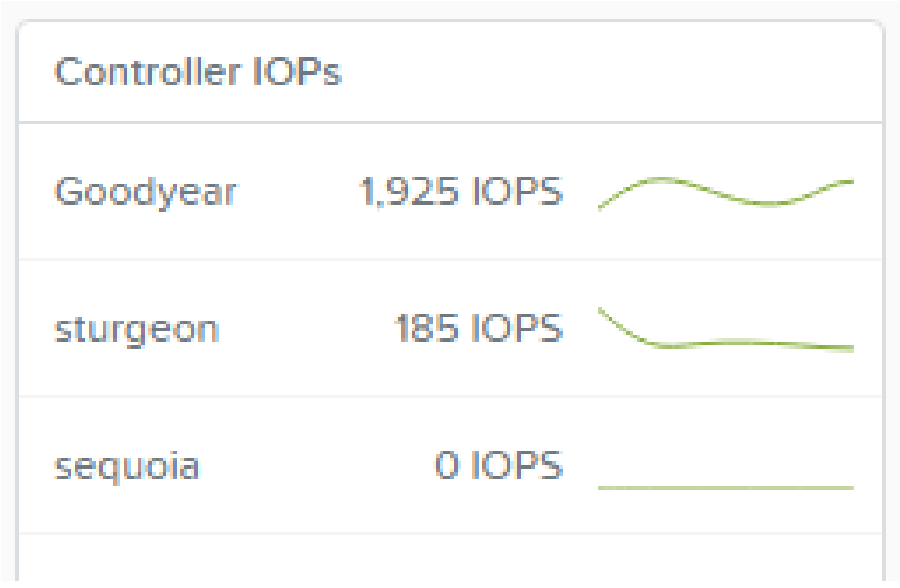

Figure 40: Controller IOPS Widget

- **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.

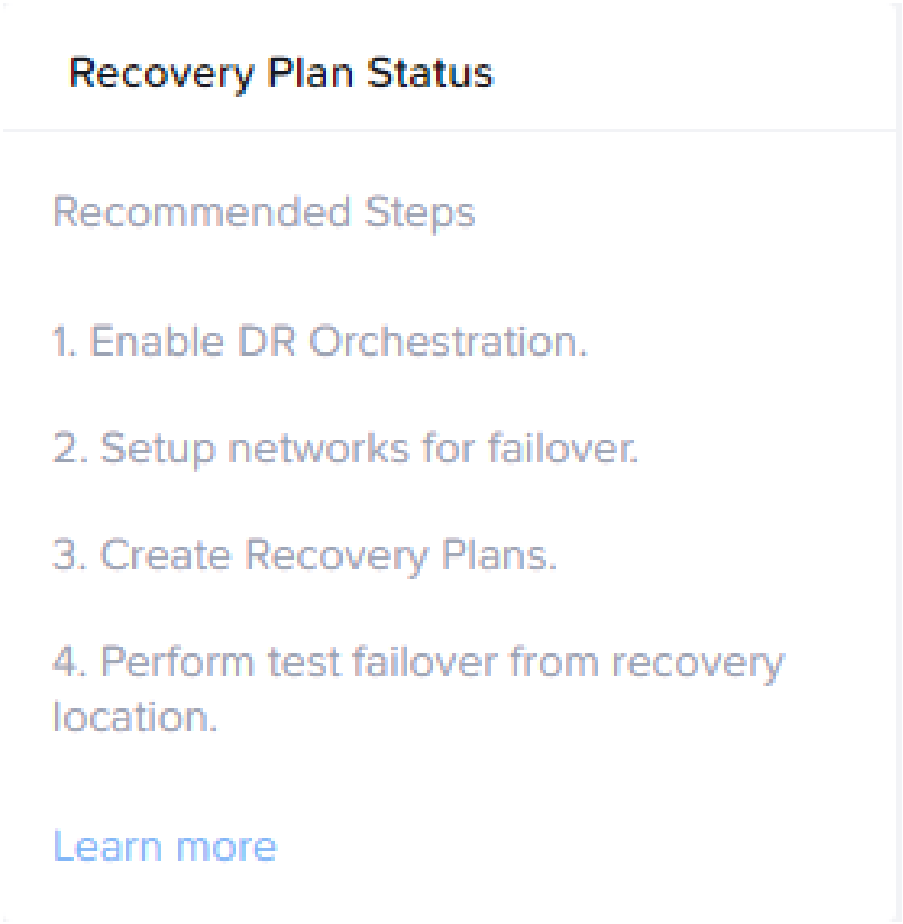

Figure 41: Recovery Plan Status Widget

- **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.

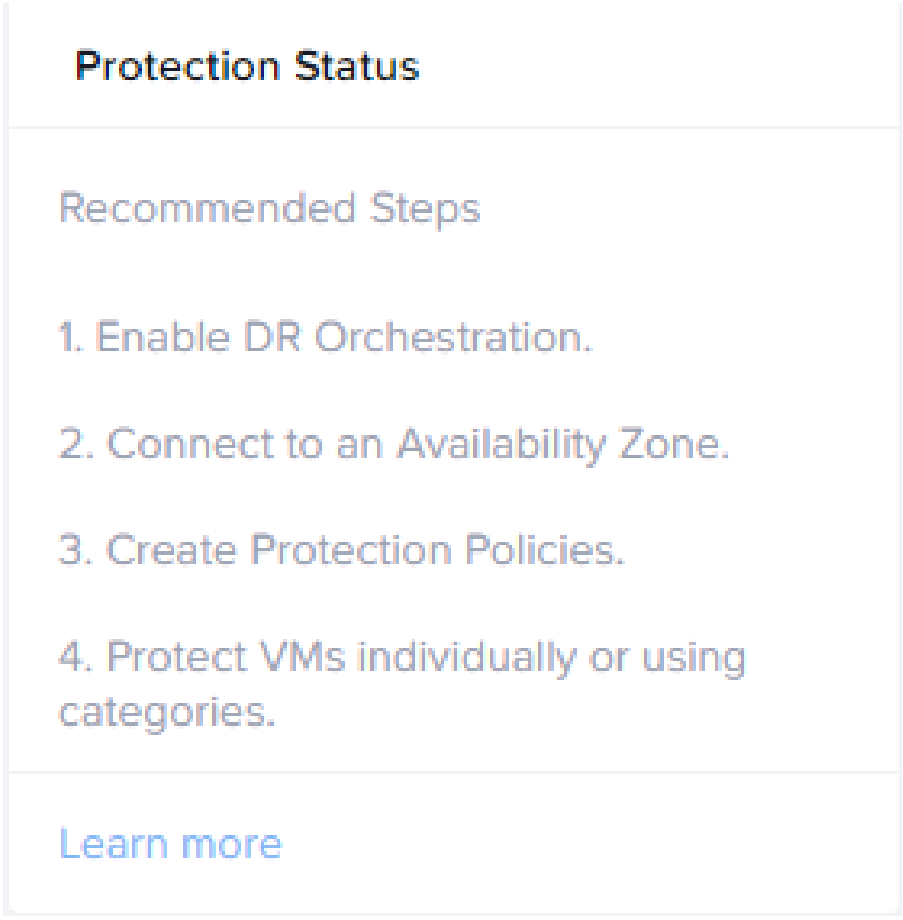

Figure 42: Protection Status Widget

- **Reports**. This widget displays a table that lists the number of total and scheduled reports with a link to the reports dashboard (see Reports View on page 159).

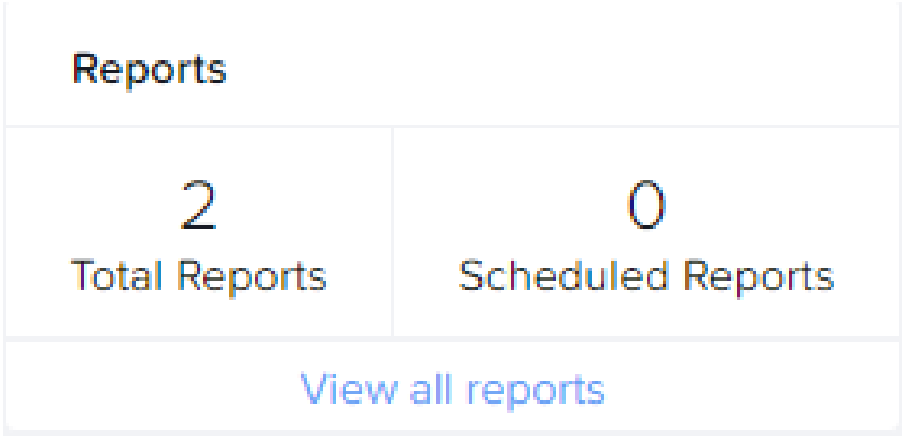

Figure 43: Reports Widget

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 new dashboards. To create a new 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 47), click the **Manage Dashboards** button.
2. In the Manage Dashboards window (see Modifying a Dashboard on page 59), 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.

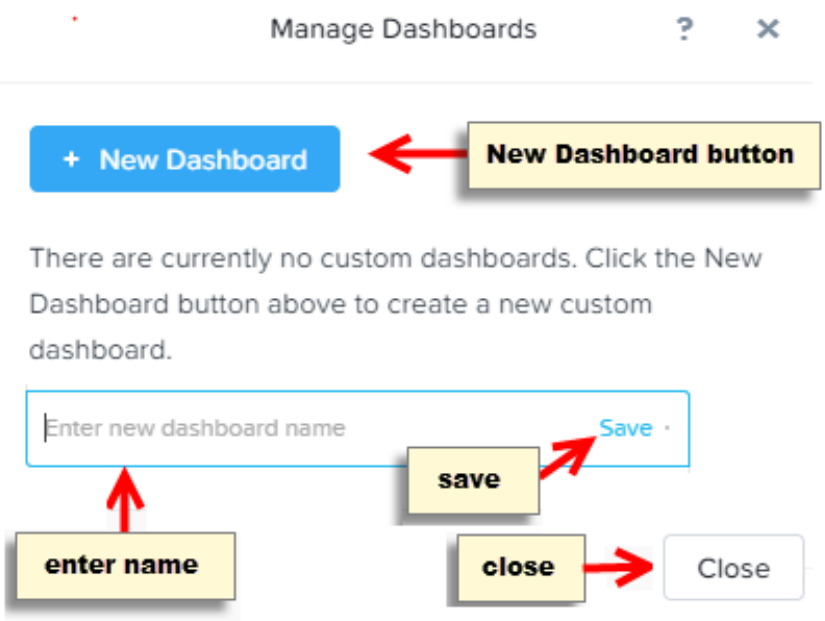

Figure 44: Manage Dashboards Windows (new dashboard)

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.

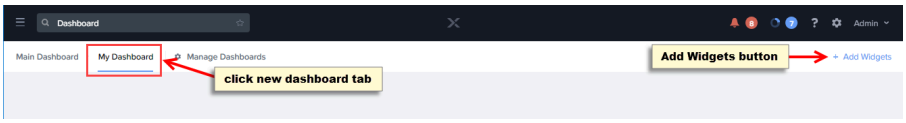

Figure 45: Custom Dashboard Display

5. In the Add Widgets screen, select a widget and add it to the dashboard (see Adding Dashboard Widgets on page 61).

   Repeat this step until all the desired widgets have been added to the dashboard.

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 47) 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 61).

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.

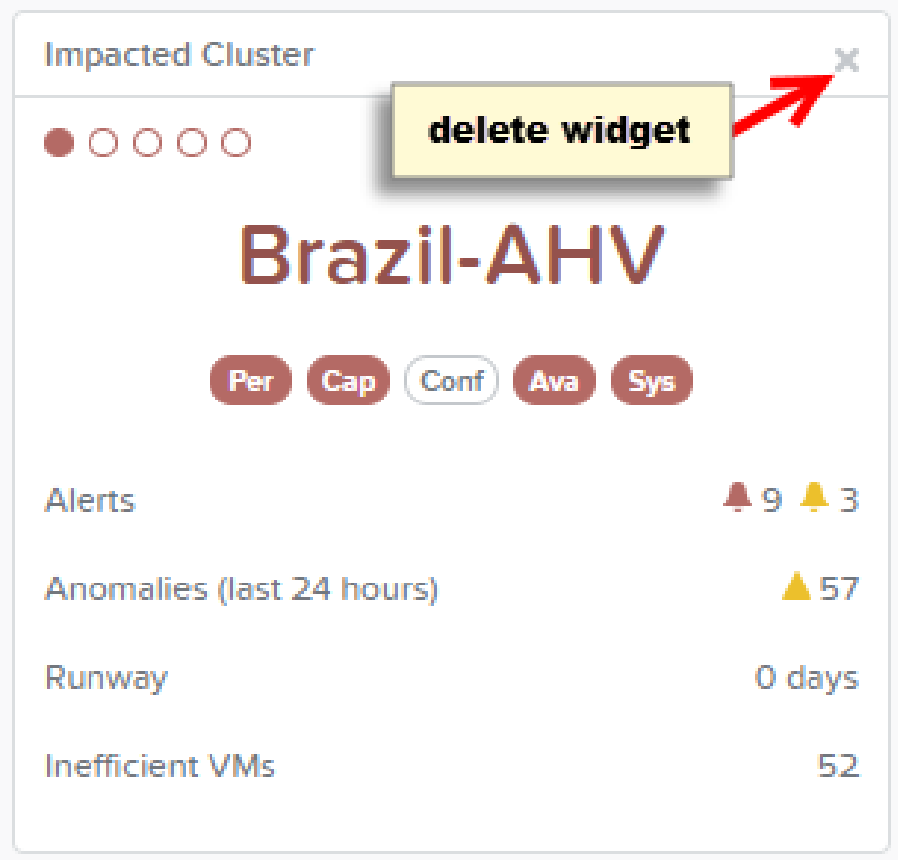

Figure 46: Deleting a Widget

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.

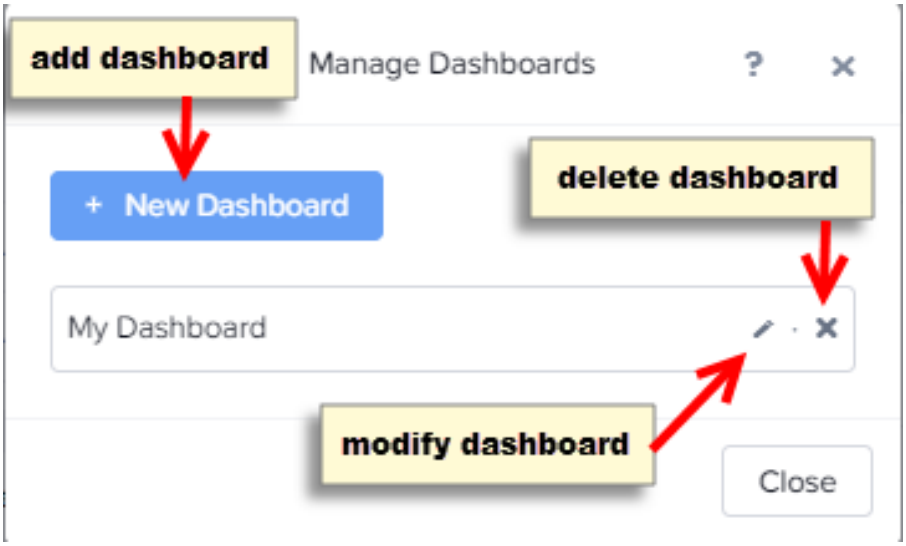

Figure 47: Manage Dashboards Window

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.

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** screen, 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.

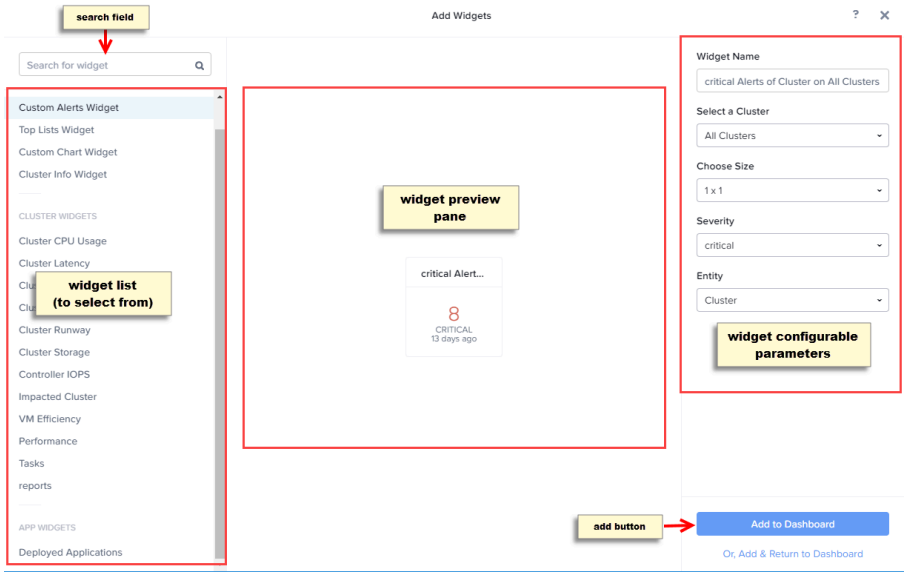

Figure 48: Add Widget Screen

The following table describes the widgets that you can add to a dashboard.

Table 7: 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
		Entity. Select the entity of interest. The default is cluster.	Cluster, Storage, VM, Hardware, DR

Widget	Description	Parameter	Values
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)
Cluster Widgets - All these widgets (except Performance) appear on the home dashboard by default (see Main Dashboard on page 47).			

Widget	Description	Parameter	Values
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 Tab" section in Cluster Details View on page 121 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
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. See Behavioral Learning Tools on page 197 for more information about VM efficiency.	(no customizable parameters)	(n/a)
Performance	Displays latency, bandwidth, and IOPS statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Tasks	Displays a list of recent tasks with the current status of each task and a link to the Tasks dashboard (see Tasks View on page 157.	(no customizable parameters)	(n/a)
Reports	Displays a table that lists the number of total and scheduled reports with a link to the reports dashboard (see Reports View on page 159).	(no customizable parameters)	(n/a)
App Widgets			
Deployed Applications	Displays a list of deployed applications. See the *Nutanix Calm Administration and Operations Guide* for more information.	(no customizable parameters)	(n/a)

4

ENTITY EXPLORING

The entities menu (see Entities Menu on page 12) 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 71 (VMs, storage containers, catalog items, images, categories, recoverable entities, and subnets)
- Policies Entities on page 106 (security policies, protection policies, recovery plans, NGT policies, and image placement policies)
- Hardware Entities on page 114 (clusters, hosts, disks, and GPUs)
- Activity Entities on page 153 (alerts, events, audits, and tasks)
- Operations Entities on page 159 (analysis, planning, reports, playbooks, action gallery, and plays)
- Administration Entities on page 161 (LCM, projects, roles, users, and availability zones)
- Services Entities on page 175 (Buckets, Calm, and Karbon)

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.

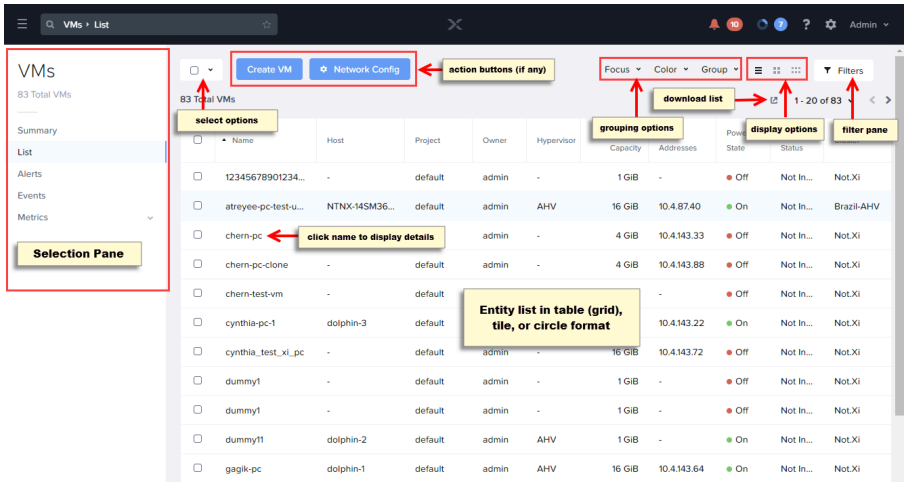

Figure 49: Dashboard Layout (example)

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 18) 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`".

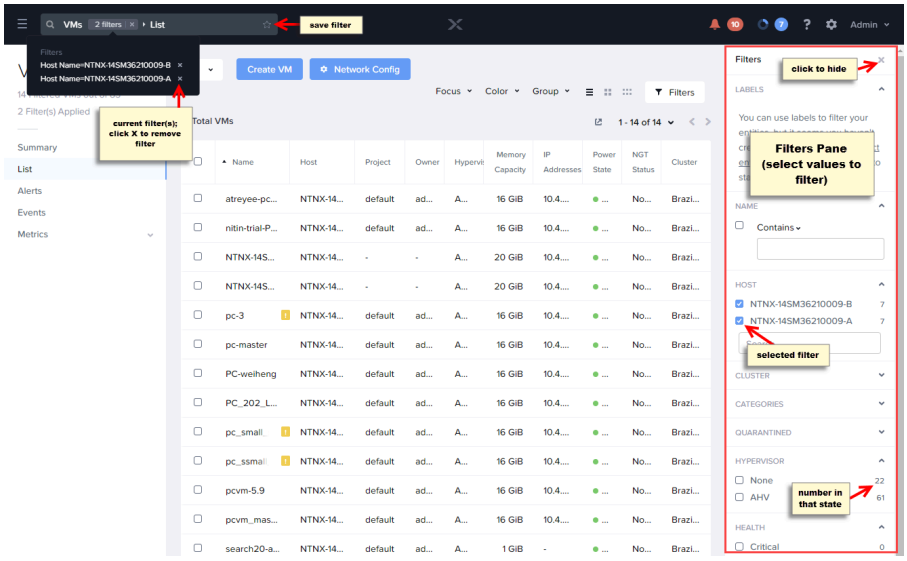

Figure 50: Filter Pane

Display Options

You can display an entity list in one of three formats by clicking the appropriate icon (grid, tiles, or circles) in the menu bar:

- The Grid (table) view appears by default. This format displays the entity list in a tabular format. Each row represents a single entity and includes basic information about that entity. The columns vary by entity type and the selected focus.
  - Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
  - Click a name to display the details screen for that entity.
- The Tiles view displays the entity list as a set of tiles. Each tile includes the same information displayed in tabular format.

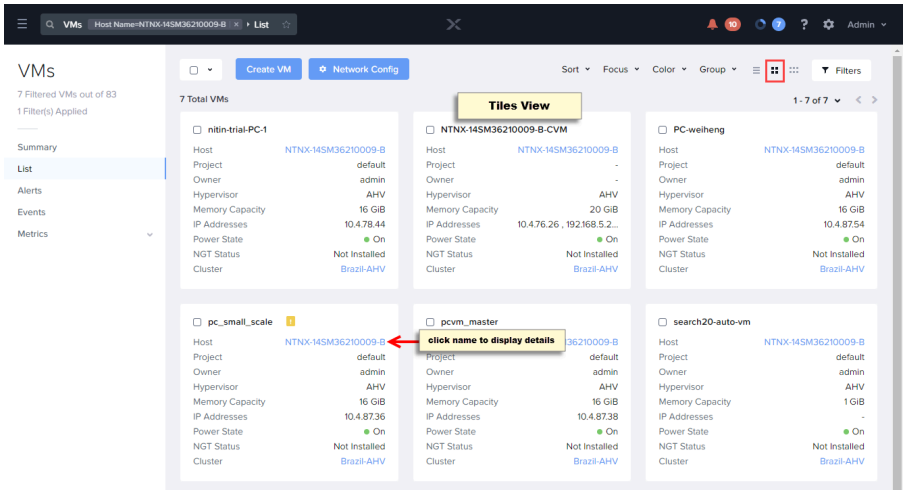

Figure 51: Tiles View

- The Circles view displays the entity list as a set of small circles. Rolling the cursor over a circle displays the entity name.

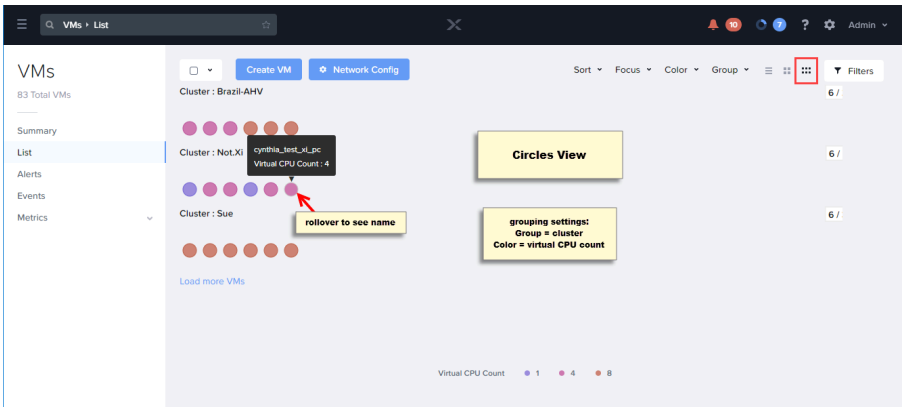

Figure 52: Circles View

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 **Focus** 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 197).
  - 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 **Focus** 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.

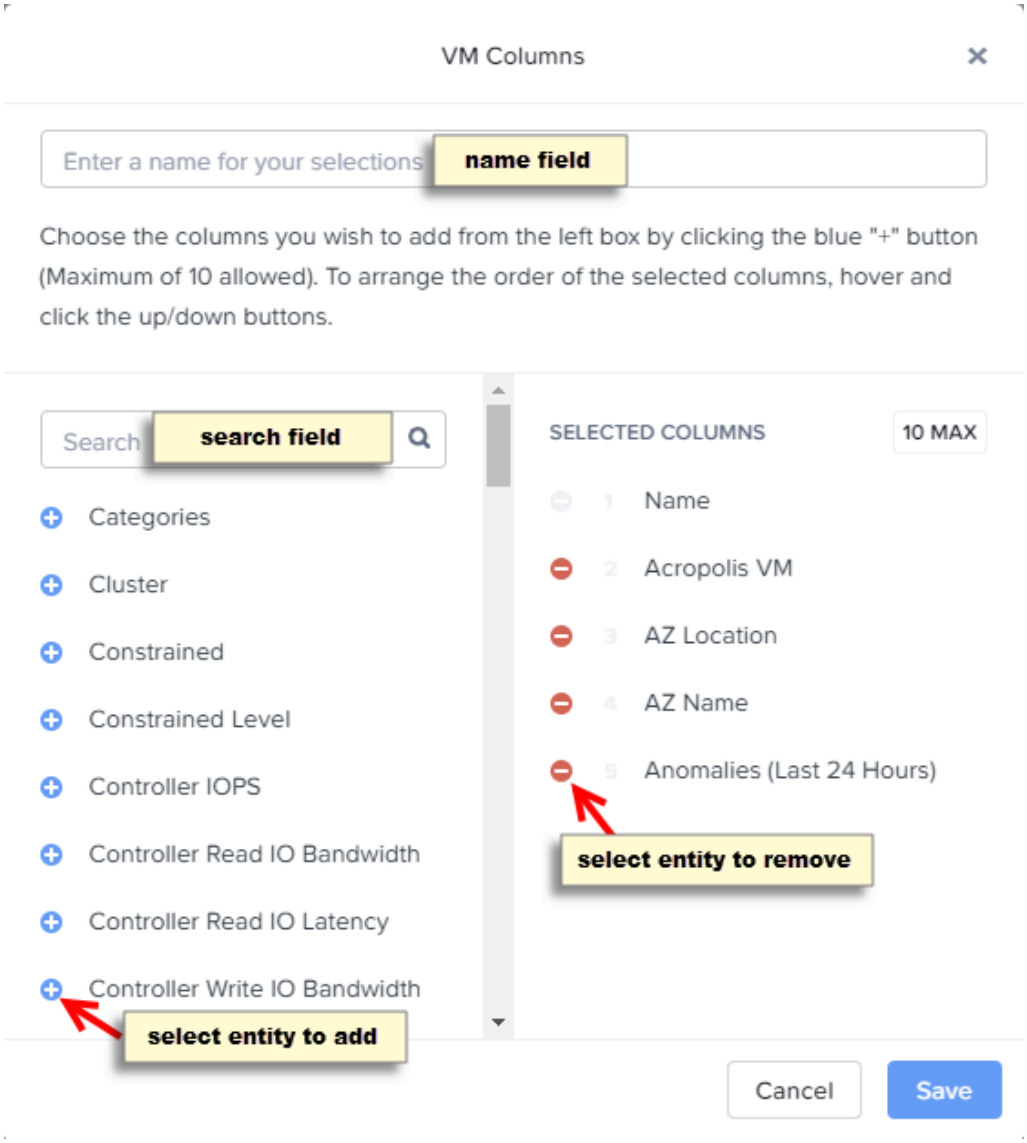

Figure 53: VM Column Window (example)

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.

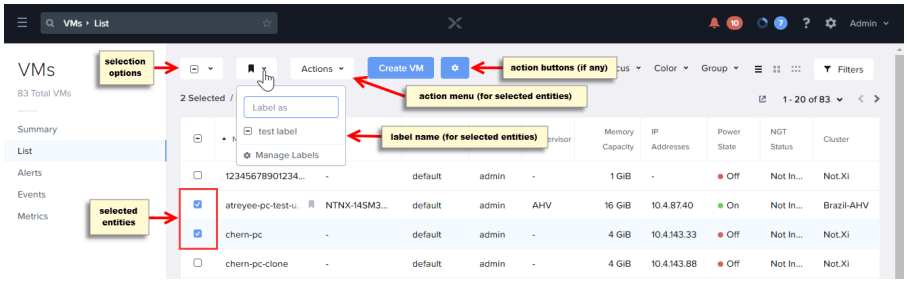

Figure 54: Action Options

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 12):

- VMs (see VMs Summary View on page 72)

- Storage Containers (see Storage Containers Summary View on page 87)
- Catalog Items (see Catalog Items View on page 97)
- Images (see Images Summary View on page 98)
- Categories (see Categories Summary View on page 102)
- Recoverable Entities (see Recoverable Entities View on page 104)
- Subnets (see Subnets View on page 105)

> 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 12). 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 to the right for the selected tab.

> Note: This section describes the information and options that appear in the VMs dashboard.
>
> - See Entity Exploring on page 66 for instructions on how to view and organize that information in a variety of ways.
> - See VM Management on page 296 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 47). See Behavioral Learning Tools on page 197 for more information about VM efficiency.

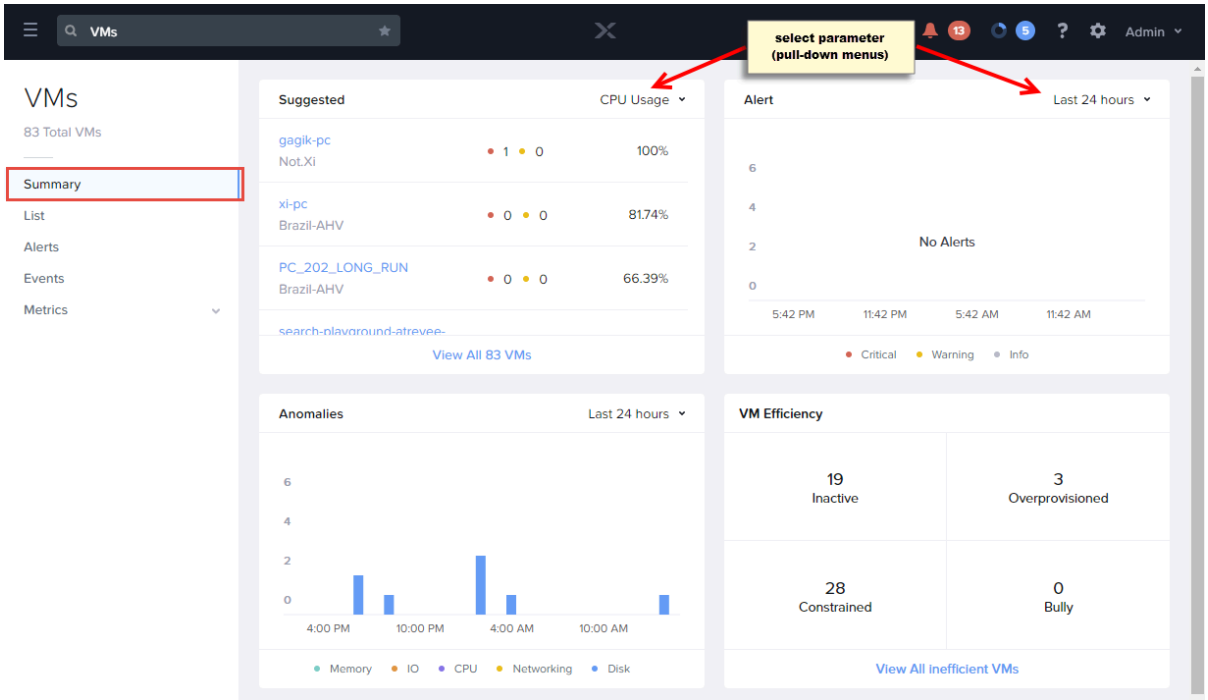

Figure 55: VMs Summary Tab

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.

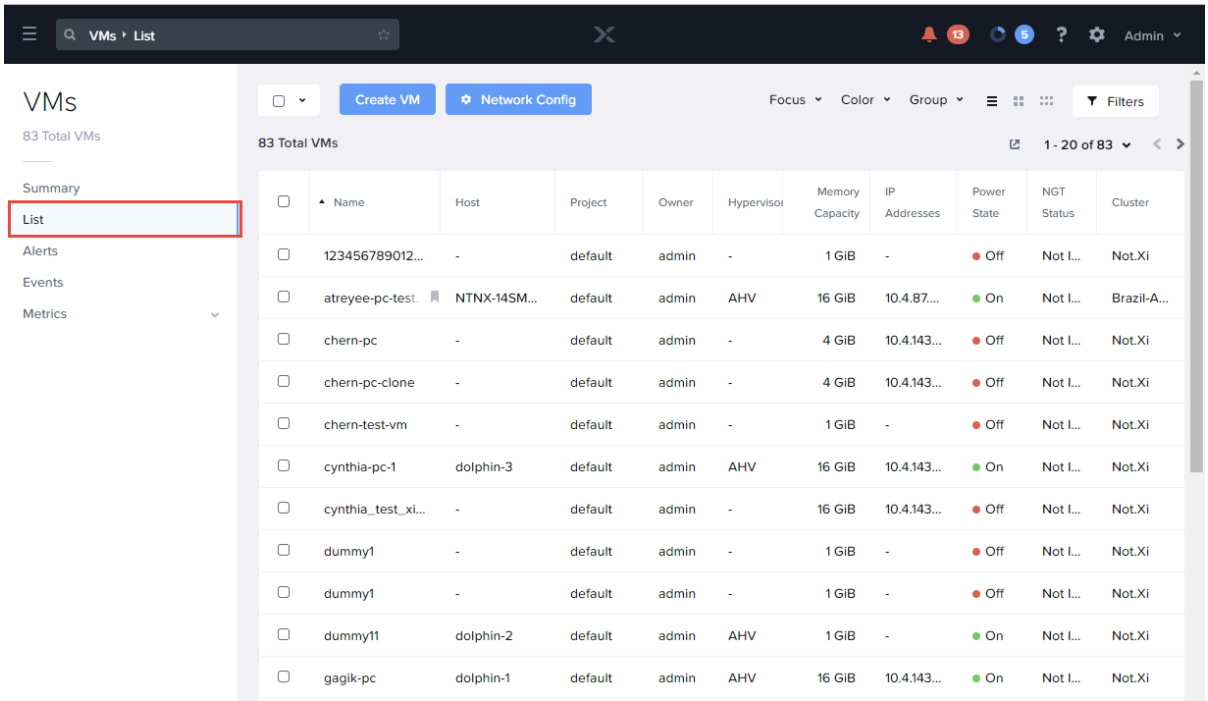

Figure 56: VMs List Tab

Table 8: 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 81).	(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)
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)

Parameter	Description	Values
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 197).	Good, Constrained, Overprovisioned, Inactive, Bully
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 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 9: 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)

Parameter	Description	Values
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)
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 197).	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)

You can group the VM list in the following ways:

- The **Color** pull-down menu allows you to color code the VM entries by virtual CPU count, power state, 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 VM entries by cluster, hypervisor, power state, virtual CPU count, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** pull-down menu allows you to group the VM entries by the information parameters (fields), which vary depending on the focus selected. (You can only choose one parameter.)
- 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:

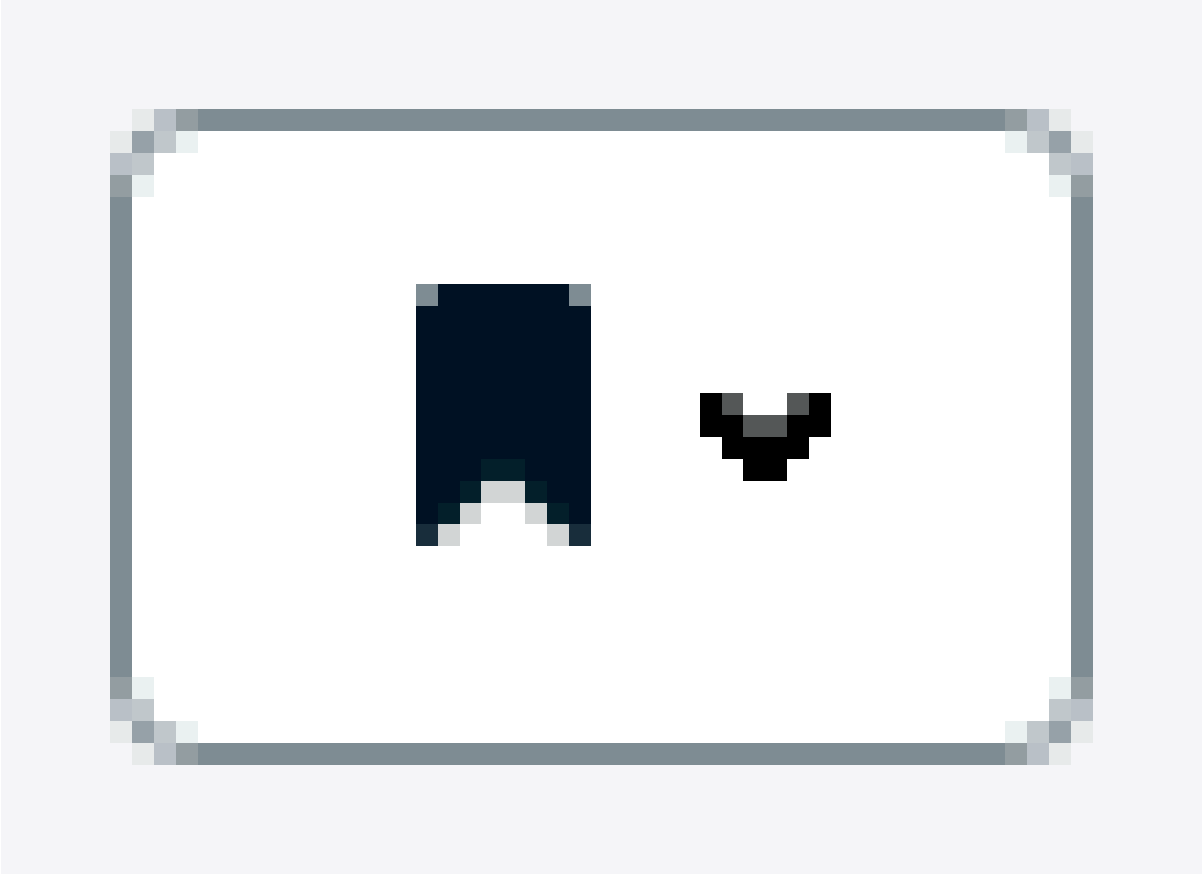

Figure 57: Label Icon

  - 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 297 or Creating a VM (ESXi) on page 318).
- To configure the network, click the **Network Config** button (see Configuring Network Connections on page 357).

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 and Self Service) on page 310 or Managing a VM (ESXi) on page 321 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 214).

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 220).

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.)

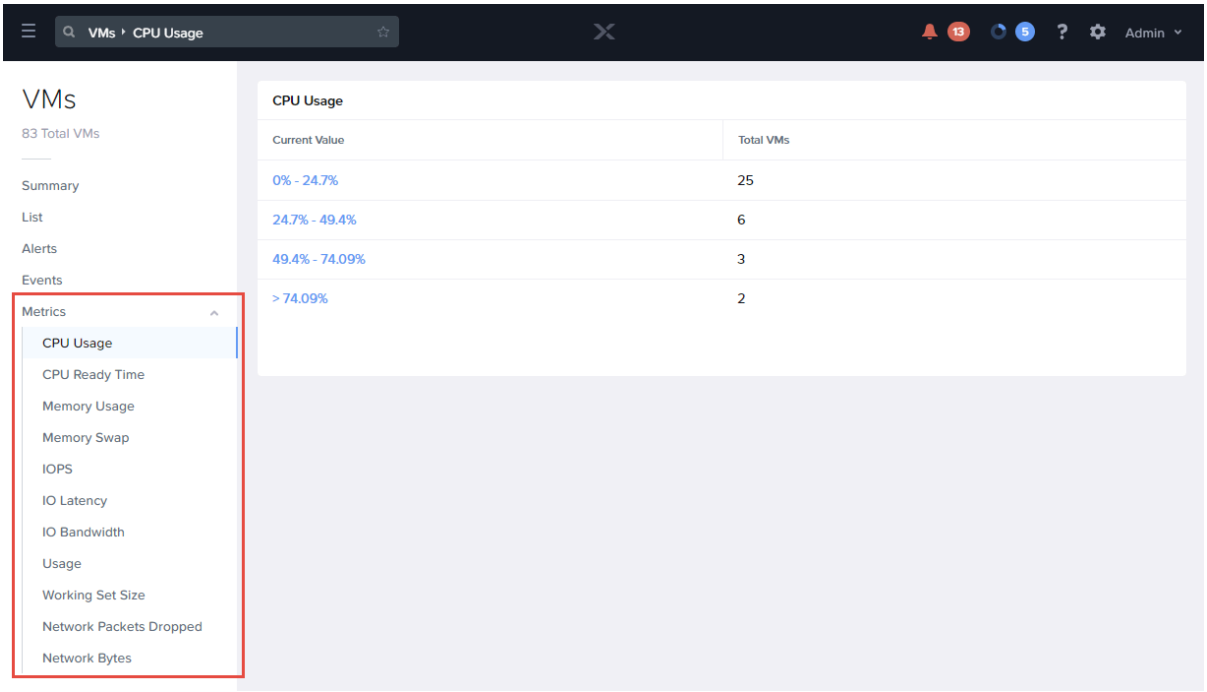

Figure 58: VMs Metrics Tab

Table 10: 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.
Network Packets Dropped	Displays tables for the number of transmitted and received packets dropped.

Metric	Description
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 72) and click the VM name. You can also access the details page by clicking the VM name wherever that name appear, 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.)

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 and Self Service) on page 310 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.

- 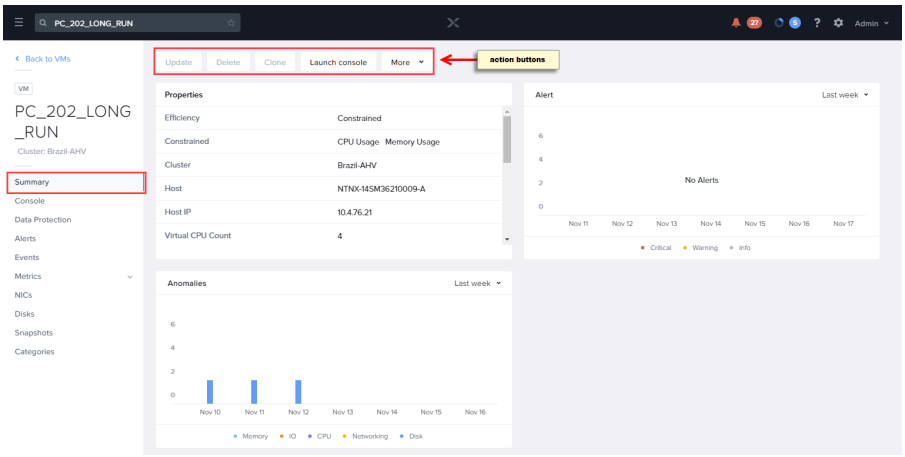

Figure 59: VM Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 11: 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 197 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 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)

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)

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.

Data Protection

The **Data Protection** 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.

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 214).

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 220).

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 197). 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 229).

The following table describes the available metrics. (Some of these metrics are not available on all hypervisors.)

Table 12: 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).

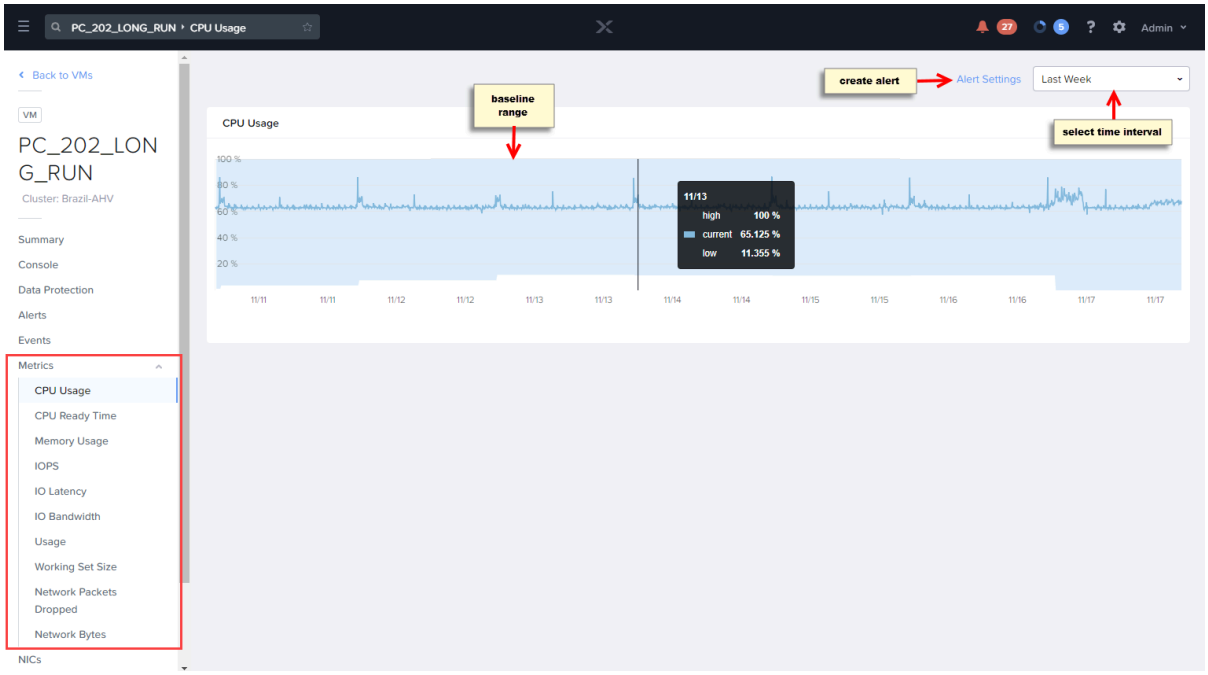

Figure 60: Metrics Tab: CPU Usage

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 13: 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 represent 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).

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:

- **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.

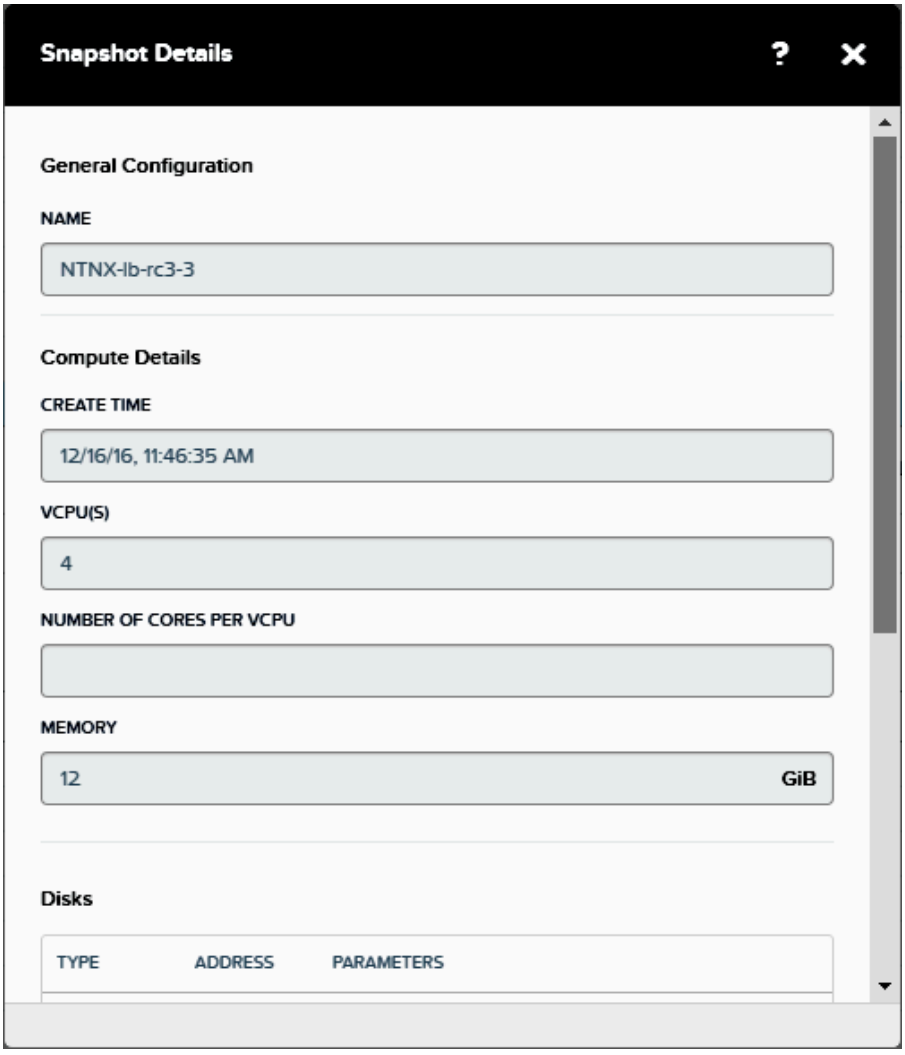

Figure 61: Snapshot Details Window

  - 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.

  See Managing a VM (AHV and Self Service) on page 310 or Managing a VM (ESXi) on page 321 for more information about these actions.

Categories Tab

Storage Containers Summary View

To access the storage containers dashboard, select **Virtual Infrastructure** > **Storage Containers** from the entities menu (see Entities Menu on page 12). 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 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 storage containers dashboard. See Entity Exploring on page 66 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 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 **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 on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

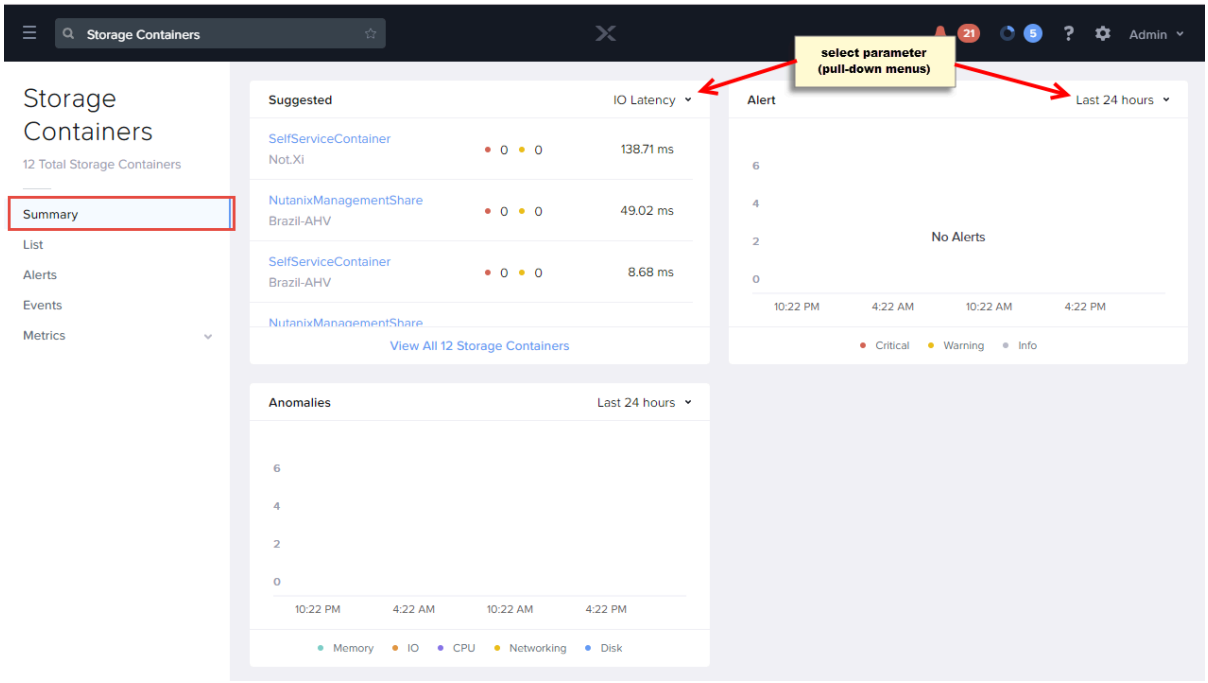

Figure 62: Storage Containers Summary Tab

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.

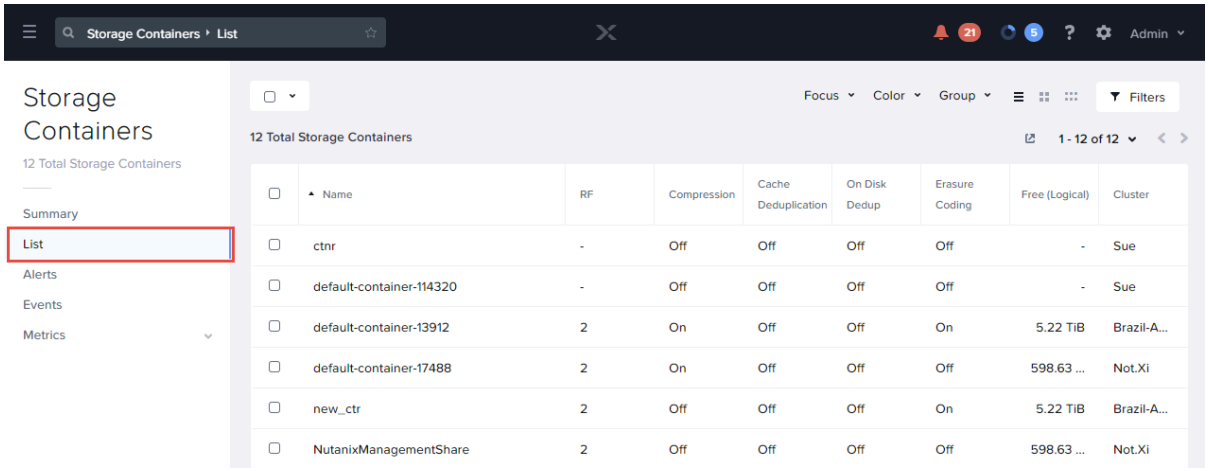

Figure 63: Storage Containers List Tab

Table 14: Storage Containers List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the name of the storage container.	(name)
RF	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]
On Disk Dedup	Displays whether on disk deduplication is enabled, that is dedup compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for on disk deduplication.	[On, Off]

Parameter	Description	Values
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Free (logical)	Displays the amount of free storage space available to the storage container.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the storage container resides.	(cluster name)
"Performance" Focus Fields		
Name	Displays the name of the storage container.	(name)
Free (logical)	Displays the amount of free storage space available to the storage container.	xxx [GB|TB]
Storage Logical Usage	Displays the amount of used storage space in the storage container.	xxx [GB|TB]
Max Capacity	Displays the total amount of storage capacity available to 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]
Cluster	Displays the name of the cluster in which the storage container resides.	(cluster name)

You can filter the storage containers list based on a variety of 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 additional context such as a constraint string or a range.

Table 15: 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.	(storage container name string)

Parameter	Description	Values
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 will return a list of storage containers that reside in the clusters which satisfy the name condition/string.	(cluster name string)
Compression	Filters on the compression setting. Select one or more settings to return a list of storage containers in that setting(s). The number of storage containers currently in each setting is displayed on the right of the line.	Low, Medium, High
Cache Deduplication	Filters on the cache deduplication setting. Select one or more settings to return a list of storage containers in that setting(s). The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
RF	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
On Disk Dedup	Filters on the on disk deduplication setting. Select one or more settings to return a list of storage containers in that setting(s). 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 in that setting(s). The number of storage containers currently in each setting is displayed on the right of the line.	On, Off
Free (logical)	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 will return 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 that state(s). The number of storage containers currently in each state is displayed on the right of the line.	Critical, Warning, Good

Parameter	Description	Values
Storage Logical Usage	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 will return a list of storage containers with used storage in that range.	([xx] to [yy] GiB range)
Max Capacity	Filters on the maximum 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 will return 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 will return 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 **Color** pull-down menu allows you to color code the storage container entries by replication factor, compression, erasure coding, or health state. 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 host entries by replication factor, compression, erasure coding, cluster, or health setting. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** 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.)

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 214).

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 220).

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.

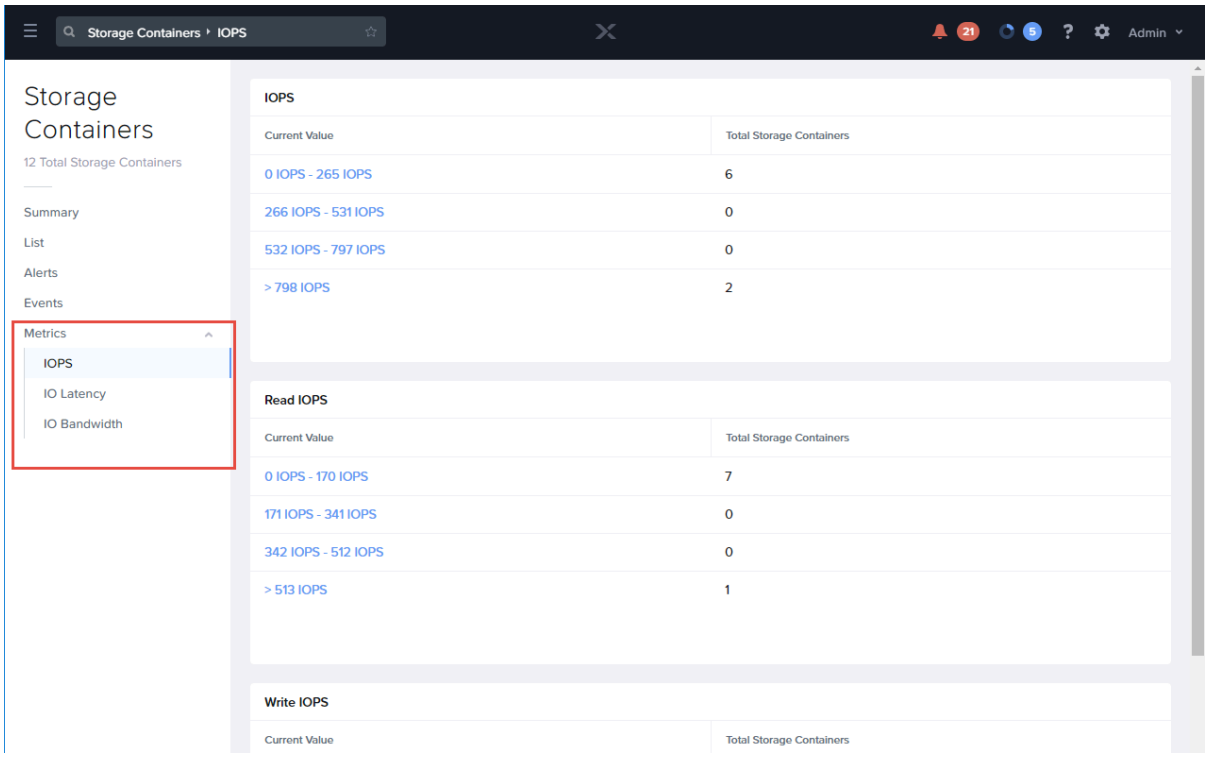

Figure 64: Storage Containers Metrics Tab

Table 16: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS tables 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). Clicking on an interval displays the Summary tab filtered to just those containers. 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.

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 87) and click the storage container name. You can also access the details page by clicking the storage container name wherever that name appear, such as in a dashboard widget or search result.

The storage container 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 Storage Containers** 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 storage container (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.

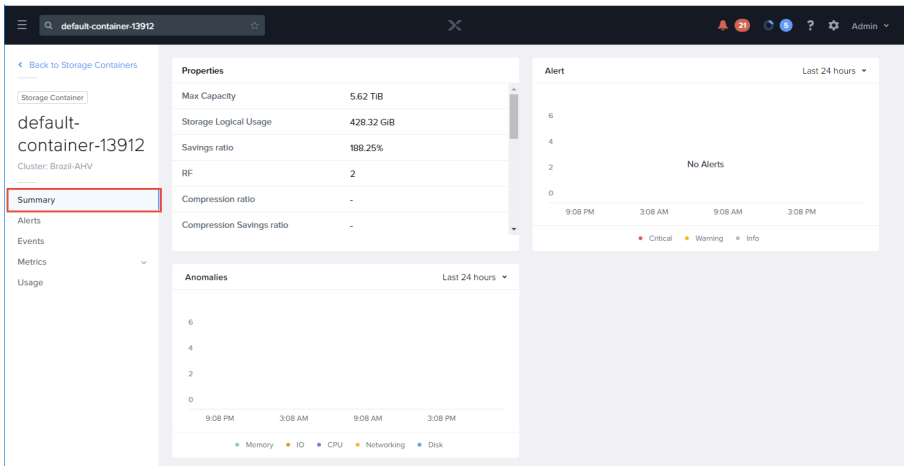

Figure 65: Storage Container Summary Tab

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 17: Storage Container Properties Fields

Parameter	Description	Values
Max Capacity	Displays the total amount of storage capacity available to the storage container (see Reserved Capacity).	xxx [TB]
Storage Logical Usage	Displays the amount of used storage space in the storage container.	xxx [GB|TB]
Saving Ratio		

Parameter	Description	Values
RF	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (normally 2 or 3) when the storage container is created.	1, 2, 3
Compression Ratio		
Compression Savings Ratio		
Storage Savings		
Free (Logical)	Displays the amount of free storage space available to the storage container.	xxx [GB|TB]
Reserved Capacity	Displays the total reserved storage capacity in the storage container. Nutanix employs a "thin" provisioning model when allocating storage space, which means space is assigned to a storage container only when it is actually needed. The maximum capacity value reflects total available storage regardless of how many storage containers are defined. Therefore, when you have two storage containers, it can appear you have twice as much capacity because the field values for both storage containers show the full amount. However, capacity can be reserved for a specific storage container, and this field displays how much storage (if any) has been reserved for this storage container.	xxx [GB|TB]
Free (Physical)		
Overall Savings Ratio		
Cluster	Displays the name of the cluster in which the storage container resides.	(cluster name)
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
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]
Compression Delay	Displays the delay (number of minutes) before data changes are compressed. A zero value indicates compression is immediate (not delayed).	xx min

Parameter	Description	Values
Thick Provisioned	Displays the reserved storage capacity for thick provisioned VMs.	xxx [GB|TB]
Datastore	Displays the name of the data store in which this storage container is located,	(datastore name)
VMs	Displays the number of VMs in this storage container.	(number)
Data Reduction Ratio	Displays how much data size has been reduced (expressed as a ratio) because of the data reduction methods employed.	(xx:yy)
Data Reduction Savings	Displays the amount of storage space saved because of the data reduction methods employed.	xxx [GB|TB]
Effective Free	Displays the amount of effective free space available for use.	xxx [GB|TB]
Overall Efficiency	Displays the overall efficiency (expressed as a ratio) because of the data reduction methods employed.	(xx:yy)
Compression Space Saved	Displays the amount of storage space saved because of compression.	xxx [GB|TB]
On Disk Deduplication	Displays whether on disk deduplication is enabled, that is dedup compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for on disk deduplication.	[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 214).

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 220).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the storage container. 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 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 197). 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 229).

Figure 66: Storage Container Metrics Tab

Usage Tab

The **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).

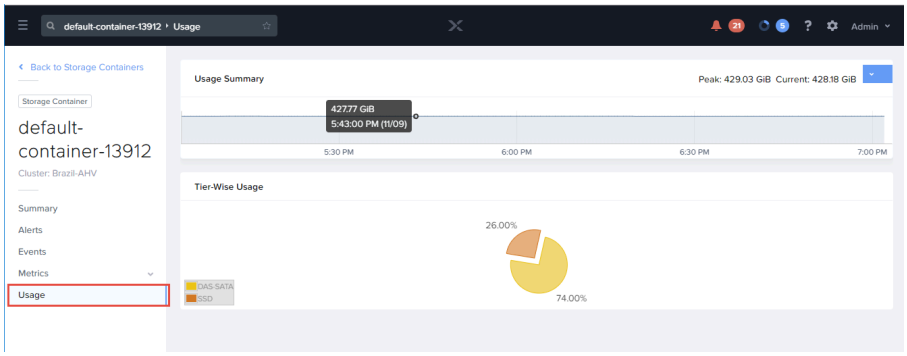

Figure 67: Storage Container Usage Tab

Catalog Items View

To access the catalog items dashboard, select **Virtual Infrastructure** > **Catalog Items** from the entities menu (see Entities Menu on page 12). 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 66 for instructions on how to view and organize that information in a variety of ways.
- See Catalog Management on page 347 for information about the catalog service and how to add (or delete) items.

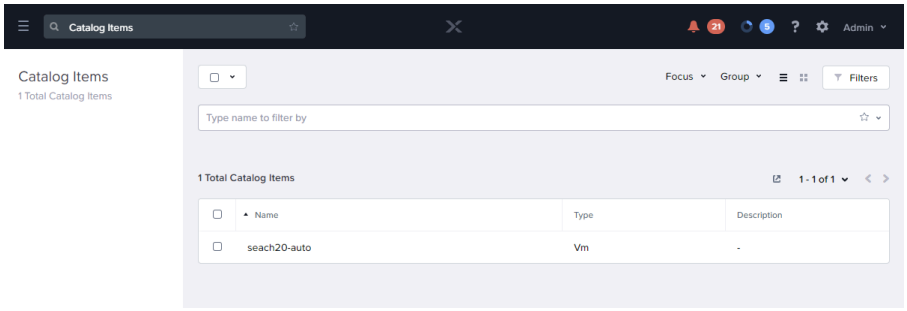

Figure 68: Catalog Items Dashboard

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 18: 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 19: 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).
- [Tiles view only; there is no Circles view] The **Sort** pull-down menu allows you to group the entries by name, type, or description. (You can only choose one parameter.)

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 350).

Images Summary View

To access the images dashboard, select **Virtual Infrastructure** > **Images** from the entities menu (see Entities Menu on page 12). 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 66 for instructions on how to view and organize that information in a variety of ways.
> - See Image Management on page 334 for information about adding and managing images through Prism Central.

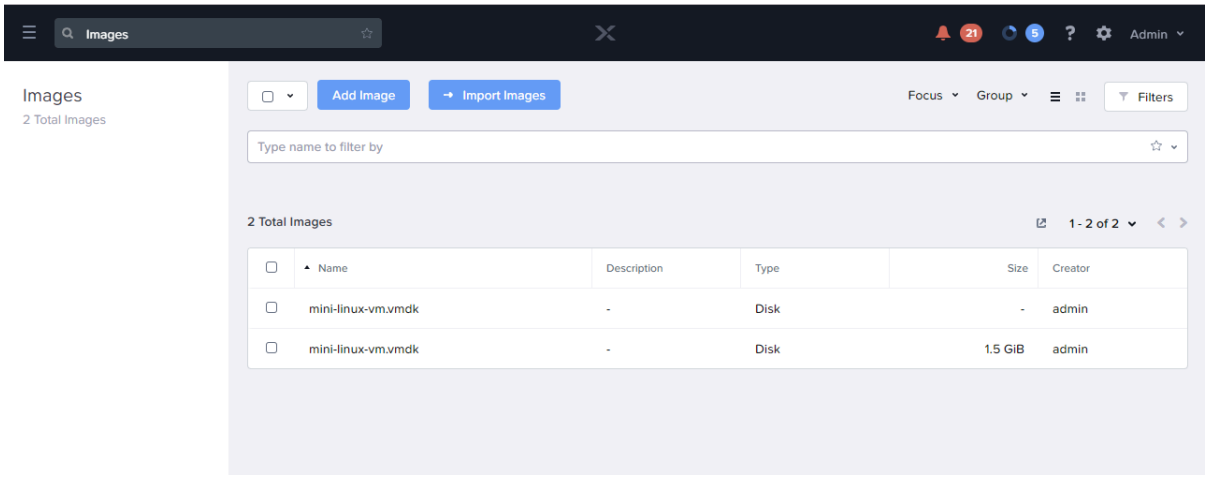

Figure 69: Images View

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 20: 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 100).	(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 21: 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)

Parameter	Description	Values
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).
- [Tiles view only; there is no Circles view] The **Sort** pull-down menu allows you to group the entries by name, description, type, or size. (You can only choose one parameter.)

The images dashboard includes the following action buttons:

- **Add Image** Click to add an image (see Adding an Image on page 336).
- **Import Images** Click to import images from registered clusters to Prism Central (see Importing Images to Prism Central on page 347.

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 346). 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 98) 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 Adding an Image on page 336).
- 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 346).
  - To add the image to the catalog, click the **Add Image to Catalog** button (see Adding a Catalog Item on page 347).
  - To assign the image to a category, click the **Manage Categories** button (see Assigning a Category on page 353).
  - 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.

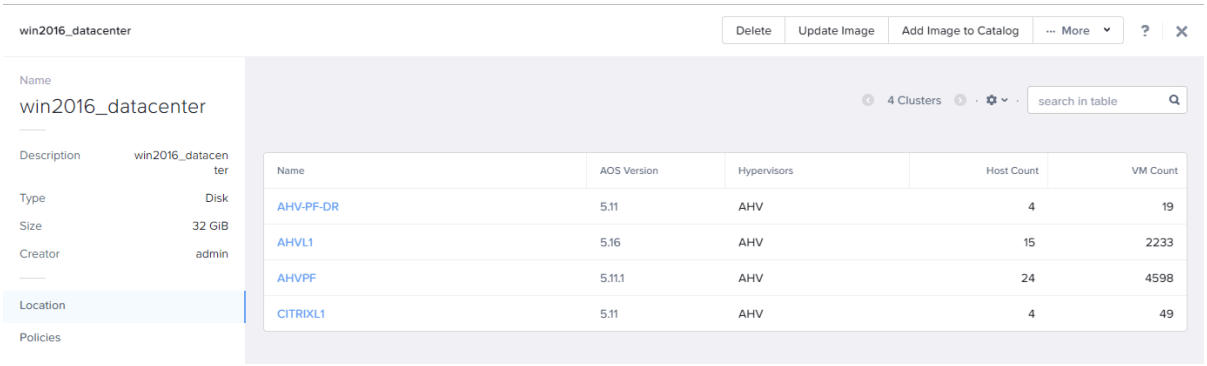

Figure 70: Image Details View (Location Tab)

Location Tab

The following table describes the information displayed in the **Location** tab, which appears by default when you first open the page.

Table 22: 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 379). The following table describes the information displayed in the **Policies** tab.

Table 23: 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)

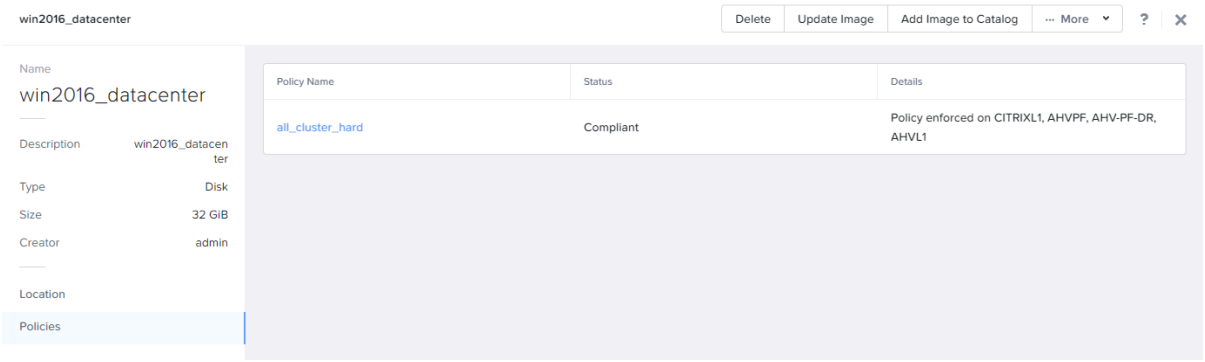

Figure 71: Image Details View (Policies Tab)

Categories Summary View

To access the categories dashboard, select **Virtual Infrastructure** > **Categories** from the entities menu (see Entities Menu on page 12). 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 66 for instructions on how to view and organize that information in a variety of ways.
- See Category Management on page 351 for information on how to create, modify, and apply categories.

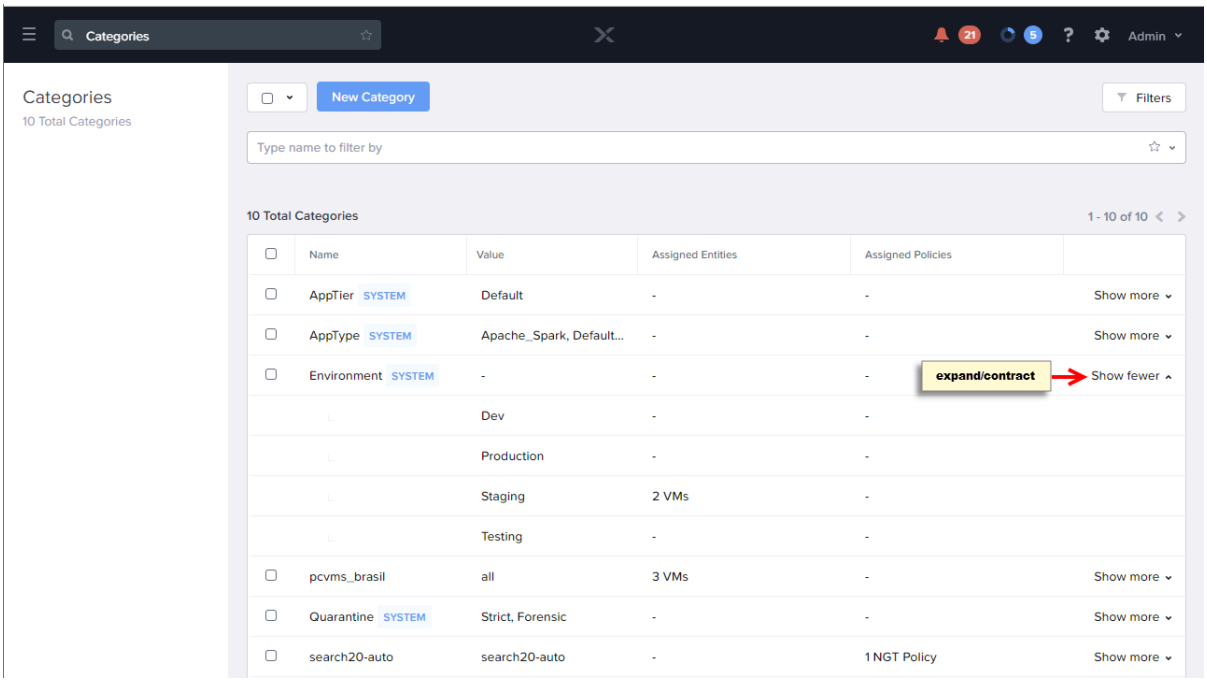

Figure 72: Categories Dashboard

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 24: 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 25: 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)
Value	Filters on the category value. 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 value condition/ string.	(value string)
Entities	Filters on the entity type. Check the box for one or more entity types.	VMs, Hosts
Policies	Filters on the policy type. Check the box for one or more entity types.	Security Policies, Affinity Policies

There is a **New Category** action button to create a new category (see Creating a Category on page 351). 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 352 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 102) 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 352); 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.

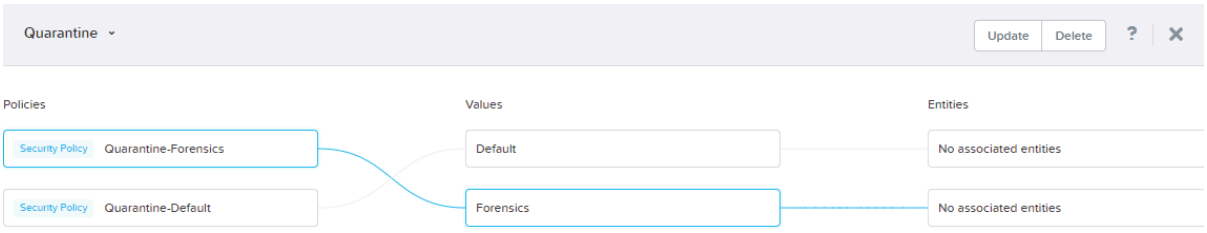

Figure 73: Category Details View

Recoverable Entities View

To access the recoverable entities dashboard, select **Virtual Infrastructure** > **Recoverable Entities** from the entities menu (see Entities Menu on page 12). 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 66 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.

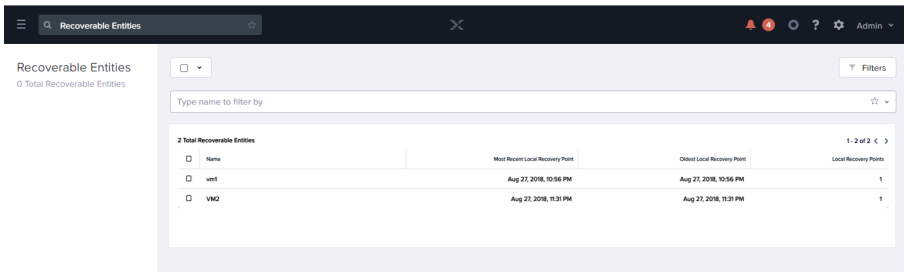

Figure 74: Recoverable Entities Dashboard

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 26: Recoverable Entities List Fields

Parameter	Description	Values
Name	Displays the recoverable entity name.	(item name)
Most Recent Local Recovery Point	Displays the time at which the most recent snapshot was taken. This is specific to the local availability zone.	(time and date)
Oldest Local Recovery Point	Displays the time at which the oldest available snapshot was taken. This is specific to the local availability zone.	(time and date)
Local Recovery Points.	Displays the number of recovery points available locally.	(integer)

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:

- **Clone**: Recover an entity from a snapshot.
- **Replicate**: Replicate a snapshot to a paired availability zone.
- **Delete**: Delete all the recovery points of a recoverable entity.

Subnets View

To access the subnets dashboard, select **Virtual Infrastructure** > **Subnets** from the entities menu (see Entities Menu on page 12). 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 subnets dashboard. See Entity Exploring on page 66 for instructions on how to view and organize that information in a variety of ways.

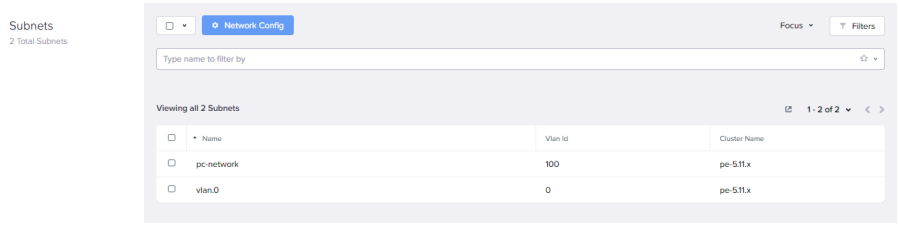

Figure 75: Subnets Dashboard

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 27: Recoverable Entities List Fields

Parameter	Description	Values
Name	Displays the subnet name.	(subnet name)
Vlan Id	Displays the VLAN identification number.	(ID number)
Cluster Name	Displays the name of the cluster for which this subnet is configured.	(cluster name)

To filter the list by subnet name, enter a string in the filter field. (Ignore the **Filters** pane as it is blank.)

There is a **Network Config** action button to configure a new network (see Configuring Network Connections on page 357). The **Actions** menu appears when one or more subnets are selected and includes a **Manage Categories** option (see Assigning a Category on page 353).

Policies Entities

You can access dashboards for the following policy types from the **Policies** category of the entities menu (see Entities Menu on page 12):

- Security Policies (see Security Policies Summary View on page 106)
- Protection Policies (see Protection Policies View on page 109)
- Recovery Plans (see Recovery Plans View on page 110)
- NGT Policies (see NGT Policies View on page 111)
- Image Placement Policies (see Image Placement Policies on page 379)

Security Policies Summary View

To access the security policies dashboard, select **Policies** > **Security Policies** from the entities menu (see Entities Menu on page 12). 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 66 for instructions on how to view and organize that information in a variety of ways.
- See “” () for information about how to create and apply security policies.

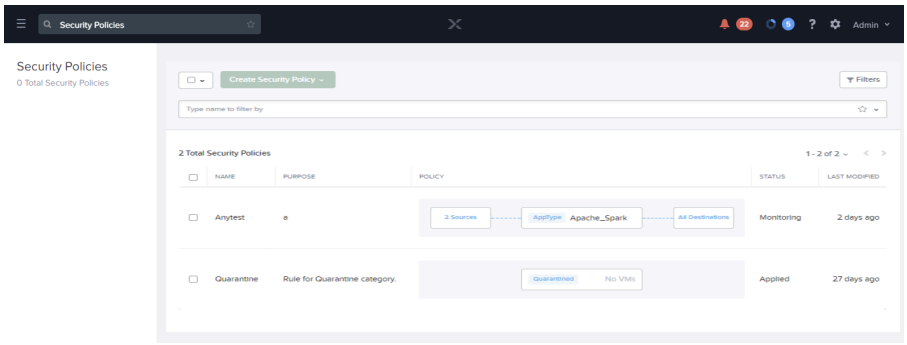

Figure 76: Security Policies Dashboard

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 28: 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 29: 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 106). 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.

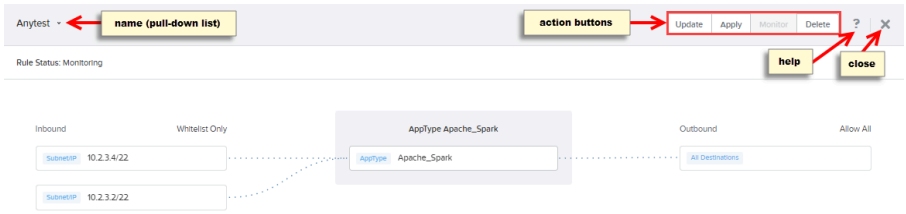

Figure 77: Security Policy Details View: Monitoring Rule Example

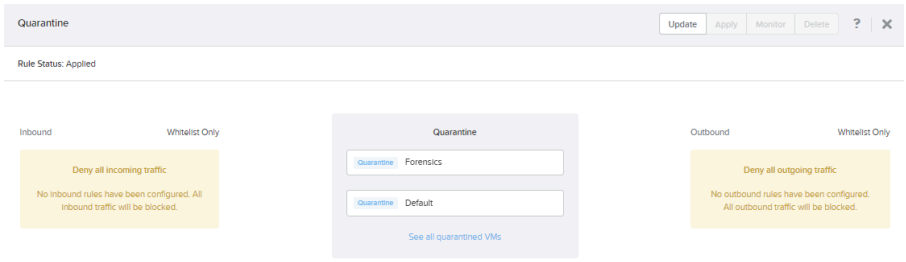

Figure 78: Security Policy Details View: Applied Rule Example

Protection Policies View

To access the protection policies dashboard, select **Policies** > **Protection Policies** from the entities menu (see Entities Menu on page 12). 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 66 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.

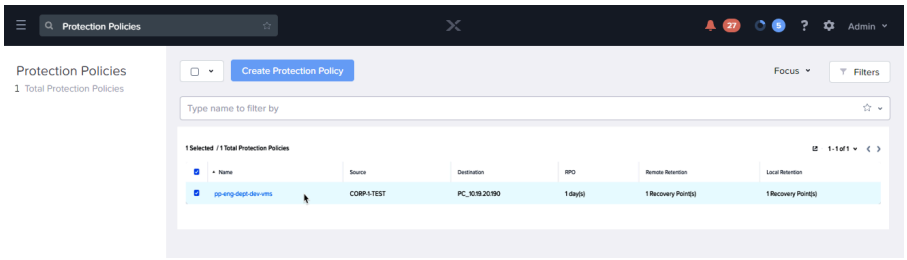

Figure 79: Protection Policies Dashboard

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 30: 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 12). 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 66 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.

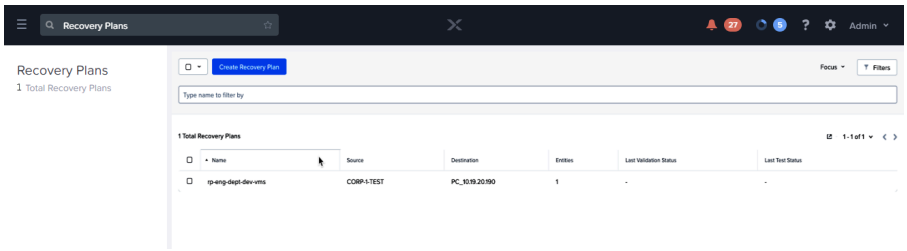

Figure 80: Recovery Plans Dashboard

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 31: 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)
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:

- **Update**: Update the recovery plan.
- **Validate**: Clone a recovery plan.
- **Test**: Test the recovery plan.
- **Failover**: Perform a failover operation.
- **Delete**: Delete the recovery plan.

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 12). 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 66 for instructions on how to view and organize that information in a variety of ways.
> - See NGT Policies on page 377 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.

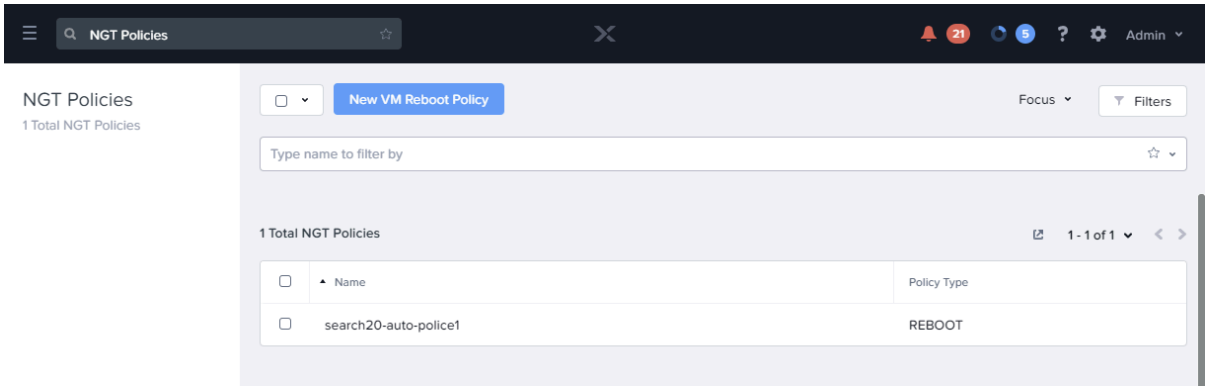

Figure 81: NGT Policies Dashboard

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 377). The **Actions** menu appears when one or more policies are selected and includes the following options:

- **Update** (see Modifying a Restart Policy on page 379)
- **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 12).

Note: This section describes the information and options that appear in the image placement policies dashboard.

- See Entity Exploring on page 66 for instructions on how to view and organize that information in a variety of ways.
- See Image Placement Policies on page 379 for information about how to create and apply image placement policies.

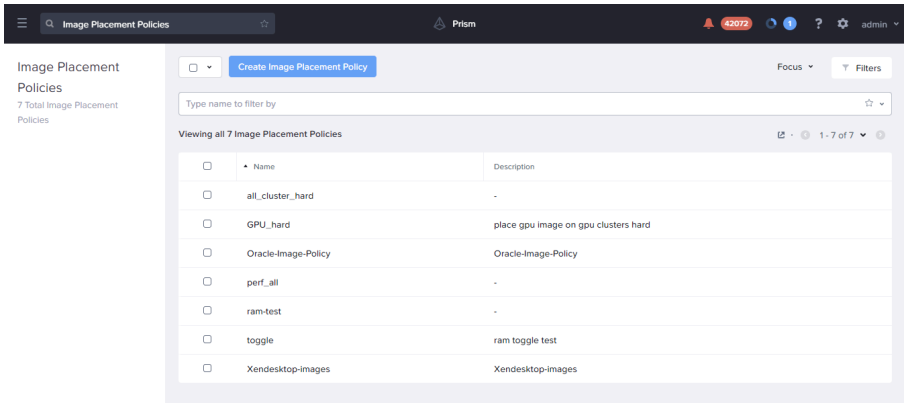

Figure 82: Image Placement Policies Dashboard

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 32: 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 381).

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 112). 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 382). 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.

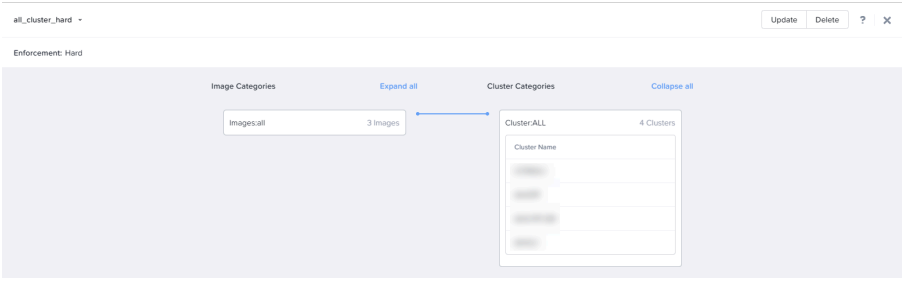

Figure 83: Image Placement Policy Details View

Hardware Entities

You can access dashboards for the following hardware components from the **Hardware** category of the entities menu (see Entities Menu on page 12):

- Clusters (see Clusters Summary View on page 114)
- Hosts (see Hosts Summary View on page 131)
- Disks (see Disks Summary View on page 141)
- GPUs (see GPUs Summary View on page 150)

Clusters Summary View

To access the clusters dashboard, select **Hardware** > **Clusters** from the entities menu (see Entities Menu on page 12). 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 to the right for the selected tab.

> Note: This section describes the information and options that appear in the clusters dashboard. See Entity Exploring on page 66 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 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**, **Memory Usage**, **IO Latency**, and **IOPS**. Click the **View all XX Clusters** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of cluster-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.

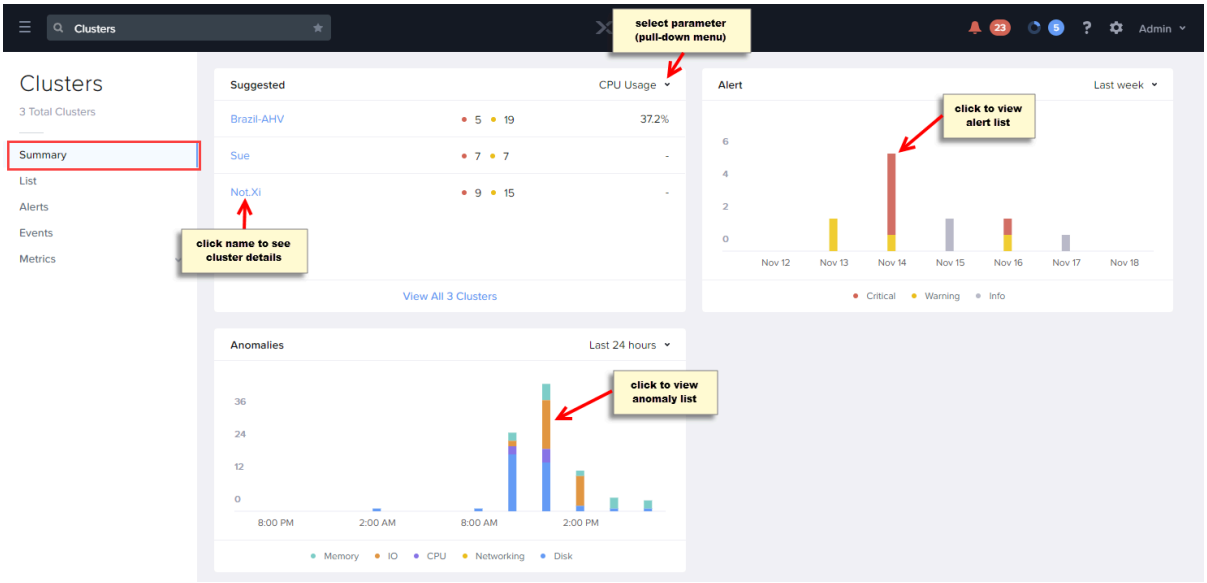

Figure 84: Clusters Summary Tab

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 **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

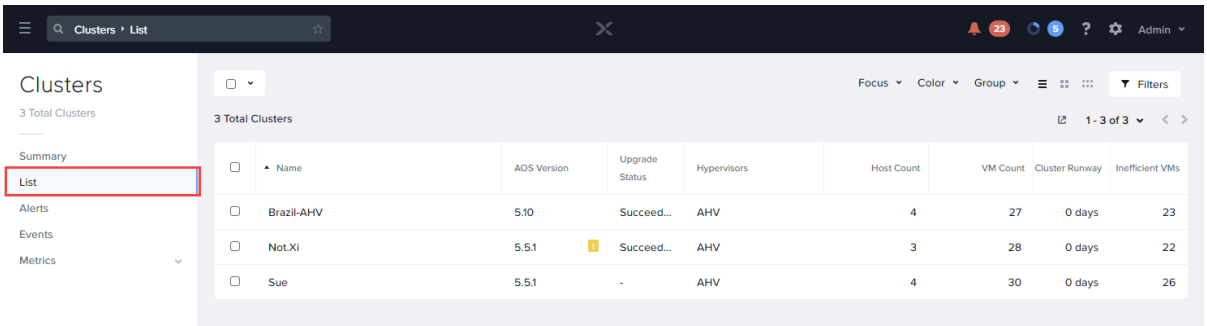

Figure 85: Clusters List Tab

Table 33: Clusters List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the cluster name. Clicking on the name displays the details page for that cluster (see Cluster Details View on page 121).	(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)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 121 for more information).	(number of days)
Inefficient VMs	Displays the number of inefficient VMs in the cluster (see Behavioral Learning Tools on page 197 for more information).	(number)
"Performance" Focus 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]

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 34: 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
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)

Parameter	Description	Values
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

You can group the clusters list in the following ways:

- The **Color** pull-down menu allows you to color code the cluster entries by AOS version 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 cluster entries by AOS version, host count, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** pull-down menu allows you to group the cluster entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

- 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:

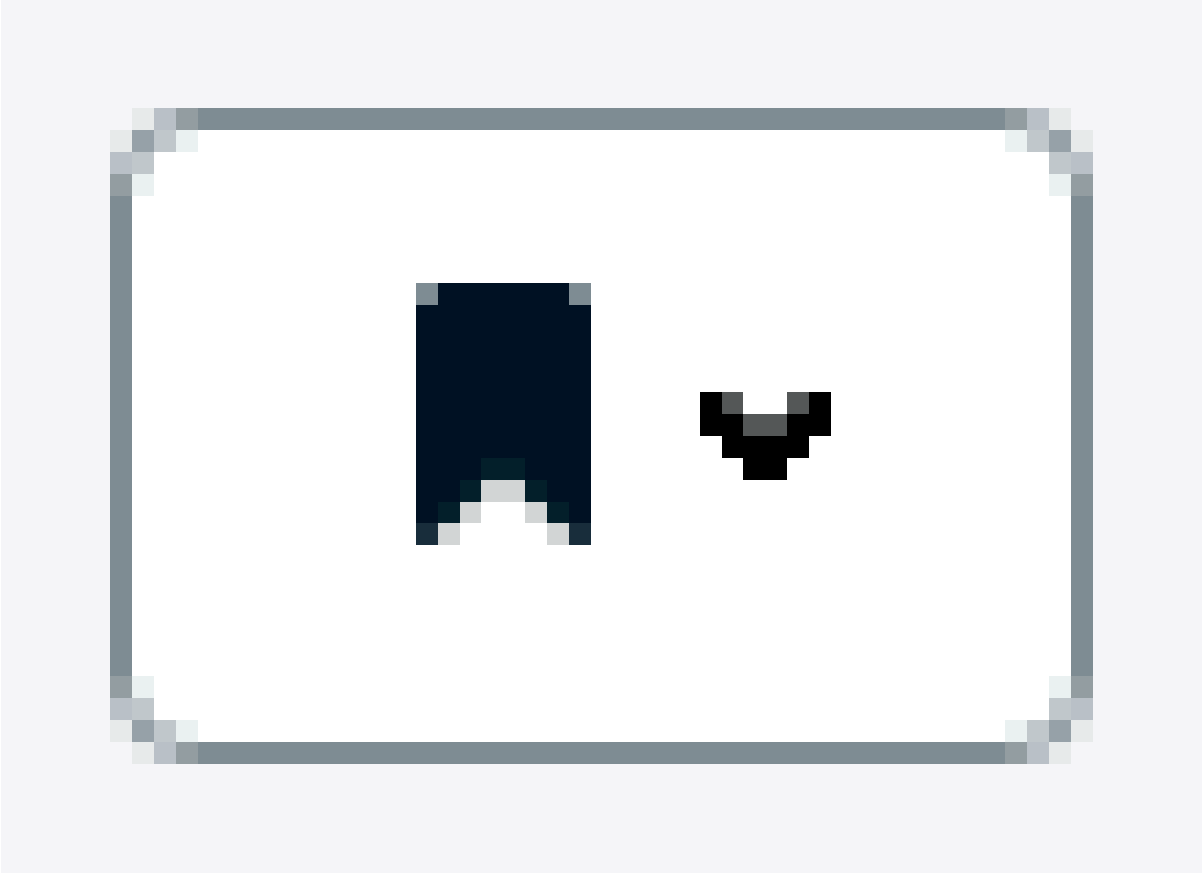

Figure 86: Label Icon

  - 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 **Unregister Cluster** to unregister this cluster from Prism Central. Unregistering means Prism Central will no longer contact the cluster to retrieve information or send commands.
- Select **Upgrade Software** to upgrade the AOS version on that cluster (see Installing or Upgrading on page 27).
- 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.

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 214).

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 220).

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 to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

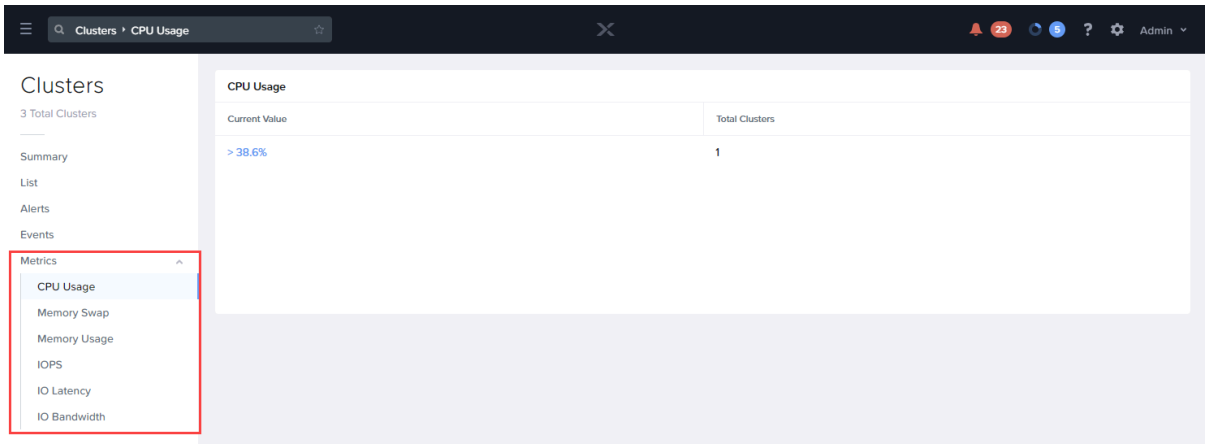

Figure 87: Clusters Metrics Tab

Table 35: 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 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.

Cluster Details View

To access the details page for a cluster, go to the clusters dashboard **List** tab (see Clusters Summary View on page 114) and click the cluster name. You can also access the details page by clicking the cluster name wherever that name appear, such as in a dashboard widget or search result.

The cluster name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, **Usage**, **Hardware**, **Entities**, and **Capacity**. Click a tab to display that information on the right. (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) 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 **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 197 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**.
  - **Unregister Cluster**: Click this button to unregister the cluster from Prism Central (see Register (Unregister) Cluster with Prism Central on page 34).
  - **Upgrade Software**: Click this button to upgrade the AOS version on this cluster (see Installing or Upgrading on page 27).
  - **Rack Configuration**: Click this button to configure the rack awareness feature (see the *Prism Web Console Guide*).

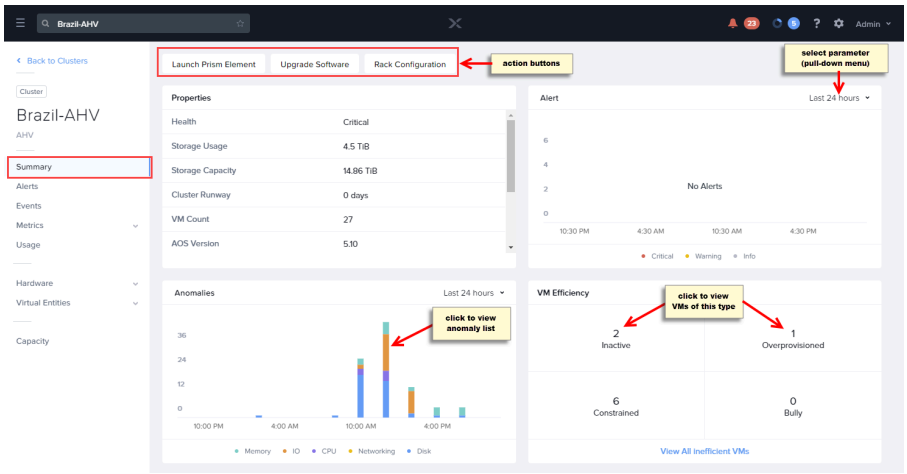

Figure 88: Cluster Summary Tab

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 36: 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 214).

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 220).

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 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 197). 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 229).

Figure 89: Cluster Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 37: 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.
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the cluster.
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 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.

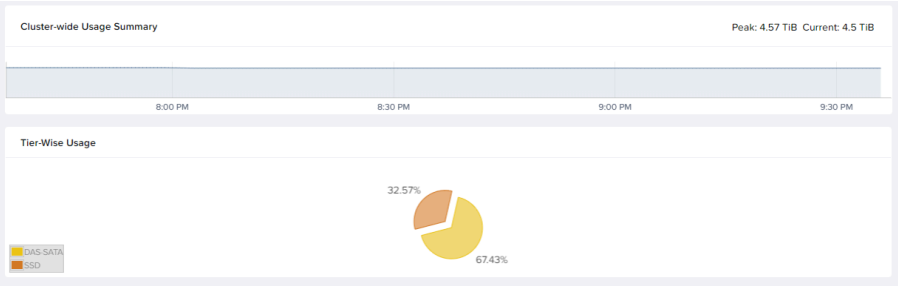

Figure 90: Cluster Usage Tab

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 114 and Virtual Infrastructure Entities on page 71 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.
> - Prism Central requires 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. Prism Central checks 90 days of past data for monthly seasonality and adjusts the baseline accordingly if it finds a seasonality pattern.

- 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, means 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 197). 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 385).

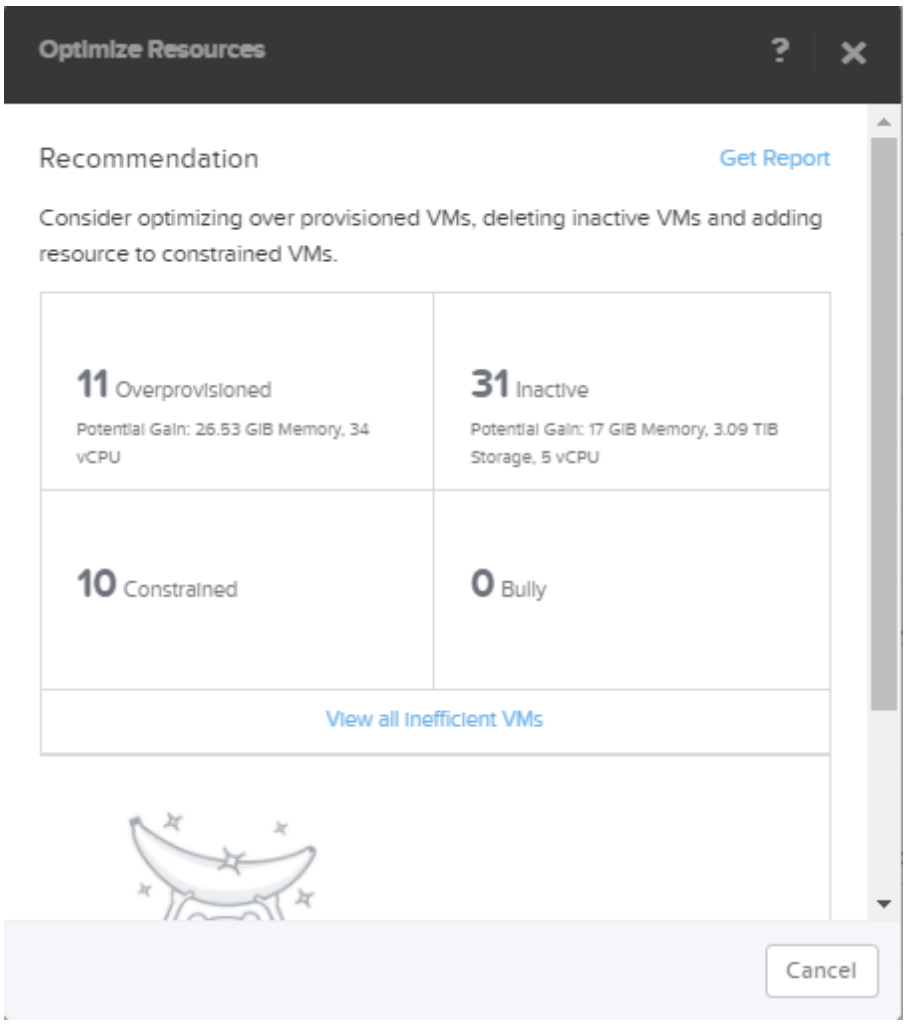

Figure 91: Optimize Resources Window

- 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 179).
- 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 on 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.

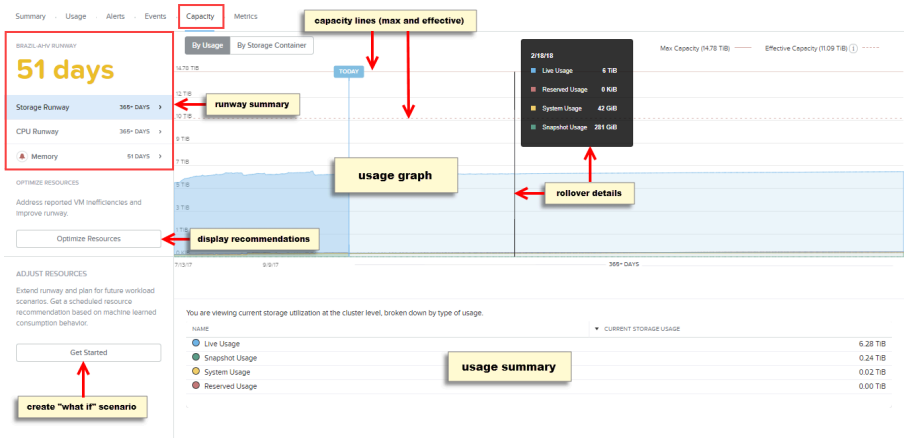

Figure 92: Cluster Capacity Tab

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. 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.

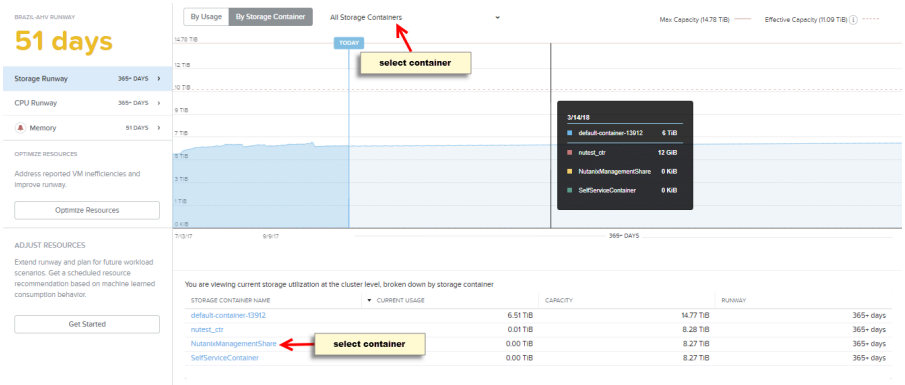

Figure 93: Storage Container View

The following table describes the fields in the storage table.

Table 38: 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
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 39: 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 40: 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]
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 197 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 229).

The following table describes the available metrics.

> Note: Some of these metrics are not available on all hypervisors.

Table 41: 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.

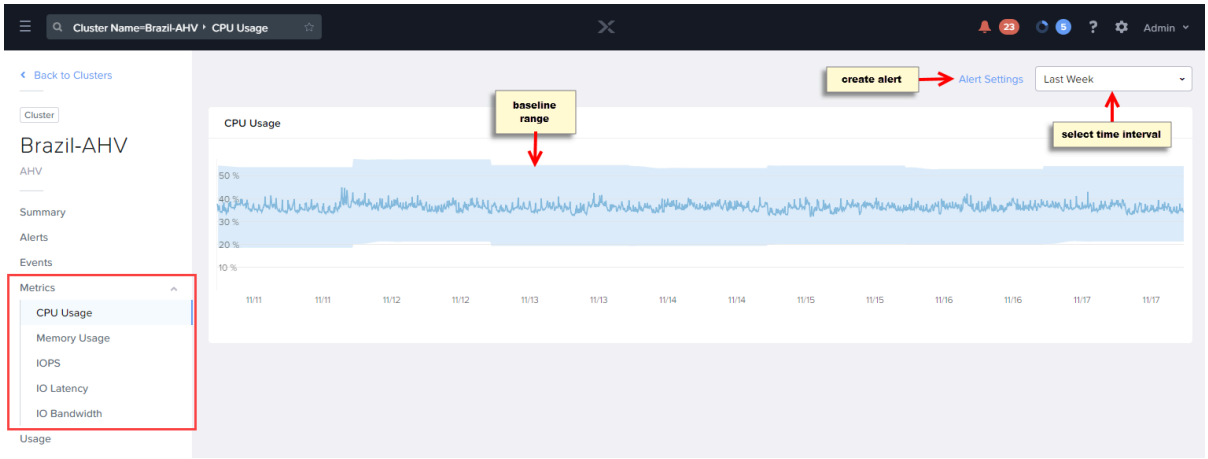

Figure 94: Metrics Tab: CPU Usage metric

Hosts Summary View

To access the hosts dashboard, select **Hardware** > **Hosts** from the entities menu (see Entities Menu on page 12). 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 66 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.

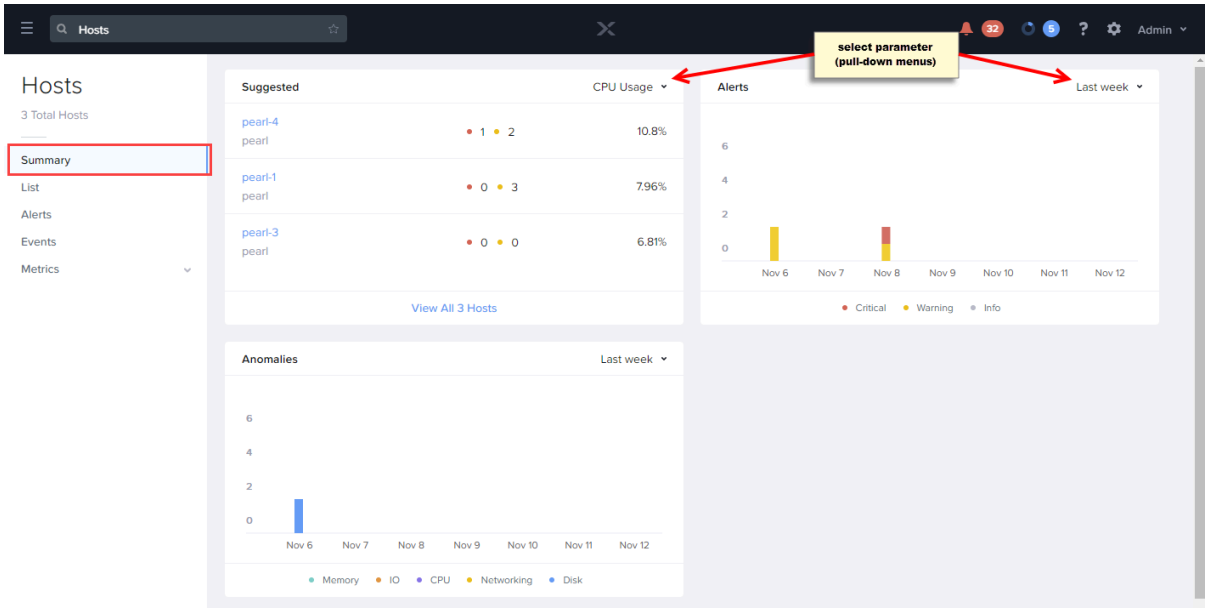

Figure 95: Hosts Summary Tab

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.

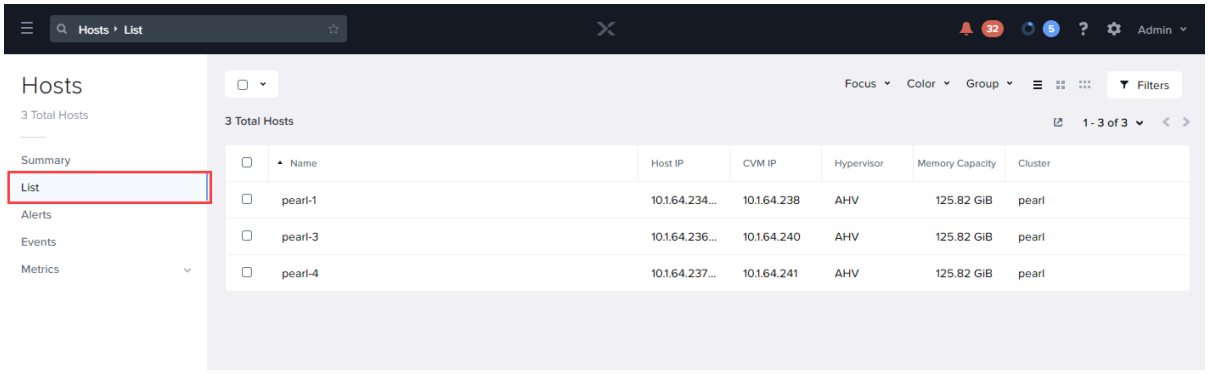

Figure 96: Hosts List Tab

Table 42: 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 137).	(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 43: 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)

Parameter	Description	Values
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)
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 **Color** pull-down menu allows you to color code the host entries by hypervisor type, memory capacity, 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 host entries by cluster, hypervisor type, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** 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 353).

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 214).

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 220).

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.)

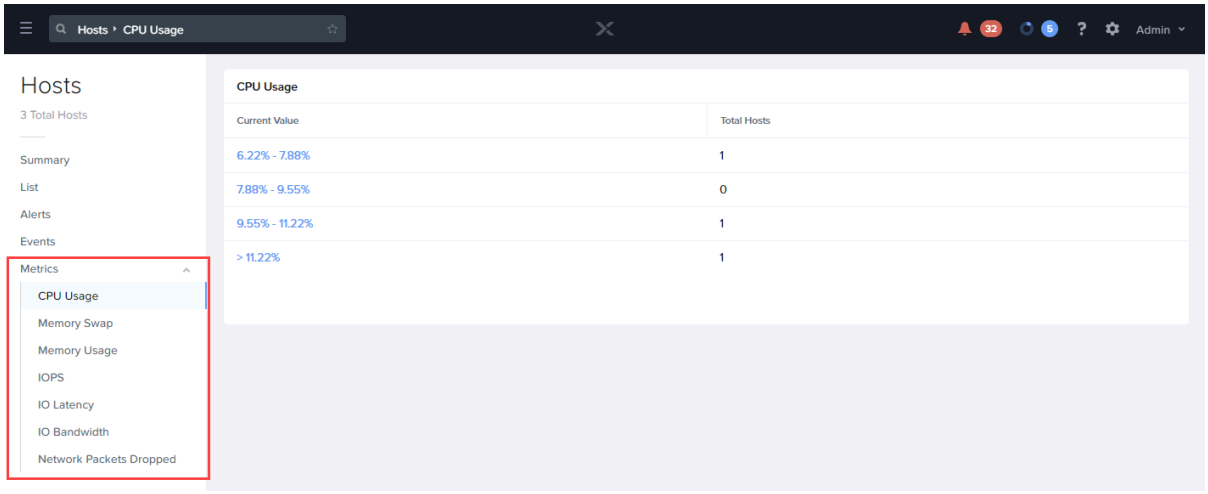

Figure 97: Hosts Metrics Tab

Table 44: 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 131) 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 351).

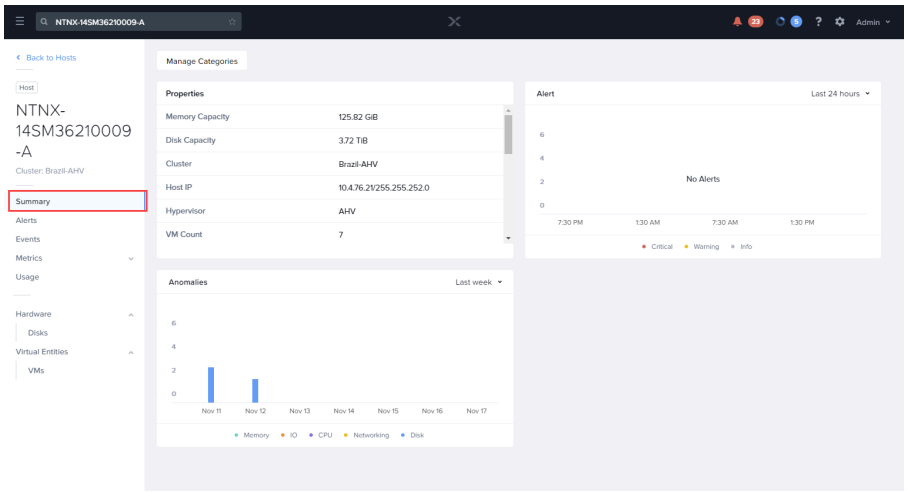

Figure 98: Host Summary Tab

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 45: 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 121).	(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)
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]

Parameter	Description	Values
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)

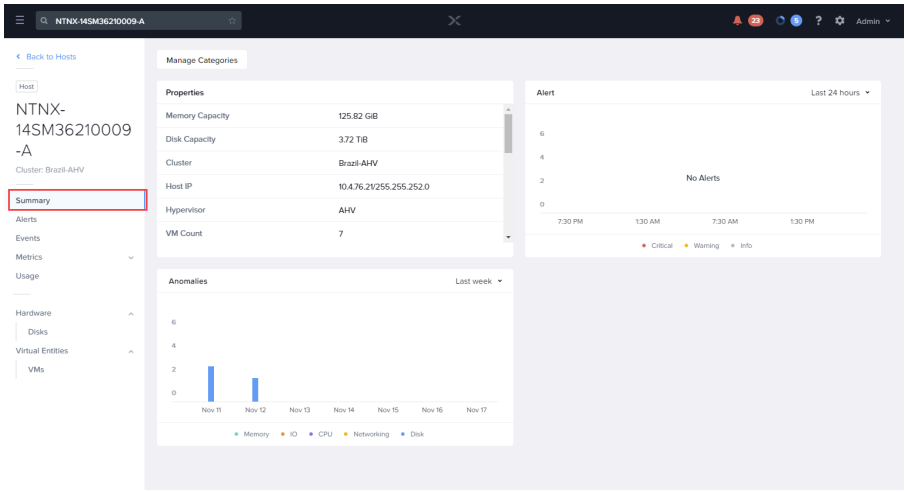

Figure 99: Host Summary Tab

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 214).

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 220).

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 197). 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 229).

Figure 100: Host Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 46: 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).

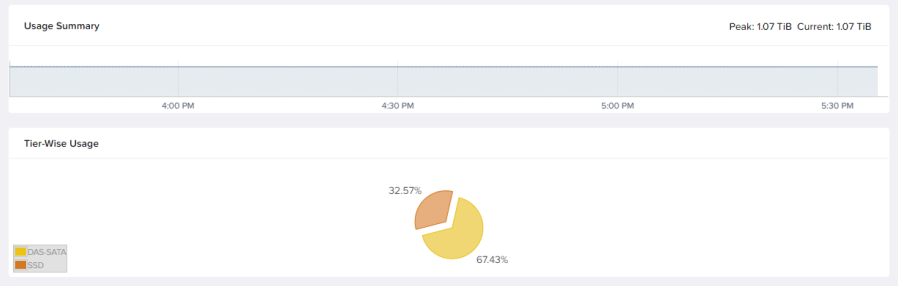

Figure 101: Host Usage Tab

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 114 and Virtual Infrastructure Entities on page 71 for more information.

Disks Summary View

To access the disks dashboard, select **Hardware** > **Disks** from the entities menu (see Entities Menu on page 12). 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 66 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.

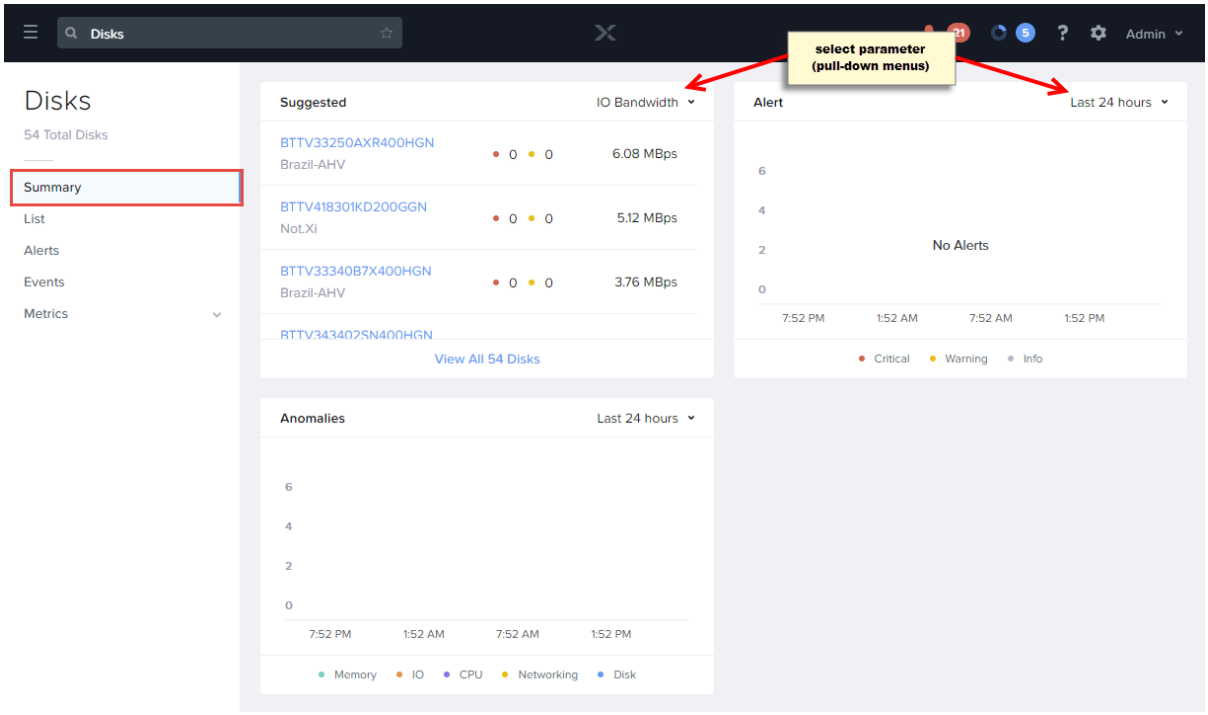

Figure 102: Disks Summary Tab

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.

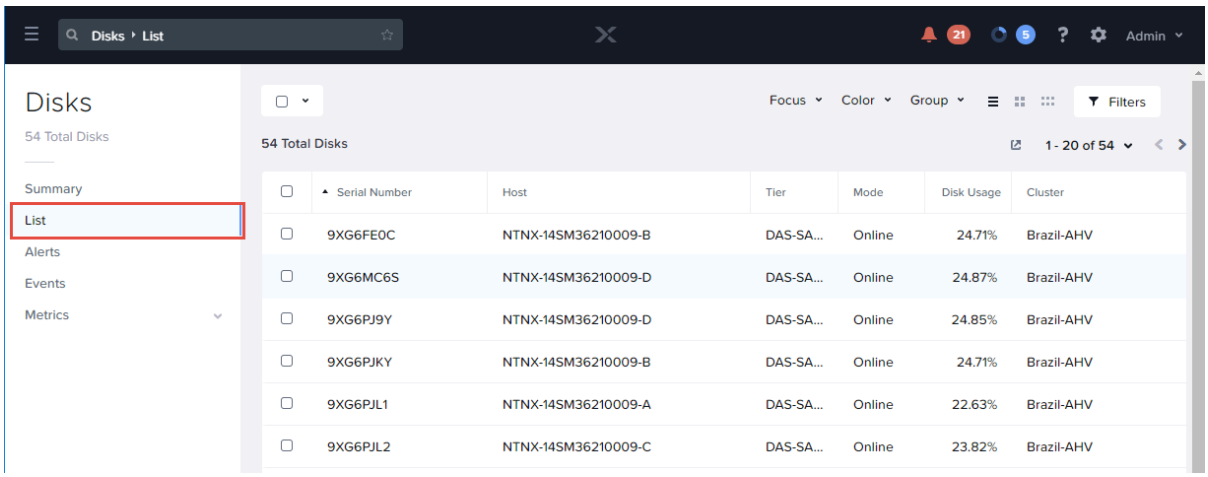

Figure 103: Disks List Tab

Table 47: Disks List Fields

Parameter	Description	Values
"General" Focus Fields		

Parameter	Description	Values
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 48: 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)

Parameter	Description	Values
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)
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 214).

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 220).

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.

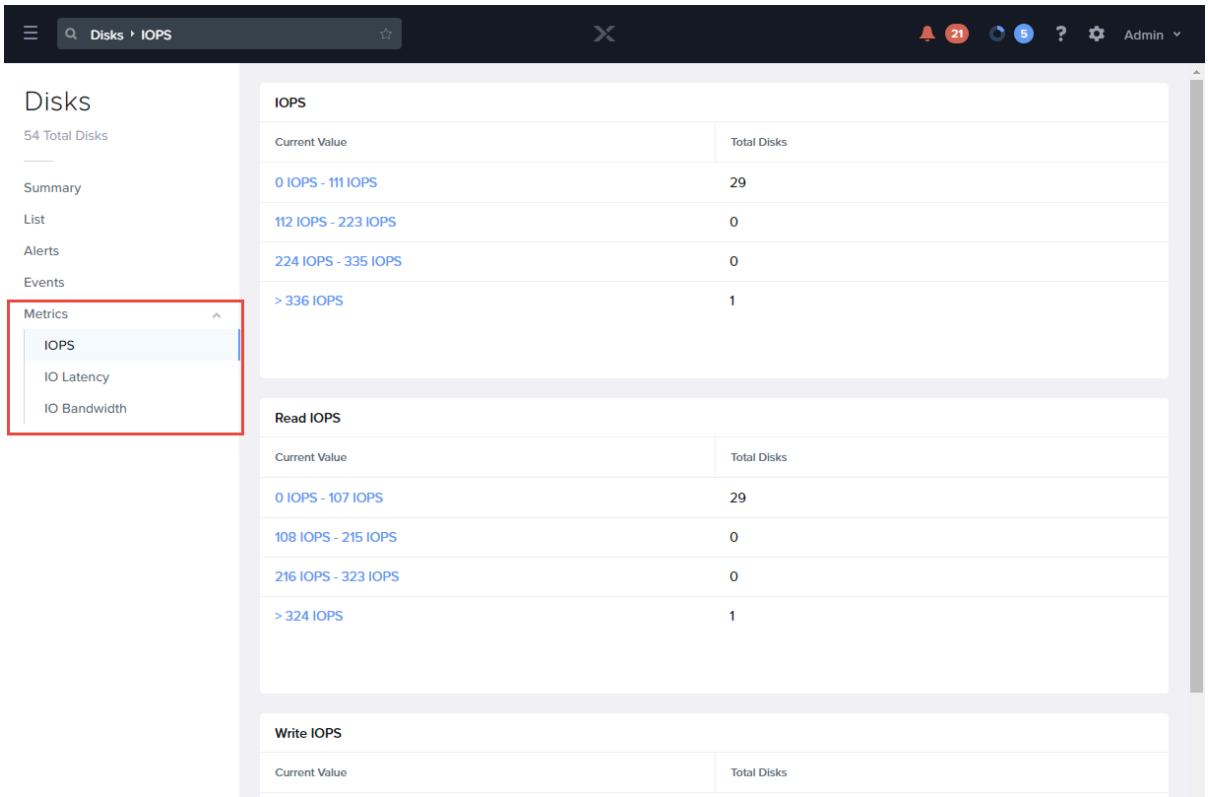

Figure 104: Disks Metrics Tab

Table 49: 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 141) 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.

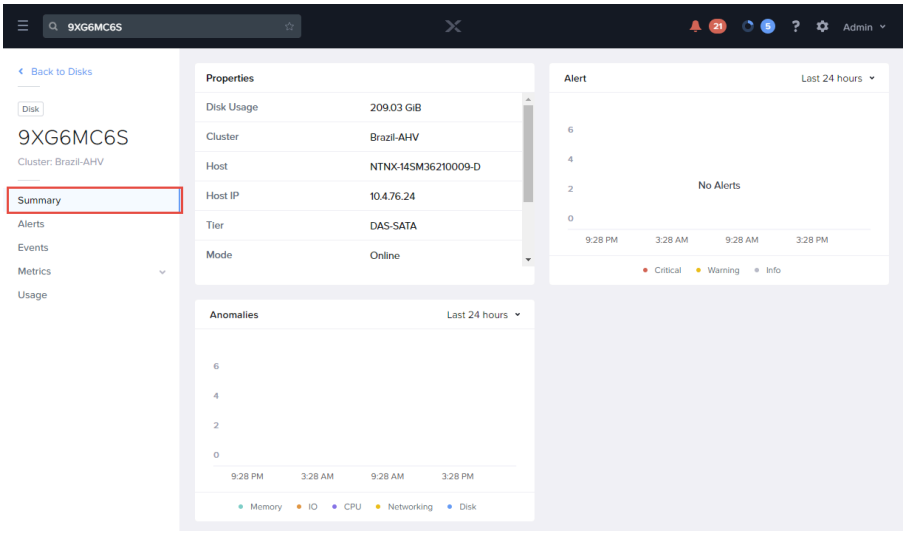

Figure 105: Disk Summary Tab

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 50: 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]

Parameter	Description	Values
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

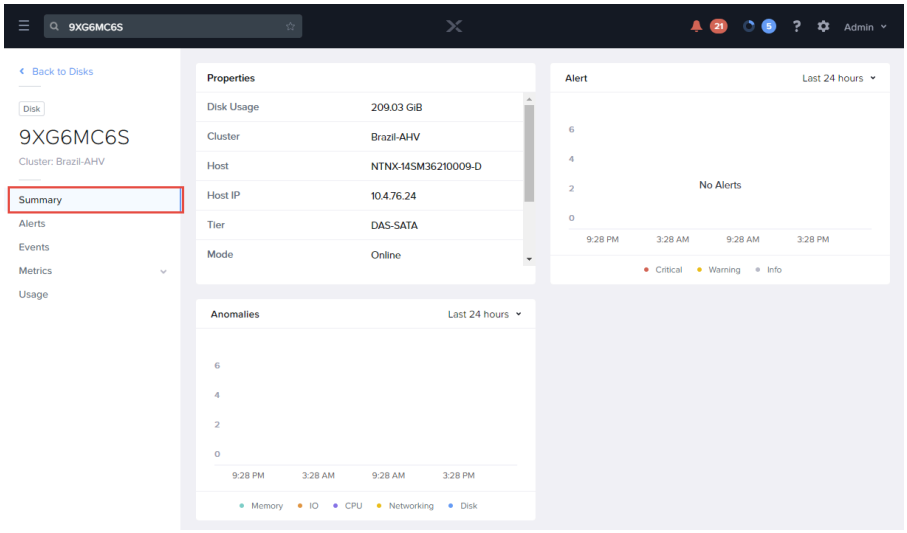

Figure 106: Disk Summary Tab

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 214).

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 220).

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 229).

Figure 107: 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.

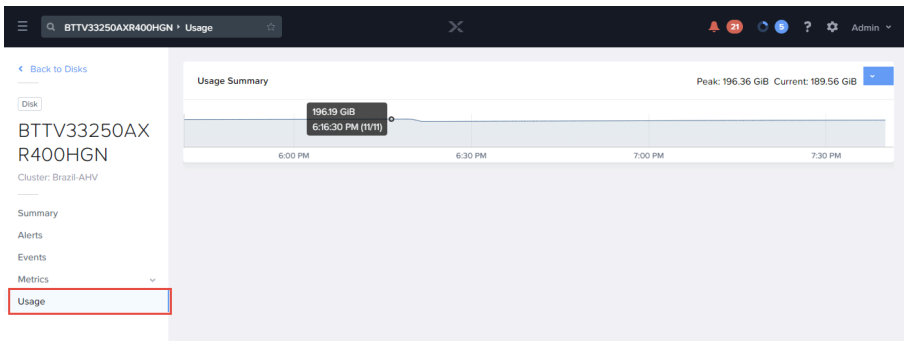

Figure 108: Disk Usage Tab

GPUs Summary View

To access the GPUs dashboard, select **Hardware** > **GPUs** from the entities menu (see Entities Menu on page 12). 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 66 for instructions on how to view and organize that information in a variety of ways.

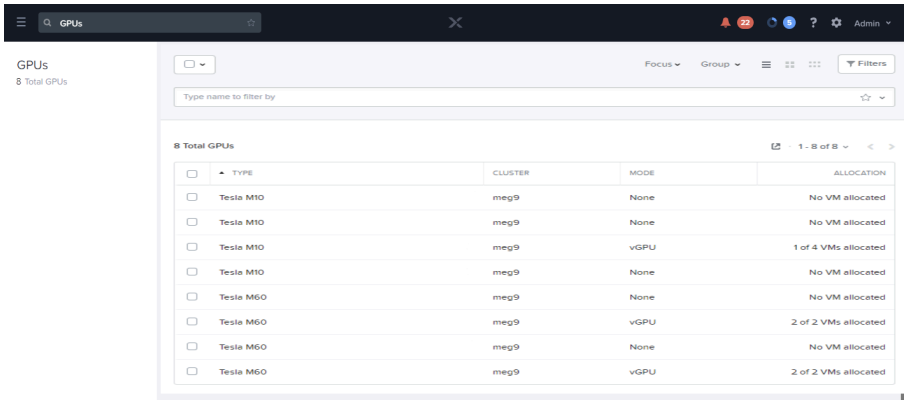

Figure 109: GPUs Dashboard

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 51: 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 52: 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.)
- [Tiles and Circles views only] 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 150). 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 53: 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)

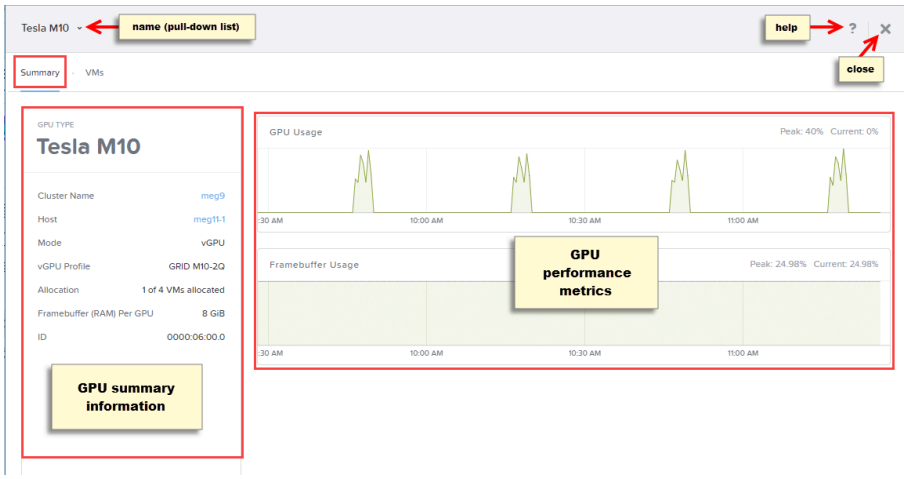

Figure 110: GPU Summary Tab

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 81).
- **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 12):

- Alerts (see Alerts Summary View on page 214)
- Events (see Events Summary View on page 220)
- Audits (see Audits Summary View on page 153)
- Tasks (see Tasks View on page 157)

Audits Summary View

To access the audits dashboard, select **Activity** > **Audits** from the entities menu (see Entities Menu on page 12). 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 66 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.

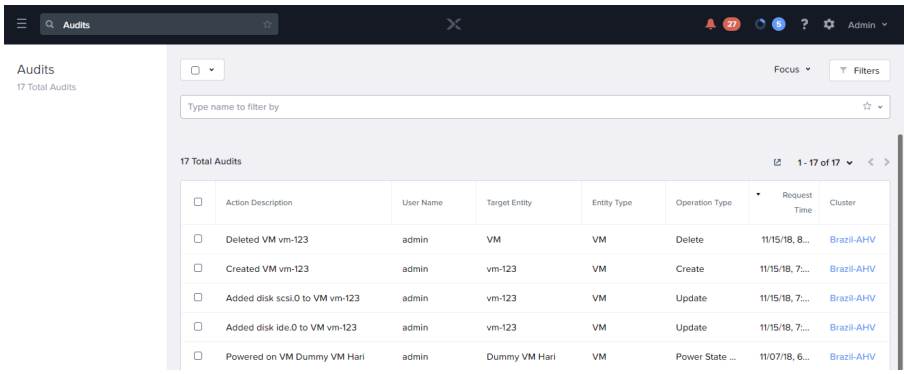

Figure 111: Audits Dashboard

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 54: 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 55: 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

Parameter	Description	Values
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 153) 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 56: 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)

Parameter	Description	Values
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 powering on a VM, so the details simply include the name of the target VM.

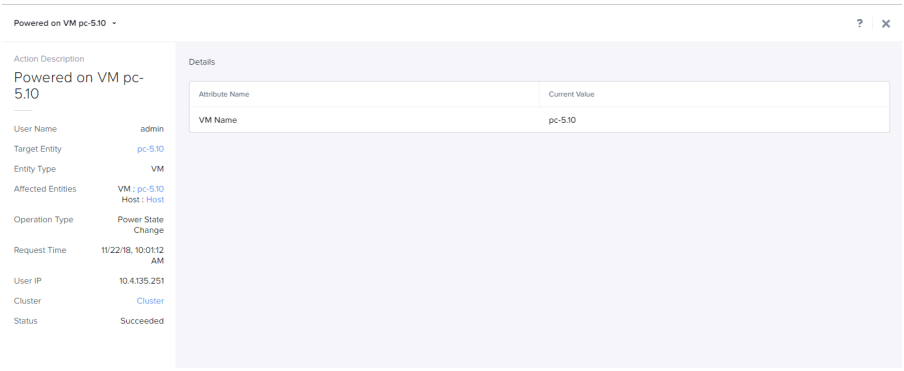

Figure 112: Audit Details Page (example 1)

In the second example the action is adding a disk, so the details include the VM name, disk address, size, and source disk UUID.

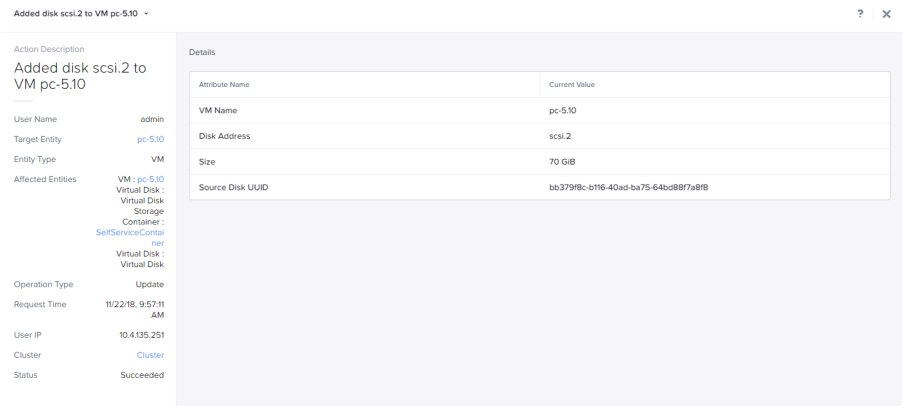

Figure 113: Audit Details Page (example 2)

In the third example the action is creating a VM, so the details include a variety of information about that VM.

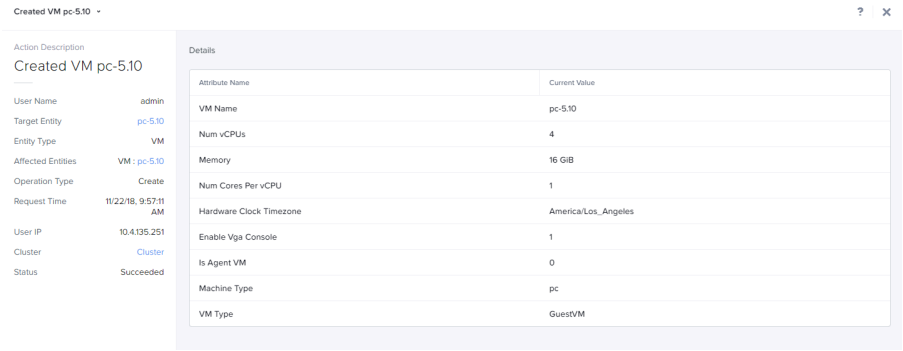

Figure 114: Audit Details Page (example 3)

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 47).
- Selecting **Activity** > **Tasks** in the entities menu (see Entities Menu on page 12).

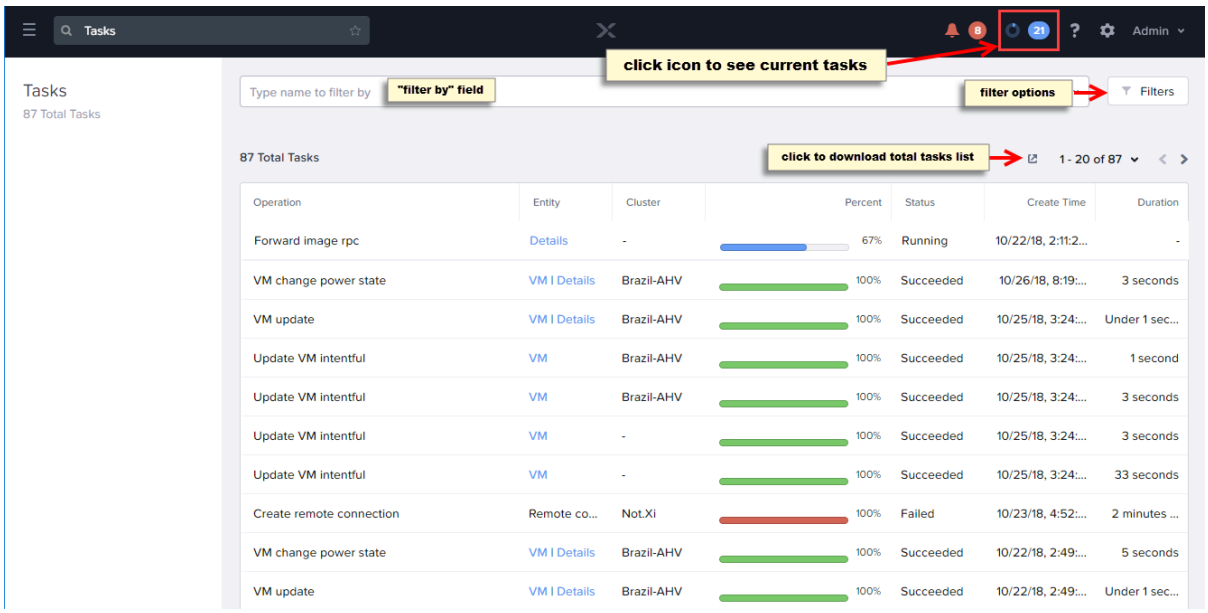

Figure 115: Tasks Screen

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.

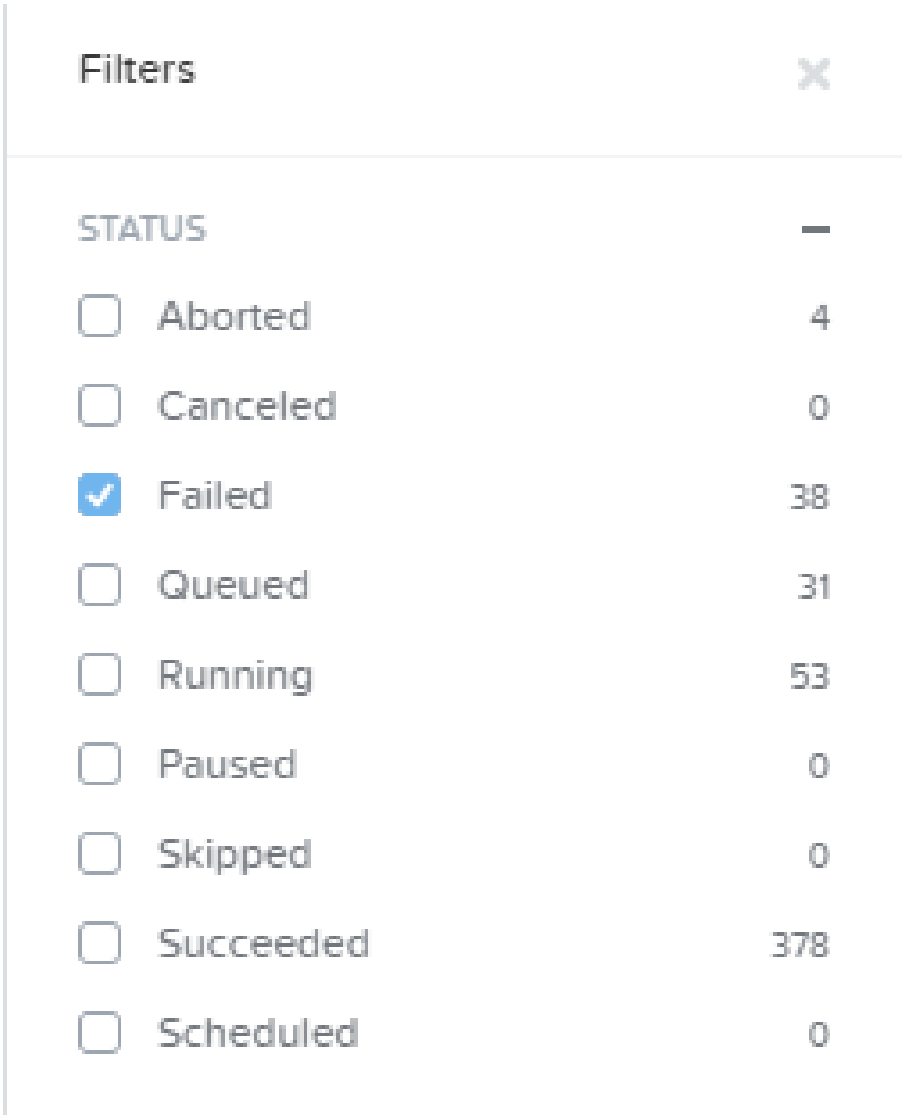

Figure 116: Filter Pane

- 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 57: Tasks List Fields

Parameter	Description	Values
Operation	Specifies which type of operation the task is performing.	(name of operation such as "upgrade")
Entity	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)
Percent	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
Create Time	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]

Operations Entities

You can access dashboards for the following tools from the **Operations** category of the entities menu (see Entities Menu on page 12):

- Analysis (see Analysis Dashboard on page 202)
- Planning (see Resource Planning on page 176)
- Reports (see Reports View on page 159)
- Codeless Task Automation (X-Play) on page 415

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 12).

Note: This section describes the information and options that appear in the Reports dashboard.

- See Entity Exploring on page 66 for instructions on how to view and organize that information in a variety of ways.
- See Reports Management on page 385 for detailed information on how to create, configure, and run reports.

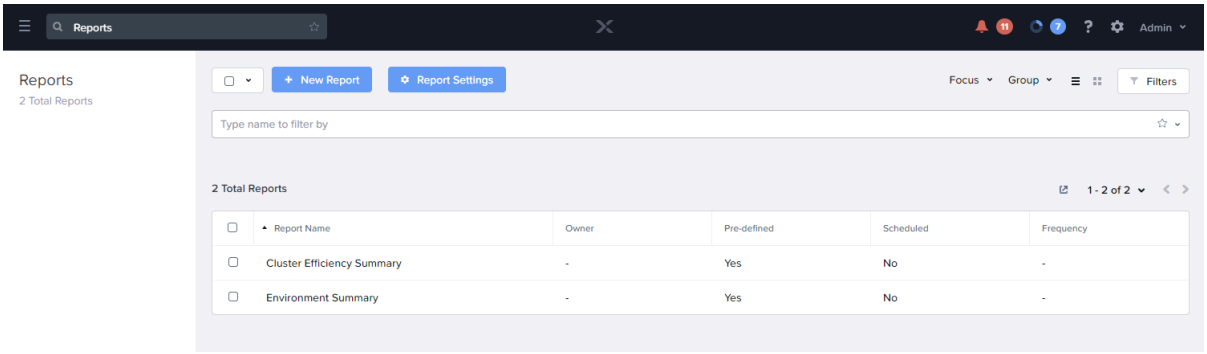

Figure 117: Reports Dashboard

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 58: Reports List Fields

Parameter	Description	Values
Report Name	Displays the name of the report.	(report name)
Owner	Displays the name of the user who created the report.	(user name)
Pre-Defined	Indicates whether this is a pre-defined report. • "Yes" means this is one of the pre-defined reports included with the product. • "No" means it is a custom report created by a user.	Yes, No
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.

You can group the reports by pre-defined status, that is the pre-defined reports appear as one group while the custom reports appear as a separate group.

When the Tiles view is selected, the **Sort** pull-down menu allows you to group the entries by report name, owner, pre-defined, scheduled, or frequency. (You can only choose one parameter.)

The Reports view includes two action buttons: a **Create Report** action button with a drop-down list to create a new report or configure report settings.

- Click the **New Report** button to create a new custom report (see Creating a New Report on page 386).
- Click the **Report Settings** button to configure the report appearance, e-mail, and retention settings (see Configuring Report Settings on page 411).

The **Actions** menu appears when one or more reports are selected. It includes the following actions (see Managing a Report on page 387):

- This action can be applied to multiple reports: **Delete**
- These actions can be applied to only one report at a time: **Run**, **Edit**, **View Instances**, **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 12):

- LCM (see LCM View on page 161)
- Projects (see Projects Summary View on page 161)
- Roles (see Roles Summary View on page 167)
- Users (see Users Summary View on page 170)
- Availability Zones (see Availability Zones View on page 174)

LCM View

To access the LCM dashboard, select **Administration** > **LCM** from the entities menu (see Entities Menu on page 12). 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.

Projects Summary View

To access the projects dashboard, select **Administration** > **Projects** from the entities menu (see Entities Menu on page 12). 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 66 for instructions on how to view and organize that information in a variety of ways.
- See Project Management on page 373 for information on how to create and manage projects.

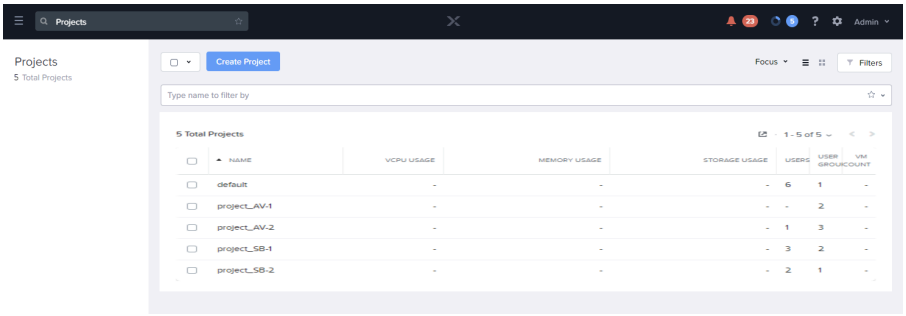

Figure 118: Projects Dashboard

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 59: 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]
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 60: 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)

Parameter	Description	Values
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 373). 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 376). 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 161) 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 “” ()).
  - 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 61: 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)
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)

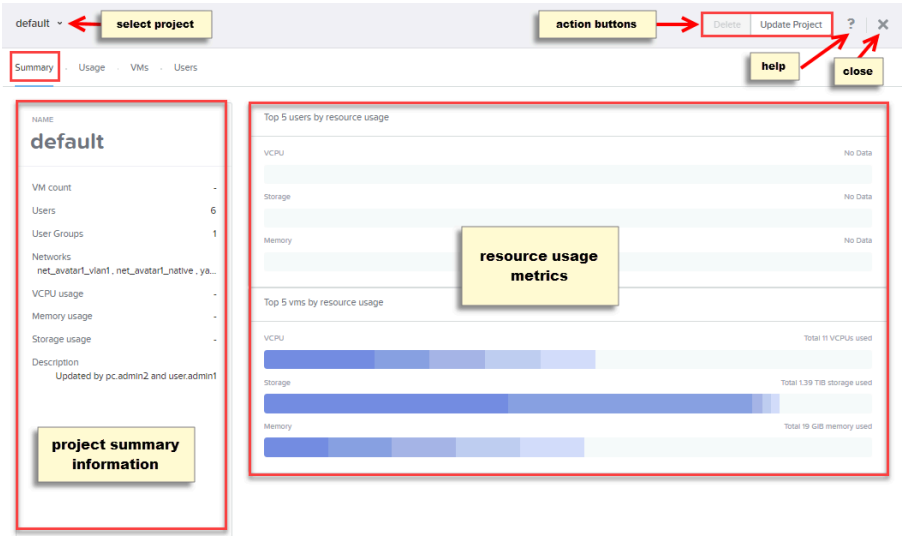

Figure 119: Project Summary Tab

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.

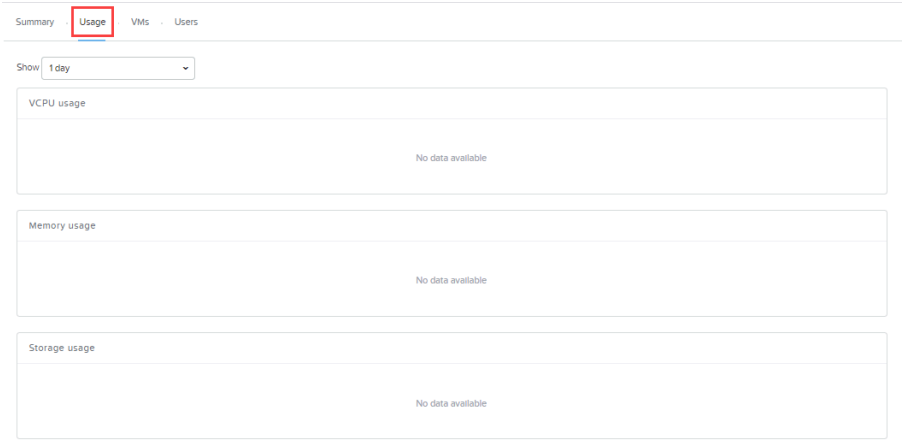

Figure 120: Project Usage Tab

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 62: 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]

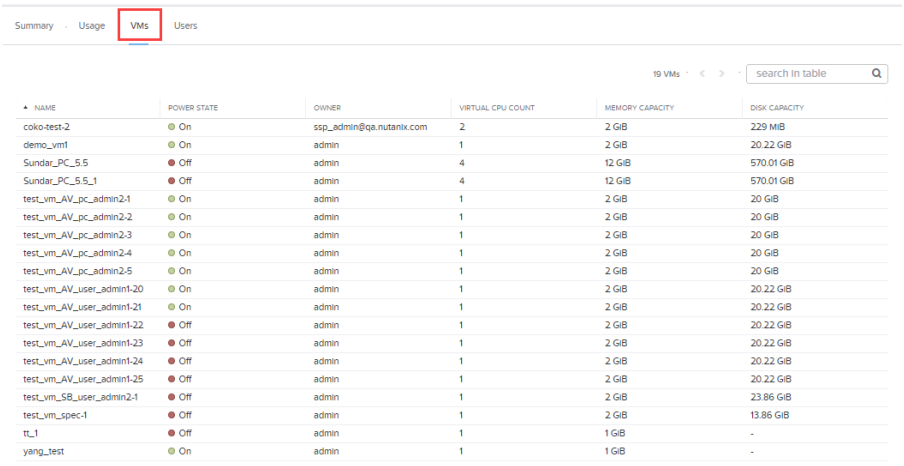

Figure 121: Project VMs Tab

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 63: Users Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
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]

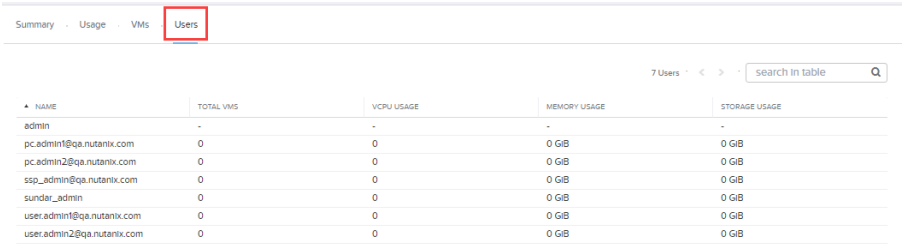

Figure 122: Project Users Tab

Roles Summary View

To access the roles dashboard, select **Administration** > **Roles** from the entities menu (see Entities Menu on page 12). 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 66 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.

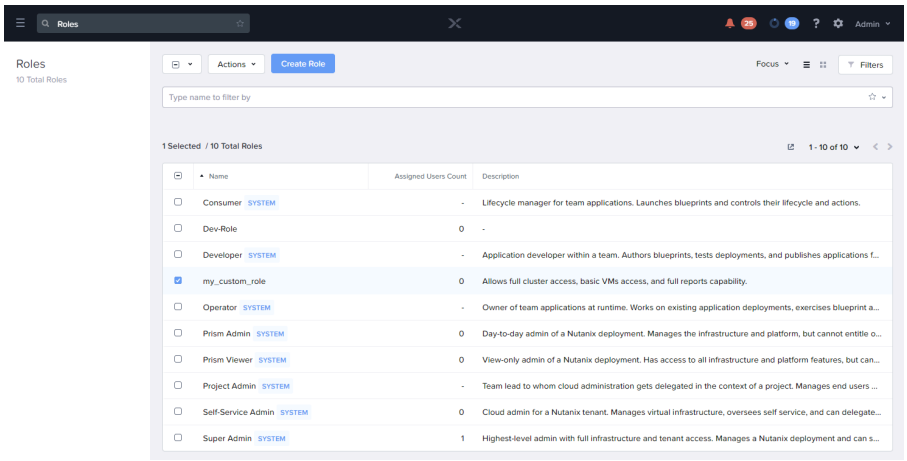

Figure 123: Roles Dashboard

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 64: 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 65: 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 168).

> 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 167) 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 66: 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)

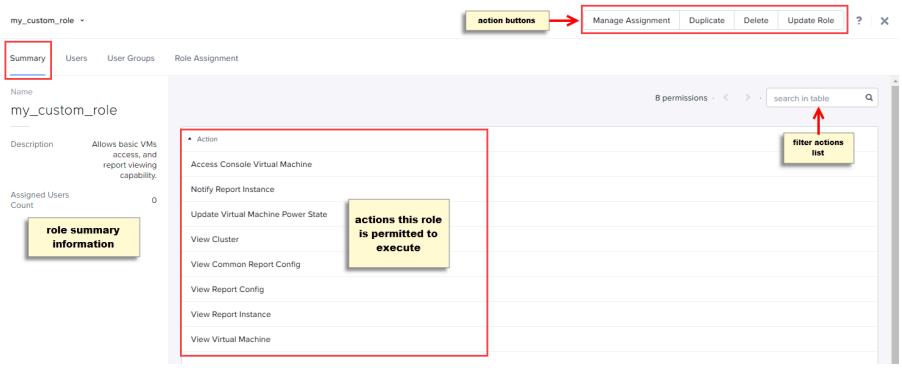

Figure 124: Role Summary Tab

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.

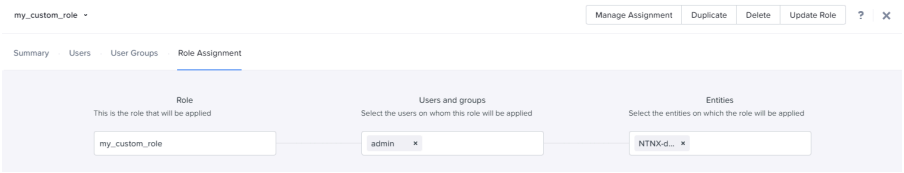

Figure 125: Role Assignment Tab

Users Summary View

To access the users dashboard, select **Administration** > **Users** from the entities menu (see Entities Menu on page 12). 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 66 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.

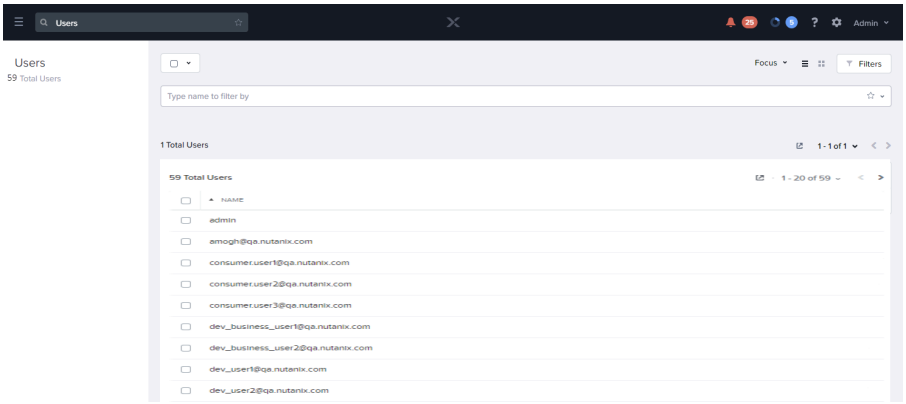

Figure 126: Users Dashboard

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 170) 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 67: 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]

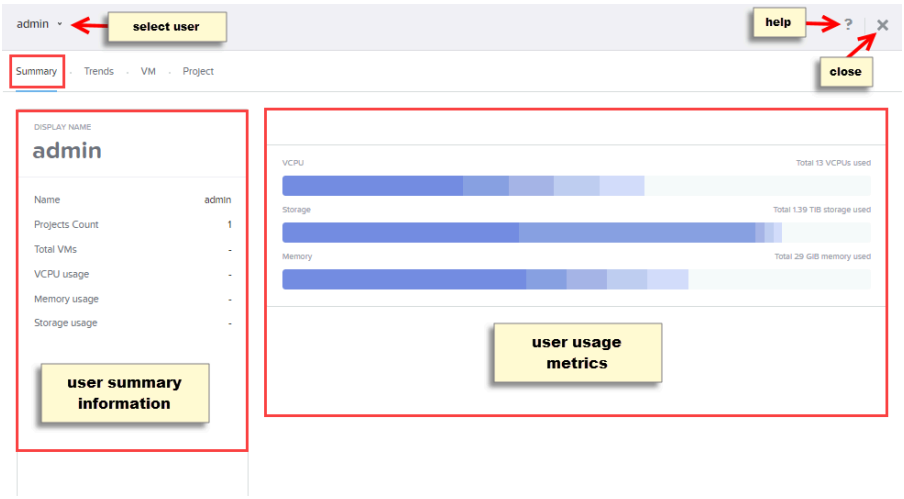

Figure 127: User Summary Tab

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.

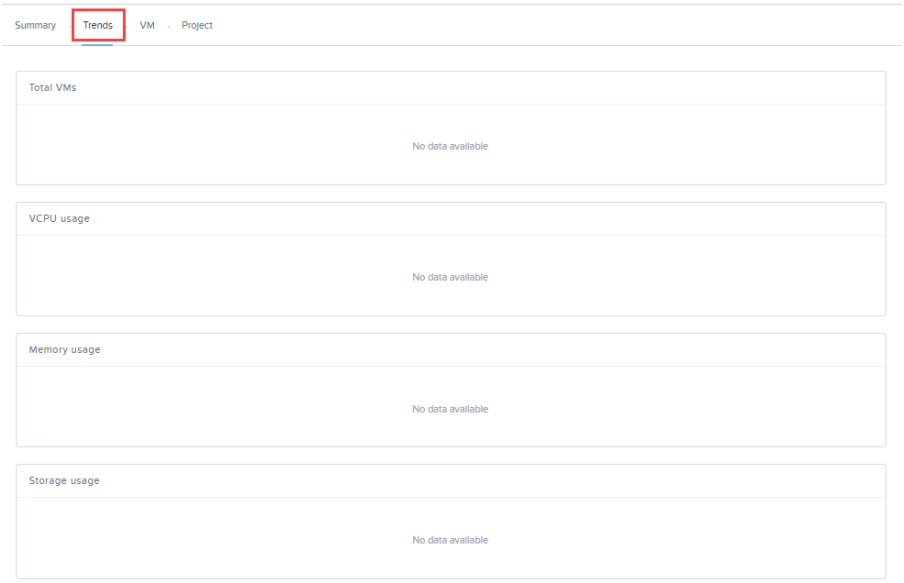

Figure 128: User Trends Tab

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 68: 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]

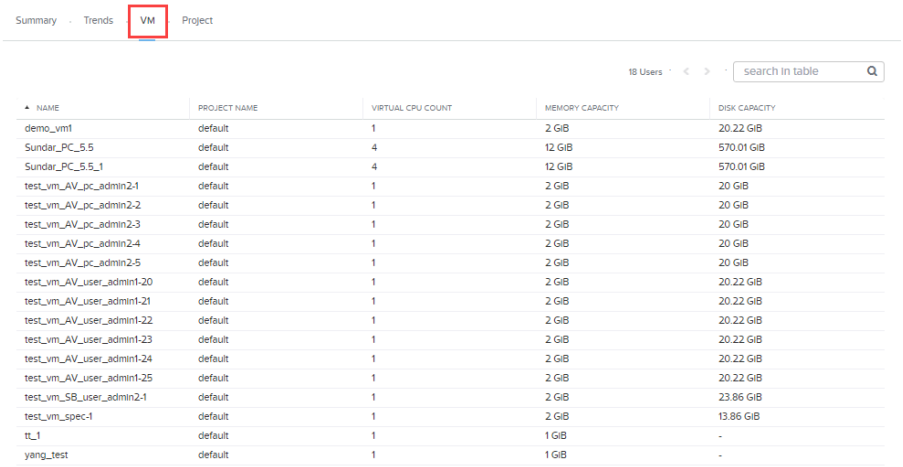

Figure 129: User VM Tab

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 69: User Project Fields

Parameter	Description	Values
Name	Displays the name of the project.	(project name)

Parameter	Description	Values
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]

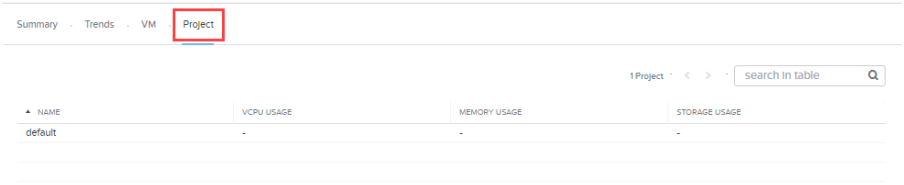

Figure 130: User Project Tab

Availability Zones View

To access the availability zones dashboard, select **Administration** > **Availability Zones** from the entities menu (see Entities Menu on page 12). 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 66 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.

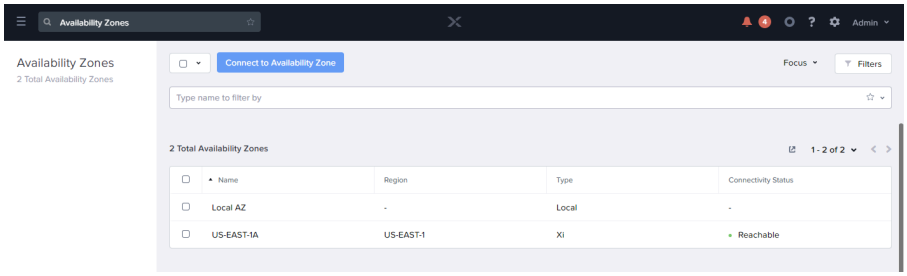

Figure 131: Availability Zones Dashboard

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 70: 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.

Services Entities

You can enable the following services from the **Services** category of the entities menu (see Entities Menu on page 12):

- Buckets (see Enabling Objects on page 213)
- Calm (see Enabling Calm on page 212)
- Karbon (see Enabling Karbon on page 213)

5

RESOURCE PLANNING

You can review and analyze current and potential resource needs through the planning dashboard. To access the planning dashboard, select **Operations** > **Planning** from the entities menu (see Entities Menu on page 12). 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 clusters (see Capacity Runway View on page 176).
- 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 177).
- 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 196.

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 197).

Capacity Runway View

Click the **Capacity Runway** tab in the planning dashboard (see Resource Planning on page 176) to view summary resource runway information for the registered clusters and to access detailed runway information about each cluster.

> Note: Prism Central requires 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. Prism Central checks 90 days of past data for monthly seasonality and adjusts the baseline accordingly if it finds a seasonality pattern.

- 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 121).

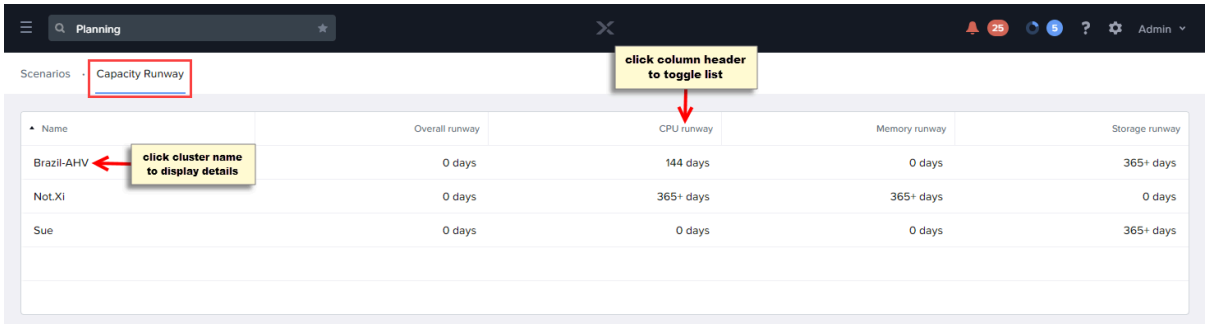

Figure 132: Capacity Runway View

The following table describes the fields that appear in the Capacity Runway view.

Table 71: Capacity Runway Tab Fields

Parameter	Description	Values
Name	Displays the cluster name. Clicking on 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 121).	(cluster name)
Overall Runway	Displays the overall runway for the cluster. The "overall" 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 176) 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 179).
- 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.

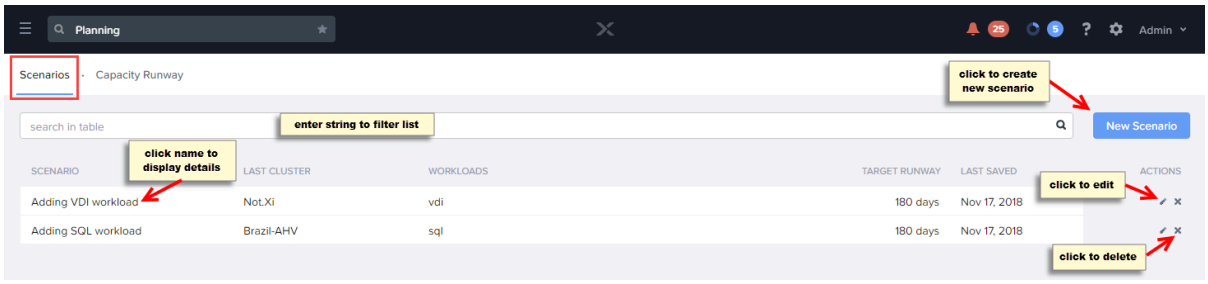

Figure 133: Scenarios View

The following table describes the fields that appear in the Scenarios view.

Table 72: 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 179).	(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),
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)

Parameter	Description	Values
Actions	• To edit the scenario, click the pencil icon (see Modifying a Scenario on page 195). • 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 177) and click the **New Scenario** button.
   The Scenario configuration page appears.

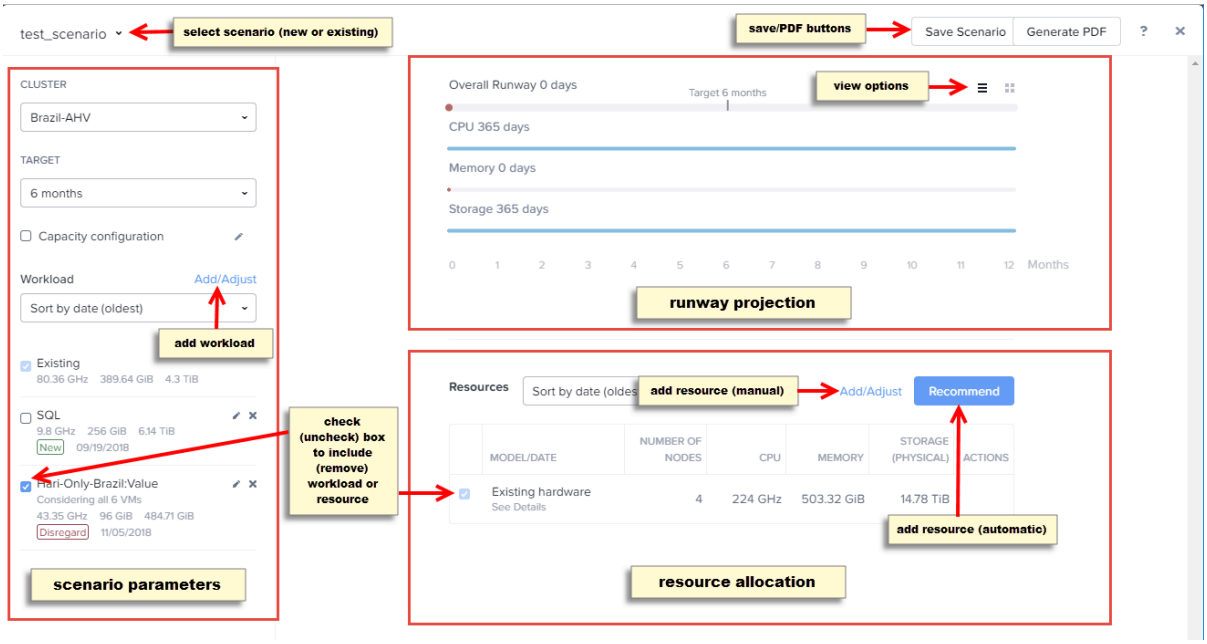

Figure 134: Scenario Configuration Page (example)

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 177), 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.

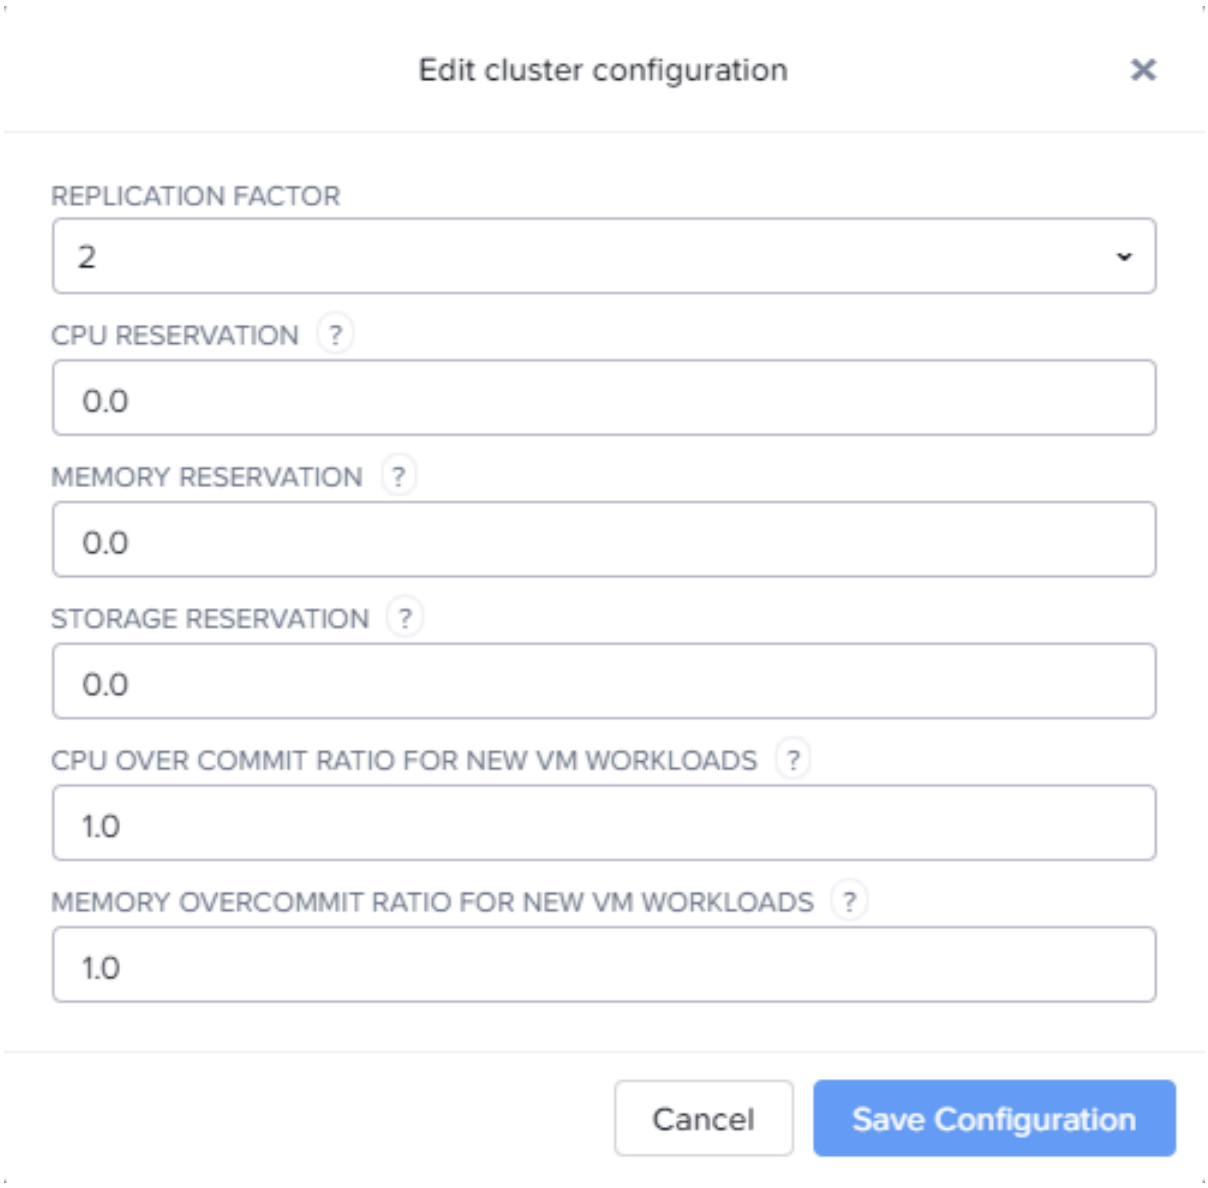

Figure 135: Edit Cluster Configuration Window

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.

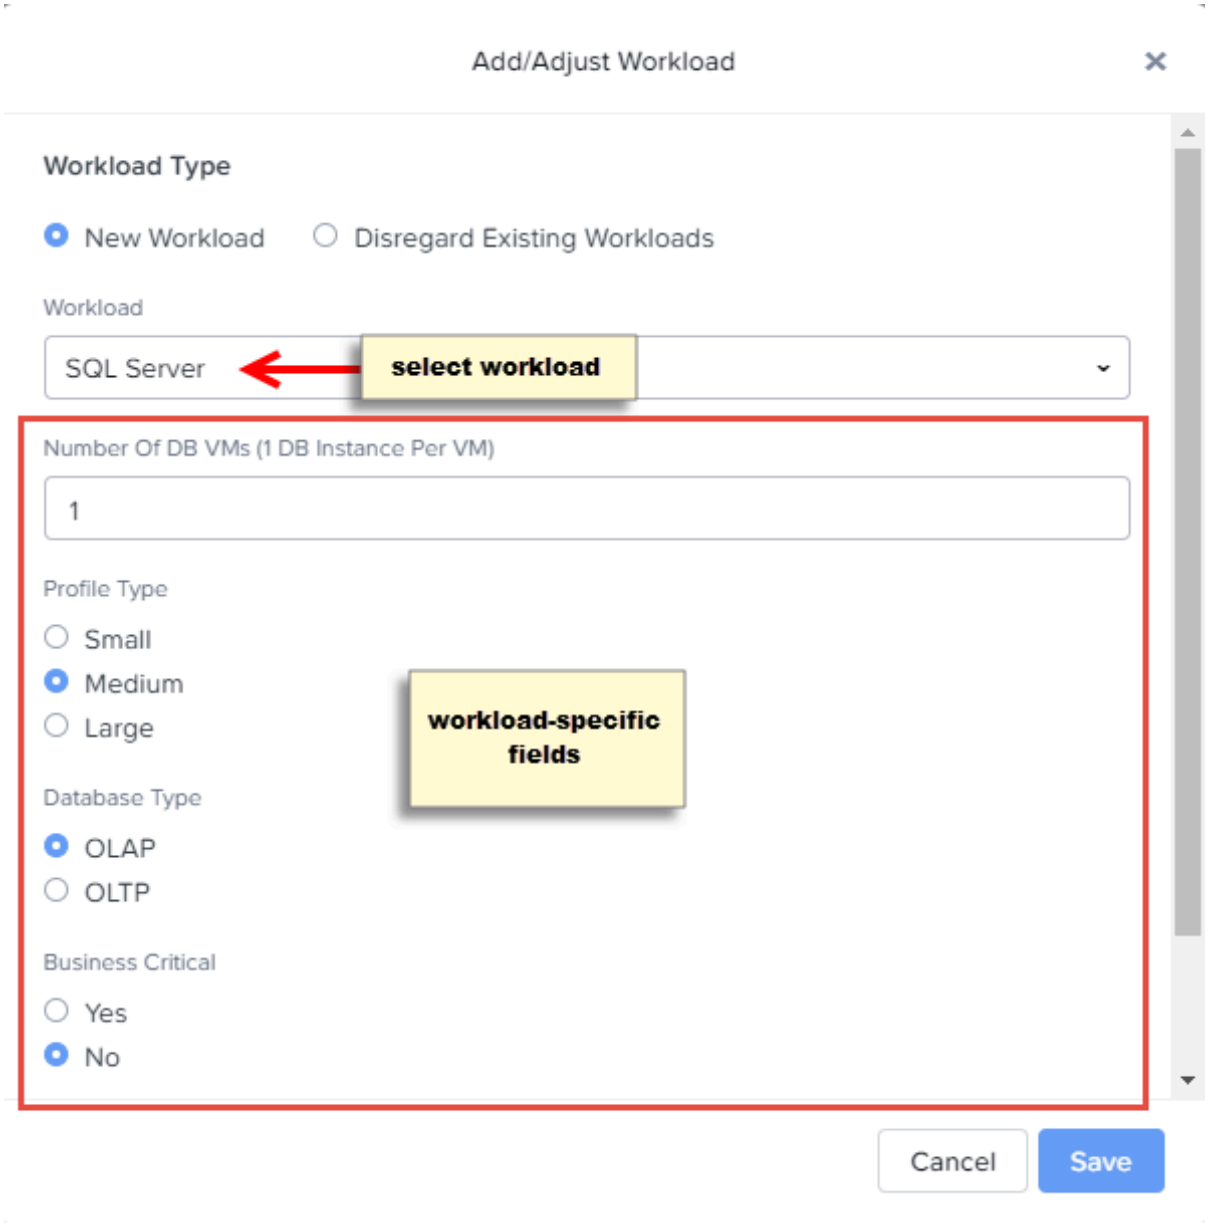

Figure 136: New Workload Fields (SQL Server example)

Table 73: 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 192 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 351).

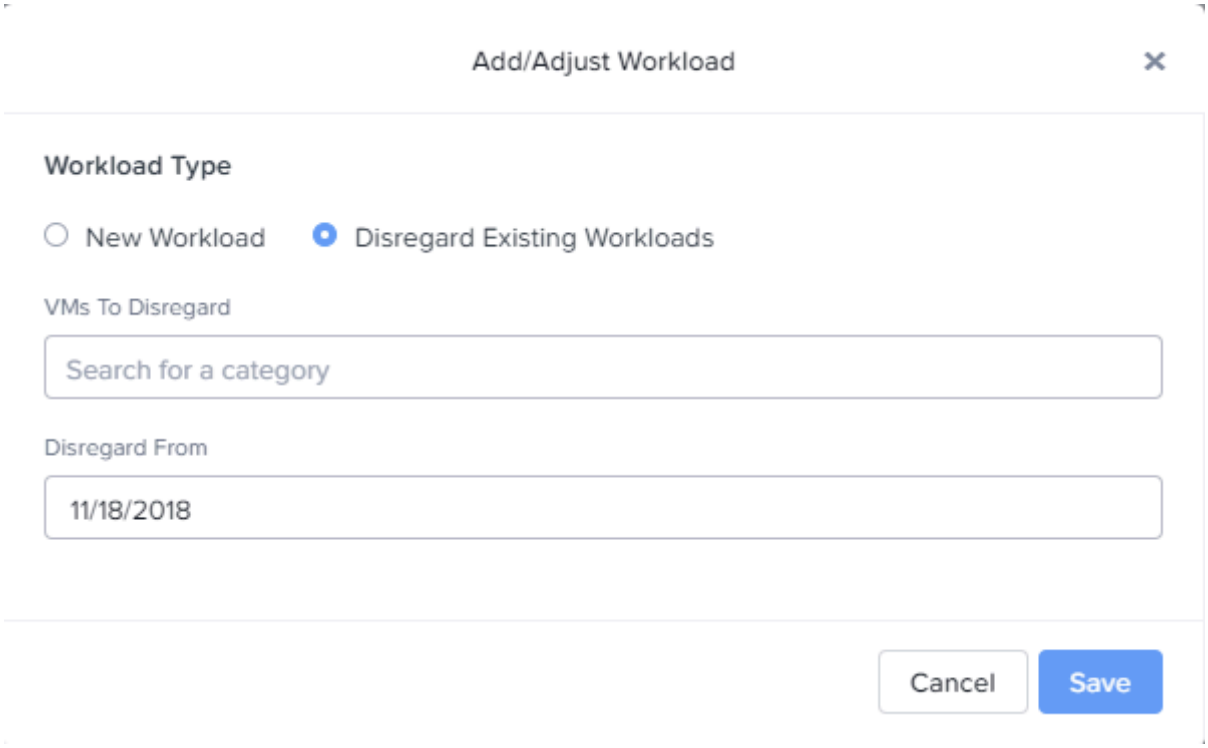

Figure 137: Disregard Existing Workloads Fields

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.

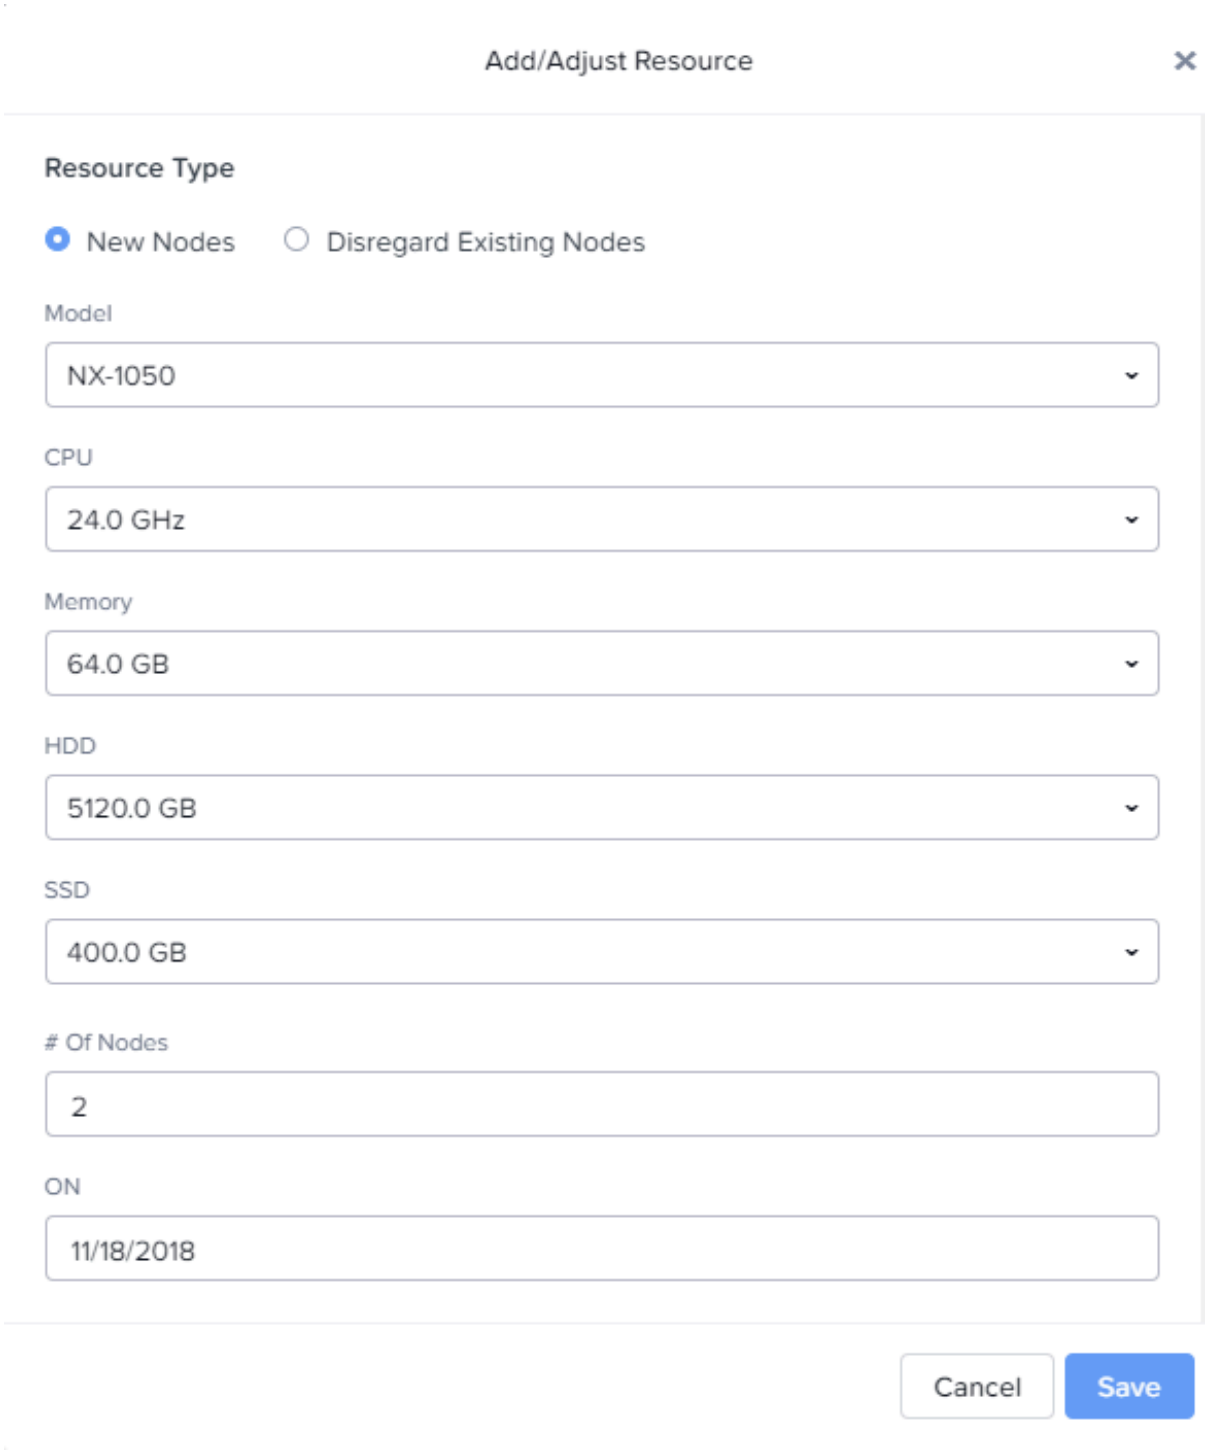

Figure 138: New Nodes Fields

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.

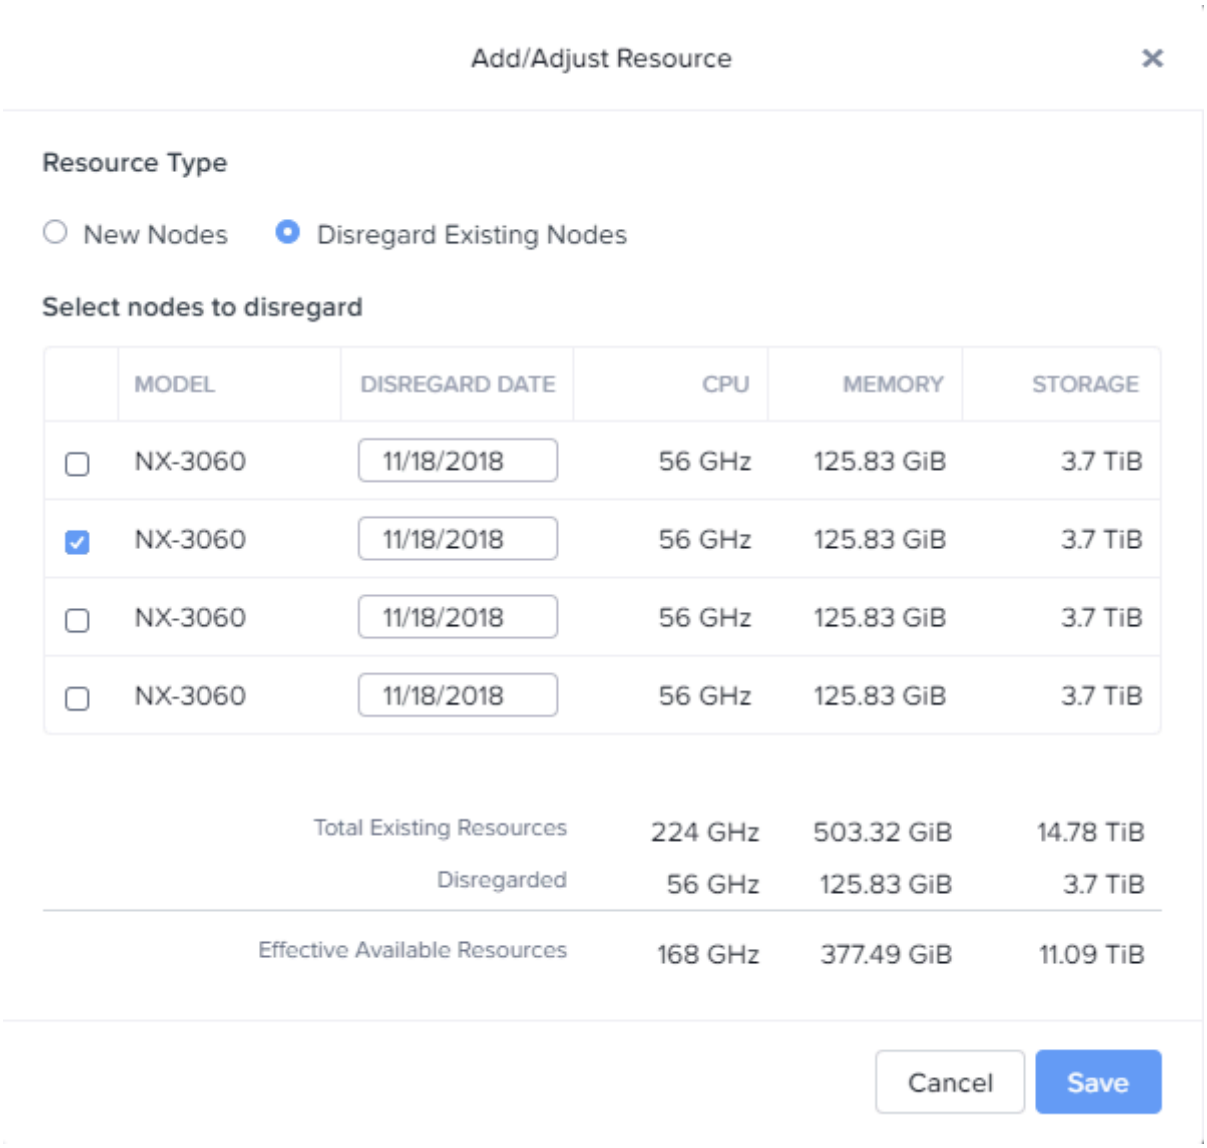

Figure 139: Disregard Existing Nodes Fields

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 177).

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.

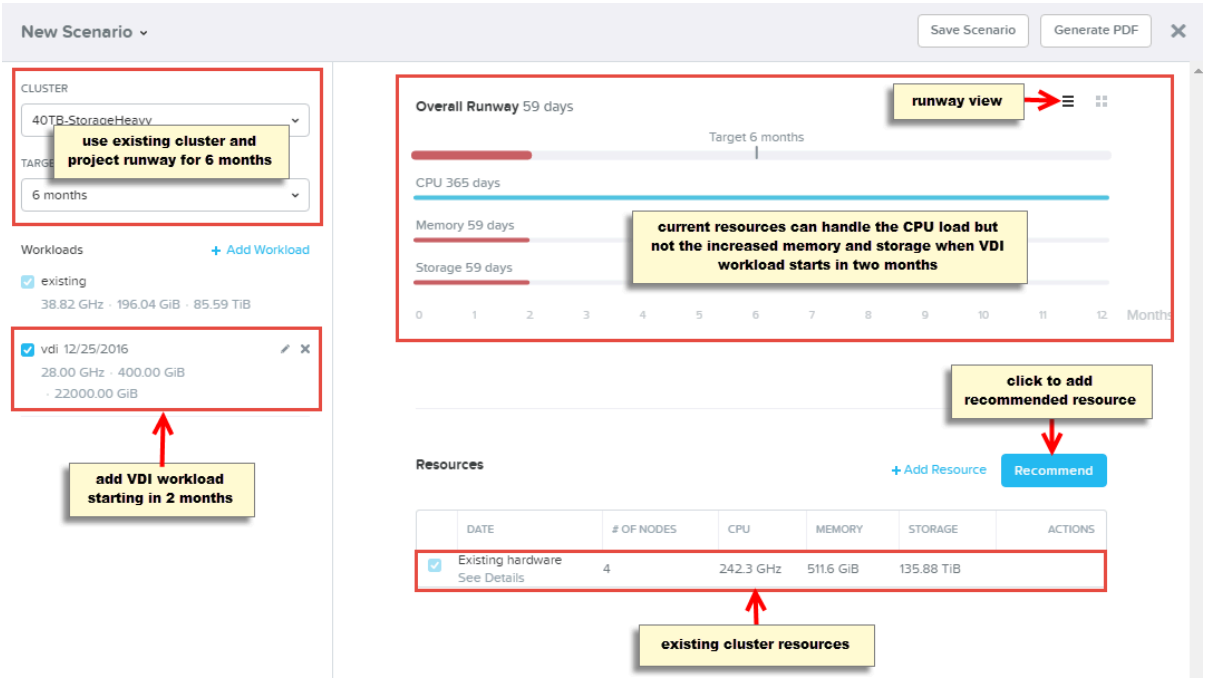

Figure 140: Scenario Screen (with added workload)

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.

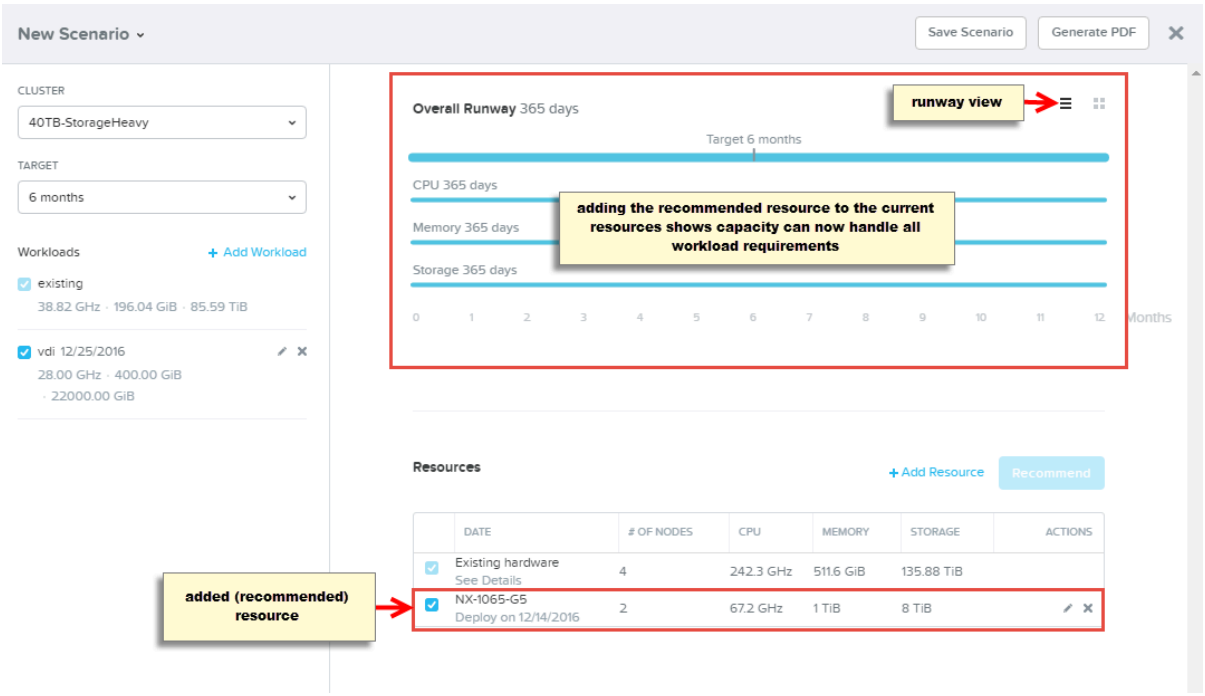

Figure 141: Scenario Screen (with added workload and recommended resource)

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.

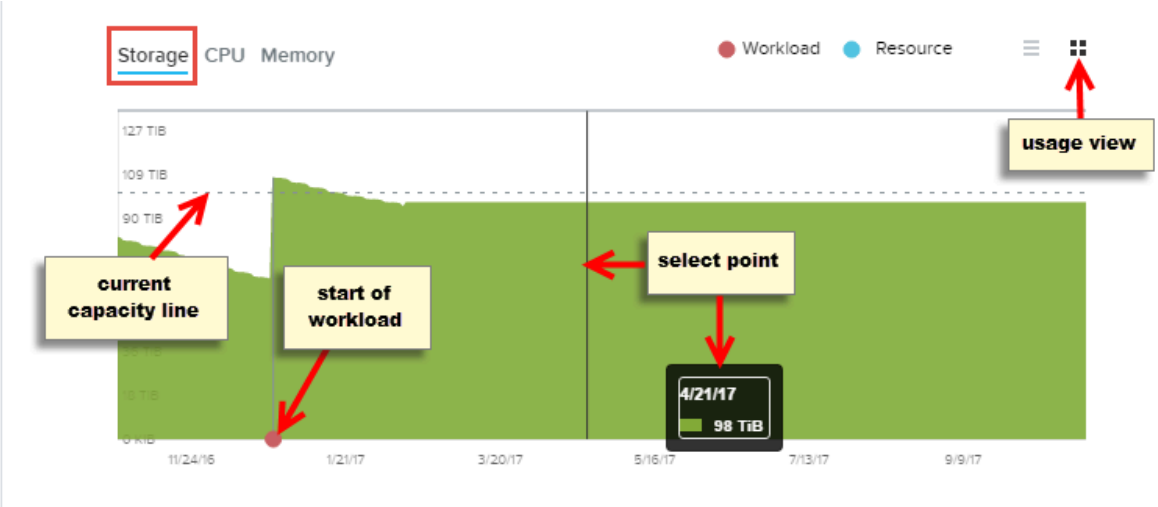

Figure 142: Storage Usage View

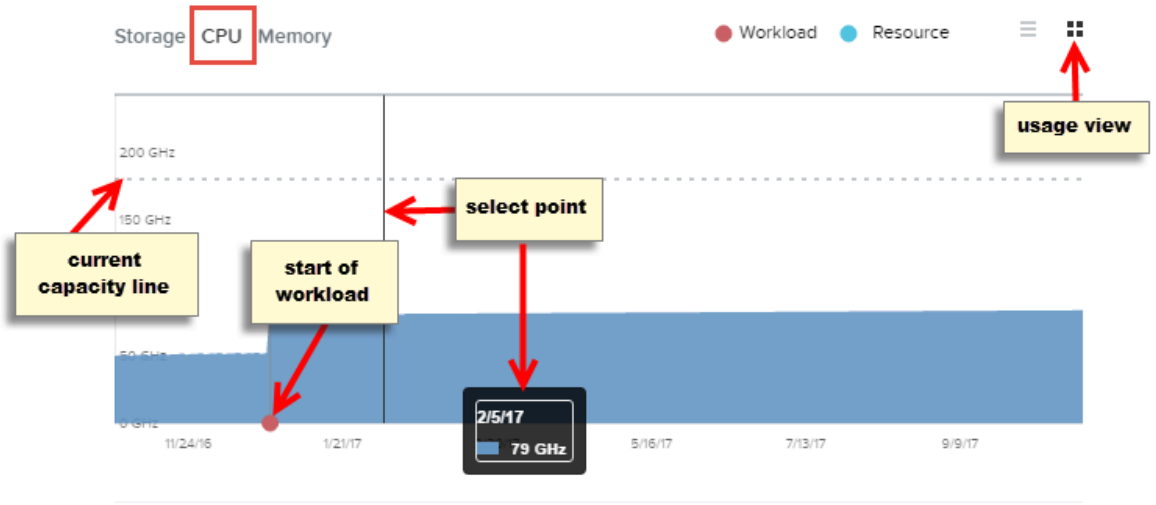

Figure 143: CPU Usage View

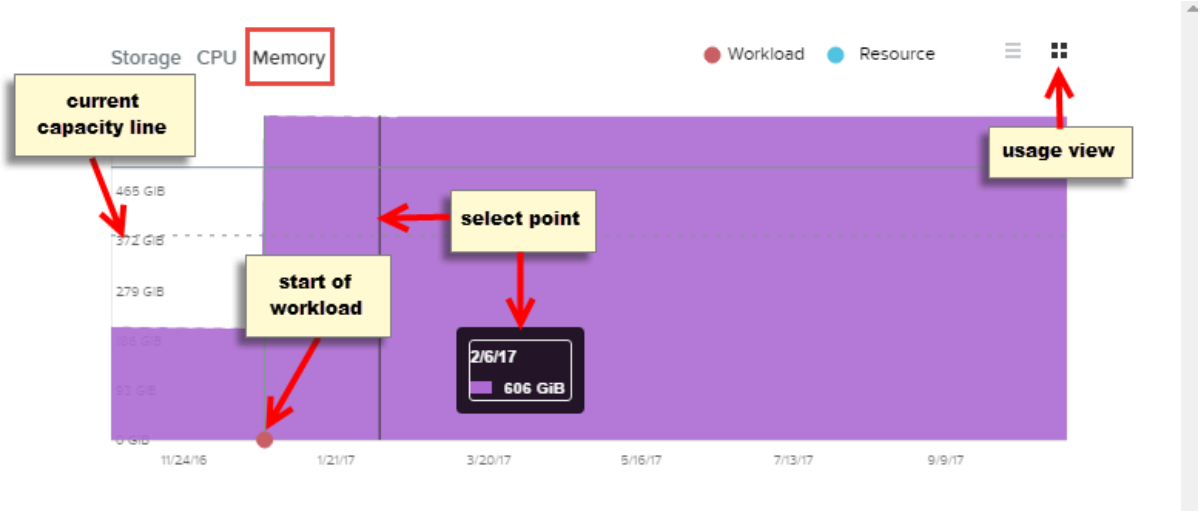

Figure 144: Memory Usage View

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 177), 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 179 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 176) 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.

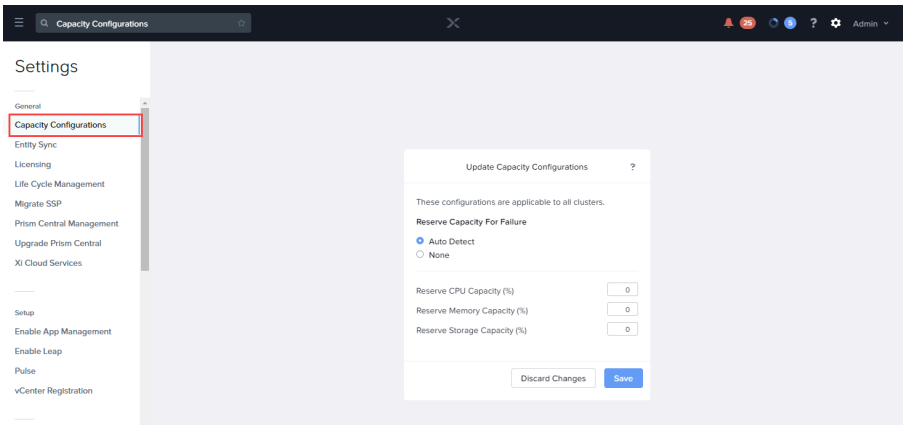

Figure 145: Update Capacity Configurations Window

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 74: 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 30 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.

Type	Description
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 30 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 30 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 30 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 47 and Adding Dashboard Widgets on page 61) and from the **Capacity** tab of a cluster details page (see Cluster Details View on page 121).

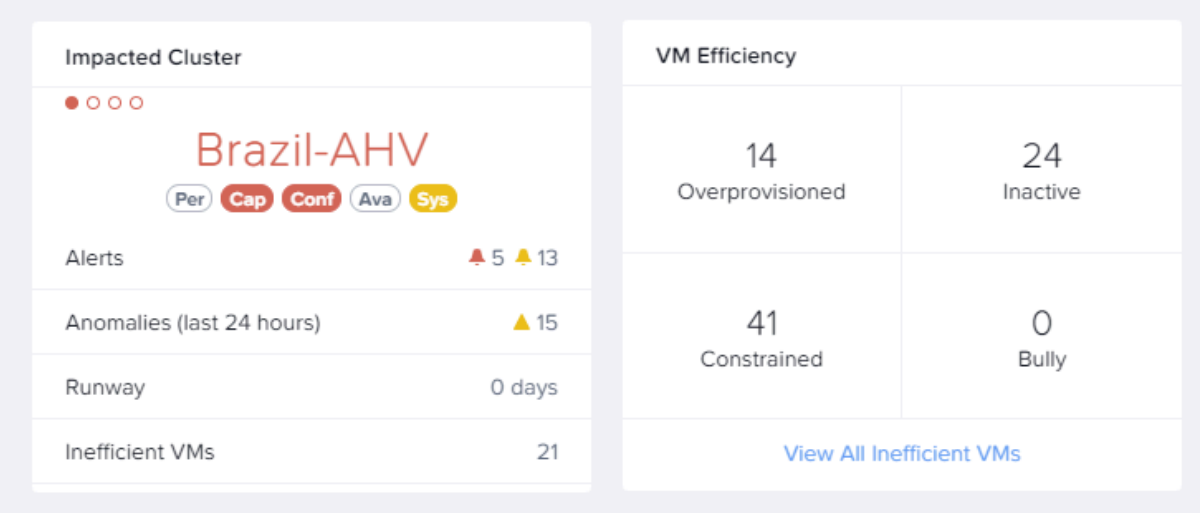

Figure 146: Impacted Cluster and VM Efficiency Widgets

You can also use a search filter from the VMs Summary View on page 72 to identify inefficient VMs.

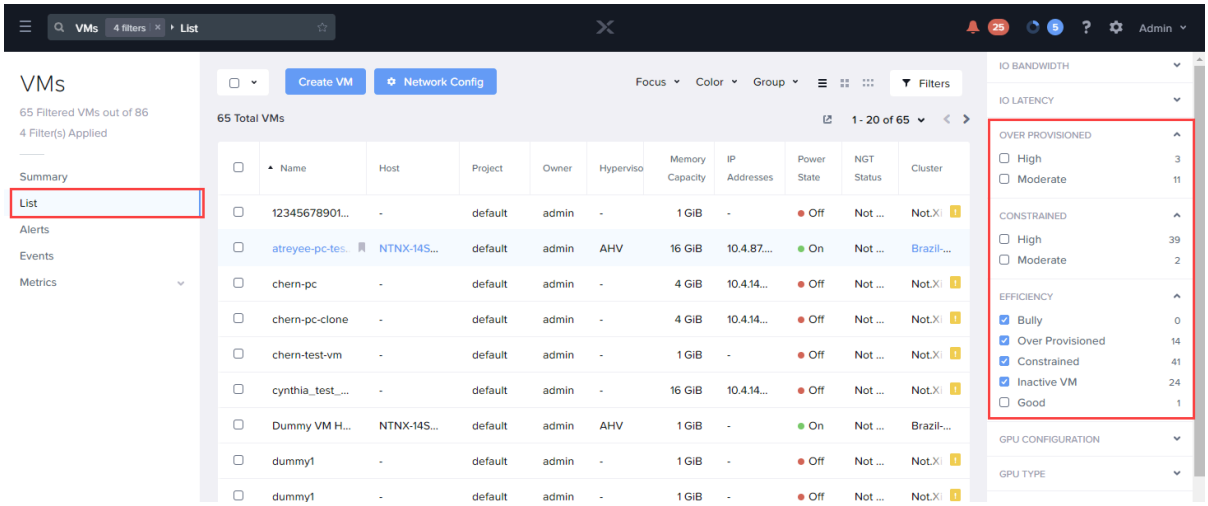

Figure 147: Filtering VMs

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 7 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 221). Anomalies appears as outliers in the behavioral anomaly event details screen and the VM details Metrics tab.

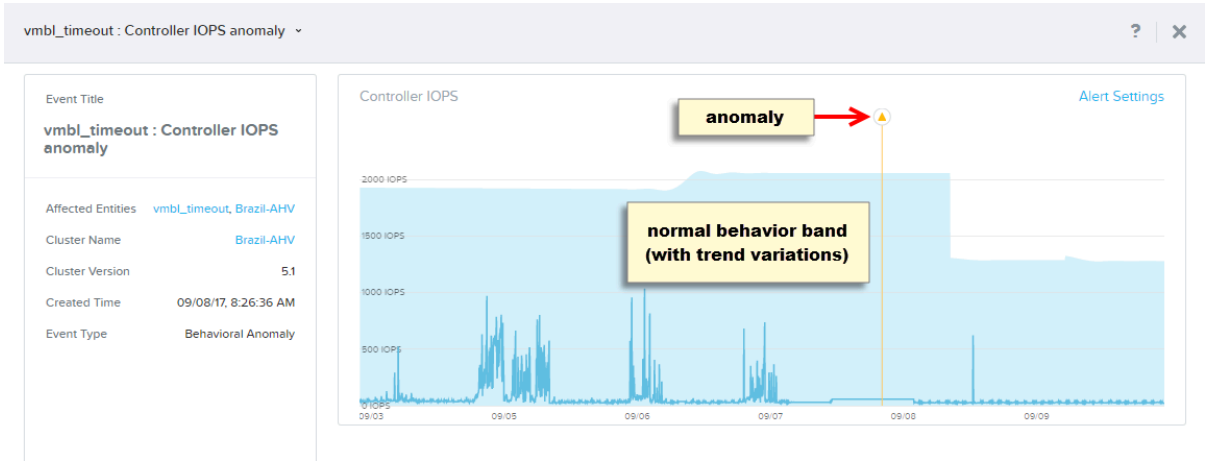

Figure 148: Anomaly Event Details Screen

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 229 for instructions on how to create a custom anomaly-based alert.

Cluster Efficiency Reports

One of the default reports (see Reports View on page 159) 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.

6

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 202).
- You can create custom entity performance charts (see Creating an Entity Chart on page 204).
- You can create custom metric performance charts (see Creating a Metric Chart on page 206).

> 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 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 12). The Analysis dashboard includes three sections.

- Chart definitions. The pane on the left lists the charts that can be run. No charts are provided by default, but you can create any number of charts. A chart defines the metrics to monitor. There are two types of charts, metric and entity. A metric chart monitors one or more entities for a single metric. An entity chart monitors one or more metrics for a single entity.

  > Note: You can change the color assigned to a metric or entity by clicking that color box in the chart (left pane) and then selecting a different color from the displayed palette.

- Chart monitors. When a chart definition is checked, the monitor appears in the middle pane. An Alerts monitor always appears first. The remaining monitors are determined by which charts are checked in the left pane. You can customize the display by selecting a time interval from the **Range** drop-down (above the charts) and then refining the monitored period by moving the time interval end points to the desired length.
- Alerts. Any alerts that occur during the interval specified by the time line in the middle pane appear in the pane on the right.

The following figure is a sample view, and the table describes each field in this view. Some fields can include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current.

> Note: See Understanding Displayed Statistics on page 25 for information about how the metrics are measured.

Figure 149: Analysis Dashboard

Table 75: Analysis Screen Fields

Name	Description
Charts	Displays the set of defined charts. Check the box next to a chart name to run that chart in the middle pane. The chart monitor appears in the middle pane shortly after checking the box. Uncheck the box to stop that monitor and remove it from the middle pane. To edit a chart definition, click the pencil icon to the right of the name. This opens the edit chart window, which is the same as the new chart window except for the title. To delete a chart, click the cross icon on the right.
New Metric Chart	Select **New** > **New Metric Chart** to create a chart that tracks one or more entities for a single metric (see Creating a Metric Chart on page 206).
New Entity Chart	**New** > **New Entity Chart** to create a chart that tracks one or more metrics for a single entity (see Creating an Entity Chart on page 204).
(range time line and monitor period)	Displays a time line that sets the duration for the monitor displays. To set the time interval, select the time period (3 hour, 6 hour, 1 day, 1 week, 1 month, 3months, 6 months) from the **Range** field pull-down menu (far right of time line). To customize the monitor period, drag the time line end points to the desired times on the time line.

Name	Description
Alerts Monitor	Displays a monitor of alert messages that were generated during the time interval. Alerts are tracked by a moving histogram with each bar indicating the number of messages generated during that time. The message types are color coded in the histogram bars (critical alert = red, warning alert = yellow, informational alert = gray).
(defined chart monitors)	Displays monitors for any enabled (checked) charts. (In the figure example, three charts are enabled.) You can edit or delete the chart data by clicking on the chart header to display an edit/ delete drop-down menu. You can export the chart data by clicking the download icon. This displays a drop-down menu (below) to save the data in CSV or JSON format. It also includes a chart link option that displays the URL to that chart, which you can copy to a clipboard and use to import the chart.
Alerts	Displays the alert messages that occurred during the time interval. Clicking the **details** link opens the details screen for that alert (see Alert Details on page 217).

Creating an Entity Chart

About this task

An entity chart monitors the performance of one or more metrics for a single entity. To create an entity chart definition, do the following:

Procedure

1. In the Analysis dashboard, select**New** > **New Entity Chart** at the top of the Charts column (left).
   The New Entity Chart dialog box appears.

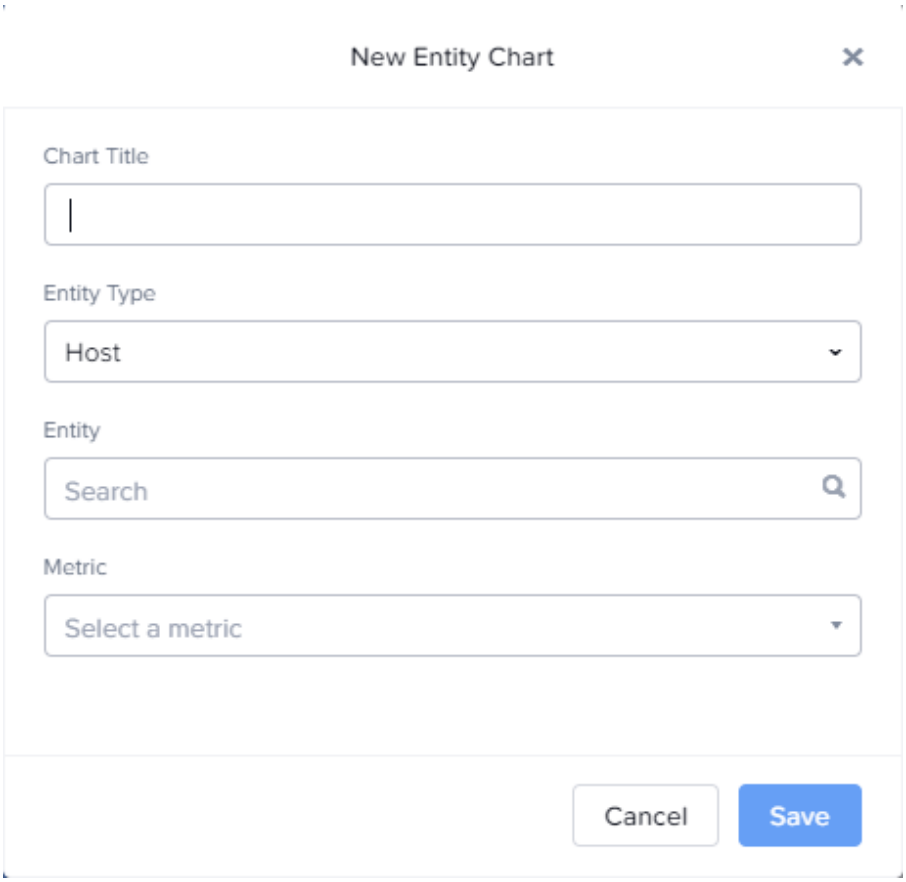

Figure 150: New Entity Chart Window

2. Do the following in the indicated fields:

   a. **Chart Title**: Enter a title for this chart.

   b. **Entity type**: Select an entity from the pull-down list.

      The entity types can include host, disk, storage pool, storage container, virtual machine, remote site, protection domain, replication link, virtual disk, and cluster.

   c. **Entity**: Enter the name of the target entity.

      As you enter characters in this field, it displays a list of matching entries of that entity type. Click the name when it appears in the search list. The list spans the registered clusters.

   d. **Metric**: Select a metric from the pull-down list. (Repeat to include additional metrics.)

      For descriptions of the available metrics, see Chart Metrics on page 207.

3. When all the field entries are correct, click the **Save** button.

   The Analysis dashboard reappears with the new chart appearing in the list of charts on the left of the screen.

Creating a Metric Chart

About this task

A metric chart monitors the performance of a single metric on one or more entities. To create a metric chart definition, do the following:

Procedure

1. In the Analysis dashboard, select**New** > **New Metric Chart** at the top of the Charts column (left).
   The New Metric Chart dialog box appears.

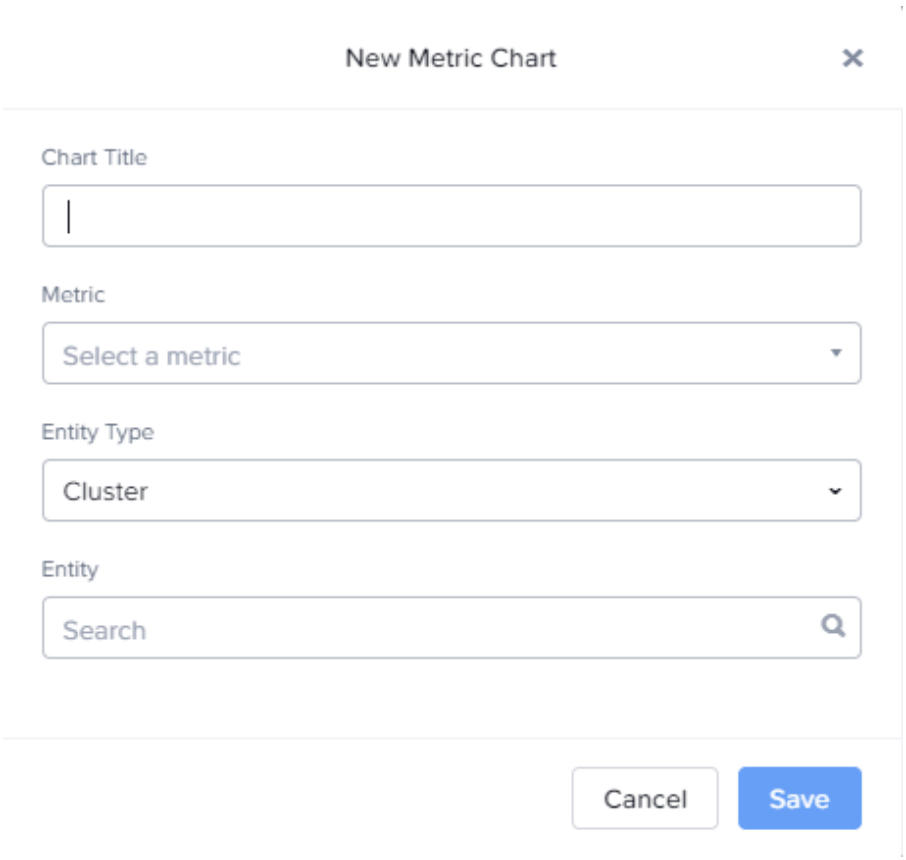

Figure 151: New Metric Chart Window

2. Do the following in the indicated fields:

   a. **Chart Title**: Enter a title for this chart.

   b. **Metric**: Select a metric to monitor from the pull-down list.

      For descriptions of the available metrics, see Chart Metrics on page 207.

   c. **Entity Type**: Select an entity type from the pull-down list.

      The entity type options depend on the metric selected.

   d. **Entity**: Enter the name of the target entity.

      As you enter characters in this field, it displays a list of matches of the entity type. Click the name when it appears in the search list. The list spans the registered clusters.

      Note: Repeat the last two steps to add additional entities.

3. When all the field entries are correct, click the **Save** button.

   The Analysis dashboard reappears with the new chart appearing in the list of charts on the left of the screen.

Chart Metrics

These metrics can be added to charts.

Metric	Description
Content Cache Hit Rate (%)	Content cache hits over all lookups. ID: `CONTENT_CACHE_HIT_PPM`
Content Cache Hits	Number of hits on the content cache. ID: `CONTENT_CACHE_NUM_HITS`
Content Cache Logical Memory Usage	Logical memory (in bytes) used to cache data without deduplication. ID: `CONTENT_CACHE_LOGICAL_MEMORY_USAGE_BYTES`
Content Cache Logical SSD Usage	Logical SSD memory (in bytes) used to cache data without deduplication. ID: `CONTENT_CACHE_LOGICAL_SSD_USAGE_BYTES`
Content Cache Lookups	Number of lookups on the content cache. ID: `CONTENT_CACHE_NUM_LOOKUPS`
Content Cache Memory Saved	Memory (in bytes) saved due to content cache deduplication. ID: `CONTENT_CACHE_SAVED_MEMORY_USAGE_BYTES`
Content Cache Physical Memory Usage	Real memory (in bytes) used to cache data by the content cache. ID: `CONTENT_CACHE_PHYSICAL_MEMORY_USAGE_BYTES`
Content Cache Reference Count	Average number of content cache references. ID: `CONTENT_CACHE_NUM_DEDUP_REF_COUNT_PPH`

Metric	Description
Content Cache SSD Usage	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	SSD usage (in bytes) saved due to content cache deduplication. ID: `CONTENT_CACHE_SAVED_SSD_USAGE_BYTES`
Deduplication Fingerprints Cleared	Number of written bytes for which fingerprints have been cleared. ID: `DEDUP_FINGERPRINT_CLEARED_BYTES`
Deduplication Fingerprints Written	Number of written bytes for which fingerprints have been added. ID: `DEDUP_FINGERPRINT_ADDED_BYTES`
Disk I/O Bandwidth	Data transferred per second in KB/second from disk. ID: `STATS_BANDWIDTH`
Disk I/O Bandwidth - Read	Read data transferred per second in KB/second from disk. ID: `STATS_READ_BANDWIDTH`
Disk I/O Bandwidth - Write	Write data transferred per second in KB/second from disk. ID: `STATS_WRITE_BANDWIDTH`
Disk I/O Latency	I/O latency in milliseconds from disk. ID: `STATS_AVG_IO_LATENCY`
Disk IOPS	Input/Output operations per second from disk. ID: `STATS_NUM_IOPS`
Disk IOPS - Read	Input/Output read operations per second from disk. ID: `STATS_NUM_READ_IOPS`
Disk IOPS - Write	Input/Output write operations per second from disk. ID: `STATS_NUM_WRITE_IOPS`
Hypervisor CPU Usage (%)	Percent of CPU used by the hypervisor. ID: `STATS_HYP_CPU_USAGE`
Hypervisor Memory Usage (%)	Percent of memory used by the hypervisor. ID: `STATS_HYP_MEMORY_USAGE`
Hypervisor CPU Ready Time (%)	CPU ready time percentage. ID: `STATS_HYP_CPU_READY_TIME`

Metric	Description
Hypervisor I/O Bandwidth	Data transferred per second in KB/second from Hypervisor. ID: `STATS_HYP_BANDWIDTH`
Hypervisor I/O Bandwidth - Read	Read data transferred per second in KB/second from Hypervisor. ID: `STATS_HYP_READ_BANDWIDTH`
Hypervisor I/O Bandwidth - Write	Write data transferred per second in KB/second from Hypervisor. ID: `STATS_HYP_WRITE_BANDWIDTH`
Hypervisor I/O Latency	I/O latency in milliseconds from Hypervisor. ID: `STATS_HYP_AVG_IO_LATENCY`
Hypervisor I/O Latency - Read	I/O read latency in milliseconds from Hypervisor. ID: `STATS_HYP_AVG_READ_IO_LATENCY`
Hypervisor I/O Latency - Write	I/O write latency in milliseconds from Hypervisor. ID: `STATS_HYP_AVG_WRITE_IO_LATENCY`
Hypervisor IOPS	Input/Output operations per second from Hypervisor. ID: `STATS_HYP_NUM_IOPS`
Hypervisor IOPS - Read	Input/Output read operations per second from Hypervisor. ID: `STATS_HYP_NUM_READ_IOPS`
Hypervisor IOPS - Write	Input/Output write operations per second from Hypervisor. ID: `STATS_HYP_NUM_WRITE_IOPS`
Logical Usage	Logical usage of storage (physical usage divided by replication factor). ID: `STATS_UNTRANSFORMED_USAGE`
Physical Usage	Actual usage of storage. ID: `STATS_TRANSFORMED_USAGE`
Read IOPS (%)	Percent of IOPS that are reads. ID: `STATS_READ_IO_PPM`
Replication Bandwidth - Received	Replication data received per second in KB/second ID: `STATS_REP_BW_RECEIVED`
Replication Bandwidth - Transmitted	Replication data transferred per second in KB/second ID: `STATS_REP_BW_TRANSFERRED`
Replication Bytes - Received	Number of bytes received. ID: `STATS_REP_NUM_RECEIVED_BYTES`

Metric	Description
Replication Bytes - Total Received	Total number of bytes received. ID: `STATS_REP_TOT_RECEIVED_BYTES`
Replication Bytes - Total Transmitted	Total number of bytes transmitted. ID: `STATS_REP_TOT_TRANSMITTED_BYTES`
Replication Bytes - Transmitted	Number of bytes transmitted. ID: `STATS_REP_NUM_TRANSMITTED_BYTES`
Storage Controller Bandwidth	Data transferred in KB/second from the Storage Controller. ID: `STATS_CONTROLLER_BANDWIDTH`
Storage Controller Bandwidth - Read	Read data transferred in KB/second from the Storage Controller. ID: `STATS_CONTROLLER_READ_BANDWIDTH`
Storage Controller Bandwidth - Write	Write data transferred in KB/second from the Storage Controller. ID: `STATS_CONTROLLER_WRITE_BANDWIDTH`
Storage Controller IOPS	Input/Output operations per second from the Storage Controller ID: `STATS_CONTROLLER_NUM_IOPS`
Storage Controller IOPS - Read	Input/Output read operations per second from the Storage Controller ID: `STATS_CONTROLLER_NUM_READ_IOPS`
Storage Controller IOPS - Read (%)	Percent of Storage Controller IOPS that are reads. ID: `STATS_CONTROLLER_READ_IO_PPM`
Storage Controller IOPS - Write	Input/Output write operations per second from the Storage Controller ID: `STATS_CONTROLLER_NUM_WRITE_IOPS`
Storage Controller IOPS - Write (%)	Percent of Storage Controller IOPS that are writes. ID: `STATS_CONTROLLER_WRITE_IO_PPM`
Storage Controller Latency	I/O latency in milliseconds from the Storage Controller. ID: `STATS_CONTROLLER_AVG_IO_LATENCY`
Storage Controller Latency - Read	Storage Controller read latency in milliseconds. ID: `STATS_CONTROLLER_AVG_READ_IO_LATENCY`
Storage Controller Latency - Write	Storage Controller write latency in milliseconds. ID: `STATS_CONTROLLER_AVG_WRITE_IO_LATENCY`

Metric	Description
Storage container own usage	Storage container's own usage + Reserved (not used). ID: `NEW_CONTAINER_OWN_USAGE_LOGICAL`
Swap In Rate	Rate of data being swapped in. ID: `STATS_HYP_SWAP_IN_RATE`
Swap Out Rate	Rate of data being swapped out. ID: `STATS_HYP_SWAP_OUT_RATE`
Virtual NIC bytes received packets with error	Virtual NIC bytes received packets with error. ID: `STATS_NETWORK_ERROR_RECEIVED_PACKETS`
Virtual NIC bytes received rate	Virtual NIC bytes received rate in kbps. ID: `STATS_NETWORK_RECEIVED_RATE`
Virtual NIC bytes transmitted rate	Virtual NIC bytes transmitted rate in kbps. ID: `STATS_NETWORK_TRANSMITTED_RATE`
Virtual NIC dropped transmitted packets	Number of dropped transmitted packets by the VNIC. ID: `STATS_NETWORK_DROPPED_TRANSMITTED_PACKETS`
Virtual NIC received packets dropped	Number of receive packets dropped by the VNIC. ID: `STATS_NETWORK_DROPPED_RECEIVED_PACKETS`
Write IOPS (%)	Percent of IOPS that are writes. ID: `STATS_WRITE_IO_PPM`

7

SERVICES ENABLEMENT

You can enable selected services through Prism Central.

- Enabling Calm on page 212
- Enabling Karbon on page 213
- Enabling Objects on page 213

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 9) and then select **Services** > **Calm** from the entities menu (see Entities Menu on page 12).
   - » Select **Enable App Management** from the Settings menu (see Settings Menu on page 14).
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 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 menu button in the main menu (see Main Menu on page 9) and then select **Services** > **Karbon** from the entities menu (see Entities Menu on page 12).
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 9) and then select **Services** > **Objects** from the entities menu (see Entities Menu on page 12).
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.

8

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 12), allow you to view and monitor alert and event activity across the registered clusters.
  - The Alerts Summary View on page 214 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 217)
  - The Events Summary View on page 220 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 221).
- 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 summary view displays a list of alert messages across the registered clusters. From this view you can customize the order and scope of listed alerts, display the details for a specific alert, and configure alert policies. The following figure is a sample view, and the table describes each field in this view. To access the Alerts dashboard, select **Activities** > **Alerts** from the entities menu (see Entities Menu on page 12).

> Note: This section describes the information and options that appear in the Alerts dashboard. See Entity Exploring on page 66 for instructions on how to view and organize that information in a variety of ways. See Alerts/Health checks on page 250 for a list of alert messages and corrective actions.

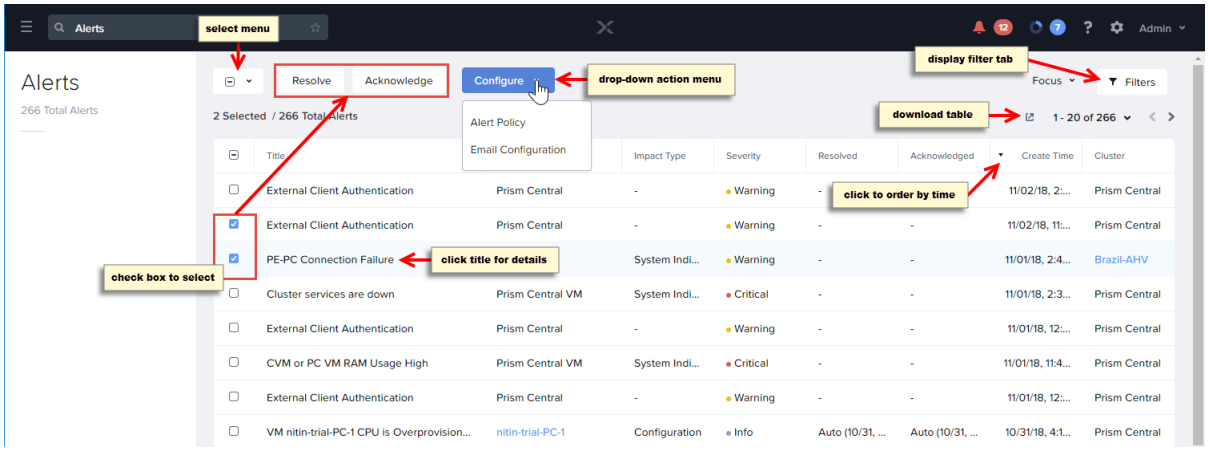

Figure 152: Alerts Dashboard

- To order the alerts chronologically, click the **Create Time** column header, which toggles the alert list latest-to-earliest or earliest-to-latest.
- To download the table of alerts in CSV format, click the download icon.
- To display the details page for an alert, click the alert title (see Alert Details on page 217).
- To configure alert policies, do one or both of the following:
  - Select the **Configure** > **Alert Policy** button to specify what events should generate an alert (see Configuring Alert Policies on page 226).
  - Select the **Configure** > **Email Configuration** button to enable alert emails and specify email addresses to which alerts should be sent (see Configuring Alert Emails on page 223).
- 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.
  - 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*
- 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 (far left above the list of alerts), 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 76: Alerts View 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. **Informational** An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational

Parameter	Description	Values
Resolved	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 "Auto (*date_time*)" entry indicates this alert was resolved 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 "By *user* (*date_time*)" entry indicates this alert was resolved manually by the specified user at that date and time. To manually resolve an alert, check the box in the far left column and then select the**Resolve** button (left of screen above the search field). The field value changes from blank to the user name and date/time the alert was resolved.	(blank, Auto (*date_time*), By *user* (*date_time*)
Acknowledged	Indicates whether the alert has been acknowledged. Acknowledging an alert means you recognize the error exists (no more reminders for this condition), but the alert status remains. The options (blank, auto, or user) are the same as for Resolved. To manually acknowledge an alert, check the box in the far left column and then click the **Acknowledge** button (left of screen above the search field). The field value changes from blank to the user name and date/time the alert was acknowledged.	(blank, Auto (*date_time*), By *user* (*date_time*)
Create Time	Displays the date and time when the alert occurred.	(date and time)
Cluster	Displays the name of the cluster in which the alert was issued.	(cluster name)

Alert Details

Clicking on an alert message in the dashboard (see Alerts Summary View on page 214) 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 77: Alert Details Fields

Parameter	Description	Values
Alert Title	Displays the title of the alert.	(alert title)
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** A "critical" alert is one that requires immediate attention, such as a failed Controller VM. **Warning** A "warning" alert is one that might need attention soon, such as an issue that could lead to a performance problem. **Informational** An "informational" alert highlights a condition to be aware of, for example, a reminder that the support tunnel is enabled.	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 121).	(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

Parameter	Description	Values
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
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)

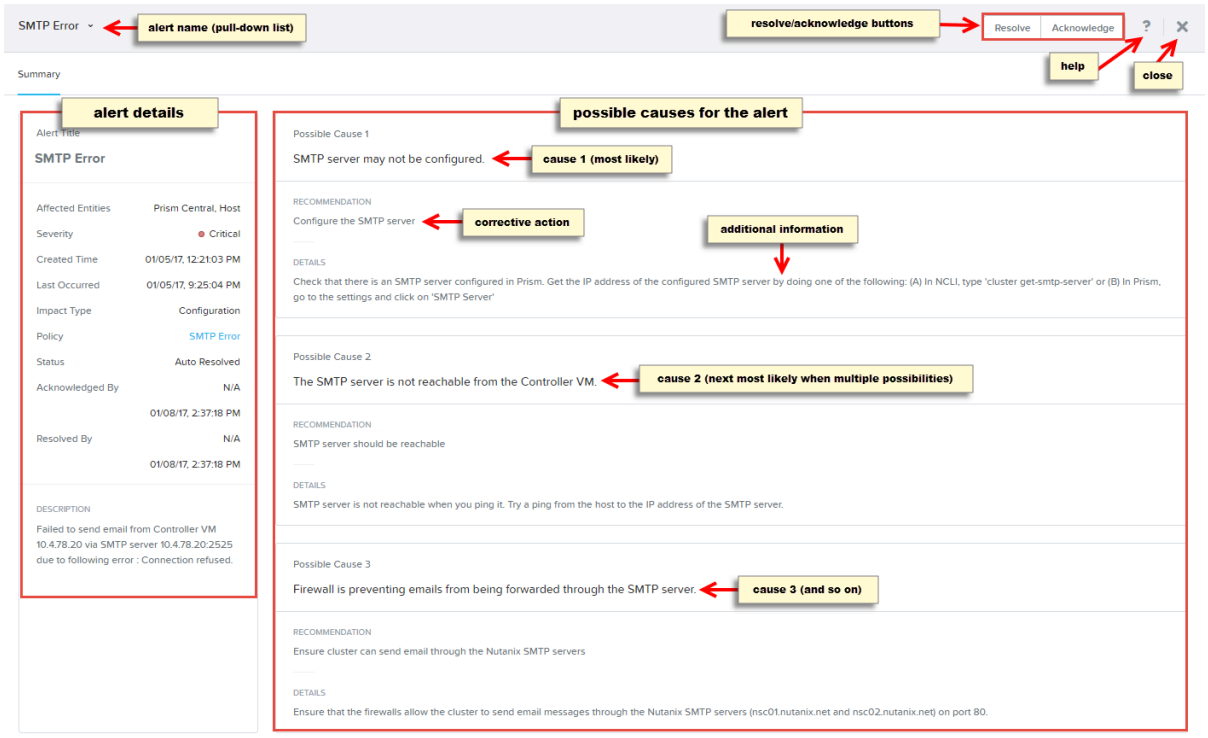

Figure 153: Alert Details Screen

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 12).

> Note: This section describes the information and options that appear in the Events dashboard. See Entity Exploring on page 66 for instructions on how to view and organize that information in a variety of ways.

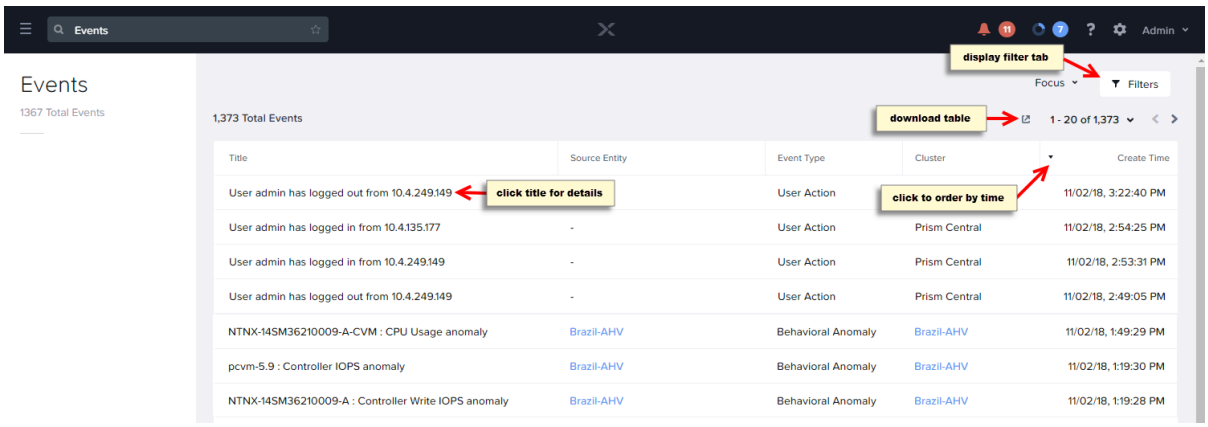

Figure 154: Events Dashboard

- To display the details page for an event, click the event title (see Event Details on page 221).
- 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 download icon.
- 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 78: 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 221).	(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 197), 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 220) 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.

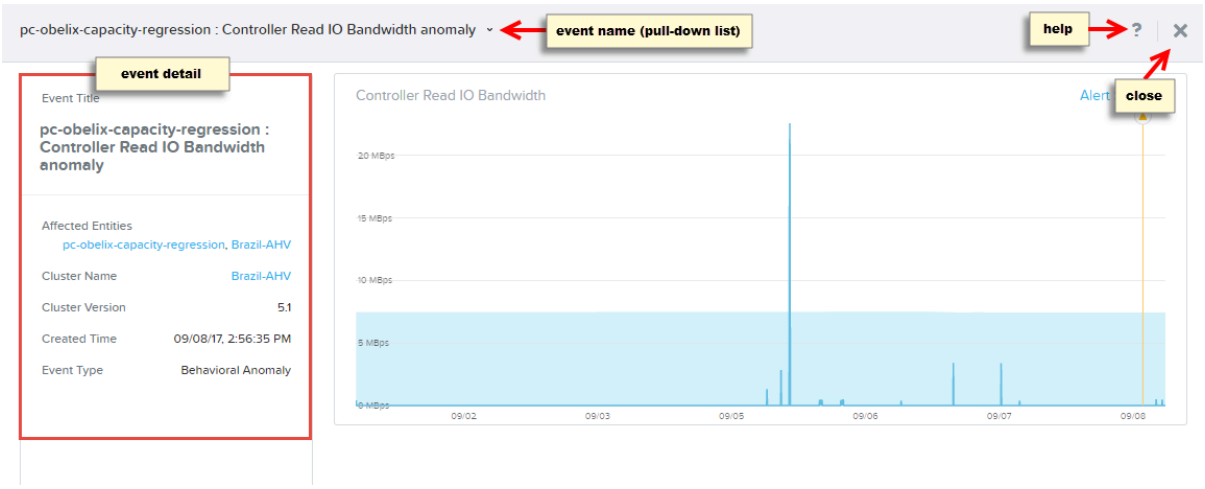

Figure 155: Event Details Screen

- 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 79: 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 121).	(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 282).

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 14).

   » 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 **Email Every Alert** box to send an email whenever the event occurs.
      - Check the **Email Daily Digest** box to send a cumulative (24 hour) list of alerts once a day.

      You can check one or both boxes. If neither box is checked, no alert emails are sent by Prism Central.

   b. Check the **Nutanix Support** box to send alert notification to Nutanix customer support.

   c. To send alert notifications to others, enter their email addresses in a comma separate list in the field below the **Nutanix Support** box.

   d. Click the **Apply** button to apply the changes and continue the configuration (or click the **Save** button to apply the changes and close the window).

   The **Connection Status** section displays mail transport status information. In this example, an SMTP server is configured for mail service. Clicking the **Details** link opens the SMTP server settings window (see Configuring an SMTP Server on page 282).

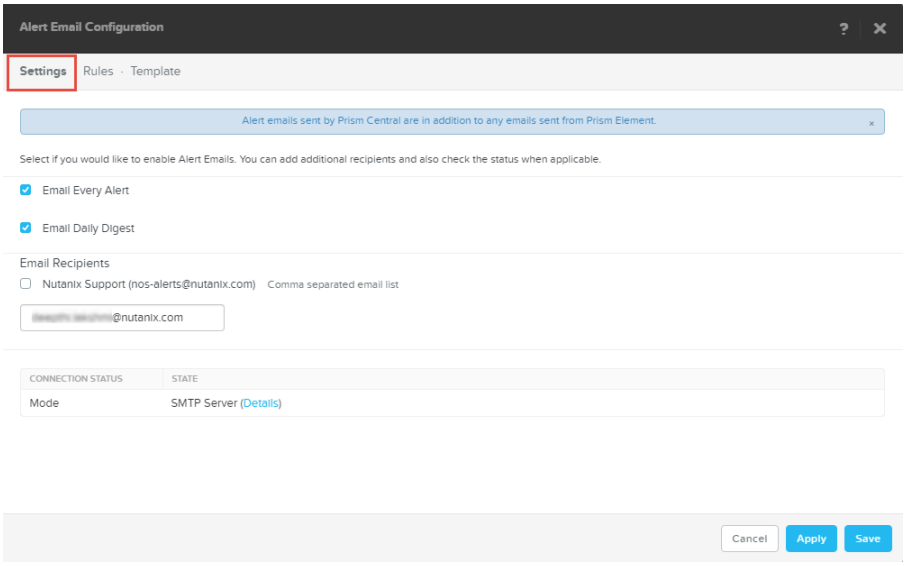

Figure 156: Alert Email Configuration Window (Settings tab)

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:

      - **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 Alert Contains**: Select one or more of the clusters from the pull-down list (*cluster_name*, All) in the top line and then enter a key phrase or word in the following line 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**, then specify the phrase **low memory**.

b. Specify who should receive the alert email.

- **Email Addresses**: Enter recipient email addresses as a comma separated list in the box.
- **Global Recipient List**: Click this box to add everyone on the global recipient list. This is in addition to any users listed in the **Email Addresses** field.

c. Click the **Apply** button to apply the rule.

d. Repeat these steps to apply additional custom rules.

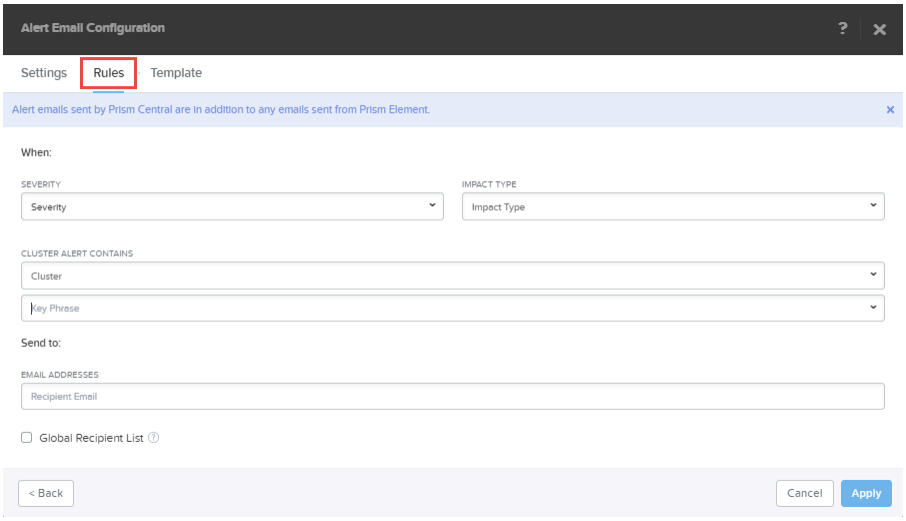

Figure 157: Alert Email Configuration Window (Rules tab)

4. To create a template for the email messages, click the **Template** tab and do the following:

   a. In the **Prepend Content to the Email Subject** field, enter the desired text.

      This text will appear at the beginning of the subject field in each alert email. If the field is left blank, no prepended text will appear in the subject.

   b. In the **Append Content to the Email Body** field, enter the desired text.

      This text will appear at the end of the message body in each alert email. If the field is left blank, no appended text will appear in the message body.

   c. Click the **Save** button to apply the changes and close the window (or click the **Apply** button to apply the changes and continue the configuration).

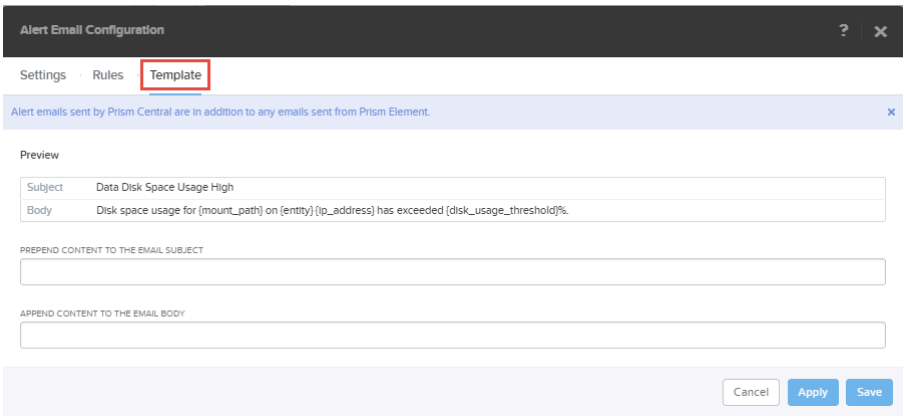

Figure 158: Alert Email Configuration Window (Template tab)

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 two categories of alert policies:

Procedure

- System Created Alert Policies. You can view and customize policies for any of the system alerts (see Modifying System Alert Policies on page 227).
- User Created Alert Policies. In addition to the system alerts, you have the option to create custom alert policies (see Adding Custom Alert Policies on page 229).

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. Open the Alert Policies page using one of the following methods:

   » Click the gear icon and select **Alert Policies** from the **Settings** menu (see Settings Menu on page 14).

   » In the **Alerts** dashboard, click the **Configure** button and select **Alert Policy** from the pull-down list.

   In the Alert Policies page, click the **System Created** tab (default view) to display a list of system alert events. Enter a string in the search box at the top or use the scroll bar on the right to search through the list. Each line includes the alert ID number (Axxxx), title, impact type, entity type to which the alert applies, global rule (default or overridden), enabled status (enabled or disabled), auto resolved status (enabled, disabled, or not applicable), number of exceptions (0 for default or positive integer for overridden), and last update time.

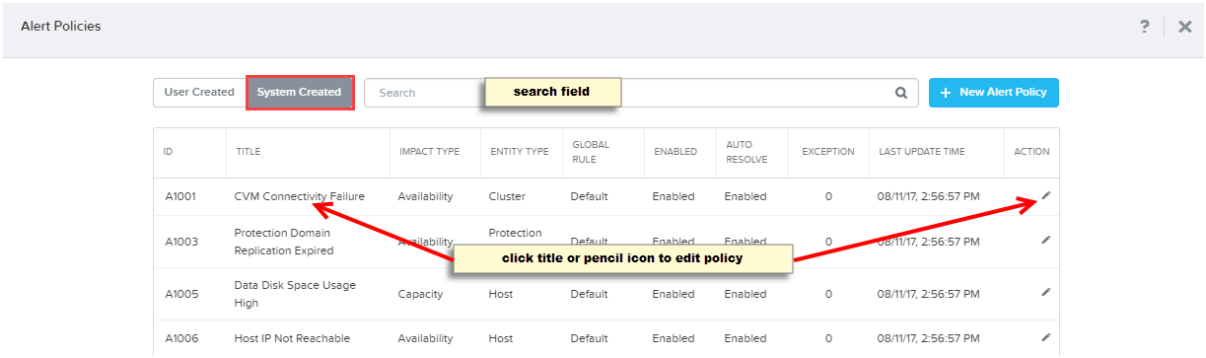

Figure 159: Alert Policies Page: System Created Tab

2. To modify the global rule for a system alert, click the title or pencil icon for that alert. The Update Policy window appears.

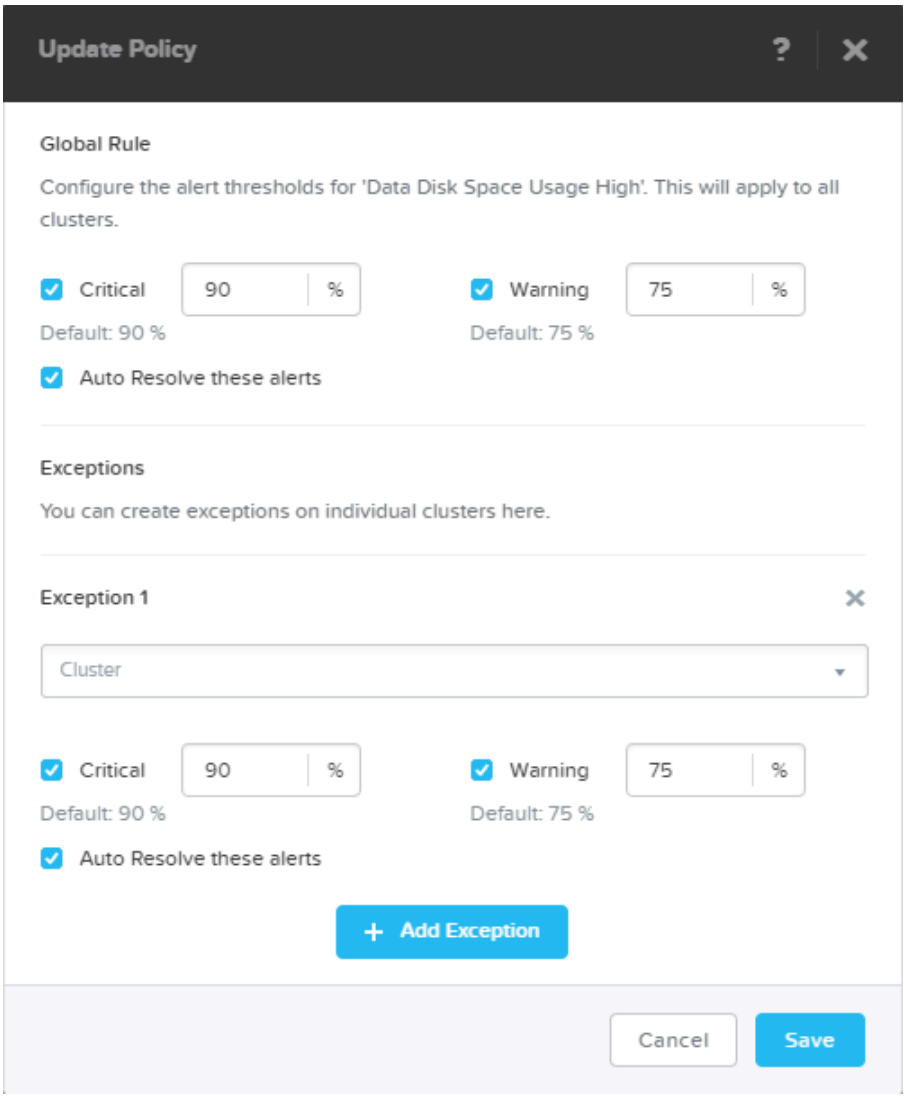

Figure 160: Update Policy window

3. 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 cluster to exempt from the pull-down list.

   Excluded clusters use the default rule for this alert. Repeat for each cluster you want to exclude.

d. Click the **Save** button to save the changes and return to the Alert Policies window.

4. Click the **Close** button to close the Alert Policies window.

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 235.

Creating Custom Alert Policies

About this task

Perform the following procedure to create a custom alert policy.

Procedure

1. Do one of the following:

   » Click the gear icon and select **Alert Policies** from the **Settings** menu (see Settings Menu on page 14).

   » In the **Alerts** dashboard, click the **Configure** button and select **Alert Policy** from the pull-down list.

   » In the VM or cluster details page (see VM Details View on page 81 or Cluster Details View on page 121), select **Alert policy** from the pull-down menu.

2. In the **Alert Policies** window, click the **User Created** tab to view the existing user created policies (if any).

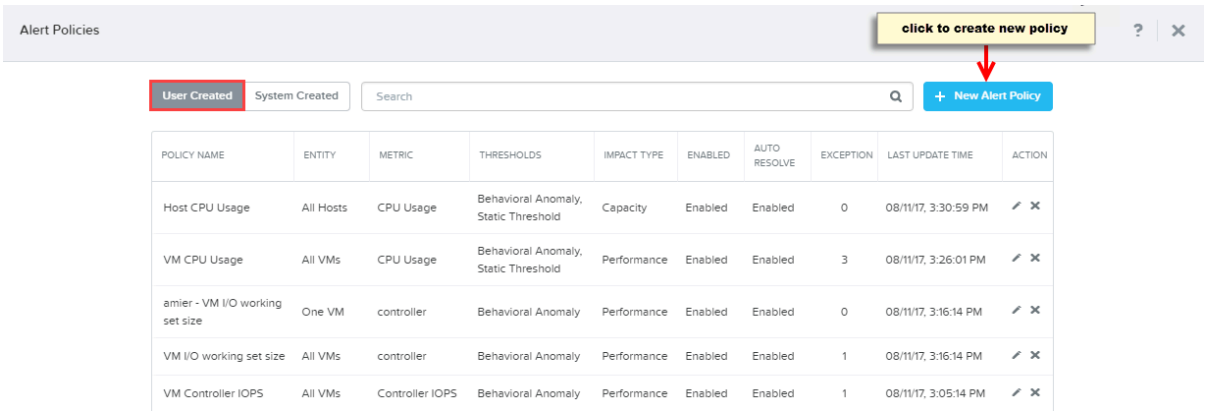

Figure 161: Alert Policies Page: System Created Tab

3. Click **+ New Alert Policy** to add a new custom alert policy.
The **Create Alert Policy** dialog box appears.

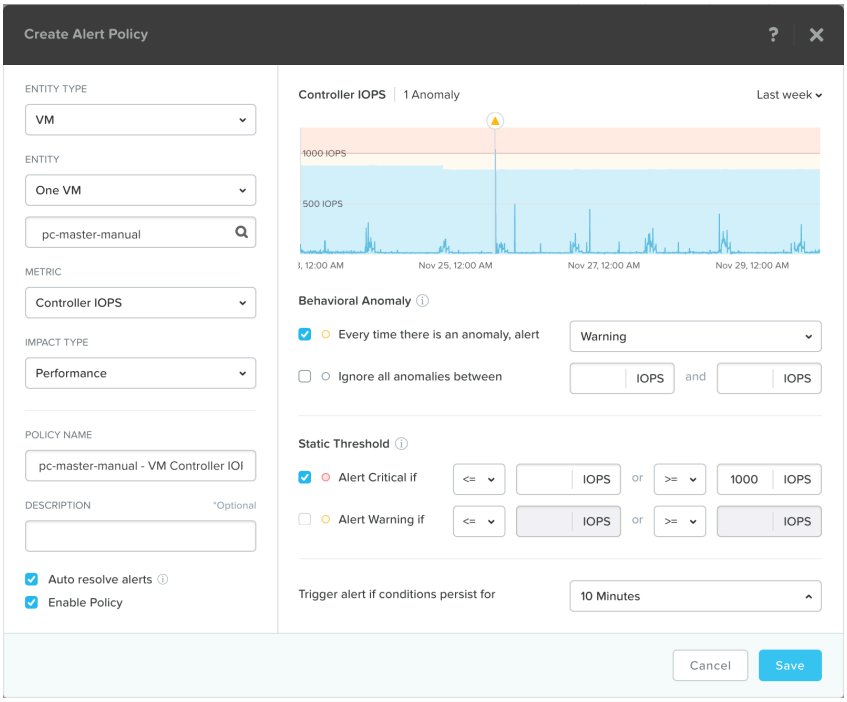

Figure 162: Create Alert Policy window

4. 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. **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**.

e. **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.

f. **Description** (optional): Enter a description for the policy.

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.

5. 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 197 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.

6. 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 229), 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:

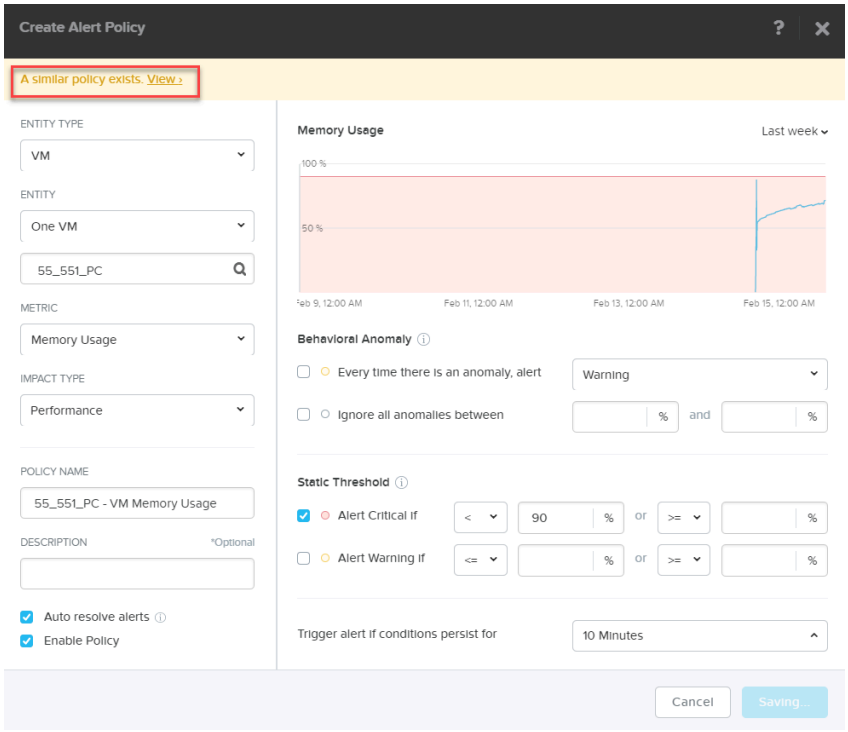

Figure 163: Overlapping Policy

Click **View** to display the overlapping policies and click **Save** to save the overlapping policy.

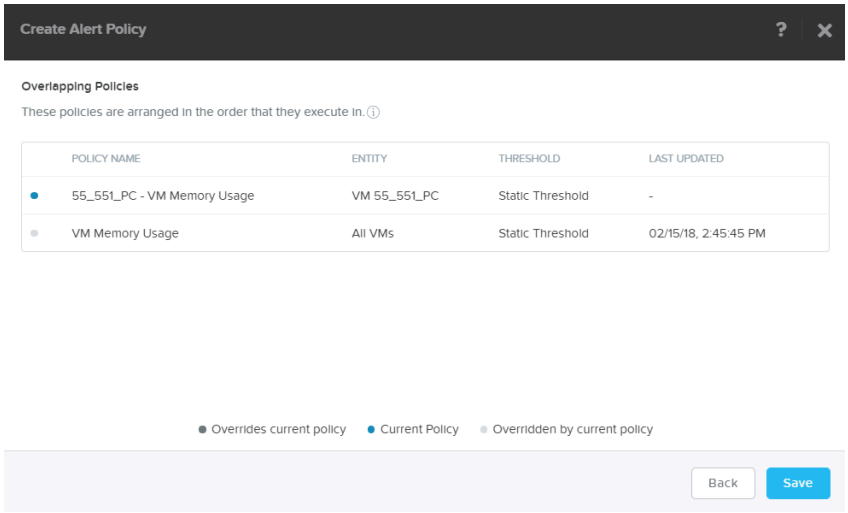

Figure 164: Overlapping Policies List

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

Perform the following procedure to modify a custom alert policy.

Procedure

1. Do one of the following:

   » In the **Alerts** dashboard, click the **Configure** button and select **Alert Policy** from the pull-down list.

   » Click the gear icon and select **Alert Policies** from the **Settings** menu (see Settings Menu on page 14).

2. In the **Alert Policies** window, click the **User Created** tab.

3. To modify a policy, click the pencil icon under **Actions** next to the policy you want to modify.

   The **Update Alert Policy** dialog box appears. Options in this dialog box are similar to the **Create Alert Policy** dialog box (see Creating Custom Alert Policies on page 229). The Overlapping Polices tab is displayed only if the policy that you want to modify has overlapping policies.

4. Make the changes and click **Save**.

Deleting Custom Alert Policies

About this task

Perform the following procedure to delete a custom (user-created) alert policy.

Procedure

1. Do one of the following:
   - » In the **Alerts** dashboard, click the **Configure** button and select **Alert Policy** from the pull-down list.
   - » Click the gear icon and select **Alert Policies** from the **Settings** menu (see Settings Menu on page 14).
2. In the **Alert Policies** window, click the **User Created** tab.
3. To delete the policy, click the **X** icon under **Actions** next to the policy you want to delete, and then click **Yes**.

Alert Metrics

The following table describes the VM metrics.

Table 80: 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 81: 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 82: 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.

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.

Table 83: 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	
cron_avahi_monitor.log	Avahi process status	
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	

Log	Contents	Frequency
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 213).	
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	
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 84: Location: **/home/nutanix/data/logs/cassandra**

Log	Contents
system.log	Cassandra system activity

Table 85: 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 86: Location: **/home/nutanix/data/serviceability/alerts**

Log	Contents
num.processed	Alerts that have been processed

Table 87: Location: **/var/log**

Log	Contents
dmesg	OS start messages
kernel	OS kernel messages
messages	OS messages after starting

Alerts/Health checks

Cluster

Table 88: 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.]

Table 89: 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.

Table 90: 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 threshold or more service restarts of service within one day across all Controller VM(s).`
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.

Table 91: 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.

Table 92: 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.

Table 93: Incorrect {vm_type} NTP Configuration [103076] [A103076]

Name	NTP configuration check
Description	Check that NTP is configured properly on the CVM and hypervisor
Alert message	`vm_type NTP is not properly configured. alert_msg`
Cause	Detected problems with NTP configuration.
Impact	Metadata operations or alerts might not work properly.

Resolution	Follow the instructions in KB 4519.

Table 94: 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

Table 95: 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.

Table 96: 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

Table 97: Checks if AOS version of the remote site configured for synchronous replication is on AOS version >= 5.12 [110021] [A110021]

Name	Remote AOS Version Check
Description	Checks if the remote site configured is on AOS version >= 5.12 to support Synchronous Replication

Alert message	`For protection rule: protection_rule_name, target cluster(s): remote_site_list have version older than 5.12`
Cause	Remote site is on AOS version < 5.12
Impact	Synchronous Replication will not be supported.
Resolution	Upgrade the cluster on the remote site to AOS version >= 5.12

Table 98: Latency to target Prism Element(s) is greater than maximum value allowed for synchronous replications. [110022] [A110022]

Name	Remote Site Latency Check
Description	Check if the the latency to target cluster is less than the maximum value allowed for synchronous replications.
Alert message	`Latency to target Prism Element(s) err_msg is greater than max_permitted_latency`
Cause	Target cluster is unreachable or latency to target cluster is greater than 5ms.
Impact	Sync Rep will be affected.
Resolution	Ensure that the target cluster is reachable or choose another taregt cluster.

Table 99: Cassandra nodes up [111002]

Name	Cassandra nodes up
Description	Check that all Cassandra nodes are running
Cause	
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.

Table 100: Cassandra tokens consistent [111003]

Name	Cassandra tokens consistent
Description	Check that all Cassandra nodes are consistent
Cause	
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.

Table 101: 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.

Table 102: 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	
Resolution	Check the proxy server setting like network, port, user credentials, etc

Table 103: 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 and verify that all cluster services are up.

Table 104: 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.

Table 105: 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.

Table 106: 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.

Table 107: 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	There may be connectivity problems between PC and Cluster
Cause	Catalog service may be down in Cluster.
Impact	Image will not be available on the Cluster.
Resolution	Check cluster connectivity to PC.
Resolution	Check Catalog service on cluster.

Table 108: 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.

Table 109: 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.

Table 110: 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.

Table 111: 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.

Table 112: Cluster Connectivity 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 from cluster cluster_name (uuid = cluster_uuid)is not up-to-date.`

Cause	Cluster network connectivity or CVM 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 CVM services are up.

Table 113: PC vCPU Availability Check [200301] [A200301]

Name	PC vCPU Availability Check
Description	Checks if the number of vCPUs is sufficient for the number of VM entities in Prism Central.
Alert message	`The PC does not have enough vCPUs for the number of VM entities it has.`
Cause	Too many VMs in Prism Central for the number of vCPUs in the Prism Central VM.
Impact	Prism Central performance may be degraded.
Resolution	Reduce the number of VM entities in Prism Central or provide more vCPUs for the Prism Central VM.

Table 114: 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.

Table 115: PC Memory Availability Check [200303] [A200303]

Name	PC 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 PC does not have enough memory for the number of VM entities it has.`
Cause	Too many VMs in Prism Central for the amount of memory in the Prism Central VM.
Impact	Services running in Prism Central may run out of memory and crash.

Resolution	Reduce the number of VM entities in Prism Central or provide more memory for the Prism Central VM.

Table 116: 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 PC cannot handle this 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.

Table 117: PC cluster version check [200310]

Name	PC cluster version check
Description	Checks if Prism Central is running GA version.
Cause	Non GA version detected on Prism Central.
Impact	Prism Central is on Non GA version.
Resolution	Review KB 6325.

Table 118: Prism Central Version EOL [200311] [A200311]

Name	PC 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.

Table 119: 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 still persists, contact Nutanix support.

Table 120: 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 still persists, contact Nutanix support.

Table 121: 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.

Controller VM

Table 122: 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 Controller VM or Prism Central VM ip_address is high, available_memory_kb KB of memory 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 utilisation of Prism Central VM or Controller VM. If abnormal behaviour is seen, please collect logs and contact Nutanix Support.

CVM Available Memory Threshold	Critical Threshold: 768000 KB

Table 123: The {vm_type} is not synchronizing time with any external servers. [3026] [A3026]

Name	CVM NTP Time Synchronized
Description	Checks to ensure that the Controller VM is synchronizing time with an NTP server.
Alert message	`The vm_type is not synchronizing time with any external servers. 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	Verify that the external NTP servers are configured and reachable.

Table 124: CVM Rebooted [3028] [A1024]

Name	CVM Rebooted Check
Description	Check that 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.

Table 125: 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 atleast 10 times in 15 minutes on Controller VM ip_address.Following service(s) have restarted : failed_components.`
Cause	Faulty behaviour in one or more service(s), leading to frequent crashes,in a short period of time.
Impact	Cluster performance may have significantly degraded. In case of multiple nodes with the same condition, the cluster may lose the ability to serve I/O requests.
Resolution	If this alert is generated frequently, please collect CVM logs and contact Nutanix support.

Table 126: 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 behaviour 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 once or infrequently, no action is necessary. If this alert occurs frequently, contact Nutanix support.

Table 127: 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.

Table 128: 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.

Table 129: 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	Increase the disk capacity of the Prism Central VM. Refer to KB 5228 for further details.
Prism Central VM disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %

Table 130: 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.

Table 131: The hypervisor is not synchronizing time with any external servers. [103090] [A103090]

Name	Hypervisor time synchronized
Description	Checks to ensure that the hypervisor is synchronizing time with an NTP server.
Alert message	`The hypervisor is not synchronizing time with any external servers. alert_msg`
Cause	External NTP servers are not configured or are not reachable.
Impact	Logs may have different timestamps in the hypervisor and the CVMs. Hypervisor may not work as expected.
Resolution	Verify if the NTP servers are configured and reachable from the hypervisor.

DR

Table 132: 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.

Table 133: Disconnected Availability Zones are affecting some entities. [110403] [A110403]

Name	Entities Affected By Disconnected Availability Zones
Description	Lists all entites 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.

Table 134: 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.

Table 135: 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 Availablity Zones in the Recovery Plan network settings to be in accordance with Availablity Zone order in the Protection Policy.

Table 136: 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.

Table 137: 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.

Table 138: 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.

Table 139: 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.

Table 140: 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.

Table 141: 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	VMs recovery will fail.
Resolution	Update the Recovery Plans to ensure that a Floating IP address is to be assigned to only one VM.

Table 142: 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.

Table 143: 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.

Guest VM

Table 144: 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 145: 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 146: 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 147: 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 148: 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 149: 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 150: 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	Network configuration/connectivity may not be proper.
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.

Network

Table 151: 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	Network configuration/connectivity may not be proper.
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.

Table 152: 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	Microseg 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 microseg service, PE acropolis service, PC to PE connection, PE to AHV connection.

Table 153: 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	Microseg 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 microseg service, Check PE acropolis service, AHV microseg module.

Table 154: 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

Other

Table 155: 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 fault_azs configured in the Protection Policy protection_rule_name is 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.

Table 156: 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

Table 157: 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

Table 158: 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

Table 159: PCVM 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 PCVMs is not set`
Cause	PCVMs do not have annotation set.
Impact	Guardrails to avoid potentially disruptive user actions on a PCVM will not be applied.
Resolution	Set the annotation for the PCVM.

Table 160: Unequal metadata partition size(s) across Prism Central VMs. [200308] [A200308]

Name	PCVM same disk size check
Description	Checks if metadata partition size(s) are the same among PC VMs.
Alert message	`The metadata partitions across PC 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.

Table 161: High time difference between PC and registered PEs. [200309] [A200309]

Name	PC-PE time drift check
Description	Checks ntp sync between the PC and the registered PEs.
Alert message	`Time differs between PC and registerd PEs.`
Cause	Time is not synchronized between PC and registered PEs.
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.

Table 162: PE-PC incompatible AOS versions. [200312] [A200312]

Name	PC-PE incompatible versions
Description	Checks if the PE and PC AOS versions are compatible.
Alert message	`The AOS version of the registered cluster is incompatible with the PC version.`
Cause	The AOS version of the registered cluster is incompatible with the PC version.
Impact	PC features will not work as expected for this cluster.
Resolution	Upgrade the AOS version of the registered cluster to a version compatible with this PC.

Table 163: PCVM active upgrade check [200313]

Name	PCVM active upgrade check
Description	Checks if there are any active upgrades

Cause	An upgrade was requested for this entity
Impact	Maintenance activities will have to wait until upgrade is finished.
Resolution	Manually upgrade PC. Contact Nutanix support for assistance.

Table 164: PC Upgrades are disabled on {cvm_ip} [200314] [A200314]

Name	PCVM upgrades are disabled check
Description	Checks if upgrades are disabled on PCVM
Alert message	`PC Upgrades are disabled on cvm_ip`
Cause	PC Upgrade has been paused manually.
Impact	Upgrades from Prism or command line won't proceed.
Resolution	Manually upgrade the PC. Contact Nutanix support for assistance.

Table 165: No Protection Policy found for VMs in Recovery Plan [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.

Table 166: 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 VMs from Test Failover operation might get affected.
Resolution	Cleanup all the VMs recovered as part of the test failover of the recovery plan.

Table 167: Recovery Plan has multiple Availability Zone Orders. [300424] [A300424]

Name	Recovery Plan Multiple Availability Zone Order Check

Description	Checks if Recovery Plan has multiple Availability Zone orders.
Alert message	`Recovery Plan recovery_plan_name has more than one Availability Zone Order.`
Cause	Recovery Plan contains more than one Availability Zone Order.
Impact	The Recovery Plan update will not be allowed.
Resolution	Update associated entities in Recovery Plan to have single Availability Zone Order.

Table 168: 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.

Table 169: Incompatible Recovery Availability Zones for 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.

Table 170: 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.

Table 171: 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.
Resolution	Ensure that different IPs are not specified for the same VM in different RPs.

Table 172: 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	Nucalm containers will be inaccessible
Resolution	Please provide correct data service ip. Review KB 5199

Table 173: 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	Nucalm windows automation will fail
Resolution	Please provide correct karan service ip and port. Review KB 5209.

Table 174: Calm Containers are unhealthy [400112] [A400112]

Name	Calm containers health check
Description	Check for Calm Container's state
Alert message	`Nucalm or Epsilon container is in unhealthy state`
Cause	Internal services of each calm containers are down, Docker plugin is not working properly
Impact	Nucalm services may be inaccessible or performs incorrectly.
Resolution	Please check internal services of each calm containers.

Storage

Table 175: Prism Central VM System Root Partition Space Usage High [200316] [A200316]

Name	Prism Central VM System Root Partition Usage
Description	Checks if PC system root partition usage is within threshold to ensure uninterrupted operations. Does not refer to cluster shared storage.
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 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 %

Table 176: 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. Does not refer to cluster shared storage.
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	Critical Threshold: 90 %

System Indicator

Table 177: 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.

9

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 277
- You can specify one or more name servers (see Configuring Name Servers on page 280).
- You can specify one or more NTP servers for setting the system clock (see Configuring NTP Servers on page 281).
- You can configure SNMP (see Configuring SNMP on page 283).
- You can configure a login banner page (see Configuring a Banner Page on page 291).
- You can change the language setting for displayed text (see Changing the Language Settings (Prism Central) on page 293).

Configuring an HTTP Proxy

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.

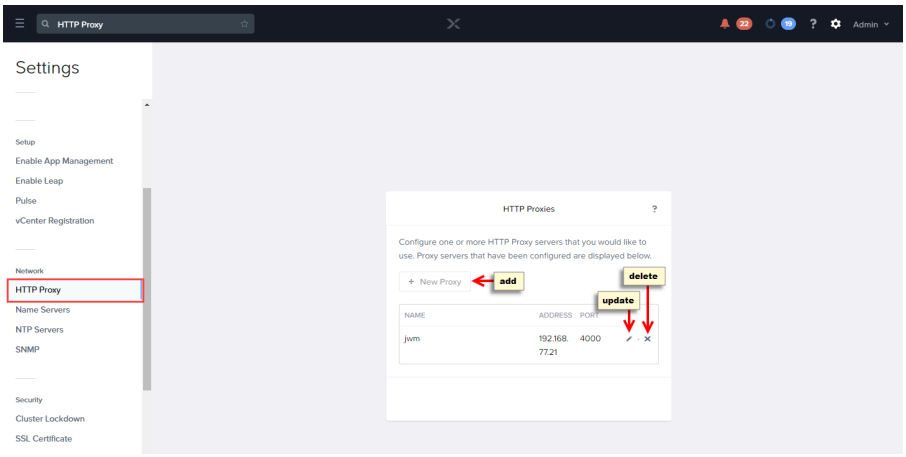

Figure 165: HTTP Proxy Window

2. To add an HTTP proxy, click the **New Proxy** button and do the following in the displayed fields:

   > Note: Only one HTTP proxy can be configured 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 or HTTPS.

   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.

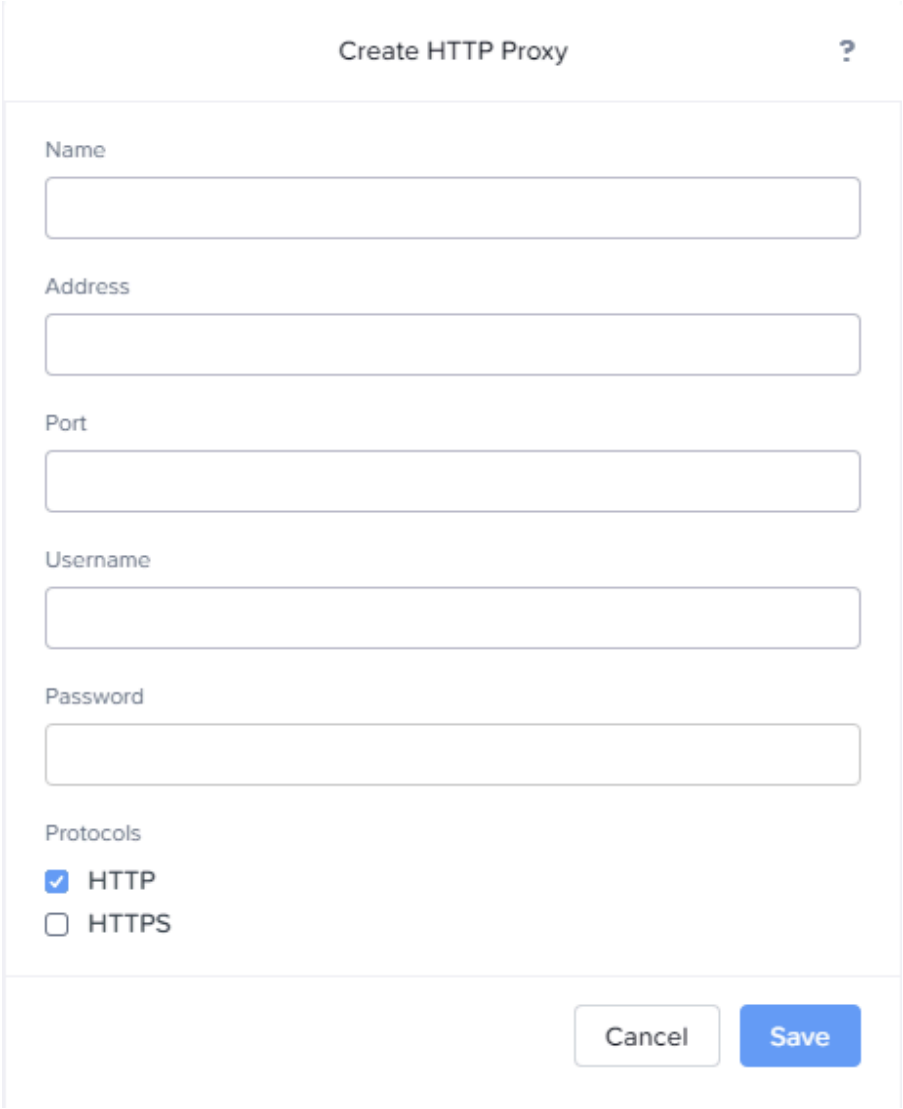

Figure 166: Create HTTP Proxy Window

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 opens a line to enter a target address. Enter the target IP address and then click the **Save** link in that field.
   - 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.

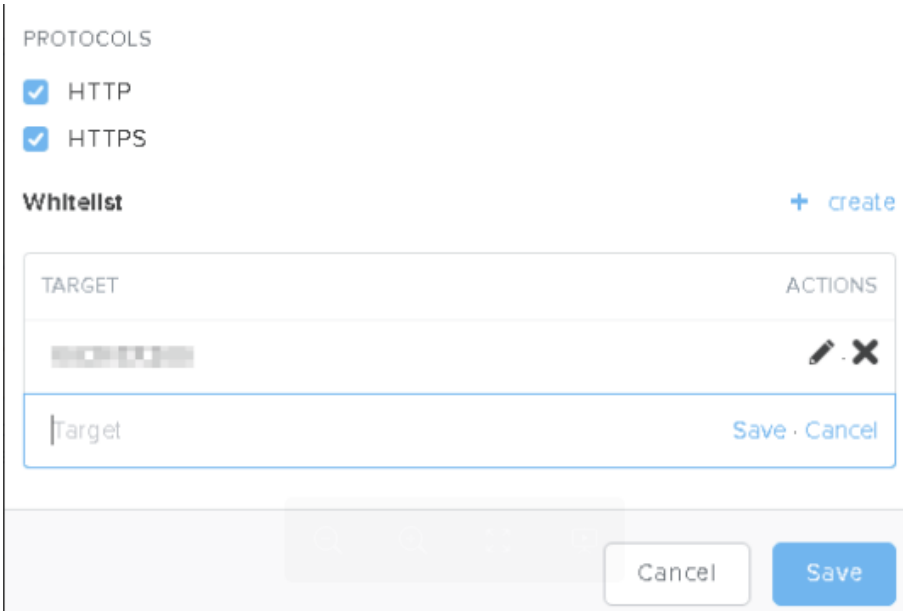

Figure 167: Whitelist Targets

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.

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.

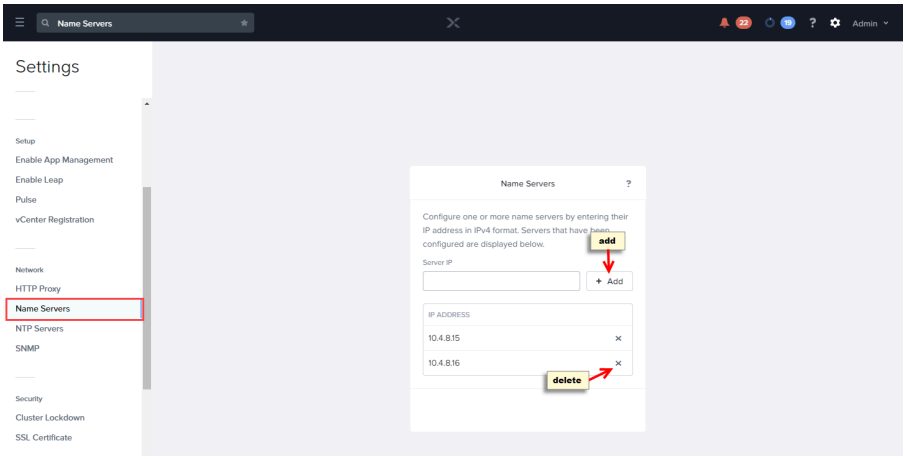

Figure 168: Name Servers Window

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.

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.

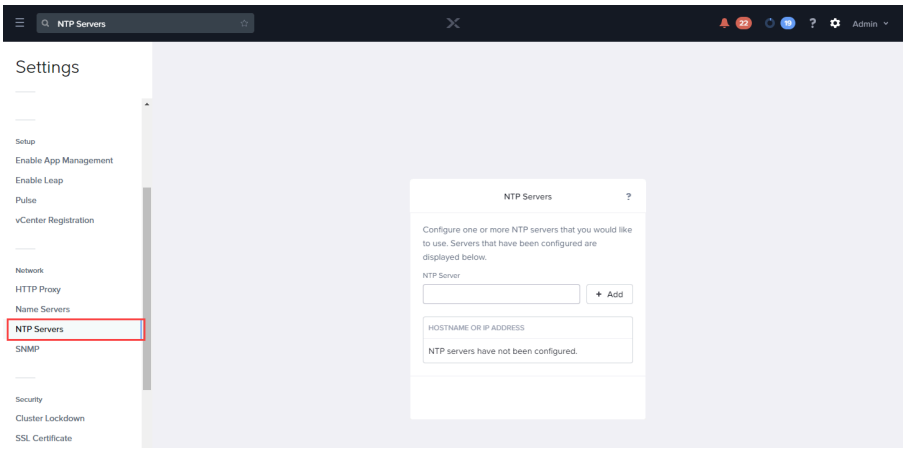

Figure 169: NTP Servers Window

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.

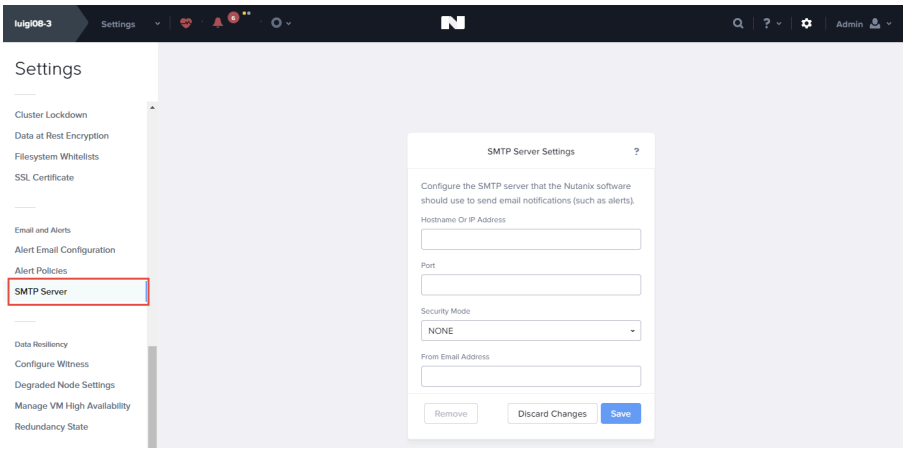

Figure 170: SMTP Server Settings Window

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).

   c. **Security Mode**: Enter the desired security mode from the pull-down list.

      The options are **NONE** (unencrypted), **STARTTLS** (use TTL 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.

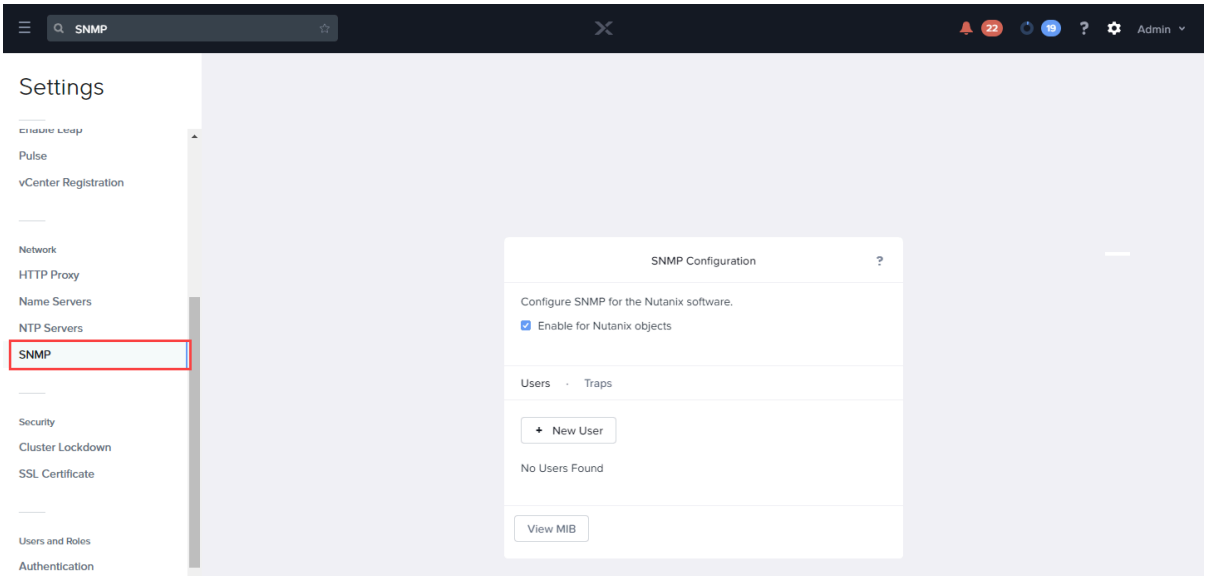

Figure 171: SNMP Configuration Window

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:

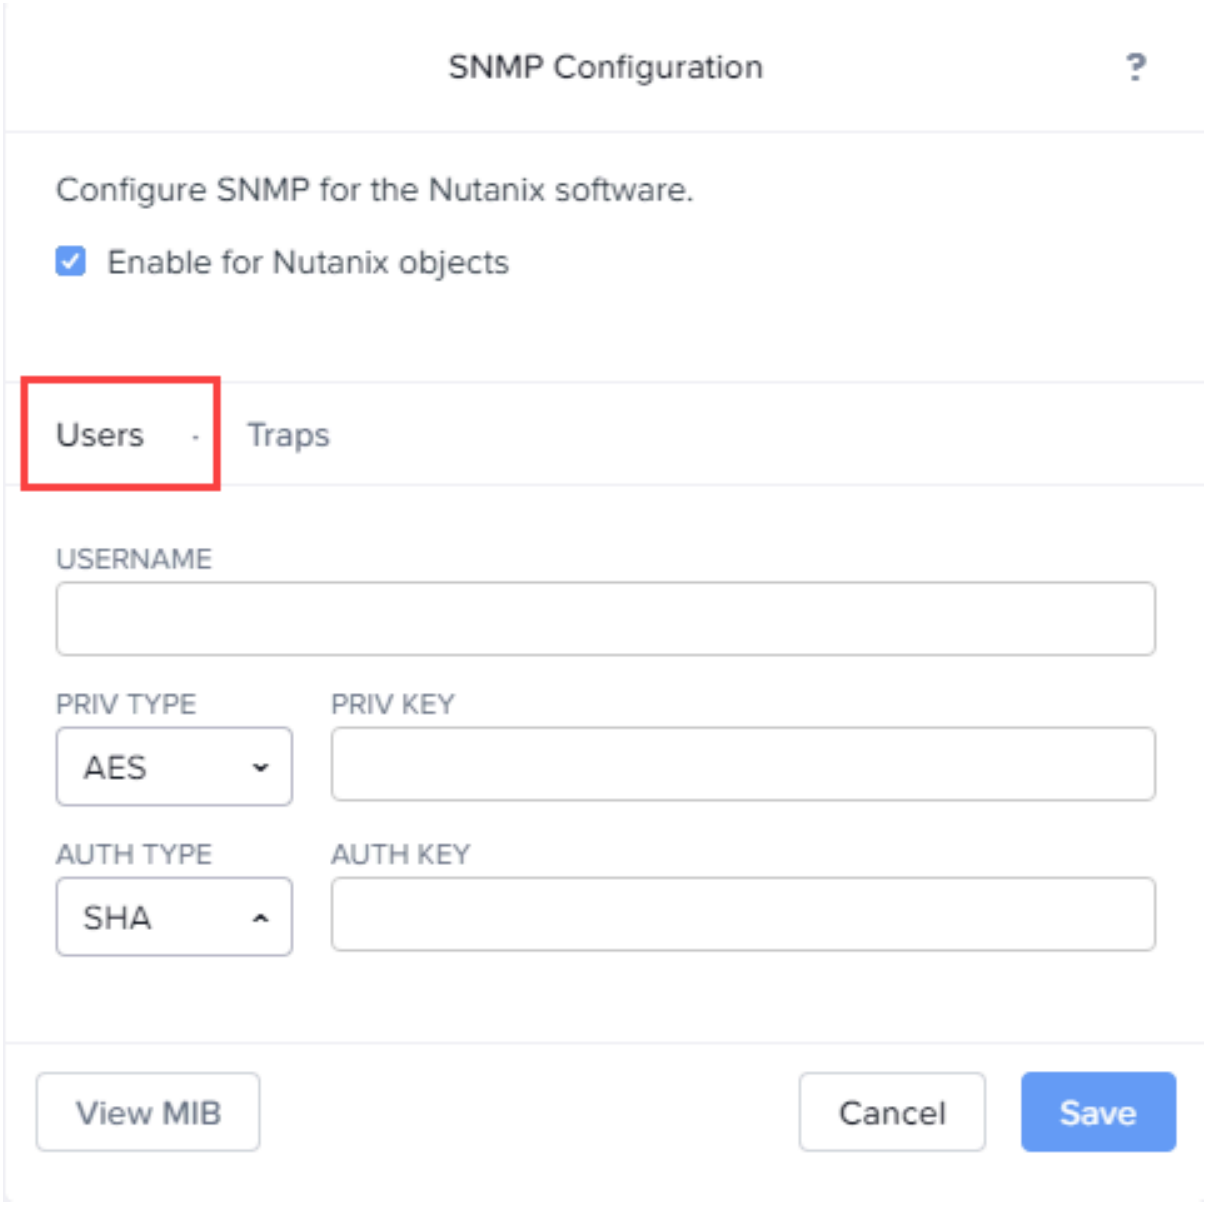

Figure 172: SNMP Configuration: Users Tab

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:

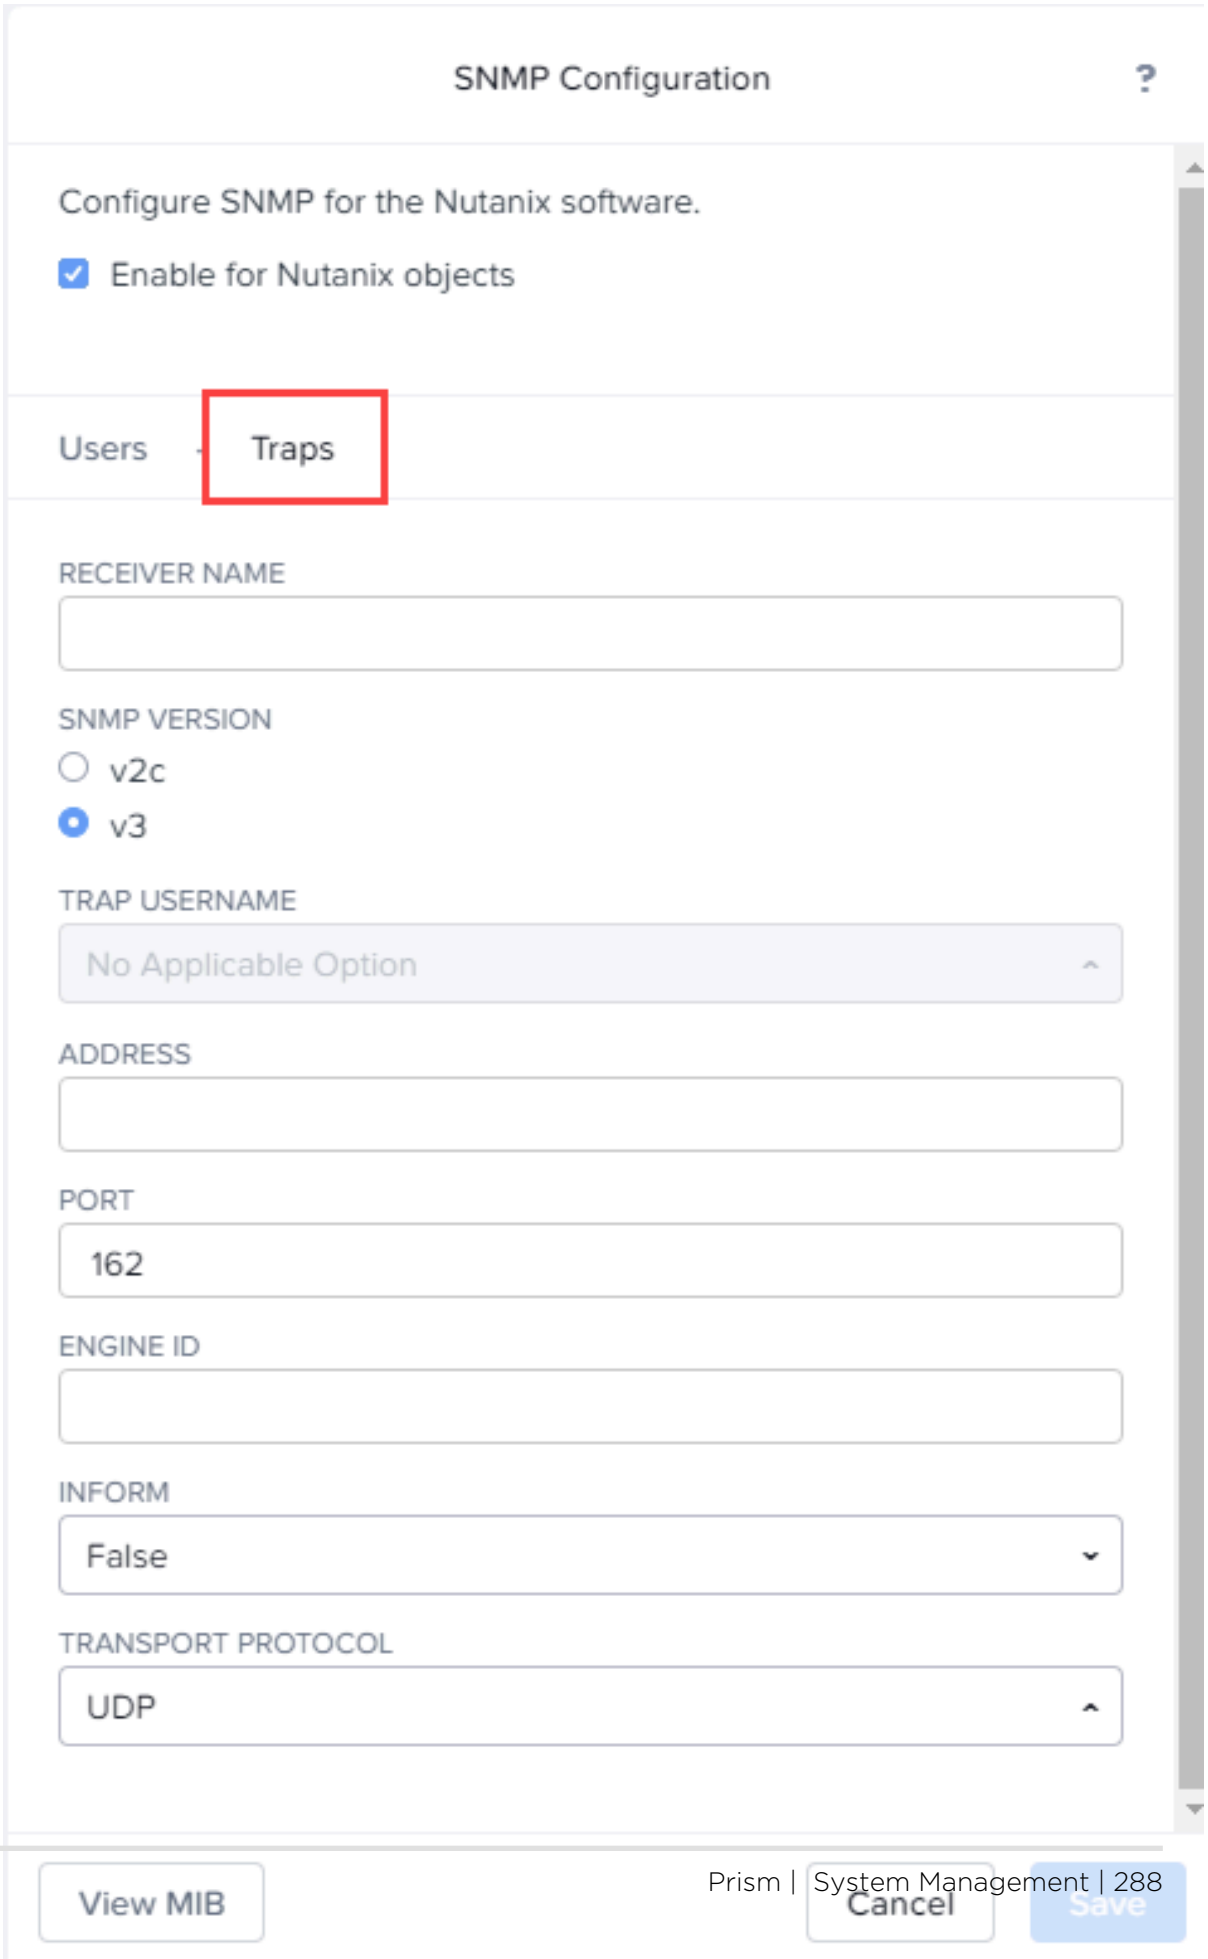
SNMP Configuration
Configure SNMP for the Nutanix software.
Enable for Nutanix objects
Users
Traps
RECEIVER NAME
SNMP VERSION
v2c
v3
TRAP USERNAME
No Applicable Option
ADDRESS
PORT
162
ENGINE ID
INFORM
False
TRANSPORT PROTOCOL
UDP
View MIB
Cancel
Save

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 standard SNMP port number is 161.

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

About this task

Using Prism Central, you can configure syslog monitoring to forward system logs (API Audit, Audit, and Flow logs) of the registered clusters to an external syslog server. 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 the log module (**API Audit**, **Audit**, and **Flow**) and the **Severity Level**. Click **Save**.

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 14).
   The UI Settings window appears.

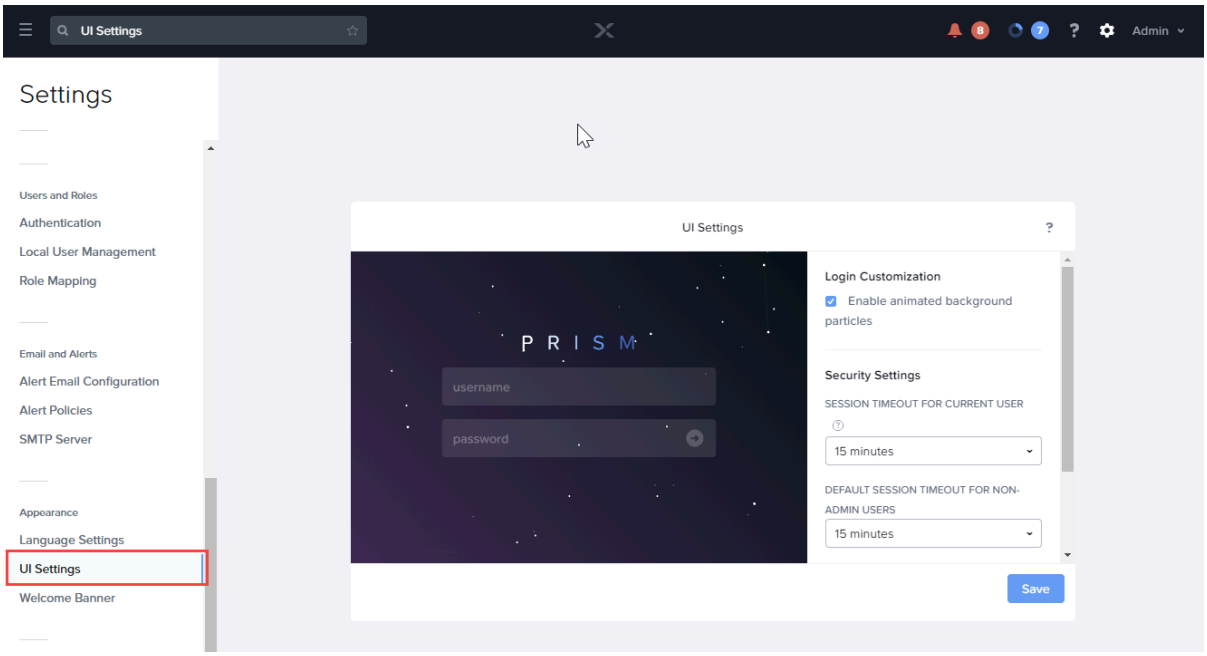

Figure 174: UI Settings Window

2. To disable the login page background animation, under **Login Customization**, clear the **Enable animated background particles** option (or select it to enable).
3. 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.
4. Clear the **Disable 2048 game** option to disable the 2048 game.
5. 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.

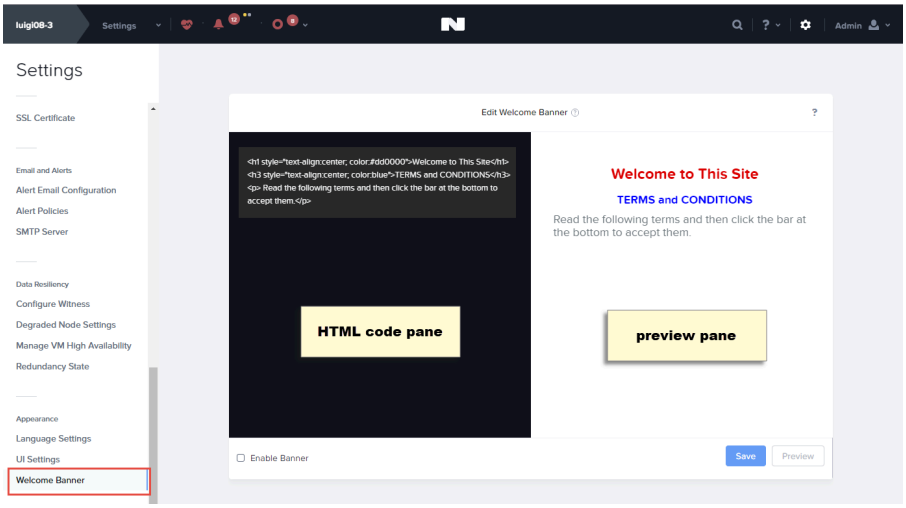

Figure 175: Welcome Banner Window

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 178: 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

Supported Entities	Unsupported Entities
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 293.

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.

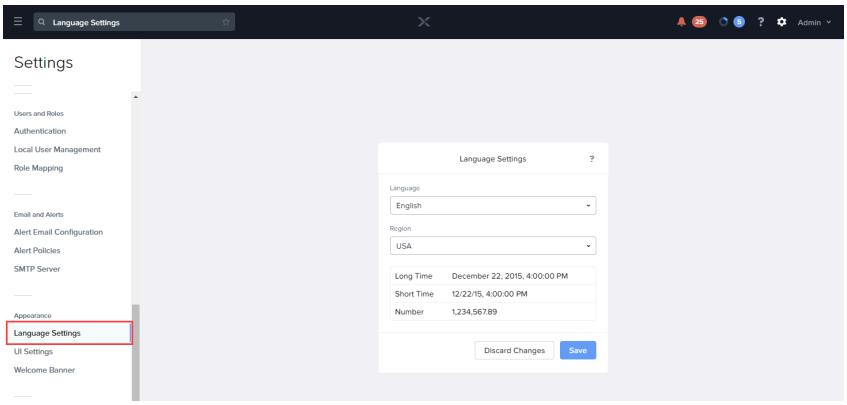

Figure 176: Language Settings Window

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 292. Also, the user interface is localized according to the selection. For more information about localization, see Localization (L10n) (Prism Central) on page 293.

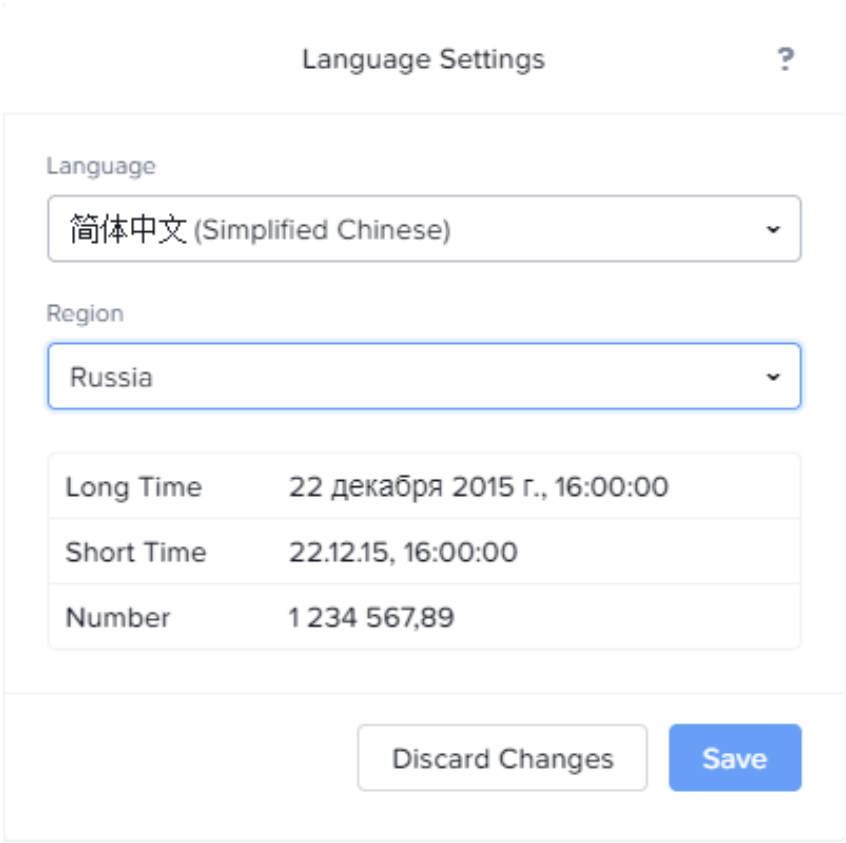

Figure 177: Localized Settings (Chinese/Russian)

10

SECURITY AND USER MANAGEMENT

Prism Central provides several mechanisms to manage security and control user access.

> Note: The following links are to topics in the "Security Management using Prism Central (PC)" chapter in the *Nutanix Security Guide*.

- Set the user authentication method to local, directory service, or both (see “” ())
- Add, edit, or delete local user accounts (see “” ())
  - Update your account (see “” ())
- Assign roles to users (see “” ())
- Install SSL certificates (see “” ())
- Control SSH access to Prism Central (see “” ())
- Create and apply security-related policies (see “” ()).

11

VIRTUAL INFRASTRUCTURE (CLUSTER) ADMINISTRATION

You can view hardware and infrastructure information about registered (individual) clusters through the Entities menu (see Hardware Entities on page 114 and Virtual Infrastructure Entities on page 71). 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 297, Creating a VM (Self Service) on page 306, and Managing a VM (AHV and Self Service) on page 310
  - ESXi: Creating a VM (ESXi) on page 318 and Managing a VM (ESXi) on page 321
- To upload disk or ISO images that can by applied when creating a VM, see Image Management on page 334.
- 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 347.
- To configure affinity polices for an AHV cluster, see Affinity Policies for AHV on page 354
- To configure network connections for a cluster, see Configuring Network Connections on page 357.
- To perform other administrative tasks for a cluster, see Performing Other Administrative Tasks on page 362.
- To register (or unregister) an ESXi cluster with vCenter, see vCenter Server Integration on page 363.

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 318.
  - To manage or modify a VM, see Managing a VM (ESXi) on page 321.
- AHV
  - To create a VM, see Creating a VM (AHV) on page 297.
  - To manage or modify a VM, see Managing a VM (AHV and Self Service) on page 310.

- Prism Self Service (AHV)
  - To create a VM, see Creating a VM (Self Service) on page 306.
  - To manage or modify a VM, see Managing a VM (AHV and Self Service) on page 310.

Creating a VM (AHV)

You can create virtual machines (VMs) in Acropolis managed clusters through Prism Central.

About this task

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 306.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 72) and click the **Create VM** button.
2. In the Cluster Selection window, select the target cluster from the pull-down list.

   A list of registered clusters appears in the window; you can select a cluster running AHV only. Clicking on a cluster name displays the Create VM dialog box for that cluster.

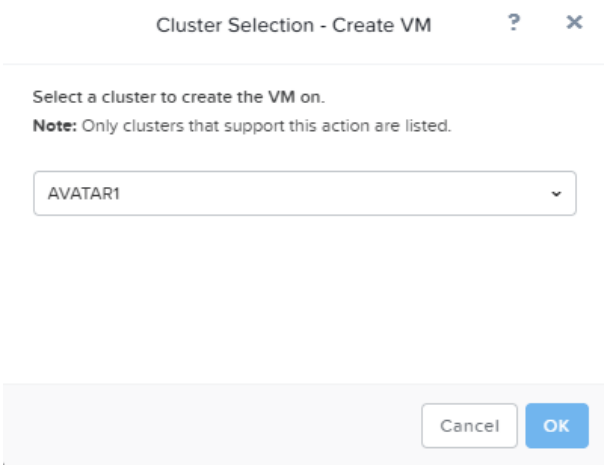

Figure 178: Cluster Selection Window

3. In the Create VM dialog box, do the following in the indicated fields:

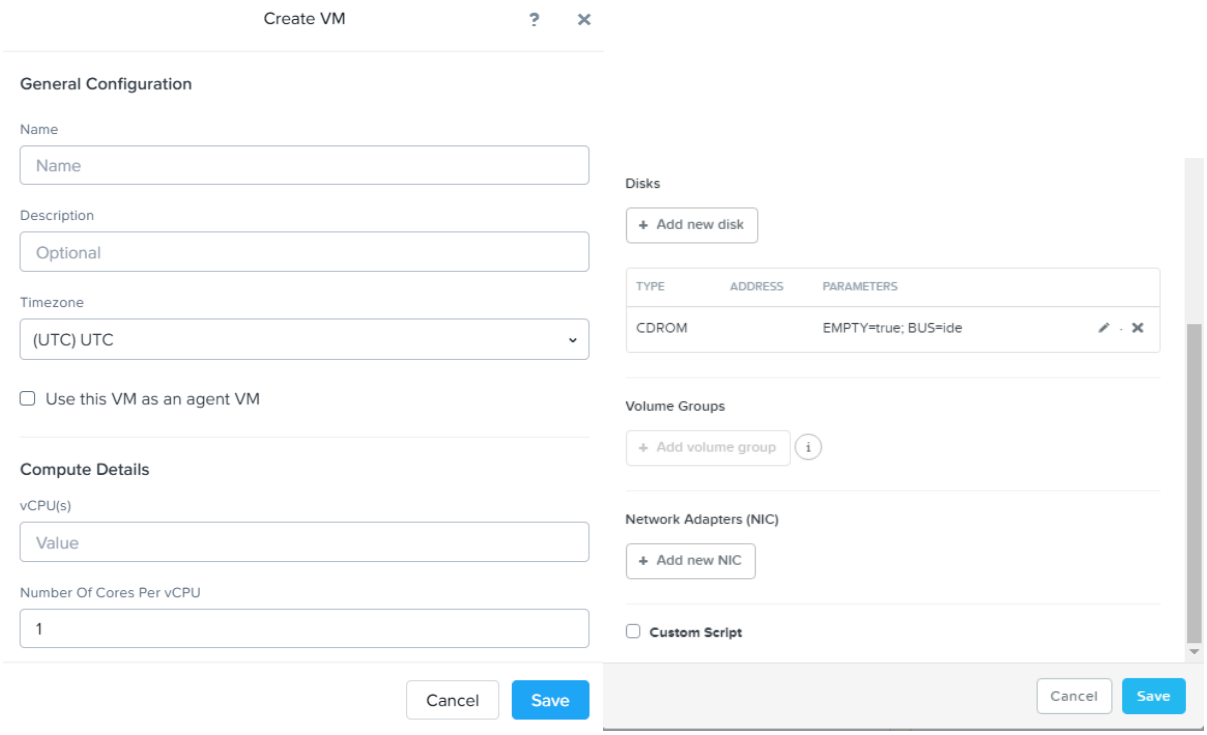

Figure 179: Create VM Dialog Box

a. **Name**: Enter a name for the VM.

b. **Description** (optional): Enter a description for the VM.

c. **Timezone**: Select the local timezone to use from the pull-down list.

d. **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).

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 MiBs) to allocate to this VM.

4. 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 pull-down list. This sets (filters the list of) available profiles. Click the help link (circled question mark) for information about the license types.
   - **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.

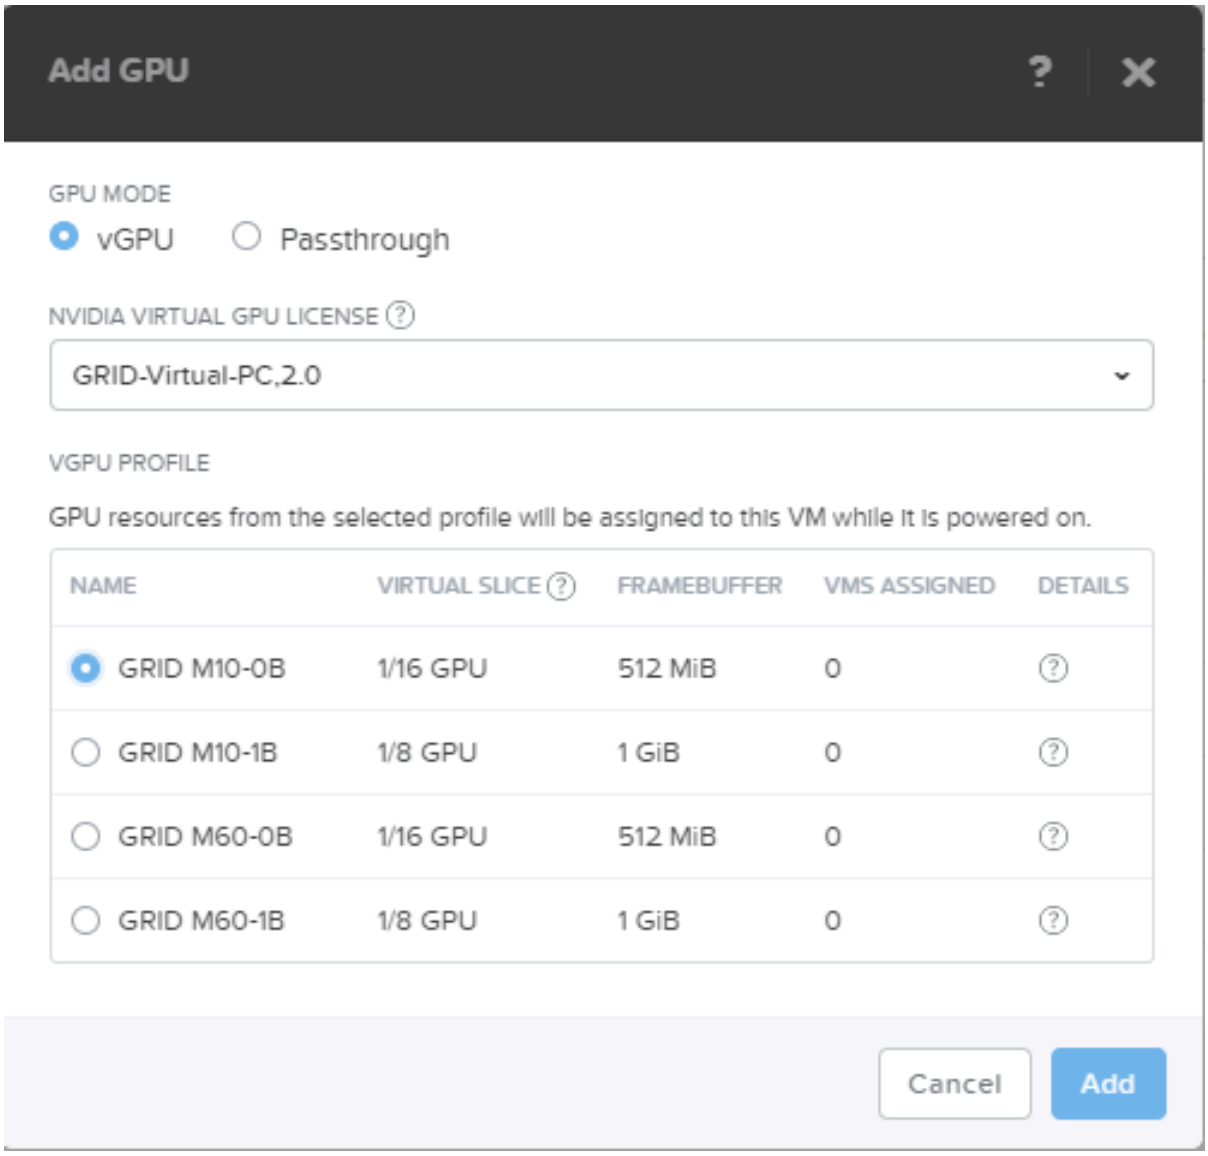

Figure 180: Add GPU Dialog Box: vGPU

   c. If you selected **Passthrough**, click the radio button for the desired type (Tesla M60 or Tesla M10).

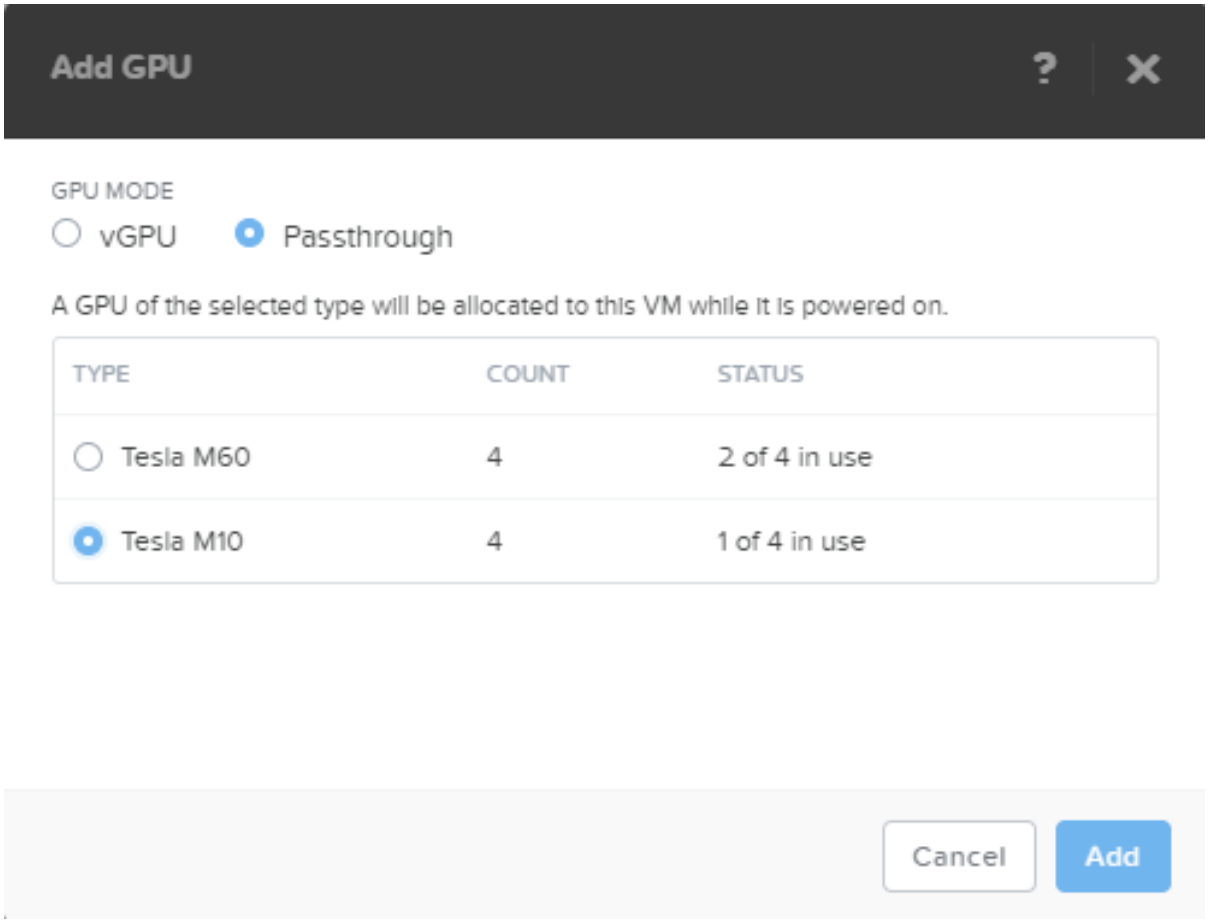

Figure 181: Add GPU Dialog Box: Passthrough

d. Click the **Add** button.

5. To attach a disk to the VM, click the **New Disk** button.

   The Add Disks dialog box appears. Do the following in the indicated fields:

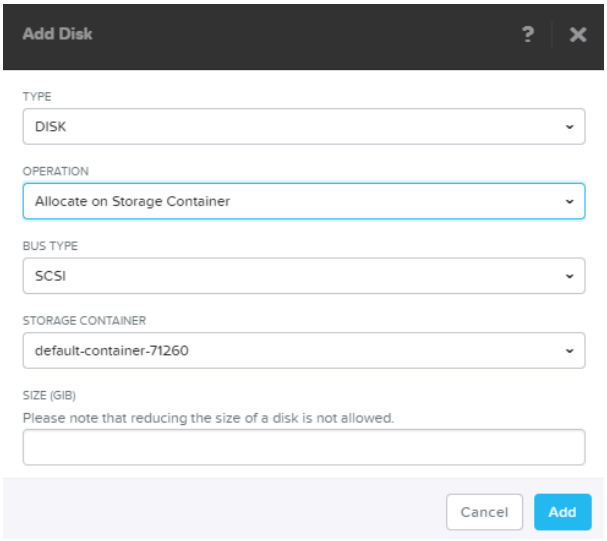

Figure 182: Add Disk Dialog Box

   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 Service** 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**, or **SATA**.

   d. **Path**: Enter the path to the desired system image.

      This field appears only when **Clone from NDSF file** is selected. It specifies the image to copy. Enter the path name as `/container_name/iso_name.iso`. For example to clone an image from `myos.iso` in a storage container named `crt1`, enter `/crt1/myos.iso`. When a user types the storage container name (`/container_name/`), a list appears of the ISO files

in that storage container (assuming one or more ISO files had previously been copied to that storage container).

e. **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 this KB article.

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. **Size**: Enter the disk size in GiBs.

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.

6. 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*.

7. To create a network interface for the VM, click the **New NIC** button.

   The Create NIC dialog box appears. Do the following in the indicated fields:

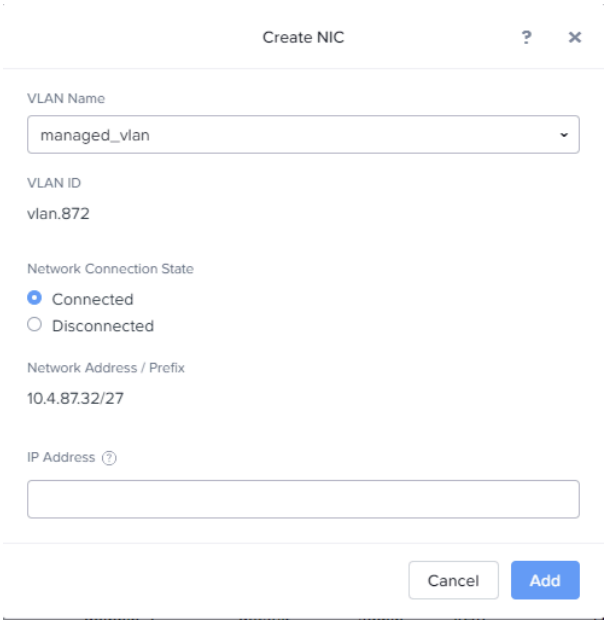

Figure 183: Create NIC Dialog Box

   a. **VLAN Name**: Select the target virtual LAN from the pull-down list.

      The list includes all defined networks (see Configuring Network Connections on page 357).

   b. **VLAN ID**: This is a read-only field that displays the VLAN ID.

   c. **Network Connection State**: Select (click the radio button for) **Connected** or **Disconnected** to connect or disconnect the NIC to the network.

   d. **VLAN UUID**: This is a read-only field that displays the VLAN UUID.

   e. **Network Address/Prefix**: This is a read-only field that displays the network IP address and prefix.

   f. **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.

   g. 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.

   h. Repeat this step to create additional network interfaces for the VM.

8. To configure affinity policy for this VM, click **Set Affinity**.
   The Set VM Host Affinity dialog box appears.

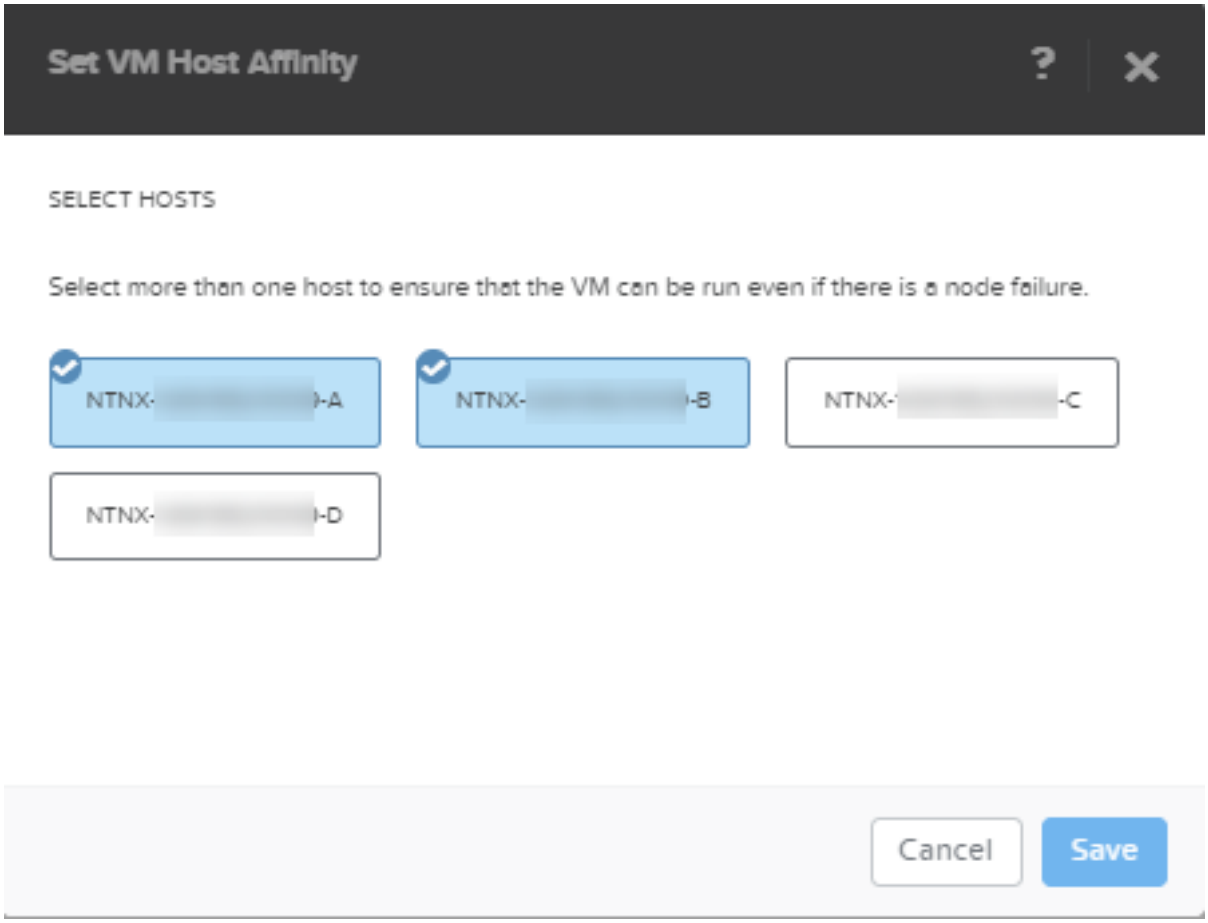

Figure 184: Set VM Host Affinity Dialog Box

   a. Select the host or hosts on which you want to configure the affinity for this VM.
   b. Click **Save**.
      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. This configuration is permanent. The VM will not be moved from this host or hosts even in case of HA event and will take effect once the VM starts. See Affinity Policies for AHV on page 354 for more information about the VM affinity policies.

9. To customize the VM by using Cloud-init (for Linux VMs) or Sysprep (for Windows VMs), select the **Custom Script** check box.

   Fields required for configuring Cloud-init and Sysprep, such as options for specifying a configuration script or answer file and text boxes for specifying paths to required files, appear below the check box.

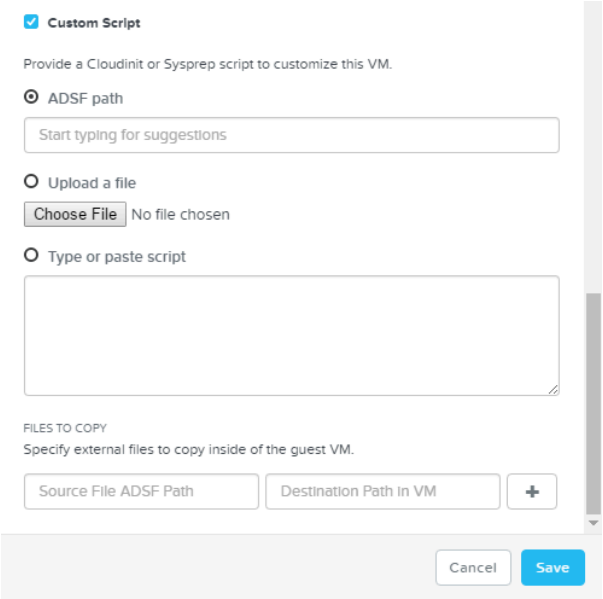

Figure 185: Create VM Dialog Box (custom script fields)

10. To specify a user data file (Linux VMs) or answer file (Windows VMs) for unattended provisioning, do one of the following:

   » If you uploaded the file to a storage container on the cluster, click **ADSF path**, and then enter the path to the file.

   Enter the ADSF prefix (`adsf://`) followed by the absolute path to the file. For example, if the user data is in `/home/my_dir/cloud.cfg`, enter `adsf:///home/my_dir/cloud.cfg`. Note the use of three slashes.

   » If the file is available on your local computer, click **Upload a file**, click **Choose File**, and then upload the file.

   » If you want to create or paste the contents of the file, click **Type or paste script**, and then use the text box that is provided.

11. To copy one or more files to a location on the VM (Linux VMs) or to a location in the ISO file (Windows VMs) during initialization, do the following:

   a. In **Source File ADSF Path**, enter the ADSF prefix (`adsf://`) followed by the absolute path to the file.

   For example, if the file is `/home/my_dir/myfile.txt`, enter `adsf:///home/my_dir/myfile.txt`. Note the use of three slashes.

   b. In **Destination Path in VM**, enter the absolute path to the target directory.

   You do not have to enter a path prefix in this field.

   c. To add another file or directory, click the button beside the destination path field. In the new row that appears, specify the source and target details.

12. When all the field entries are correct, click the **Save** 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 368), 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 347). 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 297.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 72) 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.

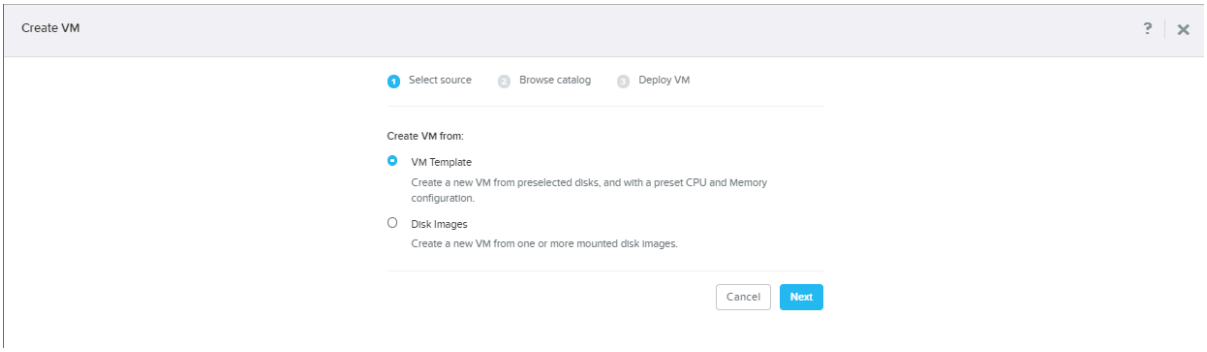

Figure 186: Create VM page: Select Source tab

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.

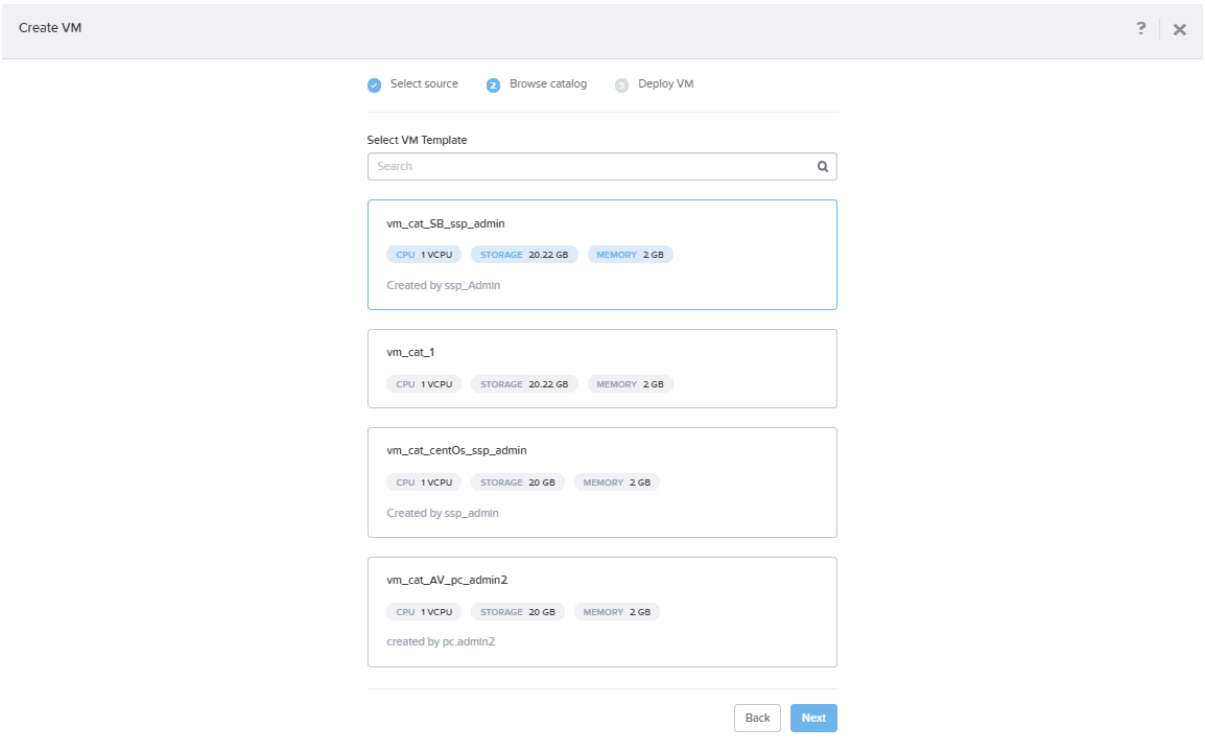

Figure 187: Create VM page: Browse Catalog tab (VM templates)

   - 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.

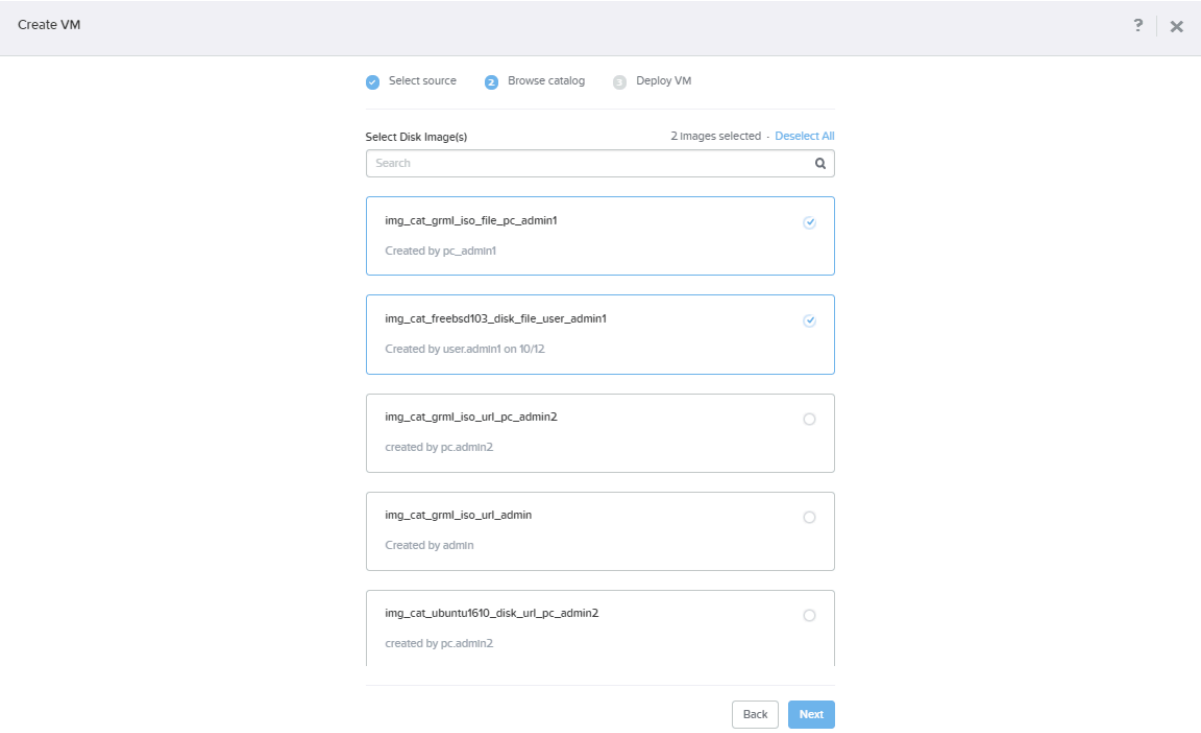

Figure 188: Create VM page: Browse Catalog tab (disk images)

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 347).

   f. **Network**: Select a network for the VM.

      If there is a default network specified for your project (see Creating a Project on page 373), 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.

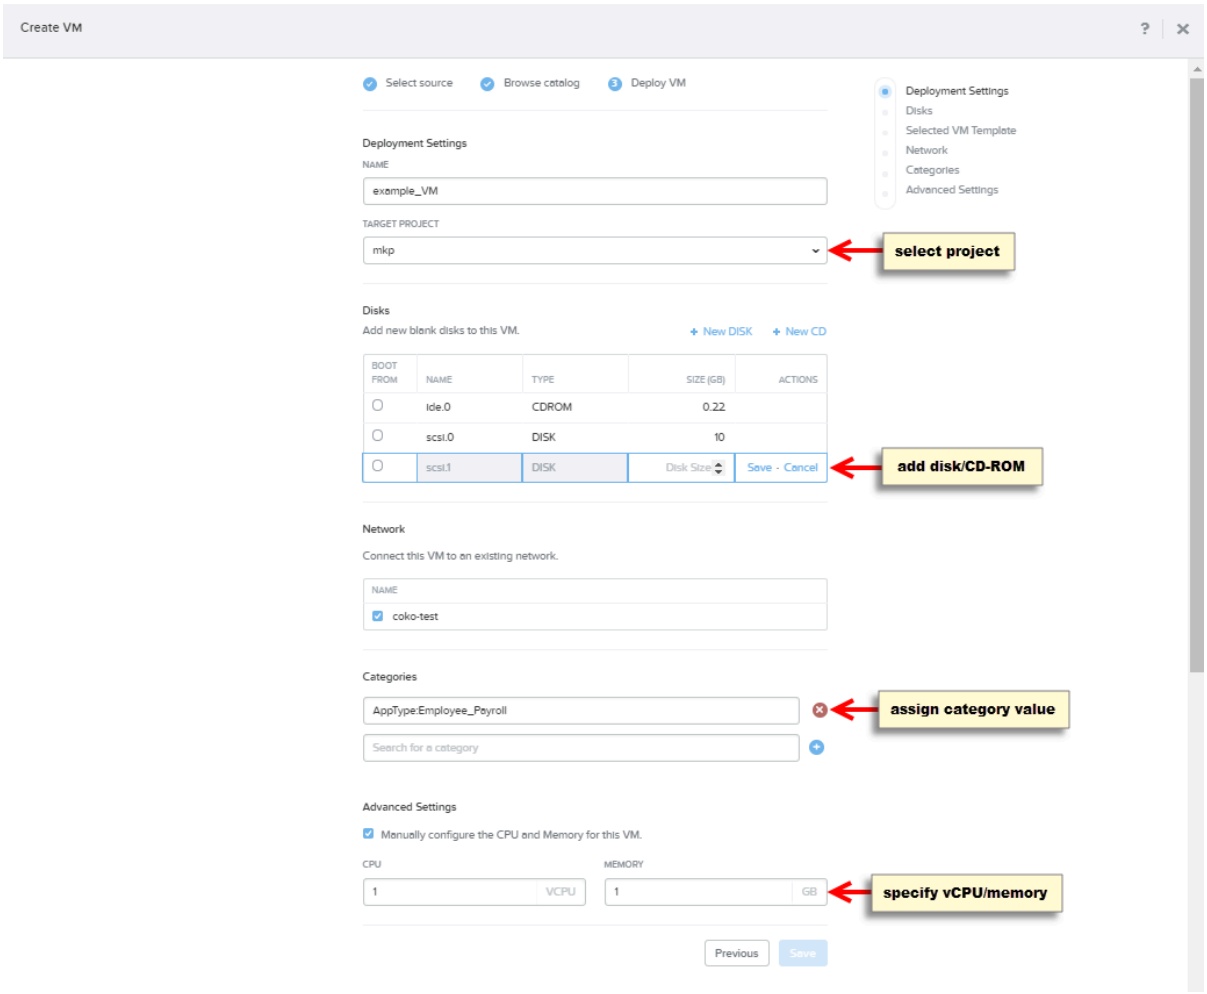

Figure 189: Create VM page: Deploy VM tab (disk images)

Managing a VM (AHV and Self Service)

About this task

After creating a VM (see Creating a VM (AHV) on page 297 or Creating a VM (Self Service) on page 306), 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, pause (or resume) the VM, assign the VM to a protection policy, take a snapshot, migrate the VM, add the VM to a recovery plan, run a playbook, manage categories, quarantine (or unquaratine) the VM, install and manage Nutanix Guest Tools (NGT), configure VM host affinity, add the VM image to the catalog, manage the VM ownership, or configure QoS settings.

> Note: You can perform only those operations for which you have permission.

To perform one of these tasks, either select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 72) or go to the details page for that VM (see VM Details View on page 81) and select the desired action.

> Note: To choose an action, select from the **Actions** menu (VMs Summary View) or click the appropriate action button in the upper right (VM Details View). 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.

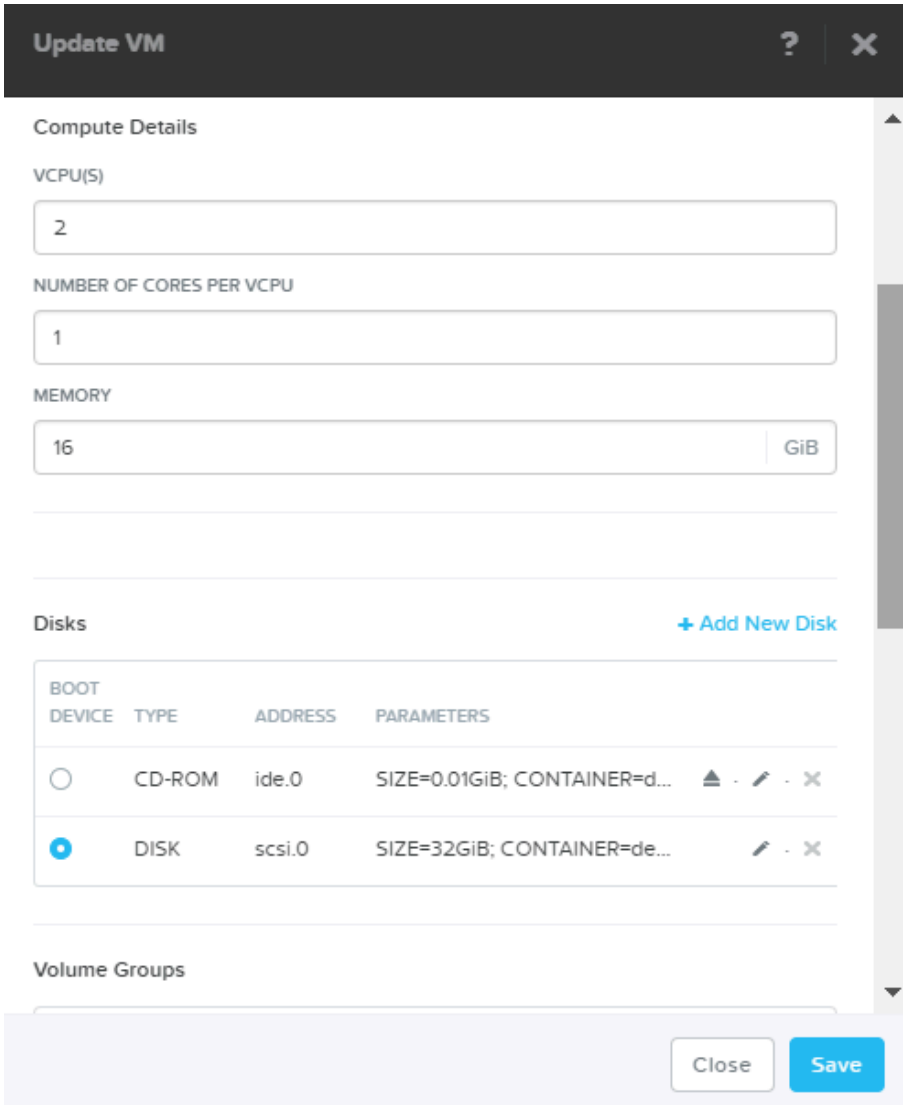

Figure 190: Update VM Window

- 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 but with all fields (except the name) filled in with the current VM settings. Enter

a name for the clone and then click the **Save** button to create the clone. You can create a modified clone by changing some of the settings before clicking the **Save** button.

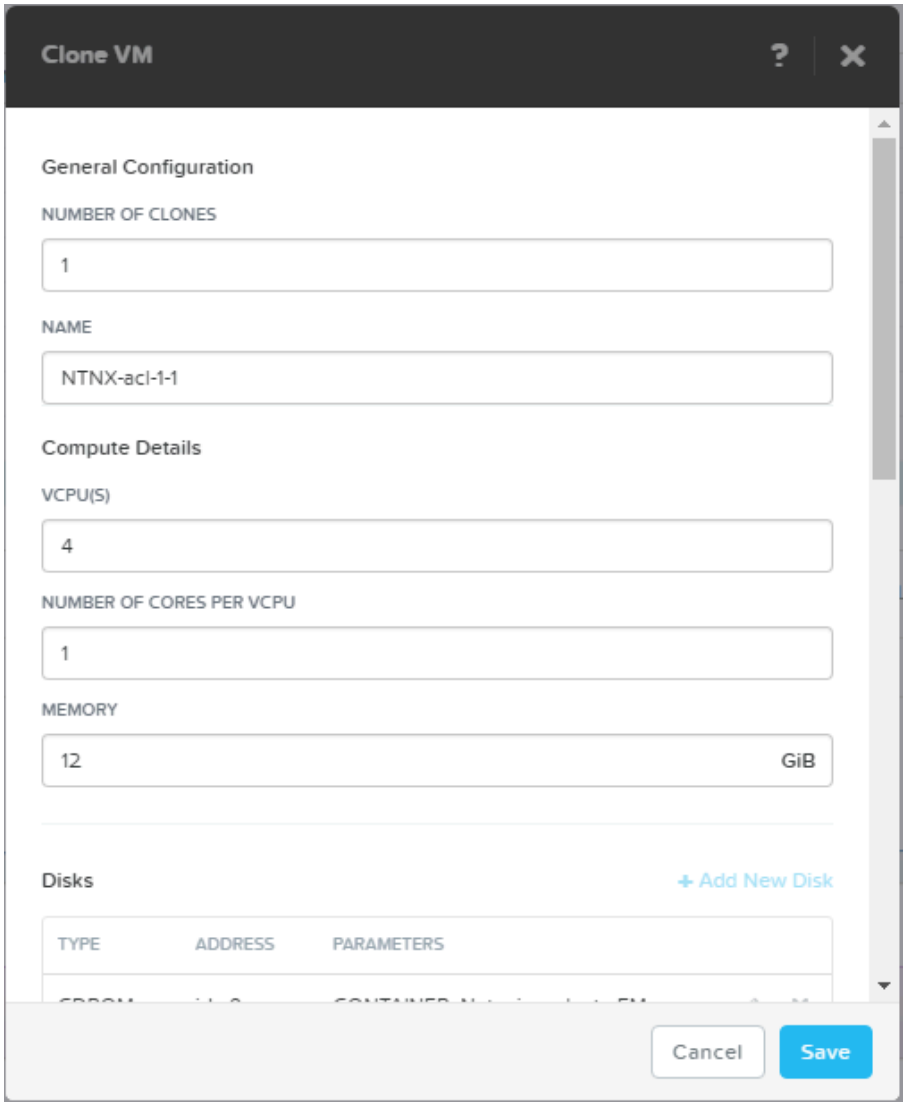

Figure 191: Clone VM Window

- 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.)

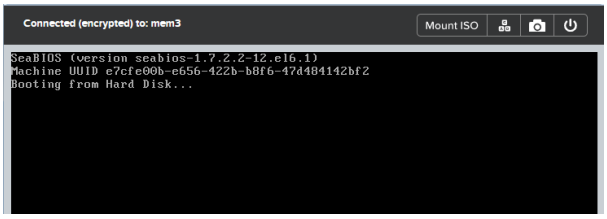

Figure 192: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To pause (or resume) the VM, select **Pause/Suspend** (or **Resume**). This option is available only when the VM is powered on (off).
- 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 377). To remove the VM from a protection policy, select **Unprotect**.
- To take a snapshot of the VM, select **Take Snapshot**.

  This displays the Take Snapshot dialog box. Enter a name for the snapshot and then click the **Submit** button to start the backup.

> Note: These snapshots (stored locally) cannot be replicated to other sites.

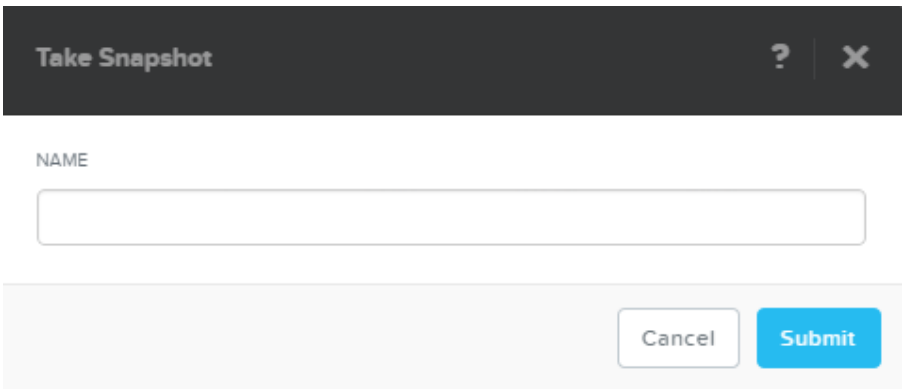

Figure 193: Take Snapshot Window

- 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.

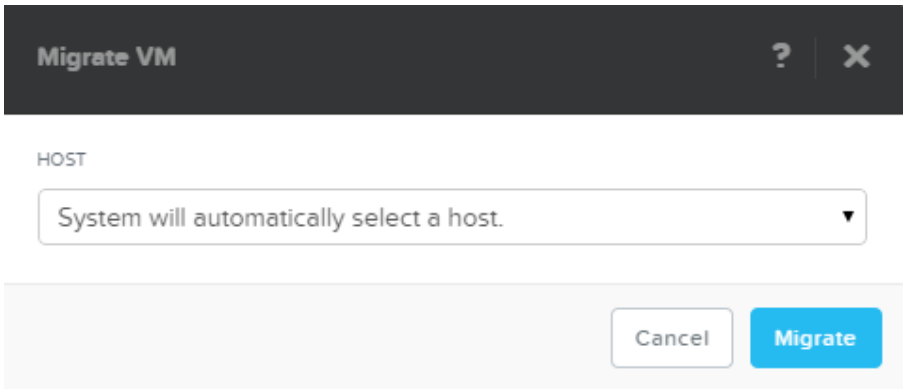

Figure 194: Migrate VM Window

> 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 426).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 353).

- To quarantine a VM, select **Quarantine VM**.

  This displays the Quarantine VMs window. Click the radio button for the desired quarantine method:

  - Select **Strict** to isolate the VM from everything on the network.
  - Select **Forensics** to isolate the VM from everything except the forensics tools specified in the quarantine policy (see “” ()).

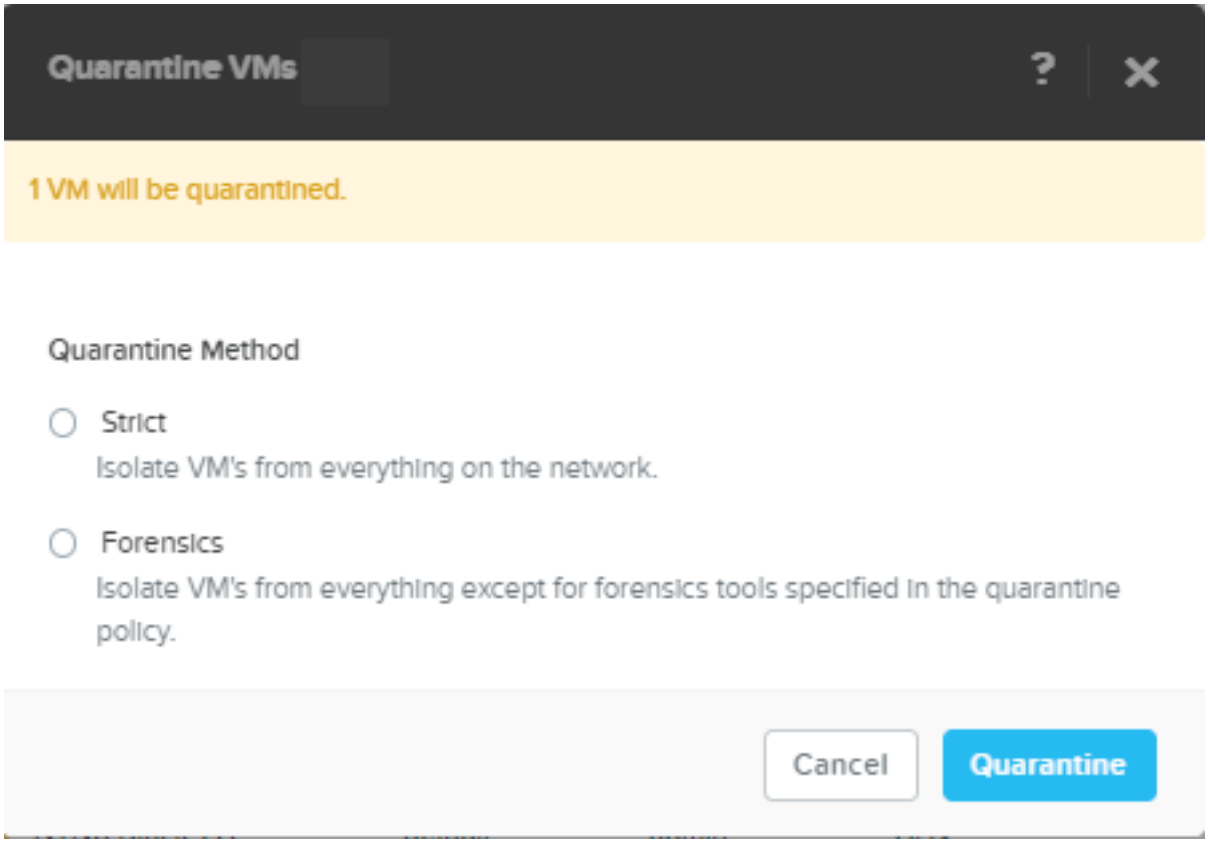

Figure 195: Quarantine VMs Window

- To unquarantine a VM, select **Unquarantine VM**.

  This displays the Unquarantine VMs window. Click the **Unquarantine** button to allow the VM access to the network.

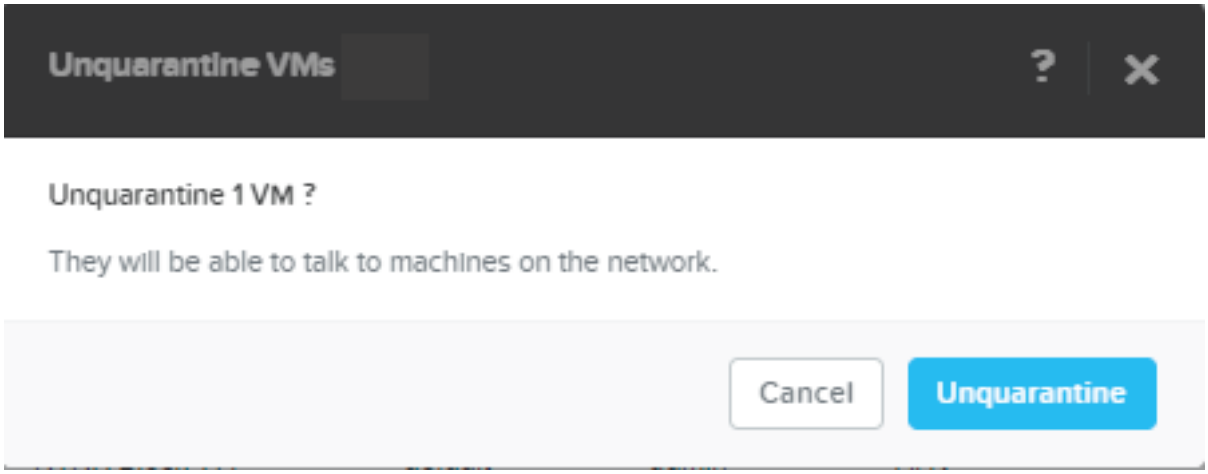

  Figure 196: Unquarantine VMs Window

- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 327).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 329).

  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 330).

- 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 354 for more information about the VM affinity policies.

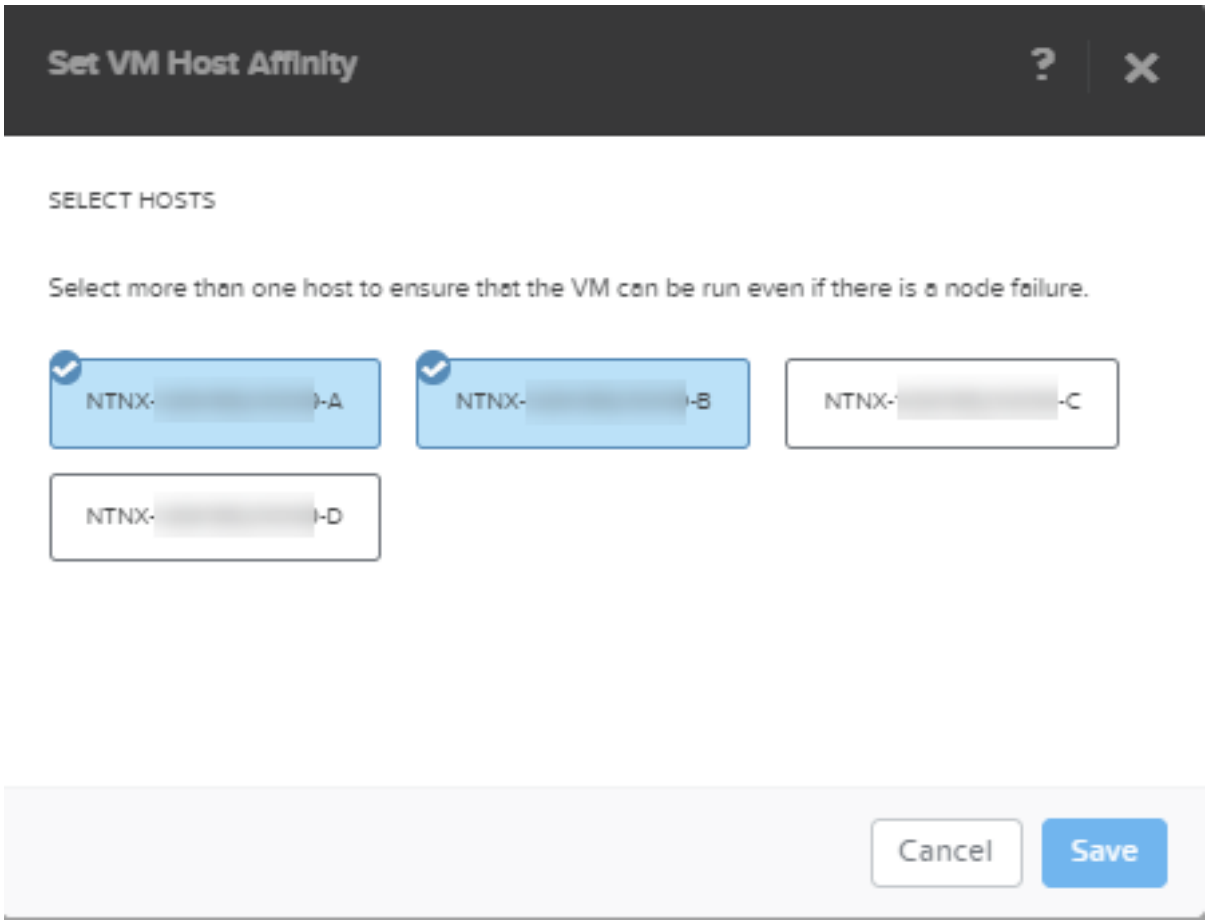

Figure 197: Set VM Host Affinity Window

- To add a VM snapshot to the catalog, select **Add to Catalog**. This displays the **Add VM to Catalog** page (see Adding a Catalog Item on page 347).

- 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.

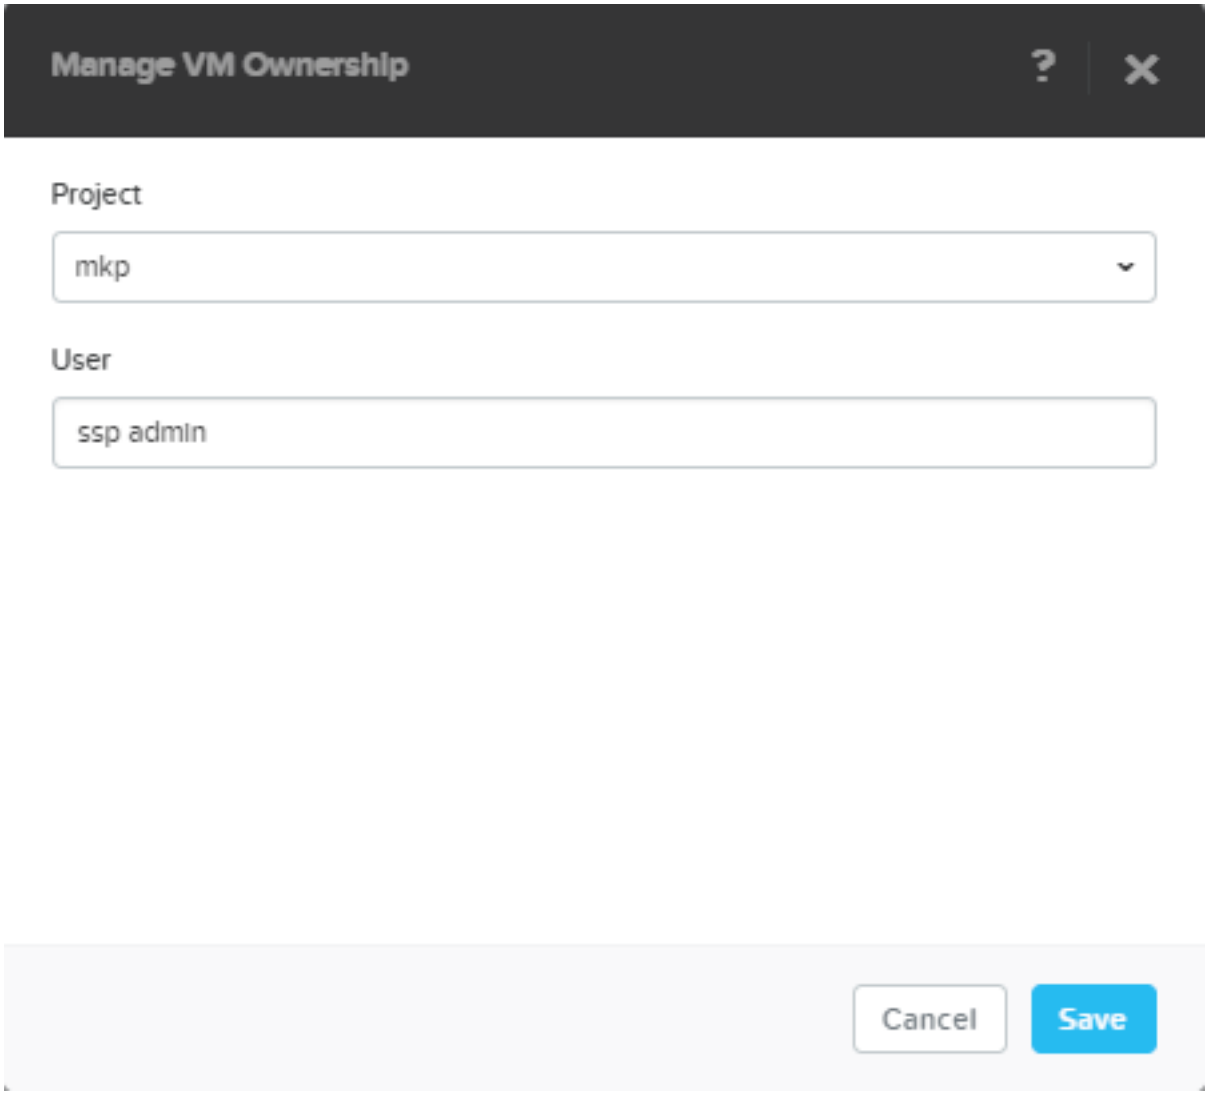

Figure 198: VM Ownership Window

- To configure quality of service (QoS) settings, select **Set QoS Attributes** (see Setting QoS for an Individual VM on page 333).

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 363 before proceeding.
- Register the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 364.

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 72) and click the **Create VM** button.
2. In the Cluster Selection window, select the target cluster from the pull-down list.

   A list of registered clusters appears in the window; you can select a cluster running AHV only. Clicking on a cluster name displays the Create VM dialog box for that cluster.

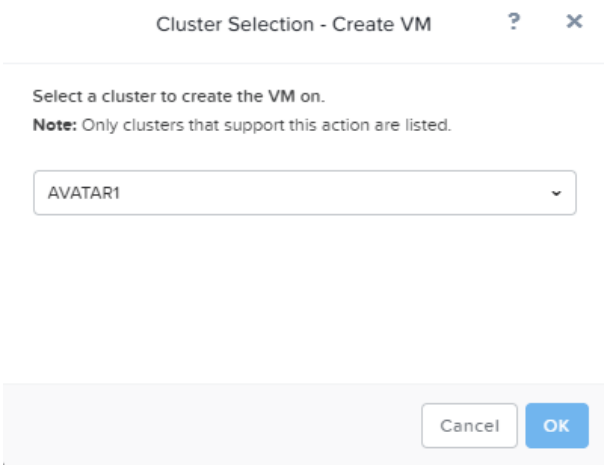

Figure 199: Cluster Selection Window

3. In the Create VM dialog box, 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. **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 363 topic.
   d. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
   e. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
   f. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

4. To attach a disk to the VM, click the **Add New Disk** button.
   The Add Disks dialog box appears. Do the following in the indicated fields:

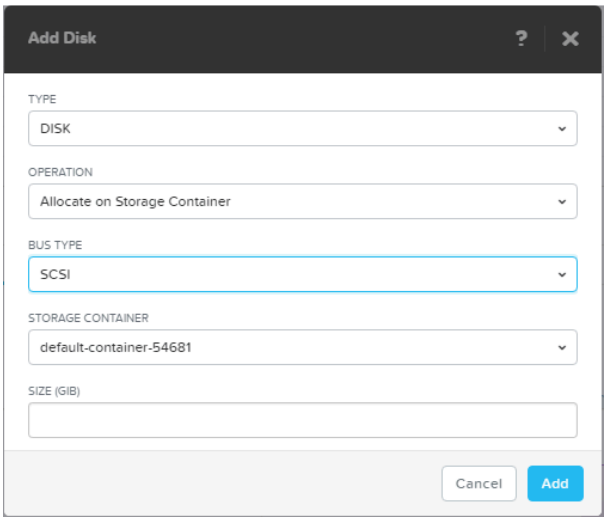

Figure 200: Add Disk Dialog Box

   a. **Type**: Select the type of storage device, **DISK** or **CD-ROM**, from the pull-down list.

      The following fields and options vary depending on whether you choose **DISK** or **CD-ROM**. You can use the **CD-ROM** type only to create a blank CD-ROM device for mounting NGT or VMware guest tools.

   b. **Operation**: Specify the device contents from the pull-down list.

      - Select **Clone from ADSF file** to copy any file from the cluster that can be used as an image onto the disk.
      - Select **Allocate on Storage Container** to allocate space without specifying an image. (This option appears only when **DISK** is selected in the previous field.) Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.

   c. **Bus Type**: Select the bus type from the pull-down list. The choices are **IDE** or **SCSI**.

   d. **ADSF Path**: Enter the path to the desired system image.

      This field appears only when **Clone from ADSF file** is selected. It specifies the image to copy. Enter the path name as `/container_name/vmdk_name.vmdk`. For example to clone an image from `myvm-flat.vmdk` in a storage container named `crt1`, enter `/crt1/myvm-flat.vmdk`. When a user types the storage container name (`/container_name/`), a list

appears of the VMDK files in that storage container (assuming one or more VMDK files had previously been copied to that storage container).

> Note: Make sure you're copying from a flat file.

e. **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.

f. **Size**: Enter the disk size in GiBs.

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.

5. To create a network interface for the VM, click the **Add New NIC** button.

   The Create NIC dialog box appears. Do the following in the indicated fields:

   a. **VLAN Name**: Select the target virtual LAN from the pull-down list.

      The list includes all defined networks (see Configuring Network Connections on page 357).

   b. **Network Adapter Type**: Select the network adapter type from the pull-down list.

      See vCenter Server Integration on page 363 for the list of supported adapter types.

   c. **Network UUID**: This is a read-only field that displays the network UUID.

   d. **Network Address/Prefix**: This is a read-only field that displays the network IP address and prefix.

   e. 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.

   f. Repeat this step to create additional network interfaces for the VM.

6. When all the field entries are correct, click the **Save** 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 clusters through Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration on page 363 before proceeding.
- Ensure that you have registered the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 364.

About this task

After creating a VM (see Creating a VM (ESXi) on page 318), 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, pause (or resume) the VM, assign the VM to a protection policy, take a snapshot, 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.

To perform one of these tasks, either select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 72) or go to the details page for that VM (see VM Details View on page 81) and select the desired action.

> Note: To choose an action, select from the **Actions** menu (VMs Summary View) or click the appropriate action button in the upper right (VM Details View). The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM.

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. You cannot change the name, number of assigned vCPUs, or memory size of the VM, but you can add or delete disks and NICs.

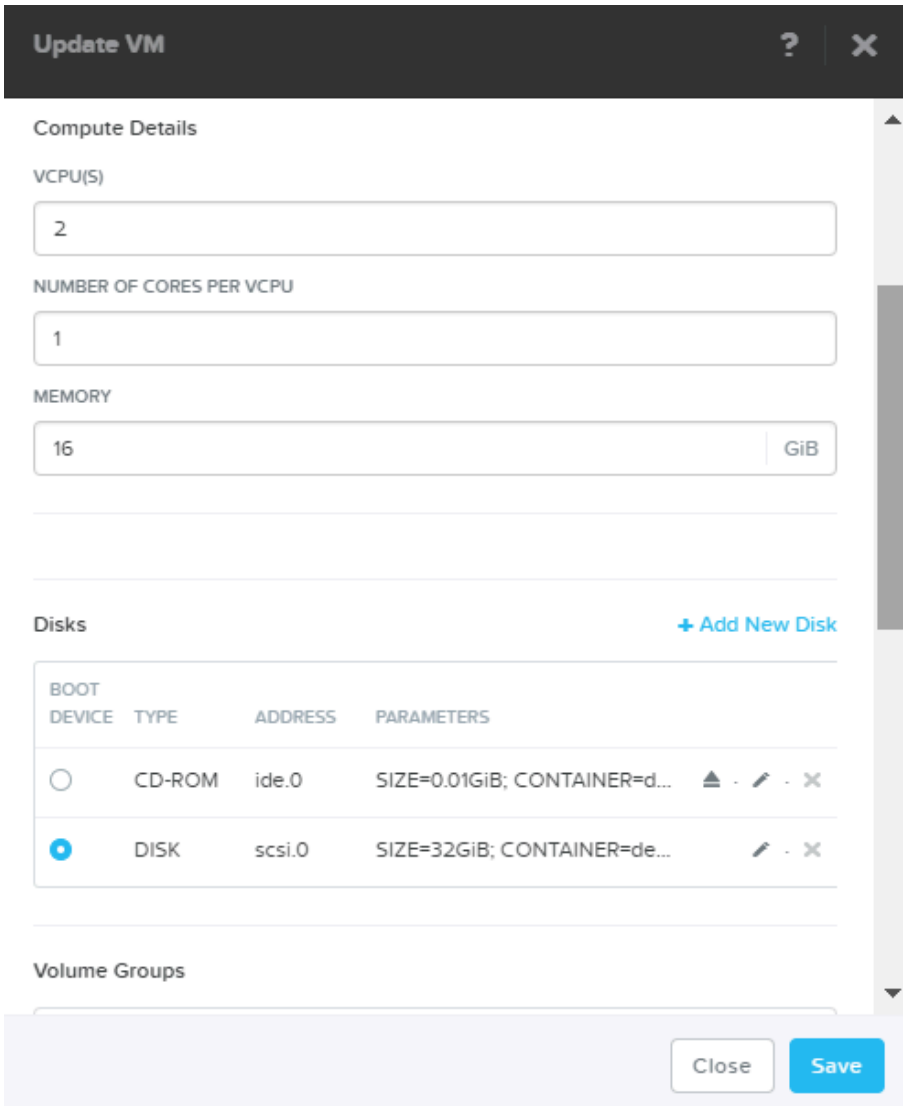

Figure 201: Update VM Window

- 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 but with all fields (except the name) filled in with the current VM settings. Enter a name for the clone and then click the **Save** button to create the clone. You can create a modified clone by changing some of the settings before clicking the **Save** button.

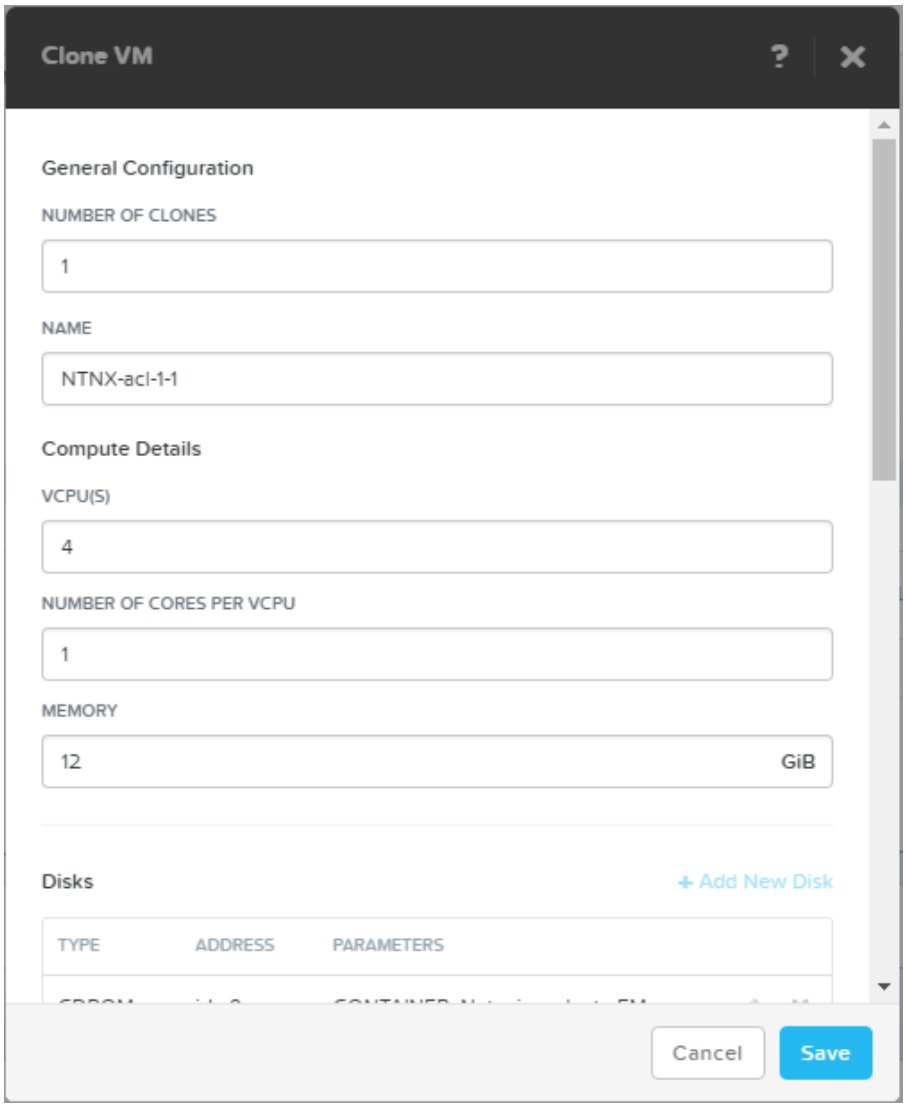

Figure 202: Clone VM Window

- 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.)

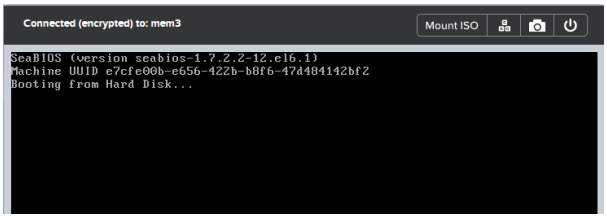

Figure 203: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To pause (or resume) the VM, select **Pause/Suspend** (or **Resume**). This option is available only when the VM is powered on (off).
- 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 377). To remove the VM from a protection policy, select **Unprotect**.
- To take a snapshot of the VM, select **Take Snapshot**.

  This displays the Take Snapshot dialog box. Enter a name for the snapshot and then click the **Submit** button to start the backup.

  > Note: These snapshots (stored locally) cannot be replicated to other sites.

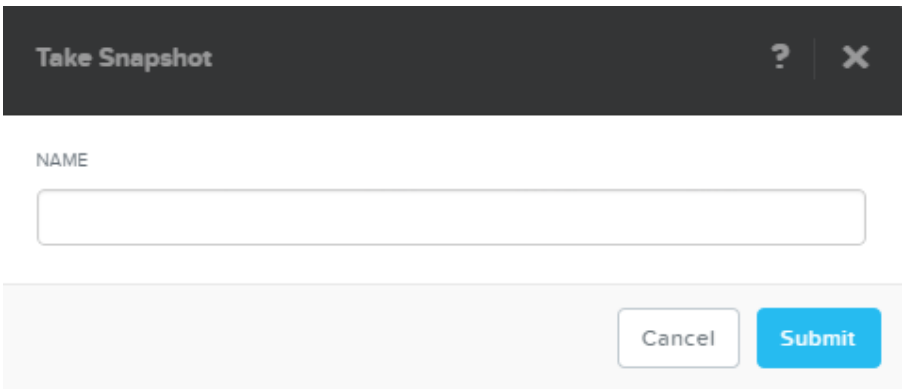

Figure 204: Take Snapshot Window

- 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 426).
- To assign the VM a category value, select **Manage Categories** (see Assigning a Category on page 353).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on Multiple VMs on page 327).

- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 329).

  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 330).
- To configure quality of service (QoS) settings, select **Set QoS Attributes** (see Setting QoS for an Individual VM on page 333).

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.
- The following ports should be open between the Controller VMs and the user VMs with NGT installed:
  - Controller VM: 9440, 2073, 2074
  - User VM: 5000 (dynamic), 23578. These ports do not have to be open to the outside, just be free for communication with the Controller VMs. In addition, in order for Prism Central to manage NGT install and upgrade, user VMs should have port 5986 open on Windows (winrm actions require them) or port 22 on Linux.
- VMs and virtual IP address of the cluster should be able to communicate with each other.
- 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 329 topic for how to enable NGT on the cloned VM.

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 command 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 72).
   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.

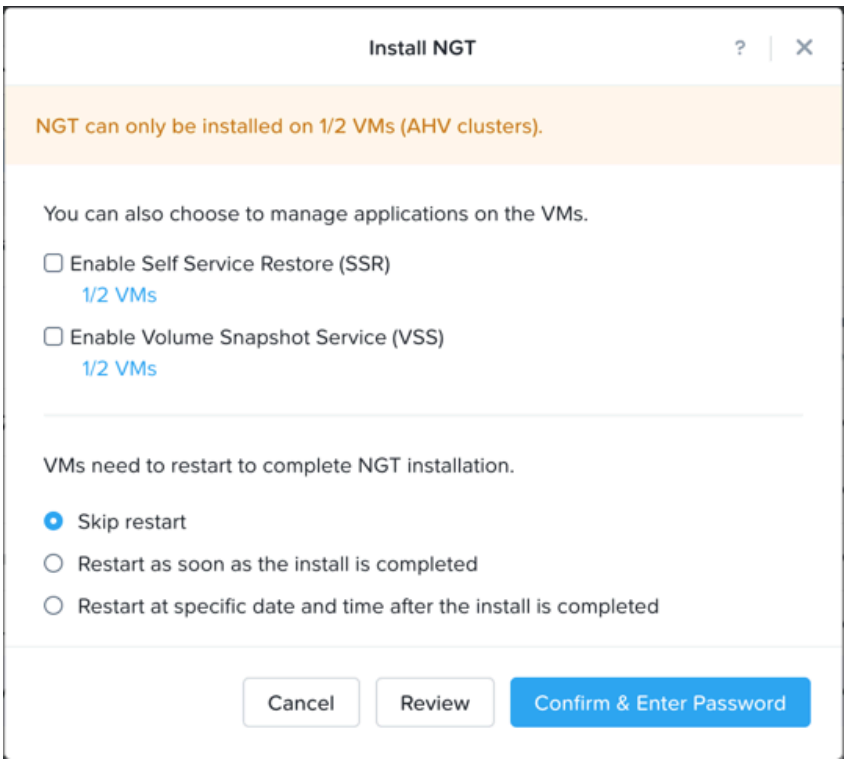

Figure 205: Installing NGT

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 329.
   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 377.
- **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.

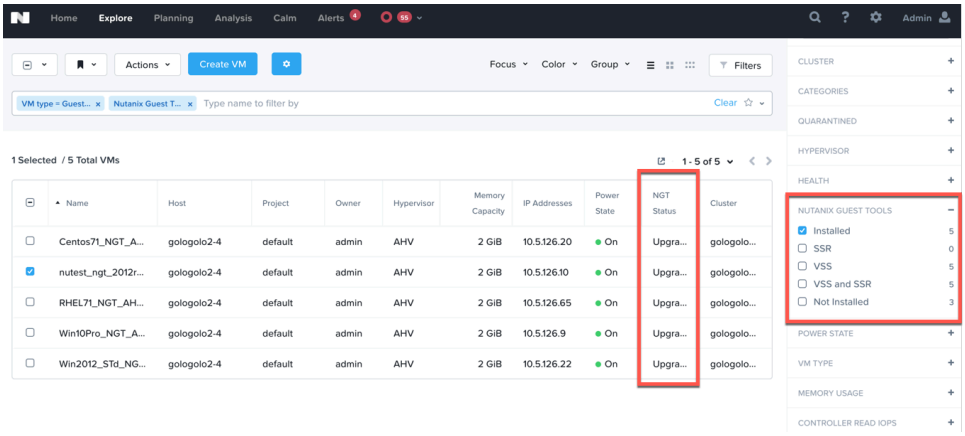

Figure 206: NGT Status and Filter Option

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 327. 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 72).

2. Select the VM or VMs on which you want to enable SSR and VSS applications by clicking **Manage NGT Applications**.

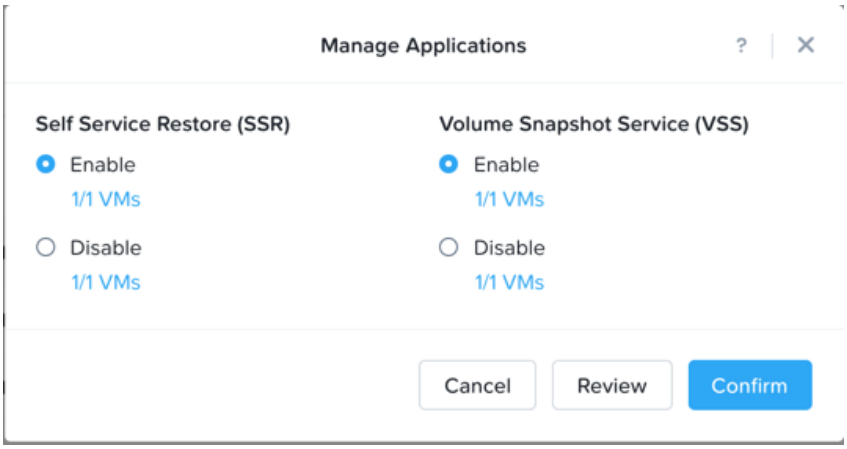

Figure 207: Manage NGT Applications

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 72).
   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.

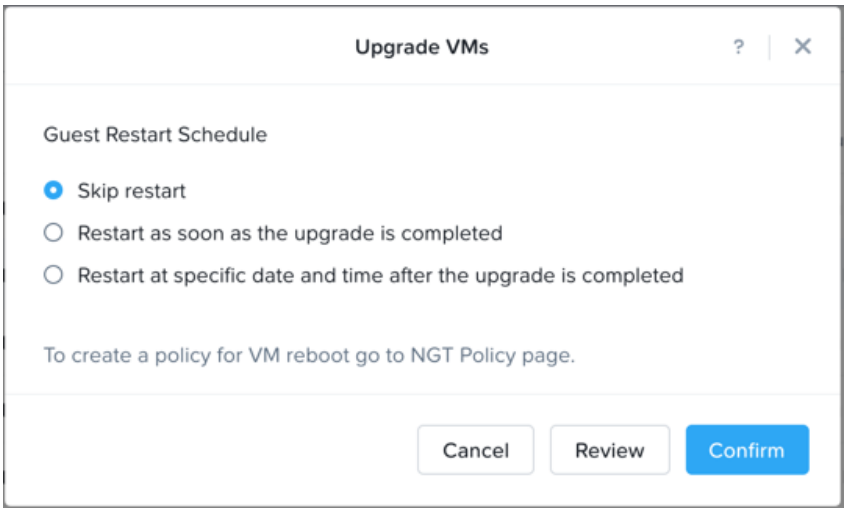

Figure 208: Upgrade NGT

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 377.
   - **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 professional or above license.
- 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 179: 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
64 KiB	400	25 MBps

Block size	IOPS	Throughput
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 AFS 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.

Procedure

1. Go to the **List** tab of the VM dashboard (see VMs Summary View on page 72) 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**.

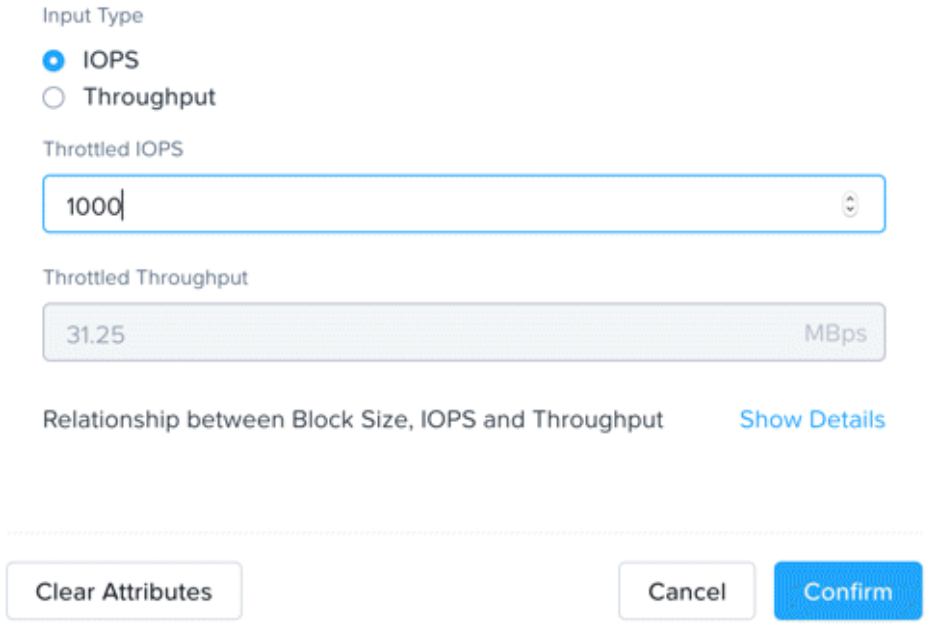

Figure 209: Setting QoS Attributes

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.

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 379).

Figure 210: 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

For information about these methods, see Adding an Image on page 336.

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 379.

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 336 and Uploading Images from a Remote Server on page 341, 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 347.

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`.
- If an image that you want to use is not active on a cluster, you might not be able 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.

Adding 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 336.
- Uploading Images from a Remote Server. See Uploading Images from a Remote Server on page 341.

Uploading Images from a Workstation

You can browse the workstation from which you are accessing Prism Central and upload images from it. 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. When a cluster must create a VM from an image that is available on another cluster, it downloads the image to its local storage.

About this task

> 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 98).
   The Add Images page appears.

2. Under **Image Source**, click **Image File**, and then click **Add File**.

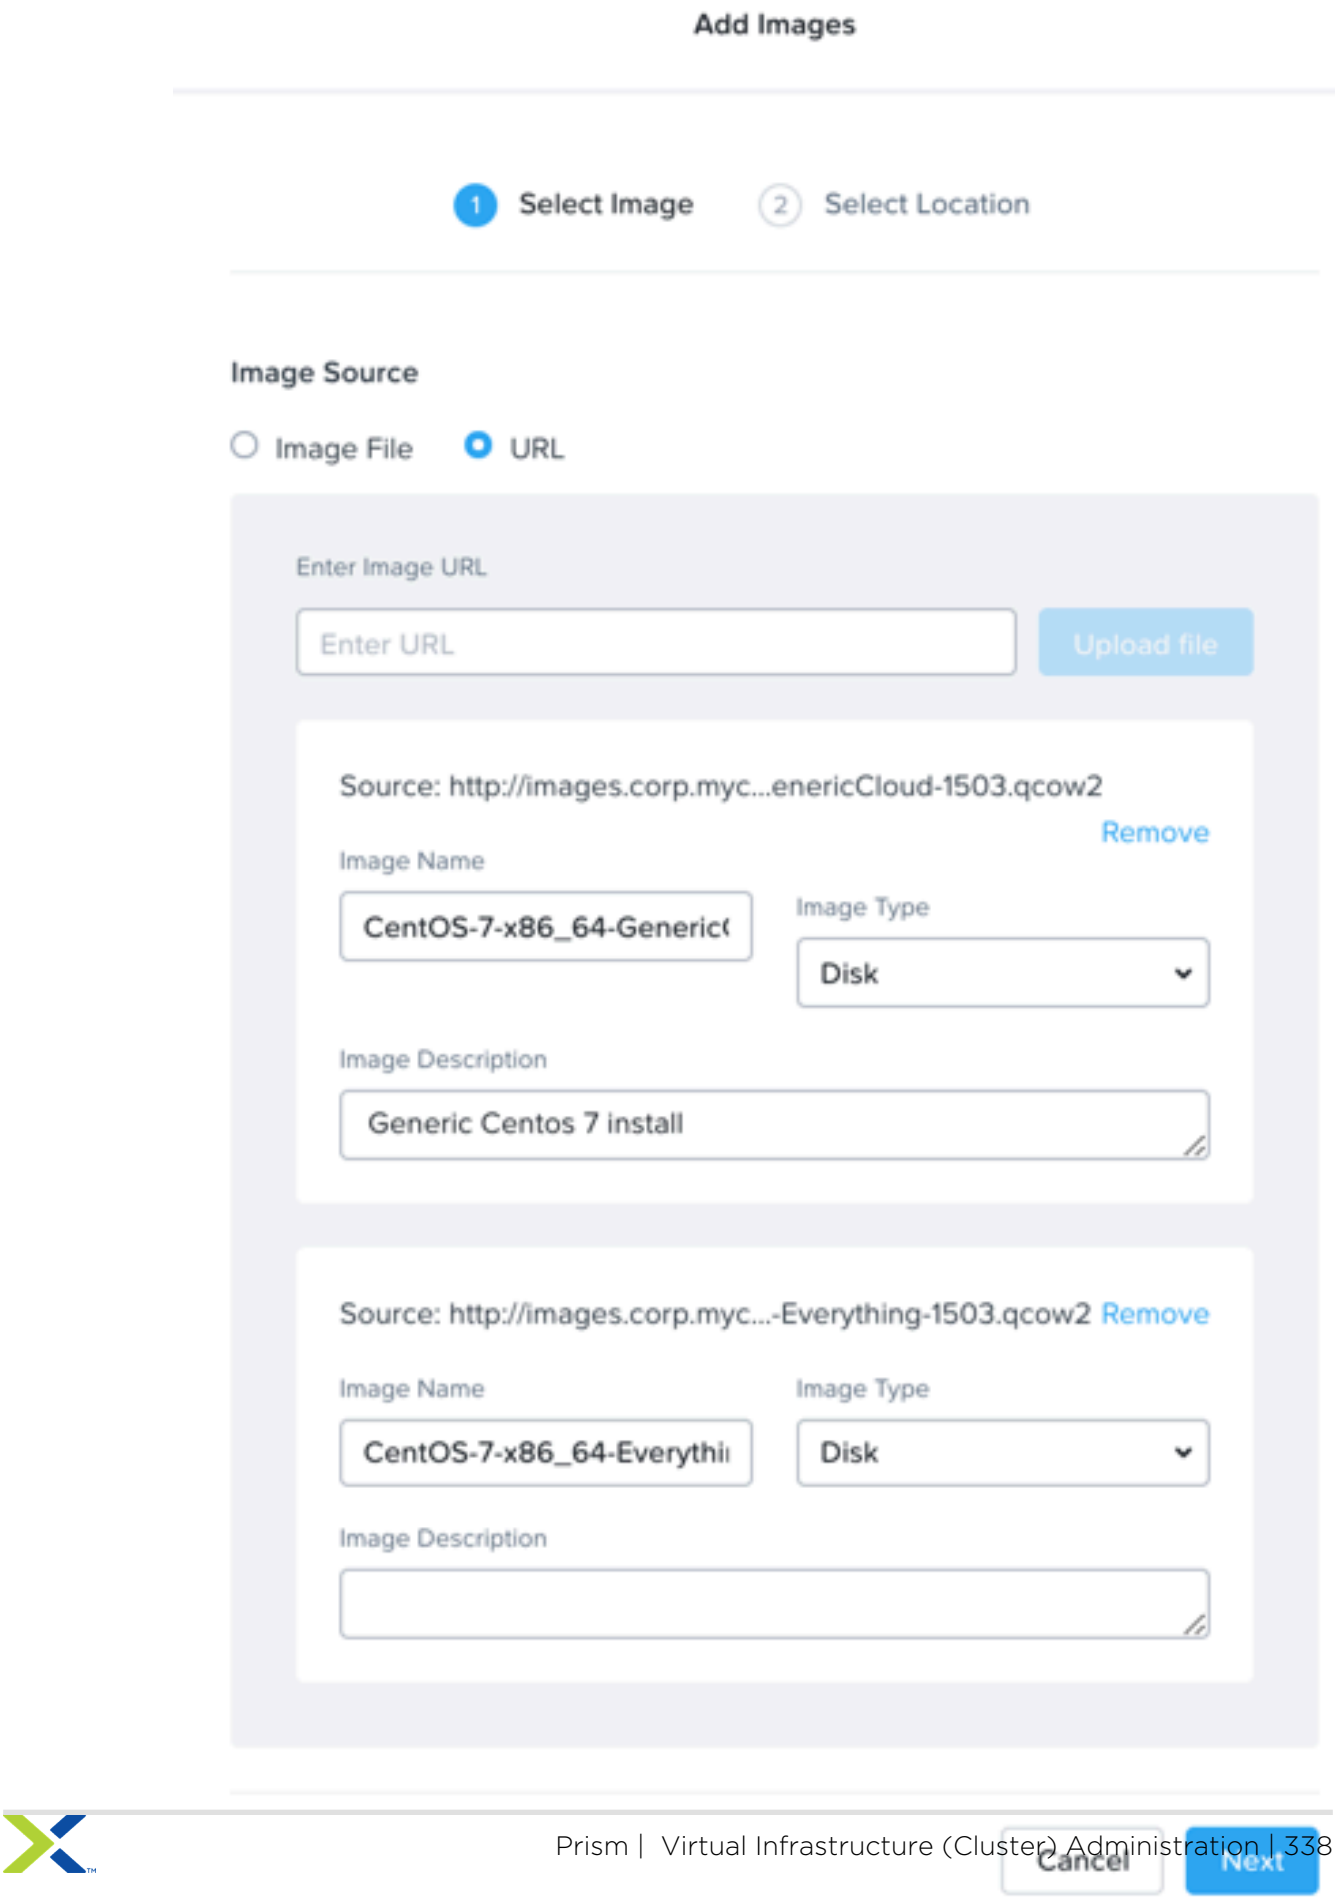
Add Images
1 Select Image
2 Select Location
Image Source
Image File
URL
Enter Image URL
Enter URL
Upload file
Source: http://images.corp.myc...enericCloud-1503.qcow2
Remove
Image Name
CentOS-7-x86_64-Generic(
Image Type
Disk
Image Description
Generic Centos 7 install
Source: http://images.corp.myc...-Everything-1503.qcow2 Remove
Image Name
CentOS-7-x86_64-Everythi
Image Type
Disk
Image Description
Cancel
Next

3. To specify one or more 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.

   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.

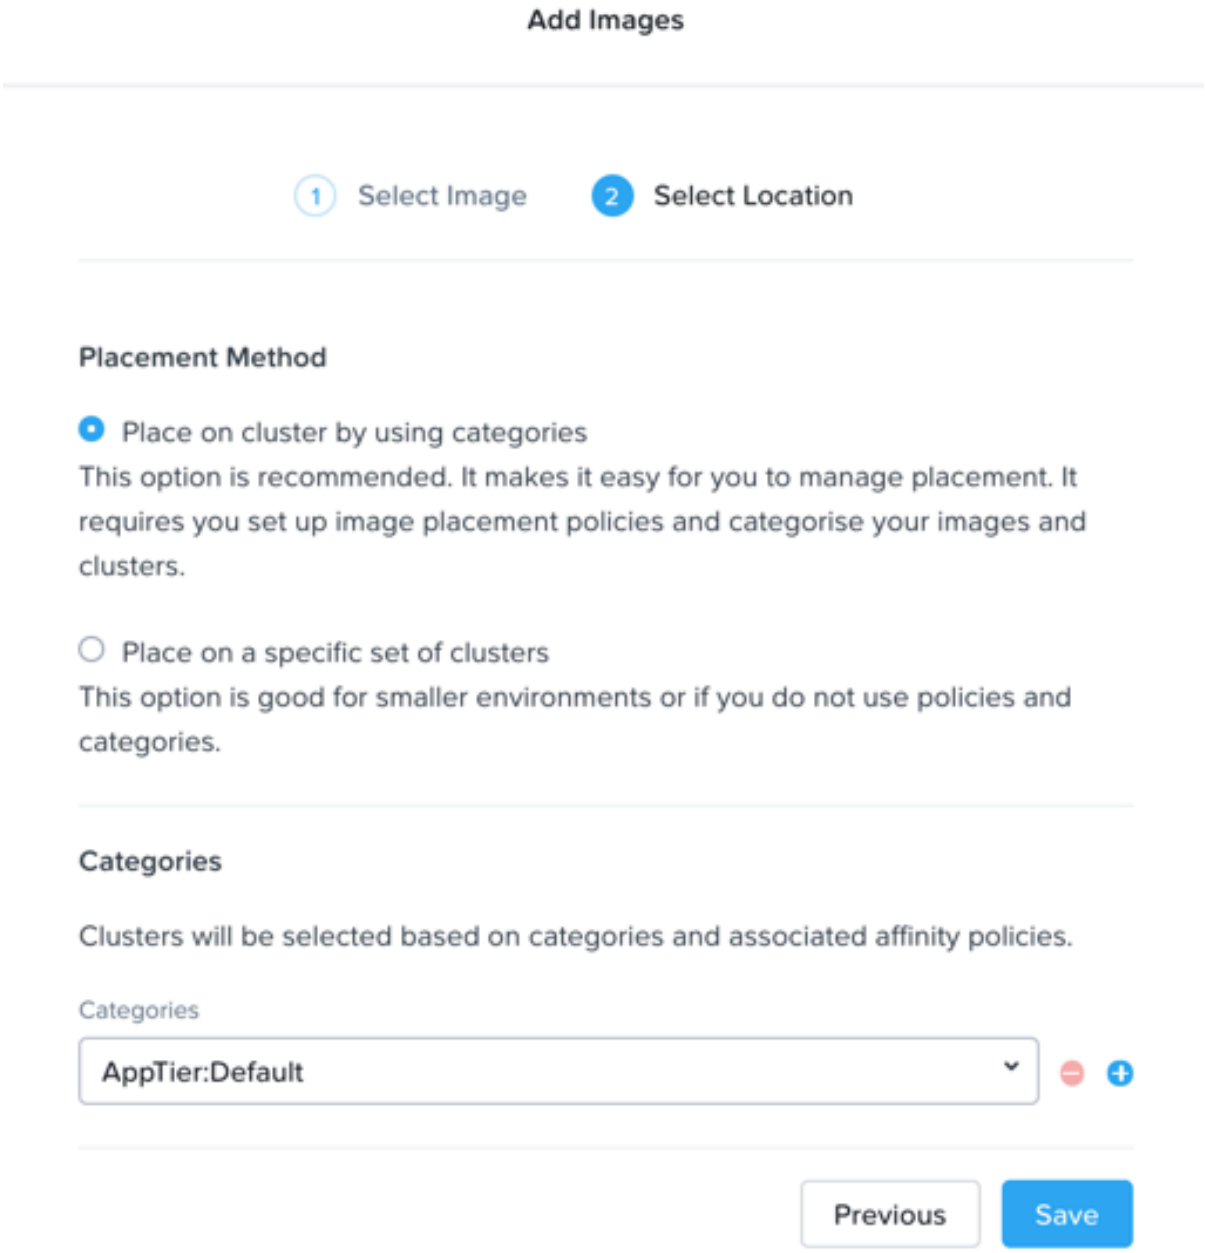

Figure 212: Add Images: Select Location

6. In **Placement Method**, do one of the following:

   - To assign categories to the images and delegate image placement decisions to configured policies, click **Place on cluster by using categories**, 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.

   - To manually select the target clusters, click **Place on a specific set of 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 **Place on all clusters registered to the PC**, and then click **Save**.
     - If you want to upload to only a subset of the registered clusters, clear **Place on all clusters registered to the PC**, select the clusters you want from the list, and then click **Save**.

   Prism Central uploads the image files in batches.

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 98).
   The Add Images page appears.

2. Under **Image Source**, click **URL**.

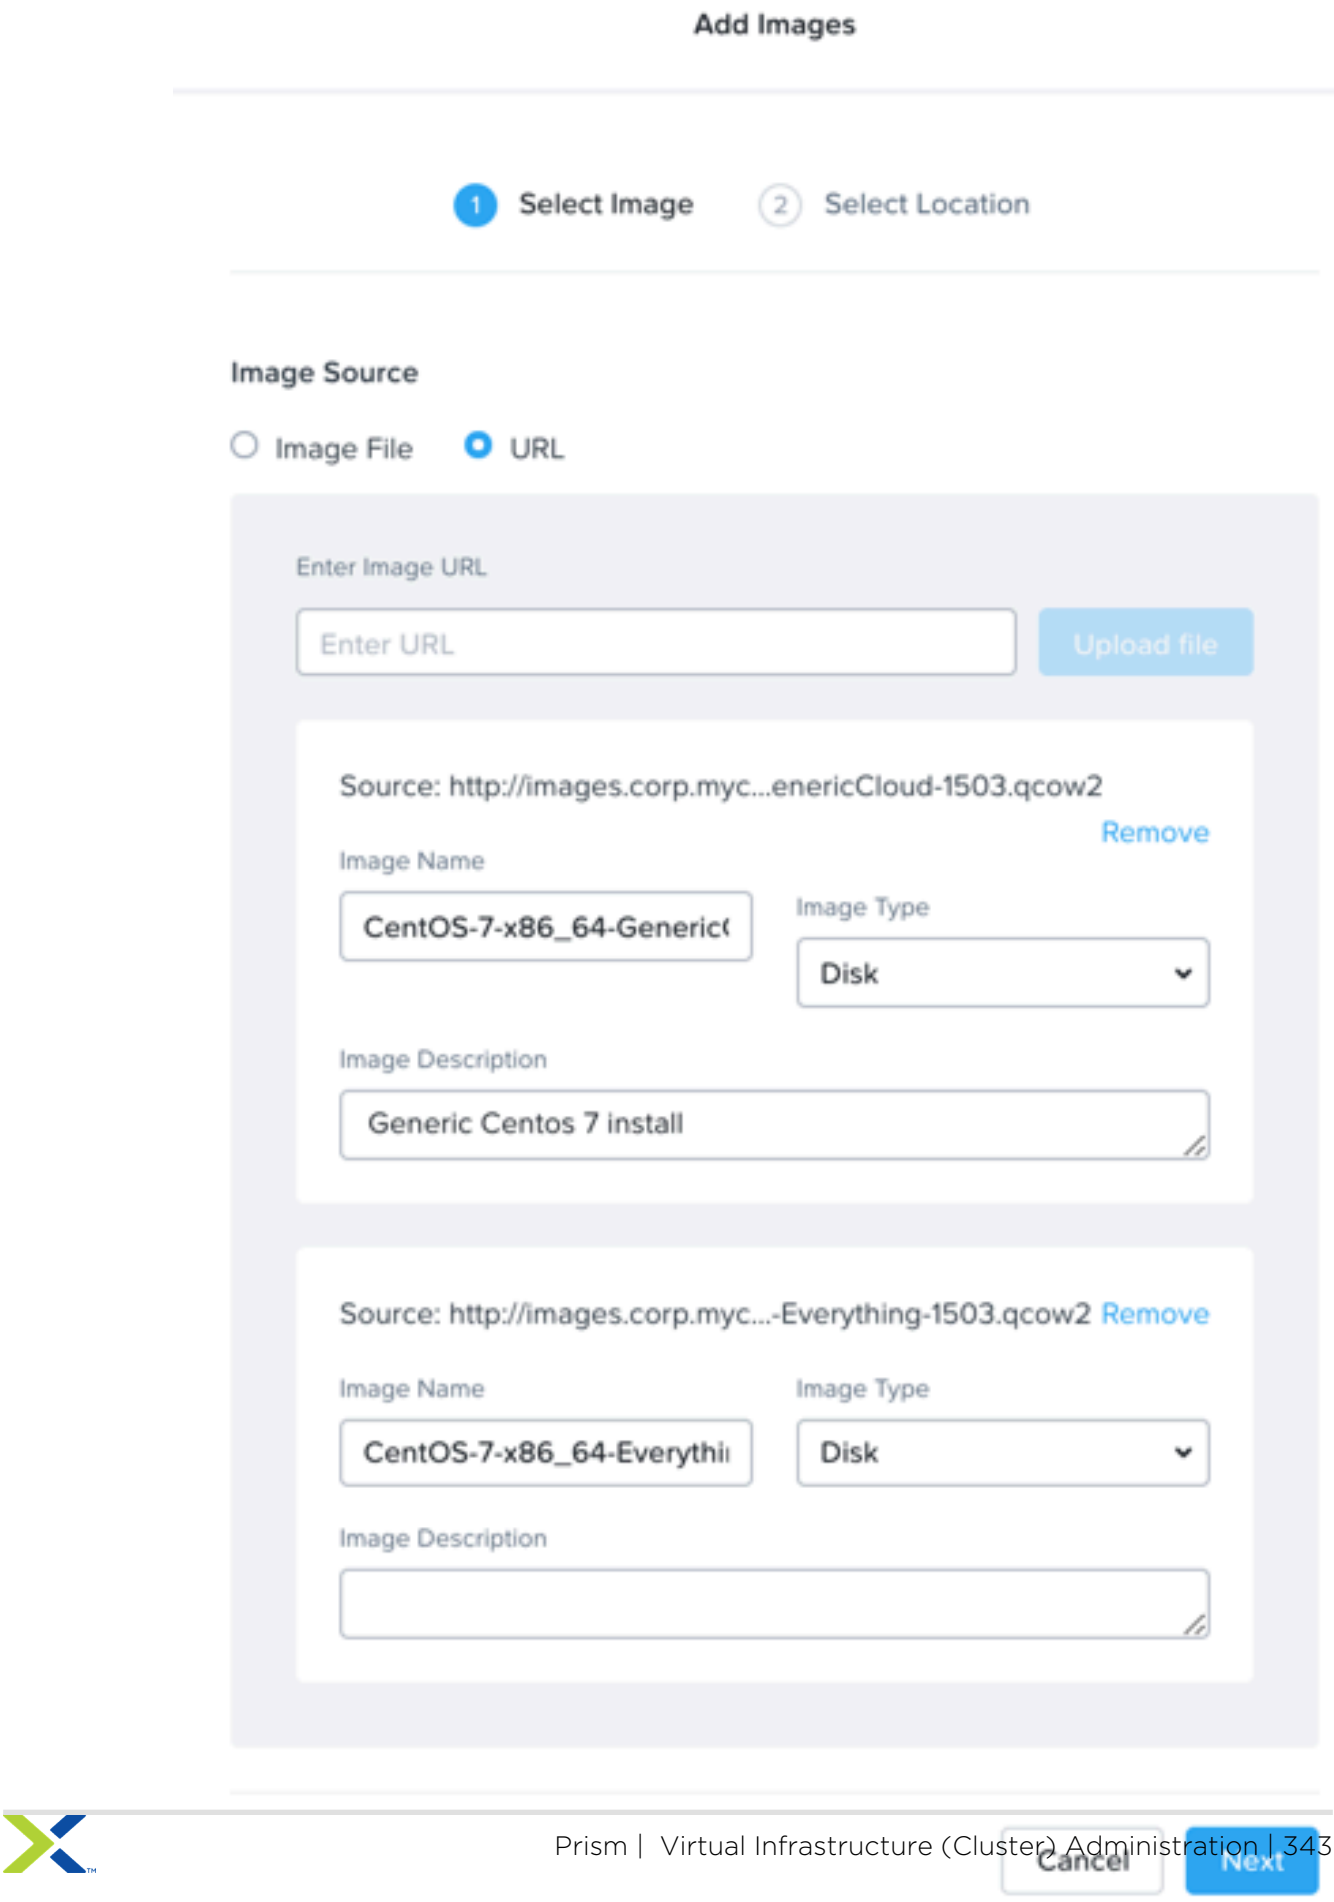
Add Images
1 Select Image
2 Select Location
Image Source
Image File
URL
Enter Image URL
Enter URL
Upload file
Source: http://images.corp.myc...enericCloud-1503.qcow2
Remove
Image Name
CentOS-7-x86_64-Generic(
Image Type
Disk
Image Description
Generic Centos 7 install
Source: http://images.corp.myc...-Everything-1503.qcow2 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 URL to the image, and then click **Upload File**.
   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.

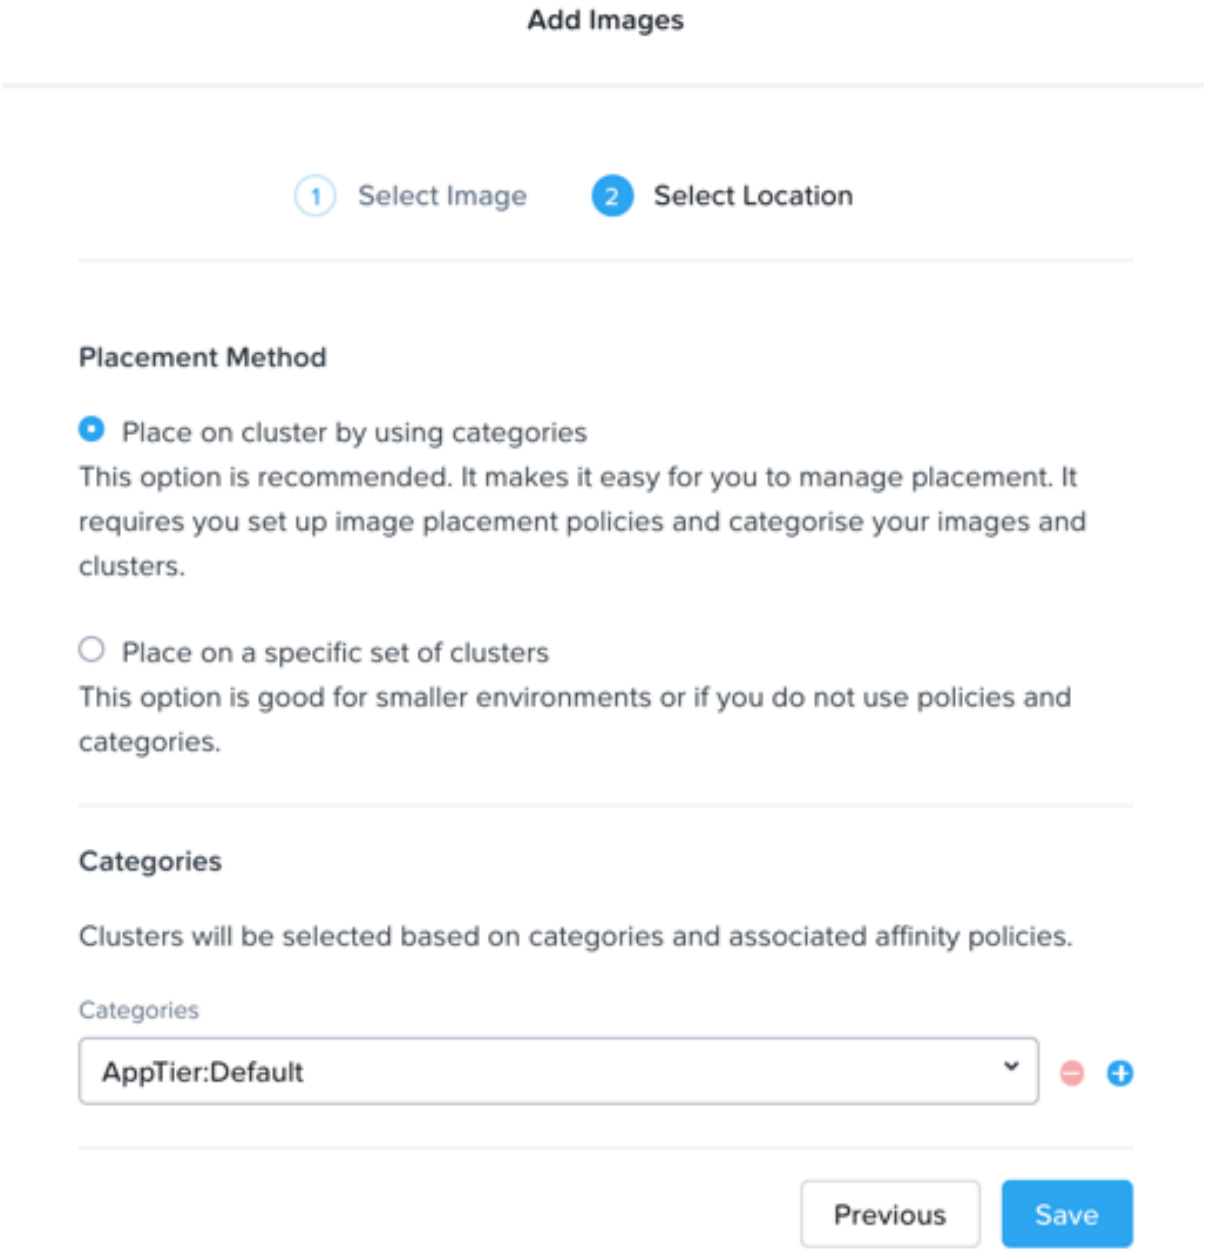

Figure 214: Add Images: Select Location

6. In **Placement Method**, do one of the following:

   - To assign categories to the images and entrust image placement to configured policies, click **Place on cluster by using categories**, 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.

   - To manually select the target clusters, click **Place on a specific set of 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 **Place on all clusters registered to the PC**, and then click **Save**.
     - If you want to upload to only a subset of the registered clusters, clear **Place on all clusters registered to the PC**, select the clusters you want from the list, and then click **Save**.

   Prism Central uploads the image files in batches.

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 98), 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.

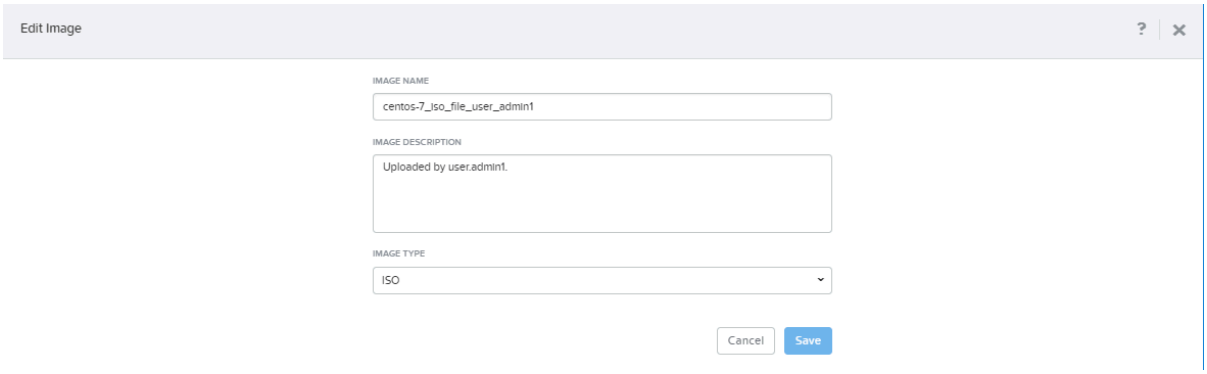

Figure 215: Update Image

- To add an image to the catalog (see Catalog Management on page 347), 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 98).
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.

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 368).

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 72).
   - » Go to the images dashboard and select the image that you want to add to the catalog (see Images Summary View on page 98).
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.

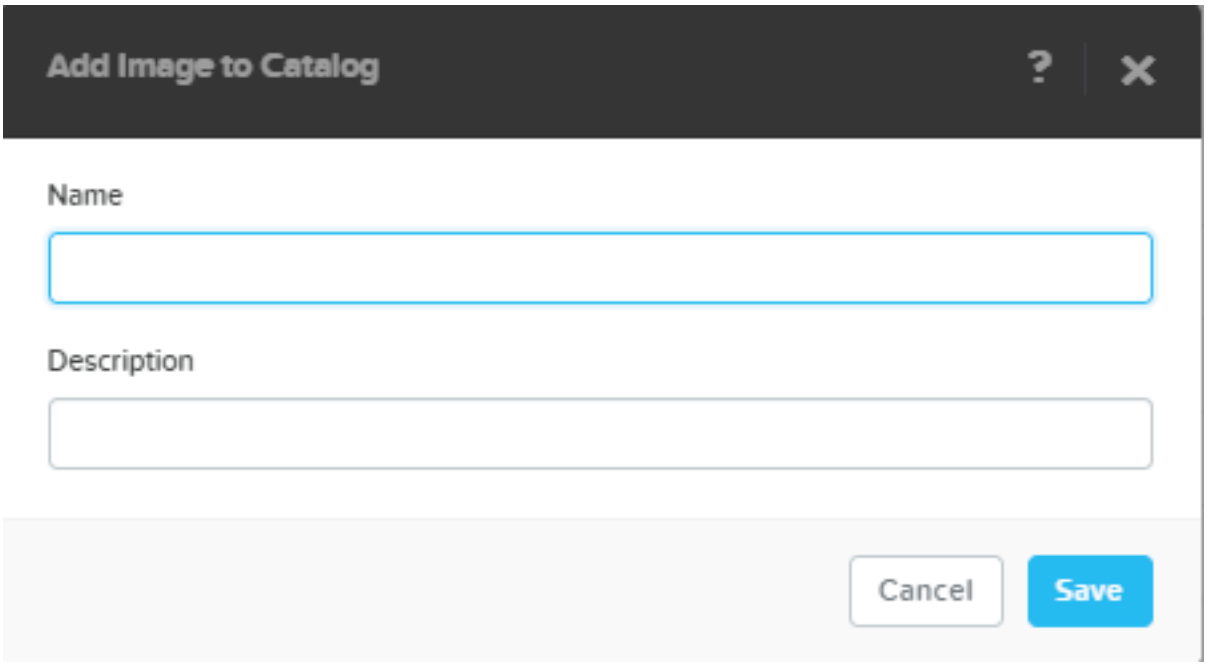

Figure 216: Add Image to Catalog Window

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.

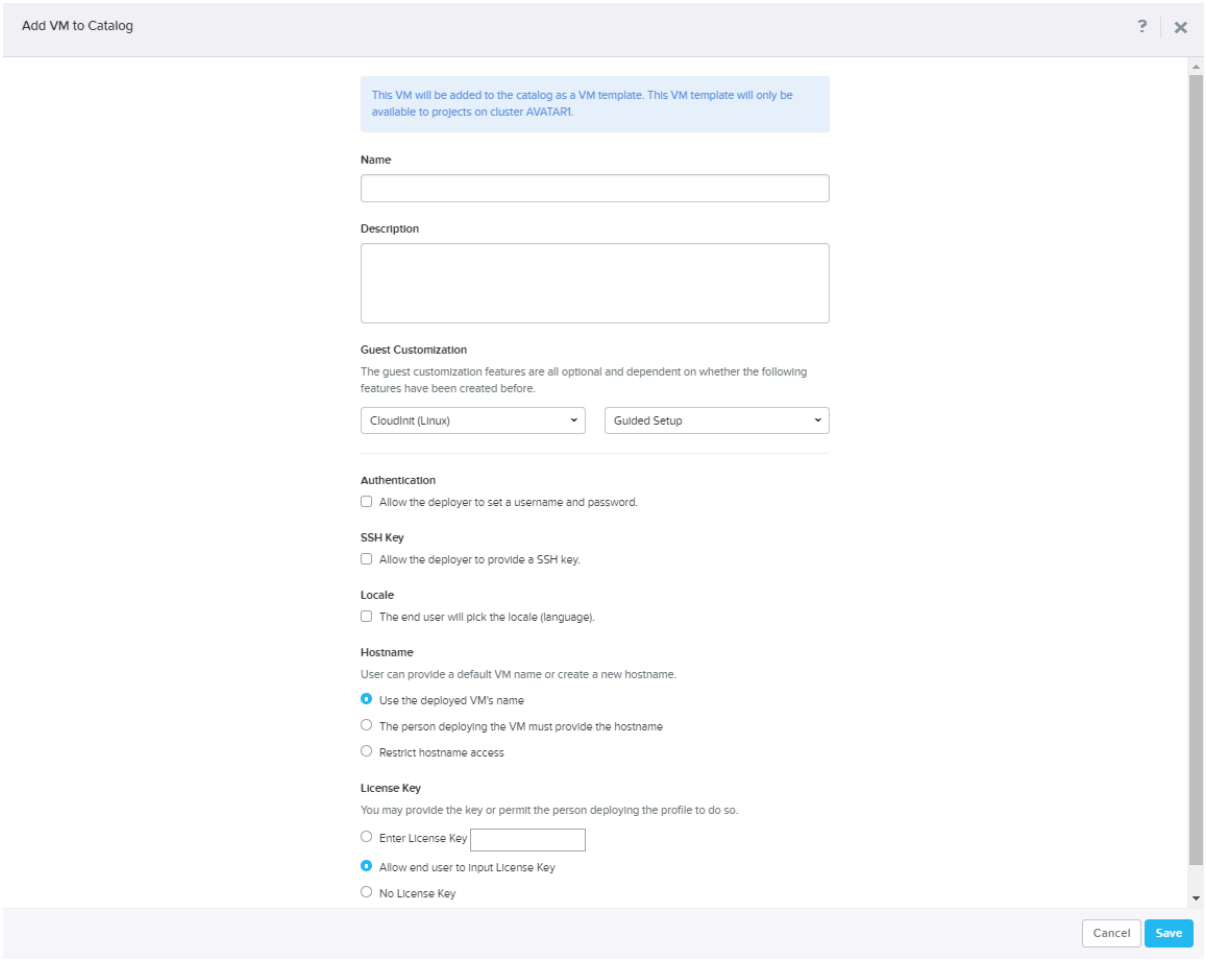

Figure 217: Add VM to Catalog Page

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 72).

   » Go to the images dashboard and select the image that you want to delete from the catalog (see Images Summary View on page 98).

2. From the **Actions** menu, select **Delete Catalog Item**.

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.

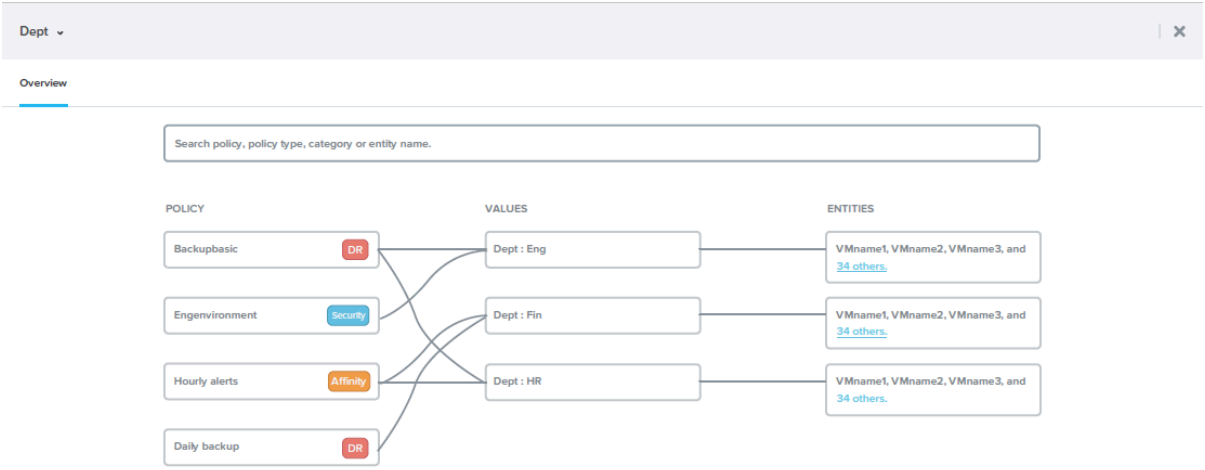

Figure 218: Dept Details View (example)

Creating a Category

About this task

To create a category, do the following:

Procedure

1. Go to the categories dashboard (see Categories Summary View on page 102) and click the **Create Category** button.
   The Create Category page appears.

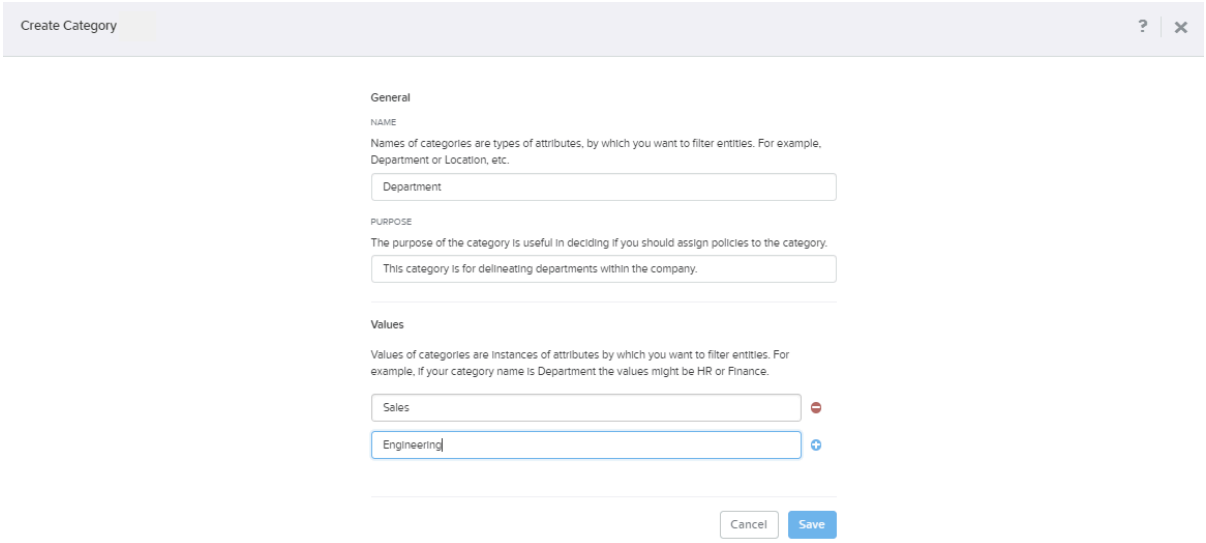

Figure 219: Create Category Page

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 102) 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 351). 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 114) 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 72) and select the target VMs.
   - Image: Go to the Images dashboard (see Images Summary View on page 98) and select the target images.
   - Subnet: Go to the Subnets dashboard (see Subnets View on page 105) 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.

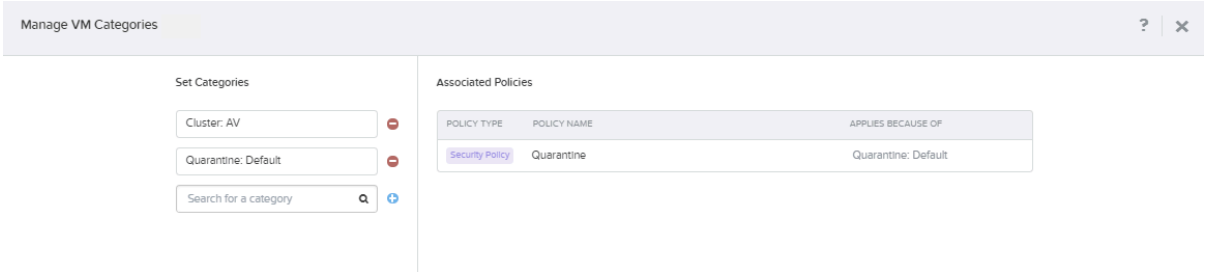

Figure 220: Manage VM Categories Page

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 297.

VM-VM Anti-Affinity Policy

You can use this policy to specify anti-affinity between the virtual machines. The VM-VM anti-affinity policy keeps the specified virtual machines apart in such a way that when a problem occurs with one host, you should not lose both virtual machines. However, this is a preferential policy. This policy does not limit the Acropolis Dynamic Scheduling (ADS) feature to take necessary action in case of resource constraints.

Note:

- Currently, you can only define VM-VM anti-affinity policy by using aCLI. For more information, see Configuring VM-VM Anti-Affinity Policy on page 355.
- Anti-Affinity policy is only applied during the initial placement of VMs (when a VM is powered on). Anti-Affinity policy is over-ridden when you manually migrate a VM to the same host or when the host is put in the maintenance mode. However, ADS will attempt to fix the violations periodically.
- The VM-VM affinity policy is not supported.

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.
- You cannot remove the VM-host affinity for a powered on VM from Prism. You can use the `vm.affinity_unset vm_list` aCLI command to perform this operation.

Configuring VM-VM Anti-Affinity Policy

Perform the following procedure to configure the VM-VM anti-affinity policy.

About this task

To configure VM-VM anti-affinity policies, you must first define a group and then add all the VMs on which you want to define VM-VM anti-affinity policy.

Note: Currently, the VM-VM affinity policy is not supported.

Procedure

1. Log in to the Controller VM in your cluster through an SSH session and access the Acropolis command line.
2. Create a group.

   ```
   acli> vm_group.create group_name
   ```

   Replace `group_name` with the name of the group.
3. Add the VMs on which you want to define anti-affinity to the group.

   ```
   acli> vm_group.add_vms group_name vm_list=vm_name
   ```

   Replace `group_name` with the name of the group. Replace `vm_name` with the name of the VMs that you want to define anti-affinity on.

4. Configure VM-VM anti-affinity policy.

   ```
   acli> vm_group.antiaffinity_set group_name
   ```

   Replace *group_name* with the name of the group.

   After you configure the group and then power on the VMs, the VMs that are part of the group are started (attempt to start) on the different hosts. However, this is a preferential policy. This policy does not limit the Acropolis Dynamic Scheduling (ADS) feature to take necessary action in case of resource constraints.

Removing VM-VM Anti-Affinity Policy

Perform the following procedure to remove the VM-VM anti-affinity policy.

Procedure

1. Log in to the Controller VM in your cluster through an SSH session and access the Acropolis command line.
2. Remove the VM-VM anti-affinity policy.

   ```
   acli> vm_group.antiaffinity_unset group_name
   ```

   Replace *group_name* with the name of the group.

   The VM-VM anti-affinity policy is removed for the VMs that are present in the group, and they can start on any host during the next power on operation (as necessitated by the ADS feature).

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. To display the Network Configuration page, click the **Network Config** button in either the VMs or Subnets dashboard (see VMs Summary View on page 72 or Subnets View on page 105).

The Network Configuration page includes two tabs.

- The **Virtual Networks** tab displays a list of the configured networks.

Figure 221: Virtual Networks Tab

- The **Internal Interfaces** tab displays a list of LAN interfaces.

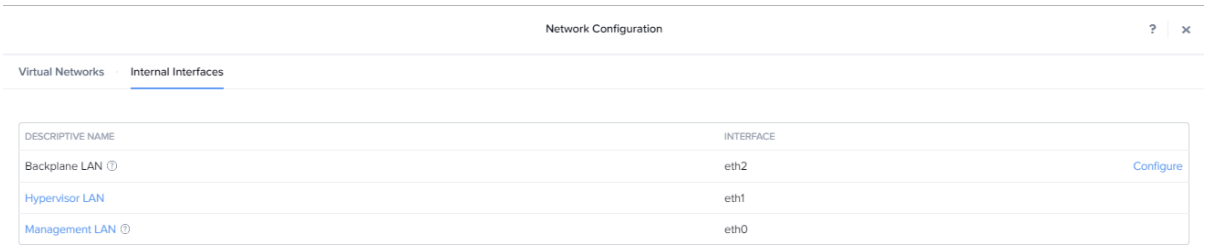

Figure 222: Internal Interfaces Tab

The following table describes the fields in each tab.

Table 180: Network Configuration Fields

Parameter	Description	Values
Virtual Networks Tab		
Name	Displays the name of the virtual network.	(name)
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)
Interface	Displays the interface designation such as eth0 or eth1.	(interface name)

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 72) and click the **Network Config** button.

   » Go to the Subnets dashboard (see Subnets View on page 105 and click the **Network Config** button.

2. In the Cluster Selection window, select the target cluster.

   A list of registered clusters appears in the window. Clicking on a cluster name displays the Network Configuration page for that cluster (see Network Connections on page 356).

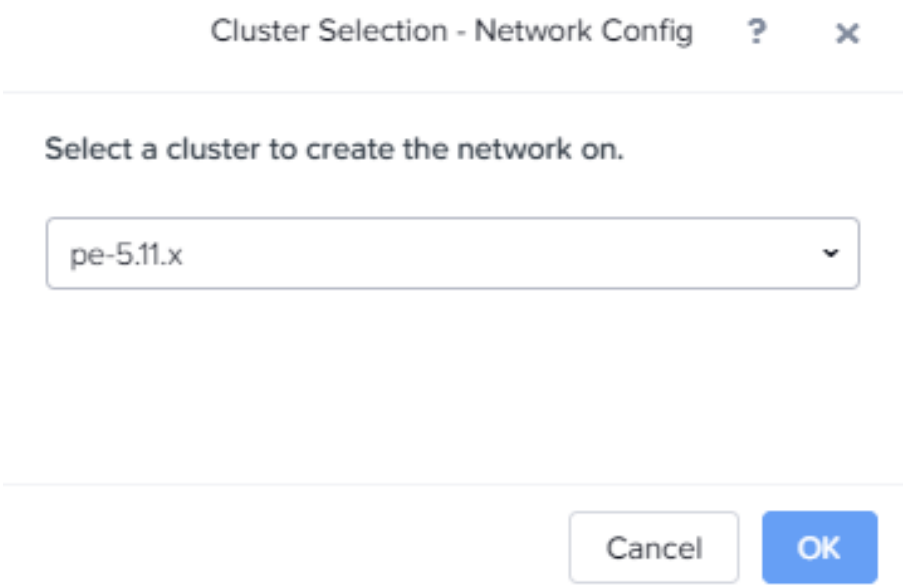

Figure 223: Cluster Selection Window

3. In the Virtual Networks tab, click the **Create Network** button.
   This displays the Create Network page.

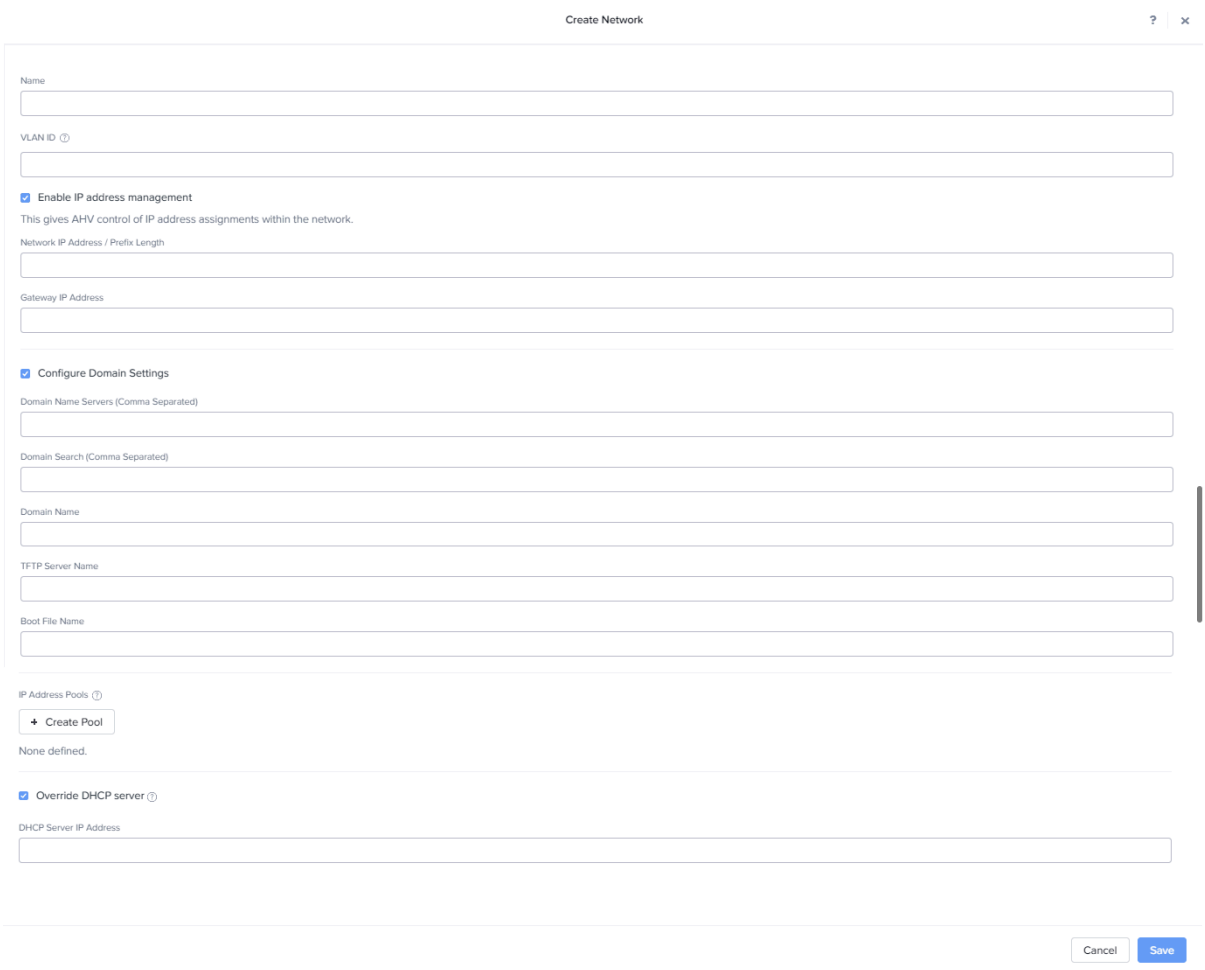

Figure 224: Create Network Page (expanded)

4. Do the following in the indicated fields:

   a. **Name**: Enter a name for the network.

   b. **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.

   c. **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.

   d. **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).

   e. **Gateway IP Address**: Enter the VLAN default gateway IP address.

   f. **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.

   g. **Domain Name Servers (comma separated)**: Enter a comma-delimited list of DNS servers.

   h. **Domain Search (comma separated)**: Enter a comma-delimited list of domains.

   i. **Domain Name**: Enter the VLAN domain name.

   j. **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).

   k. **Boot File Name**: Name of the boot file to download from the TFTP server.

5. 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.

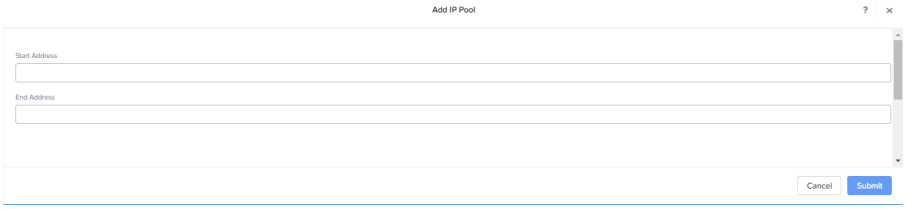

Figure 225: Add IP Pool Page

   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.

6. 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`.

7. When all the information is correct, click the **Save** button to configure the network connection, close the Create Network page, and return to the Network Configuration page.

8. Click the **Close** button to close the Network Configuration page.

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 72) and click the **Network Config** button.

   » Go to the Subnets dashboard (see Subnets View on page 105 and click the **Network Config** button.

2. In the Cluster Selection window, select the target cluster.

   A list of registered clusters appears in the window. Clicking on a cluster name displays the Network Configuration page for that cluster (see Network Connections on page 356).

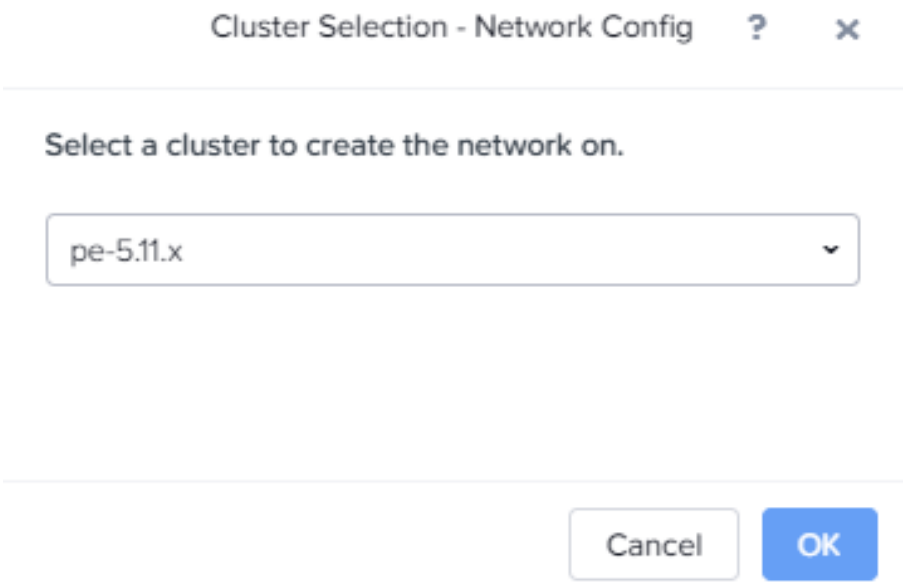

Figure 226: Cluster Selection Window

3. To modify a network configuration, open the Virtual Network tab, select the target line, and click the pencil icon.

   The Update Network page appears, which contains the same fields as the Create Network page (see Configuring Network Connections on page 357). 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.

4. To delete a network configuration, open the Virtual Network tab, select the target line, 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.

5. Click the **Close** button to close the Network Configuration page.

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 296). 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 114).

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 181: User Menu List

Name	Description
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page.

Name	Description
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 364. 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 367.

  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 14).

   - 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.

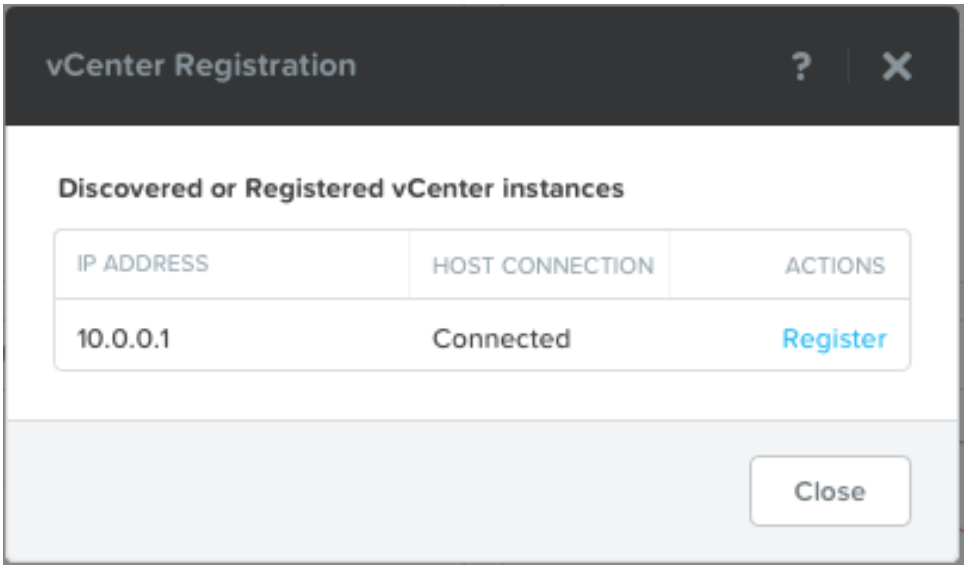

Figure 227: vCenter Registration

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.

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**.

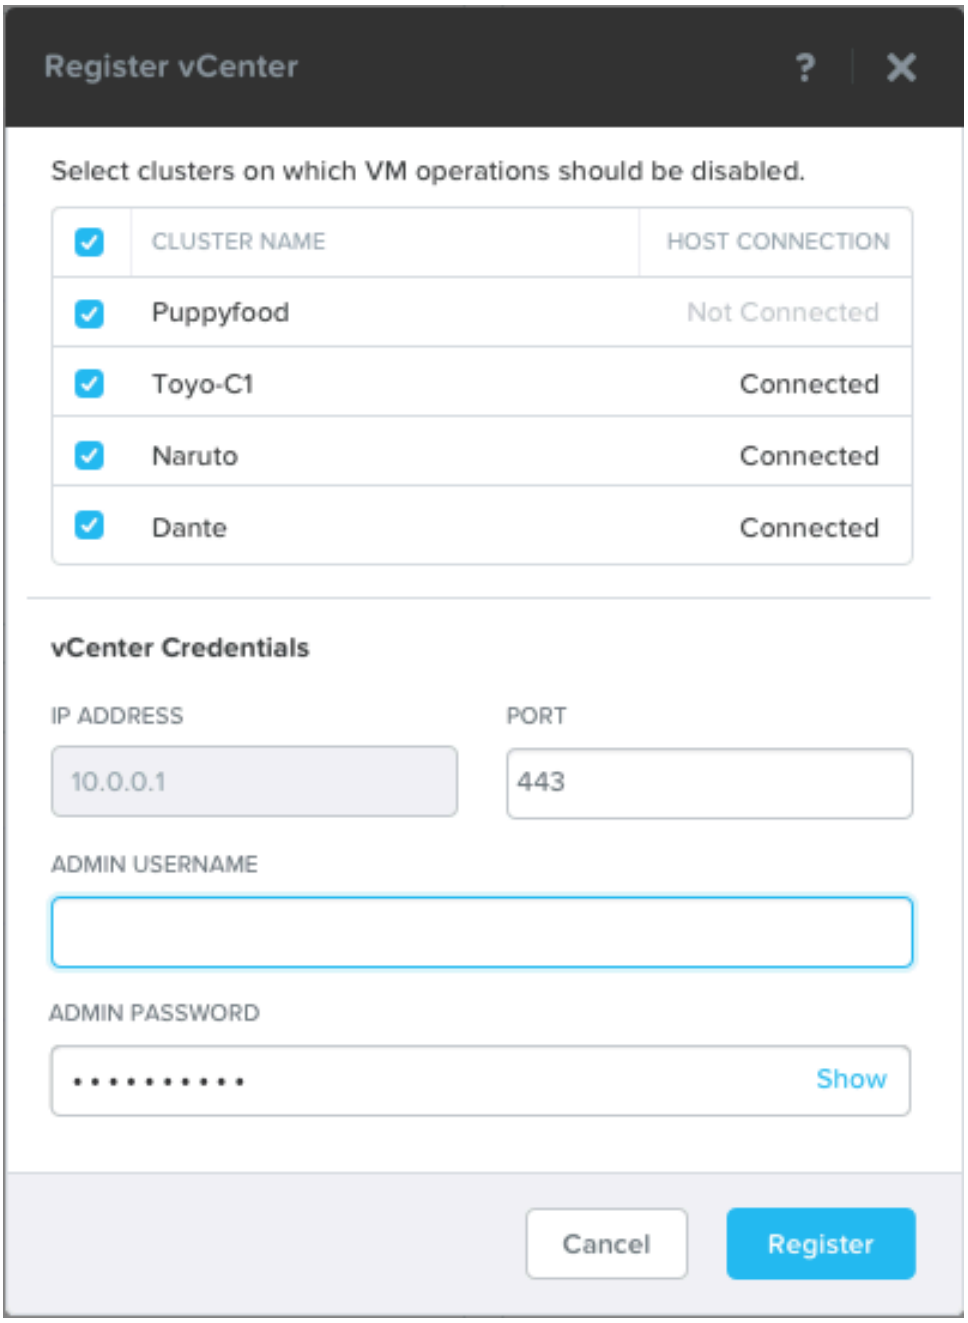

Figure 228: Cluster Selection

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 367.

Procedure

1. Click the gear icon in the main menu and select **vCenter Registration** from the Settings menu (see Settings Menu on page 14).

   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 14).

   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 367.

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 364.

12

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 87) 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 369).
- Create a project for each team that needs self service and add users and groups to the projects (see Project Management on page 373).
- 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 347).
- 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 161 and Project Details View on page 163).
- Create VMs as needed and assign them to project members as appropriate (see VM Management on page 296). 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 372.

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).

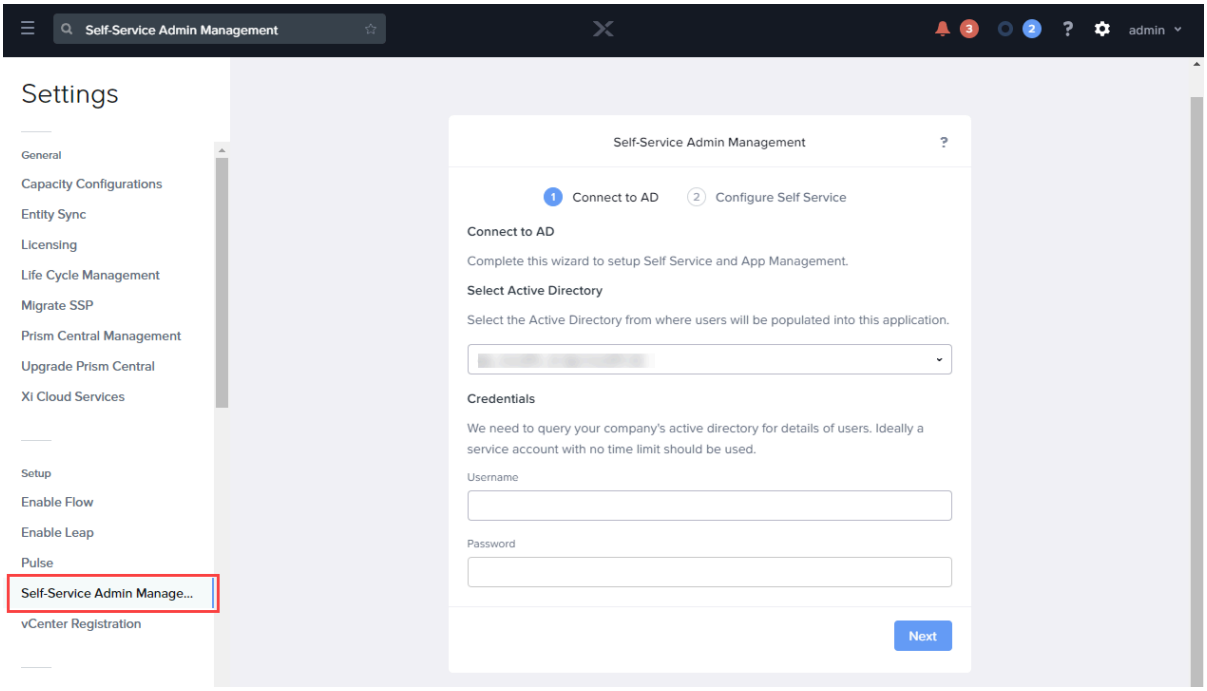

Figure 229: Self-Service Admin Management Window

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.

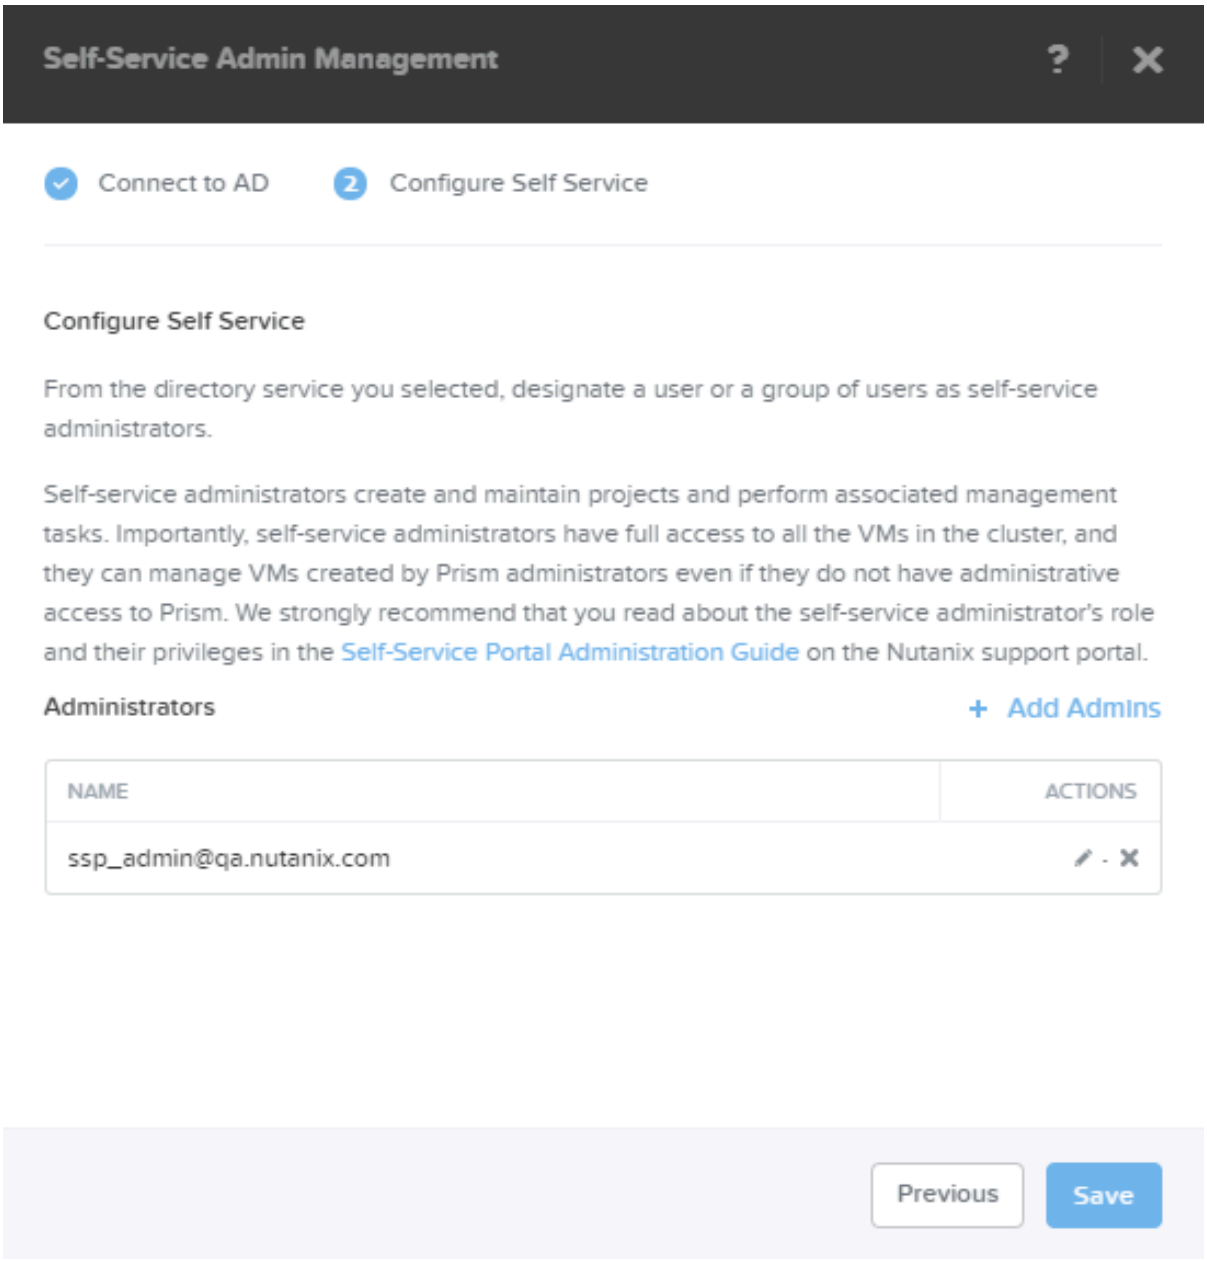

Figure 230: Configure Self Service tab

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.

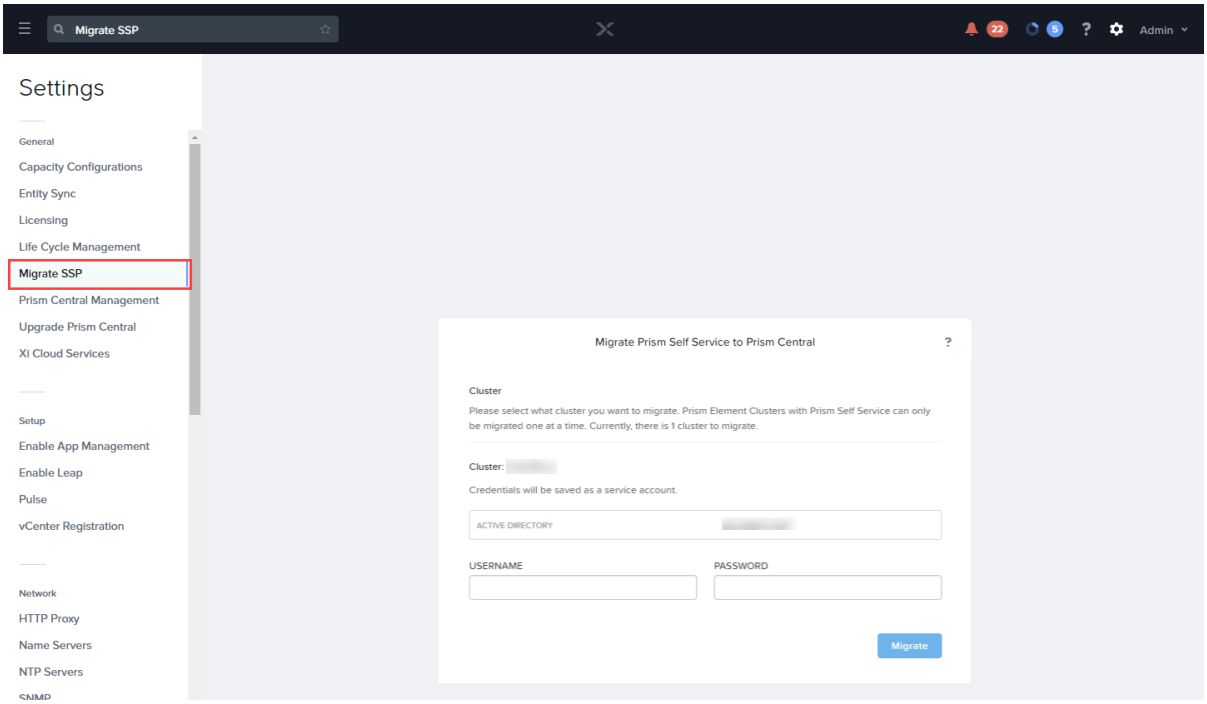

Figure 231: Migrate Prism Self Service to Prism Central Window

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 157).

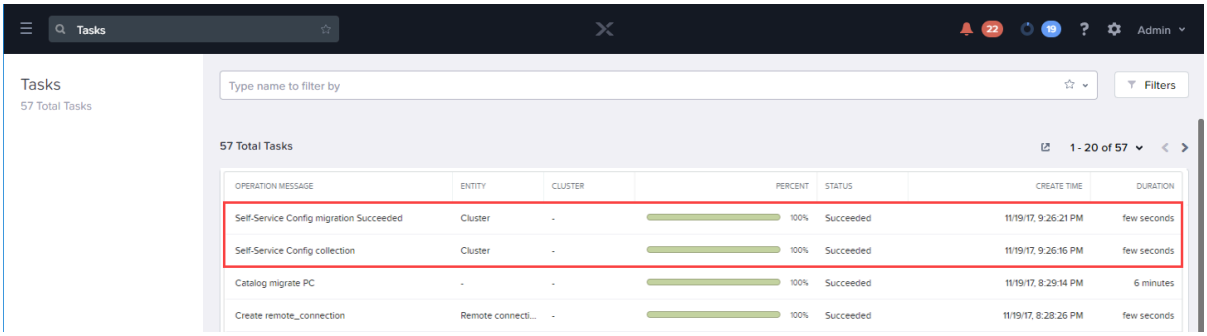

Figure 232: SSP Migration Tasks

Project Management

A project defines a set of Active Directory with a common set of requirements or a common function, such as a team of people collaborating on an engineering project. The project also specifies the roles to associate with its members, networks that they can use, and (optionally) 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 161) and click the **Create Project** button.
   The Create Project page appears.

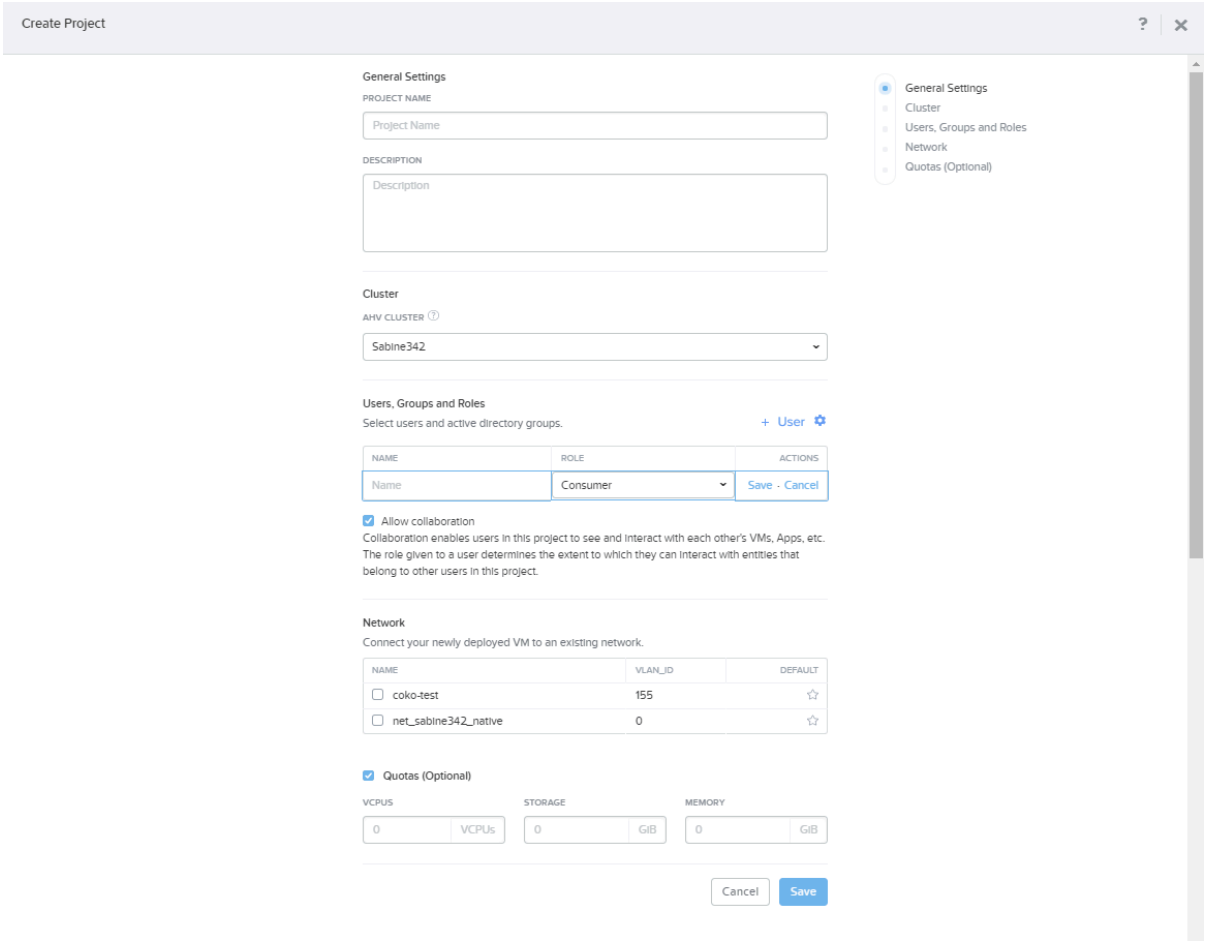

Figure 233: Create Project Page

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.

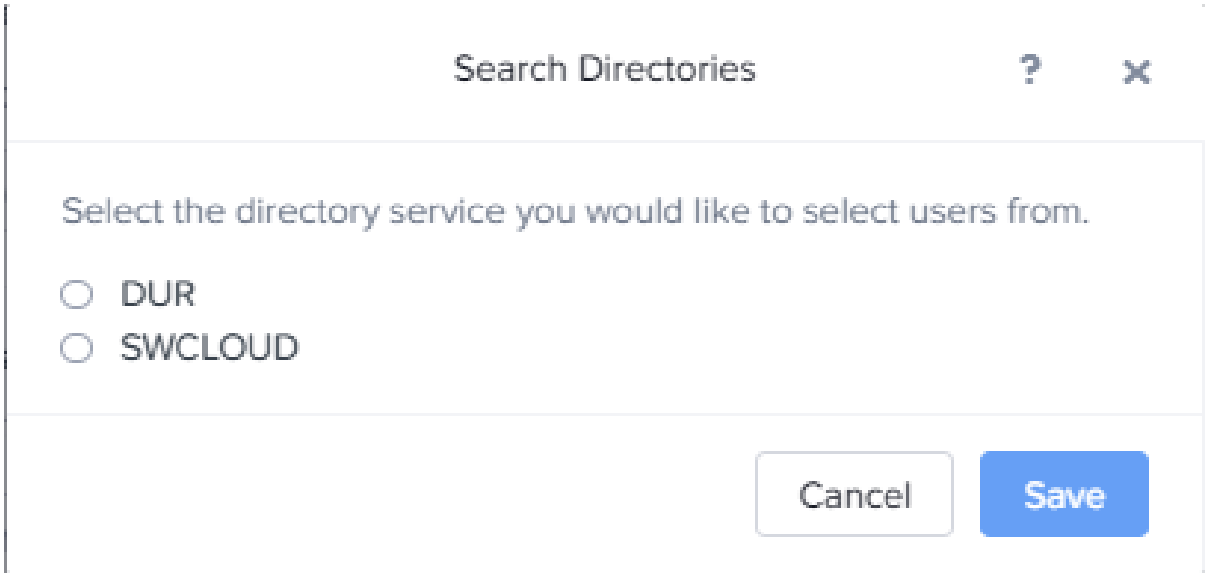

Figure 234: Search Directories Window

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 161) 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 373). 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.

13

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.

- "" () (in the *Nutanix Security Guide*) describes how to create and apply security-related policies.
- NGT Policies on page 377 describes how to create and apply NGT-related policies.
- Image Placement Policies on page 379 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 111) and click the **New VM Reboot Policy** button.

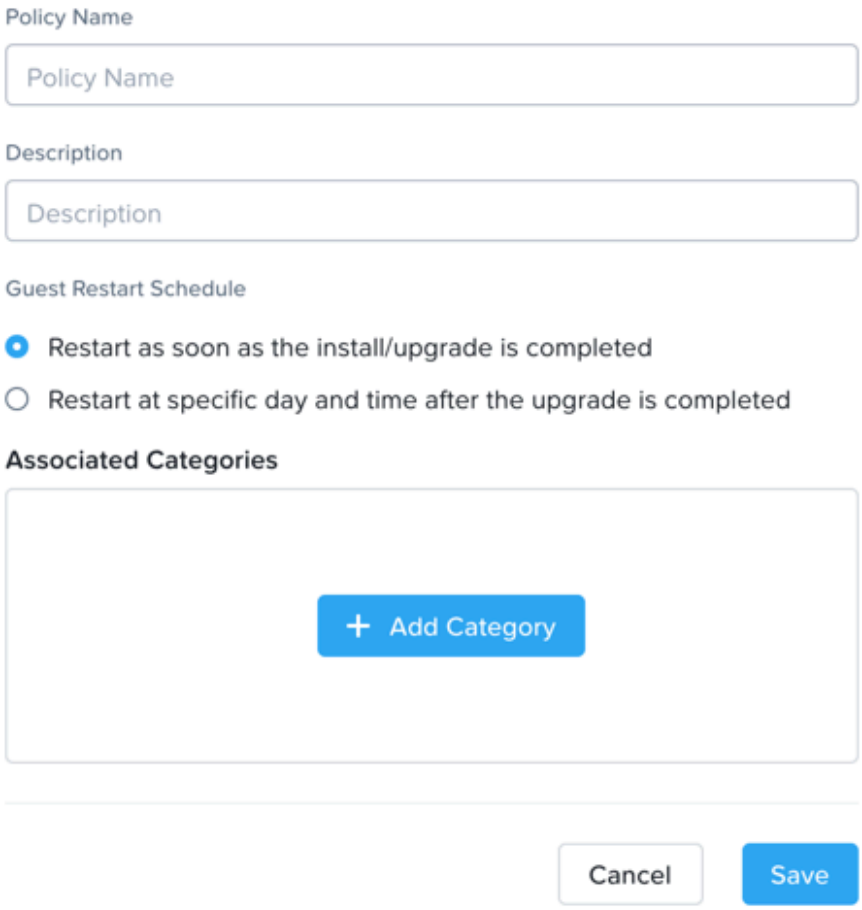

Figure 235: NGT Policy

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 111).
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 through the use of categories associated with both those entities. For example, an image placement policy might state that Linux images associated with the categories `OS:Linux` and `distribution:Centos` must be uploaded to clusters associated with the category l`ocation:SFO`.

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.

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 a registered Nutanix cluster, 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. Prism Central displays an alert highlighting the conflict.
- 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 182: 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.

14

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 etc.

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 386.

> 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, copyright information in the footer, a cover page, header, and footer.

- 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 and Limitations

- License requirements: Prism Central should be running Prism Pro license.

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 159), 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 393.
3. Configure a schedule for the report. For more information about scheduling, see Scheduling a Report on page 409.
4. Configure the settings specific for this report. For more information about configuring report settings, Configuring Report Settings on page 411.

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**.

     This option saves and generates the report. A **Run Report** window appears. For more information, see Generating a Report on page 387.

   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).

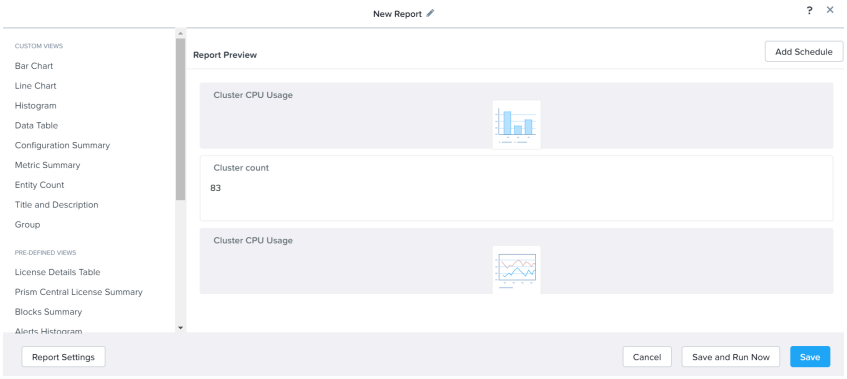

Figure 236: Create Report

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 387.
- To edit a report, see Editing a Report on page 388.
- To clone a report, see Cloning a Report on page 389.
- To delete a report, see Cloning a Report on page 389.

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 159), 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 411.

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 159), 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 386 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 405.

     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 159), 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 386 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 159), 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 159), 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.

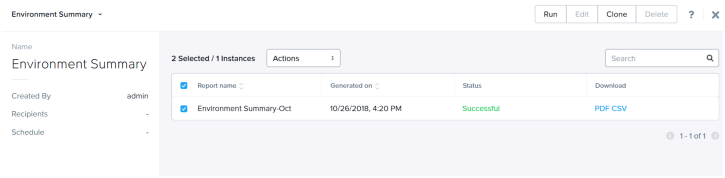

Figure 237: Viewing Instance Information

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, date when the report was generated, copyright, page number, logo, and so on).

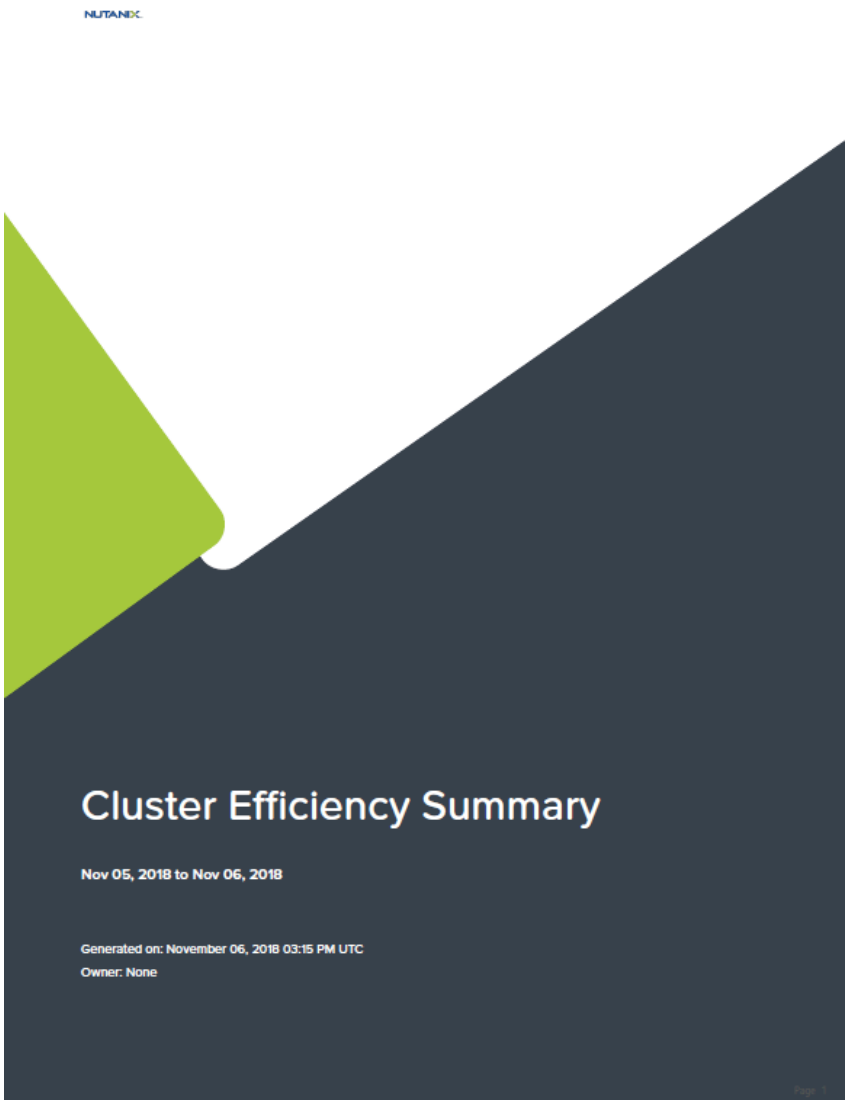

Figure 238: Stylistic Information

- 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.

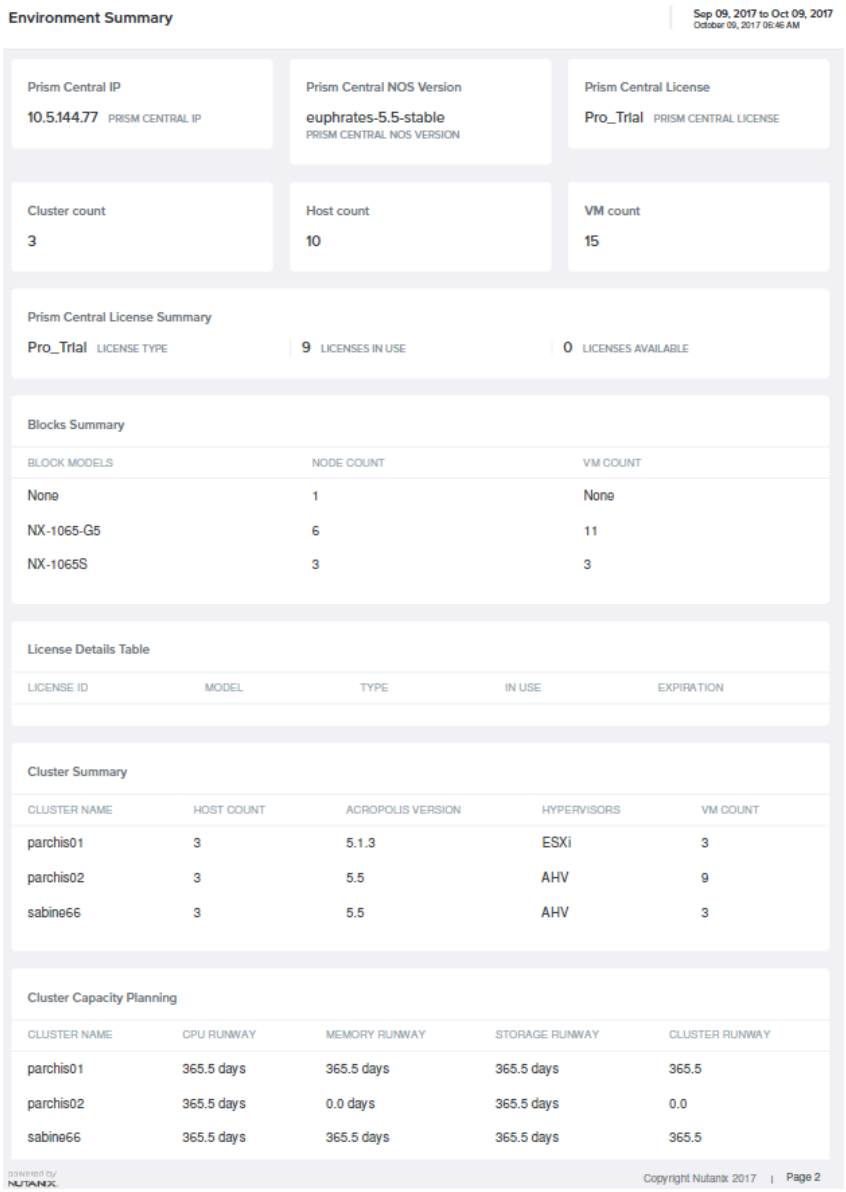

Environment Summary

Sep 09, 2017 to Oct 09, 2017
October 09, 2017 08:46 AM

Prism Central IP	Prism Central NOS Version	Prism Central License
10.5.144.77 PRISM CENTRAL IP	euphrates-5.5-stable PRISM CENTRAL NOS VERSION	Pro_Trial PRISM CENTRAL LICENSE

Cluster count	Host count	VM count
3	10	15

Prism Central License Summary

Pro_Trial LICENSE TYPE | 9 LICENSES IN USE | 0 LICENSES AVAILABLE

Blocks Summary

BLOCK MODELS	NODE COUNT	VM COUNT
None	1	None
NX-1065-G5	6	11
NX-1065S	3	3

License Details Table

LICENSE ID	MODEL	TYPE	IN USE	EXPIRATION

Cluster Summary

CLUSTER NAME	HOST COUNT	ACROPOLIS VERSION	HYPERVISORS	VM COUNT
parchis01	3	5.1.3	ESXi	3
parchis02	3	5.5	AHV	9
sabine66	3	5.5	AHV	3

Cluster Capacity Planning

CLUSTER NAME	CPU RUNWAY	MEMORY RUNWAY	STORAGE RUNWAY	CLUSTER RUNWAY
parchis01	365.5 days	365.5 days	365.5 days	365.5
parchis02	365.5 days	0.0 days	365.5 days	0.0
sabine66	365.5 days	365.5 days	365.5 days	365.5

powered by NUTANIX

Copyright Nutanix 2017 | Page 2

Figure 239: Detailed Metrics Information (Image 1)

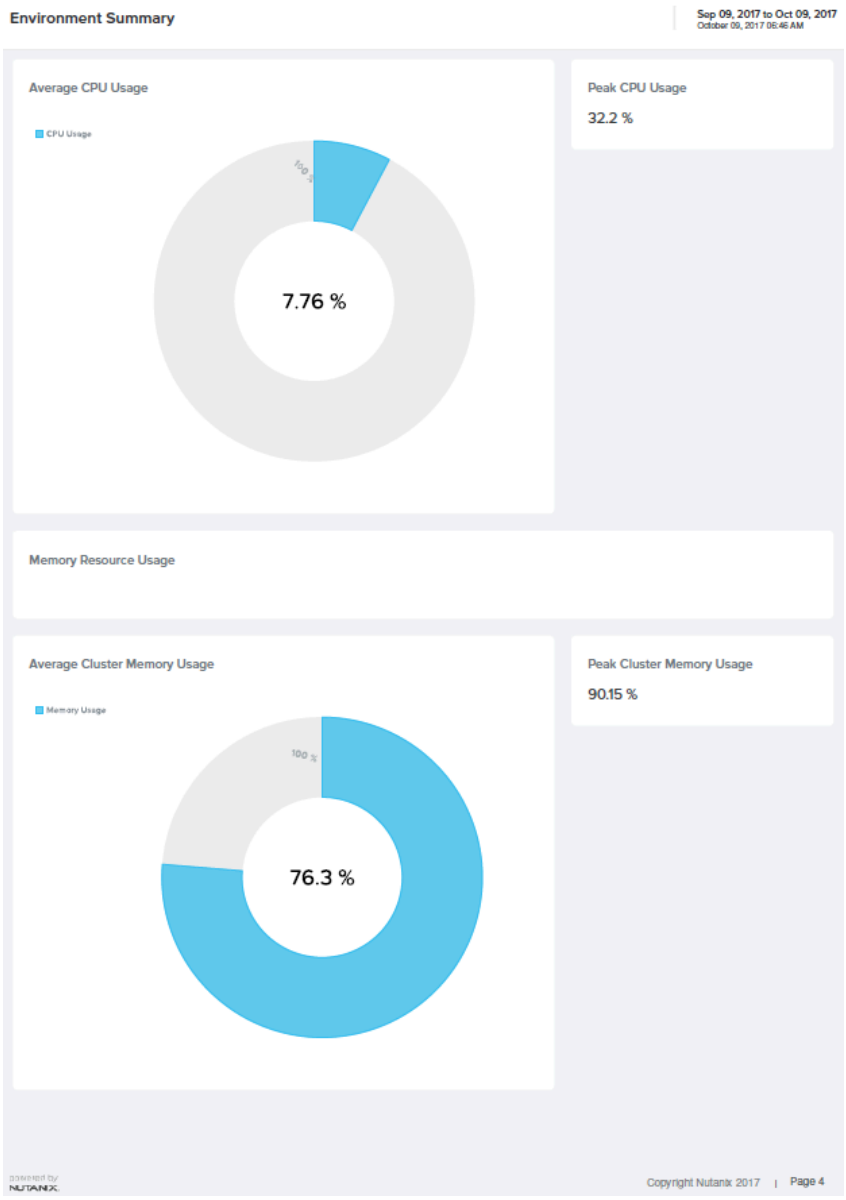

Figure 240: Detailed Metrics Information (Image 2)

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 394.
- To add a data table view, see Adding a Data Table View to a Report on page 397.
- To add a configuration summary view, see Adding Configuration Summary View on page 399.
- To add a metric summary view, see Adding Metric Summary View on page 401.
- To add an entity count view, see Adding Entity Count View on page 403.
- To add a title and description view, see Adding Title and Description View on page 405.
- To add a group view, see Adding a Group View on page 405.
- To add a pre-defined view, see Adding a Pre-Defined View on page 408.

> 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.

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.

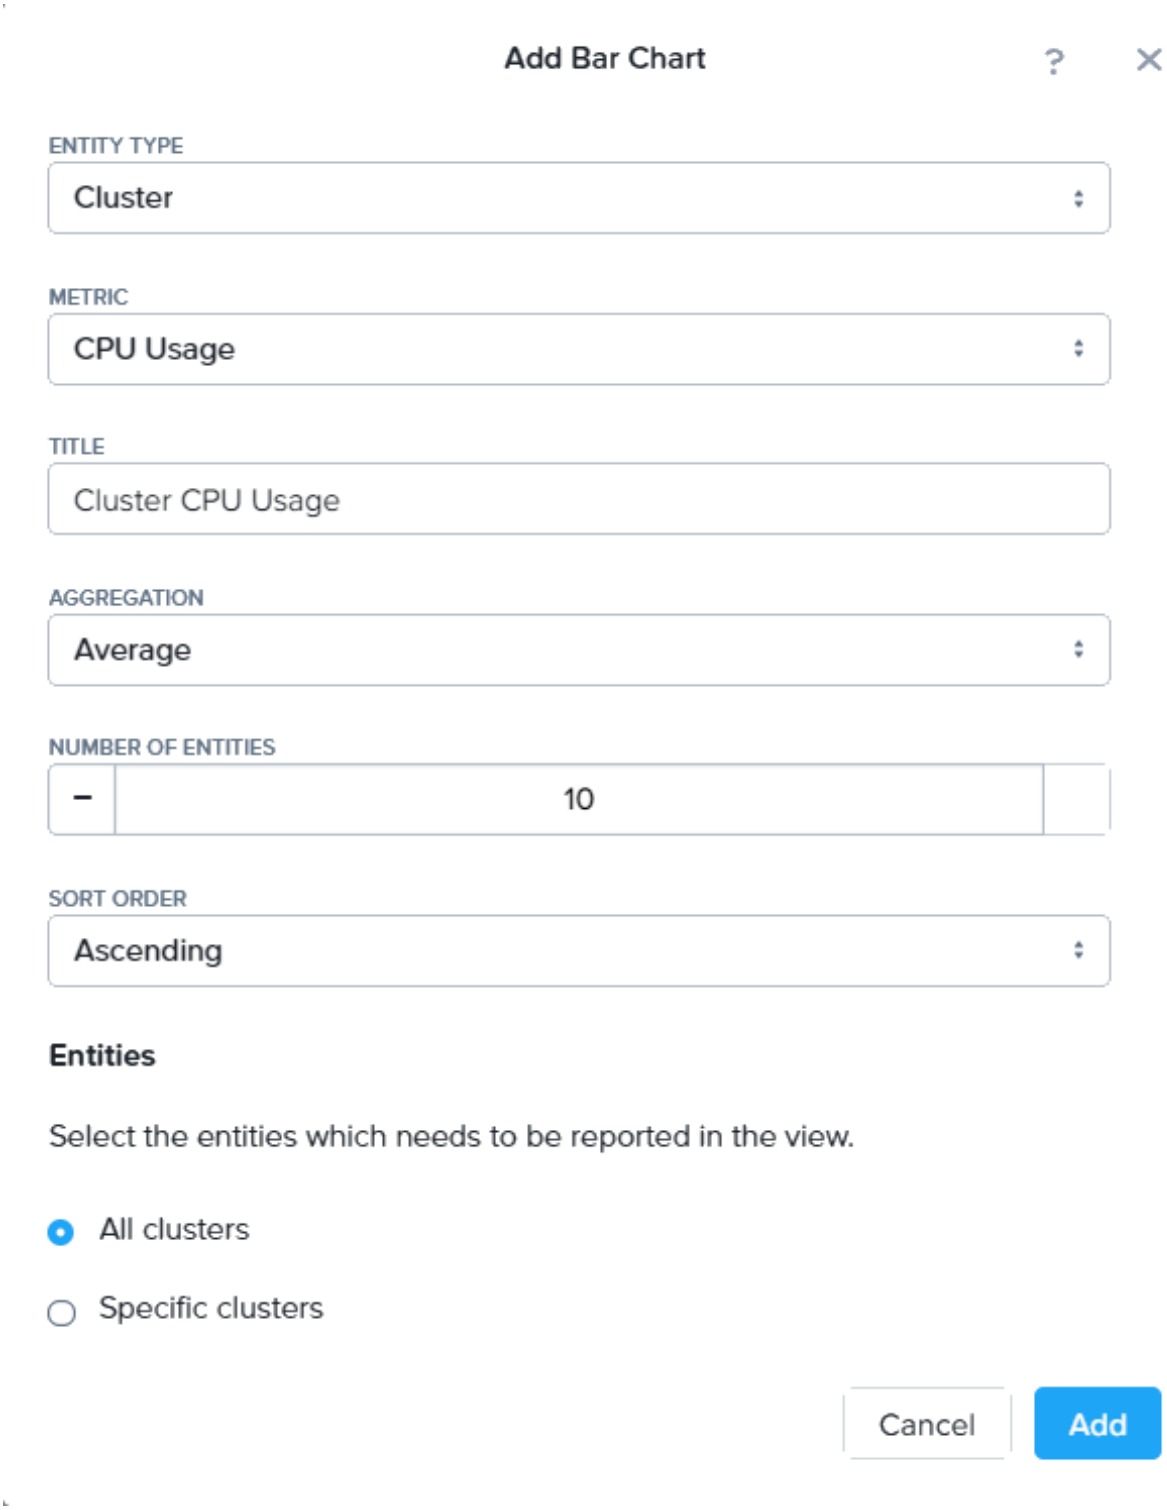

Figure 241: Bar Chart View

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**, **Controller I/O Bandwidth**, **Controller I/O Latency**, **storage.free_bytes**, **storage.capacity_bytes**, **Controller Read I/O Bandwidth**, **Controller Write I/O Bandwidth**, **Controller Read I/O Latency**, **Controller Write I/O Latency**, **storage.logical_usage_bytes**, **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 Buckets** 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.

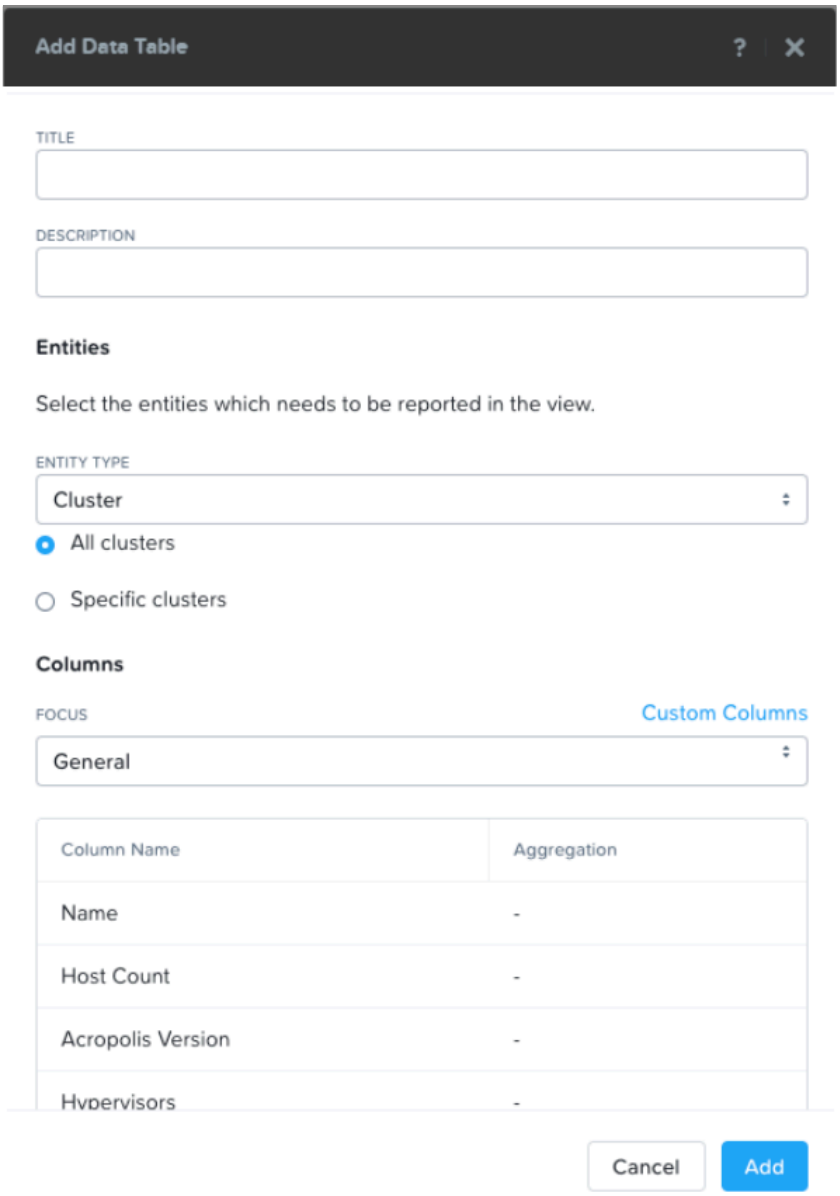

Figure 242: Data Table View

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.

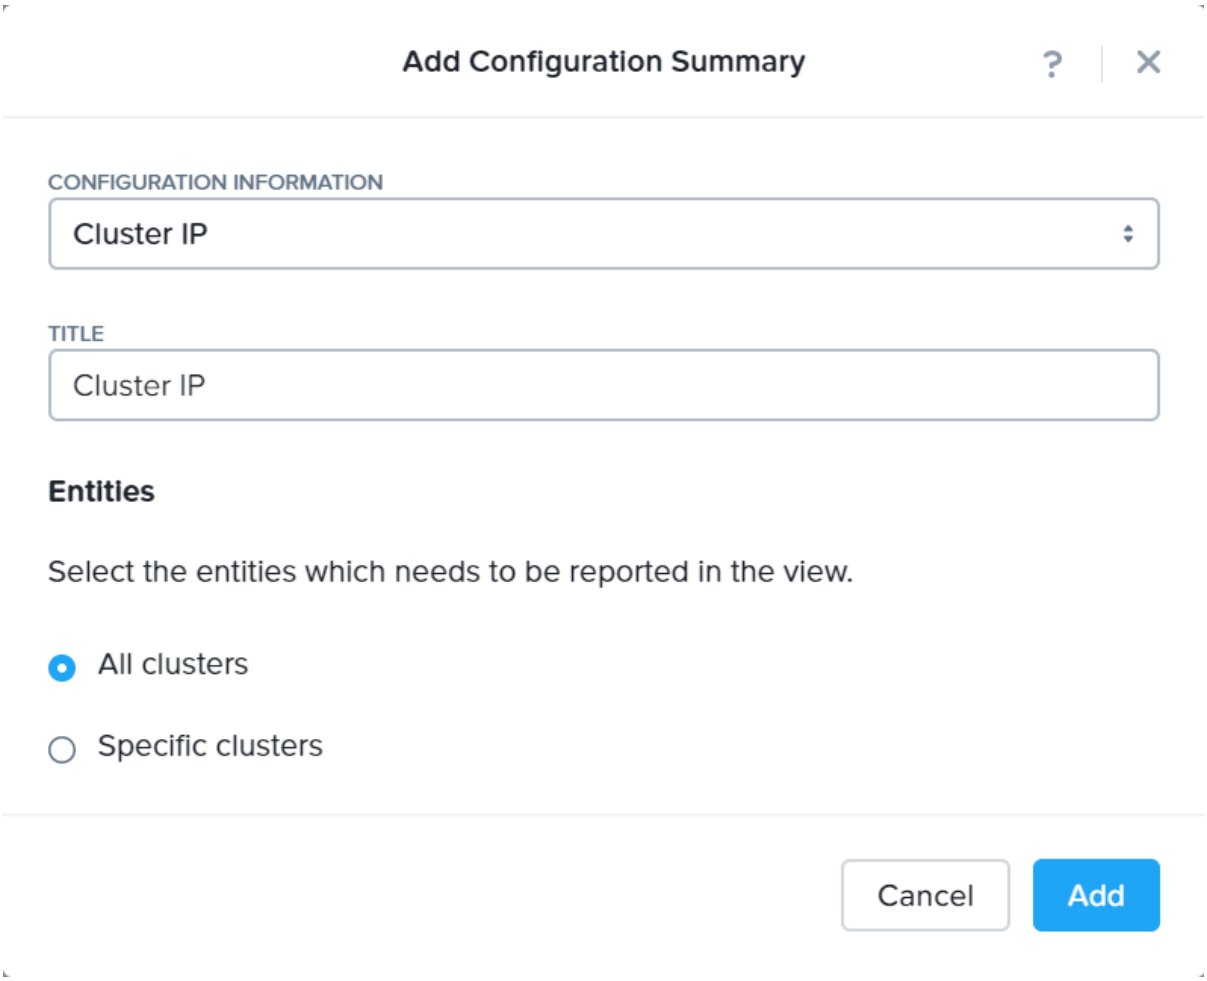

Figure 243: Configuration Summary View

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.

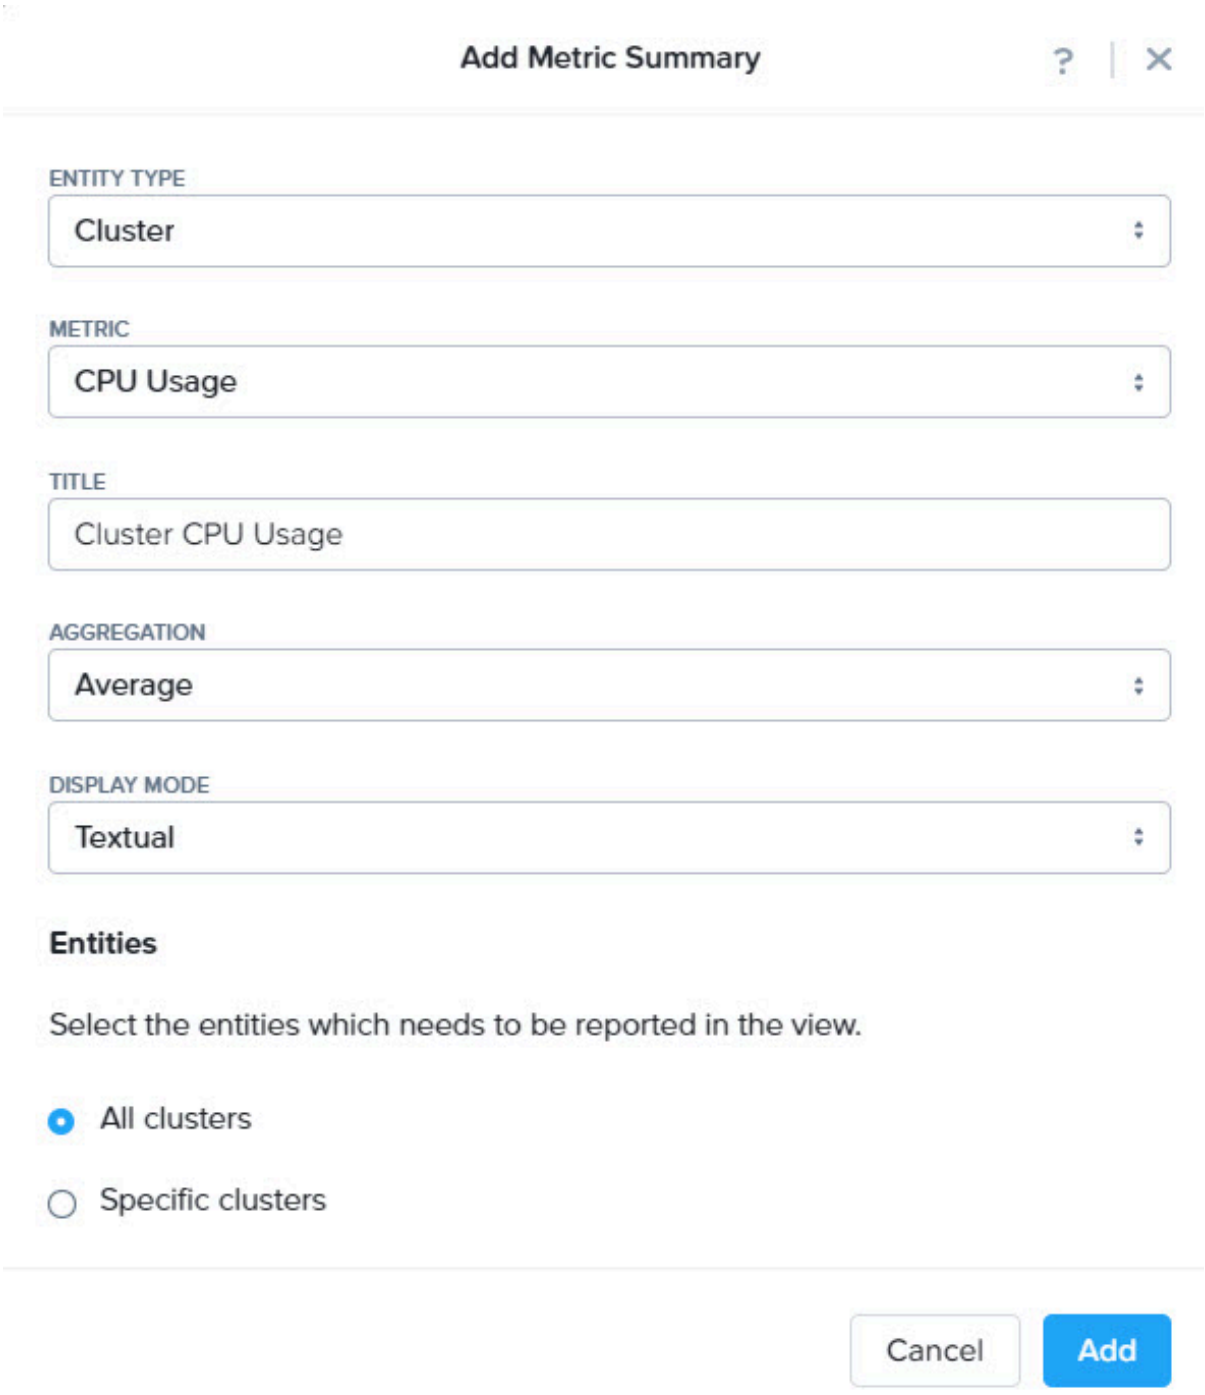

Figure 244: Metric Summary View

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.

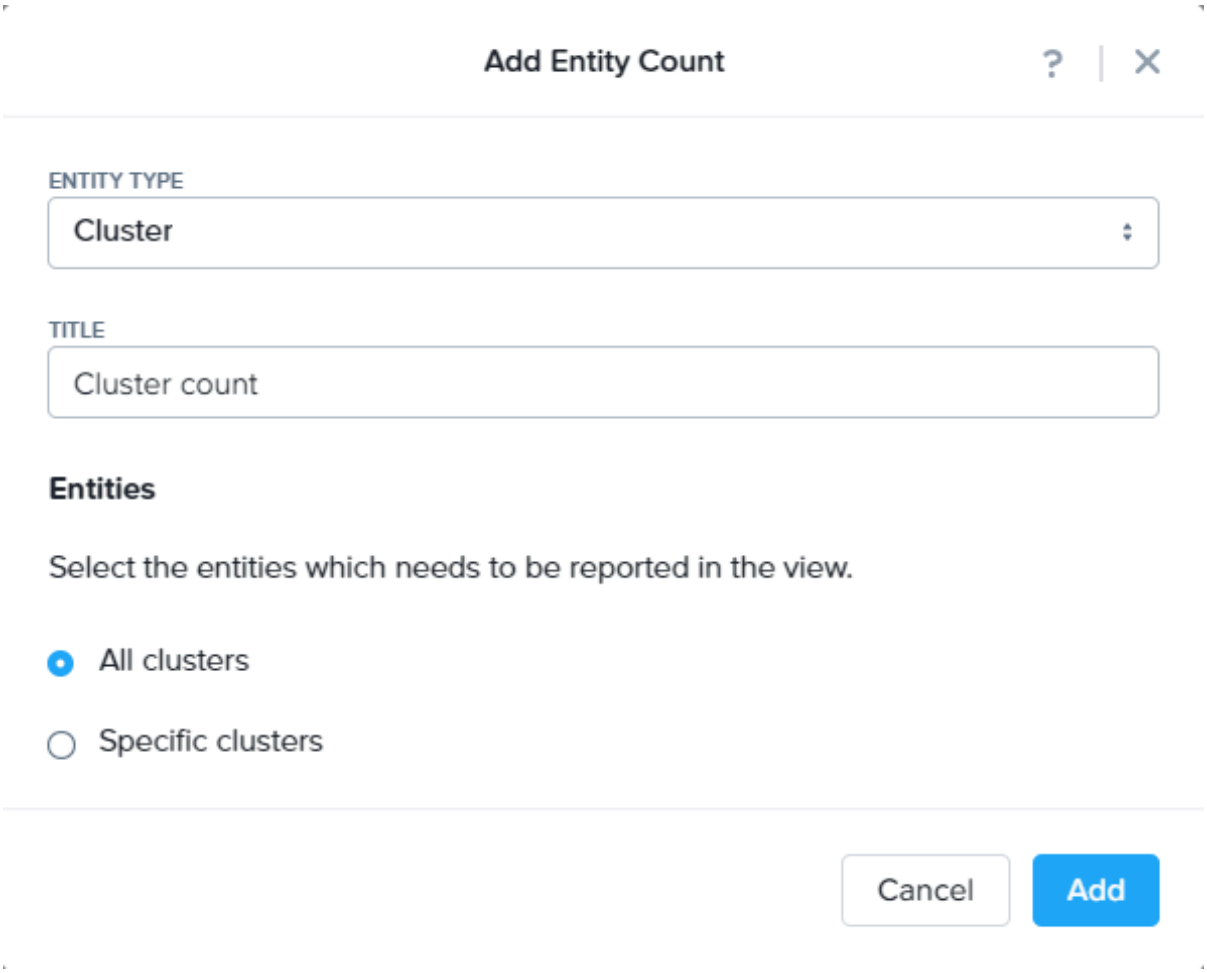

Figure 245: Entity Count View

   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.

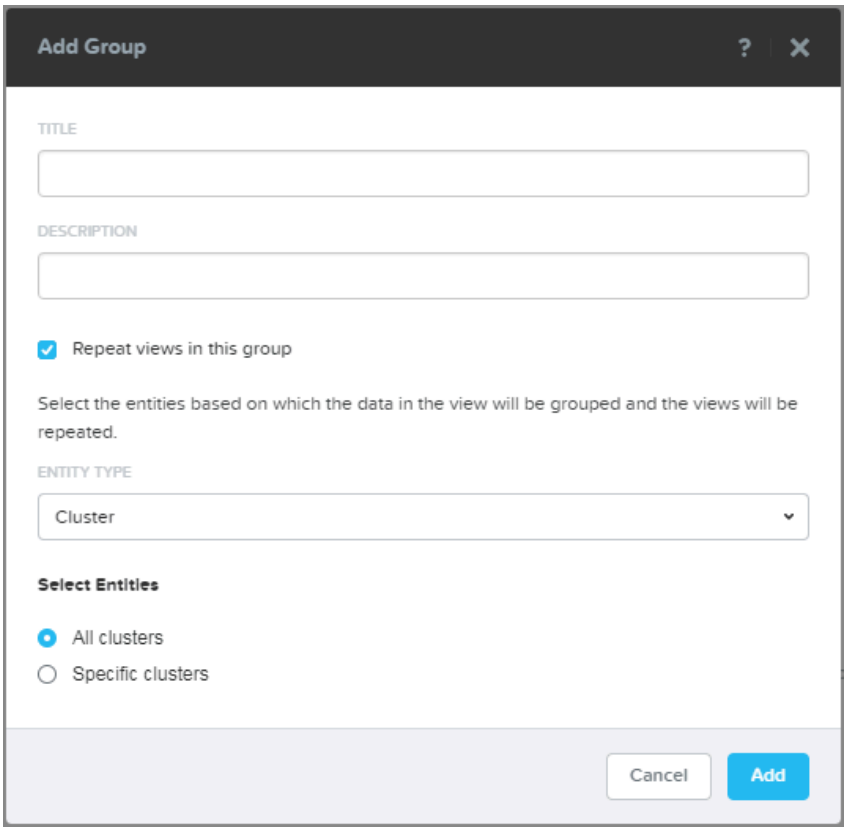

Figure 246: Group View

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 407

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 405

   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**.

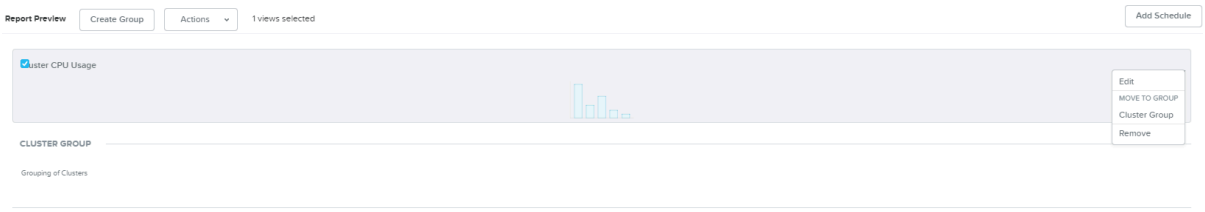

Figure 247: Adding Views to Group view

      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 405.
      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 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 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 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.

- **Potential CPU Gain:** Provides summarized information about the potential CPU resource gain on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Memory Gain:** Provides summarized information about the potential memory resource gain on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Storage Gain:** Provides summarized information about the potential storage resource gain on the cluster from inactive VMs and over-provisioned VMs.

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.

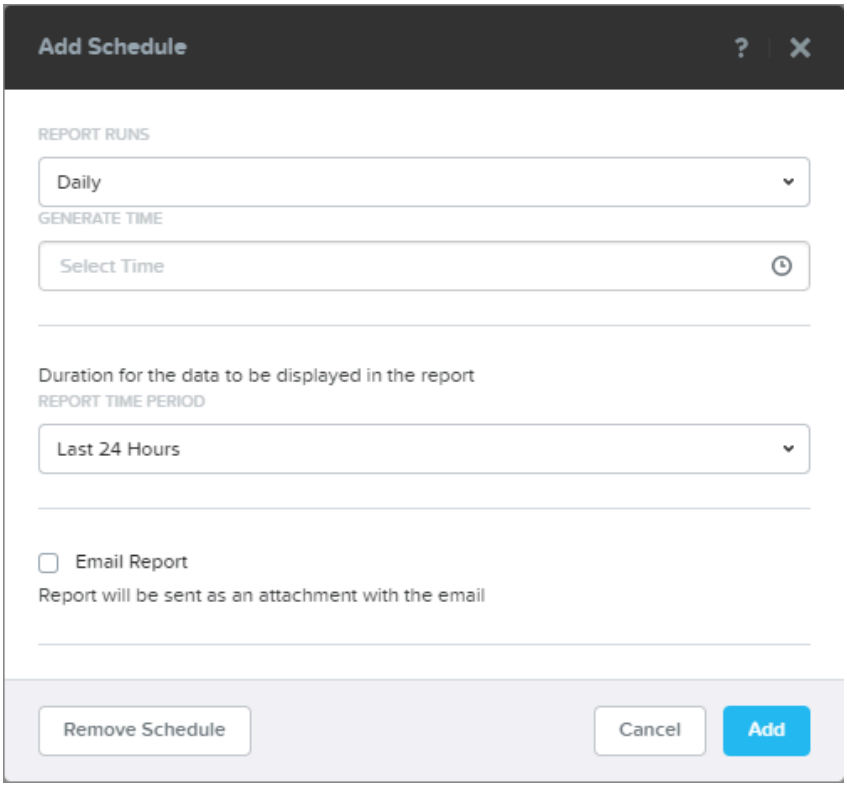

Figure 248: Report Scheduling

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 159), 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 386.

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 159).
   - » 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, background color, and inclusion of copyright

information is included in the cover page settings. Configuration of color of the header is included in the content page settings.

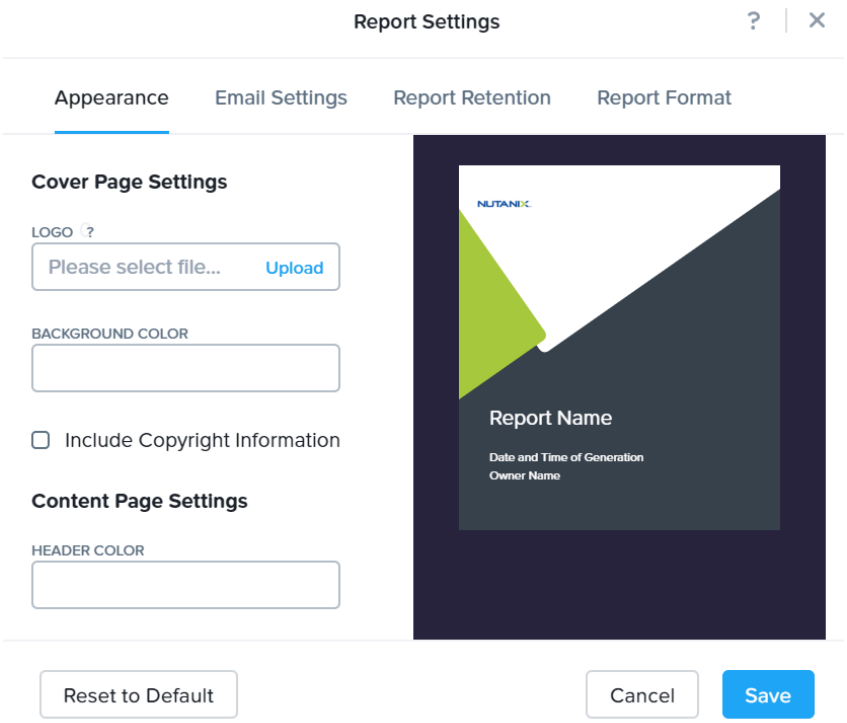

Figure 249: Report Settings Configuration

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 include the copyright information, click the **Include Copyright Information** check box and update the text required in the **Copyright Text** text box.

d. 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 183: 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	#	#	#	#		#	#	#
View Report Instance	#	#	#	#	#	#	#	#

15

CODELESS TASK AUTOMATION (X-PLAY)

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 on page 423 for details.

You can view the instances of playbook execution using the Plays on page 416 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 416 table for details on the available system actions.

To view the action gallery, go to **Operations** > **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

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.

> 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 421 for details.

Table 184: Playbook Actions

Action	Action Description	Configurable attributes
VM Actions		

Action	Action Description	Configurable attributes
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**.
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 421 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**.

Action	Action Description	Configurable attributes
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.
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.
Resolve Alert	Mark an Alert as resolved.	Select the alert from the **Target Alert** drop-down menu.
Send to ServiceNow	Send an Alert to ServiceNow [em_event] table.	Enter the **ServiceNow Instance Name**, **Username**, and **Password** for the ServiceNow instance.
Report Actions		

Action	Action Description	Configurable attributes
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.
Wait Actions		
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 185: 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.

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.

Creating Playbooks

You can create playbooks based on the following triggers.

- Alert Trigger - The **Alert** trigger option 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.
- Manual Trigger - The **Manual** trigger option 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.

The configured playbooks are available in the **Entities** menu (see Entities Menu on page 12).

Creating Playbooks Using Alerts as Triggers

About this task

To create a playbook using predefined alerts, 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 **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. 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

   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.

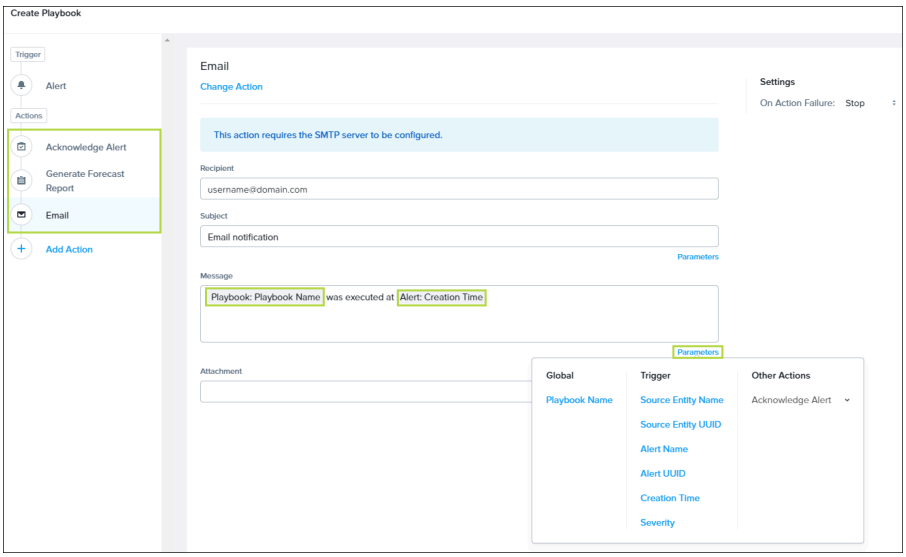

Figure 250: Parameters

To configure an action, set the action attributes based on your requirement. Refer to the Playbook Actions on page 416 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.

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**.

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**.

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 428 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.

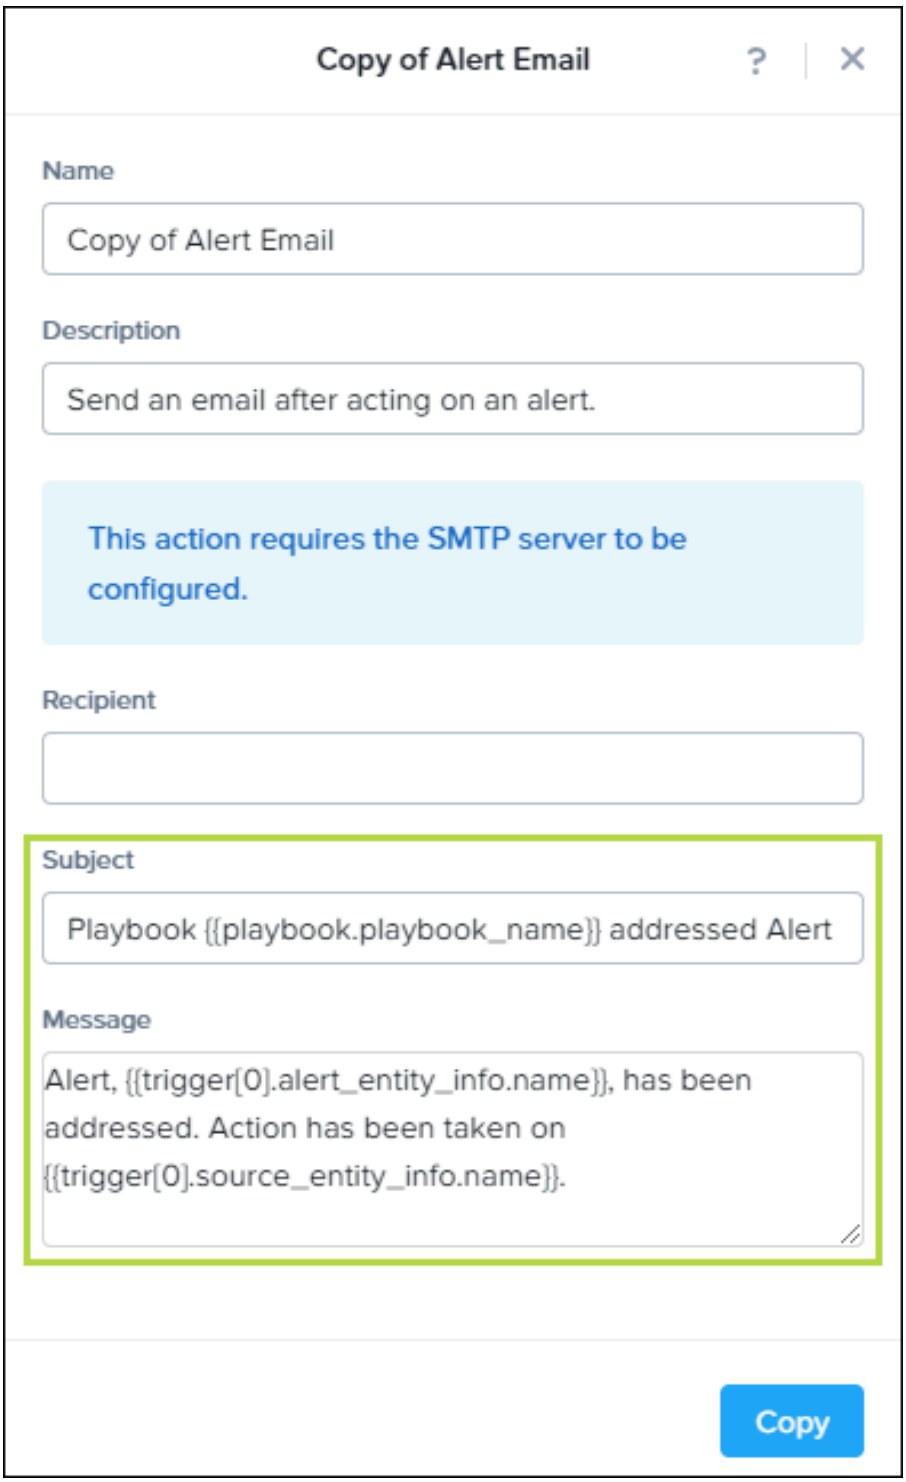

Figure 251: Configuring parameters manually during Clone operation

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 428 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 186: 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}}

Entity	Parameter	
	Report Data (bytes)	{{action[index].report.fileData}}
IP Address Powershell	Return Code (long)	{{action[index].ret_code}}
	Standard Output (string)	{{action[index].stdout}}
	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}}

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 select **Send message as <Application Name>**.
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 must be activated by ServiceNow personnel. 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 Playbooks for ServiceNow Alert Management

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** within the **Alert** option.
3. To select an alert policy, start typing the name of alert policy in the **Alert Policy** field.
4. Click **Add Action** and select **Send to ServiceNow** from the actions deck.
5. 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** as the **Trigger**.

   b. Enter the **ServiceNow Instance Name**.

      Note: Do not enter the complete URL of your ServiceNow instance in this field. If your instance URL is *https://foo.servicenow.com*, the **ServiceNow Instance Name** you must enter is *foo*.

   c. Enter the **Username** and **Password** for the ServiceNow instance.

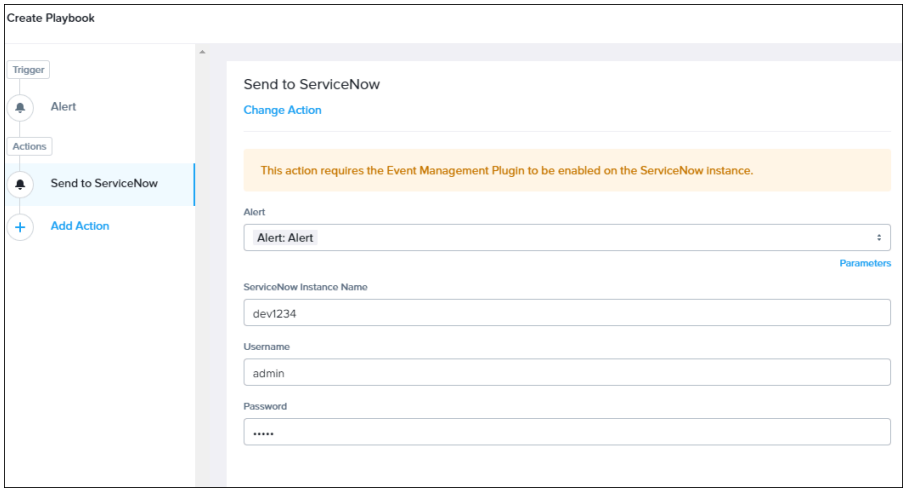

Figure 252: ServiceNow Action

6. 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**.

   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.

16

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 433 and Configuring Pulse on page 435).
- 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 445).
- If you need help, you can create a support ticket directly from Prism Central (see Creating a Support Case on page 439).

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 282).

> 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 438), 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 277). 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 433 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 282).
- 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. In the Settings menu available from the gear icon (see Main Menu on page 9), select **Pulse**.
   The Pulse dialog box appears.

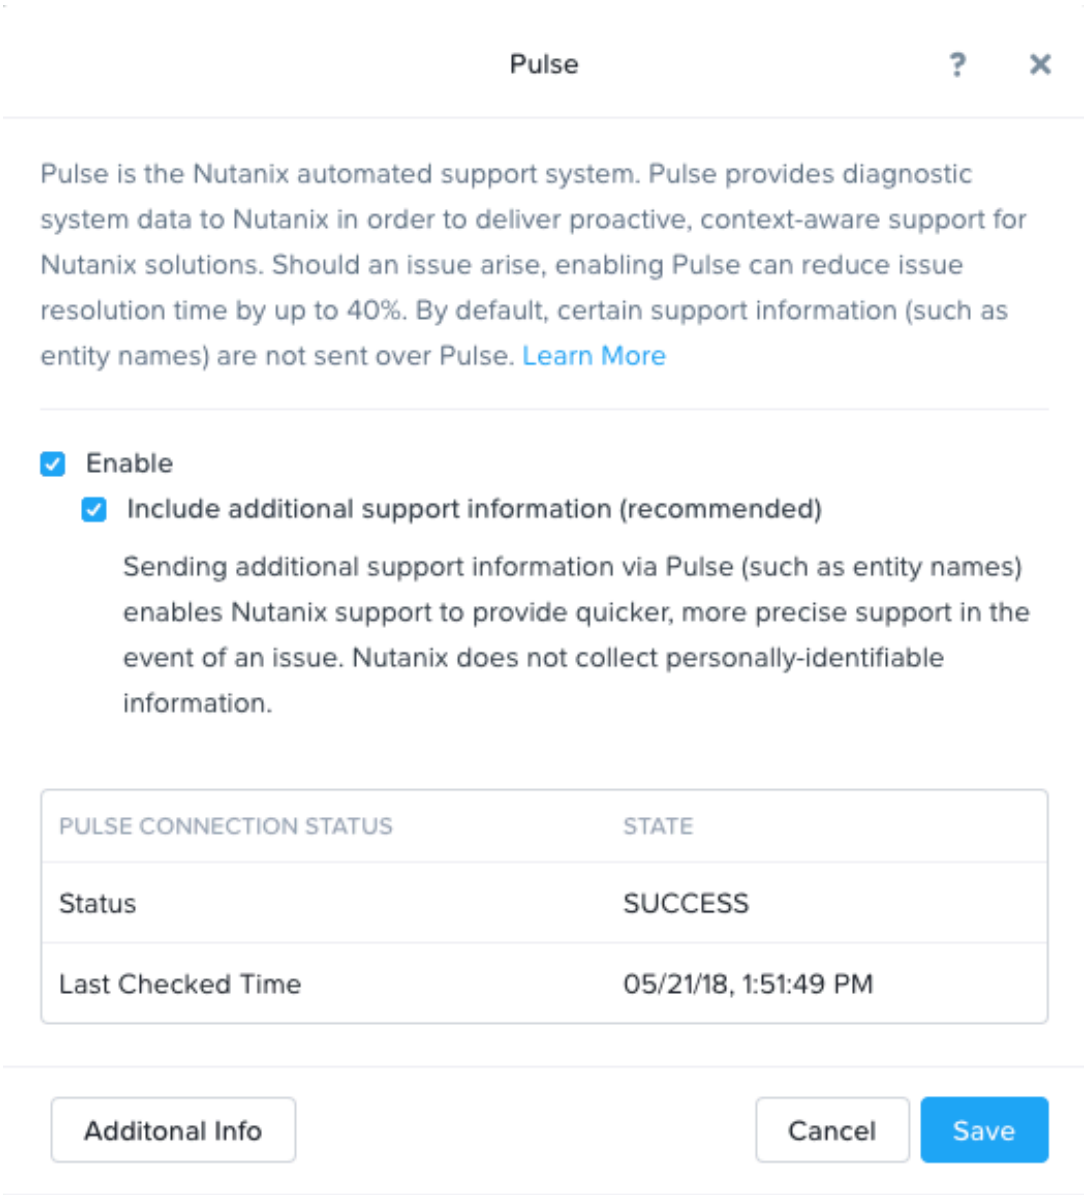

Figure 253: Pulse Window

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.

   Note: Nutanix does not collect any personally-identifiable information (PII) through Pulse.

3. Click the **Additional info** button.
   Additional connection and recipient information appears.

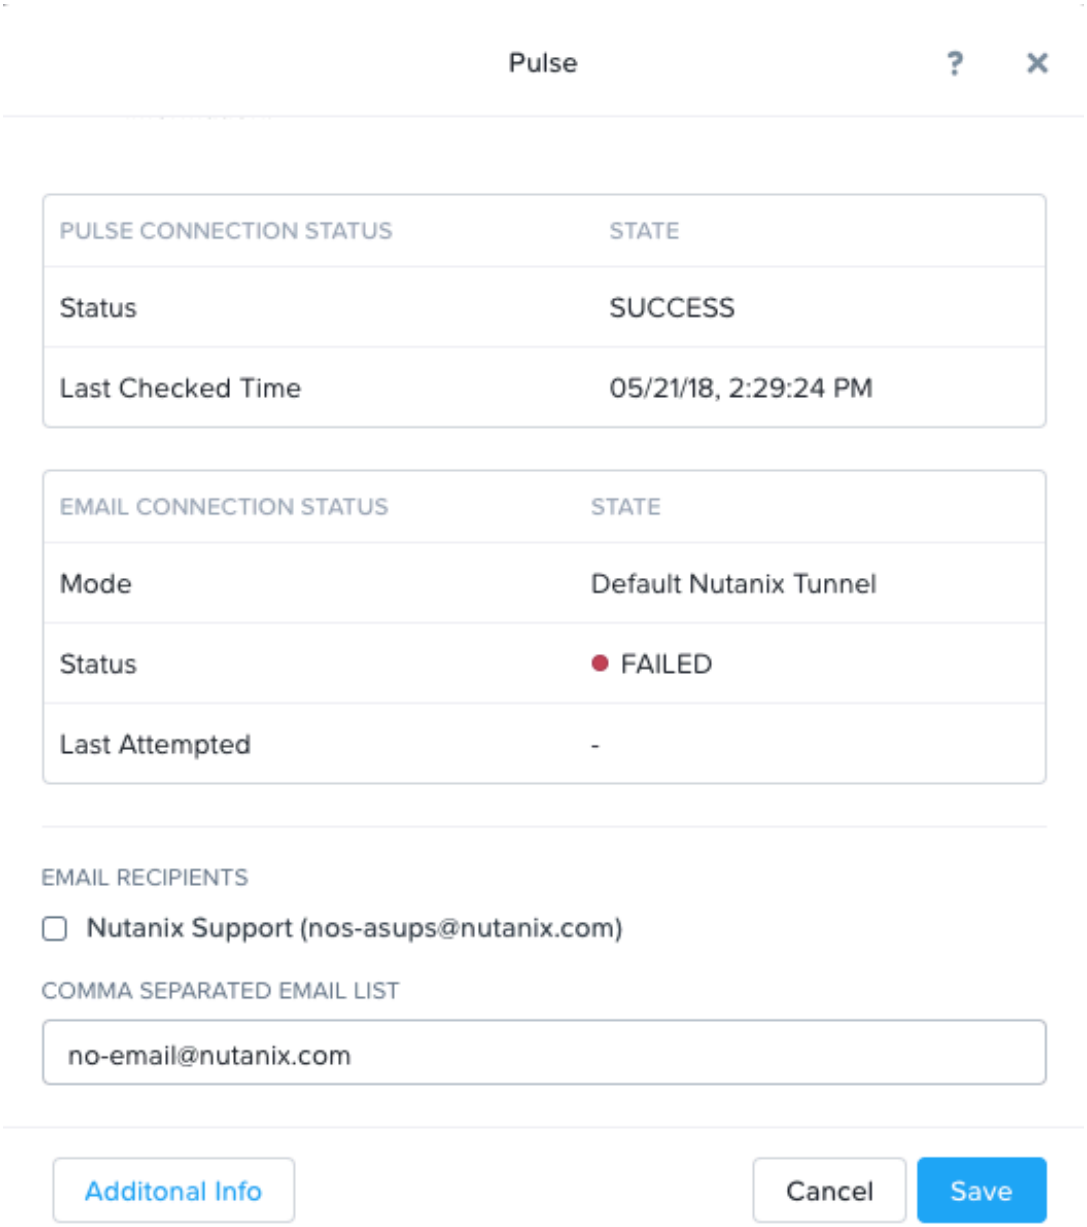

Figure 254: Pulse Window: Additional Info

4. Check the **Email Connection Status** field for connection information.

   This field appears only when the feature is enabled, and it provides the following information:

   - **Mode**: Displays the mode of e-mail transport, which is either Default Nutanix Tunnel or SMTP Server.
   - **Status**: Displays the transport mechanism status, which is **Success** when the default tunnel or an SMTP server is active or **Failure** when neither is available (or **Unknown** when the status is not known).
   - **Last Attempted**: Time of the last summary e-mail attempt.

5. Check the **Pulse Connection Status** field for connection information.

   This field appears only when the feature is enabled, and it provides the following 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.

6. 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).

7. 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.

8. 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 277 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 445). 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 27). 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 444). 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.

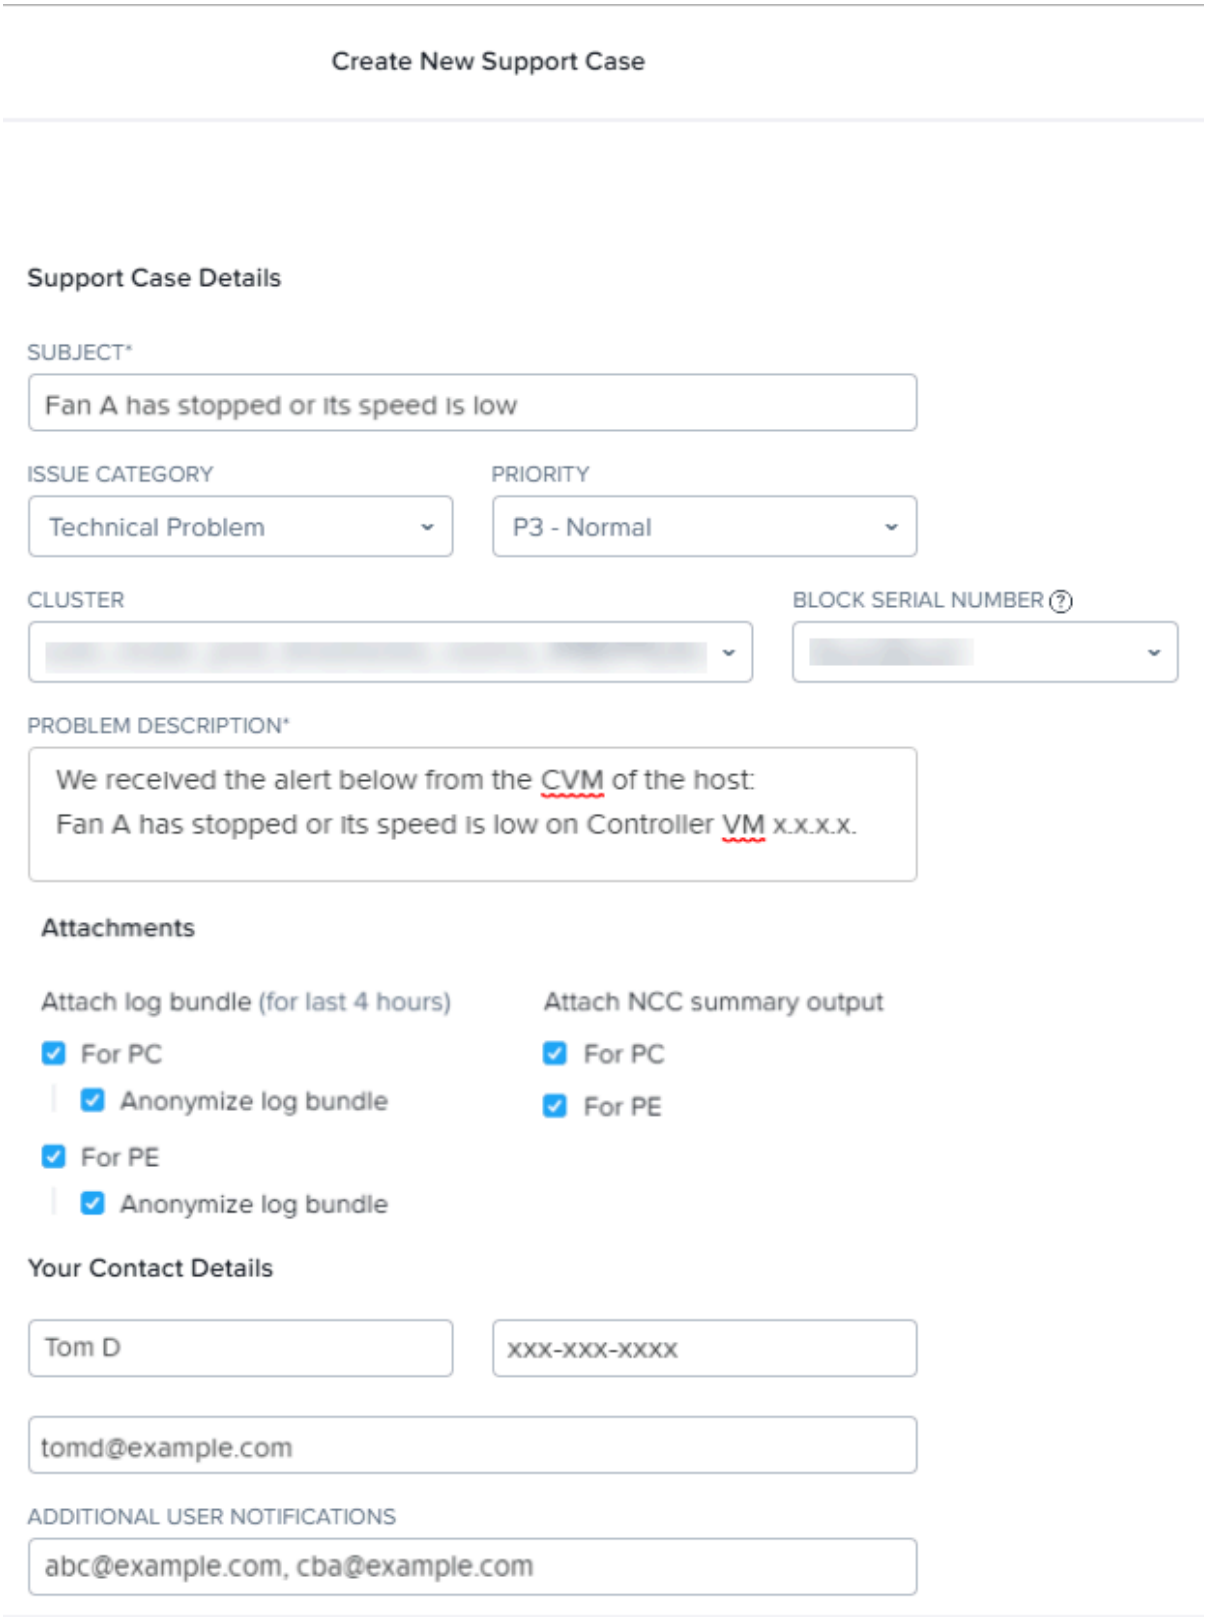

Figure 255: Create New Support Case Page

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 157). 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 444).

Operation Message	Entity	Cluster	Percent	Status	Create Time	Duration
Log collector	Cluster	-	100%	Failed	05/22/19, 11:16:43 AM	29 seconds
Health check	Cluster	-	100%	Succeeded	05/22/19, 11:16:43 AM	44 seconds
Case #00562960 has been updated	Prism Central	-	100%	Succeeded	05/22/19, 11:16:42 AM	49 seconds
Log collector	Cluster	auto_cluster_prod...	100%	Failed	05/22/19, 11:16:42 AM	29 seconds
Health check	Cluster	auto_cluster_prod...	100%	Succeeded	05/22/19, 11:16:42 AM	1 minute 42 ...
Case #00562960 has been updated	Cluster	auto_cluster_prod...	100%	Succeeded	05/22/19, 11:16:41 AM	1 minute 49 ...
Create support case upload	Support case upload	-	100%	Succeeded	05/22/19, 11:16:35 AM	57 seconds

Figure 256: Tasks Dashboard for Prism Central (create case tasks)

19 Total Tasks 1 - 19 of 19

Operation Message	Entity	Percent	Status	Create Time	Duration
Log collector	Cluster	100%	Failed	05/22/19, 11:16:42 AM	29 seconds
Health check	Cluster	100%	Succeeded	05/22/19, 11:16:42 AM	1 minute 42 secon...
Case #00562960 has been updated	Cluster	100%	Succeeded	05/22/19, 11:16:41 AM	1 minute 49 secon...
Create support case upload	Support case uplo...	100%	Succeeded	05/22/19, 11:16:35 AM	1 minute 56 secon...

Figure 257: 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.

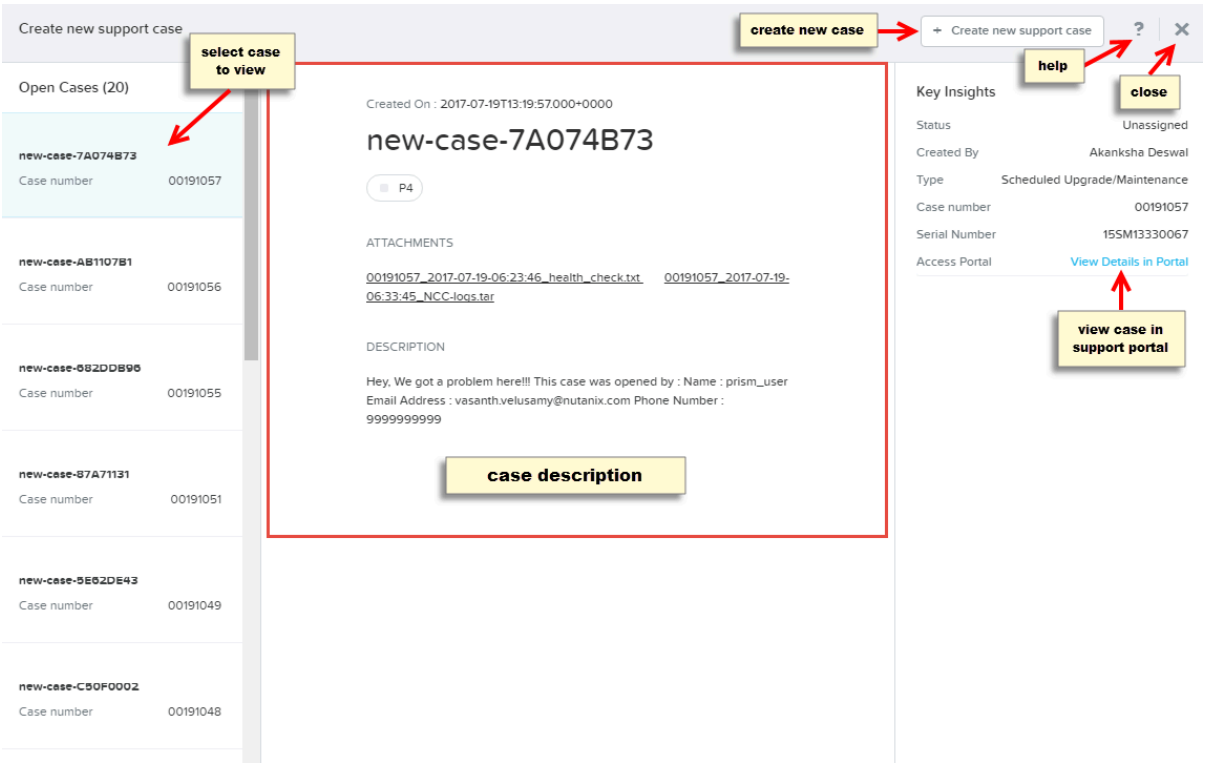

Figure 258: Create New Support Case Page (open cases)

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 187: 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.
	Tools & Firmware	Displays a table of tools that can be downloaded including the Nutanix Cluster Check (NCC).
	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.

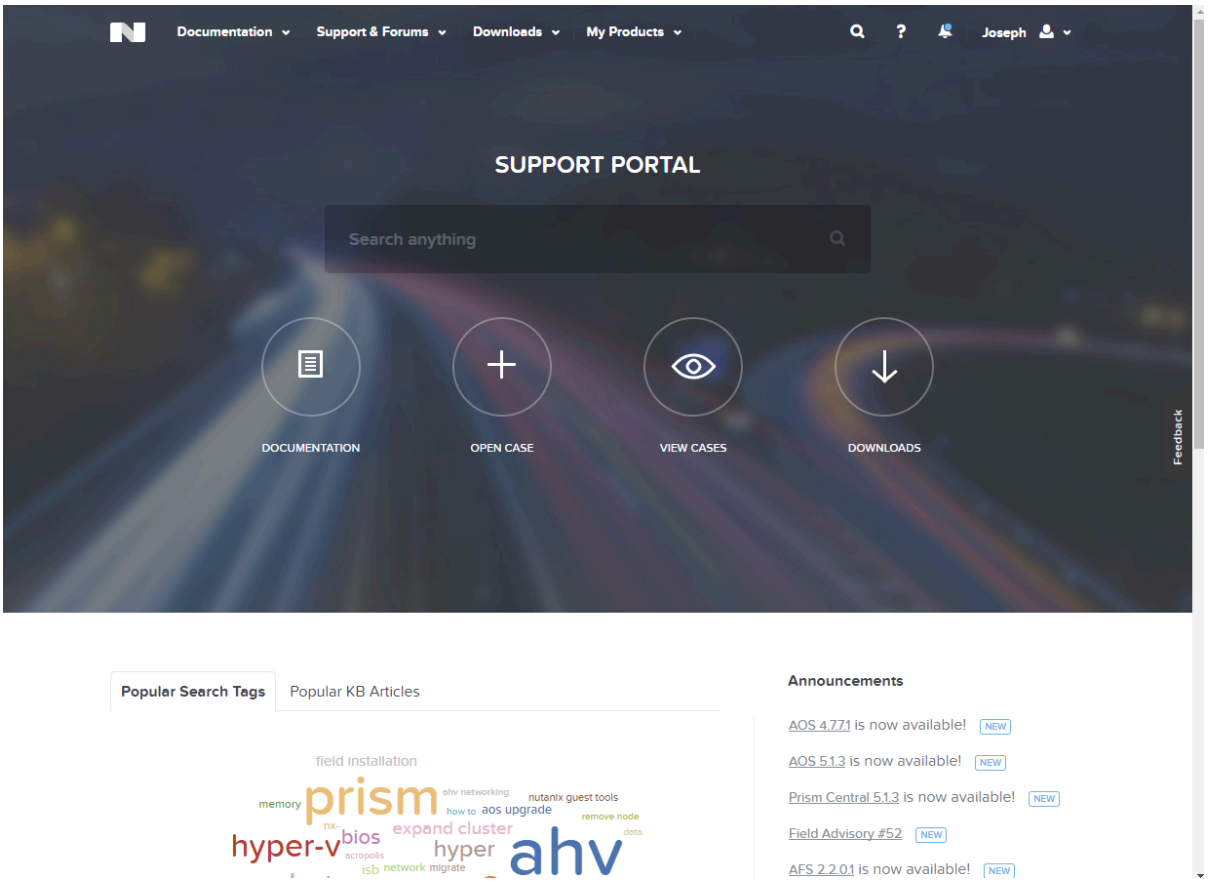

Figure 259: Nutanix Support Portal

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.

17

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 450).
- Nutanix customer support portal (see Accessing the Nutanix Support Portal on page 445).
- Nutanix community forum (see Accessing the Nutanix Next Community on page 452).
- Glossary of terms (see Glossary on page 453).

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.

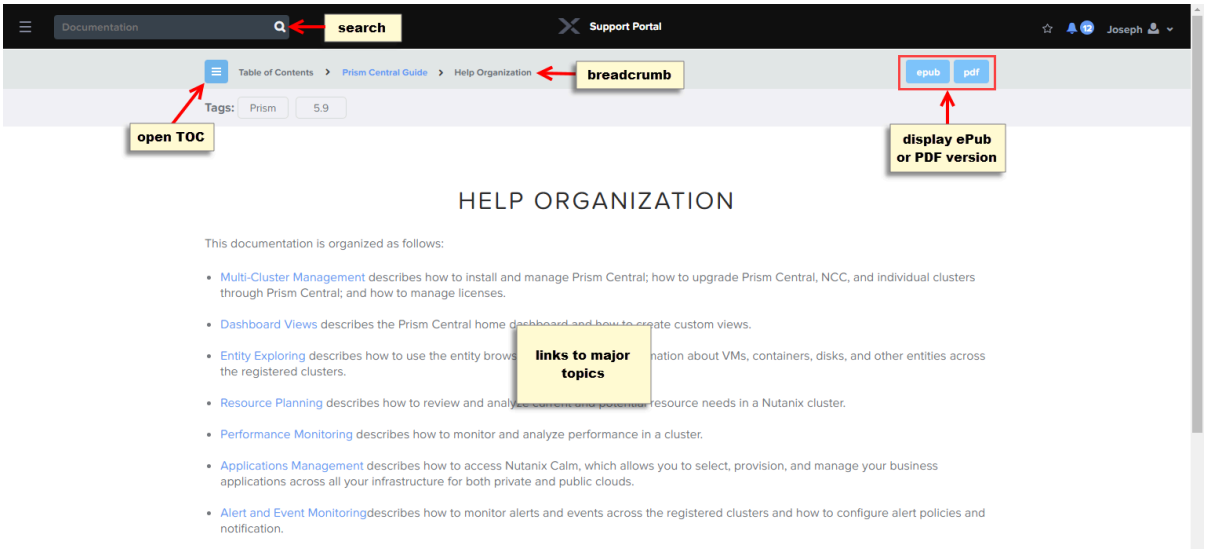

Figure 260: Help Organization Page

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.

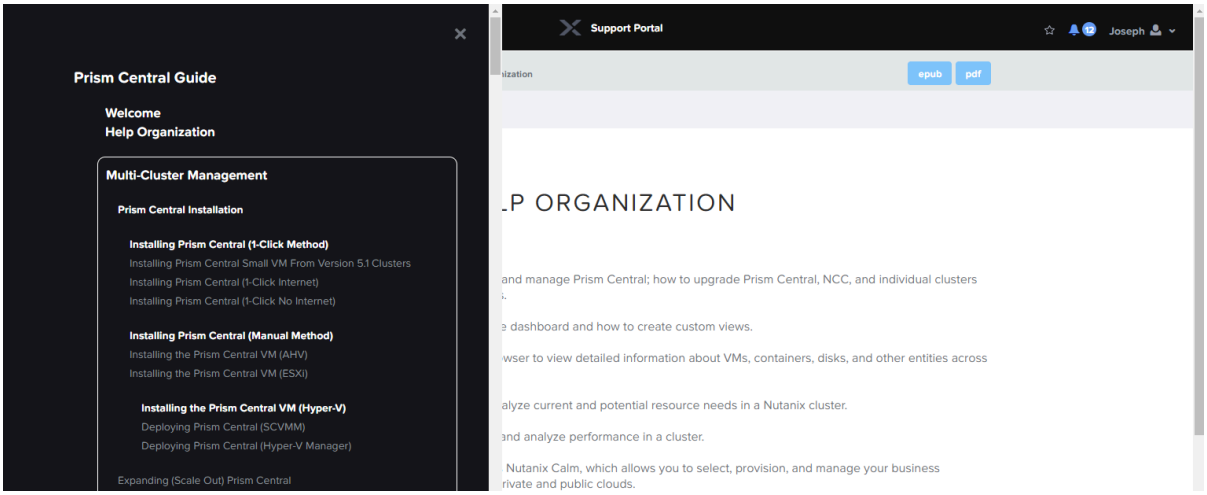

Figure 261: Table of Contents Pane

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.

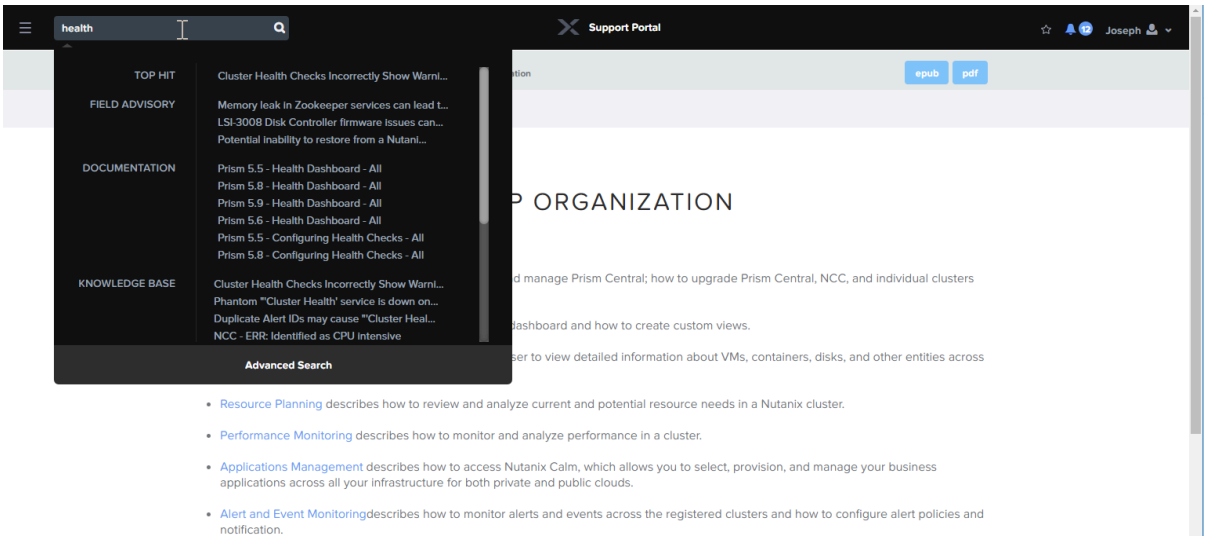

Figure 262: Search Results Example

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 9.

  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.

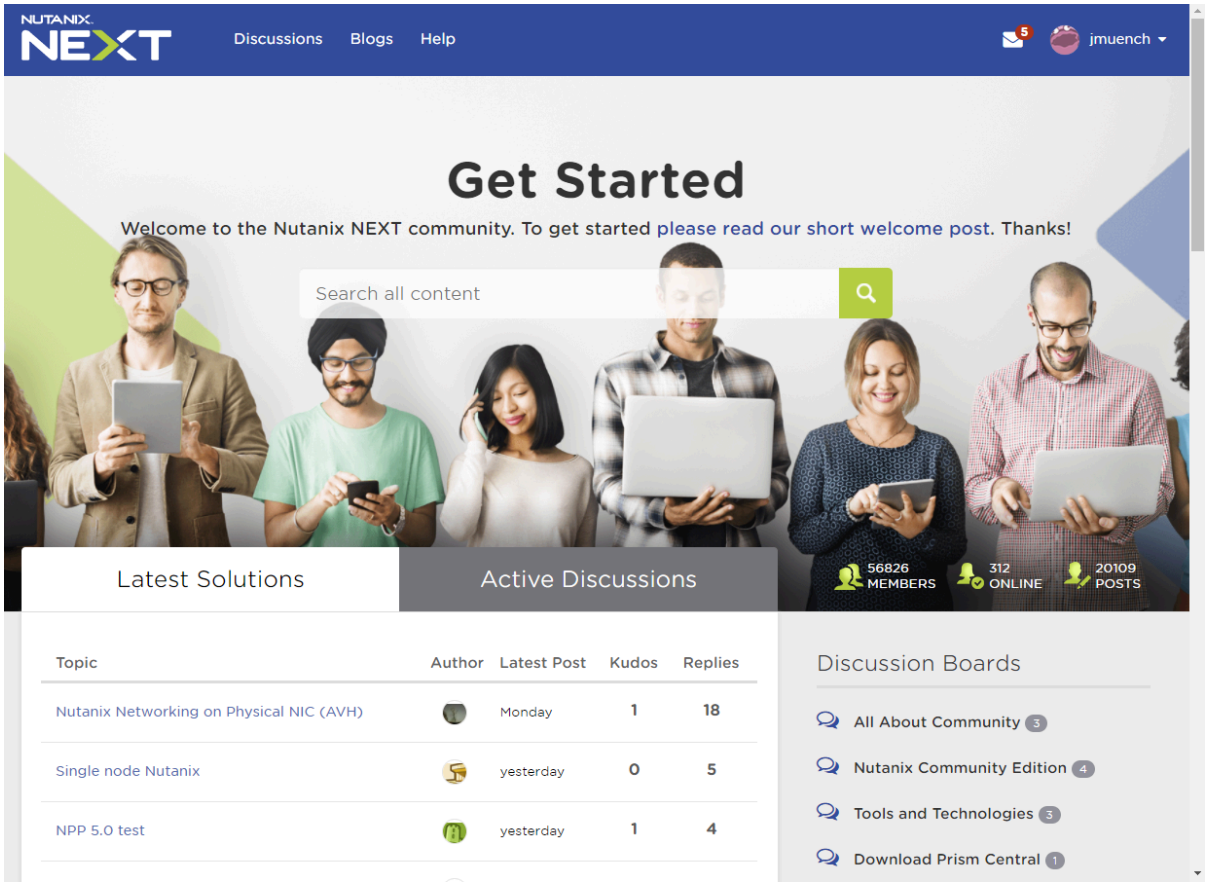

Figure 263: Next Community Screen

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) or Microsoft Azure 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. Because every workload has a different I/O profile, Nutanix has no recommended delay value. 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 or Microsoft Azure, 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 2020 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: April 7, 2020 (2020-04-07T16:38:21+05:30)

Prism 5.15

Prism Web Console Guide

May 15, 2020

Contents

Help Organization viii

1. Nutanix Platform Overview 9
- Guest VM Data Management 9
  - MapReduce Tiering 10
  - Live Migration 10
  - High Availability 10
  - Virtualization Management VM High Availability 11
  - Data Path Redundancy 11
- Cluster Components 12
  - Zeus 12
  - Zookeeper 12
  - Medusa 13
  - Cassandra 13
  - Stargate 13
  - Curator 13
  - Prism 13
- Availability Domains 14
  - Node Failure 14
  - Drive Failures 15
  - Network Link Failure 17
  - Rack Fault Tolerance 17
  - Block Fault Tolerance 22
  - Configuring Node Fault Tolerance 27
  - Redundancy Factor 3 28
  - Degraded Node 30
- Maximum System Values 35
- Three Node Cluster Considerations 36

2. Cluster Management 37
- Web Console Overview 37
  - Logging Into the Web Console 38
  - Logging Out of the Web Console 46
  - Main Menu 46
  - Settings Menu 50
  - Home Dashboard 53
  - Understanding Displayed Statistics 56
  - Modifying Cluster Details 58
  - Modifying UI Settings 60
  - Checking Version 61
- Licensing 61
- Software and Firmware Upgrades 61
- Life Cycle Management (LCM) 62
- Multi-Cluster Management 62
  - Register (Unregister) Cluster with Prism Central 62
  - Installing Prism Central (1-Click Internet) 69
- Single-Node Clusters 78

Single-Node Cluster Guidelines....78
Overriding Read-Only Mode....79
Two-Node Clusters....81
Two-Node Cluster Guidelines....82
Configuring a Witness (Two-Node Cluster)....84
Witness VM Dashboard (two-node cluster)....85
Failure and Recovery Scenarios....87
Increasing the Cluster Fault Tolerance Level....88
Increasing the Controller VM Memory Size....93
Restarting the Nodes in the Cluster (AHV and XenServer)....96

3. Storage Management....97
Storage Components....97
Compression....99
Deduplication....100
Erasure Coding....101
Capacity Reservation Best Practices....103
Storage Dashboard....103
Storage Overview View....105
Storage Diagram View....107
Storage Table View....114
Creating a Storage Pool....125
Modifying a Storage Pool....126
Creating a Storage Container....126
Modifying a Storage Container....131
Volume Group Configuration....132
Concurrent Access from Multiple Clients....133
Creating a Volume Group....133
Modifying a Volume Group....135
Flash Mode for Virtual Machines and Volume Groups....135
Removing Flash Mode for Virtual Disks of a Volume Group....137

4. Network Management....138
Network Configuration for VM Interfaces....138
Configuring a Virtual Network for Guest VM Interfaces....138
Network Segmentation....144
Configuring Network Switch Information....144
Creating SNMP Profiles....148
Modifying Switch Information....150
Network Visualization....151
Prerequisites....152
Network Visualizer....152
Viewing the Network Visualizer....153
Customizing the Topology View....153
Viewing VM Information....154
Viewing Host Information....155
Viewing Switch Information....156
Uplink Configuration (AHV Only)....156

5. Hardware Management....161
Hardware Dashboard....161
Hardware Overview View....163
Hardware Diagram View....165

Hardware Table View ... 173
Expanding a Cluster ... 185
Cluster Expansion Guidelines ... 190
Starting Host Boot Disk Repair (Proactive or Graceful Replacement) ... 194
Modifying a Cluster ... 198
Compute-Only Node Configuration (AHV Only) ... 201
Adding a Compute-Only Node to an AHV Cluster ... 203
Storage-only Node Configuration ... 204
Imaging and Adding Storage-only Nodes ... 206

6. Nutanix Volumes ... 207
Example Use Cases Supported by Volumes ... 208
If You are Already Using Volume Groups ... 208
Volumes Requirements and Supported Clients ... 208
Enabling Nutanix Volumes ... 210
Creating a Volume Group for Use with Volumes ... 211
Configuring Windows Clients ... 215
Configuring Linux Clients ... 222
Configuring AIX Clients ... 225
Booting Over iSCSI ... 227
Converting Volume Groups and Updating Clients to Use Volumes ... 231
Modifying Windows Client Settings to Use Volumes ... 232
Modifying Linux Client Settings to Use Volumes ... 232
Modifying a Volume Group for Use with Volumes ... 233
Creating a Linux Guest Cluster ... 234

7. File Server Management ... 237

8. Data Protection ... 238

9. Health Monitoring ... 239
Health Dashboard ... 239
Configuring Health Checks ... 244
Configuring NCC Frequency ... 246
Running Checks by Using Web Console ... 247
Collecting Logs by Using Web Console ... 248

10. Virtual Machine Management ... 252
VM Dashboard ... 252
VM Overview View ... 254
VM Table View ... 255
VM Management ... 268
Creating a VM (AHV) ... 268
Managing a VM (AHV) ... 275
Creating a VM (ESXi) ... 283
Managing a VM (ESXi) ... 285
Configuring Images ... 289
Virtual Machine Customization ... 292
Customizing Linux Virtual Machines with Cloud-Init ... 293
Cloud-Init Limitations and Guidelines ... 293
Customization of Windows Virtual Machines with System Preparation ... 294

VM High Availability in Acropolis....295
Enabling High Availability for the Cluster....296
Nutanix Guest Tools....297
Nutanix Guest Tools Requirements and Limitations....298
Enabling NGT and Mounting the NGT Installer in a VM....299
NGT Installation....301
Enabling NGT and Mounting the NGT Installer Simultaneously on Multiple Cloned VMs....304
NGA and Controller VM Communication....305
CD-ROM Eject Functionality of NGT....306
Nutanix Guest Tools Usage in Disaster Recovery....306
Upgrading NGT....306
Reconfiguring NGT....307
Disabling, Uninstalling, and Removing Nutanix Guest Tools....307
Affinity Policies for AHV....308
Configuring VM-VM Anti-Affinity Policy....309
Connect to Citrix Cloud....310
Connecting to the Citrix Cloud....311

11. Performance Monitoring....317
Analysis Dashboard....317
Creating an Entity Chart....319
Creating a Metric Chart....320
Chart Metrics....321
Exporting Performance Data....326

12. Alert and Event Monitoring....327
Alerts Dashboard....327
Alerts Summary View....327
Events Summary View....333
Configuring Alert Emails....335
Configuring Alert Policies....338
Alerts/Health checks....340
Availability Zone....340
Cluster....340
Controller VM....426
DR....439
Entity Sync....482
Guest VM....482
Hardware....492
License....517
Node....517
Network....539
Nutanix Files....547
Other....563
Storage....574
System Indicator....590
VM....591

13. Task Status....592
View Task Status....593

14. System Management....594
Configuring a Filesystem Whitelist....594
Configuring Name Servers....595
Cluster Time Synchronization....596
Recommendations for Time Synchronization....596
Configuring NTP Servers....597
Configuring an SMTP Server....598
Configuring SNMP....599
Nutanix MIB....608
Configuring a Banner Page....615
VM Management through Prism Element (ESXi)....616
Registering a vCenter Server....617
Managing vCenter Server Registration Changes....619
In-Place Hypervisor Conversion....620
Requirements and Limitations for In-Place Hypervisor Conversion....621
In-Place Hypervisor Conversion Process....624
Converting Cluster (ESXi to AHV)....625
Converting Cluster (AHV to ESXi)....626
Aborting Cluster Conversion....628
Internationalization (i18n)....629
Localization (L10n)....629
Changing the Language Settings....630
Hyper-V Setup....631
Joining the Cluster and Hosts to a Domain....631
Creating a Failover Cluster for Hyper-V....633
Enabling Kerberos for Hyper-V....634

15. Security and User Management....639

16. Support Services....640
Pulse Health Monitoring....640
Configuring Pulse....642
Pulse Access Requirements....644
Controlling Remote Connections....645
Configuring HTTP Proxy....646
Accessing the Nutanix Support Portal....648
Nutanix REST API....651
Accessing the REST API Explorer....652

17. Help Resources....656
Accessing Online Help....656
Accessing the Nutanix Next Community....660
Glossary....661

Copyright....667
License....667
Conventions....667

Version.....667

HELP ORGANIZATION

This documentation is organized as follows:

- Nutanix Platform Overview on page 9 describes the Nutanix architecture.
- Cluster Management on page 37 describes how to access and use the web console, how to apply a Nutanix cluster license, how to upgrade the cluster to a later AOS release, and how to upgrade other software components such as disk firmware.
- Multi-Cluster Management on page 62 describes how to install a VM that runs an application called Prism Central and how to use Prism Central to monitor and manage multiple clusters.
- Storage Management on page 97 describes how to monitor storage use in a Nutanix cluster and how to create storage pools and containers.
- Hardware Management on page 161 describes how to monitor hardware configurations in a Nutanix cluster and how to expand the cluster.
- Health Monitoring on page 239 describes how to monitor the health of VMs, hosts, and disks across a Nutanix cluster.
- Virtual Machine Management on page 252 describes how to monitor status of the VMs across a Nutanix cluster.
- Performance Monitoring on page 317 describes how to monitor and analyze performance in a Nutanix cluster.
- Alert and Event Monitoring on page 327 describes how to monitor activity in a Nutanix cluster and how to configure alert policies and notification.
- System Management on page 594 describes how to configure various system settings such as for SNMP, NTP, and SMTP.
- Security and User Management on page 639 describes how to configure various security settings including authentication method, SSL certificates, and SSH keys. It also describes how to add, edit, and delete user accounts.
- Support Services on page 640 describes how to enable (or disable) Nutanix technical support access to your cluster, how to access the Nutanix support portal, and how to access the Nutanix REST API explorer.

1

NUTANIX PLATFORM OVERVIEW

The Nutanix Enterprise Cloud Platform is a converged, scale-out compute and storage system that is purpose-built to host and store virtual machines. All nodes in a Nutanix cluster converge to deliver a unified pool of tiered storage and present resources to VMs for seamless access. A global data system architecture integrates each new node into the cluster, allowing you to scale the solution to meet the needs of your infrastructure.

The foundational unit for the cluster is a Nutanix node. Each node in the cluster runs a standard hypervisor and contains processors, memory, and local storage (SSDs and hard disks).

A Nutanix Controller VM runs on each node, enabling the pooling of local storage from all nodes in the cluster.

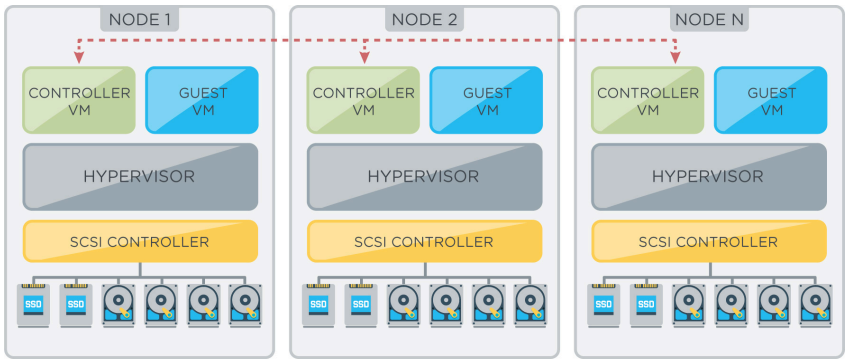

Guest VM Data Management

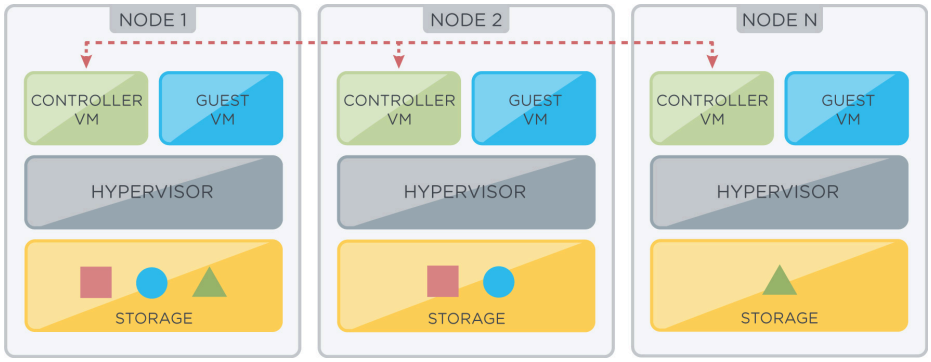

Hosts read and write data in shared Nutanix datastores as if they were connected to a SAN. From the perspective of a hypervisor host, the only difference is the improved performance that results from data not traveling across a network. VM data is stored locally, and replicated on other nodes for protection against hardware failure.

When a guest VM submits a write request through the hypervisor, that request is sent to the Controller VM on the host. To provide a rapid response to the guest VM, this data is first stored on the metadata drive, within a subset of storage called the oplog. This cache is rapidly

distributed across the 10 GbE network to other metadata drives in the cluster. Oplog data is periodically transferred to persistent storage within the cluster. Data is written locally for performance and replicated on multiple nodes for high availability.

When the guest VM sends a read request through the hypervisor, the Controller VM reads from the local copy first, if present. If the host does not contain a local copy, then the Controller VM reads across the network from a host that does contain a copy. As remote data is accessed, the remote data is migrated to storage devices on the current host, so that future read requests can be local.

MapReduce Tiering

The Nutanix cluster dynamically manages data based on how frequently it is accessed. When possible, new data is saved on the SSD tier. Frequently-accessed, or "hot" data is kept on this tier, while "cold" data is migrated to the HDD tier. Data that is accessed frequently is again moved back to the SSD tier.

This automated data migration also applies to read requests across the network. If a guest VM repeatedly accesses a block of data on a remote host, the local controller VM migrates that data to the SSD tier of the local host. This migration not only reduces network latency, but also ensures that frequently-accessed data is stored on the fastest storage tier.

Live Migration

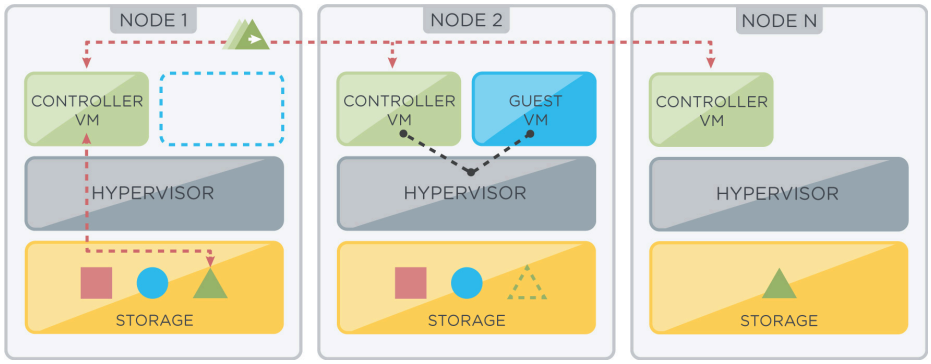

Live migration of VMs, whether it is initiated manually or through an automatic process like vSphere DRS, is fully supported by the Nutanix Enterprise Cloud Computing Platform. All hosts within the cluster have visibility into shared Nutanix datastores through the Controller VMs. Guest VM data is written locally, and is also replicated on other nodes for high availability.

If a VM is migrated to another host, future read requests are sent to a local copy of the data, if it exists. Otherwise, the request is sent across the network to a host that does contain the requested data. As remote data is accessed, the remote data is migrated to storage devices on the current host, so that future read requests can be local.

High Availability

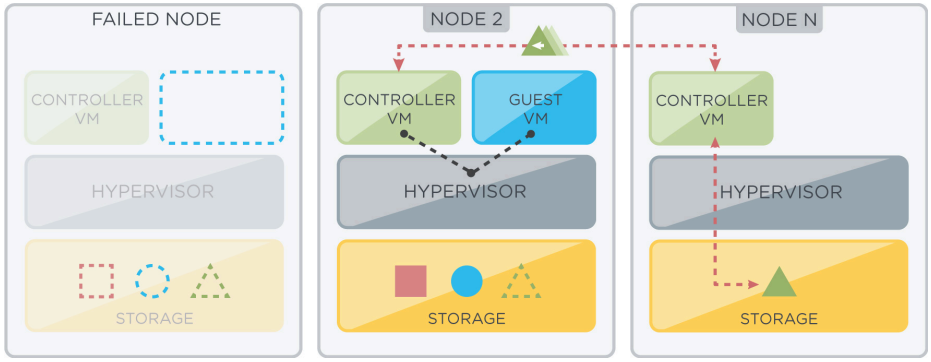

The built-in data redundancy in a Nutanix cluster supports high availability provided by the hypervisor. If a node fails, all HA-protected VMs can be automatically restarted on other nodes in the cluster. The hypervisor management system, such as vCenter, selects a new host for the VMs, which may or may not contain a copy of the VM data.

If the data is stored on a node other than the VM's new host, then read requests are sent across the network. As remote data is accessed, the remote data is migrated to storage devices on the current host, so that future read requests can be local. Write requests is sent to the local storage, and also replicated on a different host. During this interaction, the Nutanix software also creates new copies of pre-existing data, to protect against future node or disk failures.

Virtualization Management 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.

Typically, an entity failure is detected by its isolation from the network (the failure to respond to heartbeats). Virtualization management ensures that at most one instance of the VM is running at any point during a failover. This property prevents concurrent network and storage I/O that could lead to corruption.

Virtualization management VM high availability may implement admission control to ensure that in case of node failure, the rest of the cluster has enough resources to accommodate the VMs.

Data Path Redundancy

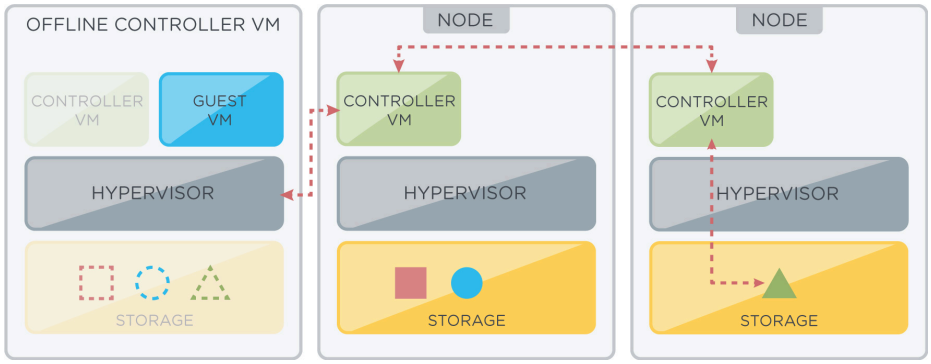

The Nutanix cluster automatically selects the optimal path between a hypervisor host and its guest VM data. The Controller VM has multiple redundant paths available, which makes the cluster more resilient to failures.

When available, the optimal path is through the local Controller VM to local storage devices. In some situations, the data is not available on local storage, such as when a guest VM was recently migrated to another host. In those cases, the Controller VM directs the read request across the network to storage on another host through the Controller VM of that host.

Data Path Redundancy also responds when a local Controller VM is unavailable. To maintain the storage path, the cluster automatically redirects the host to another Controller VM. When the local Controller VM comes back online, the data path is returned to this VM.

Cluster Components

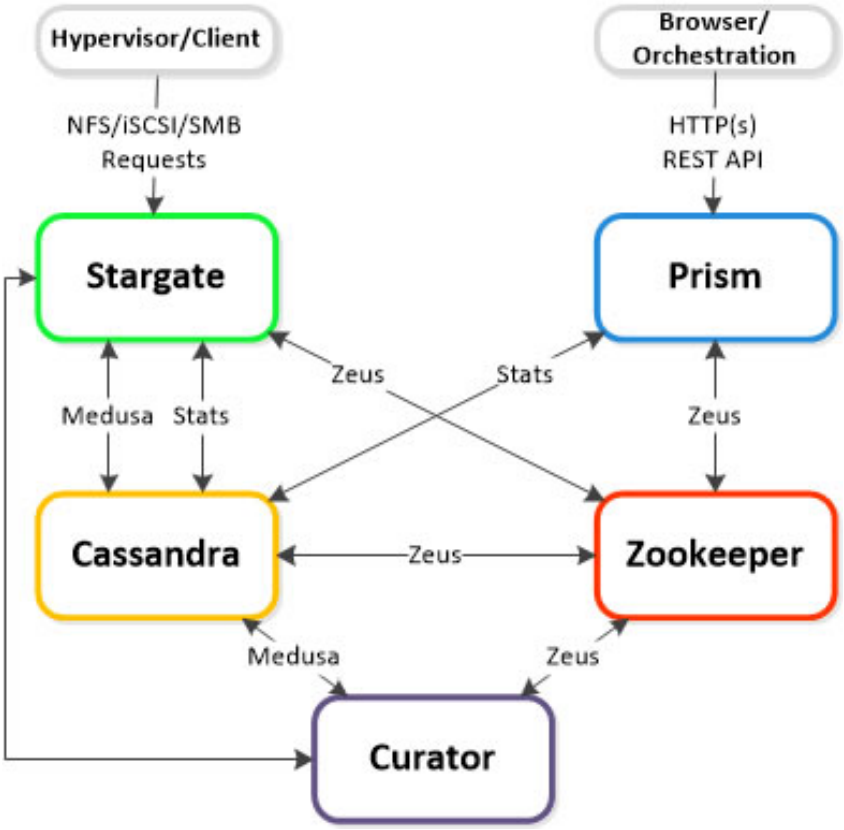

The Nutanix cluster has a distributed architecture, which means that each node in the cluster shares in the management of cluster resources and responsibilities. Within each node, there are software components that perform specific tasks during cluster operation.

All components run on multiple nodes in the cluster, and depend on connectivity between their peers that also run the component. Most components also depend on other components for information.

Zeus

A key element of a distributed system is a method for all nodes to store and update the cluster's configuration. This configuration includes details about the physical components in the cluster, such as hosts and disks, and logical components, like storage containers. The state of these components, including their IP addresses, capacities, and data replication rules, are also stored in the cluster configuration.

Zeus is the Nutanix library that all other components use to access the cluster configuration, which is currently implemented using Apache Zookeeper.

Zookeeper

Zookeeper runs on either three or five nodes, depending on the redundancy factor that is applied to the cluster. Using multiple nodes prevents stale data from being returned to other components, while having an odd number provides a method for breaking ties if two nodes have different information.

Of these three nodes, one Zookeeper node is elected as the leader. The leader receives all requests for information and confers with the two follower nodes. If the leader stops responding, a new leader is elected automatically.

Zookeeper has no dependencies, meaning that it can start without any other cluster components running.

Medusa

Distributed systems that store data for other systems (for example, a hypervisor that hosts virtual machines) must have a way to keep track of where that data is. In the case of a Nutanix cluster, it is also important to track where the replicas of that data is stored.

Medusa is a Nutanix abstraction layer that sits in front of the database that holds this metadata. The database is distributed across all nodes in the cluster, using a modified form of Apache Cassandra.

Cassandra

Cassandra is a distributed, high-performance, scalable database that stores all metadata about the guest VM data stored in a Nutanix datastore. In the case of NFS datastores, Cassandra also holds small files saved in the datastore. When a file reaches 512K in size, the cluster creates a vDisk to hold the data.

Cassandra runs on all nodes of the cluster. These nodes communicate with each other once a second using the Gossip protocol, ensuring that the state of the database is current on all nodes.

Cassandra depends on Zeus to gather information about the cluster configuration.

Stargate

A distributed system that presents storage to other systems (such as a hypervisor) needs a unified component for receiving and processing data that it receives. The Nutanix cluster has a large software component called Stargate that manages this responsibility.

From the perspective of the hypervisor, Stargate is the main point of contact for the Nutanix cluster. All read and write requests are sent across vSwitchNutanix to the Stargate process running on that node.

Stargate depends on Medusa to gather metadata and Zeus to gather cluster configuration data.

> Tip: If Stargate cannot reach Medusa, the log files include an HTTP timeout. Zeus communication issues can include a Zookeeper timeout.

Curator

In a distributed system, it is important to have a component that watches over the entire process. Otherwise, metadata that points to unused blocks of data could pile up, or data could become unbalanced, either across nodes, or across disk tiers.

In the Nutanix cluster, each node runs a Curator process that handles these responsibilities. A Curator master node periodically scans the metadata database and identifies cleanup and optimization tasks that Stargate or other components should perform. Analyzing the metadata is shared across other Curator nodes, using a MapReduce algorithm.

Curator depends on Zeus to learn which nodes are available, and Medusa to gather metadata. Based on that analysis, it sends commands to Stargate.

Prism

A distributed system is worthless if users can't access it. Prism provides a management gateway for administrators to configure and monitor the Nutanix cluster. This includes the nCLI and web console.

Prism runs on every node in the cluster, and like some other components, it elects a leader. All requests are forwarded from followers to the leader using Linux iptables. This allows administrators to access Prism using any Controller VM IP address. If the Prism leader fails, a new leader is elected.

Prism communicates with Zeus for cluster configuration data and Cassandra for statistics to present to the user. It also communicates with the ESXi hosts for VM status and related information.

Availability Domains

Hardware failures are a part of any datacenter lifecycle. The Nutanix architecture was designed with this inevitability in mind. A cluster can tolerate one or two failures (depending on the replication factor of the cluster or container) of a variety of hardware components while still running guest VMs and responding to commands through the management console. Many of these failures also trigger an alert through that same management console in order to give the administrator a chance to respond to the situation.

Nutanix provides the ability to tolerate rack failures for extended data availability, in addition to drive, node, block, and network link failure.

- Rack Fault Tolerance on page 17
- Block Fault Tolerance on page 22
- Configuring Node Fault Tolerance on page 27

Node Failure

A Nutanix node is comprised of a physical host and a Controller VM. Either component can fail without impacting the rest of the cluster.

Controller VM Failure

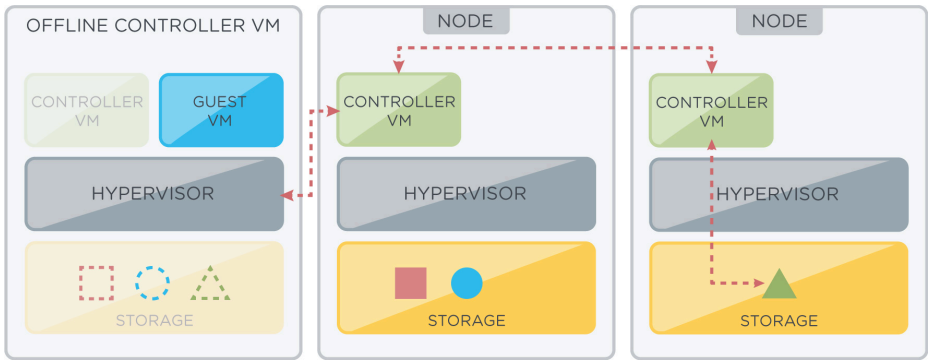

The Nutanix cluster monitors the status of Controller VMs in the cluster. If any Stargate process fails to respond two or more times in a 30-second period, another Controller VM redirects the storage path on the related host to another Stargate. Reads and writes occur over the 10 GbE network until the missing Stargate comes back online.

To prevent constant switching between Stargates, the data path is not restored until the original Stargate has been stable for 30 seconds.

What Will Users Notice?

During the switching process, the host with a failed Controller VM may report that the shared storage is unavailable. Guest VMs on this host appear to "hang" until the storage path is restored. Although the primary copy of the guest VM data is unavailable because it is stored on disks mapped to the failed Controller VM, the replicas of that data are still accessible.

As soon as the redirection takes place, VMs can resume reads and writes. Performance may decrease slightly, because the IO is traveling across the network, rather than across the internal network. Because all traffic goes across the 10 GbE network, most workloads do not diminish in a way that is perceivable to users.

What Happens If Another One Fails?

A second Controller VM failure has the same impact on the VMs on the other host, which means there will be two hosts sending IO requests across the network. More importantly, however, is the additional risk to guest VM data. With two Controller VMs unavailable, there are now two sets of physical disks that are inaccessible. In a cluster with a replication factor 2, there is now a chance that some VM data extents are missing completely, at least until one of the Controller VMs resumes operation.

Host Failure

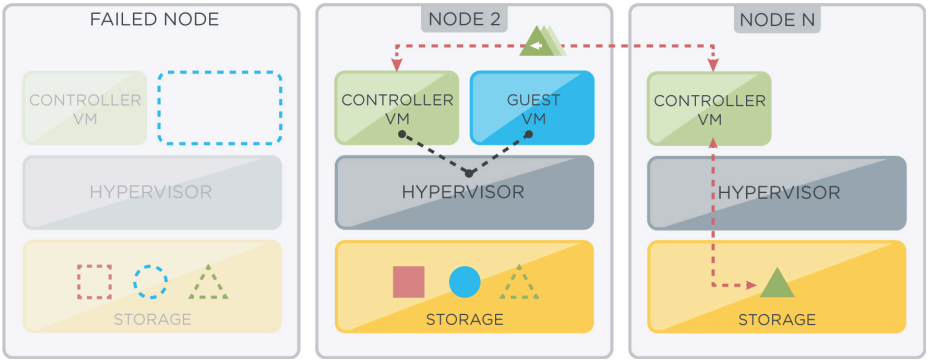

The built-in data redundancy in a Nutanix cluster supports high availability provided by the hypervisor. If a node fails, all HA-protected VMs can be automatically restarted on other nodes in the cluster.

Curator and Stargate responds to two issues that arise from the host failure. First, when the guest VM begins reading across the network, Stargate begins migrating those extents to the new host. This improves performance for the guest VM. Second, Curator notices that there is a missing replica of those extents, and instruct Stargate to begin creating a second replica.

What Will Users Notice?

Users who are accessing HA-protected VMs will notice that their VM is unavailable while it is restarted on the new host. Without HA, the VM will need to be manually restarted.

What Happens If Another One Fails?

Depending on how fully loaded the cluster is, a second host failure could leave the remaining hosts with insufficient processing power to restart the VMs from the second host. However, even in a lightly-loaded cluster, the bigger concern is additional risk to guest VM data. With two sets of physical disks that are inaccessible, there would be a chance that some VM data extents are missing completely and the IO requests are not served.

Drive Failures

Drives in a Nutanix node store four primary types of data: persistent data (hot-tier and cold-tier), storage metadata, oplog, and Controller VM boot files. Cold-tier persistent data is stored in the hard-disk drives of the node. Storage metadata, oplog, hot-tier persistent data, and Controller VM boot files are kept in the SATA-SSD in drive bay one. SSDs in a dual SSD system are used for storage metadata, oplog, hot-tier persistent data according to the replication factor of system and in a RAID-1 configuration for CVM files. In all-flash nodes, data of all types is stored in the SATA-SSDs.

Note: On hardware platforms that contain PCIe-SSD drives, the SATA-SSD holds the only the Controller VM boot files. Storage metadata, oplog, and hot-tier persistent data reside on the PCIe-SSD.

Boot Drive Failure

Each Controller VM boots from a SATA-SSD. During cluster operation, this drive also holds component logs and related files.

A boot drive failure eventually causes the Controller VM to fail. The host does not access the boot drive directly, so other guest VMs can continue to run. Data Path Redundancy redirects the storage path to another Controller VM. For more information, see Controller VM Failure on page 14.

Note: The Controller VM restarts if a boot drive fails, or if you remove a boot drive without marking the drive for removal and the data has not successfully migrated.

Metadata Drive Failure

The metadata "drive" (metadata spans multiple drives) serves many purposes. It holds the oplog for each host, which provides a fast response to VMs that send write requests. It is used as a persistent data tier. It is also used by the Cassandra database to provide fast read and write responses for cluster metadata. To protect against potential failure, the metadata is replicated to metadata drives in other hosts in the cluster.

Metadata is sharded, and Cassandra uses up to four SSDs to store metadata. If any metadata drive fails on a host, the Controller VM restarts. Once the Cassandra process restarts, the missing metadata is retrieved from the other Controller VMs and sharded across the remaining metadata drives. When the faulty drive recovers or is replaced, metadata is stored on that drive again.

If the Cassandra process is down on a node for more than 30 minutes, the surviving Cassandra nodes detach the node from the Cassandra database so that the unavailable metadata can be replicated on other nodes. The process of healing the database takes about 30-40 minutes. If the Cassandra process restarts and remains up for 30 minutes, the procedure to detach the node is canceled if the healing procedure is still running. If the process resumes and is stable after the healing procedure is complete, the node can be manually added to the database using the nCLI command:

```
ncli> host enable-metadata-store id=host_id
```

What Will Users Notice?

Performance may decrease slightly for user VMs on the host due to the reboot and the fact that some I/O is traveling across the network. However, most workloads should not diminish in a way that is perceivable to users (other than during the reboot).

What Happens If Another One Fails?

Multiple drive failures in a single selected domain (node, block, or rack) is tolerated.

Data Drive Failure

Each node contributes data drives to the cluster storage pool. Cold-tier data is stored in HDDs, while hot-tier data is stored in SSDs for faster performance. Because the HDDs have moving parts, and outnumber any other hardware component, this is the most likely component to experience a failure. Data is replicated across the cluster, so a single hard-disk drive failure does not result in data loss. In all-flash nodes, data of all types is stored in SSDs.

The cluster software receives a hardware alert from the host that a data drive (HDD or SSD) has failed, and immediately begin working to reduce the impact of a second failure. Curator instructs Stargate to create a second replica of any guest VM data that was stored on the drive.

> Note: The Controller VM restarts if a data drive fails, or if you remove a data drive without marking the drive for removal and the data has not successfully migrated.

What Will Users Notice?

For a brief period of time, guest VMs with files on the failed data drive will need to read across the network. Curator will eventually instruct Stargate to migrate the relevant data to another drive on the current host.

What Happens If Another One Fails?

In a cluster with a replication factor 2, losing two drives on different different domains (node, block, or rack) means that some VM data extents could lose both replicas. Although a single drive failure does not have the same impact as a host failure, it is important to replace the failed drive as soon as possible.

Network Link Failure

The physical network adapters on each host are grouped together on the external network . Failure of a network link is tolerated with no impact to users if both 10 GbE ports are connected to the customer's network. If the cluster must fail over to a 1 GbE link, then write performance decreases.

What Will Users Notice?

Each Nutanix node is configured at the factory to use one 10 GbE port as the primary pathway for vSwitch0. The other 10 GbE port is configured in standby mode. In this configuration, guest VM performance does not decrease. If a 10 GbE port is not configured as the failover path, then traffic will fail over to a 1 GbE port. This failover reduces the throughput of storage traffic, and decreases the write performance for guest VMs on the host with the failed link. Other hosts may experience a slight decrease as well, but only on writes to extents that are stored on the host with the link failure.

What Happens if Another One Fails?

If both 10 GbE links are down, then the host will fail over to a 1 GbE port, if it is configured as a standby path. This failover reduces the throughput of storage traffic, and decreases the write performance for guest VMs on the host with the failed link. Other hosts may experience a slight decrease as well, but only on writes to extents that are stored on the host with the link failure.

Rack Fault Tolerance

Rack fault tolerance is the ability to provide rack level availability domain. With rack fault tolerance, redundant copies of data are made and placed on the nodes that are not in the same rack.

Rack failure can occur in the following situations:

- All power supplies fail within a rack
- Top-of-rack (TOR) switch fails
- Network partition; where one of the racks becomes inaccessible from other racks

When rack fault tolerance is enabled, the cluster has rack awareness and the guest VMs can continue to run with failure of one rack (RF2) or two racks (RF3). The redundant copies of guest VM data and metadata exist on other racks when one rack fails.

> Note:
>
> - If sufficient number of working racks are not present after failure, the data is rebuilt in a non-rack-aware manner to get node fault-tolerance level to one. The rack-fault-tolerance stays 0.
> - Rack Fault Tolerance is supported for AHV and ESXi only.

Configuring Rack Fault Tolerance

To configure rack fault tolerance, do the following.

About this task

To enable rack fault tolerance, you must specify the mapping of the blocks to the racks based on the actual placement of the blocks in the datacenter.

Before you begin

- You must have information on the actual physical mapping of racks and blocks in the datacenter.
- Minimum cluster requirements:
  - Replication factor 2 - 3 racks (4 with Erasure Coding)
  - Replication factor 3 - 5 racks (6, with Erasure Coding), 1 node in each rack

Procedure

1. Go to **Settings** > **Setup** > **Rack Configuration.**.

> Note: Directory List (AD and OpenLDAP) users can view and configure Rack Awareness only if a service account is configured for the directory service. To configure a service account for the directory service, see “” () for details.

2. Select **Rack** for **Choose the fault tolerance domain**.

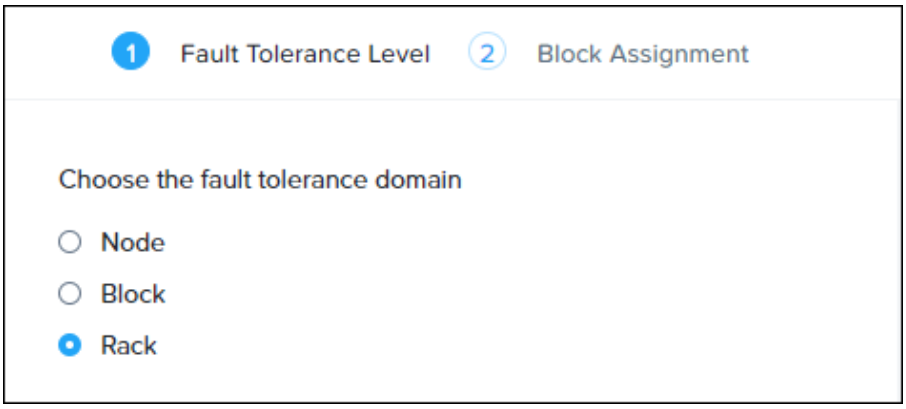

Figure 1: Select Fault Tolerance Level

Click **Add New Rack** and click **Next**.

3. Click **Add Blocks** to assign the discovered blocks to the rack. Click the plus icon to add the block to the rack based on the actual placement of the block in the datacenter.

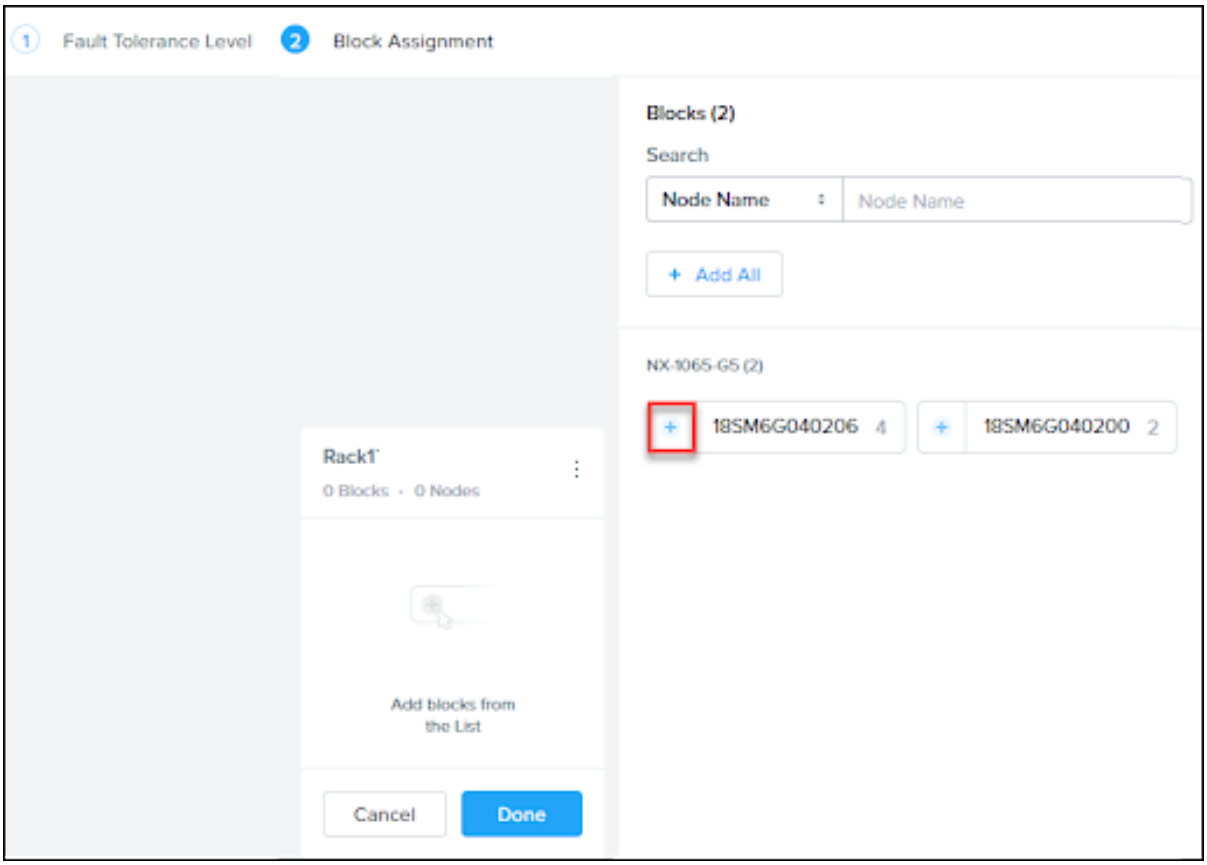

Figure 2: Add blocks to rack

Click **Add New Rack** and click **Next**.

Repeat steps 3 to 5 for adding more rack configurations. Verify the configuration and click **Save**.

Note:

- Once rack fault tolerance is enabled; create, edit, delete operations are not allowed. To make any new modifications, you must re-configure the fault tolerance level as **Node**, see Configuring Node Fault Tolerance on page 27.
- Enabling rack fault tolerance may trigger cassandra ring change operations and trigger curator scans to redistribute replicas for rack fault tolerance. The time

taken for these depend on existing data on cluster, workload on the cluster, and configurations.

The **Data Resiliency** widget on the Prism dashboard shows the current state of the fault tolerance.

Figure 3: Data Resiliency widget

Data Resiliency Levels for Rack Fault Tolerance

Table 1: Data Resiliency Levels

The following table shows the level of data resiliency (simultaneous failure) provided for the following combinations of replication factor, minimum number of nodes, minimum number of blocks, and minimum number of racks.

Fault Domain	Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Minimum Number of Racks	Data Resiliency
Node	2	3*	1	1	1 node or 1 disk failure
	3	5*	2	1	2 nodes or 2 disk failures
Block	2	3	3*	1	1 node or 1 block or 1 disk failure
	3	5	5*	1	2 nodes or 2 blocks or 2 disk failures
Rack	2	3	3	3*	1 node or 1 block or 1 rack or 1 disk

Fault Domain	Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Minimum Number of Racks	Data Resiliency
	3	5	5	5*	2 nodes or 2 blocks or 2 racks or 2 disks

* Erasure Coding with Rack Awareness - Erasure coding is supported on a rack ware cluster. You can enable erasure coding on new containers in rack aware clusters provided certain minimums are met that are shown in the table above.

Block Fault Tolerance

Block fault tolerance lets a Nutanix cluster make redundant copies of data and metadata and place the copies on nodes in different blocks.

A block is a rack-mountable enclosure that contains one to four Nutanix nodes. All nodes in a block share power supplies, front control panels (ears), backplane, and fans.

Nutanix offers block fault tolerance as an opt-in procedure, as in Configuring Block Fault Tolerance on page 22, or a best-effort procedure, as in Block Fault Tolerance in Best Effort mode on page 23. The opt-in block fault tolerance feature offers guaranteed data resiliency when required conditions are met. For best-effort fault tolerance mode, data copies remain on the same block when there is insufficient space across all blocks.

With block fault tolerance enabled, guest VMs can continue to run after a block failure because redundant copies of guest VM data and metadata exist on other blocks.

Configuring Block Fault Tolerance

To configure block fault tolerance, do the following.

About this task

Procedure

1. Go to **Settings** > **Setup** > **Rack Configuration.**.
2. Select **Block** for **Choose the fault tolerance domain**.

Figure 4: Select Fault Tolerance Level

3. Optionally, click **Add New Rack** to specify rack and block mapping.

Block Fault Tolerance in Best Effort mode

When certain conditions are met, Nutanix clusters become block fault tolerant. Block fault tolerance is applied automatically when:

- Every storage tier in the cluster contains at least one drive on each block.
- Every storage container in the cluster has replication factor of at least two.
- For replication factor 2, there are a minimum of three blocks in the cluster.
- There is enough free space in all the tiers, in at least replication factor number of blocks in the cluster. For example, if the replication factor of storage containers in the cluster is two, then at least two blocks require free space.
- A minimum of four blocks for RF2 or six blocks for RF3 is required to maintain block awareness if erasure coding is enabled on any storage container. (If the cluster has fewer blocks, block awareness is lost when erasure coding is enabled.)

> Note: This is not applicable for single-node replication target clusters. For more information on how single-node replication target clusters handle failures, see Single-Node Replication Target Clusters in the guide titled *Data Protection and Recovery with Prism Element*.
>
> Once block fault tolerance conditions are met, the cluster can tolerate a specific number of block failures:
>
> - A replication factor two or replication factor three cluster with three or more blocks can tolerate a maximum failure of one block.
> - A replication factor three cluster with five or more blocks can tolerate a maximum failure of two blocks.

Block fault tolerance is one part of a resiliency strategy. It does not remove other constraints such as the availability of disk space and CPU/memory resources in situations where a significant proportion of the infrastructure is unavailable.

Metadata Block Awareness Requirements for Block Fault Tolerance

Metadata block awareness is required for block fault tolerance. Metadata block awareness can be enabled for replication factor 2 and replication factor 3. To enable metadata block fault tolerance, your metadata must meet the following requirements.

Table 2: Minimum Cluster Requirements

Replication Factor	Minimum Number of Blocks and Nodes Per Block Required
Replication factor 2	• 3 blocks • 1 node per block
Replication factor 3	• 5 blocks • 1 node per block

Table 3: Additional Requirements when Adding Nodes to an Existing Block Aware Cluster

Replication Factor	Requirement	Example
Note: Be sure your cluster has met the previous minimum cluster requirements.		
Replication factor 2	There must be at least 3 blocks populated with a specific number of nodes to maintain block fault tolerance. To calculate the number of nodes required to maintain block fault tolerance when the cluster RF=2, you need twice the number of nodes as there are in the block with the most or maximum number of nodes.	If a block contains 4 nodes, you need 8 nodes distributed across the remaining (non-failing) blocks to maintain block fault tolerance for that cluster. X = number of nodes in the block with the most nodes. In this case, 4 nodes in a block. 2X = 8 nodes in the remaining blocks.
Replication factor 3	There must be at least 5 blocks populated with a specific number of nodes to maintain block fault tolerance. To calculate the number of nodes required to maintain block fault tolerance when the cluster replication factor 3 you need four times the number of nodes as there are in the block with the most or maximum number of nodes.	If a block contains 4 nodes, you need 16 nodes distributed across the remaining (non-failing) blocks to maintain block fault tolerance for that cluster. X = number of nodes in the block with the most nodes. In this case, 4 nodes in a block. 4X = 16 nodes in the remaining blocks

Block Fault Tolerant Data Placement

Stargate is responsible for placing data across blocks, and Curator makes data placement requests to Stargate to maintain block fault tolerance.

New and existing clusters can reach a block fault tolerant state. New clusters can be block fault tolerant immediately after being created if the configuration supports it. Existing clusters that were not previously block fault tolerant can be made tolerant by reconfiguring the cluster in a manner that supports block fault tolerance.

New data in a block fault tolerant cluster is placed to maintain block fault tolerance. Existing data that was not in a block fault tolerant state is moved and scanned by Curator to a block fault tolerant state.

Depending on the volume of data that needs to be relocated, it might take Curator several scans over a period of hours to distribute data across the blocks.

Block fault tolerant data placement is on a best effort basis but is not guaranteed. Conditions such as high disk usage between blocks may prevent the cluster from placing guest VM redundant copy data on other blocks.

Guest VM Data

Redundant copies of guest VM data are written to nodes in blocks other than the block that contains the node where the VM is running. The cluster keeps two copies of each write stored in the oplog.

Redundant copies of the guest VM data (designated by ■ ● ▲) are placed on different blocks.

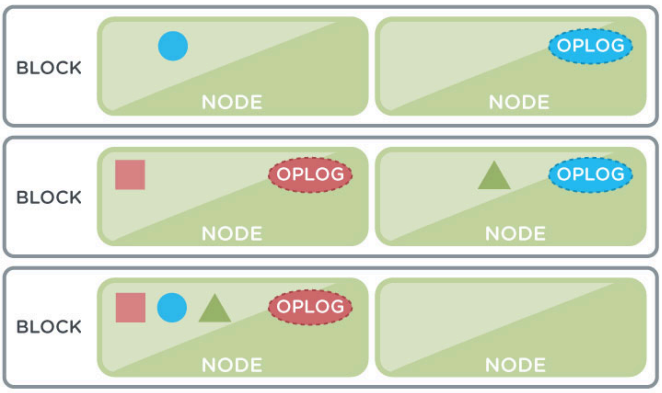

Figure 5: Block-aware placement of guest VM data

In the case of a block failure, the under-replicated guest VM data is copied to other blocks in the cluster, and one copy of the oplog contents is available.

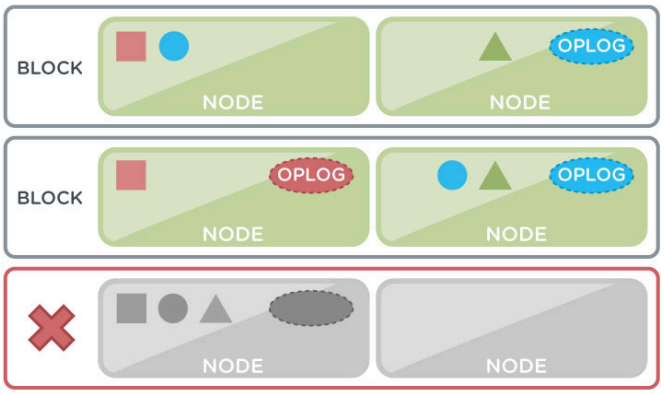

Figure 6: Block-aware placement of guest VM data with block failure

Metadata

The Nutanix Medusa component uses Cassandra to store metadata. Cassandra uses a ring-like structure where data is copied to peers within the ring to ensure data consistency and availability. The cluster keeps at least three redundant copies of the metadata, at least half of which must be available to ensure consistency.

With block fault tolerance, the Cassandra peers are distributed among the blocks to ensure that no two peers are on the same block. In the event of a block failure, at least two copies of the metadata is present in the cluster .

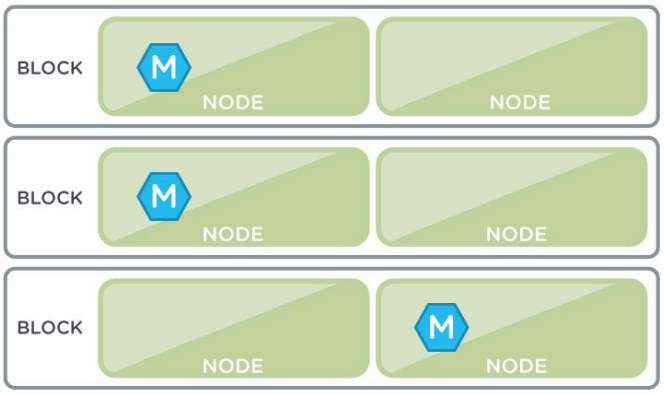

Figure 7: Block-aware placement of metadata

Configuration Data

The Nutanix Zeus component uses Zookeeper to store essential configuration data for the cluster. The Zookeeper role is distributed across blocks to ensure availability in the case of a block failure.

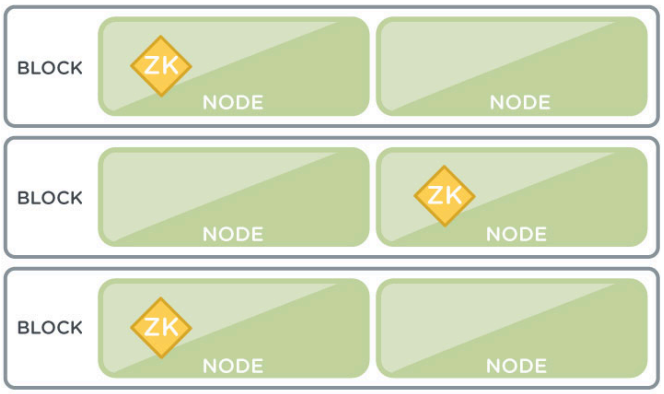

Figure 8: Block-aware placement of configuration data

Data Resiliency Levels for Block Fault Tolerance

Table 4: Data Resiliency Levels

The following table shows the level of data resiliency (simultaneous failure) provided for the following combinations of replication factor, minimum number of nodes, and minimum number of blocks.

Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Data Resiliency
2	3	1	1 node or 1 disk failure
2	3	3 (minimum 1 node each)	1 block or 1 node or 1 disk failure
3	5	2	2 nodes or 2 disk failures

Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Data Resiliency
3	5	5 (minimum 1 node each)	2 blocks or 2 nodes or 2 disks
3	6	3 (minimum 2 nodes each)	1 block or 2 nodes or 2 disks
Metro Cluster (see See "Metro Availability" in the "Data Protection" chapter of the *Prism Web Console Guide*.)	3 nodes at each site	2	1 cluster failure

The state of block fault tolerance is available for viewing through the Prism Web Console and Nutanix CLI. Although administrators must set up the storage tiers or storage containers, they cannot determine where data is migrated. AOS determines where data is migrated.

Web Console	Data Resiliency Status view on the **Home** screen
nCLI	`ncli> cluster get-domain-fault-tolerance-status type="rackable_unit"`

Configuring Node Fault Tolerance

To configure node fault tolerance, do the following.

About this task

Procedure

1. Go to **Settings** > **Setup** > **Rack Configuration.**.
2. Select **Node** for **Choose the fault tolerance domain**.

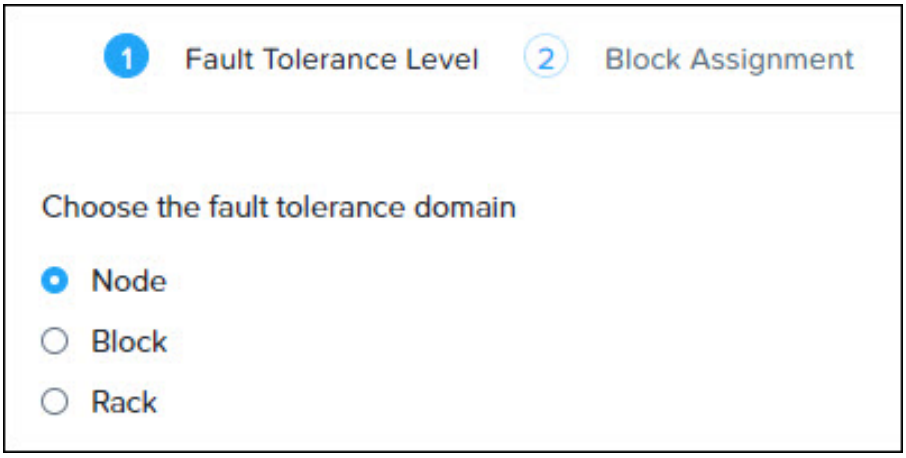

Figure 9: Select Fault Tolerance Level

3. Optionally, click **Add New Rack** to specify rack and block mapping.
   It recommended to configure rack and block mapping if you want to enable rack awareness at a later stage.

Redundancy Factor 3

Redundancy factor 3 is a configurable option that allows a Nutanix cluster to withstand the failure of two nodes or drives in different blocks.

By default, Nutanix clusters have redundancy factor 2, which means they can tolerate the failure of a single node or drive. The larger the cluster, the more likely it is to experience multiple failures. Without redundancy factor 3, multiple failures cause cluster unavailability until the failures are repaired.

Redundancy factor 3 has the following requirements:

- A cluster must have at least five nodes, blocks, racks for redundancy factor 3 to be enabled.
- For guest VMs to tolerate the simultaneous failure of two nodes or drives in different blocks, the data must be stored on storage containers with replication factor 3.

The state of fault tolerance is available to view through the management interfaces.

Web Console	Data Resiliency Status view on the **Home** screen
nCLI	`ncli> cluster get-redundancy-state`

Guest VM Data

For storage containers with replication factor 3, the cluster stores three redundant copies of guest VM data and the oplog.

Redundant copies of the guest VM data (designated by ■ ● ▲) are stored on different nodes in the cluster.

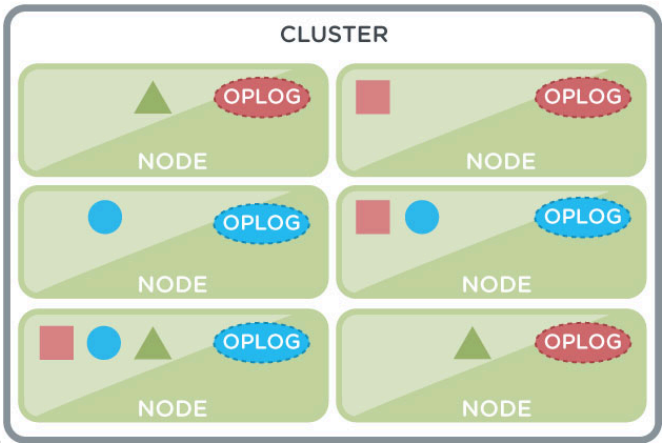

Figure 10: Replication factor 3 placement of guest VM data

In the case of two nodes failing, at least one copy of all guest VM data, including the oplog, is available. Under-replicated VM data is copied to other nodes.

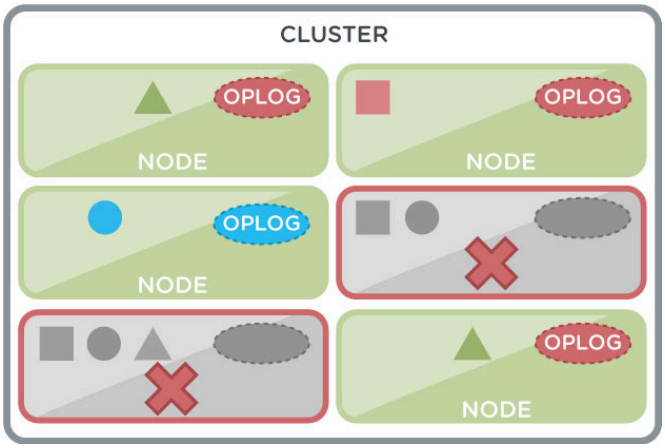

Figure 11: Replication factor 3 placement of guest VM data with failure of 2 nodes

Metadata

At least half of the redundant copies of metadata must be available to ensure consistency. Without redundancy factor 3, the cluster keeps three copies of metadata. With redundancy factor 3, the cluster keeps five copies of metadata so that if two nodes fail, at least three copies are available.

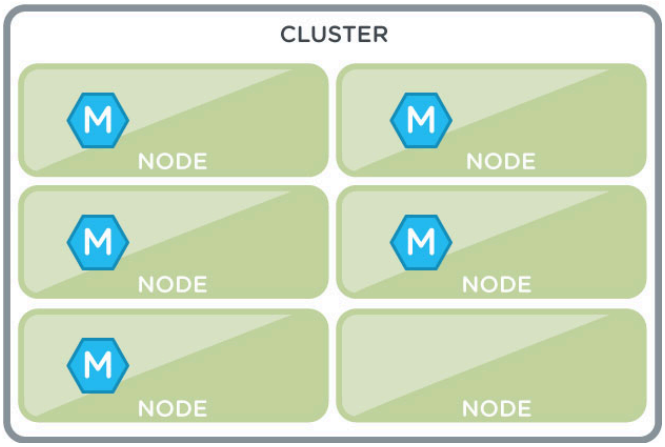

Figure 12: Redundancy factor 3 metadata

Configuration Data

The Nutanix Zeus component uses Zookeeper to store essential configuration data for the cluster. Without redundancy factor 3, the cluster keeps three copies of configuration data. With redundancy factor 3, the cluster keeps five copies of configuration data so that if two nodes fail, at least three copies are available.

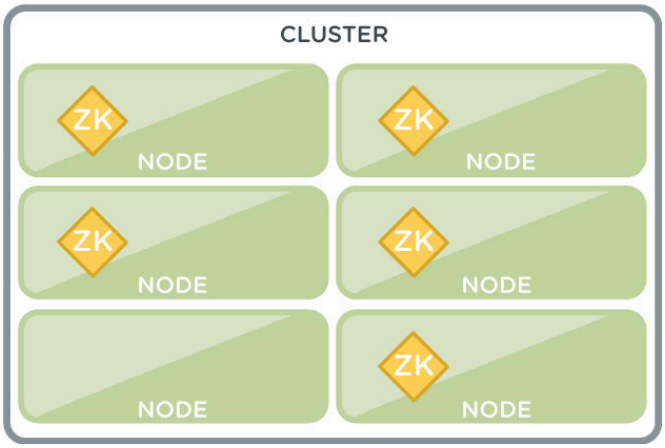

Figure 13: Redundancy factor 3 configuration data

Degraded Node

A degraded node contains a CVM that is not entirely in an unresponsive state, but cannot reliably perform normal cluster operations. A degraded (or partially available) node can adversely affect the performance of an entire cluster. Degraded node detection is introduced to mitigate the results of a degraded node. Continuing to run a degraded node can affect overall cluster and user VM performance. By enabling degraded node detection, ZooKeeper places the node into maintenance mode, which moves it from the leadership position and distributes tasks automatically.

The alert for a degraded node appears in the Alerts page of the web console.

The following events might contribute to node degradation.

- Network bandwidth reduction
- Network packet drops
- Soft lockups
- Partially bad disks
- Hardware issues (such as unreliable DIMM with ECC errors)

> Note: In the event of a degraded node, contact Nutanix Support for help investigating the issue. A support case will be automatically generated if Pulse is enabled.

Overview

Degraded node detection consists of a peer health global database. These services run on each node of the cluster. Each node publishes scores or votes for the service running on other nodes. The peer scores evaluate metrics such as:

- Remote Procedure Call (RPC) latency
- RPC failures or timeouts
- Network latency

If one node consistently receives poor scores for approximately 10 minutes then the peers mark that node as a degraded node. Clustering algorithms are used to identify outlier scores.

> Note:

- A node is not marked as *Degraded* if the current cluster fault tolerance is less than the desired value.
- Containers or data stores might be unavailable for 10 minutes until the node is marked as degraded.
- Upgrades and break fixes are not allowed until the degraded node is fixed.

Detection

The web console alerts the user when a node begins to receive poor peer scores. The node is marked as degraded after a few minutes of poor scores.

The degraded node detection software mitigates the effects of failures and reduces the overall impact to the cluster. To mitigate the impact, this software can perform one of the following actions. When prompted, select one of the following.

- Prevent components on the degraded node from acquiring leadership roles
- Place the degraded node CVM in maintenance mode and reboots the CVM to stop the services
- Shut down the host

Enabling Degraded Node Detection

Enable the degraded node detection settings in the web console.

About this task

Procedure

1. Log into the web console for your cluster.
2. Click the gear icon in the main menu and then select **Degraded Node Settings** in the **Settings** page.
   The Degraded Node Settings dialog box appears.

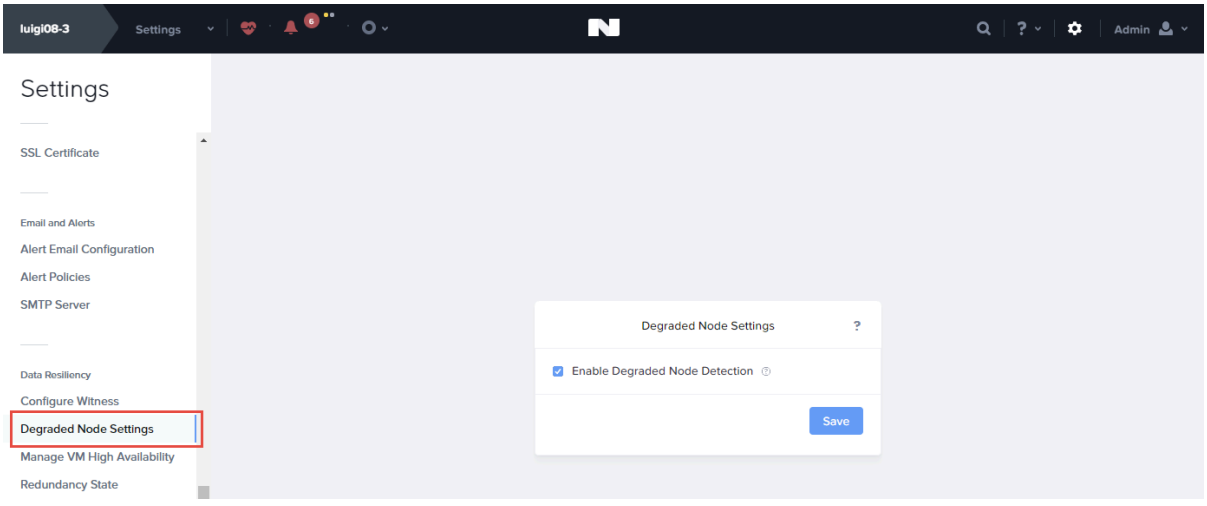

Figure 14: Degraded Node Settings Window

3. Set the degraded node detection status, **Enable Degraded Node Detection**.
   a. Be sure to check the box to enable degraded node detection.
   b. Otherwise, be sure to clear the box to disable degraded node detection.
4. Click **Save** to confirm the settings.
   The Degraded Node Settings window closes and the web console displays a pop-up confirmation.

Managing a Degraded Node

Manage a node if it is marked as degraded.

About this task

Once a node is detected as degraded, the leadership and critical services will not be hosted on that node.

Until a degraded node is unmarked from the degraded state, Cassandra services remain in the forwarding state, see Cluster Components on page 12. Once the degraded node is marked as fixed in Prism, Cassandra restarts its services on the node.

The alert for a degraded node appears in the Alerts page of the web console.

Procedure

1. Open the degraded node window.
   a. Click the **Alerts** icon at the top of the page.
   b. Click the alert for the degraded node.

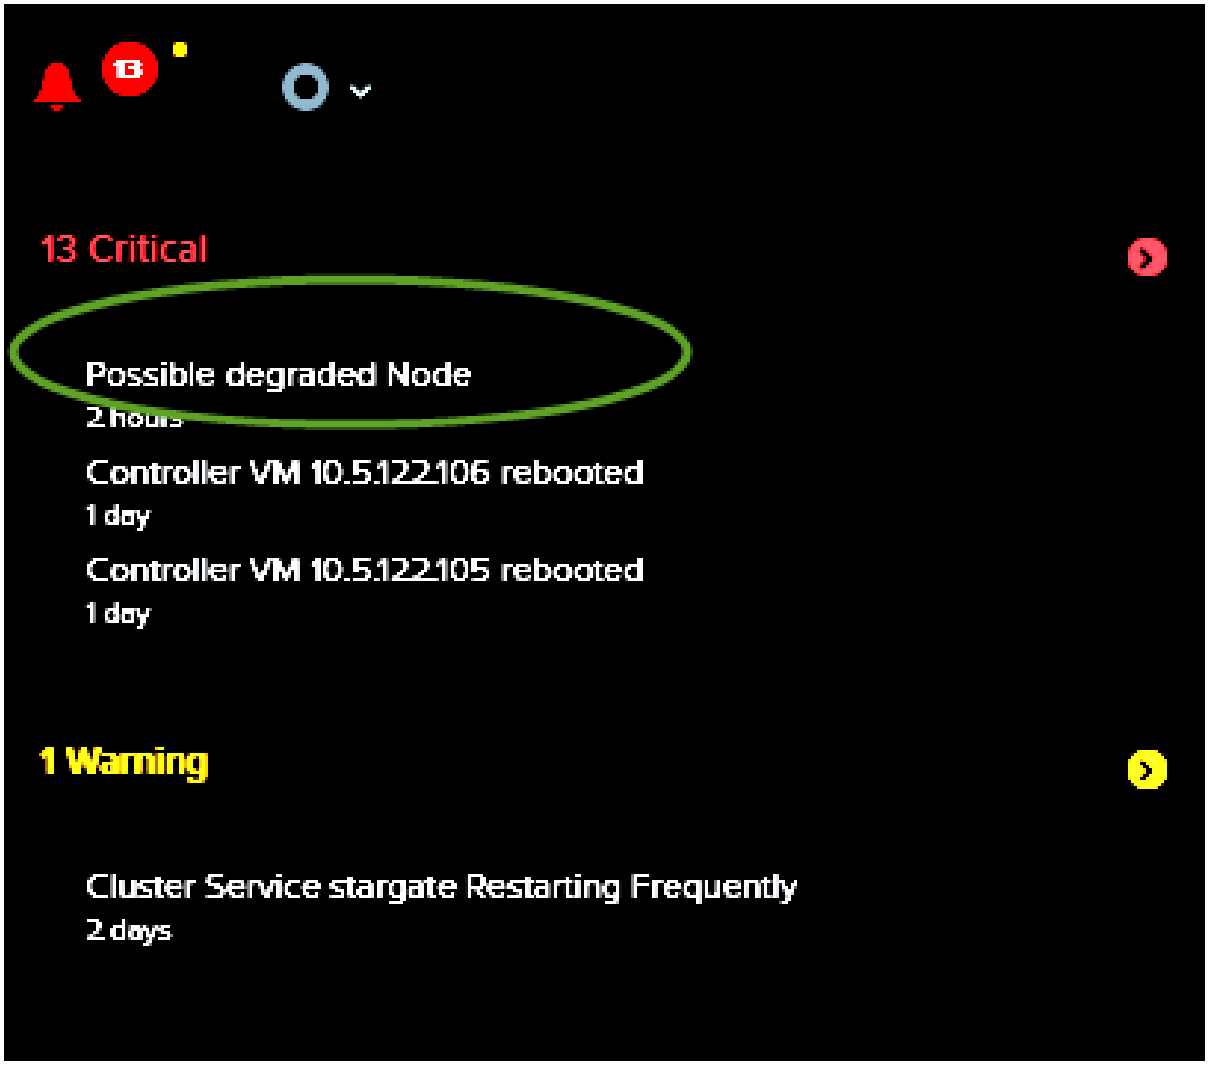

Figure 15: Possible Degraded Node in Alerts

The degraded node window is displayed.

2. Set the status for the degraded node.

   a. **Resolve**: Select this option after you have fixed the underlying issue.

   > Note: Clicking **Resolve** does not resolve the degraded node issue. It merely marks the alert status as **Resolved**, indicating you or Nutanix Support have resolved the underlying issue that triggered the alert.

   b. **Acknowledge**: Select this option to leave the degraded node in its current state, keep the alert unresolved, and mark the alert as acknowledged.

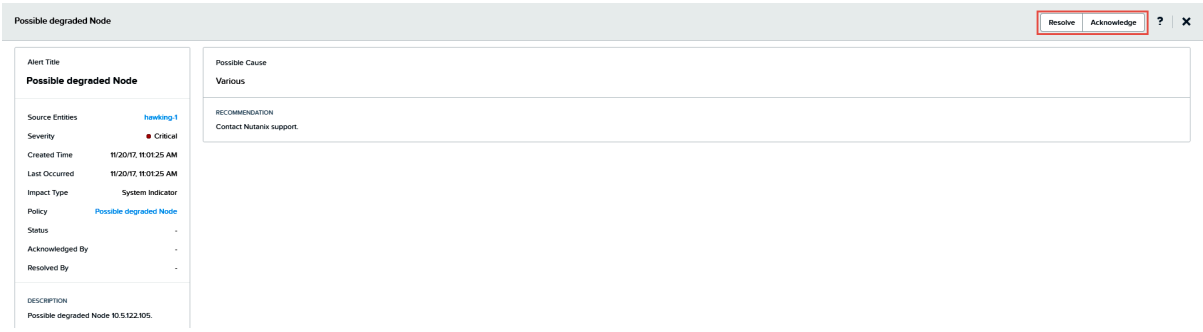

Figure 16: Possible Degraded Node Window

3. You can also view the degraded node in the Hosts table.

   a. Go to **Home** > **Hardware** > **Host**.

   b. Select the degraded node which displays a critical alert.

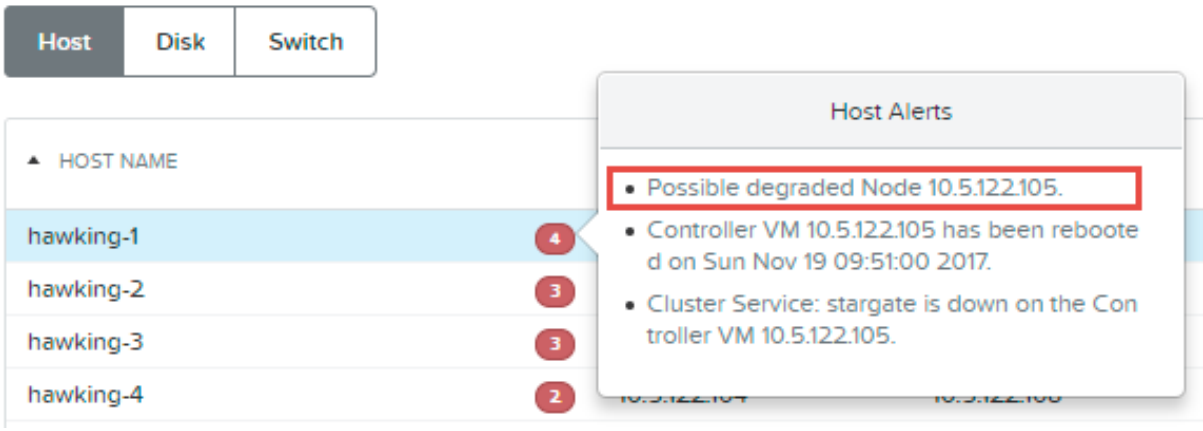

Figure 17: Degraded Node Pop-up Alert

4. To run an NCC check for the degraded node, go to the command line interface and enter

```
nutanix@cvm$ ncc health_checks system_checks degraded_node_check
```

   The web console does not remove the warning sign for the degraded node.

5.

> Note: Do not unmark the degraded node if you have not confirmed that the issue causing the node or CVM degradation has been fixed. Unmarking a node can lead to production impacting issues in the cluster. Contact Nutanix Support for assistance.

To unmark the degraded node, in the **Summary** view, select the fixed node and click **Mark as Fixed**.

Maximum System Values

The figures listed here are the maximum tested and supported values for entities in a Nutanix cluster.

Entity	Supported Maximum
VMs or volume groups	Up to available hardware resources and hypervisor limitations (no known logical limit) • Async DR: 200 VMs or volume groups for each protection domain, or consistency group • Metro Availability: See the Maximum Limits for Metro Availability table in the "Data Protection Guidelines (Metro Availability)" section of the *Data Protection and Recovery with Prism Element* guide. • 1200 files for any single consistency group and 3600 files for all entities per protection domain
vDisks (including snapshots)	600,000
vDisk size (limited by hypervisor or guest OS)	ESXi 5.5 and above = 62 TB
	ESXi 5.1 and below = 2 TB minus 512 B
	Hyper-V vhdx format = 64TB
	HyperV vhd format = 2TB
	AHV = No known limit
vCPUs per VM	Constrained by the number of logical processors available to the VM.
Memory per VM	Constrained by physical RAM capacity.
Memory per host	AHV = 3 TB
	ESXi = Limit as stated by VMware.
	Hyper-V = Limit as stated by Microsoft.
RDM vDisks	256 per vZone or ESXi host
NFS datastores	Nutanix = 256 per host
	ESXi 5.x and later = 256 per host

Entity	Supported Maximum
NFS datastore size	Available disk capacity in the Nutanix cluster (no known logical limit)
Small files (<512 KB) on NFS datastores	2.5 GB per Nutanix node
Storage per node	120 TB
Storage pools	2 per Nutanix cluster
Storage pool size	Available disk capacity in the Nutanix cluster (no known logical limit)
Physical disks per storage pool	Number of disks in the Nutanix cluster (no known logical limit)
Storage Containers (created on Nutanix cluster)	256 per Nutanix cluster
Storage Container size	Available disk capacity in a storage pool (no known logical limit)
Replication factor	3
Nutanix management interface (Prism) simultaneous connections	20 clients per Controller VM
iSCSI	128 LUNs per node
	256 LUNS per Volume Group
	No limit for the number of volume groups if you do not exceed maximum number of LUNs per node

Nutanix clusters are also subject to the vSphere maximum values documented by VMware. For a list of the vSphere maximums, see *Configuration Maximums* for the version of vSphere you are running.

Nutanix supports up to 50 VMs for each storage container if you are using Microsoft VSS-based backups in Hyper-V. If the number of VMs in each storage container increases to more than 50, backup jobs start to get into NFS timeouts.

Three Node Cluster Considerations

A Nutanix cluster must have at least three nodes. Minimum configuration (three node) clusters provide the same protections as larger clusters, and a three node cluster can continue normally after a node failure. However, one condition applies to three node clusters only.

When a node failure occurs in a cluster containing four or more nodes, you can dynamically remove that node to bring the cluster back into full health. The newly configured cluster still has at least three nodes, so the cluster is fully protected. You can then replace the failed hardware for that node as needed and add the node back into the cluster as a new node. However, when the cluster has just three nodes, the failed node cannot be dynamically removed from the cluster. The cluster continues running without interruption on two healthy nodes and one failed node, but the failed node cannot be removed when there are only two healthy nodes. Therefore, the cluster is not fully protected until you fix the problem (such as replacing a bad boot disk) for the existing node.

2

CLUSTER MANAGEMENT

Managing a Nutanix cluster involves configuring and monitoring the entities within the cluster, including virtual machines, storage containers, and hardware components. You can manage a Nutanix cluster through a web-based management console or a command line interface (nCLI).

- The web console is a graphical user interface (GUI) that allows you to monitor cluster operations and perform a variety of configuration tasks (see Web Console Overview on page 37).
- Nutanix employs a license-based system to enable your entitled Nutanix features, and you can install or regenerate a license through the web console (see Licensing on page 61).
- You can upgrade a cluster when a new AOS release is available through the web console (see Software and Firmware Upgrades on page 61.
- If you have multiple clusters, you can manage them all through a single web interface (see Multi-Cluster Management on page 62).

> Note: You can perform most administrative actions using either the web console or nCLI. However, some tasks are only supported in the nCLI either because a new feature has not yet been incorporated into the web console or the task is part of an advanced feature that most administrators do not need to use. See the Nutanix *Command Reference* for information about using the nCLI. See the *AHV Administration Guide* and hypervisor-specific guides for information about platform configuration and hypervisor-specific tasks that are not performed through the web console.

Web Console Overview

The web console, also known as Prism Element, provides a graphical user interface to manage most activities in a Nutanix cluster.

Display Features

The web console screens are divided into the following sections:

- Main menu bar. The main menu bar appears at the top of every screen in the web console. The cluster name appears on the far left of the main menu bar. To the right of the cluster name, you can select an entity from the pull-down list (Home, Health, VM, Storage, Hardware, DR, Analysis, or Alerts) to display information about that entity. You can also search for specific topics or select various tasks from the pull-down lists on the right side of the main menu bar (see Main Menu on page 46). In addition, the main menu bar includes status icons for quick access to heath, alert, and event information.
- Entity views. There is a dashboard view for each entity. Some entities (VM, Storage, Hardware, and Data Protection) include additional views such as a diagram or table view that you can select from the dashboard of that entity.
- Screen menu bar. Some entity dashboards include another menu bar below the main menu that provides options specific to that screen. In the following example from the Storage

dashboard, three view tabs (**Overview**, **Diagram**, and **Table**) and three task buttons (**+ Storage Container**, **+ Volume Group**, and **+ Storage Pool**) appear on this menu bar.

- Usage and performance/health statistics. Most views include fields that provide usage and either performance or health (or both) statistics. The usage and performance/health statistics vary based on the entity that you are viewing. For example, virtual machine usage statistics are displayed in terms of CPU and memory, while disk usage statistics show disk capacity in TB. In most cases, performance statistics are displayed in IOPS, latency, and bandwidth.
- Alert and event messages. Several views include fields that list current event and alert messages. The listed messages are context specific, so for example only storage-related alerts and events appear on the Storage screen. Clicking on a message opens the alert or event view at that message (see Alerts Dashboard on page 327).

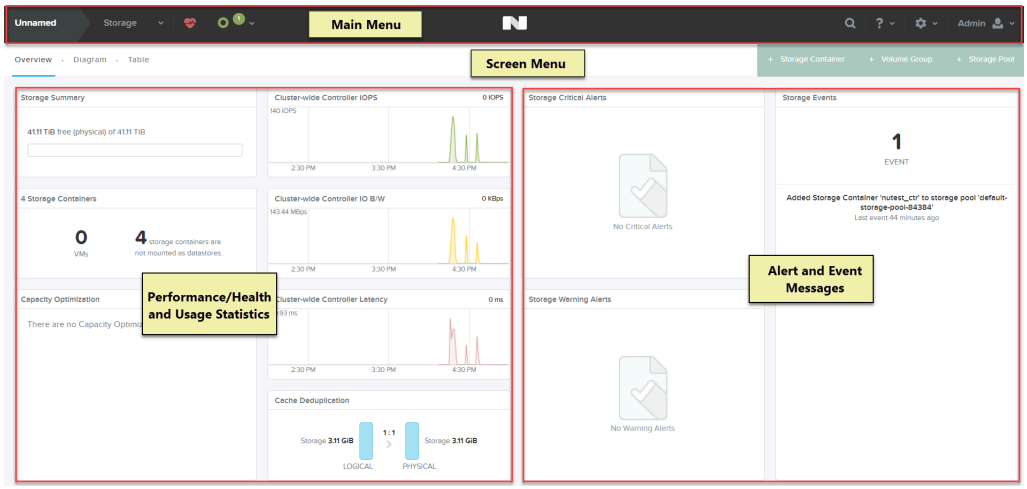

Figure 18: Overview Screen Sections

Logging Into the Web Console

About this task

To log into the web console, do the following. Knowledge base article KB 1661 lists default cluster credentials.

Procedure

1. Open a web browser.

   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.

2. Enter `http://`*`management_ip_addr`* in the address field and press **Enter**. Replace *`management_ip_addr`* with the cluster virtual IP address (if configured) or the IP address of any Nutanix Controller VM in the cluster.

   Note: If you are logging into Prism Central, enter the Prism Central VM IP address.

   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 (see “” ()).

3. If a welcome screen appears (see Configuring a Banner Page on page 615), read the message and then click the "Accept terms and conditions" bar at the bottom.

4. In the login screen, enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

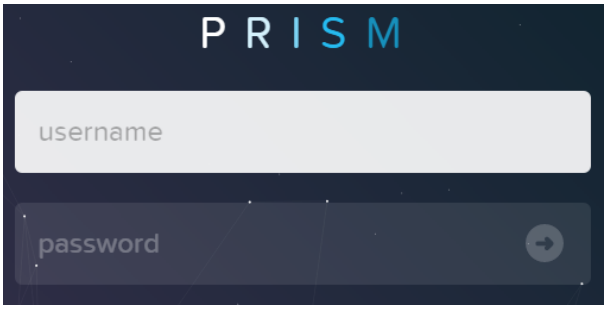

Figure 19: Login Screen

Note: If LDAP authentication is used, enter the user name in the ***samAccountName@domain*** format; the ***domain\user*** format is not supported. (Authentication does not use the user principle name [UPN]; ***user@domain*** is simply a concatenation of the user and domain names specified in “” () .)

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 60.

5. If you are logging in as an administrator (**admin** user name and password) for the first time, which requires that the default password (Nutanix/4u) be changed, enter a new password in the **password** and **re-type password** fields and then press **Enter** or click the right arrow icon.

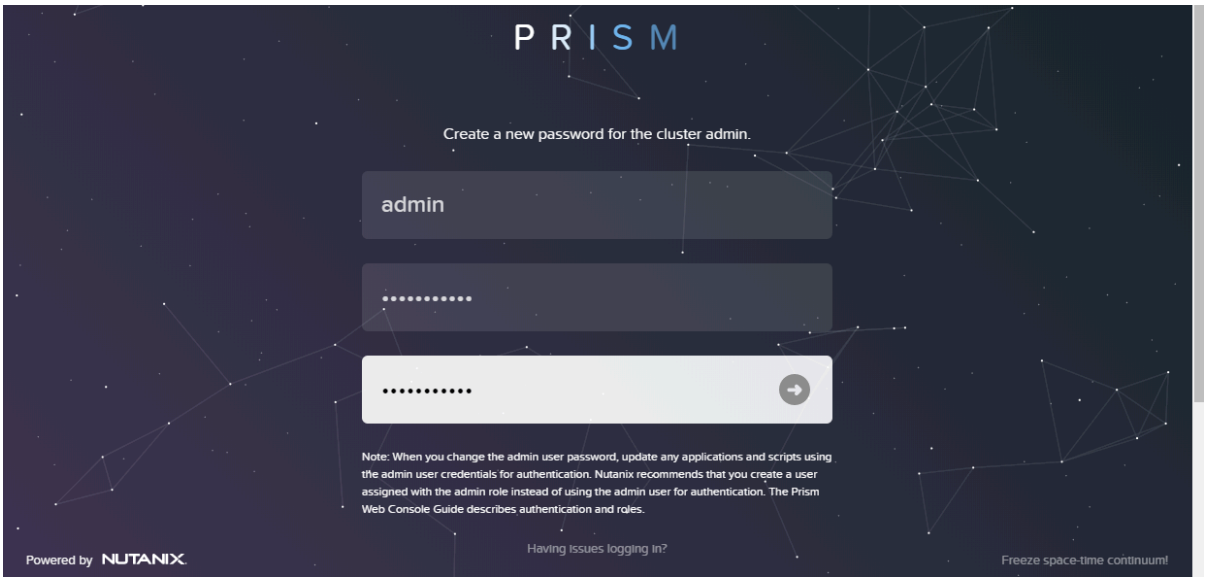

Figure 20: Login Screen (first admin login)

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
- Must not be among the last 5 passwords
- Must not have more than 2 consecutive occurrences of a character
- Must not be longer than 199 characters

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:

- You are prompted to change the password when logging in as the **admin** user for the first time after upgrading AOS. If the first login after upgrade is to the Controller VM through SSH (instead of Prism), you must log in using the default **admin** user password (Nutanix/4u) and then change the password when prompted.
- 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.

6. 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.

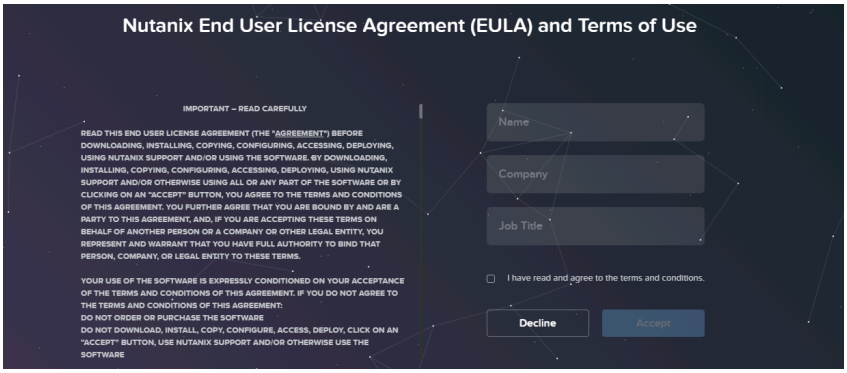

Figure 21: EULA Screen

7. If a "Pulse will be Enabled" screen appears (typically on the first login or after an upgrade), which refers to the Pulse feature that alerts Nutanix customer support regarding the health of the cluster (see Configuring Pulse on page 642), read the statement and then do one of the following:

   » Click the **Continue** button to enable the Pulse feature (recommended).

   » Click the **Disable Pulse (not recommended)** button to disable the Pulse feature.

   CAUTION: If Pulse is not enabled, alerting Nutanix customer support when problems occur is disabled.

   Pulse provides Nutanix customer support with analytic information that allows them to dynamically monitor your cluster and identify potential issues before they become problems

(see Configuring Pulse on page 642). Enabling Pulse is recommended unless providing cluster information to Nutanix customer support violates your security policy.

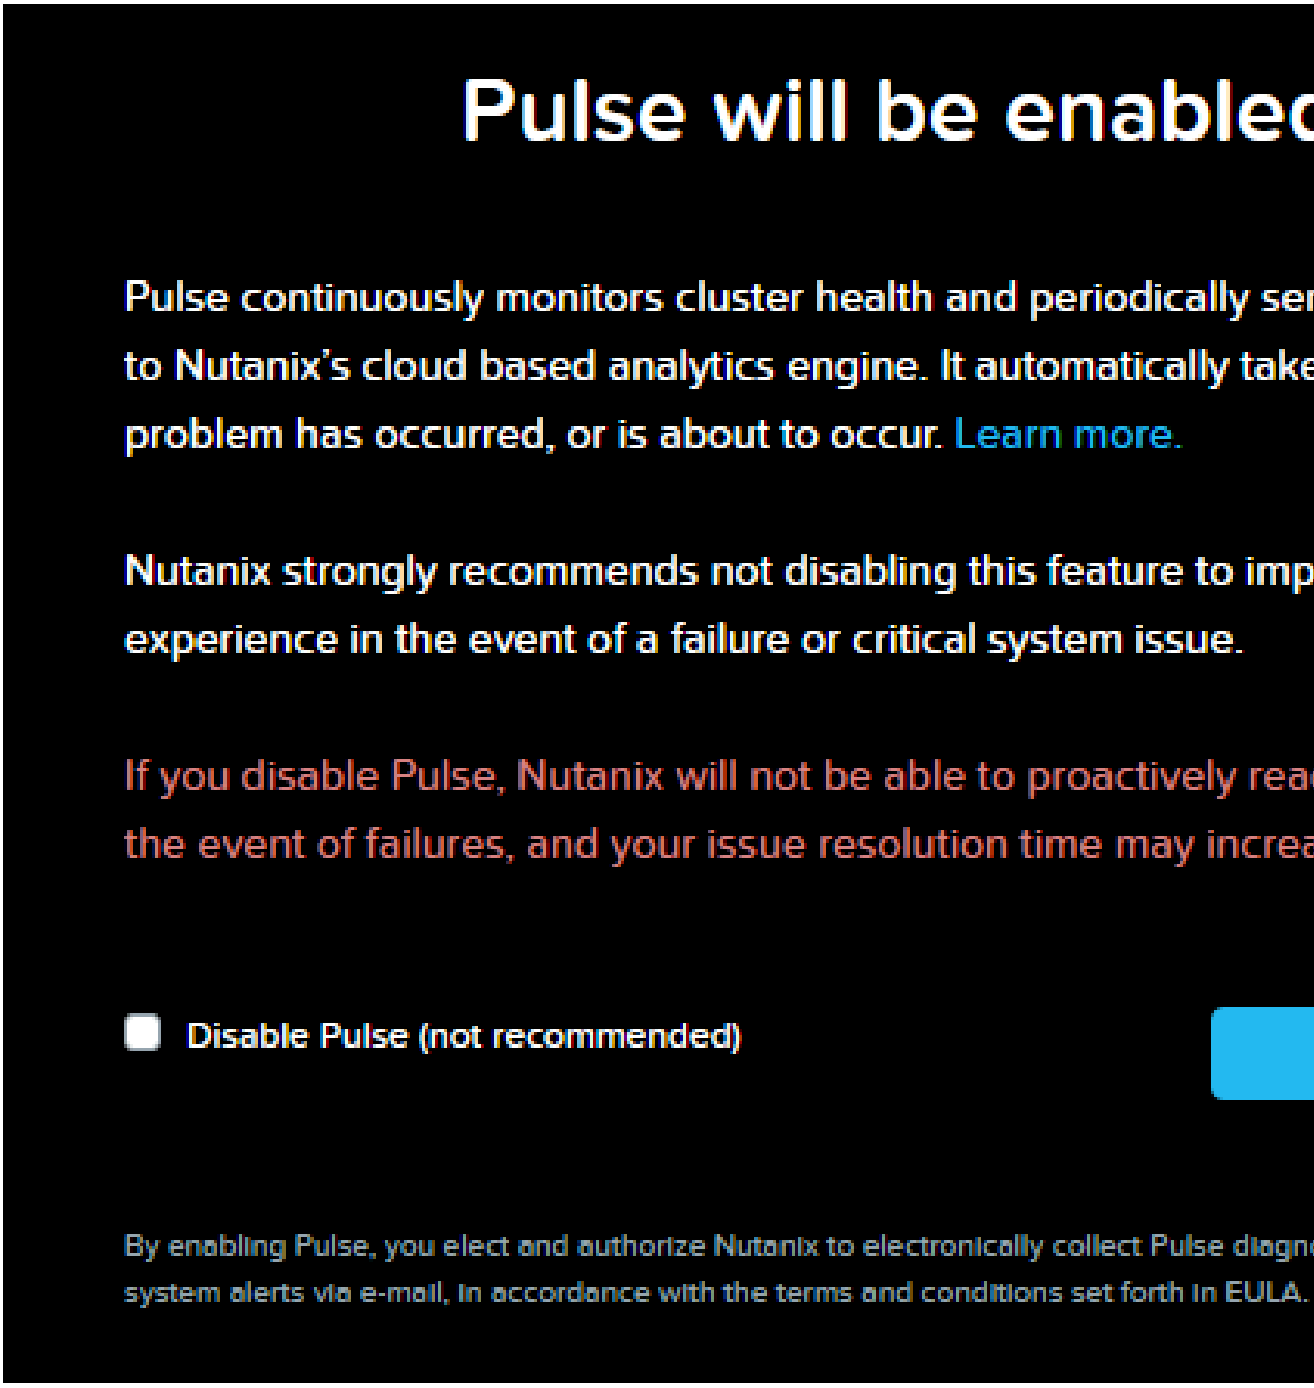

Figure 22: Pulse Screen

8. If a screen about enhanced cluster health monitoring appears (typically either after enabling Pulse in the previous step or after upgrading a cluster), read the statement and then do one of the following:

   » Click the **Yes (recommended)** button to enable enhanced cluster health monitoring.

   » Click the **Not Now** button to disable this feature.

   The enhanced (on top of standard Pulse) cluster health monitoring provides Nutanix customer support with more detailed (but more transparent) information that allows them to better monitor the health of your cluster. Enhanced cluster health monitoring is

recommended unless providing information such as entity names violates your security policy.

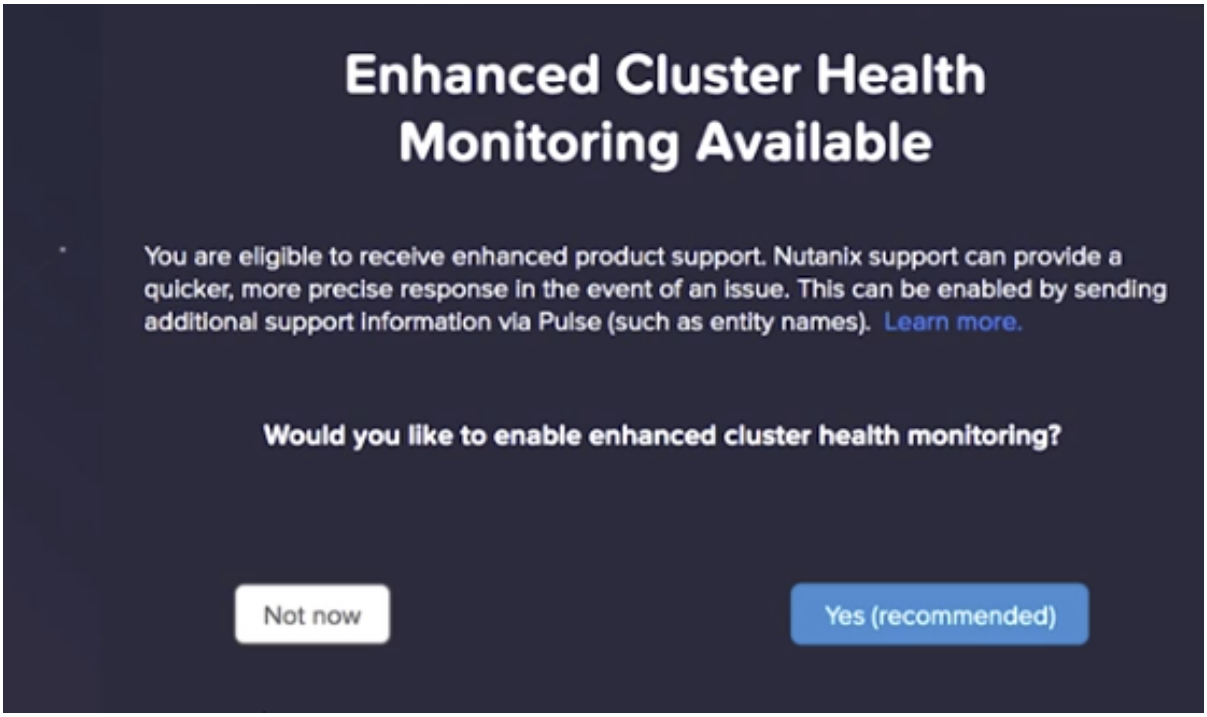

Figure 23: Enhanced Monitoring Screen (upgrade example)

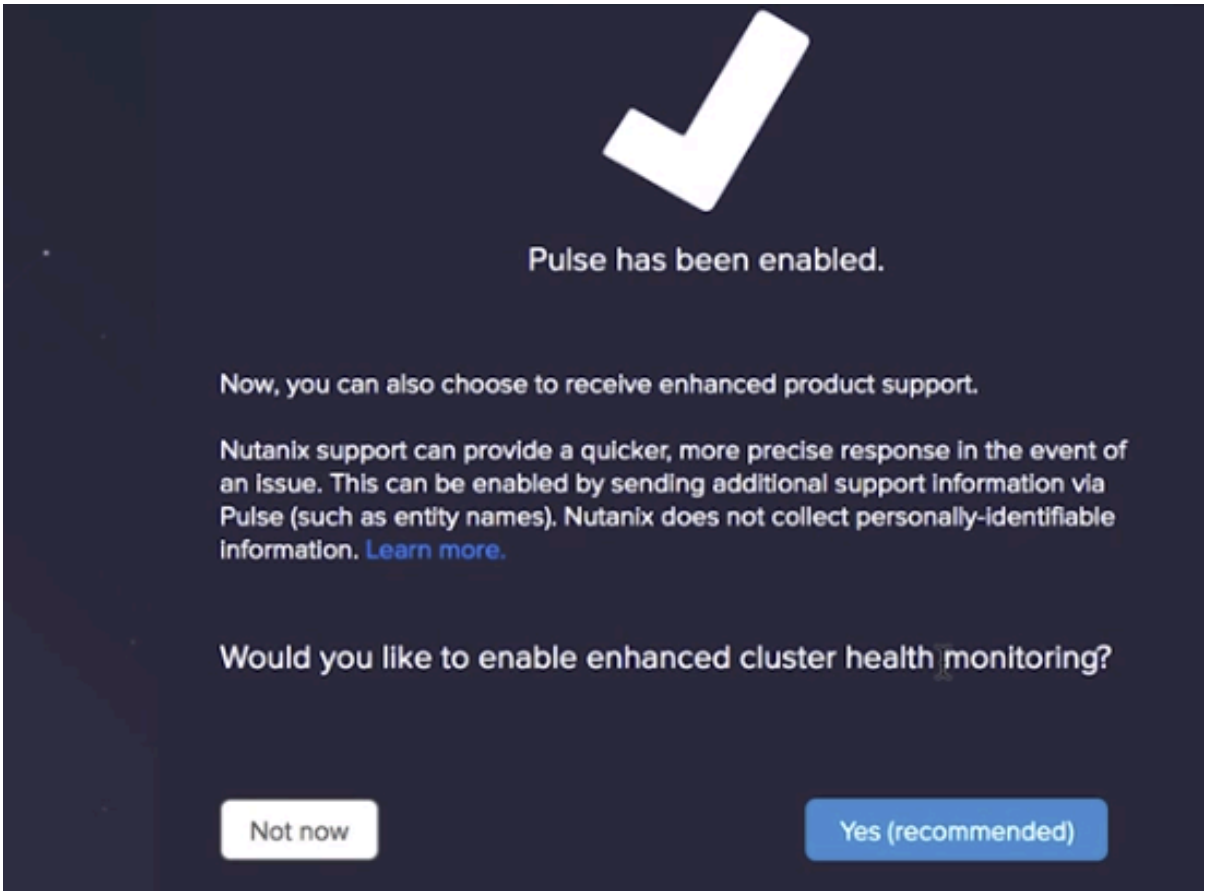

Figure 24: Enhanced Monitoring Screen Example (initial example)

9. [2-node clusters only] If a screen about registering with a Witness appears, see Configuring a Witness (Two-Node Cluster) on page 84.

Logging Out of the Web Console

Procedure

To log out from the web console, click the user icon 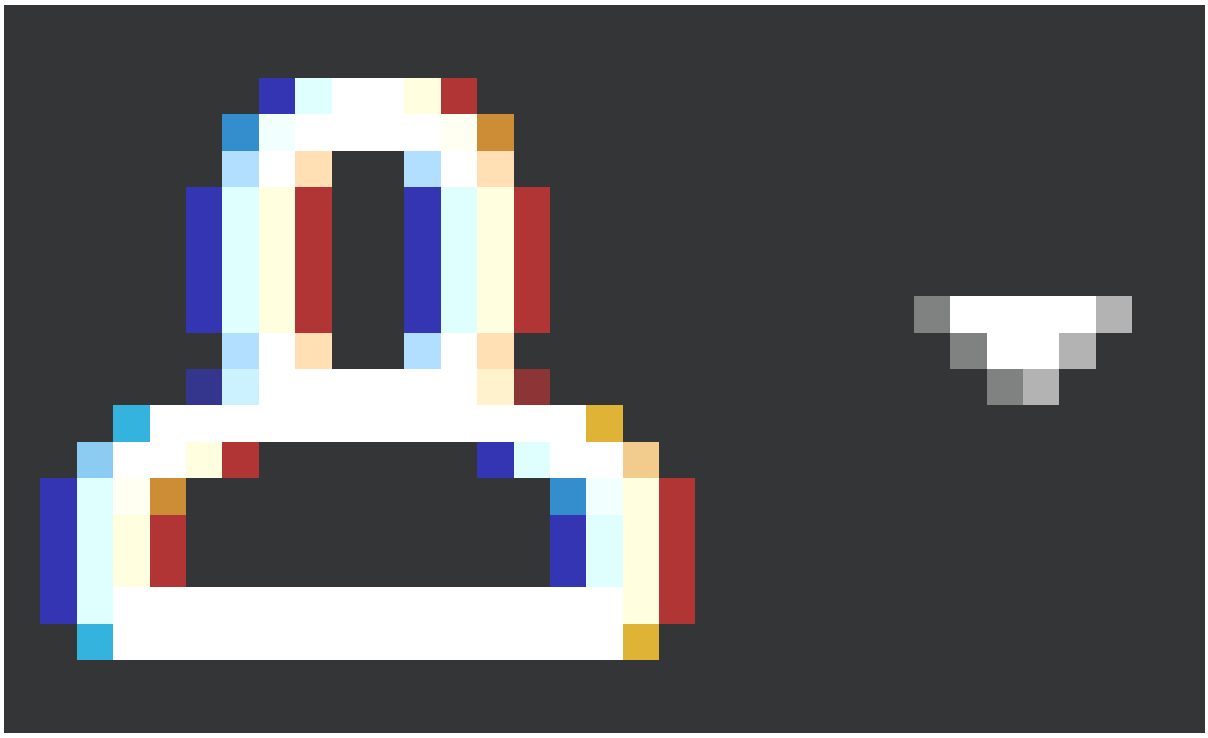 in the main menu and then select the **Sign Out** option from the pull-down list. You are logged out immediately after selecting the option (no prompt or message).

Main Menu

The main menu at the top of every screen provides access to all the features of the web console. This section describes each of the main menu options.

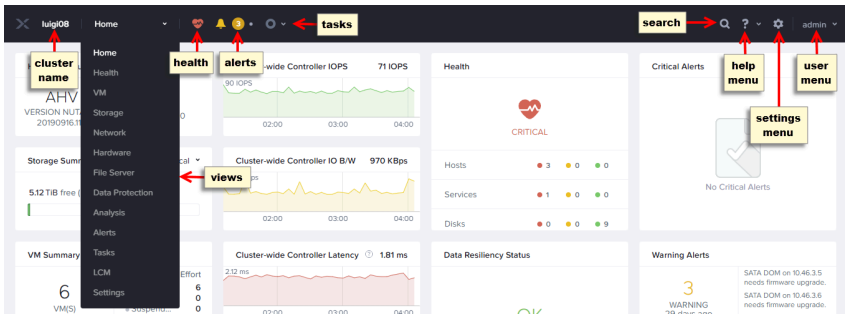

Figure 25: Main Menu with Expanded View List

Cluster Information

The cluster name appears on the far left of the main menu bar. Clicking the cluster name opens the Cluster Details window. This window displays the cluster ID number, cluster name, and cluster virtual IP address (if set). You can modify the name or virtual IP address at any time (see Modifying Cluster Details on page 58).

View Options

Selecting a view (entity name) from the pull-down list on the left displays information about that entity. Select from the following options:

- **Home**. Displays the main dashboard (see Home Dashboard).
- **Health**. Displays the health dashboard (see Health Dashboard on page 239).
- **VM**. Displays a virtual machine information dashboard (see VM Dashboard).
- **Storage**. Displays a storage information dashboard (see Storage Dashboard on page 103).
- **Network** (AHV only). Displays the network visualiser.
- **Hardware**. Displays a hardware information dashboard (see Hardware Dashboard on page 161).
- **File Server**. Displays a file server dashboard (see Nutanix Files Guide).
- **Data Protection**. Displays a data protection information dashboard (see Data Protection and Recovery with Prism Element).
- **Analysis**. Displays a screen to create and run performance monitors (see Analysis Dashboard on page 317).
- **Alerts**. Displays a screen of alert and events messages (see Alerts Dashboard).
- **Tasks**. Displays a screen of task messages (see Task Status on page 592).
- **LCM**. Displays the life cycle manager (LCM) dashboard (see Life Cycle Management (LCM) on page 62).
- **Settings**. Displays the Settings menu, as does clicking the gear icon on the right of the main menu (see Settings Menu on page 50).

Informational and Search Features

There are multiple ways to access information from the main menu:

- A health (heart) icon appears on the left of the main menu. It can be green (healthy), yellow (warning), or red (unhealthy) indicating the current heath status. Clicking the icon displays the heath details view (see Health Dashboard on page 239).
- An alerts (bell) icon appears on the left of the main menu when critical (red), warning (yellow), or informational (gray) alert messages have been generated and have not been marked as resolved. The number of unresolved alerts is displayed in the icon. Click the icon to display a drop-down list of the most recent unresolved alerts. Click a message or the right arrow link to open the alerts view (see Alerts Dashboard).
- A tasks (circle) icon appears to the right of the alerts icon. The number of active tasks (running or completed within the last 48 hours) is displayed in the icon. Click the icon to display a drop-down list of the active tasks, and then click the **View All Tasks** button to open the tasks view (see Task Status on page 592).
- A search (magnifying glass) icon appears on the right side of the main menu. Click this icon to display a search field. You can search for information about entities or actions by entering a string in this field. For example, you can enter an action name such as `add` that returns a list of add actions or an entity name such as `MyVM` that returns a list of links related to that entity.

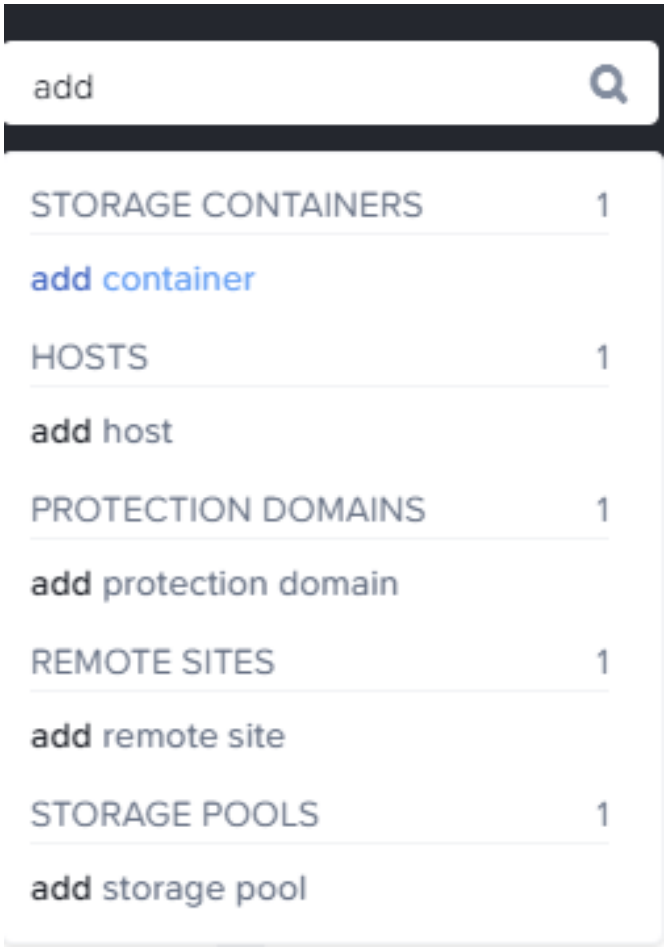

Figure 26: Search Box Example

Help Menu

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 list.

Table 5: User Menu List

Name	Description
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help on page 656).
Health Tutorial	Opens the Health dashboard tutorial that takes you through a guided tour of the health analysis features (see Health Dashboard on page 239).
General Help	Opens the online help at the introduction page.
Support Portal	Opens a new browser tab (or window) at the Nutanix support portal login page (see Accessing the Nutanix Support Portal on page 648).

Name	Description
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 660. This is an online community site for customers and partners to exchange ideas, tips, and information about Nutanix technologies and related datacenter topics.

User Menu

A user icon appears on the far right side of the main menu with the current user login name. Clicking the user icon displays a list of options to update your user account, log out from the web console, and other miscellaneous tasks. The following table describes each option in the pull-down list.

Table 6: User Menu List

Name	Description
Update Profile	Opens the **Update Profile** window to update your user name and email address (see “” ()).
Change Password	Opens the **Change Password** window to update your password (see “” ()).
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 652).
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.
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 Prism Central	Opens a new browser tab (or window) at the Support Tools page of the Nutanix support portal from which you can download the files to install Prism Central. If a login page appears, enter your Nutanix support portal credentials to access the portal. See the *Prism Central Guide* for more information.
About Nutanix	Opens a window that displays Nutanix operating system (AOS) and other version information (see Checking Version on page 61).
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 out the user from the web console (see Logging Out of the Web Console on page 46).

Settings Menu

The web console 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 on the right of the main menu (see Main Menu on page 46).
- Select **Settings** from the pull-down list on the left of the main menu.

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.

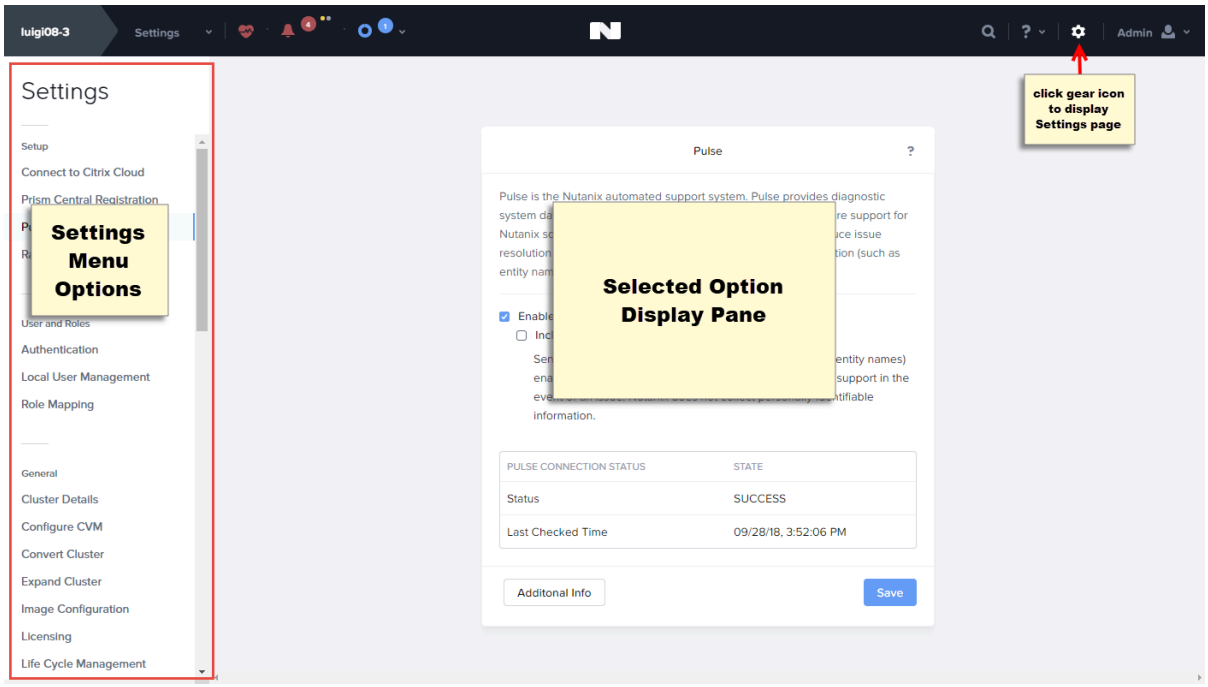

Figure 27: Settings Page

Table 7: Settings Menu List

Name	Description
General	
Cluster Details	Opens the **Cluster Details** window to view or modify certain cluster parameters (see Modifying Cluster Details on page 58).
Configure CVM	Opens the **Configure CVM** window to increase the Controller VM memory size (see Increasing the Controller VM Memory Size on page 93).
Convert Cluster	Opens the **Convert Cluster** window to convert the cluster from ESXi to AHV and then from AHV to ESXi (see In-Place Hypervisor Conversion on page 620).

Name	Description
Expand Cluster	Opens the **Expand Cluster** window to add new nodes to the cluster (see Expanding a Cluster on page 185).
Image Configuration [AHV only]	Opens the **Image Configuration** window to import and mange image files that can be used to create VMs (see Configuring Images on page 289).
Licensing	Opens the **Licensing** window to install or update the cluster license that enables entitled Nutanix features (see Licensing on page 61).
Reboot [AHV and XenServer only]	Opens the **Request Reboot** window to gracefully restart the nodes in the cluster one after the other. You can select the nodes you want to restart (see Restarting the Nodes in the Cluster (AHV and XenServer)).
Remote Support	Opens the **Remote Support Services** window, which enables (or disables) Nutanix remote support access (see Controlling Remote Connections on page 645).
Upgrade Software	Opens the **Upgrade Software** window to upgrade the cluster to a newer AOS version, or update other upgradeable components (see Software and Firmware Upgrades on page 61).
vCenter Registration [ESXi only]	Opens the **vCenter Registration** window to register (or unregister) the cluster with the vCenter instance (see Registering a vCenter Server on page 617).
Setup	
Connect to Citrix Cloud [AHV and XenServer only]	Opens the **Connect to Citrix Cloud** window to connect to the Citrix Workspace Cloud (see Connect to Citrix Cloud on page 310).
Prism Central Registration	Opens the **Prism Central Registration** window to add the cluster into a central registration for multicluster connection and support (see Register (Unregister) Cluster with Prism Central on page 62).
Pulse	Opens the **Pulse** window to enable the sending of cluster information to Nutanix customer support for analysis (see Configuring Pulse on page 642).
Rack Configuration	Opens the **Rack Configuration** page to configure the fault tolerant domain for node, block, and rack awareness (see Rack Fault Tolerance on page 17).
Network	
HTTP Proxy	Opens the **HTTP Proxy** window to configure an HTTP proxy to which the Nutanix software can connect (see Configuring HTTP Proxy on page 646).
Name Servers	Opens the **Name Servers** window to configure name servers for the cluster (see Configuring Name Servers on page 595).
Network Configuration [AHV only]	Opens the **Network Configuration** window to configure network connections for the cluster (see Network Configuration for VM Interfaces on page 138).

Name	Description
Network Switch	Opens the **Network Switch Configuration** window to configure network switch information needed for collecting network traffic statistics (see Configuring Network Switch Information on page 144). This option does not appear when running a hypervisor that does not support this feature.
NTP Servers	Opens the **NTP Servers** window to specify which NTP servers to access (see Configuring NTP Servers on page 597).
SNMP	Opens the **SNMP Configuration** window to enable and configure SNMP for the cluster (see Configuring SNMP on page 599).
Security	
Cluster Lockdown	Opens the **Cluster Lockdown** window, which allows you to delete (or add) public authentication keys used for SSH access into the cluster (see “” ()). Removing all public keys locks down the cluster from external access.
Data at Rest Encryption [SEDs only]	Opens the **Data-at-Rest Encryption** screen to configure key management for self encrypting drives (SEDs) and enable data encryption across the cluster (see “” ()). This menu option appears only when the data drives in the cluster are SEDs.
Filesystem Whitelists	Opens the **Filesystem Whitelist** window to specify whitelist addresses (see Configuring a Filesystem Whitelist on page 594).
SSL Certificate	Opens the **SSL Certificates** window to create a self-signed certificate (see “” ()).
Users and Roles	
Authentication	Opens the **Authentication Configuration** window to configure authentication for the cluster (see “” ()).
Local User Management	Opens the **Local User Management** window. This window lists current users and allows you to add, update, and delete user accounts. See “” () .
Role Mapping	Opens the **Role Mapping** window to configure role mappings that apply in the user authentication process (see “” ()).
Email and Alerts	
Alert Email Configuration	Opens the **Alert Email Configuration** window, which enables (or disables) the e-mailing of alerts (see Configuring Alert Emails on page 335).
Alert Policies	Opens the **Alert Policies** window, which allows you to specify what events should generate an alert and how frequently the system should check for each event type (see Configuring Alert Policies on page 338).
SMTP Server	Opens the **SMTP Server Settings** window to configure an SMTP server (see Configuring an SMTP Server on page 598).
Data Resiliency	

Name	Description
Configure Witness [ESXi and AHV only]	Opens the **Configure Witness** window to add a witness VM for Metro Availability and two-node clusters (see Configuring a Witness (Two-Node Cluster) on page 84).
Degraded Node Settings	Opens the **Degraded Node Settings** window to enable or disable Degraded Node Detection (see Degraded Node on page 30).
Manage VM High Availability [AHV only]	Opens the **Manage VM High Availability** window to enable high availability for guest VMs in the cluster (see Enabling High Availability for the Cluster on page 296).
Redundancy State	Opens the **Redundancy Factor Readiness** window to configure the redundancy factor of the cluster (see Increasing the Cluster Fault Tolerance Level on page 88).
Appearance	
Language Settings	Opens the **Languages** window, which allows you to select the language of the web console (see Localization (L10n) on page 629).
UI Settings	Opens the **UI Settings** window to disable (or re-enable) the login screen (see Modifying UI Settings on page 60).
Welcome Banner	Opens the **Edit Welcome Banner** window to create a welcome banner message that appears before users login to the web console (see Configuring a Banner Page on page 615).

Home Dashboard

The Home dashboard is the opening screen that appears after logging into the web console. It provides a dynamic summary view of the cluster status. To view the Home dashboard at any time, select **Home** from the pull-down list on the far left of the main menu.

Menu Options

The Home dashboard does not include menu options other than those available from the main menu (see Main Menu on page 46).

Home Screen Details

The Home dashboard displays cluster-level performance and usage statistics on the left, heath status information in the middle, and the most recent alert and event messages on the right. The following figure is a sample dashboard, and the following table describes each field in this screen. Several fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current.

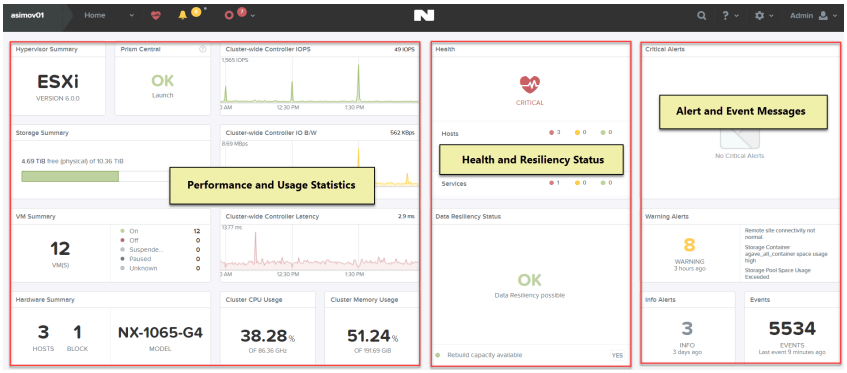

Figure 28: Home Dashboard

Note: See Understanding Displayed Statistics on page 56 for information about how the statistics are derived.

Table 8: Home Screen Fields

Name	Description
Hypervisor Summary	Displays the name and version number of the hypervisor.
Prism Central	Displays whether you have registered the cluster to a Prism Central instance or not. Click **Register** to register the cluster to a Prism Central instance. If you have already registered, you can click **OK** to launch the Prism Central instance in a new tab of your browser.
Storage Summary	Displays the total, used, and unused storage space in the cluster.
VM Summary	Displays the total number of VMs in the cluster broken down by on, off, and suspended states.
Hardware Summary	Displays the number of hosts and blocks in the cluster, plus one or more Nutanix block model numbers.
Cluster-wide Controller IOPS	Displays I/O operations per second (IOPS) in the cluster. The displayed time period is a rolling interval 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. (These display features also apply to the I/O bandwidth and I/O latency monitors.) For more in depth analysis, you can add this chart (and any other charts on the page) to the analysis page by clicking the blue link in the upper right of the chart (see Analysis Dashboard on page 317).
Cluster-wide Controller IO Bandwidth	Displays I/O bandwidth used per second in the cluster. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume.
Cluster-wide Controller Latency	Displays the average I/O latency (in milliseconds) in the cluster.
Cluster CPU Usage	Displays the current CPU utilization percentage along with the total available capacity (in GHz).

Name	Description
Cluster Memory Usage	Displays the current memory utilization percentage along with the total available capacity (in GB).
Health	Displays health status for the cluster as a whole (good, warning, critical) and summary health status for the VMs, hosts, and disks in the cluster. Clicking on the VMs, hosts, or disks line displays detailed information about that object in the Health page (see Health Dashboard on page 239).
Data Resiliency Status	Displays information indicating whether the cluster is protected currently from potential data loss due to a component failure. Click anywhere in this field to display a dialog box. Data Resiliency Status window with more information (see the following "Data Resiliency Status" section). • **Resiliency Status**. Indicates whether the cluster can safely handle a node failure, that is whether a copy exists somewhere in the cluster of all data in any node. If the status is not **OK**, the Data Resiliency Status window includes a message about the problem. Note: The resiliency status for single-node backup cluster is at the disk level and not at the node level. For more information, see Single-Node Replication Target Clusters in the guide titled *Data Protection and Recovery with Prism Element*. • **Rebuild Capacity Available**. Indicates whether there is sufficient unused storage in the cluster to rebuild a data copy after a node is lost. If the status is not **Yes**, the Data Resiliency Status window includes a message about the problem. Note: This option does not appear for single-node replication target clusters.
Critical Alerts	Displays the most recent unresolved critical alert messages. Click a message to open the Alert screen at that message (see Alerts Dashboard).
Warning Alerts	Displays the most recent unresolved warning alert messages. Click a message to open the Alert screen at that message.
Info Alerts	Displays a summary of informational alerts.
Events	Displays a summary of events.

Data Resiliency Status

The Data Resiliency Status window displays more detailed cluster resiliency status information. (Click the **OK** button at the bottom or **X** icon at the top to close the window.) This window provides information about the number and type of failures the cluster can withstand safely. The following figure is an example. The **Failures Tolerable** column indicates the number of simultaneous failures of that component type that can be tolerated (0, 1, or 2). When no failure can be tolerated, a message is displayed to highlight the limitation, as in this example where there are not enough blocks in the cluster to support a block failure.

Note: When a node goes down, extent group (egroup) fault tolerance status remains unchanged as the node is assumed (initially) to be unavailable just temporarily. However, Stargate (see

Cluster Components on page 12) fault tolerance goes down by one until all data has been migrated off that node. The egroup fault tolerance status goes down only when a physical copy of the egroup replica is permanently bad.

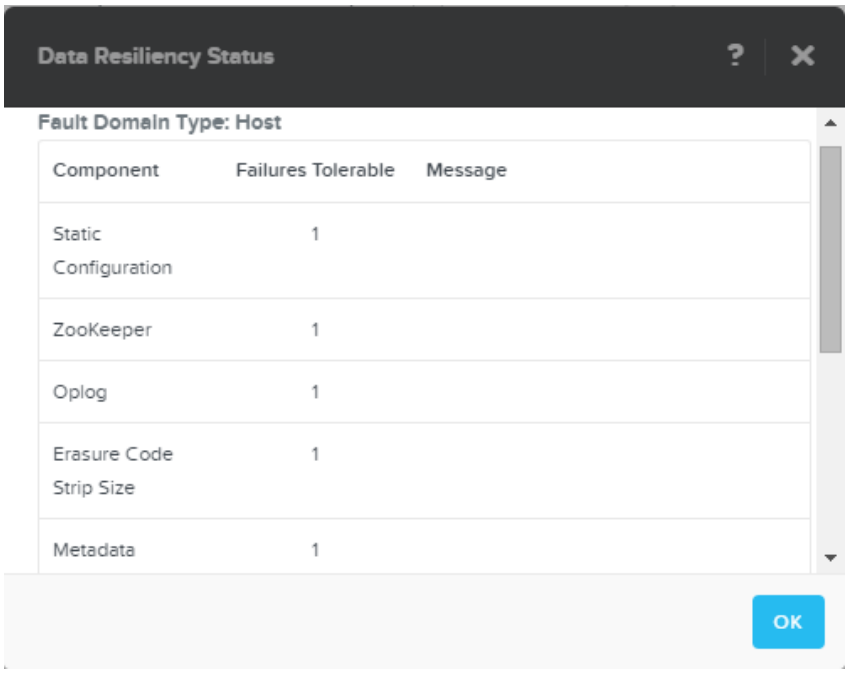

Figure 29: Data Resiliency Status Window

Fault Domain Type: Block

Component	Failures Tolerable	Message
Static Configuration	0	Not enough rackable units (blocks) in the cluster

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

Modifying Cluster Details

About this task

You can add a cluster name or virtual facing IP address at any time. To create or modify either value, do the following:

Procedure

1. In the main menu, either click the cluster name at the far left or click the gear icon in the main menu and then select **Cluster Details** in the **Settings** page.
   The Cluster Details window appears. It displays the cluster UUID (universally unique identifier), ID, and incarnation ID values, along with the cluster name, virtual IP address, and iSCSI data services IP address (if configured). The cluster ID remains the same for the life of the cluster, but the incarnation ID is reset (typically to the wall time) each time the cluster is re-initialized.

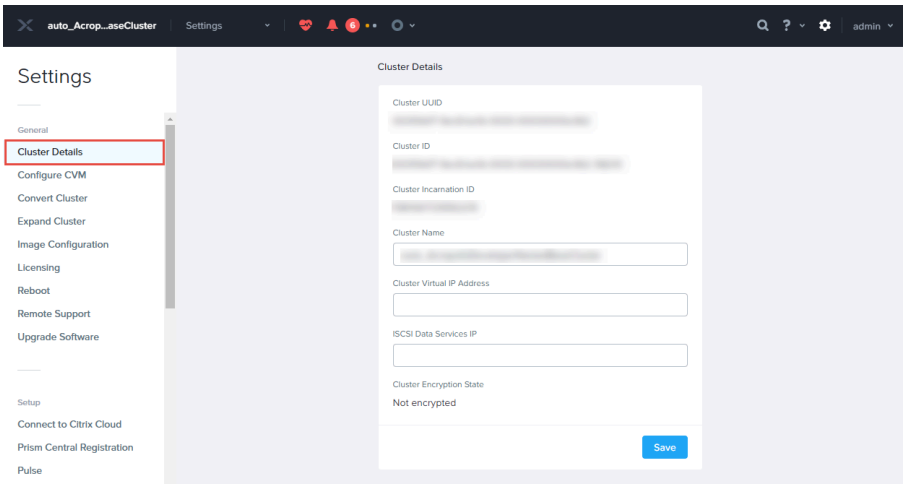

Figure 30: Cluster Details Window

2. Enter (or update) a name for the cluster in the **Cluster Name** field.
   The default name is simply Unnamed. Providing a custom name is optional but recommended.

3. Enter (or update) an IP address for the cluster in the **Cluster Virtual IP Address** field.
   A Controller VM runs on each node and has its own IP address, but this field sets a logical IP address that always connects to one of the active Controller VM in the cluster (assuming at least one is active), which removes the need to enter the address of a specific Controller VM. The virtual IP address is normally set during the initial cluster configuration (see the *Field Installation Guide*), but you can update the address at any time through this field. Providing an virtual IP address is optional but recommended.

   CAUTION: All the features that use the cluster virtual IP address will be impacted if you change that address. See Virtual IP Address Impact on page 59 for more information.

4. Enter (or update) an IP address for use with Nutanix Volumes in the **ISCSI Data Services IP Address** field.

   CAUTION: For certain features, changing the external data services IP address can result in unavailable storage or other issues. The features in question include Volumes, Calm, Leap,

Karbon, Objects, and Files. See KB 8216 and iSCSI Data Services IP Address Impact on page 59 for more information.

5. When the field entries are correct, click the **Save** button to save the values and close the window.

Virtual IP Address Impact

Any Nutanix feature that uses the virtual IP address might be adversely affected if you change the cluster virtual IP address.

- You can no longer manage a Nutanix cluster (Hyper-V) through the System Center Virtual Machine Manager (SCVMM).
- Nutanix data protection service might lose connection if you configured the remote site using the virtual IP address of that cluster.
- All the VMs running the NGT (Nutanix guest tools) instance will be affected. See Reconfiguring NGT on page 307.
- External machines mounting shares from Nutanix might fail if the virtual IP address is used for HA functionality (as recommended).
- Some products such as Nutanix Objects use the virtual IP address for access. Check your product-specific documentation before changing the virtual IP address.

iSCSI Data Services IP Address Impact

To provide access to cluster storage, Nutanix Volumes utilizes an iSCSI data services IP address to clients for target discovery which simplifies external iSCSI configuration on clients. This iSCSI data services IP address acts as an iSCSI target discovery portal and initial connection point.

Nutanix does not recommend configuring iSCSI client sessions to connect directly to Controller VM IP addresses. The data services address is owned by one Controller VM at a time. If the owner Controller VM becomes unavailable, the address moves to another Controller VM, ensuring that it is always available.

This IP address is also used as a cluster-wide address by clients configured as part of Nutanix Files and other products. This IP address:

- Should be in the same subnet as the cluster Controller VM IP `eth0` network interface addresses
- Helps load balance storage requests
- Enables path optimization in the cluster, preventing bottlenecks
- Eliminates the need for configuring a multipathing service such as Microsoft multipath I/O (MPIO)

Nutanix recommends setting the iSCSI data services IP address once for each cluster, but you can change it if needed through the Cluster Details window in the Prism web console. See Modifying Cluster Details on page 58.

If you change the iSCSI data services IP address, you will need to reconfigure any clients to use the new IP address.

- Log out of or disconnect from existing iSCSI or Nutanix Files sessions.
- Delete the old IP address if necessary.
- Re-discover the new target or reestablish Nutanix Files sessions.

CAUTION: For certain features, changing the external data services IP address can result in unavailable storage or other issues. The features in question include Volumes, Calm, Leap, Karbon, Objects, and Files. See KB 8216 for more information.

Modifying UI Settings

About this task

By default, the login page includes background animation, and users are logged out automatically after being idle for 15 minutes. You can disable the background animation, change the session timeout for users, and configure an override to the session timeout.

Procedure

1. Click the gear icon in the main menu and then select **UI Settings** in the **Settings** page. The UI Settings dialog box appears.

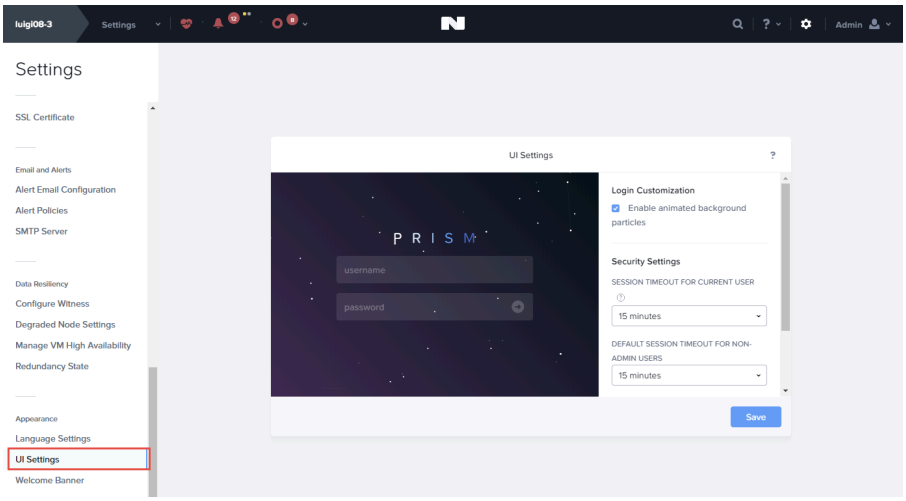

Figure 31: UI Settings Window

2. To disable the login page background animation, under **Login Customization**, uncheck the **Enable animated background particles** box (or check it to enable).
3. 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 users (except an administrator) 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.

     Note: The timeout interval for an administrator cannot be set for longer than 1 hour.
4. Uncheck the **Disable 2048 game** box to disable the 2048 game.
5. Click **Save** to save your changes.

Checking Version

Procedure

1. To view the Nutanix version running in the cluster, 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 and Nutanix cluster check (NCC) version numbers. It also includes a link to Nutanix patent information.

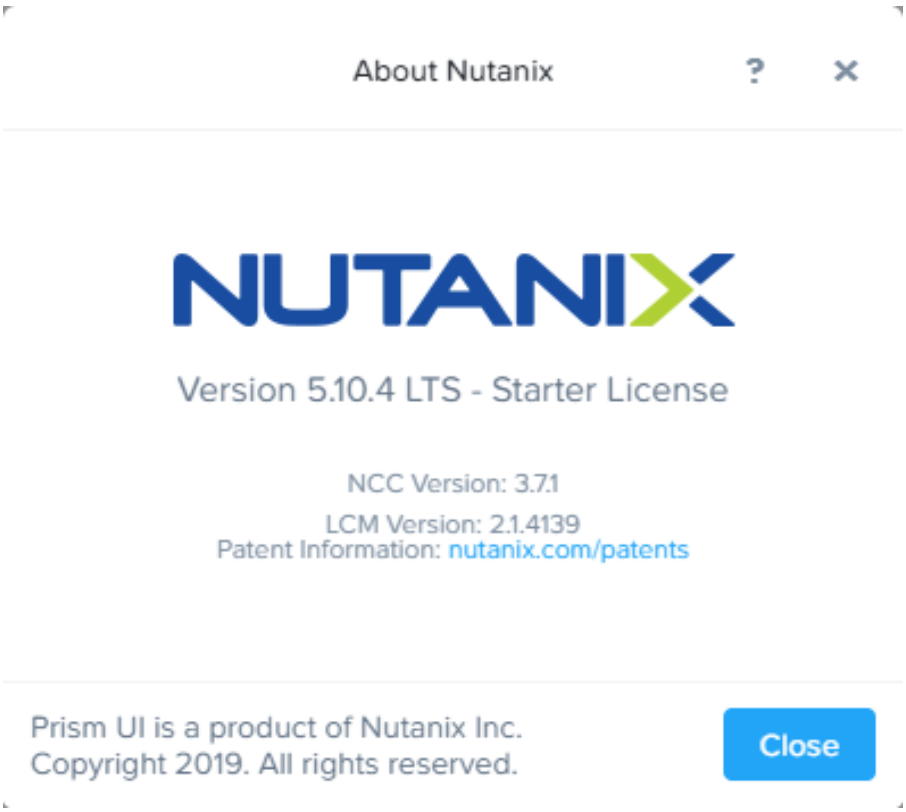

Figure 32: About Nutanix Window

2. Click the **OK** button to close the window.

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the Nutanix Licensing Guide for information about how to license your cluster, apply add-on licenses, and manage your licenses in a variety of situations.

Software and Firmware Upgrades

You can upgrade the AOS, hypervisor, NCC, or Foundation version through the Prism web console **Upgrade Software** feature (also known as 1-click upgrade).

Before you attempt a software upgrade, review the “” () for Nutanix cluster and IT administrators. This guide describes how to upgrade your Prism Element or Prism Central cluster through the web console. It includes upgrade best practices and upgrade work flows as recommended by Nutanix, including the recommended order of installation.

Life Cycle Management (LCM)

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. To view the LCM dashboard, select **LCM** from the pull-down list on the left of the main menu. See the Life Cycle Manager Guide for more information.

Multi-Cluster Management

Nutanix provides an option to monitor and manage multiple clusters through a single web console. The multi-cluster view, known as Prism Central, is a centralized management tool that runs as a separate VM. From the Prism web console, you can either register the cluster with an existing Prism Central instance or deploy a Prism Central instance and register the cluster with it.

Prism Central provides the following features:

- Single sign on for all registered clusters
- Summary dashboard across clusters that can be customized as desired
- Summary views for major entity types with drill-down options to view detailed information about individual entities
- Multi-cluster analytics capabilities
- Multi-cluster alerts summary with drill-down options to view information about possible causes and corrective actions for each alert
- Ability to configure individual clusters through direct Prism Central actions (for selected administrative tasks) or through one-click access to the web console for a cluster

> Note: See the *Prism Central Guide* for more information about Prism Central.

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 (see Logging Into 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.
- 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. Run NCC as described in Run NCC 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.

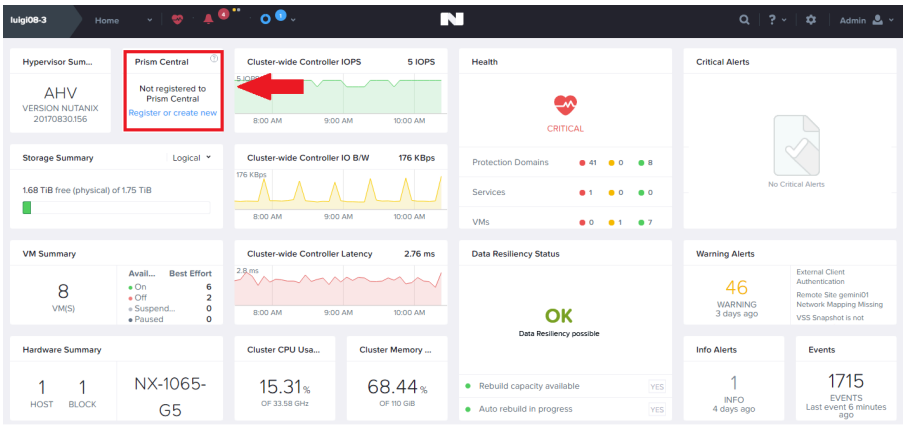

Figure 33: Home Dashboard

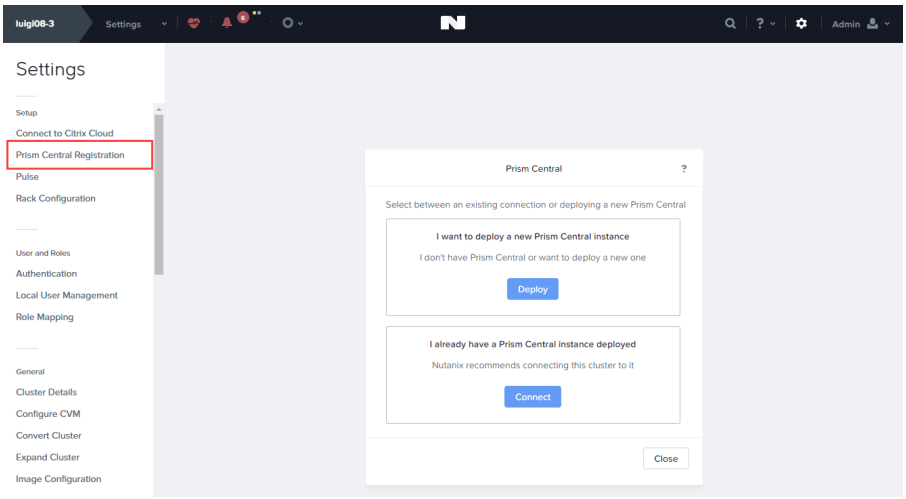

Figure 34: Settings Menu

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 “” () for instructions on how to deploy a new Prism Central instance.

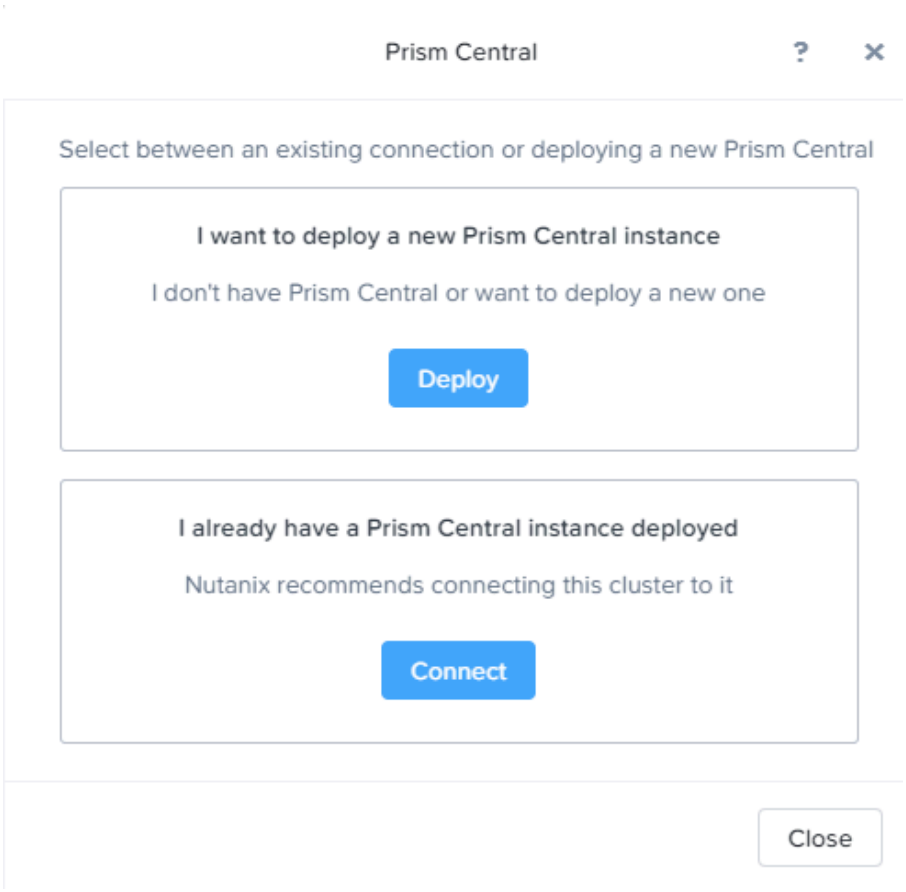

Figure 35: Prism Central Dialog Box (deploy/connect screen)

5. A screen appears outlining what services are available through Prism Element and Prism Central. After reviewing the message, click the **Next** button.

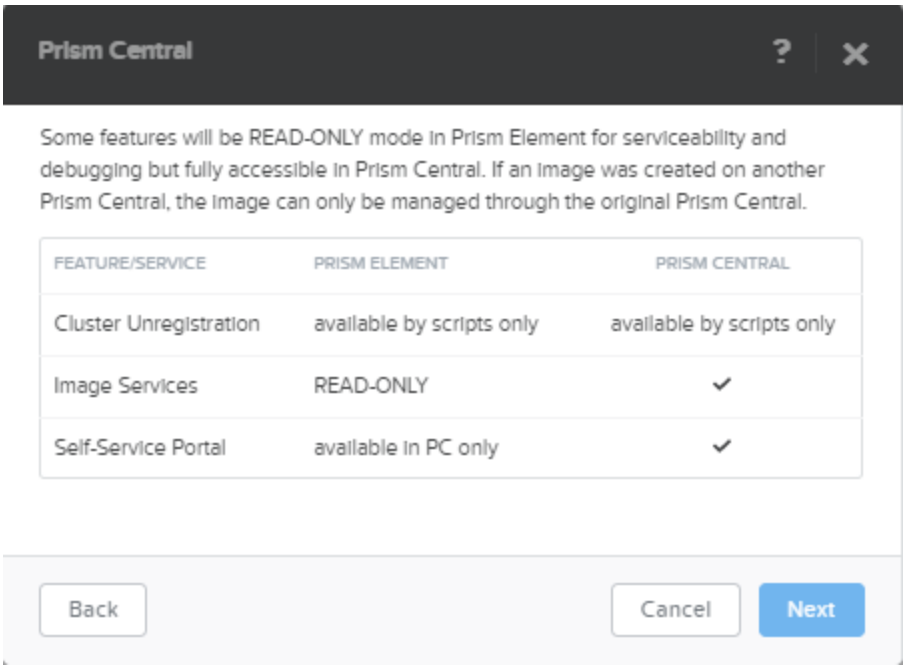

Figure 36: Prism Central Dialog Box (services screen)

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. **Username**: Enter `admin` as the Prism Central user name.
   c. **Password**: Enter the Prism Central `admin` user password.

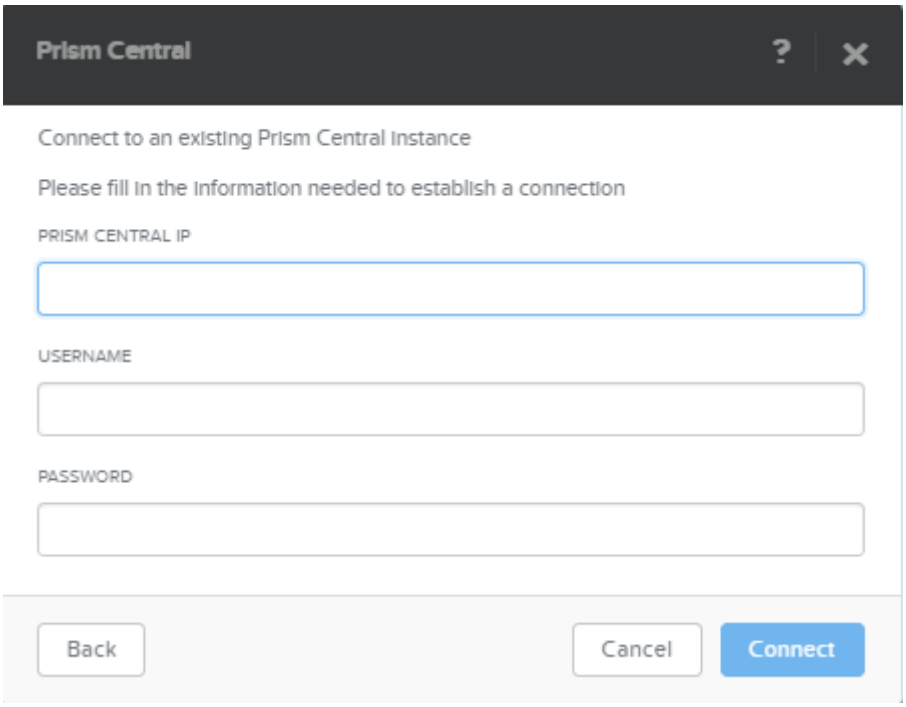

Figure 37: Prism Central Dialog Box (registration screen)

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: If the Prism Central IP address changes after registering, you must unregister and then re-register the cluster(s) with the new IP address. However, changing the Prism Central administrator password after registering any clusters does not require re-registration.

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.

```
<acropolis> exit
nutanix@NTNX-[redacted]-A-CVM:10.5.25.30:~$
nutanix@NTNX-[redacted]-A-CVM:10.5.25.30:~$ cluster status
2017-11-19 21:31:20 INFO zookeeper_session.py:110 cluster is attempting to connect to Zookeeper
2017-11-19 21:31:20 INFO cluster:2420 Executing action status on SVMs 10.5.25.30
The state of the cluster: start
Lockdown mode: Disabled

        CVM: 10.5.25.30 Up, ZeusLeader
                                Zeus   UP       [1966, 1995, 1996, 1997, 2043, 2061]
                           Scavenger   UP       [2708, 2736, 2737, 2738]
                      SSLTerminator    UP       [2104, 2136, 2137, 2138]
                     SecureFileSync    UP       [2108, 2162, 2163, 2164]
                             Medusa    UP       [2278, 2312, 2313, 2318, 2421]
                  DynamicRingChanger   UP       [2520, 2660, 2661, 2665]
                             Pithos    UP       [2524, 2586, 2587, 2615]
                             Mantle    UP       [2528, 2590, 2591, 2620]
                               Hera    UP       [2533, 2608, 2609, 2610]
                            Stargate   UP       [2823, 2893, 2894, 3250, 3258]
                          InsightsDB   UP       [2826, 2876, 2877, 3027]
                InsightsDataTransfer   UP       [2834, 2916, 2917, 3001, 3004, 3005, 3006]
                               Ergon   UP       [2842, 2963, 2964, 2965]
                             Cerebro   UP       [2901, 2984, 2985, 3448]
                             Chronos   UP       [2920, 2999, 3000, 3074]
                             Curator   UP       [2942, 3047, 3048, 3506]
                               Prism   UP       [3024, 3094, 3095, 3201, 3707, 3776]
                                 CIM   UP       [3037, 3154, 3155, 3214]
                        AlertManager   UP       [3046, 3283, 3286, 3422]
                            Arithmos   UP       [3083, 3215, 3216, 3488]
                             Catalog   UP       [3106, 3189, 3190, 3191]
                           Acropolis   UP       [3173, 3294, 3295, 3296]
                               Uhura   UP       [3192, 3389, 3390, 3396]
                                Snmp   UP       [3198, 3391, 3392, 3395]
                   SysStatCollector    UP       [3208, 3445, 3446, 3447]
                              Tunnel   UP       [3231, 3582, 3583]
                               Janus   UP       [3304, 3485, 3486]
                  NutanixGuestTools    UP       [3517, 3617, 3618, 3683]
                          MinervaCVM   UP       [4029, 4064, 4065, 4066, 4213]
```

Figure 38: Cluster Status Command

3. Unregister the cluster from Prism Central.

   ```
   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.

   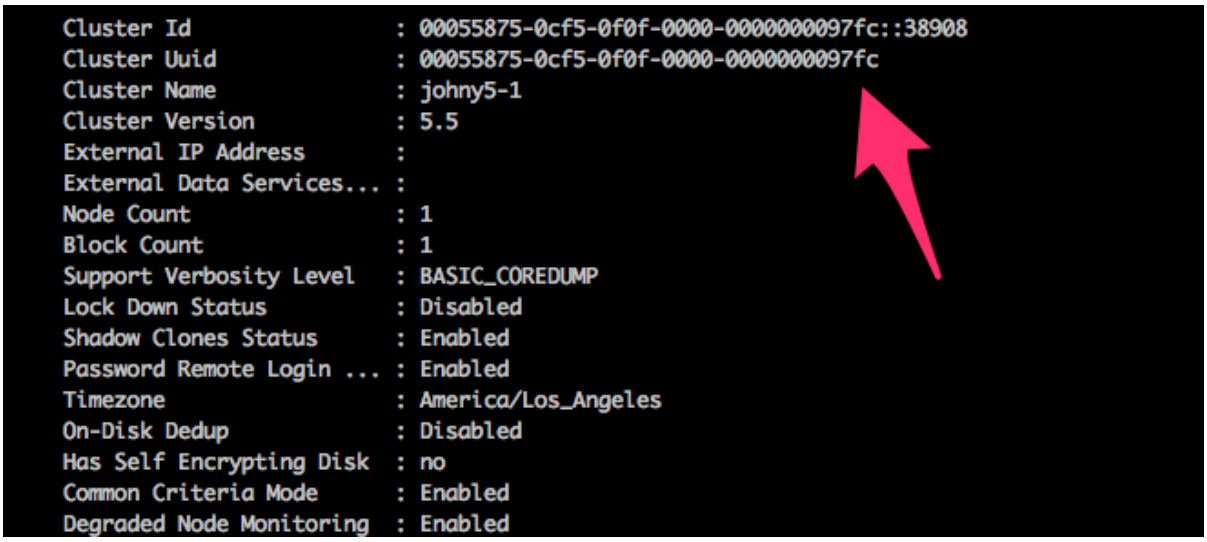

   Figure 39: Cluster Info Output

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 5. 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 5), 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 6c.

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.

Installing Prism Central (1-Click Internet)

Before you begin

The following requirements must be met before you can install Prism Central (this procedure), expand Prism Central, or add a Prism Central VM:

- The specified gateway must be reachable.
- No duplicate IP addresses can be used.
- The storage container used for deployment is mounted on all 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

To create a new Prism Central instance from an Internet-connected cluster, do the following:

Procedure

1. Log into the Prism web console for a cluster as the user `admin`.
2. Run NCC as described in Run NCC 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 in the main menu and then select **Prism Central Registration** from the **Settings** menu.

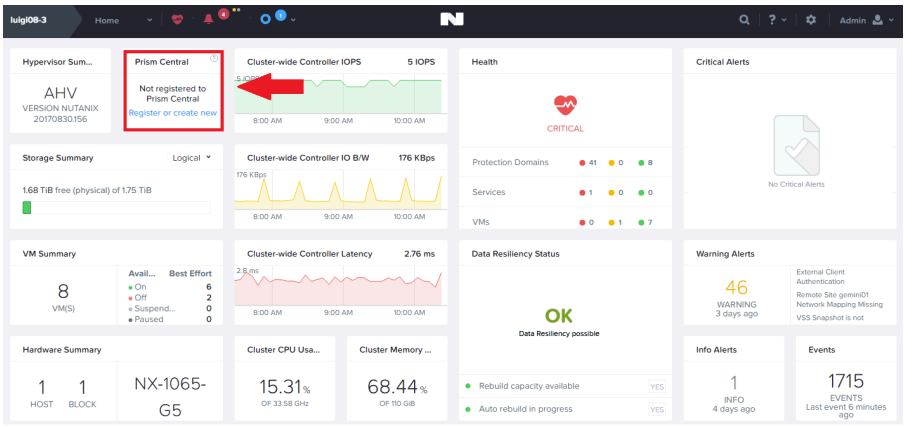

Figure 40: Home Dashboard

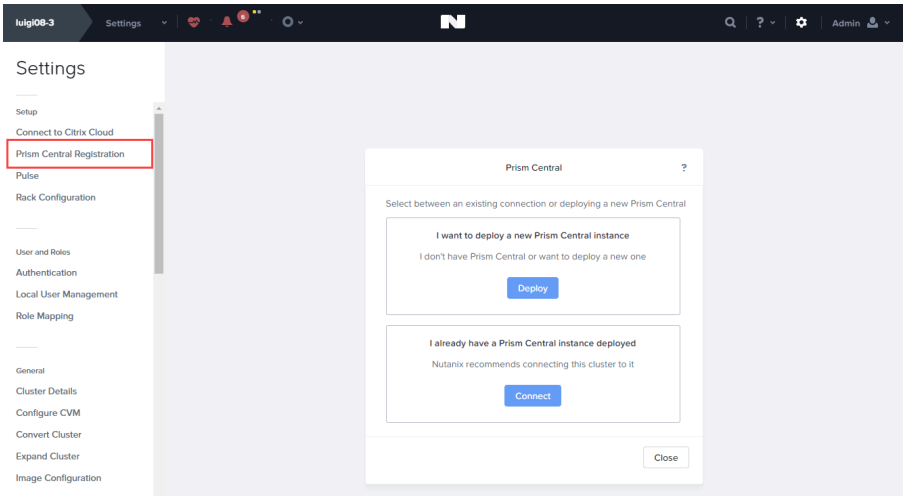

Figure 41: Settings Menu Launched from Gear Icon

4. In the first screen of the Prism Central dialog box, click the **Deploy** button.

   This screen includes two options, **Deploy** and **Connect**. This procedure describes how to deploy a new Prism Central instance. The topic *Register (Unregister) Cluster with Prism*

Central in the *Prism Central Guide* describes how to connect to an existing Prism Central instance.

> Note: On an ESXi cluster, you must first register a vCenter Server before you deploy a new Prism Central instance.

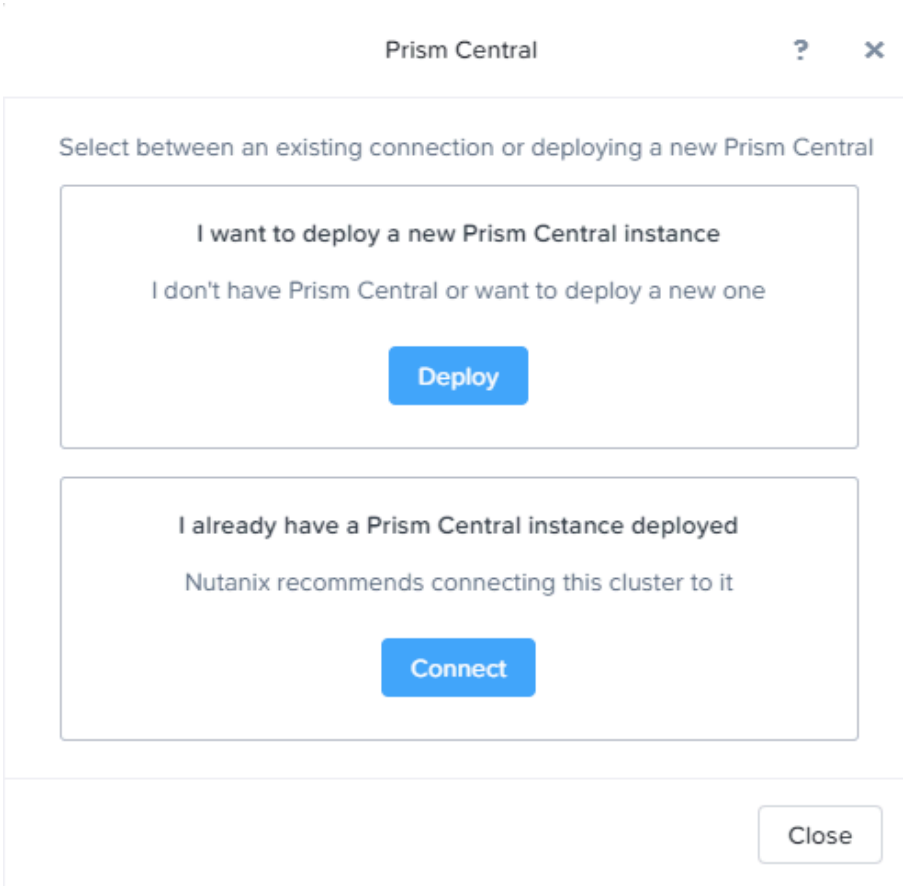

Figure 42: Prism Central Dialog Box (deploy/connect screen)

5. In the installation image screen, do one of the following:

   » Click the **Install** button to install that Prism Central version.

   » Click the **Download** button to download that Prism Central version from the Nutanix support portal. After it downloads, the button changes to **Install**. You can then install that Prism Central version.

   Note: If the Prism Central version you want to install does not appear in the list, typically because you do not have Internet access (such as at a dark site), you can click the **upload the Prism Central binary** link to upload an image from your local media.

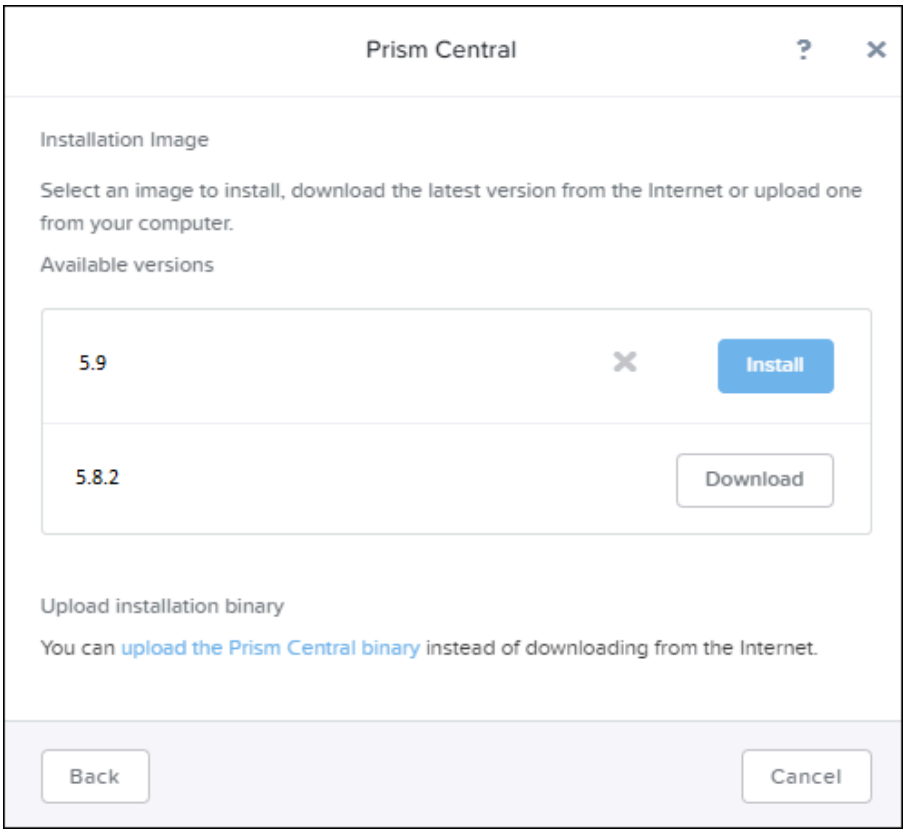

Figure 43: Prism Central Dialog Box (installation image screen)

6. In the deploy screen, do one of the following:

   » To deploy a 1-VM instance of Prism Central, click the **Deploy 1 VM PC** button.

   » To deploy a 3-VM instance of Prism Central, click the **Deploy 3 VM PC** button and then skip to step 8.

   A Prism Central instance can consist of either a single VM or a set of three VMs. A 3-VM instance increases both the capacity and resiliency of Prism Central at the cost of

maintaining the additional VMs. In addition, each VM can be either "large" or "small" in size. Thus, you may choose from four configurations.

Table 10: Prism Central Instance Configurations

VM Size	1-VM Instance	3-VM Instance
large	8 vCPUs, 32 GB of memory, and 2500 GiB of storage	24 vCPUs, 96 GB of memory, 7500 GiB of storage (combined)
small	4 vCPUs, 16 GB of memory, and 500 GiB of storage.	12 vCPUs, 48 GB of memory, and 1500 GiB of storage (combined)

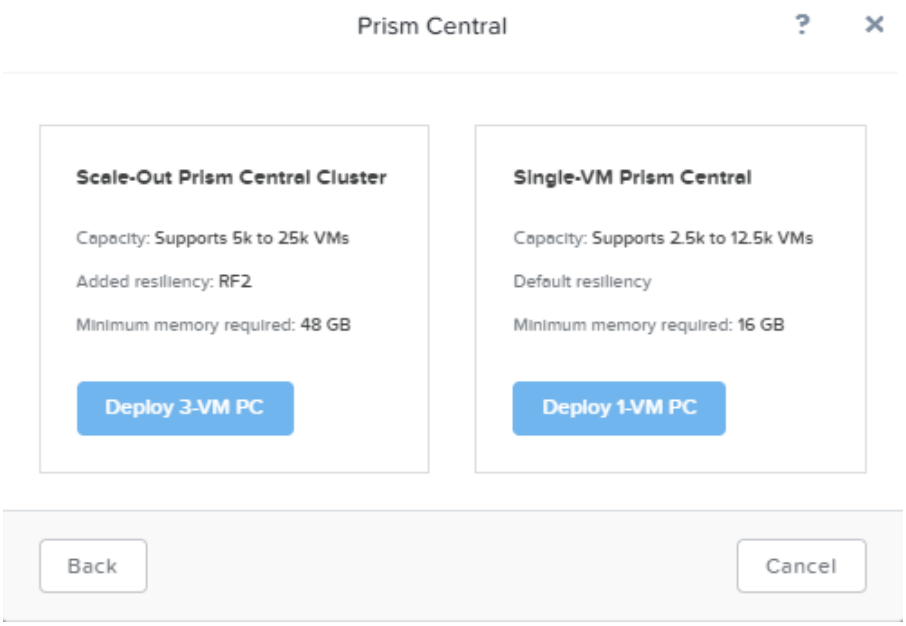

Figure 44: Prism Central Dialog Box (deploy screen)

7. [1-VM instance only] In the general configuration screen, do the following in the indicated fields:

   a. **VM Name**: Enter a name for the Prism Central VM.

   b. **Select a Container**: Select a container for the Prism Central VM from the pull-down list.

   c. **VM Sizing**: Select (click the radio button for) the Prism Central VM size based on the number of guest VMs it must manage across all the registered clusters:

      - **Small** - manage up to 2500 VMs
      - **Large** - manage up to 12500 VMs

   d. **Network**: Select an existing network for this Prism Central instance from the list.

      If the target network is not listed, click the **Create Network** link to create a new network. See the "Network Management" section in the *Prism Web Console Guide* for more information.

   e. **IP Address**: Enter a static IP address for the Prism Central VM.

   f. **Subnet Mask**: Enter the subnet mask value.

   g. **Default Gateway**: Enter the IP address of the gateway.

   h. **DNS Address(es)**: Enter the IP address for one or more DNS servers. Enter multiple addresses in a comma separated list.

   i. When all the fields are correct, click the **Deploy** button.

      This begins the deployment process. On the Home page, the Prism Central widget displays **Deploying** until the installation is completed, then it displays **OK**. Click **OK** to launch the Prism Central web console in your browser.

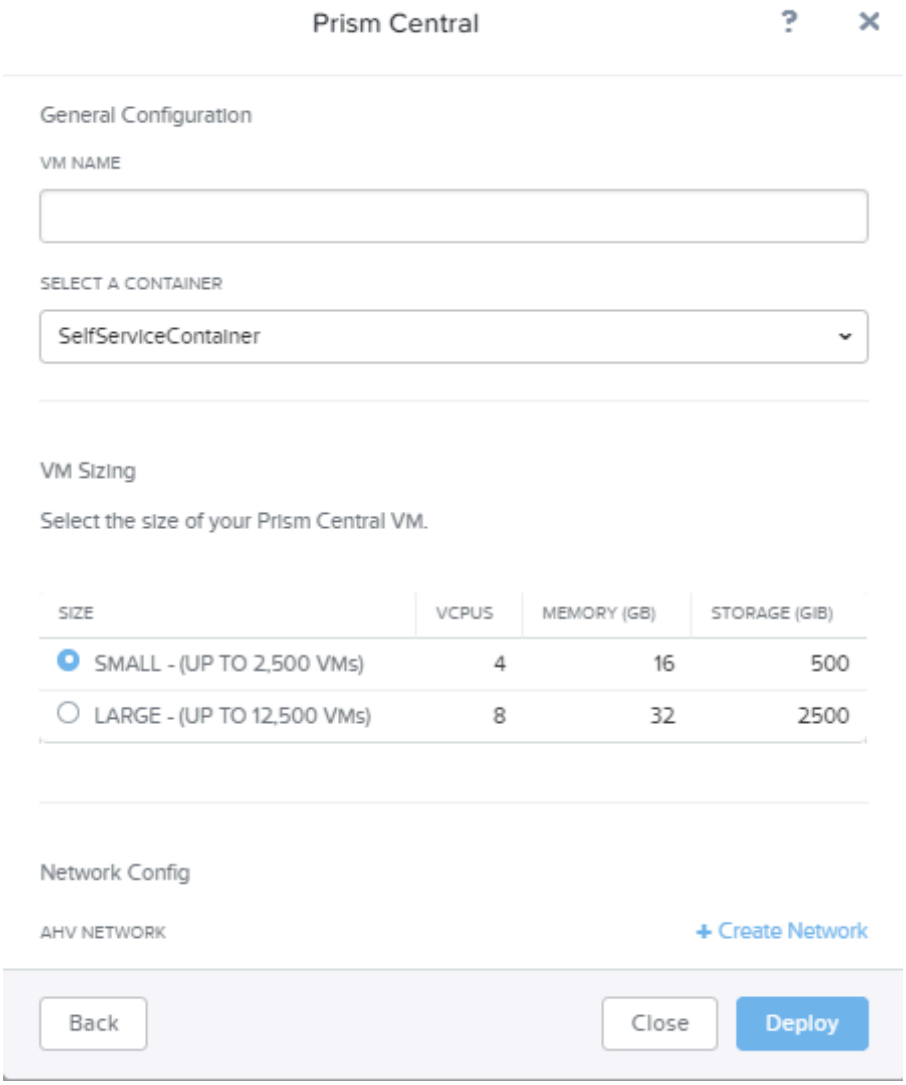

Figure 45: Prism Central Dialog Box (1-VM general configuration screen)

8. [3-VM instance only] In the deploy Prism Central screen, do the following in the indicated fields:

   a. **How many VMs will you support**: Select (click the radio button for) the Prism Central VMs size based on the number of guest VMs they must manage across all the registered clusters:

      - **Up to 5K VMs** - manage up to 5000 VMs
      - **Up to 25K VMs** - manage up to 25000 VMs

   b. **Network**: Select an existing network for this Prism Central instance from the list.

      If the target network is not listed, click the **Create New Network** option to create a new network. See the "Network Management" section in the *Prism Web Console Guide* for more information.

   c. **Subnet Mask**: Enter the subnet mask value.

   d. **Gateway**: Enter the gateway IP address.

   e. **Virtual IP** (optional): Enter a virtual IP address for the Prism Central instance.

      A virtual IP can be used as a single point of access for Prism Central.

   f. **Select a Container**: Select a container for the Prism Central VMs from the pull-down list.

   g. **<*IP address*>**: Specify IP addresses for the three new Prism Central VMs.

      > Note: If you entered a virtual IP, the IP addresses for the three VMs are populated automatically. You can keep those addresses by change them as desired.

      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 46: Edit IP Field

      2. Repeat this step for the other two VMs.

      The vCPU count and memory size in the new VMs are determined by whether you chose the small or large VM support size (5K or 25K radio button).

   h. **DNS Address(es)** (optional): Enter the IP addresses to one or more DNS servers. Enter multiple addresses in a comma separate list.

   i. When all the fields are correct, click the **Deploy** button.

      This begins the deployment process. On the Home page, the Prism Central widget displays **Deploying** until the installation is completed, then it displays **OK**. Click **OK** to launch the Prism Central web console in your browser.

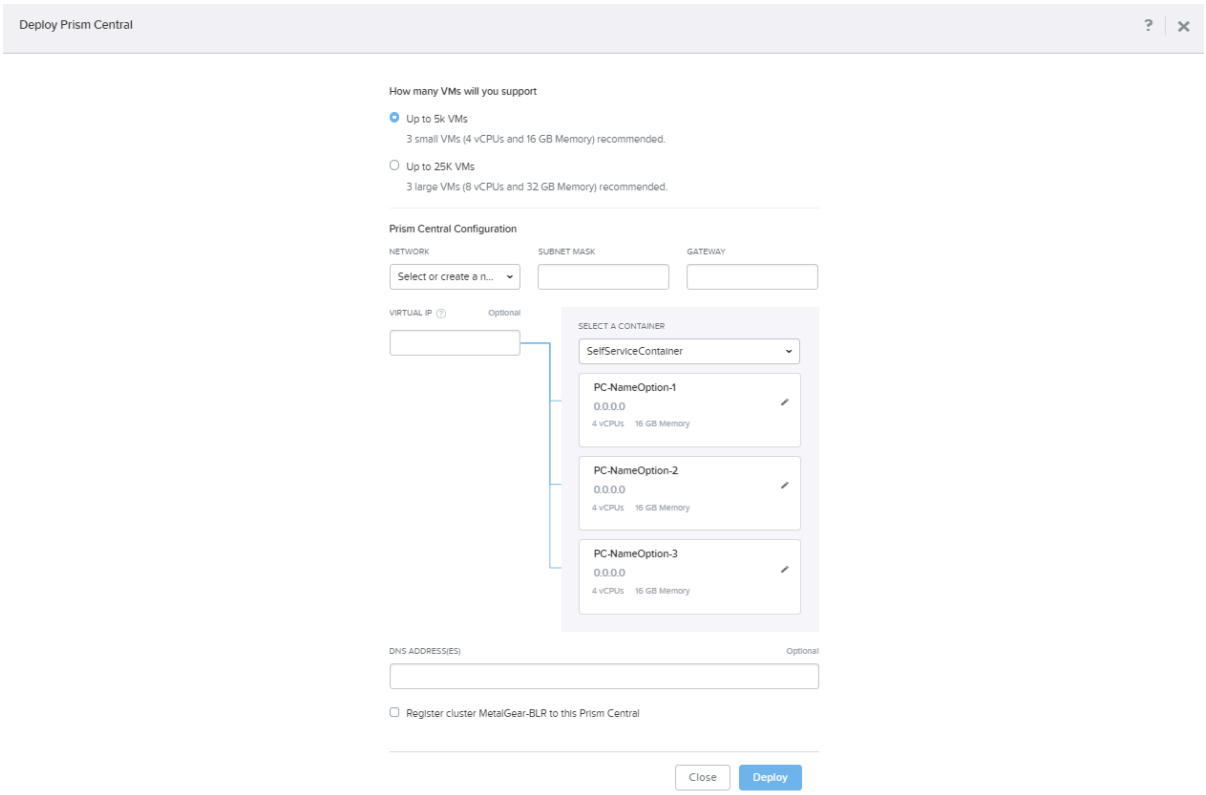

Figure 47: Deploy Prism Central Page (3 VM)

9. You can monitor the deployment progress (1-VM or 3-VM Prism Central instances) from the Tasks page and view information about the deployed VMs through the VM dashboard.

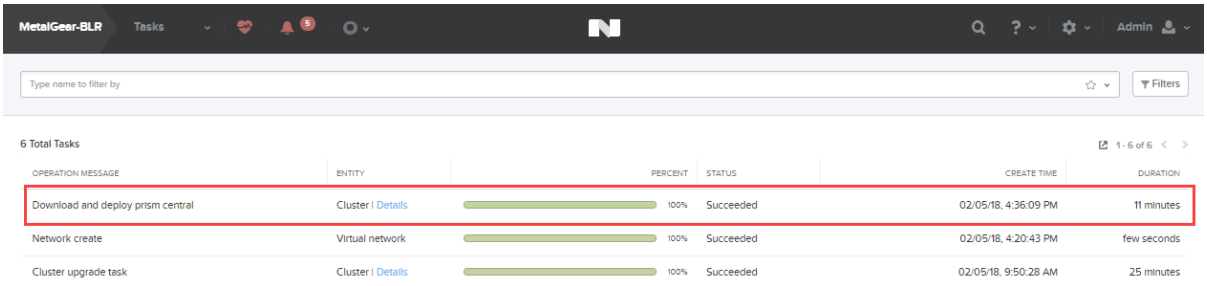

Figure 48: Tasks Page (Prism Central deployment)

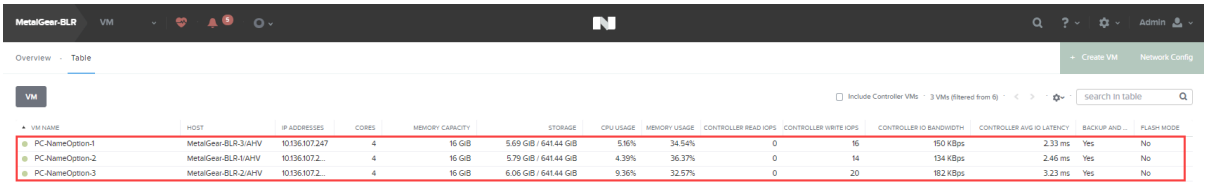

Figure 49: VM Dashboard (Table View)

What to do next

Register this cluster with Prism Central. The management features are not available until Prism Central registers the cluster in which it is located. The topic *Register (Unregister) Cluster with Prism Central* in the *Prism Central Guide* describes how to connect to an existing Prism Central instance.

Single-Node Clusters

A traditional Nutanix cluster requires a minimum of three nodes, but Nutanix also offers the option of a single-node cluster for ROBO implementations and other situations that require a lower cost option and can accept lowered resiliency protections. Single-node clusters are capable of running a limited number of user VMs. These are unlike single-node replication targets which are for replication and backup purposes. The following table outlines features and limitations of a single-node cluster.

Note:

- Single-node clusters are supported only on a limited set of hardware models.
- This is an overview video about a single-node cluster. (Link does not appear in PDF format.)

Figure 50: Single-Node Cluster Overview

A video on this topic is available https://youtu.be/qPys36C87mU on YouTube.

Table 11: Single-Node Cluster Features

Feature	Description
Controller VM minimum requirements	6 vCPU and 20 GB memory
Replication factor (RF)	RF 2 across drives within the same node
Hypervisors supported	AHV and ESXi
Unsupported features	• cluster expansion • deduplication • erasure coding • Volumes • near sync DR and Metro Availability (asynchronous DR is supported) • capacity analytics (in Prism)

Single-Node Cluster Guidelines

A single-node cluster is configured like a regular (three-node or more) cluster in many ways, but note the following for a single-node cluster:

- Do not exceed a maximum of 5 user VMs.

- Do not exceed a maximum of 1000 IOPS.
- Do not create a Prism Central instance (VM) in the cluster. There is no built-in resiliency for Prism Central in a single-node cluster, which means that a problem with the node takes out Prism Central with limited options to recover.
- LCM is not supported because firmware updates usually require services to be stopped (in a single-node cluster there is no other node to take over the workload from the node being updated).
- Async DR is supported for 6 hour RPO only.
- If an SSD fails, the cluster goes into a read-only state. In the read-only state, there are no create, read, update, or delete operations and no cluster changes. If the cluster experiences an unresponsive SSD, data loss is a risk. The Prism Central console offers an override option that allows users to continue writing to single-node clusters after it has entered read-only mode (see Overriding Read-Only Mode on page 79). Otherwise, be sure there is more than 1 SSD available for the system to place metadata.
- All user VMs must be shut down before upgrading a single-node cluster. If any user VMs are still running, a warning box appears.

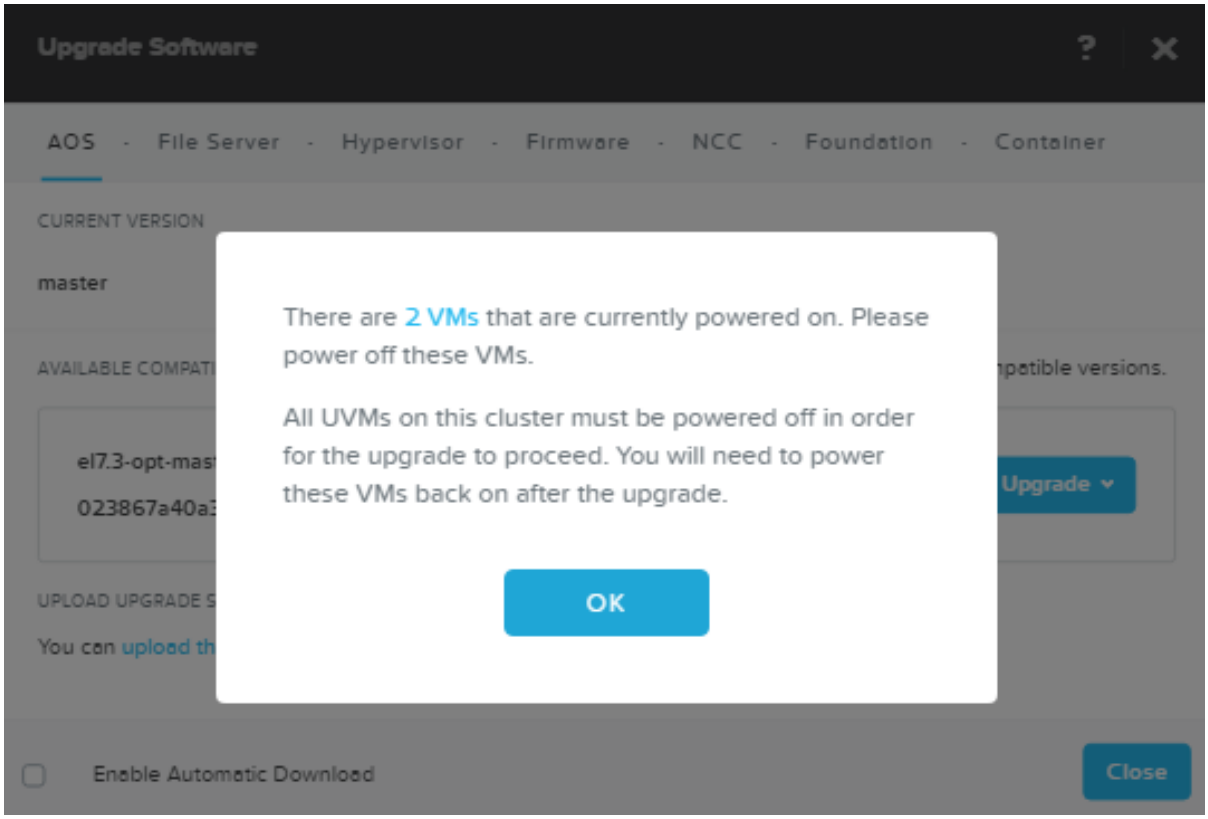

Figure 51: Upgrading Single-Node Cluster Warning

Overriding Read-Only Mode

About this task

Single-node clusters enter read-only mode when certain requirements are not met or a disk fails. An alert (a yellow exclamation mark) is generated in the web console that the single-node cluster has entered read-only mode.

Disk failure can be categorized into two categories:

- SSD Failure: In an event on an SSD failure or when an SSD is marked for removal, the single-node replication target cluster becomes a Read Only cluster and stops accepting any incoming replication traffic. An alert is displayed in the primary and single-node replication target clusters. Also, the resiliency status of the cluster is changed to the Critical state. The recovery process also slows down. Hence, it is recommended to get the cluster back into the Normal state by adding a new SSD to the cluster. After you remove the faulty disk and replace with the new SSD, this new SSD is automatically detected and the cluster is moved back into the normal state (read/write mode).
- HDD Failure: The HDD failure is handled in the same way as that of a normal cluster and none of the situations that occur when an SSD fails (as described above) occur in case of HDD failure.

> Note: If there is insufficient rebuild capacity when an HDD is unresponsive, then the cluster can fill up and IO can fail. For example, if you had an NX-1175S hybrid with 2 HDDs and 2 SSDs that has an HDD failure, the other disks do not have enough available space to bring up the data to disk fault tolerance.

Procedure

1. To perform this action through Prism, do the following:
   a. Click the yellow exclamation mark to re-enable write mode on the cluster.

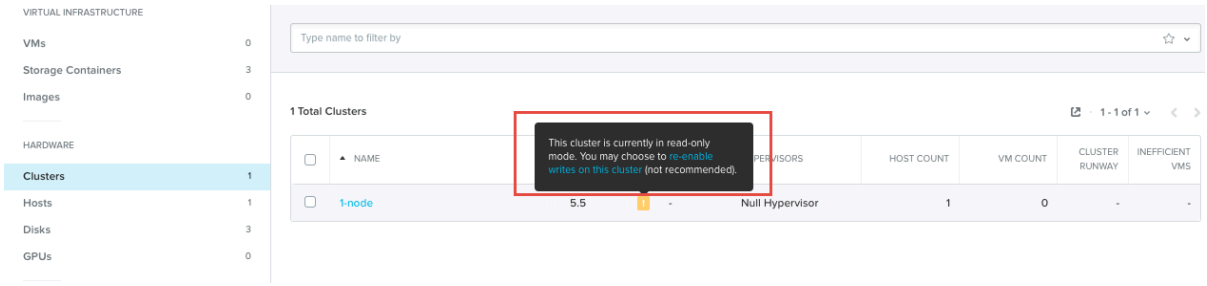

Figure 52: Re-enable Write Mode Message

   A dialog box appears to confirm that you want to override read-only mode.

   b. Click **Override** to confirm.

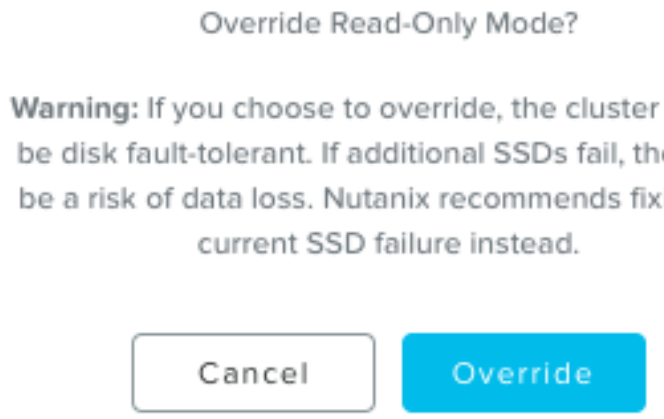

Figure 53: Confirm Override Mode

2. To perform this action through the nCLI, do the following:
   a. Log into the cluster using your Nutanix credentials.
   b. Enter the following command.

```
nutanix@cvm$ ncli cluster set-operation-mode operation-mode=override
```

Two-Node Clusters

A traditional Nutanix cluster requires a minimum of three nodes, but Nutanix also offers the option of a two-node cluster for ROBO implementations and other situations that require a lower cost yet high resiliency option. Unlike a one-node cluster (see Single-Node Clusters on page 78), a two-node cluster can still provide many of the resiliency features of a three-node cluster. This is possible by adding an external Witness VM in a separate failure domain to the

configuration (see Configuring a Witness (Two-Node Cluster) on page 84). Nevertheless, there are some restrictions when employing a two-node cluster. The following table outlines the features and limitations of a two-node cluster.

Note:

- Two-node clusters are supported only on a limited set of hardware models.
- This is an overview video about a two-node cluster. (Link does not appear in PDF format.)

Figure 54: Two-Node Cluster Overview

A video on this topic is available https://youtu.be/a-UpHAVEF60 on YouTube.

Table 12: Two-Node Cluster Features

Feature	Description
Controller VM minimum requirements	6 vCPU and 20 GB memory
Replication factor (RF)	RF2 spanned over two nodes and RF4 for metadata on SSDs over two nodes. RF4 for metadata helps during a node failure scenario to quickly transition the healthy node to run in single-node mode with the metadata remaining disk fault tolerant. (Metadata in a two-node cluster is typically small, so the storage need for four copies is modest.)
Single node failure effects	50% resource loss. Plan for 40% maximum resource usage to avoid read-only state on the remaining node. Data is made RF2 in the background so that data is resilient.
Drive failure effects	one node + one SSD failure (on other node) = read-only mode.
Hypervisors supported	AHV and ESXi
Unsupported features	• cluster expansion • dedupe (compression is supported) • erasure coding • near sync DR and Metro Availability (asynchronous DR is supported) • network segmentation

Two-Node Cluster Guidelines

A two-node cluster is configured and upgraded like a regular (three-node or more) cluster in most ways, but note the following for a two-node cluster:

- Size your implementation for N + 1 so that in the event of a node loss (50% loss of resources) the remaining node will have sufficient resources to allow the cluster to continue functioning.
- There is a heartbeat check (ping) between the nodes every two seconds. If a successful ping does not occur within 10 seconds (5 consecutive failed tries), a failover is initiated (see Failure and Recovery Scenarios on page 87). When the cluster recovers, it must remain in healthy status for at least 15 minutes before it will failback.
- The upgrade process in a two-node cluster may take longer than the usual process because of the additional step of syncing data while transitioning between single and two node state. Nevertheless, the cluster remains operational during upgrade.
- Witness VM considerations
  - A Witness VM for two-node clusters requires a minimum of 2 vCPUs, 6 GBs of memory, and 25 GBs of storage.
  - The same Witness VM can be used for both Metro Availability and two-node clusters, and a single Witness VM can support up to 50 instances (any combination of two-node and Metro Availability clusters).
  - You can bring up a two-node cluster without a Witness VM being present initially, but it is recommended that the Witness VM be alive and running before starting the cluster.
  - The Witness VM may reside on any supported hypervisor, and it can run on either Nutanix or non-Nutanix hardware.

    Note: Nutanix does not support the deployment of a Witness VM on the AWS and Azure cloud platforms.
  - The Witness VM must reside in a separate failure domain, which means independent power and network connections from any cluster it is managing. This separate platform can be deployed either on premise (including as an option the Nutanix one-node replication target NX-1155) or off premise (centrally, typically where Prism Central is hosted).
  - During a node failure, the transition of a healthy node to single-node mode can take 30-60 seconds. The client VMs can experience I/O timeouts during this period. Therefore, it is recommended that the SCSI timeout of the client VM disks should be at least 60 seconds.
  - The minimum recovery point objective (RPO) for a two-node (or one-node) cluster is six hours.
  - Network latency between a two-node cluster and the Witness VM should not exceed 500 ms. (RPC timeouts are triggered if the network latency is higher.) During a failure scenario, nodes keep trying to reach (ping) the Witness VM until successful. Nodes ping the Witness VM every 60 seconds, and each Witness request has a two-second timeout, so it can tolerate up to one second of link latency.
  - Node removal is not supported, but node replacement is supported (where the node remains part of the cluster). When replacing a node, the cluster remains active but will transition to standalone mode for a brief period of time.
  - All node maintenance work flows (software upgrades, life cycle manager procedures, node and disk break-fix procedures, boot drive break-fix procedure) require that the cluster be registered with a Witness VM.

Configuring a Witness (Two-Node Cluster)

About this task

A two-node cluster requires a Witness VM that is located in a separate failure domain either off premise or in a different physical platform on premise. You can configure a Witness either through the **Configure Witness** option in the main menu (any cluster) or when logging in to the Prism web console (two-node cluster only).

> Note: If the Witness VM does not exist yet, you must first create it. See “” () for instructions on how to install or upgrade a Witness VM. Two separate features, two-node clusters and metro availability, use a Witness VM. The installation and upgrade instructions (and the configuration instructions in this section) apply regardless of whether you are using a Witness VM for a two-node cluster or metro availability.

Procedure

- To configure a witness when logging in to the Prism web console, do the following:
  a. Get the credentials of the target witness VM (IP address, user name, and password).
  b. Click the gear icon in the main menu and then select **Configure Witness** in the **Settings** page (see Settings Menu on page 50).
  c. In the Configure Witness window, enter the Witness VM IP address, user name, and password in the indicated fields, and then click the **Save** button.

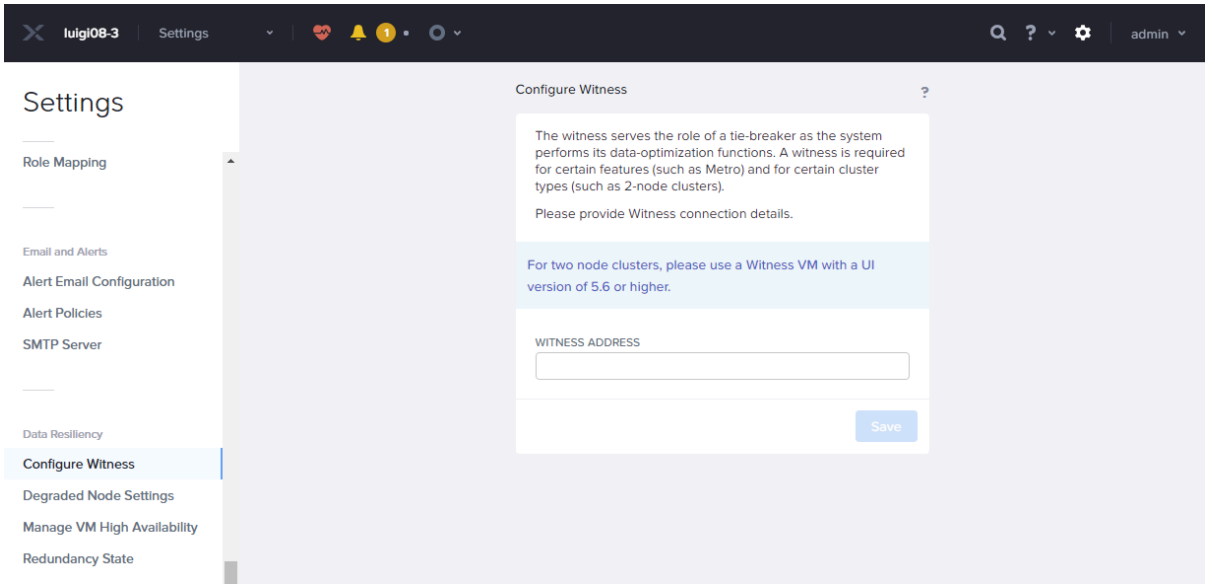

Figure 55: Configure Witness Window