the categories you want to add to the **Associated Categories** list. You can select multiple categories for one storage policy.

> Note: When you associate a category with one storage policy, you cannot associate the same category with another storage policy.

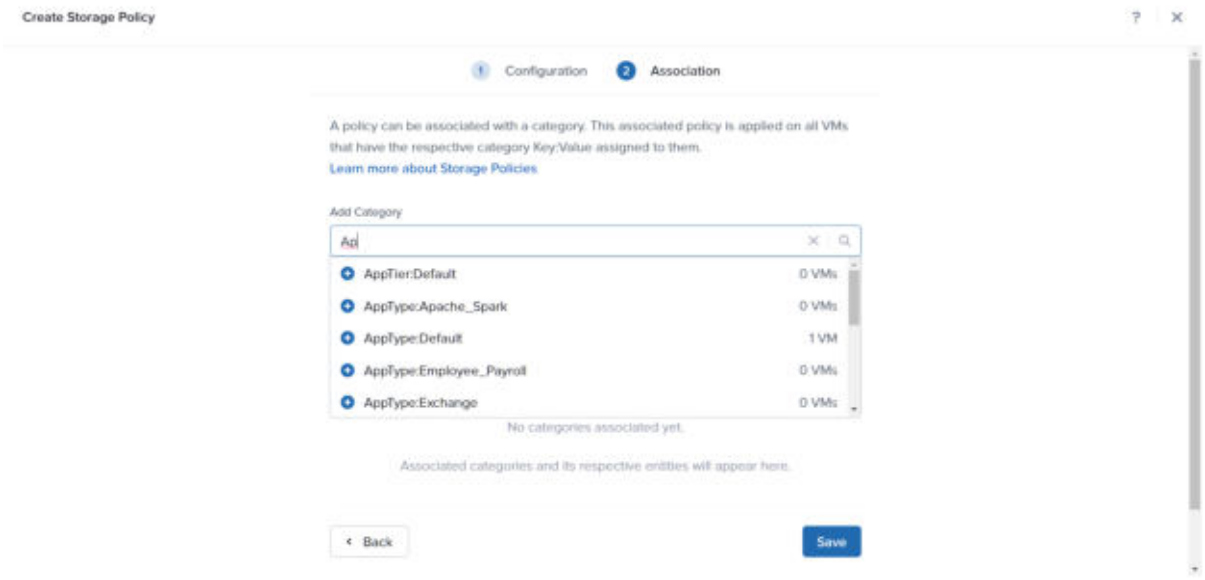

Figure 323: Create Storage Policy - Association Tab

For information about managing category associations like associating more categories during an update of the storage policy or removing a category association, see Managing Associations on page 653.

The category is added to the **Associated Categories** list when you click the category.

4. Click **Save** to save the storage policy.

Managing Associations

You can update a storage policy to add or remove a category association.

About this task

Storage policies use categories to associate with entities like VMs. A storage policy cannot associate directly with an entity like a VM. To associate an entity such as a VM, add an appropriate category to the VM and then associate the category with the storage policy. See *Limitations* section in Storage Policy Management on page 647.

> Note: You can associate multiple categories for one storage policy. When you associate a category with one storage policy, you cannot associate the same category with another storage policy.

Add or remove associations of categories with the storage policy in the Categories or Entities tabs of the storage policy details page. You could also select the storage policy in the storage policy summary **List** and click **Update** in the **Actions** menu, then add or remove the associations in the **Associations** tab of the Update Storage Policy page.

To add or remove associations through the storage policy details page, do the following:

Procedure

1. Click the name of the storage policy. On the storage policy details page, click the **Categories** (or Entities) tab.
2. On the **Categories** page, click **Manage Associations**.
   On the Manage Associations page, you can add or remove category associations.
3. In the **Add Category** field, click the category or categories you want to associate with the storage policy.
   Enter the first 2-3 characters of the category name you want to associate with in the **Add Category** field. This opens a drop-down list of categories that you can select from. Click the plus icon of the categories you want to add to the **Associated Categories** list.

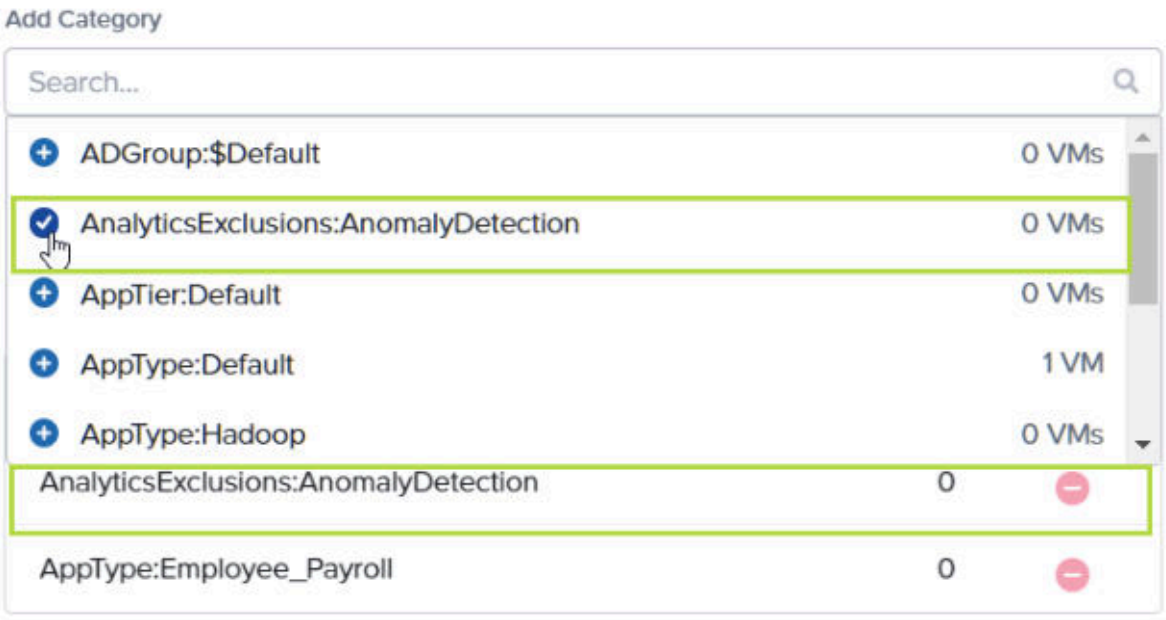

Figure 324: Add Category

The plus icon changes to a check mark icon indicating that the category is added to the **Associated Categories** list.

4. In the **Associated Categories** list, click the remove icon of the category (or categories) you want to remove.

Figure 325: Remove Categories

The category is removed from the **Associated Categories** list, thus removing the association of the category with the storage policy.

Deleting a Storage Policy

You can delete a storage policy. When you do so, the entities associated with the deleted storage policy return to their default state or the configurations from any other storage policy that is associated with the category are applied to the entities.

Before you begin

Nutanix recommends that you remove all the categories and entities associated with the storage policy. See Managing Associations on page 653 for more information about removing associations.

About this task

When a storage policy is deleted, it is marked for removal in the database. Prism Central stops displaying that storage policy. The process stops applying the storage attributes of the policy and instead applies the storage attributes configured for the container. However, the VM continues to display the storage policy as if it still applies, until the next system scan occurs and removes the policy from VM information. (A Kanon service scan occurs every 1800 seconds and refreshes the information.)

Remove associations of the storage policy to categories in the Categories or Entities tabs of the storage policy details page. See Managing Associations on page 653 for steps to remove associations of a storage policy.

Delete a storage policy from the Storage Policies dashboard or the Summary tab of the storage policy details page.

Do the following to delete a storage policy:

Procedure

1. Click the name of the storage policy. On the storage policy details page, click **Delete**.
   The Delete <storage-policy-name> dialog box is displayed as follows:

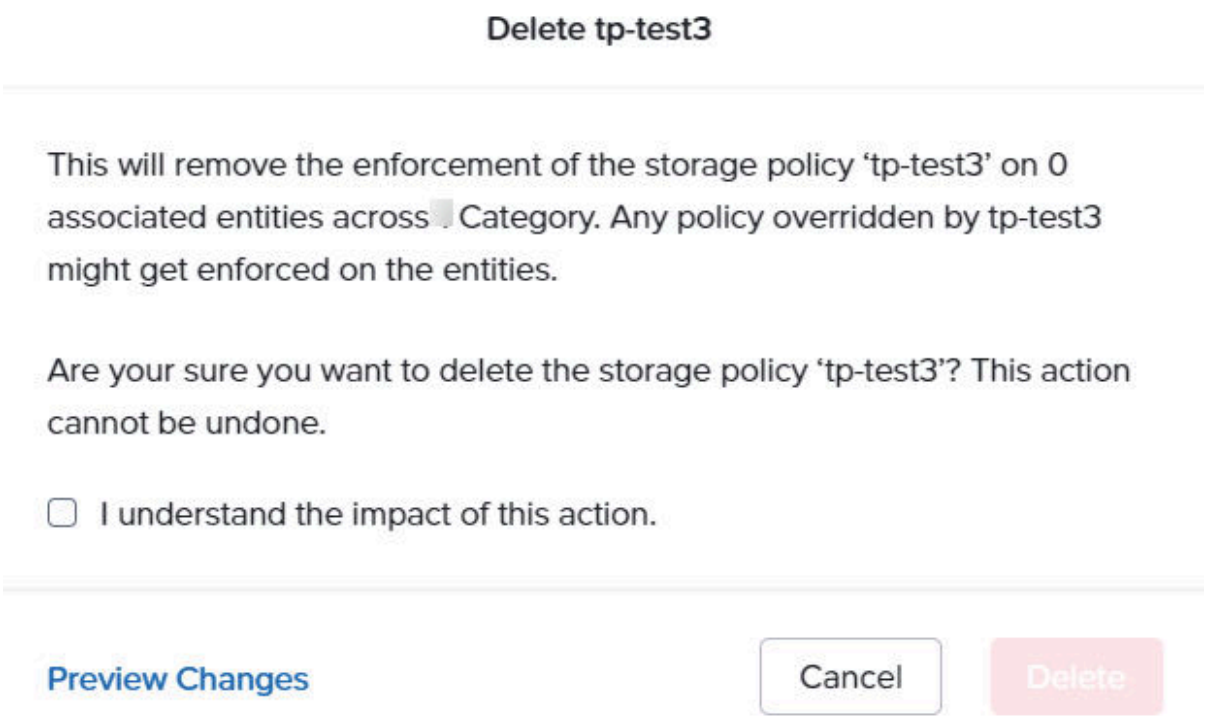

Figure 326: Delete Storage Policy

2. (Optional) Click **Preview Changes** to view the **Changes Preview** dialog box.
   The **Changes Preview** dialog box lists the categories and entities that are affected by the deletion of the storage policy.

   Tip: If you have already removed all the associations of the storage policy, then the list of categories and entities is empty.

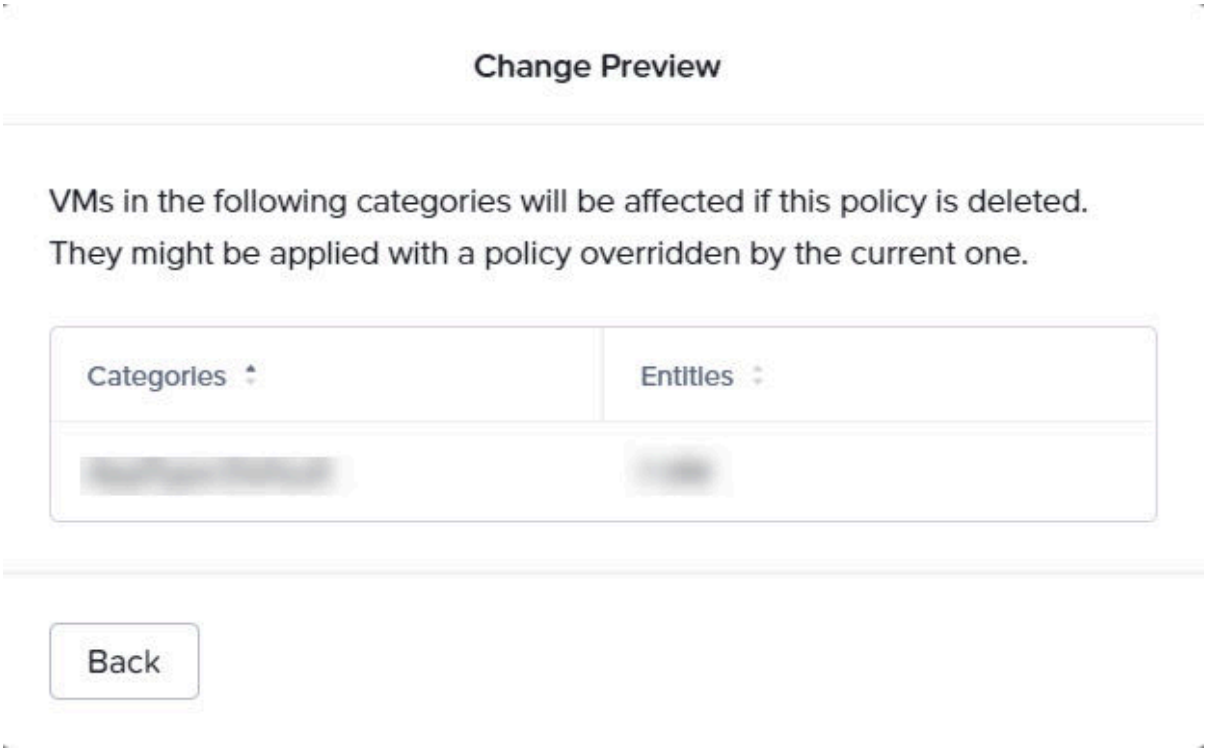

   Figure 327: Change Preview

   After you preview the impact, click **Back** to go back to the Delete <storage-policy-name> dialog box.

3. Check the **I understand the impact of this action.** box.
   The **Delete** button is actionable only after you check the **I understand the impact of this action.** box.

4. Click **Delete**.

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 713).
- You can manage the Catalog Item entity permissions by enabling or disabling the Marketplace Item permissions while creating a custom role. For example, if you

select **No Access** permission for the Marketplace Item entity while creating the custom role, the custom role will not have access to the Catalog Item entity as well. See the *Creating A Custom Role* section of the *AOS Security Guide* on how to create a custom role with access to catalog items entity.

Adding a Catalog Item

About this task

When you add a VM to the catalog, a snapshot of the VM is created. The snapshot is available to users across all the self-service projects. Users who have the requisite permissions can create VMs from the snapshot. You can continue to use the VM or delete the VM after you add it to the catalog. These actions do not affect the snapshot.

Adding an image to the catalog makes the image available to self-service users who have permissions to create VMs. A copy of the image is added to the catalog, so you can delete the image without affecting the copy in the catalog.

To add a VM snapshot or image to the catalog:

Note: Nutanix recommends that you power off the VM before adding a snapshot of it to the catalog.

Procedure

1. Do one of the following:
   - » Go to the **List** tab of the VMs dashboard and select the VM that you want to add to the catalog (see VMs Summary View on page 86).
   - » Go to the images dashboard and select the image that you want to add to the catalog (see Images Summary View on page 117).

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.

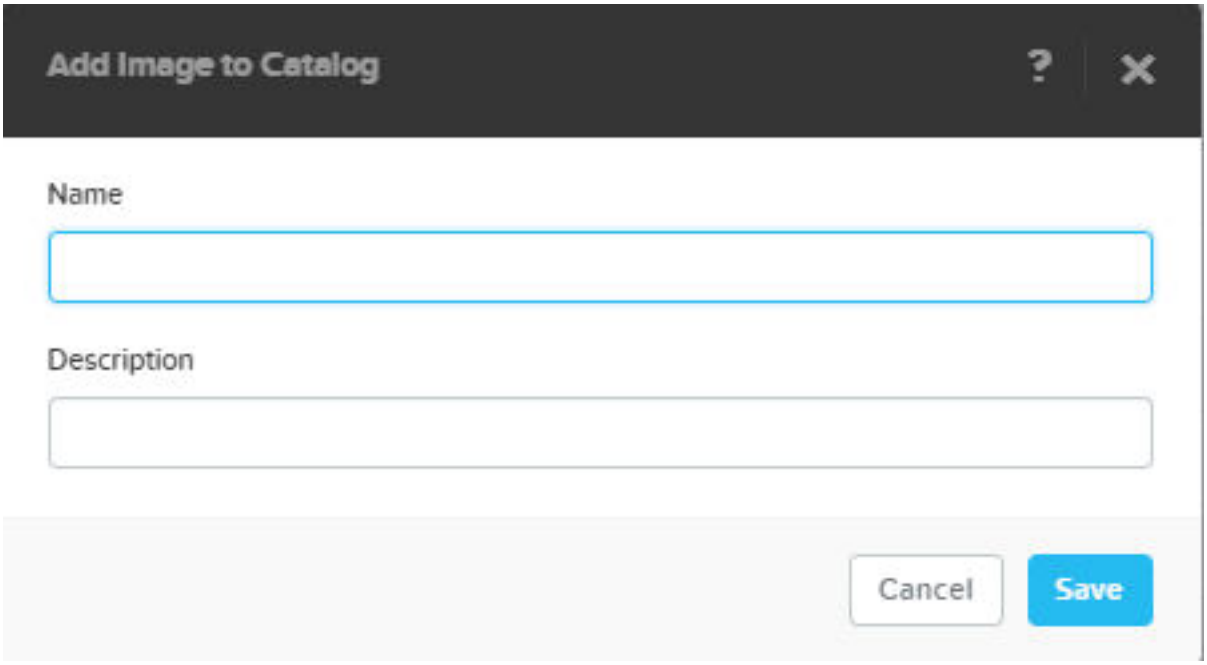

Figure 328: Add Image to Catalog Window

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 329: 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 86).

   » Go to the images dashboard and select the image that you want to delete from the catalog (see Images Summary View on page 117).

2. From the **Actions** menu, select **Delete Catalog Item**.

Image Management

Prism Central provides a centralized location to manage the images you require on registered AHV clusters. Prism Central enables you to upload images to the clusters and maintains an inventory of the images on them.

> Note: You can configure policies that govern which clusters receive the images that you upload (see Image Placement Policies on page 680).

Figure 330: Image Management Demo

A video on this topic is available https://youtu.be/UlqM8jjvr4s on YouTube.

Uploading Images from Prism Central to Registered Clusters

From Prism Central, you can upload images by using the following methods:

- Uploading images from a workstation
- Specifying URLs to images on a remote server
- Uploading images from a VM disk

For information about these methods, see Creating an Image on page 662.

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 680.

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 663 and Uploading Images from a Remote Server on page 667, 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 679.

Requirements

The image management feature in Prism Central has the following requirements:

- The version of both Prism Central and the version of AOS on registered clusters must be 5.8 or later.
- Clusters must be registered with Prism Central and must have AHV installed.
- The image service uses port 2007, so that port must be open.

Limitations

Image management from Prism Central has the following limitations:

- Prism Central cannot update images that another Prism Central instance managed earlier. However, the images are displayed in Prism Central, and you can use the images to create VMs or delete the images you no longer need.
- You cannot choose a container when uploading images from Prism Central. Prism Central uploads images to the container named `SelfServiceContainer`.
- When uploading images from a workstation, If multiple clusters are selected during the initial image upload then the uploaded image becomes active (in Prism Element) on only one of the selected clusters and is inactive on the other clusters registered to Prism Central.
- If an image is not active on a cluster, as an admin user it is possible that you are unable to use Prism Central to create that VM on that cluster. The cluster on which you want to create the VM cannot check out the image from a remote cluster. In this scenario, you must upload the image manually to the cluster on which you want to create the VM.

  When you create a VM, Prism Central uses API v2 workflows. API v2 does not have the check-out workflow. The image checkout workflow is used to copy an image on demand from cluster B to cluster A if the image does not exist or is inactive on cluster A.

  > Note:
  >
  > This limitation is removed by applying an image placement policy and a category to the image that is available on multiple remote clusters. When you apply an image placement policy to an image, Prism Central propagates the image to all the clusters that are included in the image placement policy and category. Prism Central can then check-out the image from the remote cluster where the image is active and create the VM.
  >
  > Prism Central can take up to 15 minutes to switch the image propagation from an alternative cluster. For example, if an image is located on clusters A and B. As a result of an image placement policy you applied to that image, cluster A starts copying the image to cluster C. During the copy operation, cluster A becomes unavailable. In such a case, Prism Central takes up to 15 minutes to switch to cluster B and start copying the image from cluster B to cluster C.

Creating an Image

You can add an image in one of two ways.

About this task

To add an image, do one of the following:

Procedure

- Upload Images from a workstation. See Uploading Images from a Workstation on page 663.
- Uploading Images from a Remote Server. See Uploading Images from a Remote Server on page 667.
- Uploading Images from a VM disk. See Uploading Images from a VM Disk on page 672.

Uploading Images from a Workstation

You can browse the workstation from which you are accessing Prism Central and upload images from it.

About this task

You can select multiple images and upload them as part of a single operation. With this method of upload, Prism Central identifies a single registered cluster for each specified image and uploads the image to it. The choice of cluster is random. The image becomes active (in Prism Element) on the selected cluster and inactive on other clusters registered to Prism Central. Image placement policies take effect after a short duration.

> Note: Most modern browsers impose file size limitations that affect this upload method. If you must upload images larger than 2 GB, upload the images from a remote server instead. Also, the browser type and CPU and RAM utilization on the workstation limit the number of concurrent uploads. Concurrent uploads exceeding the default limit of the browser are queued or throttled by the browser and can take more time. Large file uploads and high CPU and memory utilization can slow down the browser.

To upload an image from a workstation, do the following:

Procedure

1. In the **Entities** menu, go to **Compute & Storage** > **Images**, and then click the **Add Image** button (see Images Summary View on page 117).
   The Add Images page appears.

2. Under **Image Source**, click **Image File**, and then click **Add File**.

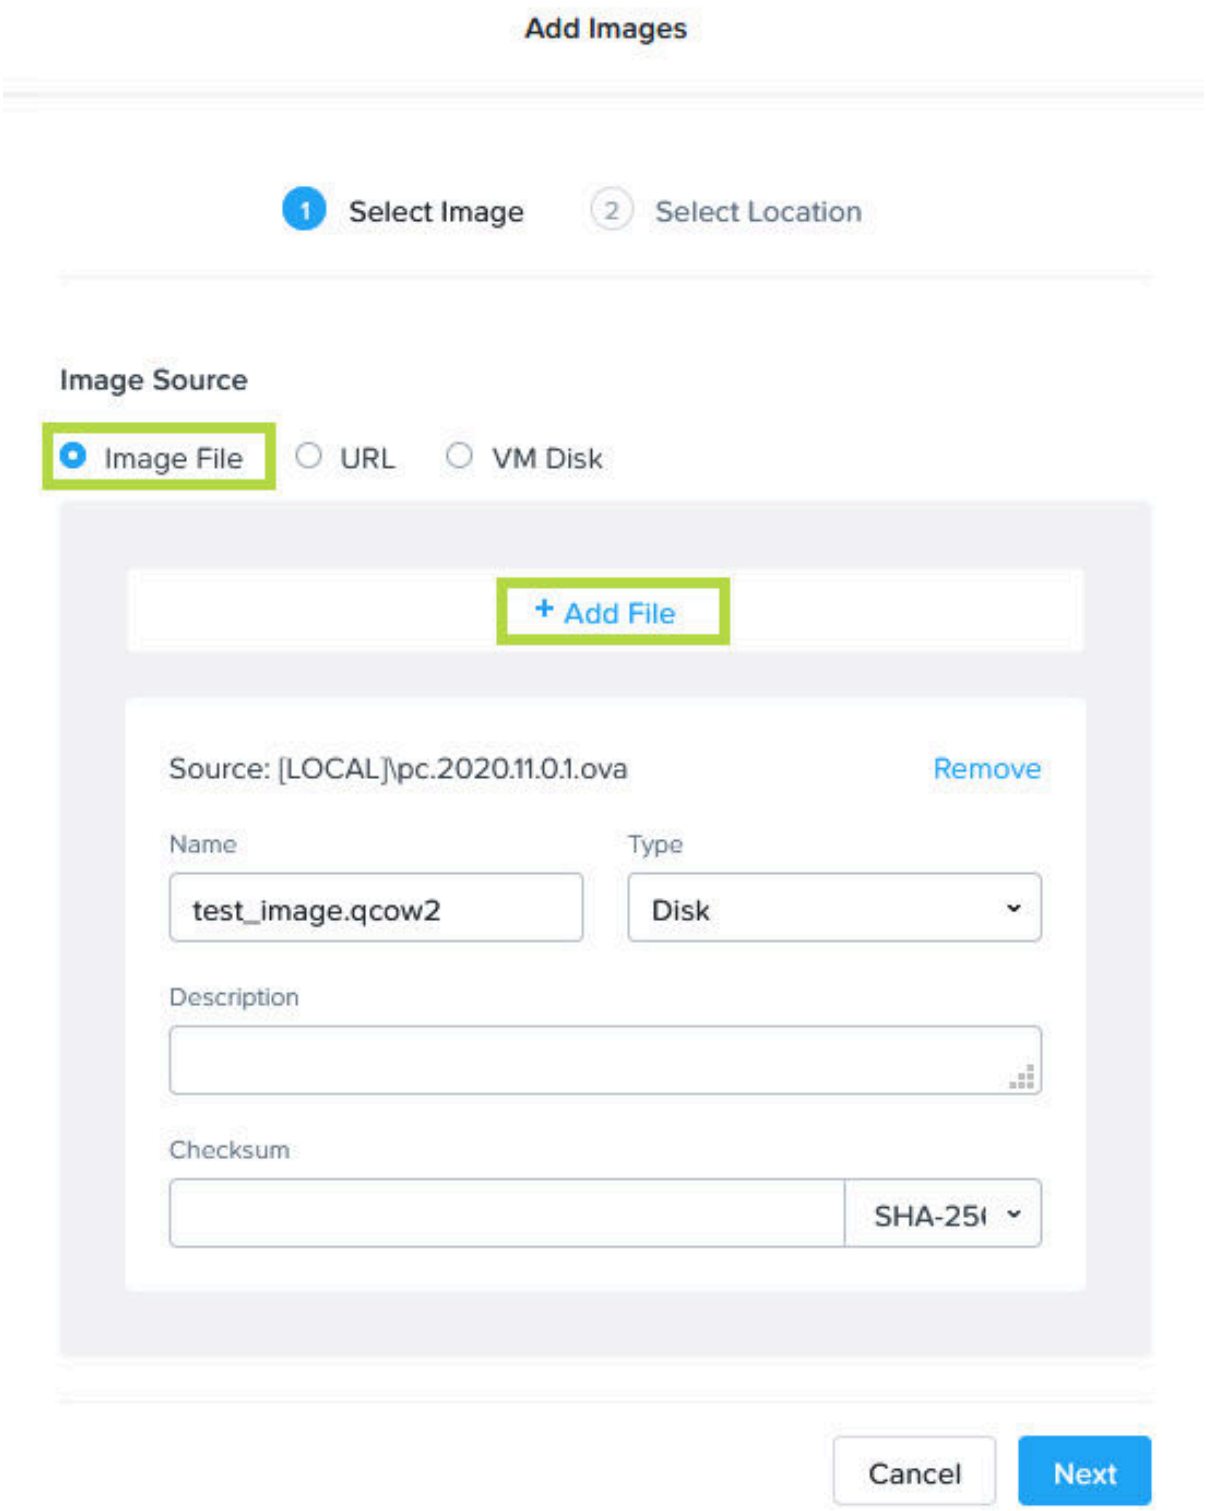

Figure 331: Add Images: File

3. To specify one or more files for upload, do the following:

   a. Browse to the location of the image file, and then click **Open**.
      Prism Central adds the image file to a list on the page.

   b. In **Image Name**, make sure that the name of the image is unique across all the images in Prism Central.
      The name defaults to the name of the file you selected, but you can change that name here.

   c. In **Image Type**, select the type of image.

   d. In **Image Description**, describe the image file.

   e. In **Checksum**, select the hashing algorithm.

   Repeat this step to add as many image files as you want.

4. To remove an image file entry, locate the entry and click **Remove**.

5. After you have added all the image files you want, click **Next**.
   The **Select Location** page appears.

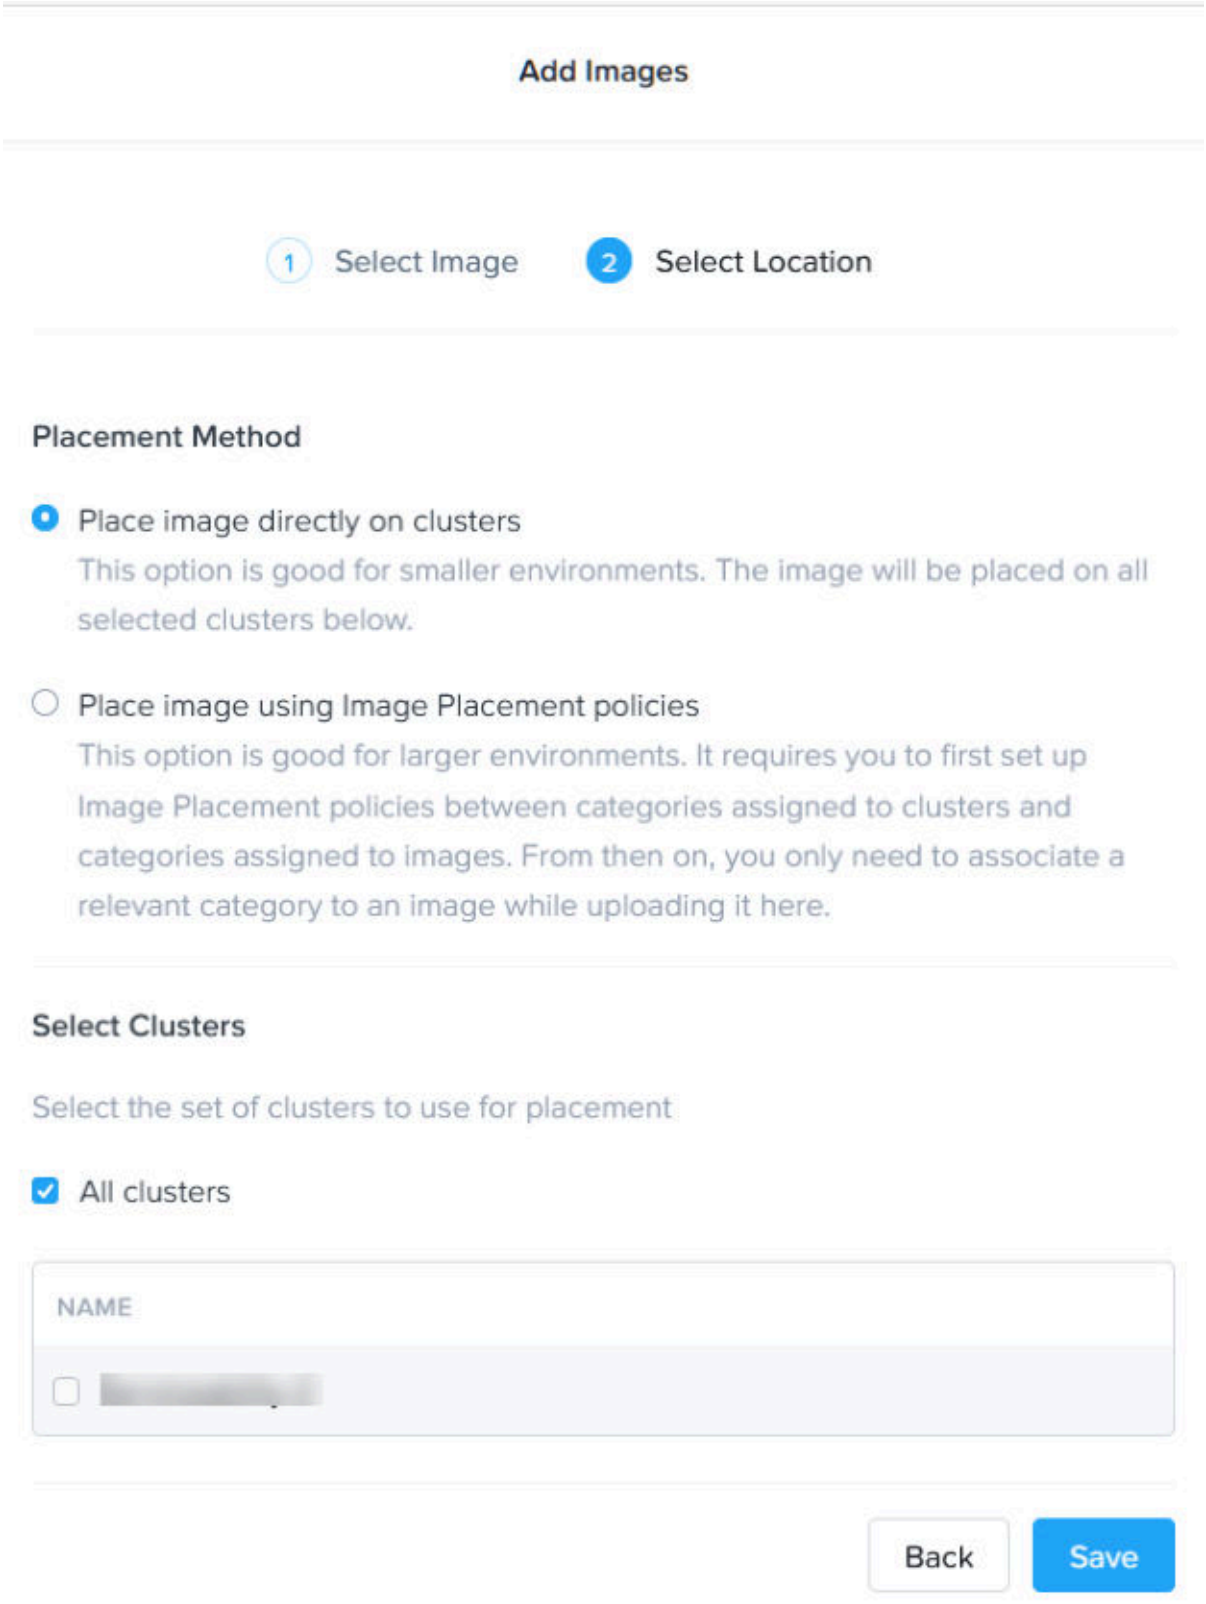

Figure 332: Add Images: Select Location

6. In **Placement Method**, do one of the following:

   - To manually select the target clusters, click **Place image directly on clusters**, and then do one of the following in the **Cluster Details** section:

     - If you want to upload the images to all registered clusters, make sure to select **All clusters**, and then click **Save**.
     - If you want to upload to only a subset of the registered clusters, clear **All clusters**, select the clusters you want from the list, and then click **Save**.

   - To delegate image placement decisions to configured policies and assign categories to the images, click **Place image using Image Placement policies**, and then do the following in the **Categories** section:

     1. Click inside the **Categories** search box and select the category you want from the list. You can also start typing the name of the category to reduce the list to matching names.
     2. To specify another category, click the add icon beside the search box. Repeat this step to add as many categories as you need.

   Prism Central uploads the image files in batches and takes some time to enforce the image placement policies.

Uploading Images from a Remote Server

You can upload an image from a remote server to registered Nutanix clusters. To upload an image from a remote server, you need the URL of the image. You can also specify URLs to multiple images as part of a single operation. When you specify image URLs, Prism Central uploads the images to all registered clusters. Uploading from a remote server also allows you to overcome file size limitations imposed by modern browsers. The file size limitation is usually 2 GB.

About this task

To upload an image from a remote server, do the following:

Procedure

1. In the **Entities** menu, go to **Compute & Storage** > **Images**, and then click the **Add Image** button (see Images Summary View on page 117).
   The Add Images page appears.

2. Under **Image Source**, click **URL**.

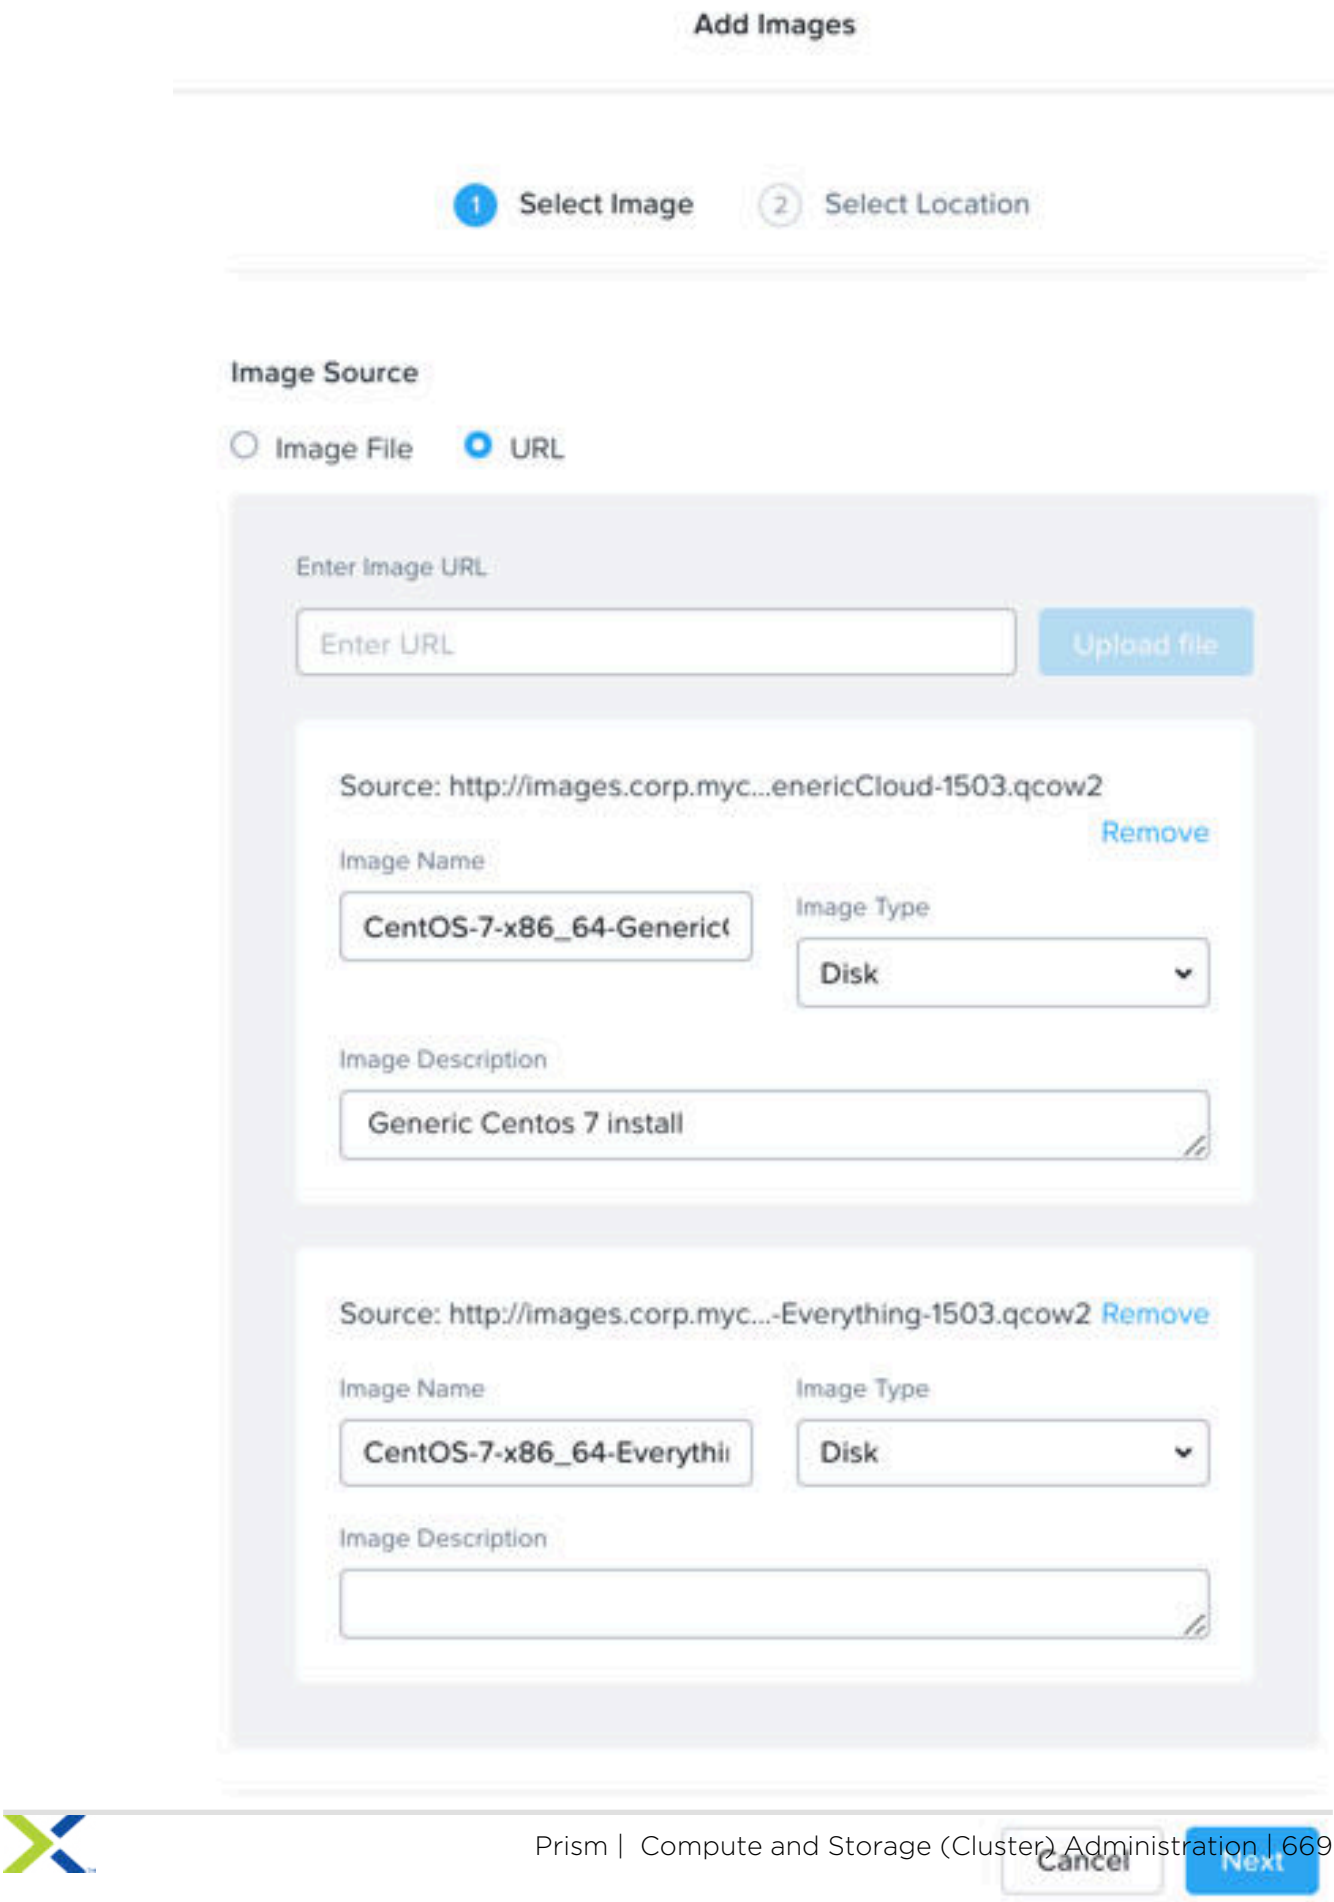
Add Images
1 Select Image
2 Select Location
Image Source
Image File
URL
Enter Image URL
Enter URL
Upload file
Source: http://images.corp.myc...enericCloud-1503.qcow2
Remove
Image Name
CentOS-7-x86_64-Generic(
Image Type
Disk
Image Description
Generic Centos 7 install
Source: http://images.corp.myc...-Everything-1503.qcow2 Remove
Image Name
CentOS-7-x86_64-Everythi
Image Type
Disk
Image Description
Cancel
Next

3. To specify one or more image URLs, do the following:

   a. In **Enter Image URL**, enter the appropriate URL address in the field using the following syntax for either NFS or HTTP, and then click **Upload File**. (NFS and HTTP are the only supported protocols.)

      ```
      nfs://[hostname|IP_addr]/path
      http://[hostname|IP_addr]/path
      ```

      Enter either the name of the host (*hostname*) or the host IP address (*IP_addr*) and the path to the file. If you use a *hostname*, the cluster must be configured to point at a DNS server that can resolve that name (see Configuring Name Servers on page 514). A file uploaded through NFS must have 644 permissions.

      Prism Central adds the URL to a list on the page and clears the **Enter Image URL** text box for another URL.

   b. In **Image Name** in the list entry, make sure that the name of the image is unique across all the images in Prism Central.

      The name defaults to the name of the file you selected, but you can change that name here.

   c. In **Image Type**, select the type of image.

   d. In **Image Description**, describe the image file.

   Repeat this step to specify as many URLs as you want.

4. To remove an image file entry, locate the entry and click **Remove**.

5. After you have added all the URLs you want, click **Next**.
   The **Select Location** page appears.

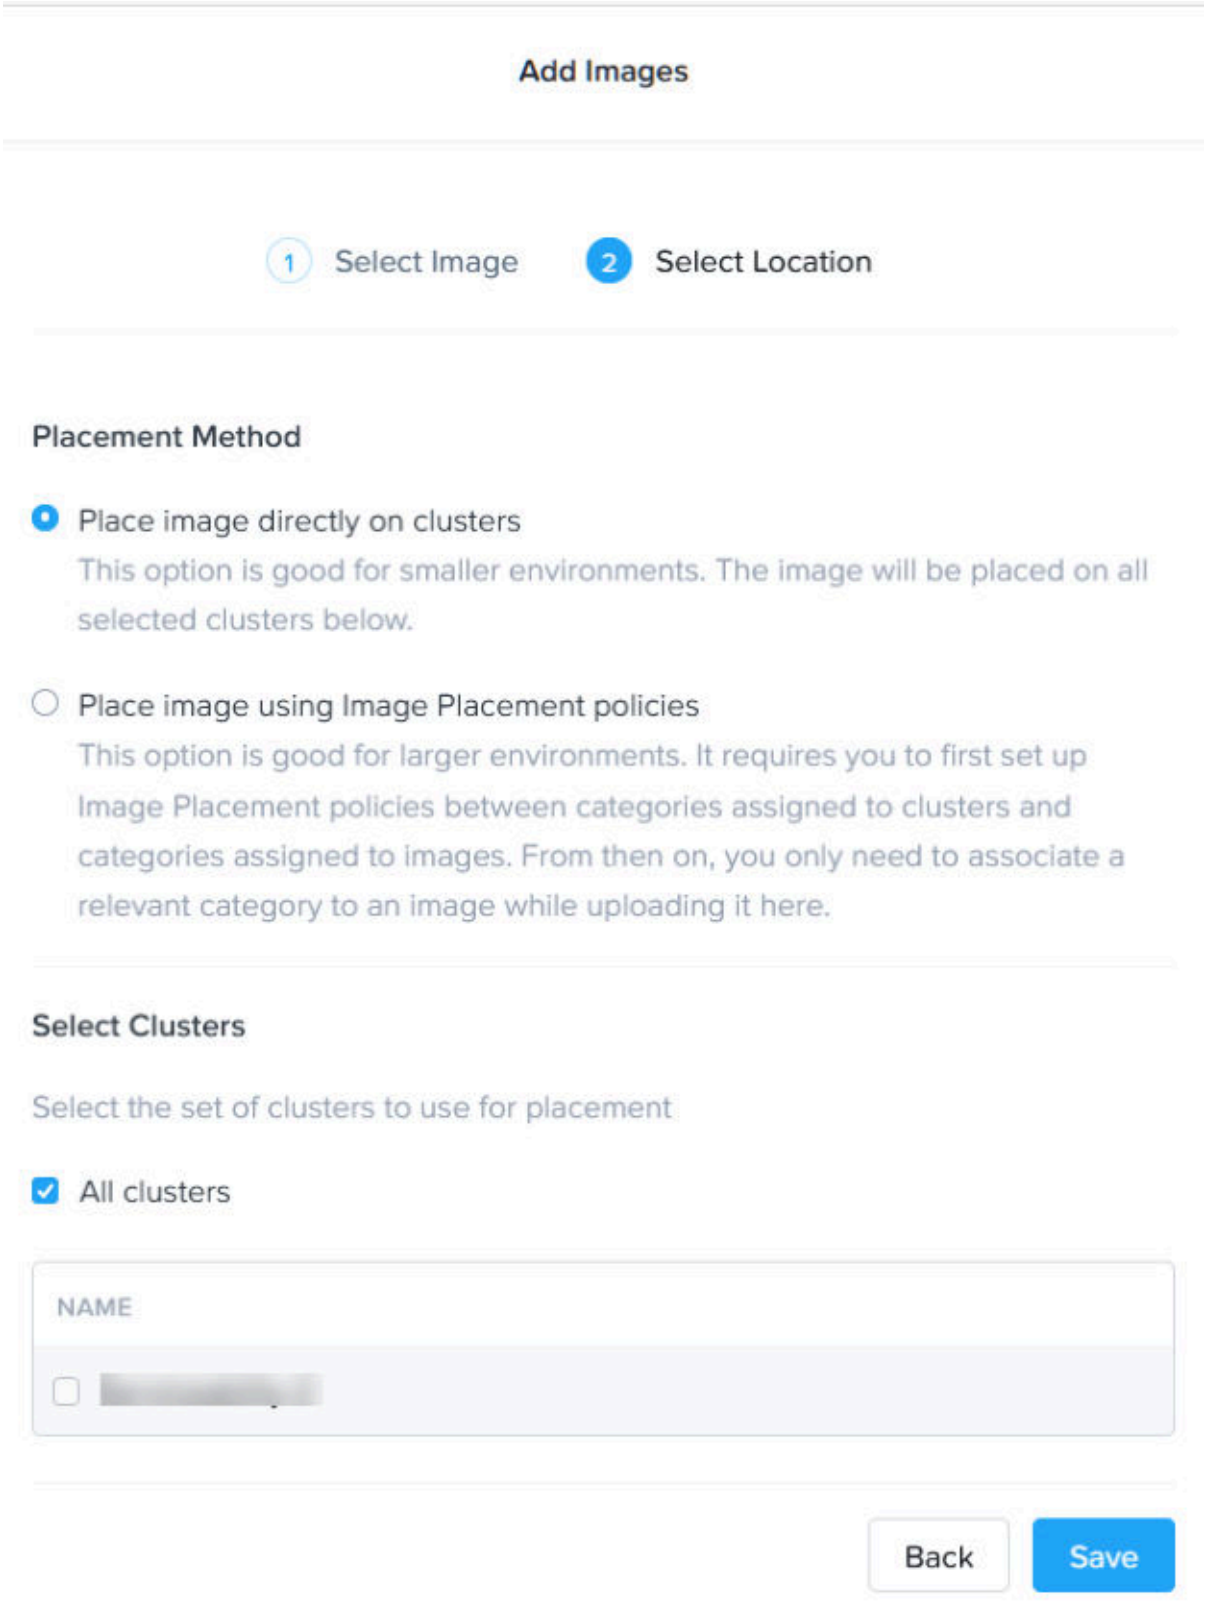

Figure 334: Add Images: Select Location

6. In **Placement Method**, do one of the following:

   - To manually select the target clusters, click **Place image directly on clusters**, and then do one of the following in the **Cluster Details** section:

     - If you want to upload the images to all registered clusters, make sure to select **All clusters**, and then click **Save**.
     - If you want to upload to only a subset of the registered clusters, clear **All clusters**, select the clusters you want from the list, and then click **Save**.

   - To delegate image placement decisions to configured policies and assign categories to the images, click **Place image using Image Placement policies**, and then do the following in the **Categories** section:

     1. Click inside the **Categories** search box and select the category you want from the list. You can also start typing the name of the category to reduce the list to matching names.
     2. To specify another category, click the add icon beside the search box. Repeat this step to add as many categories as you need.

   Prism Central uploads the image files in batches.

Uploading Images from a VM Disk

You can clone an existing VM disk to add it to the image list.

About this task

Figure 335: Uploading Images from a VM Disk

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6271311948001 on Brightcove.

- You select a powered-off source VM on an AHV cluster.
- You then select a disk attached to this source VM to add it as an image.

You can also upload the added image to the registered clusters based on the image placement method. You can either:

- Place the image on the cluster corresponding to where the source VM resides.
- Associate relevant categories to the image so that the corresponding clusters are internally selected based on categories and associated policies.

Before you begin

Ensure that the following requirements are met before you add and upload an image from a VM disk.

- The source VM belongs to an AHV cluster.
- The source VM on which the disk resides is in the powered-off state.

- To place the image on clusters associated with relevant categories, you must first set up the image placement policies between categories assigned to clusters and categories assigned to images.

  See the Creating a Category on page 735 and Assigning a Category on page 737 sections for information on how to create a category and assign a category value to an image respectively.

Procedure

1. In the **Entities** menu, go to **Compute & Storage** > **Images**, and then click the **Add Image** button (see Images Summary View on page 117).
   The Add Images page appears.

2. Under **Image Source**, click **VM Disk**.

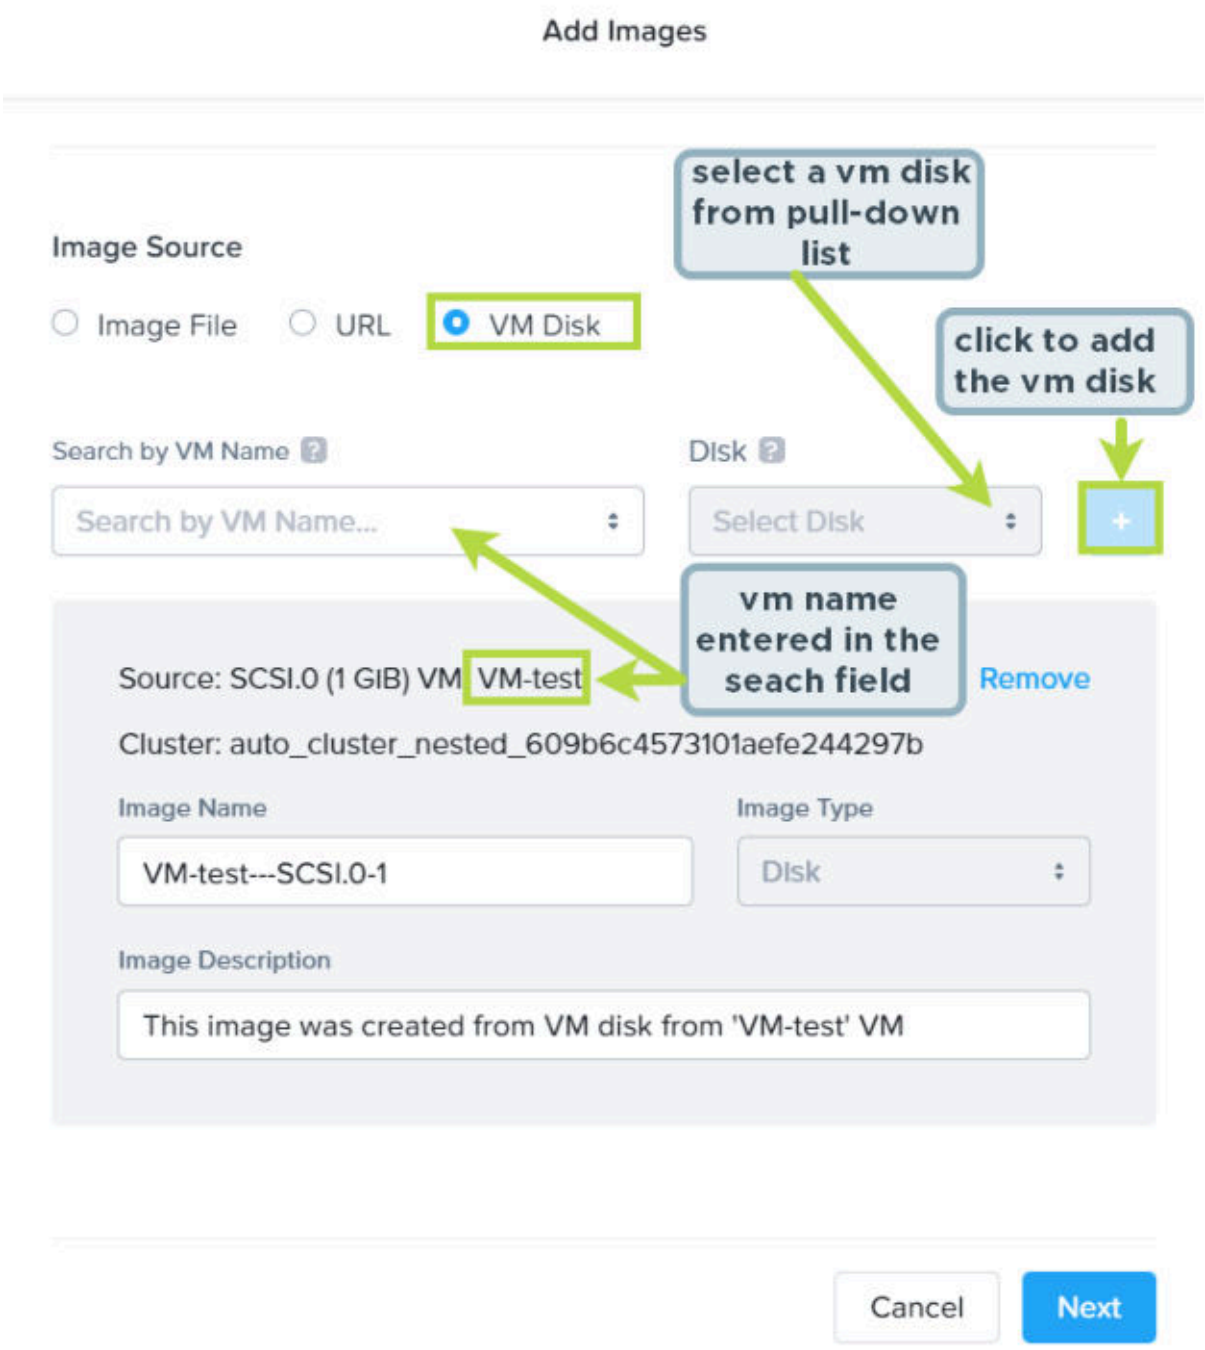

Figure 336: Add Images: VM Disk

3. To specify the VM disk, do the following:

   a. In **Search by VM Name**, click this field to populate the list of VMs from which you intend to clone the disk. The list displays all the powered-off VMs on the cluster. Alternatively, you can enter the VM name from which you intend to clone the disk.
   Prism Central adds the VM to a list on the page and populates the attached disks to this VM in the **Disks** pull-down list.

   b. In **Disk**, select the disk from the pull-down list and click the **+** button.

      The name defaults to the name of the file you selected, but you can change that name here.

      Prism Central displays the disk, and cluster information of the VM source.

4. To specify the image information, do the following:

   a. In **Image Name**, you can optionally change the auto-generated image name that defaults to the combination of the selected VM and disk name.

   b. In **Image Type**, Prism Central displays the type as ***Disk*** by default. You cannot modify this field.

   c. In **Image Description**, enter a description for the image.

5. To remove an image file entry, locate the entry and click **Remove**.

6. After you have added all the VM disks, click **Next**.
   The **Select Location** page appears.

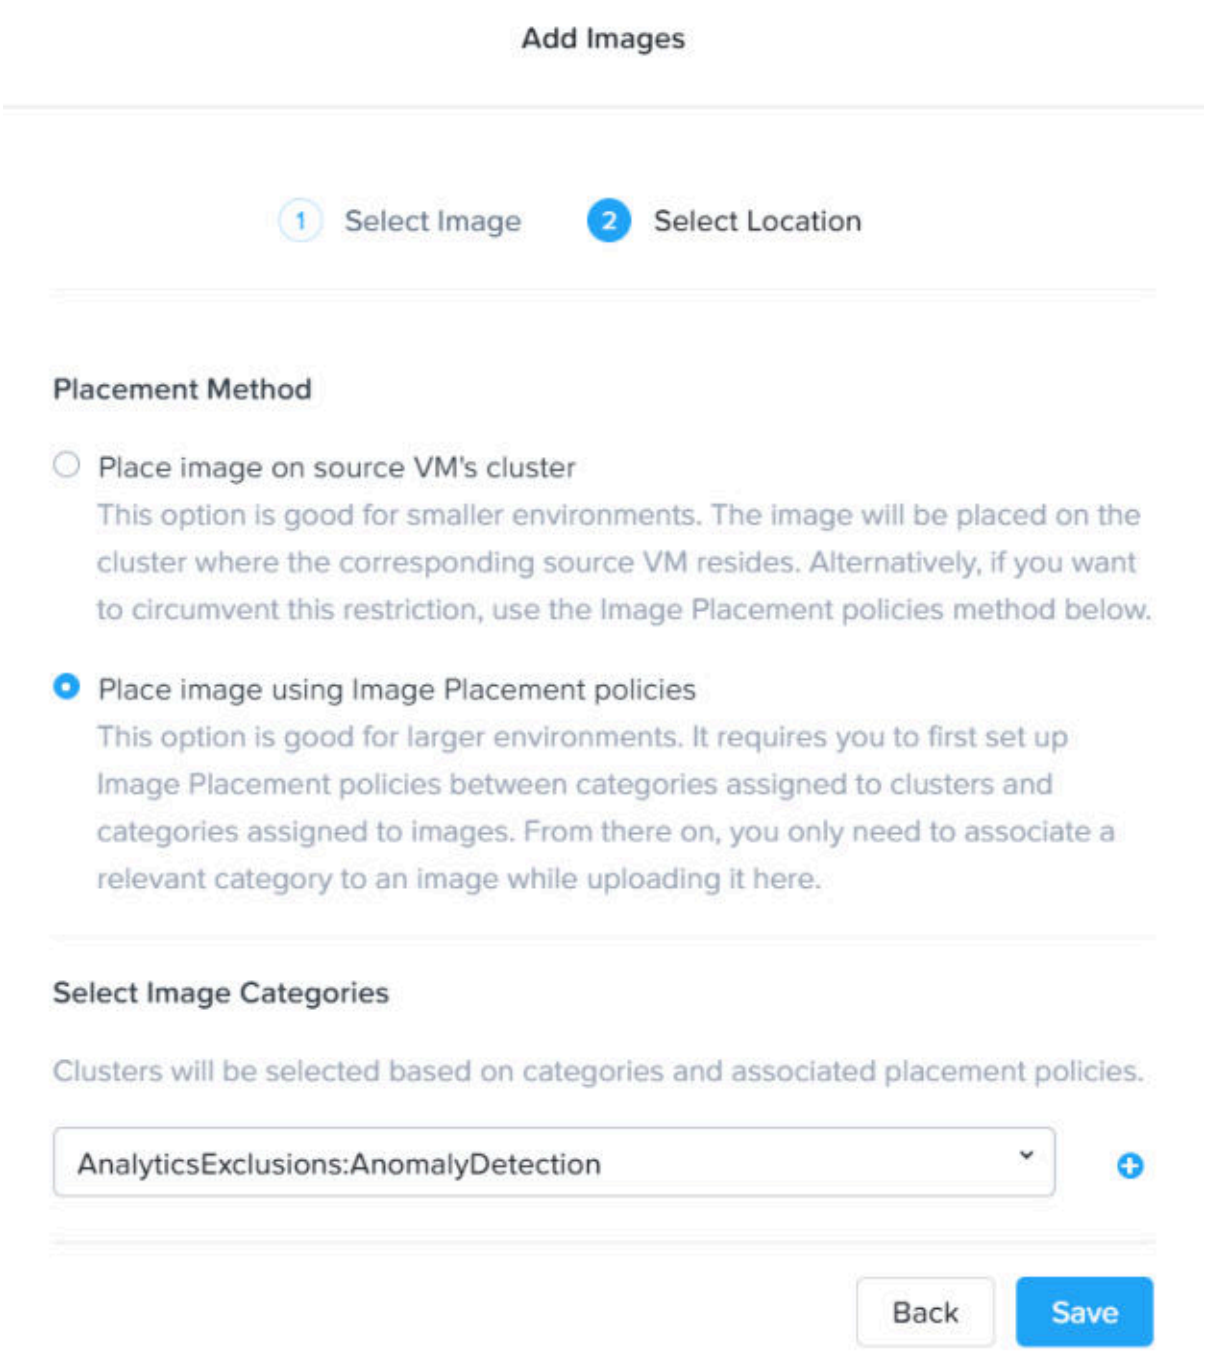

Figure 337: Add Images: Select Location

7. In **Placement Method**, do one of the following:

   » To select the source VM cluster, click **Place image on source VMs cluster** .

   » To delegate image placement decisions to configured policies and assign categories to the images, click **Place image using Image Placement policies**, and then do the following in the **Categories** section.

8. If you selected **Place image using Image Placement policies**, do the following:
   a. To assign a category, click the **Select Image Categories** search box.
   b. Select the desired category from the list.
      You can also start typing the name of the category to reduce the list to matching names.
   c. Click the add icon beside the search box.
   d. To specify another category, repeat these steps.
9. Click **Save**.
   Prism Central uploads the image files in batches.

Image Upload Verification

Verify whether an image upload to multiple clusters succeeded.

When you add an image and select multiple clusters for initial placement of the image (using the **Place image directly on clusters**), the status message box shows that the task succeeded. The status message displays the successful completion of the task even if the image is successfully uploaded to only one cluster. Therefore, go to **Tasks** and check whether the Image upload and Image update subtasks succeeded for all the clusters that you selected. It is possible that the Image upload or update subtasks for some of the clusters failed.

Nutanix recommends that you use Categories and Image Placement Policy based image placement option when you add images to multiple clusters. Image Placement policy based placement ensures that images propagate successfully to multiple clusters in a single operation. See Creating an Image on page 662 and Associating Hosts with Categories on page 681 for more details about adding an image.

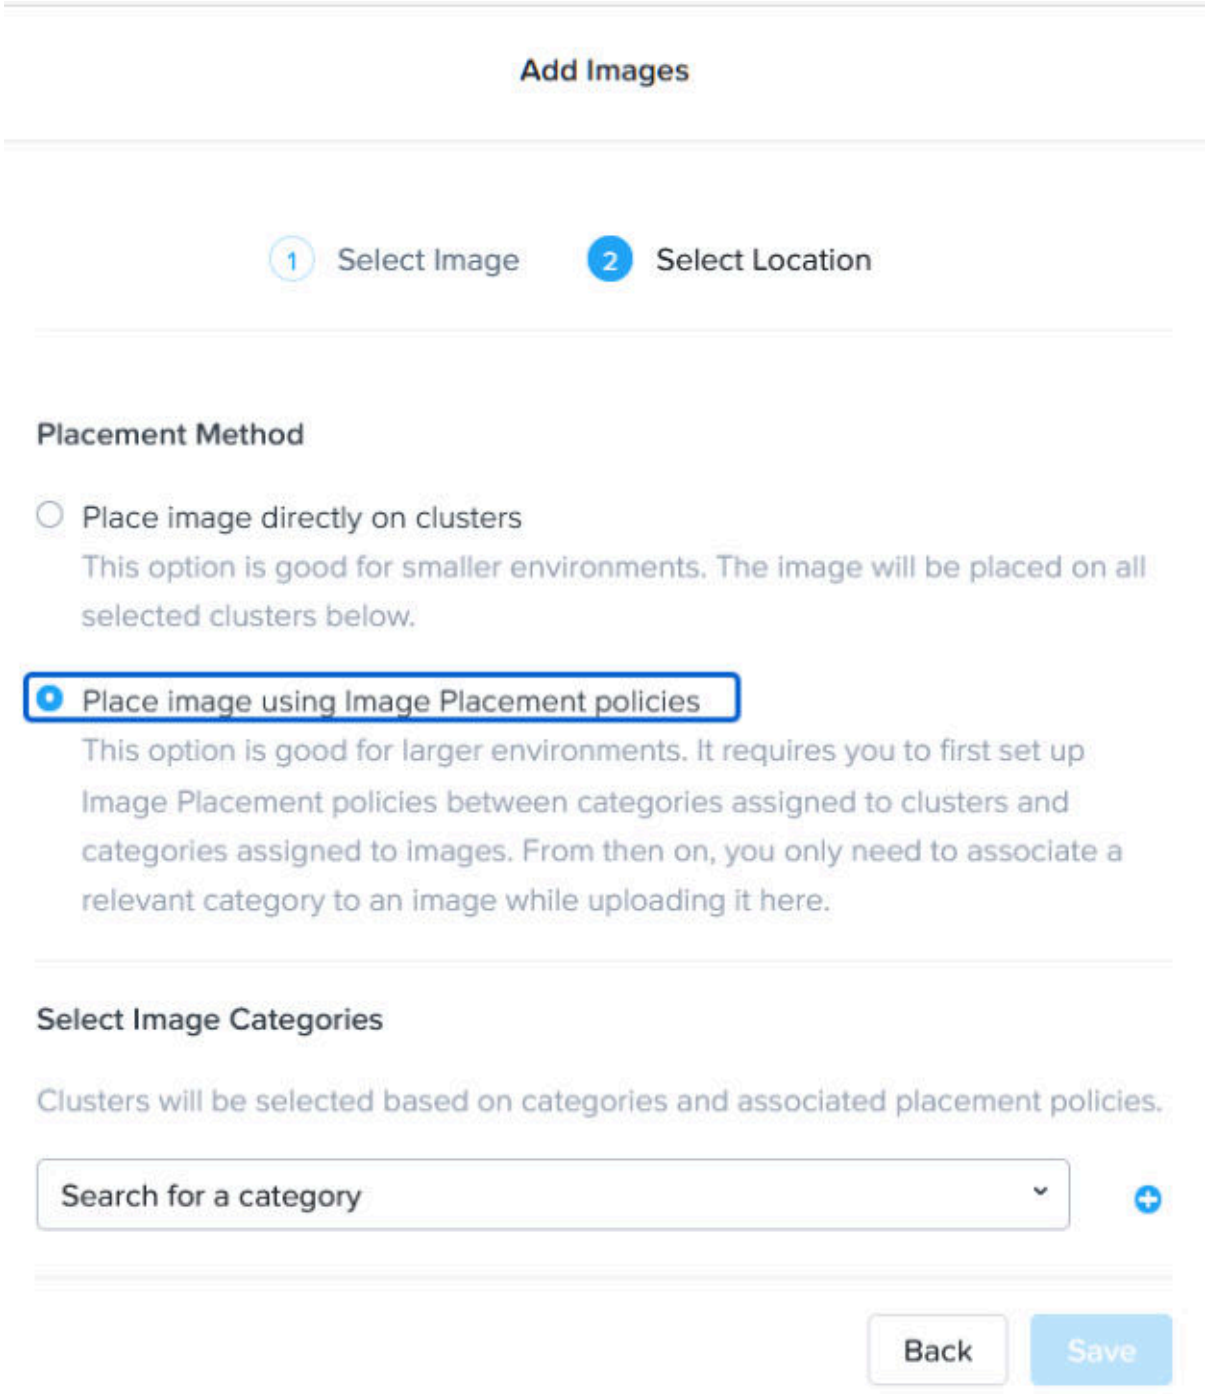

Figure 338: Image Placement Method selection for multiple clusters

Modifying an Image

About this task

To modify an image (delete, update, or add to catalog), in the **Entities** menu, go to **Compute & Storage** > **Images** (see Images Summary View on page 117), 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.

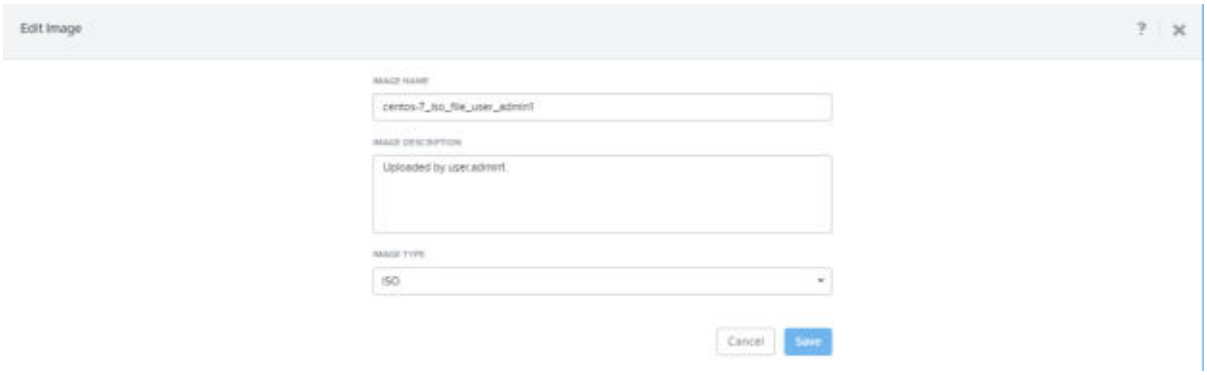

Figure 339: Update Image

- To add an image to the catalog (see Catalog Management on page 657), 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 **Compute & Storage** > **Images**, and then click the **Import Images** button (see Images Summary View on page 117).

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.

Image Policy Management

In Prism Central, you can create and manage the following types of policies for the images:

- Image Placement Policies (see Image Placement Policies on page 680)
- Bandwidth Throttling Policies (see Bandwidth Throttling Policies on page 685

Image Placement Policies

Prism Central enables you to upload VM images from your workstation or from a remote server. It also enables you to manually specify a subset of registered clusters as targets for those images. However, image placement decisions are often driven by compliance requirements, policies, and regulations. For example, you might have to confine a VM image to a Nutanix cluster in a specific location. Or, a cluster in one region might require a Windows VM image with a specific set of applications while a cluster in another region might require a Linux VM with its own set of applications.

Prism Central enables you to configure policies that govern which clusters receive the images that you upload. These policies, called image placement policies, map images to target clusters using categories associated with both those entities. For example, an image placement policy states that the Linux images associated with the `OS:Linux` and `distribution:Centos` categories must be uploaded to the clusters associated with the `Location:SFO` category.

Image placement policies also specify how strictly the policy must be enforced. Soft enforcement allows clusters not identified by the policy to use the images, when required, by checking them out. Hard enforcement disallows such usage. By using both soft and hard enforcement on the same category of clusters, you can address more specific use cases. For more information, see Sample Scenarios.

Image placement policies maintain a state of compliance across the clusters registered with Prism Central. They detect and correct violations introduced by changes in category assignments and updates to their configuration (including a switch from soft enforcement to hard enforcement). For more information, see Sample Scenarios.

To view the location of, the category and Image Placement policy associated with an image, see Images Summary View on page 117 and Image Details View on page 119.

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 VMs 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 Hosts with Categories

About this task

To associate categories with an Image, do the following:

Procedure

1. In Prism Central, in the Entities menu, go to **Compute & Storage** > **Images**.
2. Select the image that you want to associate with categories.
3. In the **Actions** menu that is displayed, click **Manage Categories**.
4. In the **Search for a category** field, type the name of the category or value that you want to add, and then select the category and value from the list that is displayed.
5. Click the add button to add as many search fields and repeat this step for each category that you want to add. To remove a category, click the remove button beside the field.
6. Click **Save**.

Configuring an Image Placement Policy

About this task

To configure an image placement policy, do the following:

Before you begin

Complete the following tasks:

- Configure the categories that you need to associate with the images.
- Configure categories for the clusters to which you want Prism Central to upload images, and then associate the categories with the clusters.
- Review existing image placement policies to avoid causing conflicts. For information about issues caused by conflicting policies, see How Prism Central Handles Conflicting Policies.

Procedure

1. In the entity menu, go to **Policies** > **Image Placement Policies**, and then click **Create Image Placement Policy**.
2. On the **Create Image Placement Policy**, enter a name and description for the policy.
3. Click inside the **Assign Images From The Following Categories** search field and select the category you want to associate with the images.

   To associate additional categories with the images, click the add icon beside the search field, and select the categories you need.
4. Click inside the **To The Clusters From The Following Categories** search field and select the category you associated with the target clusters.

   To specify additional cluster categories, click the add icon beside the search field, and select the categories you need.
5. In the **Policy Enforcement** section, from the **Enforcement** list, select one of the following:

   » **Soft**. The images are uploaded to the set of clusters identified by the image placement policy. In addition, no restriction is placed on using those images on clusters outside the identified set. For example, if clusters A, B, and C are registered with Prism Central, and an image placement policy with a soft enforcement policy uploads an image to clusters A and B, you can use the image to create a VM on cluster C. The checkout to cluster C is not blocked.

   » **Hard**. The images are uploaded to the set of clusters identified by the image placement policy. You cannot use the VM on any cluster outside the identified set. For example, if clusters A, B, and C are registered with Prism Central, and an image placement policy with a hard enforcement policy uploads an image to clusters A and B, the image is available only on clusters A and B. Cluster C cannot check out the image.

   For various scenarios involving these enforcement policies, see Sample Scenarios.
6. Click **Save**.

Updating an Image Placement Policy

You can update an Image Placement policy. Updates to an image placement policy can result in a policy violation. Prism Central attempts to correct the violation by executing a series of actions.

About this task

For example, if you change the cluster category such that the policy now identifies a different set of clusters as targets, Prism Central first copies the images to the clusters in the new category, and then removes the images from the clusters in the previously mentioned category. If you change the policy enforcement setting from Soft to Hard, and if any clusters that are not identified as targets had previously checked out the images, Prism Central corrects the violation by deleting the images from those clusters. An update to a policy might also result in a conflict with one or more of the other policies that apply to the same category of images. For more information about how Prism Central handles conflicting policies, see How Prism Central Handles Conflicting Policies.

To update an image placement policy, do the following:

Procedure

1. In the entity menu, go to **Policies** > **Image Placement Policies**, and then click the image placement policy that you want to update.
2. From the **Actions** menu that is displayed, click **Update**.
3. The **Edit Image Placement Policy** page that is displayed includes the same options and settings as the **Create Image Placement Policy** page.
4. Update the settings that you want, and then click **Save**.

How Prism Central Handles Conflicting Policies

Policies are said to be in conflict if they identify the same category of images but their resulting cluster placements are different.

For example, the following policies are conflicting policies if they apply to the same category of images:

- Policy P1, which identifies clusters C1 and C2 as target clusters, and the policy enforcement setting is Soft.
- Policy P2, which identifies clusters C3 and C4 as target clusters, and the policy enforcement setting is Hard.

Prism Central handles such conflicts in the following manner:

- Prism Central ignores conflicting policies and does not perform any action on the images. If policy P1 was configured first, Prism Central stops enforcing P1 as soon as P2 is configured, and it does not enforce P2 either. When there is a policy conflict, the conflict is noted in the **Policies** tab of the details view for that image.
- Prism Central continues to enforce other policies that are not in conflict. For example, if policies P3 and P4 apply to the same categories of images as P1 and P2, Prism Central continues to enforce P3 and P4 as long as they are not in conflict with P1 and P2.

Sample Scenarios and Configurations

The following table describes how to configure image management policies to achieve a desired result. The examples use a Prism Central instance to which three clusters A, B, and C and registered (possibly among other clusters).

Table 713: 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.
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.

Desired Result	Image Placement Policy Configuration
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.

Bandwidth Throttling Policies

Prism Central uses the bandwidth throttling policy feature to manage the usage of bandwidth when you create a new image or transfer an image from one cluster to another cluster using Prism Central. The bandwidth throttling policy limits the bandwidth consumed during image creation using the URL option in the specific clusters (see Uploading Images from a Remote Server on page 667 for more information about URL option). Without bandwidth throttling policy, an image creation from the remote server consumes as much bandwidth as available. High bandwidth consumption for creating a new image limits the bandwidth availablity for the other cluster operations.

Requirements for Bandwith Throttling Policy

The bandwidth throttling policy feature in Prism Central has the following requirements:

- The bandwidth throttling policy feature works on Prism Central version 2021.9 and the clusters must have AOS version version 6.0 or later.
- The clusters must be registered with Prism Central and must have AHV installed.

> Note: The bandwidth throttling policy is not supported on ESXi and Hyper-V hypervisors.

Limitations with Bandwidth Throttling Policy

The bandwidth throttling policy has the following limitations.

- Prism Central enforces bandwidth throttling policy only for new images created by using the URL option on clusters registered with Prism Central.
- You cannot enforce a bandwidth throttling policy if you create an image directly from a Prism Element.
- The least applicable bandwidth throttling policy on the source and destination cluster is enforced when you transfer an image from one cluster to another cluster.

Creating Bandwidth Throttling Policy

Create a new bandwidth policy and associate it with a cluster. This ensures that the new bandwidth throttling policy is implemented on the cluster when it creates or adds a new image.

About this task

Perform the following procedure to create a new bandwidth throttling policy.

Procedure

1. In the **Entity** menu, go to **Compute & Storage** > **Images** > **Policies** > **Bandwidth Throttling Policies**, and then click **Create Policy**.

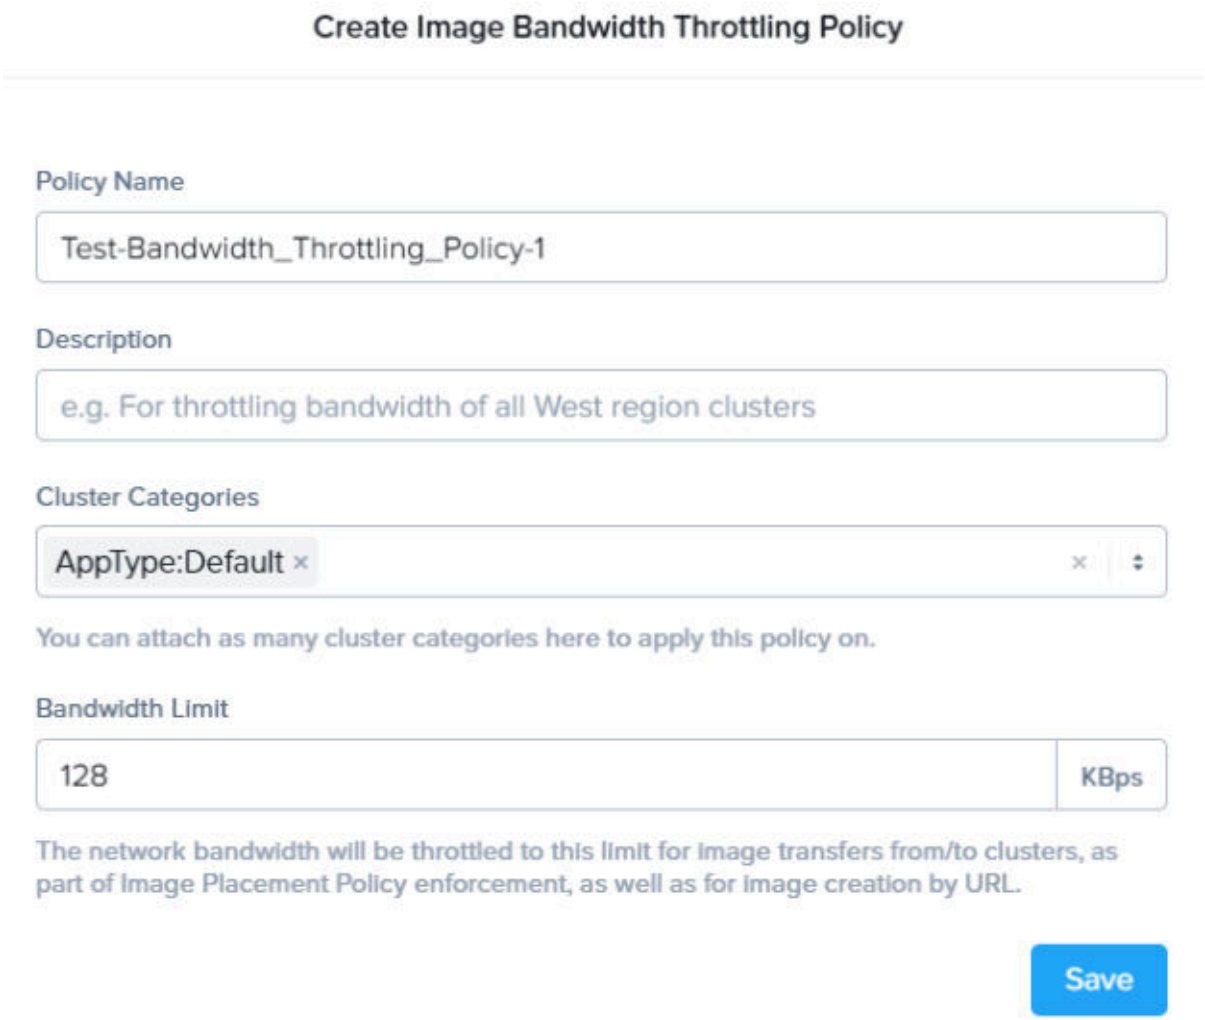

Figure 340: Bandwidth Throttling Policy

2. In the **Policy Name** text box, type a name for the bandwidth throttling policy.

3. In the **Cluster Categories** drop-down box, select the categories for which you want to apply this new bandwidth throttling policy.

   For more information about Categories, see Category Management on page 735.

   > Note: In a scenario, where you have to add multiple cluster categories in the same policy, then the bandwidth throttling policy is enforced on the Prism Element that has all the specified categories.

4. In the **Bandwidth Limit** text box, enter the bandwidth limit for use with this bandwidth throttling policy.

   The network bandwidth is throttled to the bandwidth limit set in this field during the new image creation using the URL option, and for image transfer from a cluster to another cluster as part of image placement policy enforcement.

   > Note: The bandwidth limit task is serialized on each cluster to strictly enforce the bandwidth limit on each cluster. For example, when you create a new image, only one image create task runs on the cluster at a time. The second image creation task configured with another bandwidth limit can start only after completion of the first image creation task.

5. Click **Save** to create a new bandwidth throttling policy for the selected category.
6. Add a cluster and associate it with the same categories used with the bandwidth throttle policy.

   For more information, see Associating Hosts with Categories on page 681.

7. In the **Entity** menu, go to **Compute & Storage** > **Images** > **Policies** > **Bandwidth Throttling Policies** to view the list of bandwidth throttling policies.
8. Click on a bandwidth throttling policy to view the cluster associated with it.

   The selected cluster uses the **Effective Bandwidth Limit** configured to add a new image next time.

   > Note: In a scenario where multiple bandwidth throttling policies are applied to the same cluster, the minimum of all the applicable bandwidth throttling policies is enforced on the cluster.
   >
   > For example, you have a cluster named PE-123, and you have created multiple bandwidth throttling policies with different **Bandwidth Limit**, then during the cluster association with the bandwidth throttling policy (step 6), the cluster is associated with the bandwidth throttling policy which has the least **Bandwidth Limit** configured on it.

Updating Bandwidth Throttling Policy

You can update the bandwidth throttling limit on a cluster to ensure faster or more controlled image downloads using the URL operation.

About this task

Perform the following procedure to update the bandwidth limit on a cluster.

Procedure

1. In the entity menu, go to **Compute & Storage** > **Images** > **Policies** > **Bandwidth Throttling Policies** to view the list of bandwidth throttling policies.

2. Click on a bandwidth throttling policy in the table to view the clusters associated with the bandwidth throttling policy.
3. Click **Update** to modify the bandwidth throttling policy.
   You can update the **Policy Name**, add additional **Cluster Categories** or delete the selected **Cluster Categories**, and change the **Bandwidth Limit**.

   > Note: For image create operation, you cannot update the bandwidth limit during image creation process.

   > Note: For image copy operation, you can update the bandwidth limit during the image transfer process.

4. Click **Save** to commit the changes.

Deleting Bandwidth Throttling Policy

Deleting a bandwidth throttling policy removes the bandwidth limit enforced on an image associated with cluster and category.

About this task

Perform the following procedure to delete the bandwidth throttling policy on a cluster.

Procedure

1. In the entity menu, go to **Compute & Storage** > **Images** > **Policies** > **Bandwidth Throttling Policies** to view the list of bandwidth throttling policies.
2. To delete a bandwidth throttling policy, select the checkbox associated with policy and select **Actions** > **Delete**.
   Select the checkbox in the table header to delete all the bandwidth throttling policies configred.
3. Click **Delete** in the **Delete Policy ?** dialog to delete the policy.

OVA Management

You can export a VM as an OVA, upload and manage OVAs only from Prism Central.

Of the following OVA operations, you can export a VM only from the VMs dashboard (see VMs Summary View on page 86). You can perform the other actions from the OVAs dashboard (see OVAs View on page 113).

- To export a VM in the VM dashboard as an OVA , see Exporting a VM as an OVA on page 587.
- To upload an OVA, see Uploading an OVA on page 690.
- To deploy an OVA file as a VM, see Deploying an OVA as VM on page 698.
- To download an OVA file to your local machine, see Downloading an OVA on page 704.
- To rename an OVA file, see Renaming an OVA on page 704.
- To delete an OVA file, see Deleting an OVAs on page 705 .

To perform any action an OVA in the OVAs dashboard, select the target OVA, then select the desired action in the **Actions** menu.

The task progress for these actions are displayed in the Tasks page.

> Note:
>
> The **Actions** menu appears only if you select an OVA in the list. If the OVAs list is empty, then the **Actions** menu is available only after you create an OVA by uploading it or exporting an existing VM as an OVA.

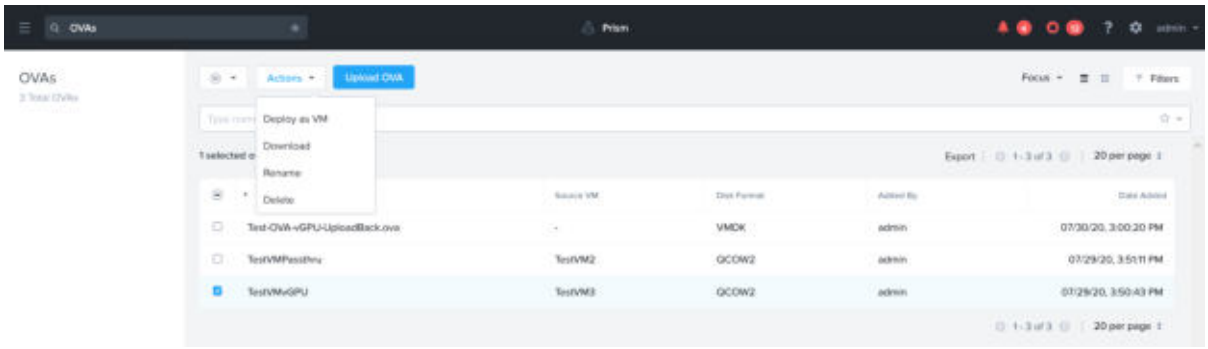

Figure 341: Actions menu

Exporting a VM as an OVA

Before you begin

If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the destination hypervisor.

About this task

As a Super Admin, you can export a VM as an OVA by selecting the **Export** as OVA action from the **Actions** dropdown list. You can export a VM as an OVA with a specified disk format for the disks configured on the VM.

> Note:
>
> - The Super Admin role was formerly called as User Admin.
> - If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the target destination hypervisor.

To export a VM as an OVA, do the following in the VMs dashboard:

Procedure

1. Select the VM in List tab.

   You can click the VM and open the VM details page.

2. In the **Actions** list, select **Export as OVA**.

   Alternatively, on the details page of the VM, select **More** and click **Export as OVA**.

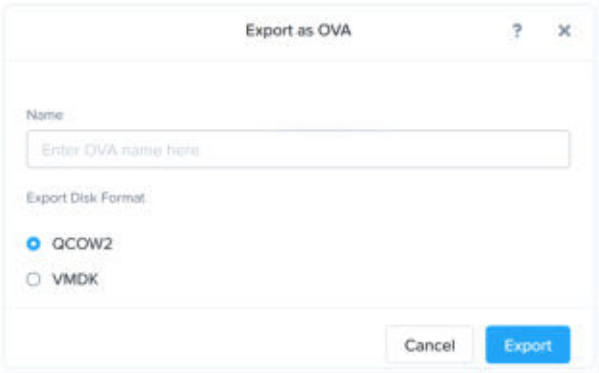

Figure 342: Export as OVA dialog box

3. In the Export as OVA window, do the following in the indicated fields:

   a. Enter a name for the OVA, in the **Name** field.

   b. Select the disk format that you want to format the disks in the VM to, in the **Export Disk Format** section.

      The default format is QCOW2. If you want to use the OVA to deploy a VM with disks formatted as VMDK, then select VMDK as the disk format.

4. Click **Export**.

   You can check the progress of the *Export as OVA* task in Tasks.

   The exported OVA is available on the OVA dashboard.

Uploading an OVA

About this task

You can upload an OVA from the following sources:

- From a local folder: When you upload an OVA from a local source, you can upload it to only one target cluster.
- From a URL: When you upload an OVA from a URL, you can upload it to multiple clusters. The upload operation runs concurrently on all the selected clusters.

Select the appropriate upload source considering the multiple-cluster upload option.

> Note:
>
> Prism Central processes up to three OVA file uploads concurrently for OVA file uploads from local folders. Prism Central queues other OVA files that you upload until one or more of the concurrent uploads are completed.
>
> The is no restriction on the number of concurrent uploads for OVA uploads from URL.

For OVA upload from local folders:

- The upload process can be interrupted due to the following reasons:
  - Network issues when the upload is running.
  - You close the browser window or tab in which the upload is running.
- The concatenate and validate operations of the OVA upload operation fail when an error occurs in the upload services because the checksum provided for the OVA file when it was created is incorrect.

  > Note:
  >
  > If the upload process is interrupted due to any reason, you can resume the process. When the validation operation is interrupted due to a concatenation error then you must run the concatenate API to resume concatenation and validation. See Concatenating Upload using APIs on page 696.

- You can also check the status of the upload in Tasks. Prism Central displays the following three sub-tasks in Tasks:
  - OVA Create
  - OVA Upload
  - OVA Validate
- Prism Central uses indicators to display the status of OVAs that you may have tried to upload. The status message is displayed when you hover the mouse or pointing device cursor on the indicator. See *OVA Alerts* table in OVAs View on page 113.

To upload an OVA, do the following:

Procedure

1. To upload an OVA, click **Upload OVA**.

2. In the Upload OVA dialog box, do the following:

   a. Select the **OVA Source**. As an example, select *OVA File* to upload from a local folder.

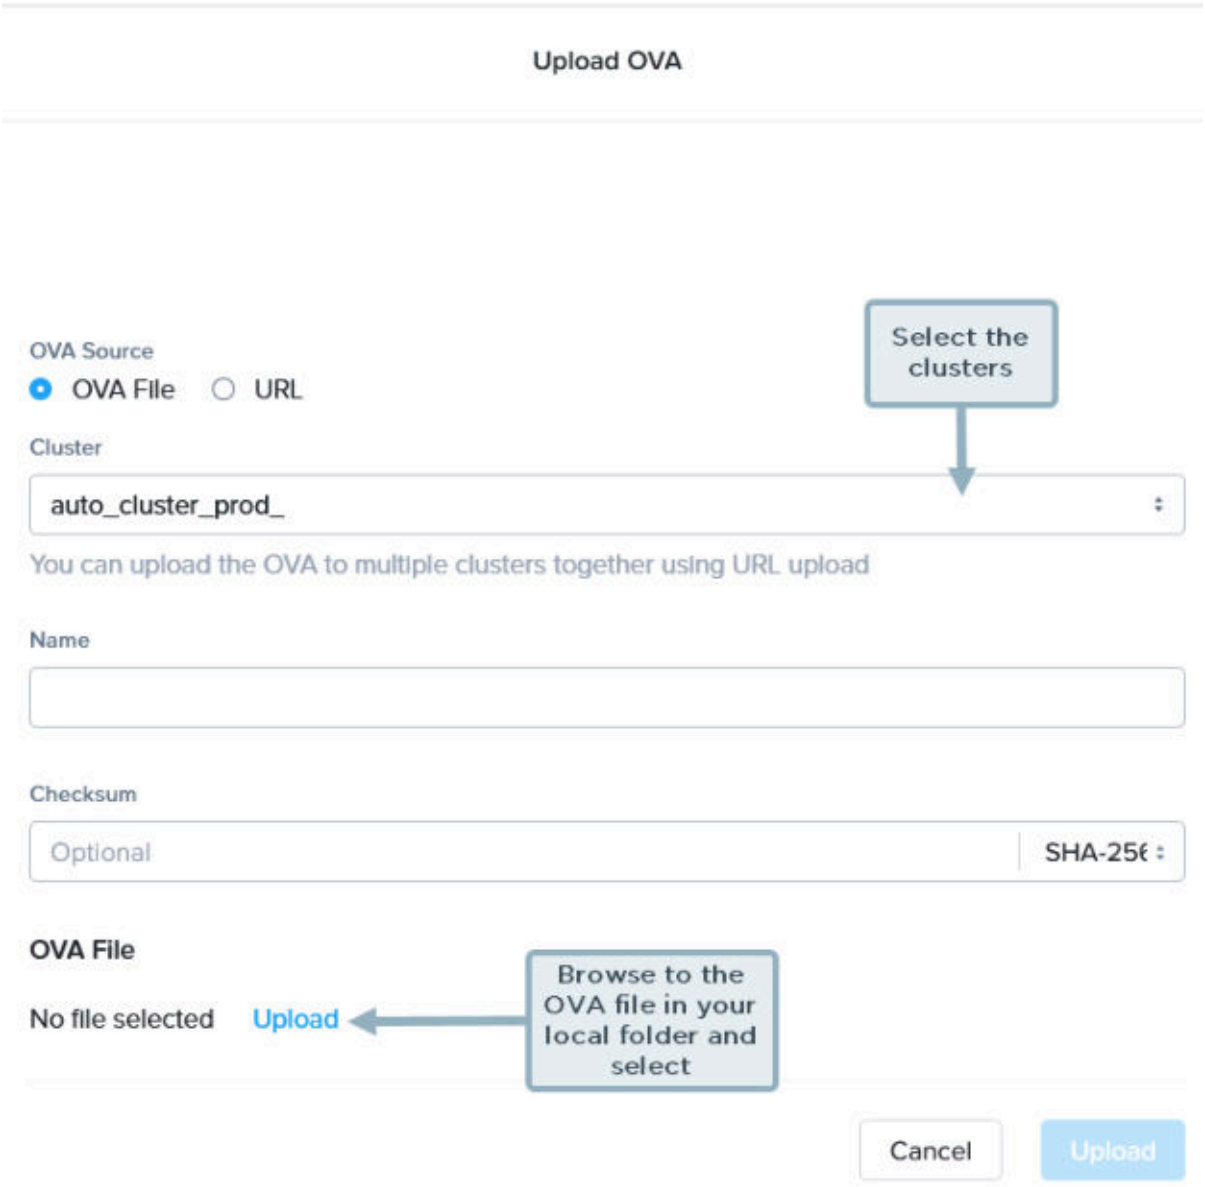

Figure 343: Upload OVA dialog box

   b. One cluster is already selected in the **Cluster** field.

      If you choose *URL* as **OVA Source** then you can select multiple clusters from the **Cluster** dropdown list that is displayed.

   c. Enter a **Name** for the OVA you are uploading.

      By default, Prism uses the file name of the OVA file that you upload, in the **Name** field if you do not enter a name and directly select the file for upload.

   d. Optionally, enter the **Checksum** value of the OVA file after selecting the encryption as SHA-1 or SHA-256 hash algorithm.

e. For **OVA File**, click **Select File** to navigate to the location of the OVA file in your local folder and open it.

   In case of upload from URL, in the **OVA URL** field, enter the source URL from where you want to upload the OVA file.

Figure 344: OVA URL field

> Note:
>
> Ensure that the URL you provide is not redirected to another URL. The upload fails for redirected URLs.

f. Click **Upload**.

3. The OVA upload starts. Wait until the **Continue in Background** button is displayed.

Note: The **Continue in Background** button is not displayed for uploads from URL.

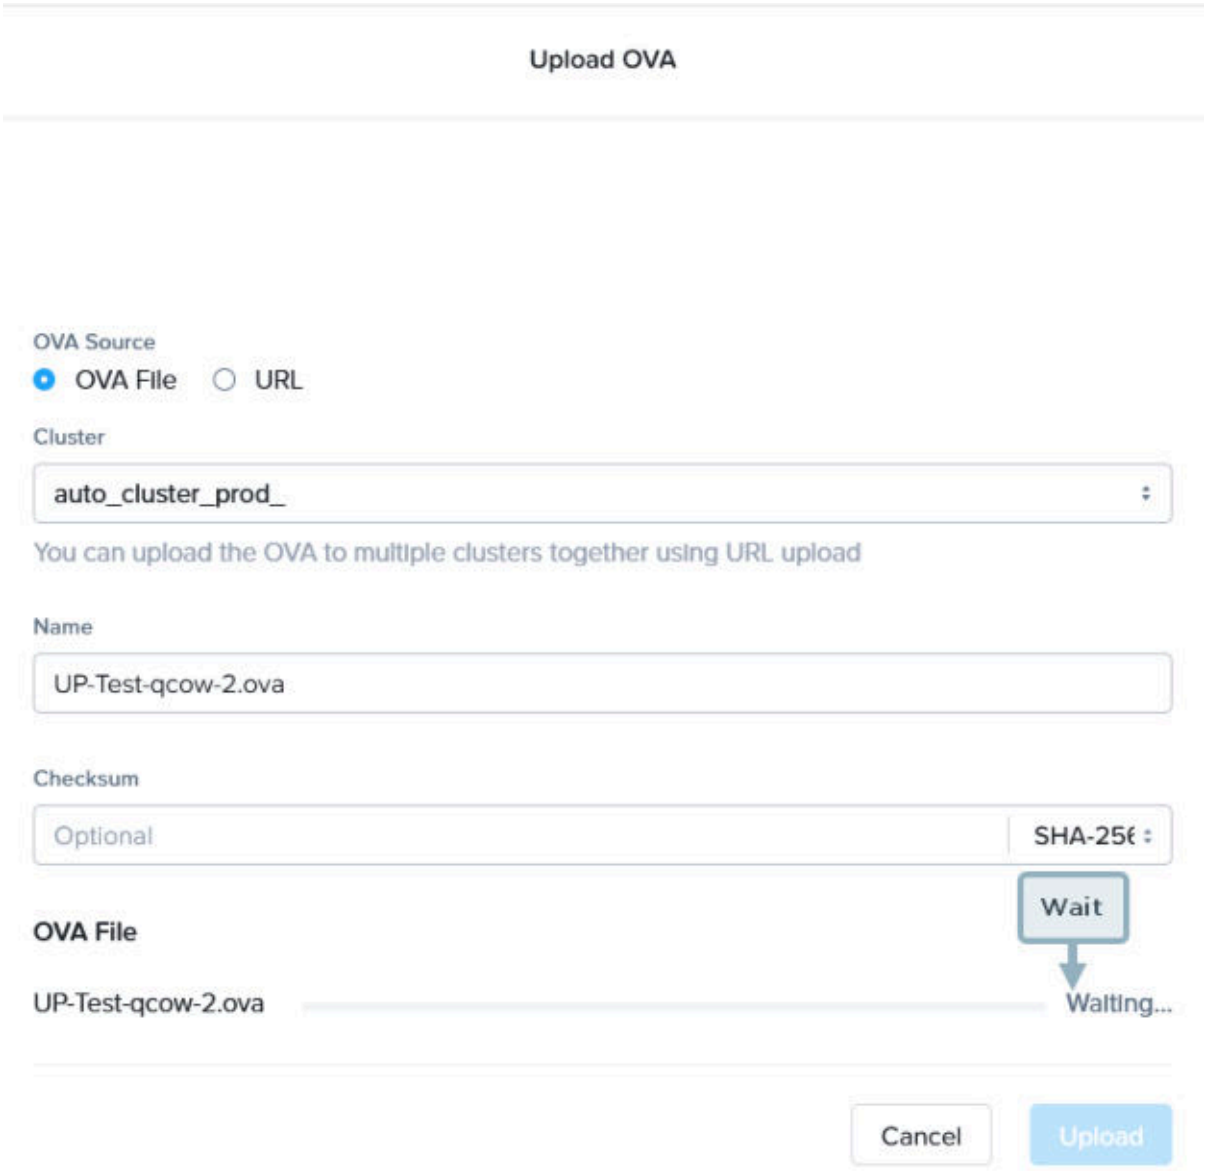

Figure 345: Upload OVA - Waiting

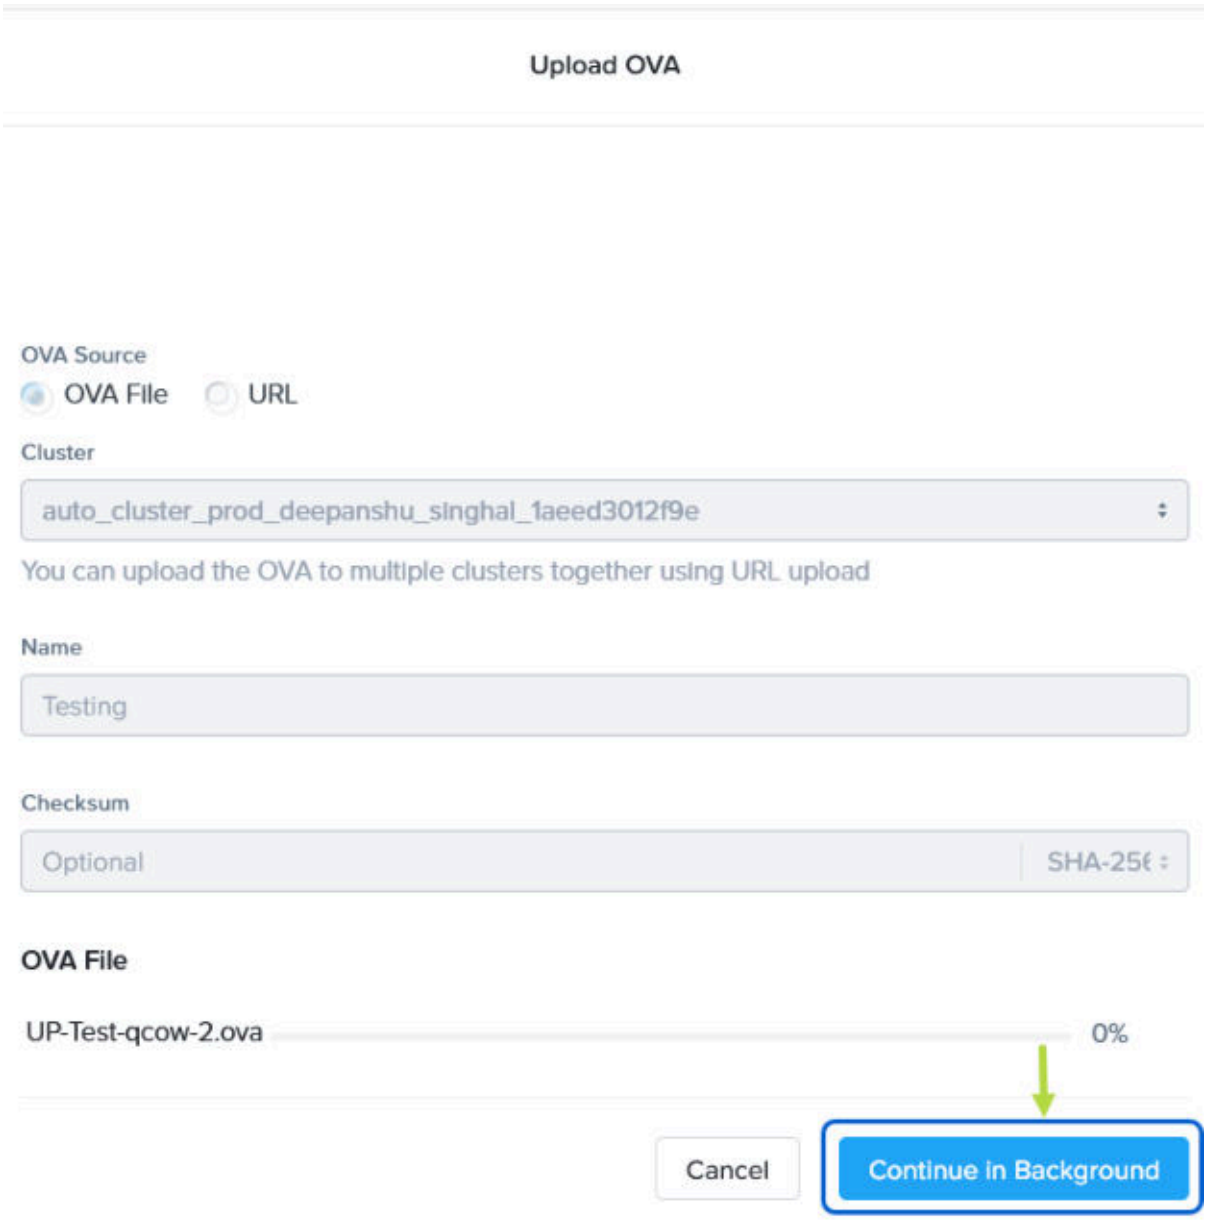

Figure 346: Continue in Background button

4. When the upload task is successfully completed, the following message is displayed:

```
Upload of the file is successfully completed. Verification of file is progressing in
 background. Please track it from Tasks page.
```

5. To resume an upload process that was interrupted, select the OVA in the OVAs dashboard. In the **Actions** menu and click the **Resume Upload** option.

   The status message that appears when you hover the mouse over the status icons for each OVA indicates whether you can use the **Resume Upload** option to resume the upload.

   Note:

   Check the reason for failure of **OVA Validate** task in the Tasks page. When the upload process is interrupted because of a concatenation and validation failure due

an incorrect checksum, the **Resume Upload** option is not available. Upload the OVA again in case the upload process was interrupted due to a checksum error.

Concatenating Upload using APIs

About this task

During an OVA upload, the OVA file is chunked and the chunks are uploaded. After upload, during the validation phase, the chunks are concatenated and validated.

When the OVA upload process is interrupted because of a concatenation error, you need to run the concatenation API to resume the upload process.

To run the API successfully, you need the UUID of the OVA for which the upload is interrupted. The UUID is derived using the name of the OVA.

Procedure

1. Click the user icon in the upper-right corner of the web console, and click **REST API Explorer**.

   The REST API Explorer displays a list of the objects that can be managed by the API.

2. In the REST API Explorer browser tab, select **ovas** in the list and click **List Operations**.

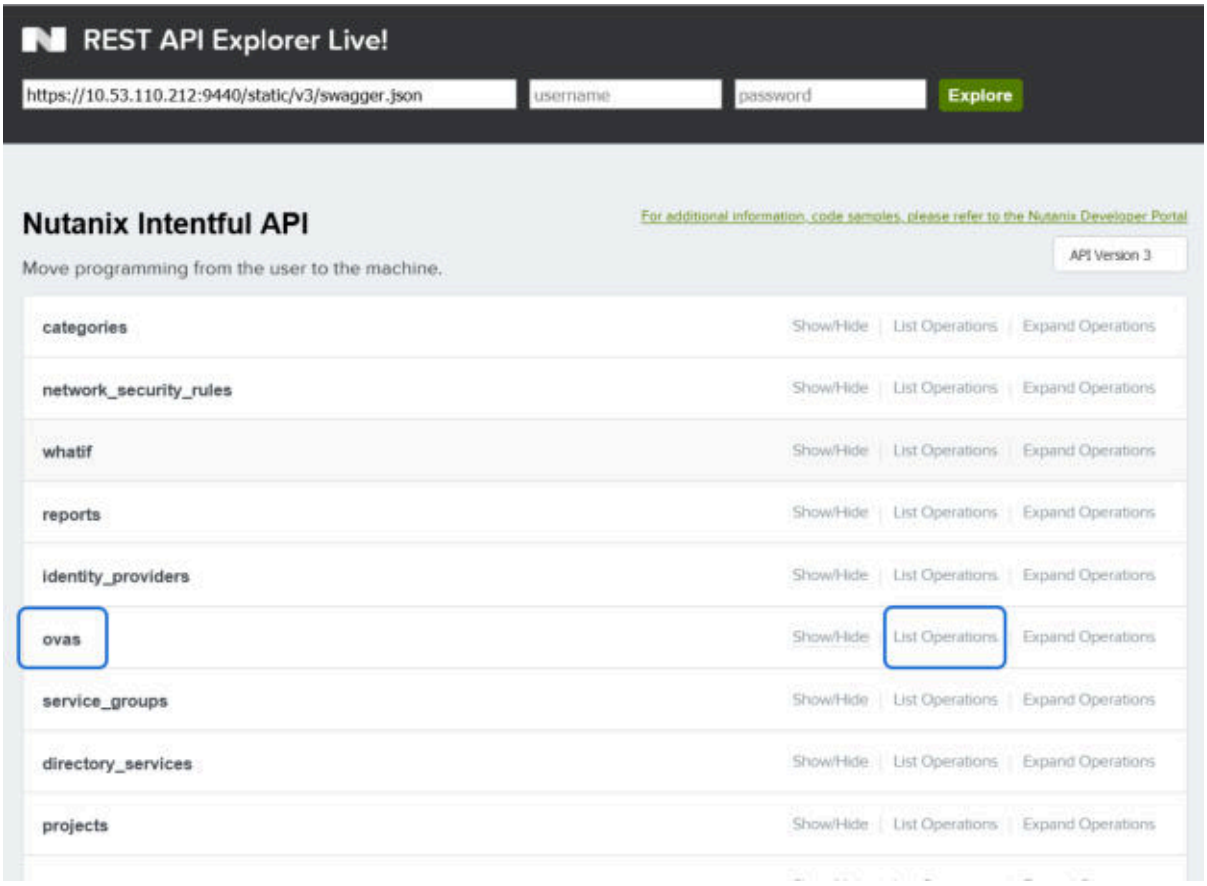

Figure 347: ovas List Operations

3. In the expanded **ovas** list of operations, click **Post /ova/list**.

   To use the concatenate API, you need the UUID of the OVA. **Post /ova/list** helps you filter the OVAs and get the details of the details of the interrupted OVA.

4. Enter the name of the OVA in the **get_entities_request** parameter field in the following format and click **Try it out!**:

```
{
  “filter”:”name==<name of ova>”
}
```

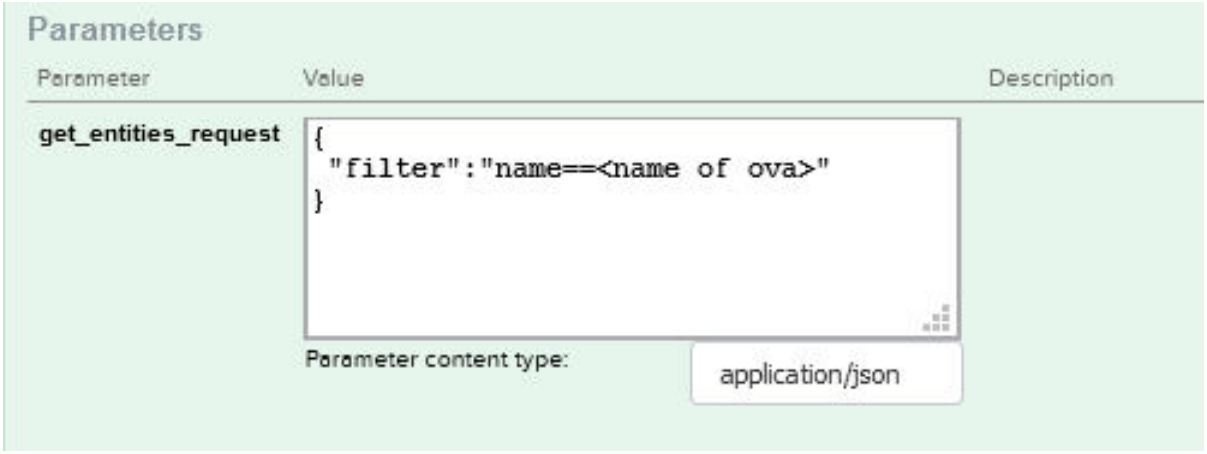

Figure 348: Parameter Filter

5. In the **Response Body** box, the UUID of the queried OVA is listed. Note or copy the UUID.

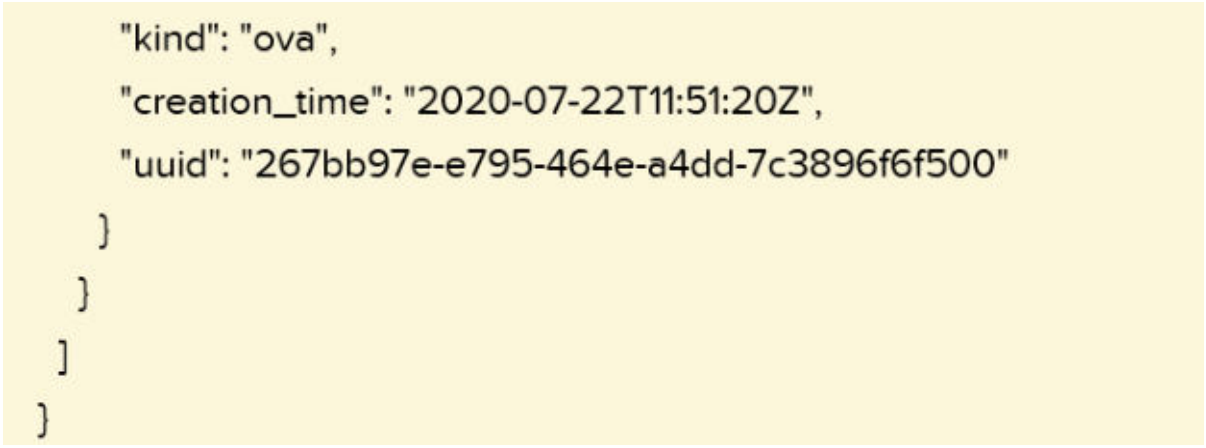

Figure 349: OVA Upload UUID

6. In the expanded **ovas** list of operations, click **Post /ova/{uuid}/chunks/concatenate** and enter the UUID of the OVA in the UUID parameter value field.

7. Click **Try it out!**.

   The **Response Code** field displays 202 indicating that the concatenation and validation request has been accepted.

Deploying an OVA as VM

About this task

To perform this task in the OVAs dashboard, select the target OVA. Some of the steps that you need to complete in this procedure are the same as the steps provided in Creating a VM (AHV) on page 536. The Deploy as VM dialog box has four tabs. Ensure that information in the first three tabs are appropriate for the VM you want to deploy. The Review tab helps you review the information in the first three tabs in one tab. You can navigate back and forth between the tabs using the **Back** and **Next** buttons.

> Note:
>
> In case of OVAs that are not created from AHV, ensure that the OVA has pre-installed VirtIO drivers.

To deploy an OVA as a VM, do the following:

Procedure

1. In the **Actions** menu, select **Deploy as VM**.

   You need to use the tabs in the Deploy as VM dialog box to configure the deployment. Use the **Back** and **Next** buttons to move between the tabs in the dialog box.

   > Note:
   >
   > The Actions menu is displayed only if you select an OVA in the list.

2. In the Deploy as VM dialog box, complete the indicated details in the Configuration tab:

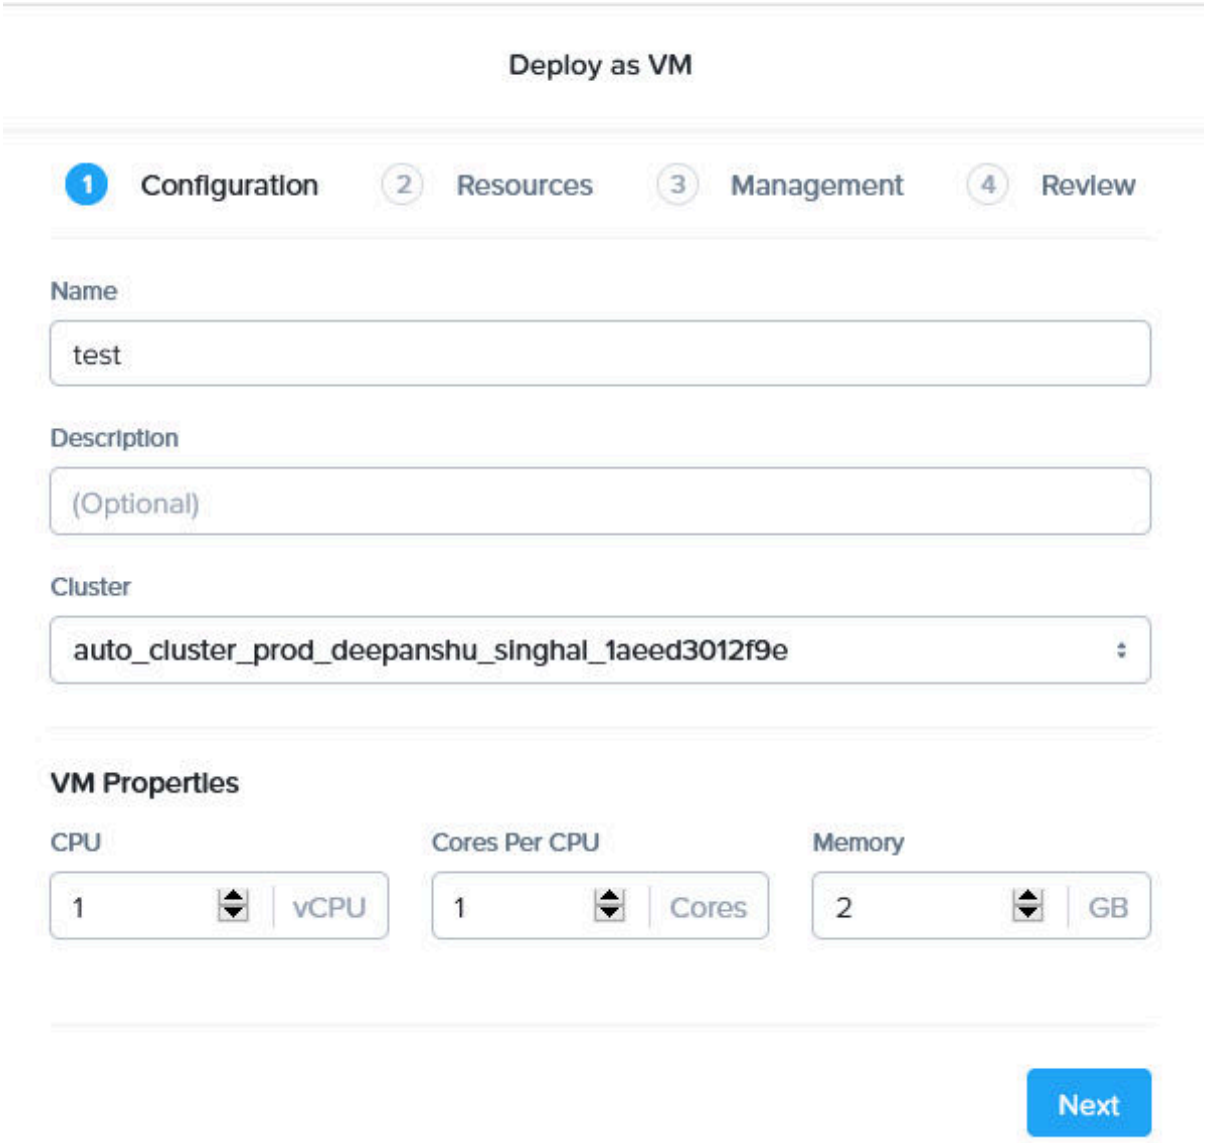

Figure 350: Deploy as VM dialog box

a. **Name** for the VM that you are deploying.

b. **Description** that are informative like Backup VM for Prism.

c. **Cluster** that you want to deploy the VM on. Select the cluster from the drop down list.

   Only clusters, where OVAs are available and you have access to, are displayed in the dropdown list. Prism Central also recommends a cluster in the list for the deployment.

d. The number of **CPU**, **Cores Per CPU** and **Memory** in **VM Properties**.

e. Select **Next**.

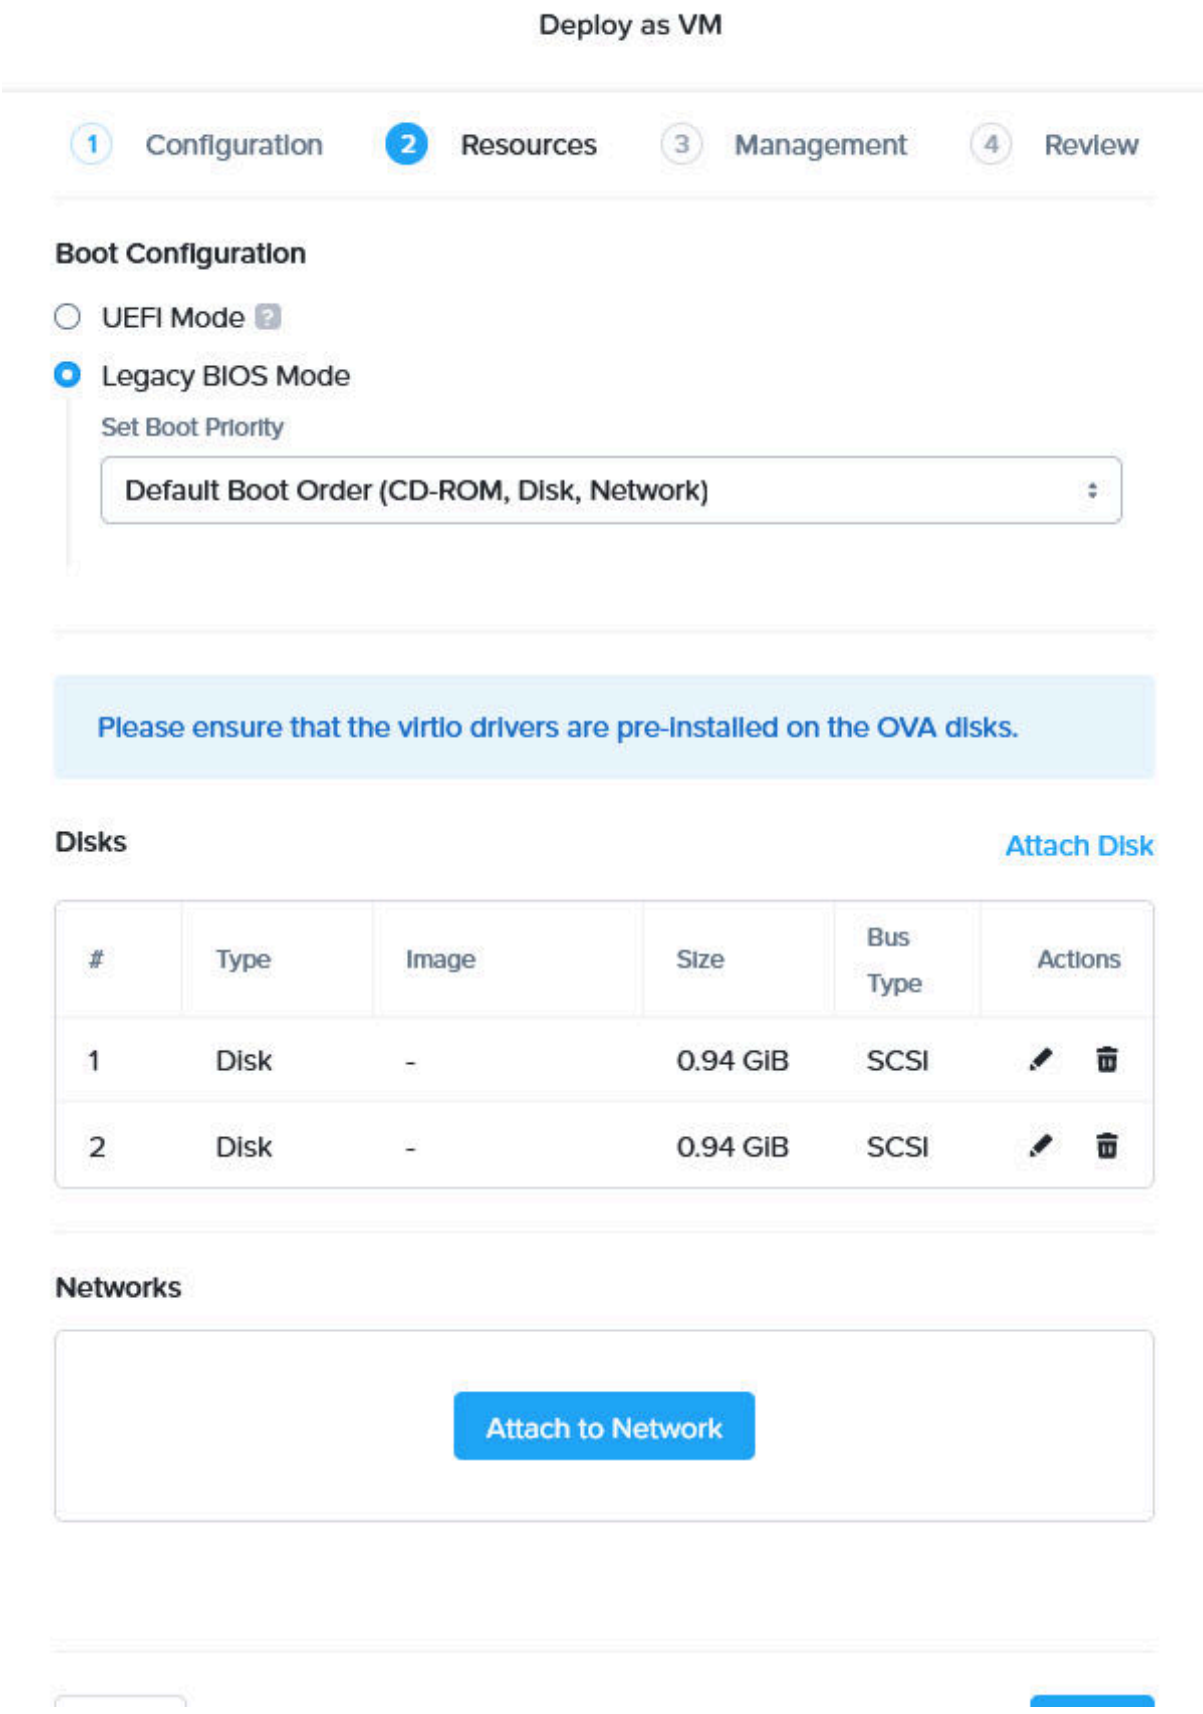

Figure 351: Deploy A VM - Resources tab

3. Complete the indicated details in the **Resources** tab.

   Note:

   You can add new disks, NICs and GPUs.

   a. **Boot Configuration** mode.
      Select **UEFI Mode** only if your hardware supports UEFI boot mode.

   b. **Set Boot Priority**.
      The default priority is already set.

   c. **Attach Disk** in addition to the disks already available in the OVA, if so required.

      In the Attach Disk dialog box, complete the indicated details and select **Save**.

      See Creating a VM (AHV) on page 536 for more information about attaching disks.

   d. Configure the network by selecting **Attach to Network**.

      In the Attach to Network dialog box, select the **Subnet** and **Network Connection State** from the drop down lists.

      See Creating a VM (AHV) on page 536 for more information about configuring networks.

   e. If the selected cluster has GPU resources, then you can add multiple vGPUs. Select **Add GPU**.

      See Creating a VM (AHV) on page 536 and Adding Multiple vGPUs to the Same VM on page 570 for more information about configuring networks..

      Note:

      You can add multiple vGPUs to the same VM. Addition of multiple vGPUs is based on the installed GPU resources. See *Multiple Virtual GPU Support*in *AHV Admin Guide*.

   f. Select **Next**.

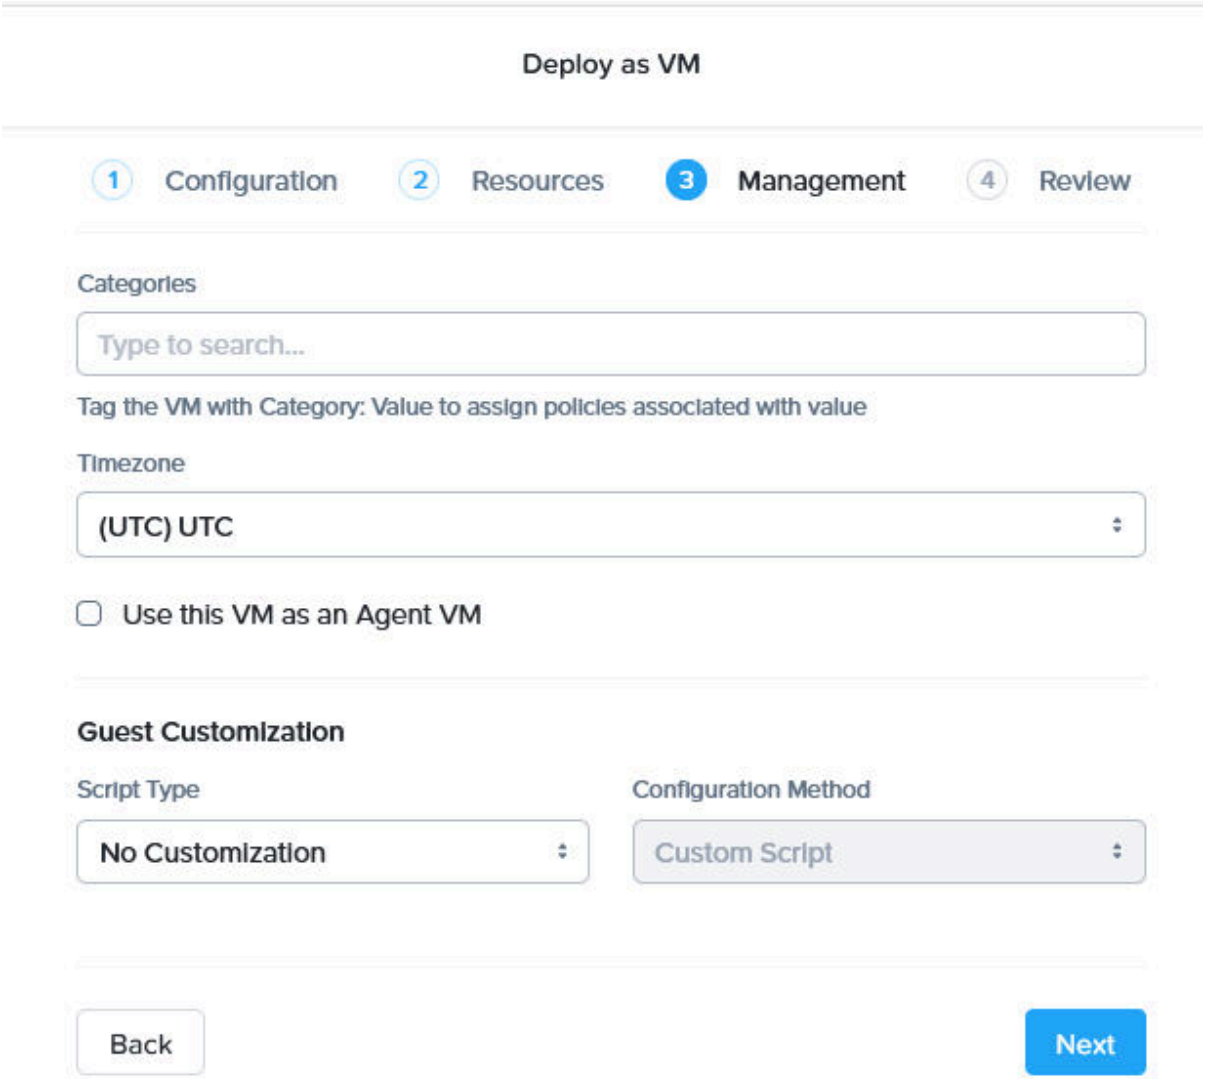

Figure 352: Deploy A VM - Management tab

g. Select the **Category**. Type in this field to display a list of available categories. For more information about categories, see Category Management on page 735.

h. Set the **Timezone**.

i. If you want to use the deployed VM as an agent VM, check the **Use this VM as an Agent VM** checkbox.

j. Configure **Guest Customizations** if you want to run scripts for configurations on Microsoft Windows or Linux VMs.

   You can choose from `No Customization`, `Sysprep (Windows)` and `Cloud-init (Linux)` as **Script Type**. You can choose from `Custom Script` or `Guided Script` as **Configuration Method**.

   If you choose `No Customization` as **Script Type**, then the **Configuration Method** field is grayed out.

k. Select **Next**.

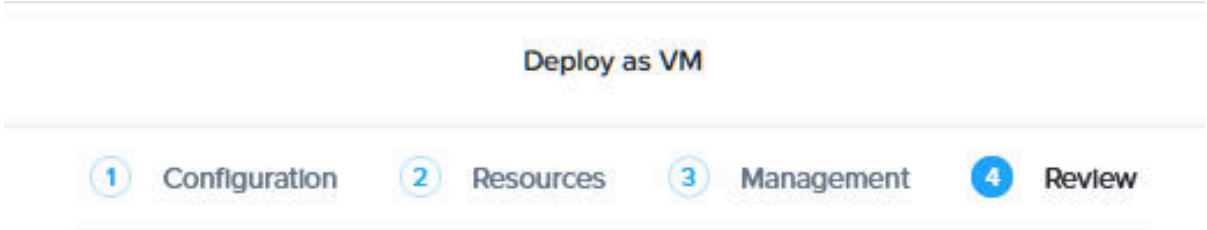
Deploy as VM
1 Configuration
2 Resources
3 Management
4 Review
Configuration
Name test
CPU 1 vCPU
Cores Per CPU 1 Cores
Memory 2 GB
Resources
Disks
Type Image Size Bus Type
1 Disk - 0.94 GiB SCSI
2 Disk - 0.94 GiB SCSI
Boot Configuration Default Boot Order (CD-ROM, Disk, Network)
Networks
Subnet Private IP
pc_network Auto-Assign
Management
Categories None
Timezone UTC
Guest Customization No Customization

4. Review the deployment configuration in the Review tab.
5. Click **Launch** to deploy the VM.

   You can check the deployment tasks in Tasks. The deployed VM is displayed in the VMs dashboard.

Downloading an OVA

About this task

To perform this task in the OVAs dashboard, select the target OVA.

> Note:
> You can download upto five OVAs in parallel.

To download an OVA, do the following:

Procedure

1. In the **Actions** menu, select **Download**.

   > Note:
   > The Actions menu is displayed only if you select an OVA in the list.

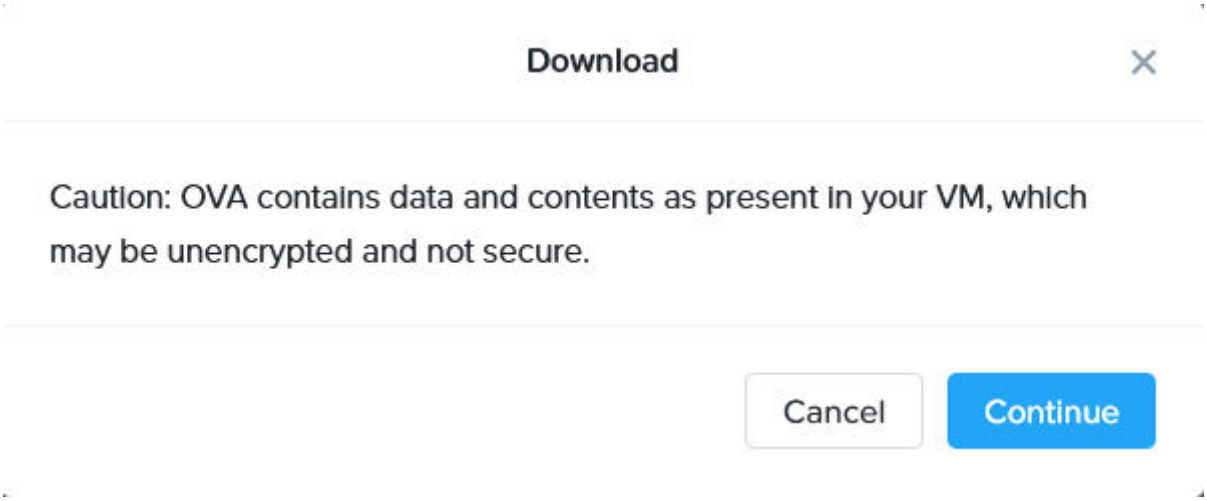

Figure 354: Download message box

2. To confirm the download action, select **Continue**.

   To cancel the download action, select **Cancel**.

   The downloaded file is saved in a local folder based on the download settings of your browser.

Renaming an OVA

About this task

To perform this task in the OVAs dashboard, select the target OVA.

To rename an OVA, do the following:

Procedure

1. In the **Actions** menu, select **Rename**.

   Note:
   The Actions menu is displayed only if you select an OVA in the list.

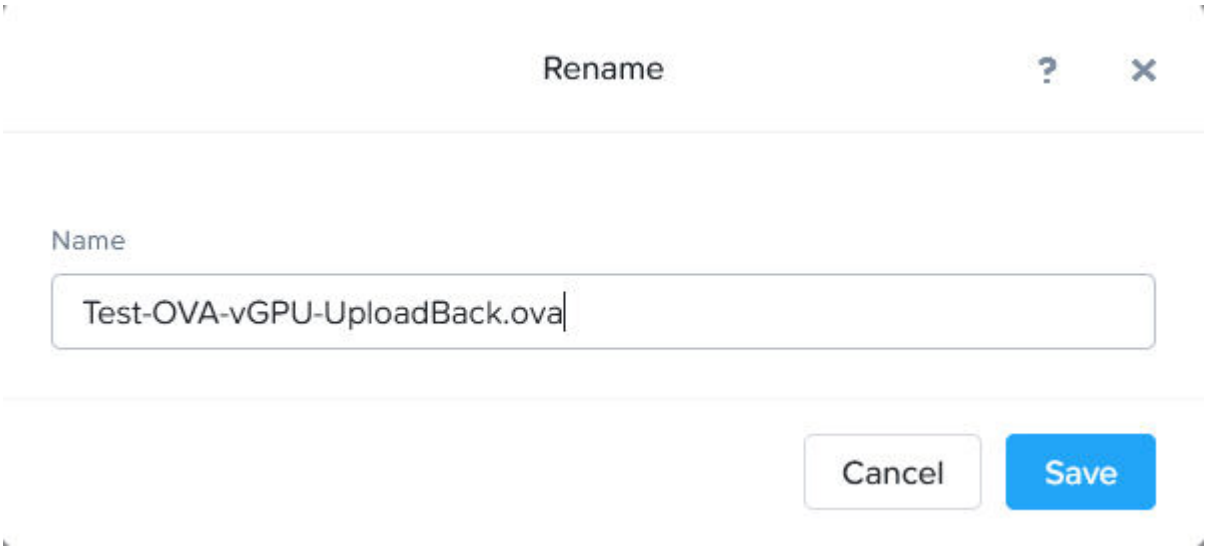

   Figure 355: Rename dialog box

2. In the **Name** field of the Rename dialog box, delete the current name and enter the new name for the OVA.
3. To confirm the rename action, select **Save**.
   To cancel the delete action, select **Cancel**.

Deleting an OVAs

About this task

To perform this task in the OVAs dashboard, select the target OVA.

To delete an OVA, do the following:

Procedure

1. In the **Actions** menu, select **Delete**.

   Note:
   The Actions menu is displayed only if you select an OVA in the list.

   The **Delete** action deletes the OVA permanently. You cannot recover a deleted OVA.
   Therefore, in the Delete dialog box, confirm if you want to delete the OVA <name>.

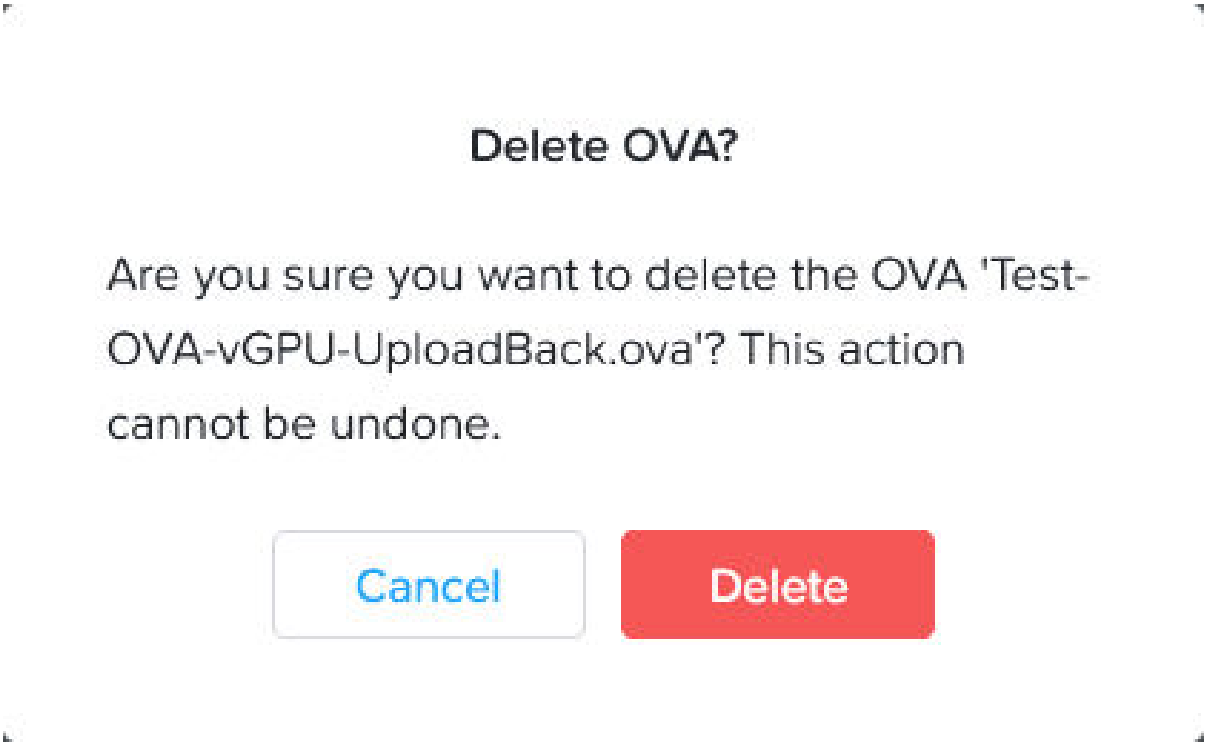

   Figure 356: Delete OVA dialog box

2. To confirm the delete action, select **Delete**.
   To cancel the delete action, select **Cancel**.

Performing Other Administrative Tasks

About this task

You can perform some administrative tasks for a selected cluster directly from Prism Central (see Compute and Storage (Cluster) Administration on page 535). 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 197).

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 714: User Menu List

Name	Description
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page.
About Nutanix	Opens a window that displays the AOS version running on the cluster.
Support Portal	Opens a new browser tab (or window) at the Nutanix support portal login page.
Help	Opens the online help.
Nutanix Next Community	Opens a new browser tab (or window) at the Nutanix Next Community entry page.
Sign Out	Logs out the user from the web console.

vCenter Server Integration

The VM management through Prism Central for ESXi provides a unified management interface for all of the vCenter Server instances. All the vCenter Server instances that are registered or not registered with any of the clusters are auto-discovered and displayed. If you have not registered the vCenter Server in Prism Element, you can register the clusters to the vCenter Server instances directly from Prism Central.

During the registration process, you have an option to select the ESXi clusters that you want to manage using Prism Central. For more information about registering vCenter Server, see Registering vCenter Server (Prism Central) on page 709. 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 712.

  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 16).

   - 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.

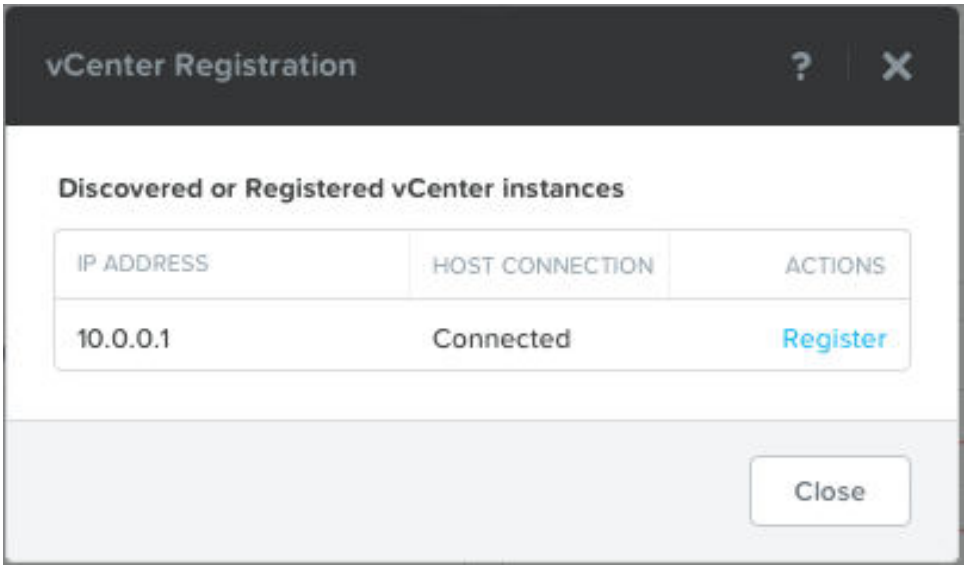

Figure 357: vCenter Registration

2. To register the vCenter Server, click **Register**.

   - The vCenter Server that is managing the hosts in the cluster is auto-discovered and its IP address is auto-populated in the **Address** field.
   - The port number field is also auto-populated with 443. Do not change the port number. For the complete list of required ports, see Port Reference.

3. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.

4. If you do not wish to manage all your ESXi clusters from Prism Central, select the check box of just the clusters that you want to manage and click **Save**.

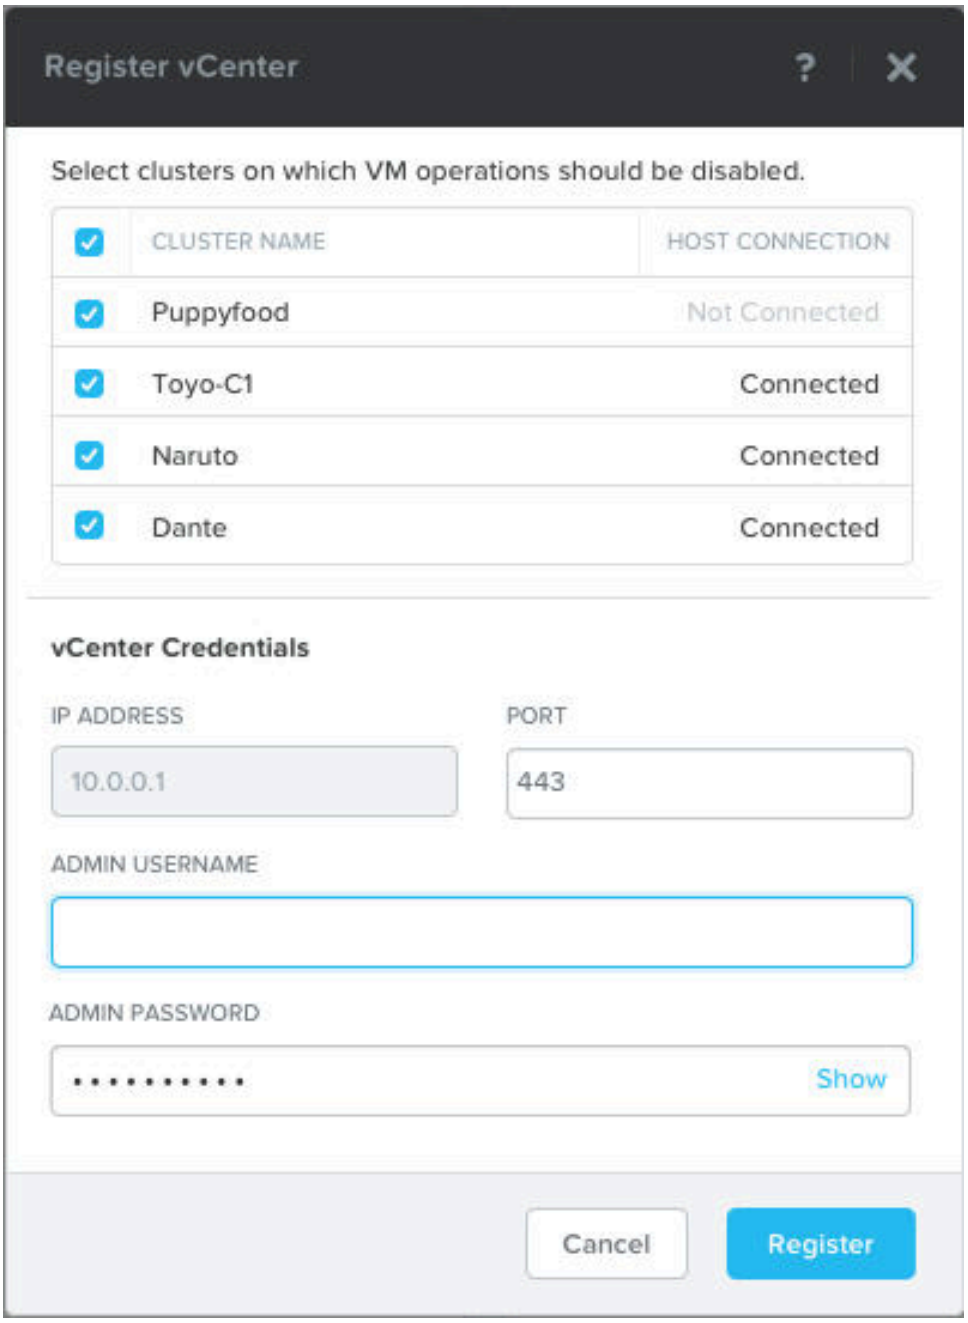

Figure 358: Cluster Selection

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 712.

Procedure

1. Click the gear icon in the main menu and select **vCenter Registration** from the Settings menu (see Settings Menu on page 16).

   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 16).

   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 712.

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 709.

9

PRISM SELF SERVICE ADMINISTRATION

The Prism Self Service feature allows you to create projects where consumers of IT infrastructure within an enterprise—individual users or teams such as development, test, and DevOps—can provision and manage VMs in a self-service manner, without having to engage IT in day-to-day operations.

> Note: This chapter describes how to configure and administer Prism Self Service. See the *Prism Self Service Administration Guide* for instructions on the role and functions of a self-service administrator. See the *Prism Self Service User Guide* for information about the self-service capabilities available to a project user.

Prism Self Service Overview

Prism Self Service represents a special view within Prism Central. While Prism Central enables infrastructure management across clusters, Prism Self Service allows end users to consume that infrastructure in a self-service manner. Prism Self Service uses the resources provided by a single AHV cluster. (Other hypervisors are not supported platforms for Prism Self Service.)

> Note: A SelfServiceContainer storage container (see Storage Containers Summary View on page 122) 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 714).
- Create a project for each team that needs self service and add users and groups to the projects (see Project Management on page 718).
- 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 657).
- 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 254 and Project Details View on page 256).
- Create VMs as needed and assign them to project members as appropriate (see VM Management on page 535). 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 717.

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).

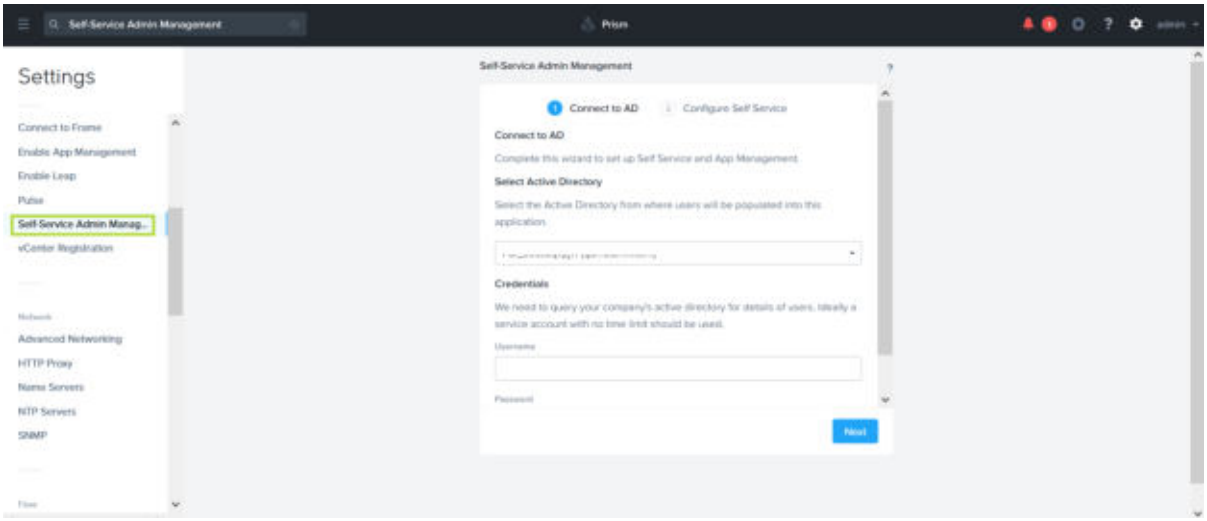

Figure 359: Self-Service Admin Management Window

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.

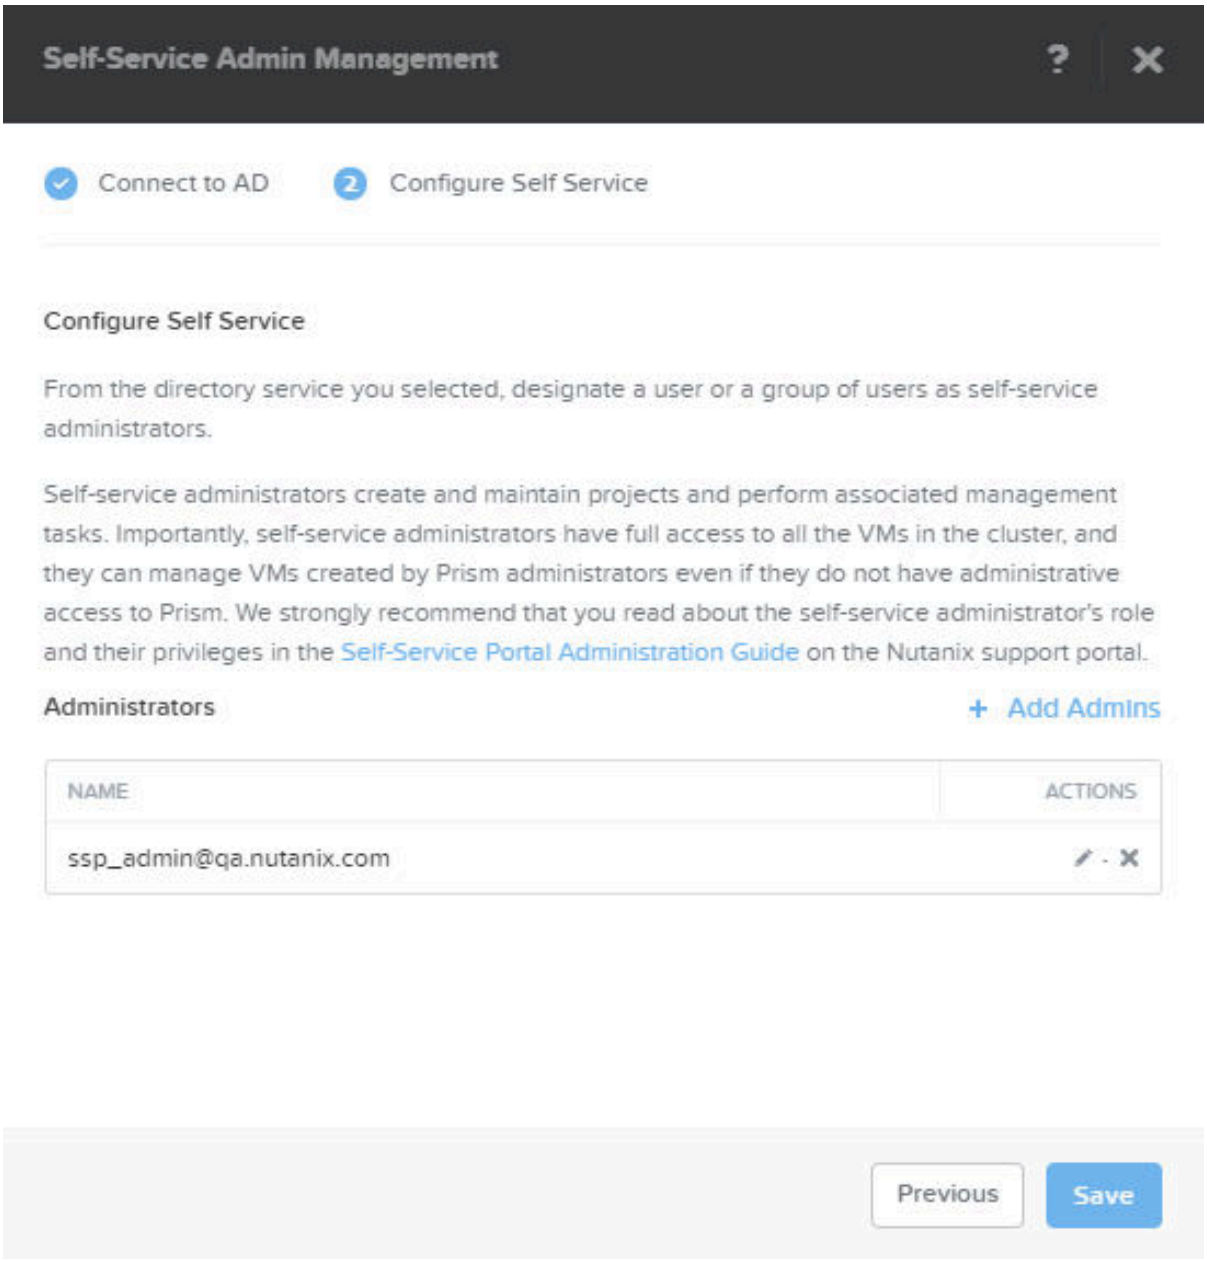

Figure 360: Configure Self Service tab

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.

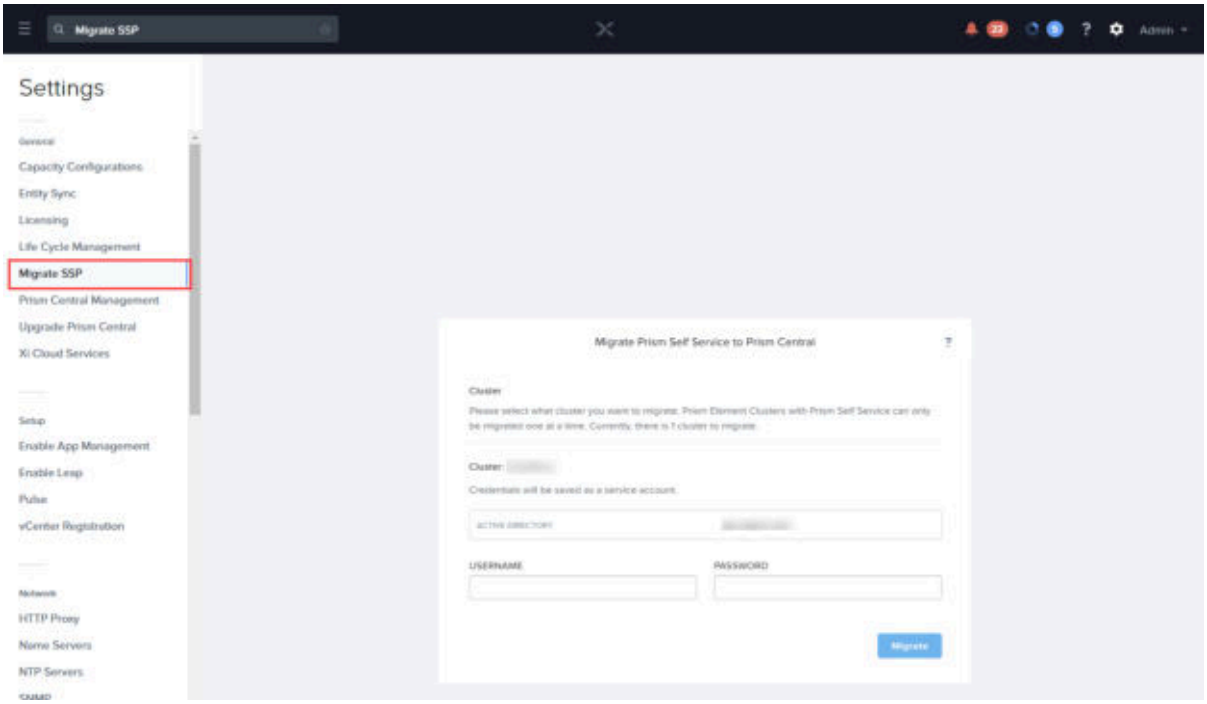

Figure 361: Migrate Prism Self Service to Prism Central Window

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 241).

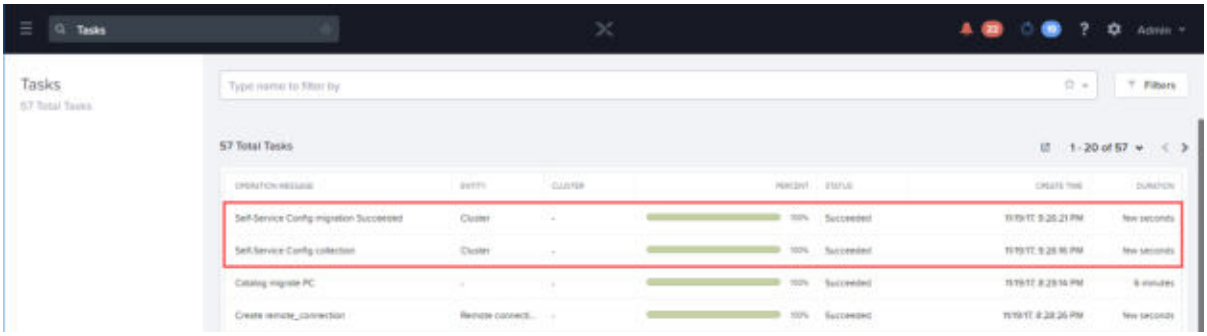

Figure 362: SSP Migration Tasks

Project Management

A project defines Active Directory users or groups to manage common set of requirements or functions. For example, a project can define a team collaborating on an engineering project. A project specifies roles to associate its members, select existing networks that the deployed VMs can use, and (optionally) set usage limits on infrastructure resources.

Creating a Project

Before you begin

Authentication must be configured before you can create a project (see “” ()).

About this task

To create a project, do the following:

Procedure

1. Go to the projects dashboard (see Projects Summary View on page 254) and click the **Create Project** button.
   The Create Project page appears.

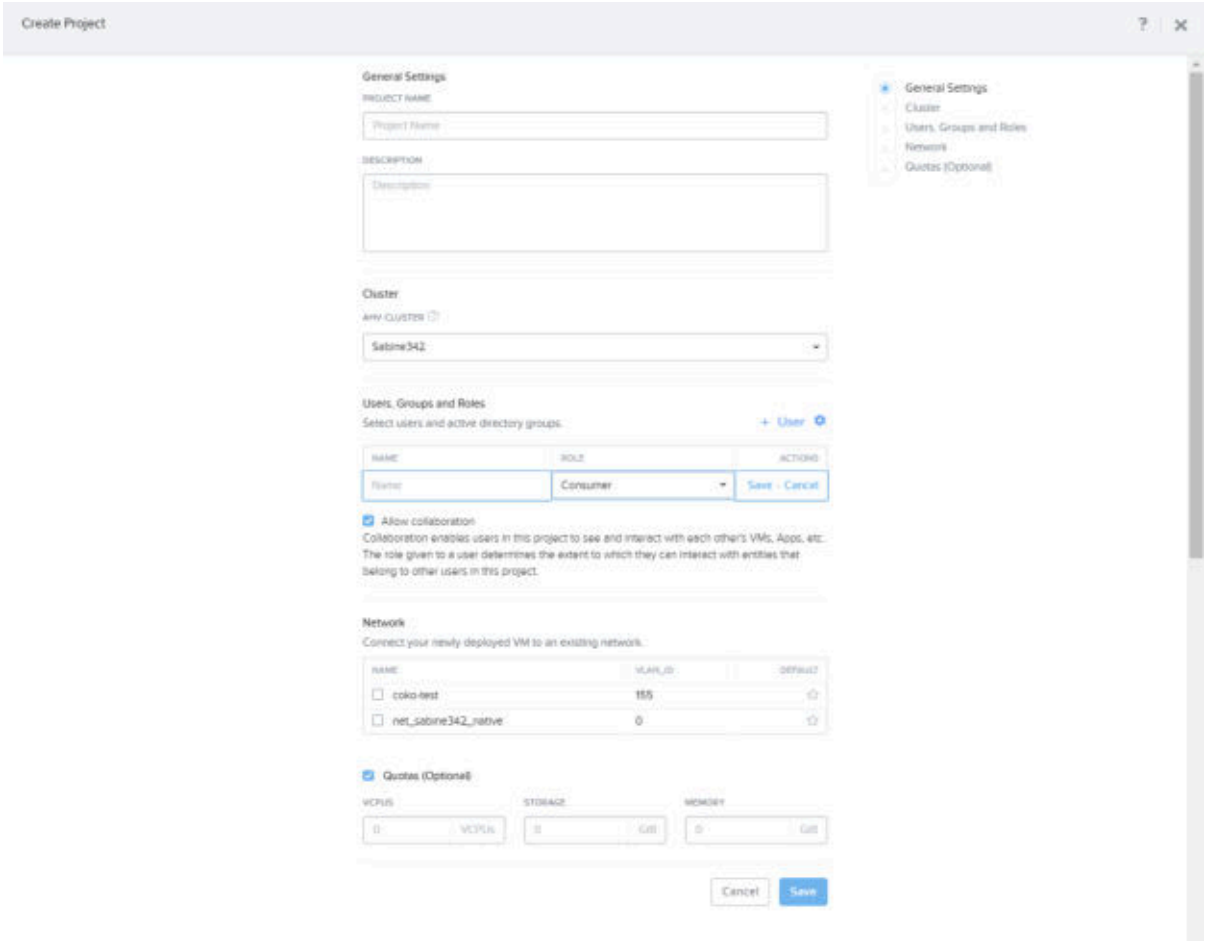

Figure 363: Create Project Page

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.

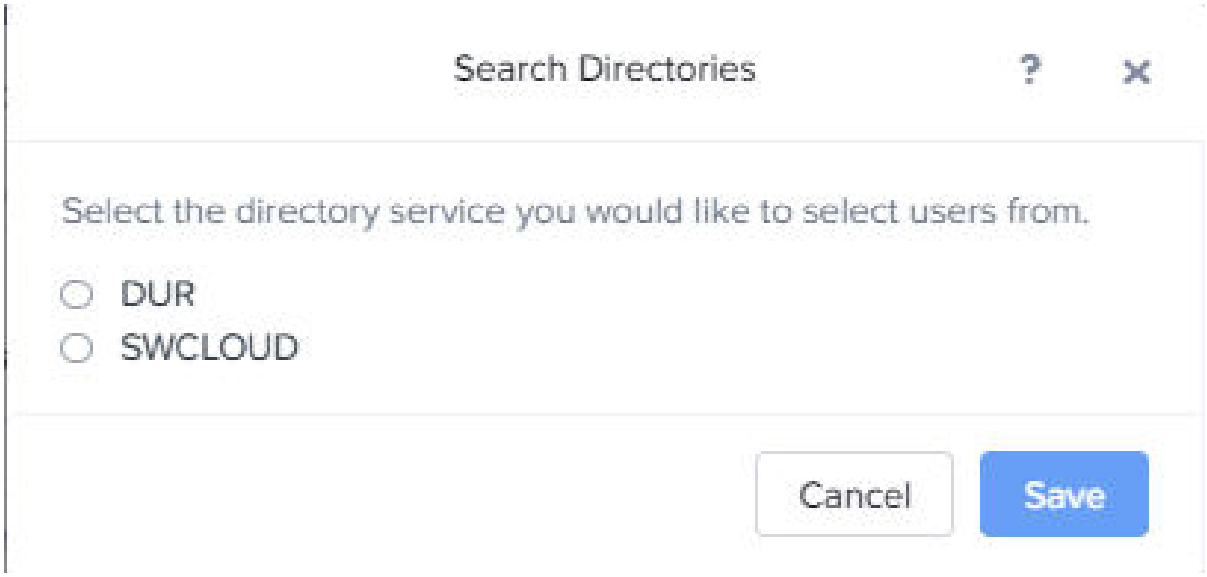

Figure 364: Search Directories Window

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.

> Important: Enabling the Nutanix Calm feature alllows you to define **Quotas** only from the Calm projects. You can set the usage restrictions for quota policy enabled Calm projects for Nutanix (AHV) and non-Nutanix (ESXi) clusters.
>
> For information on managing the Policies for Calm projects, see the *Nutanix Calm Administration and Operations Guide*.

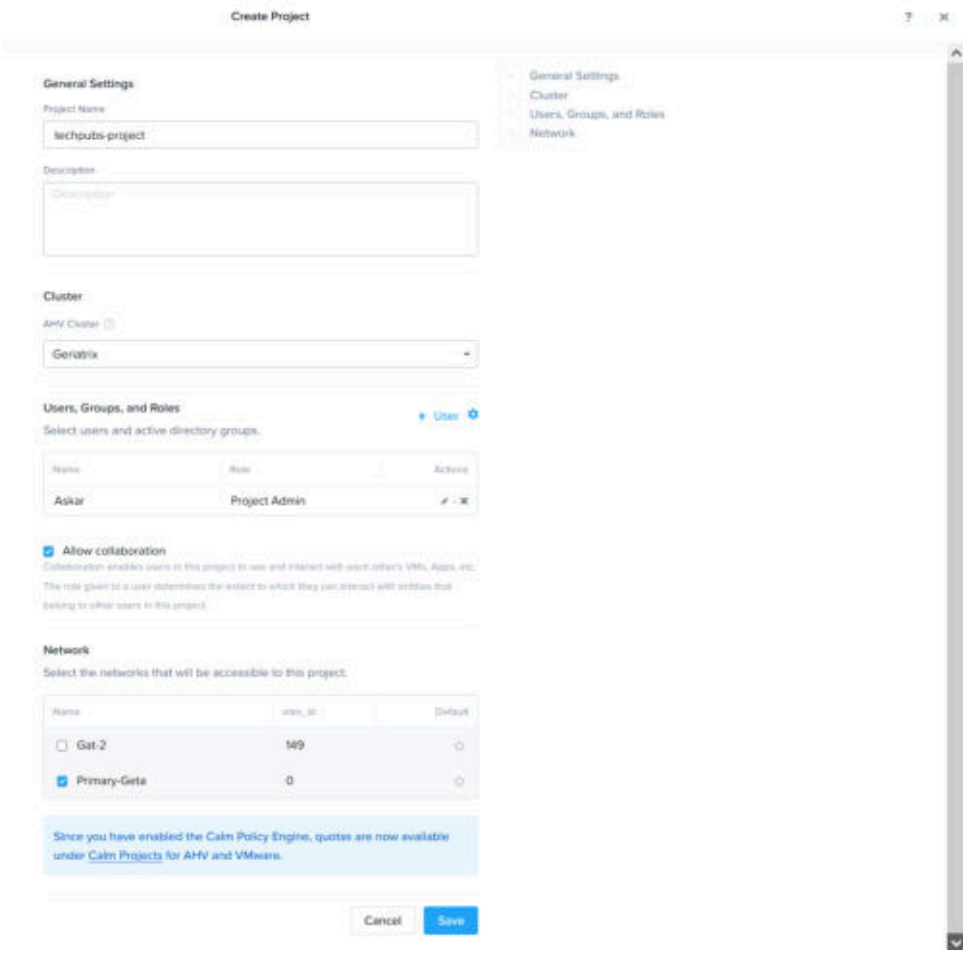

Figure 365: Create Project Page (with Nutanix Calm Policies Enabled)

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 254) 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 718). Update the field values as desired and then click the **Save** button.

   » To delete that project, select **Delete** from the **Actions** pull-down menu. You are prompted to verify the delete (click the **OK** button). The project is then deleted and removed from the list.

   Note: Before you can delete a project, you must first remove any VMs and networks, in that order, from the project. You cannot delete the **default** project.

10

NETWORK AND SECURITY MANAGEMENT

To configure network connections for a cluster, see Configuring Network Connections on page 725.

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 86 or Subnets View).

The Network Configuration page includes three tabs.

- The **Subnets** tab displays a list of the configured networks.

Figure 366: Networks Tab

- The **Internal Interfaces** tab displays a list of LAN interfaces.

Figure 367: Internal Interfaces Tab

- The **Virtual Switch** tab displays a list of virtual switches configured, including the default system-generated virtual switch vs0.

Figure 368: Virtual Switch Tab

The following table describes the fields in each tab.

Table 715: Network Configuration Fields

Parameter	Description	Values
Networks Tab		
Name	Displays the name of the virtual network.	(name)
Virtual Switch	Displays the name of the virtual switch in the form **vs#**, for example **vs0** for virtual switch 0 which is the default virtual switch.	(vs<number>)
VLAN ID	Displays the VLAN identification number for the network in the form **vlan.#**, for example **vlan.27** for virtual LAN number 27.	(ID number)
Used IP Addresses		N/A
Free IPs in Subnets		N/A
Free IPs in Pool		N/A
Actions		(Edit | Delete)
Internal Interfaces Tab		
Descriptive Name	Displays a name for the LAN.	(LAN name)
Subnet (Gateway IP / Prefix Length)	Displays the subnet that the internal interface belongs to in the form *<IP Address>*/<number (prefix)>	(IP Address/prefix number)
Features	Displays the features available on the internal interfaces.	
Interface	Displays the interface designation such as eth0 or eth1.	(interface name)
Virtual Switch		
Name	Displays the name of the switch in the form **vs#**	(vs<number>)

Parameter	Description	Values
Bridge	Displays the name of the bridge associated with the virtual switch in the form `br#`, for example `br0` for the default bridge.	(br<number>)
MTU (bytes)	Displays the MTU set for the virtual switch in bytes. The default MTU is 1500.	(number)
Bond Type	Displays the uplink bond type associated with the virtual switch. See the *Bond Type* table in . For example, `Active-Backup`	(<bond_type>)

LAG and LACP on the ToR Switch

For more information, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 733.

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 86) and click the **Network Config** button.
   - » Go to the Network and Security dashboard (see Subnets View and click the **Network Config** button.

2. In the list Subnets tab of the Network Configuration window, click the **Create Subnet** button. This displays the Create Subnet page.

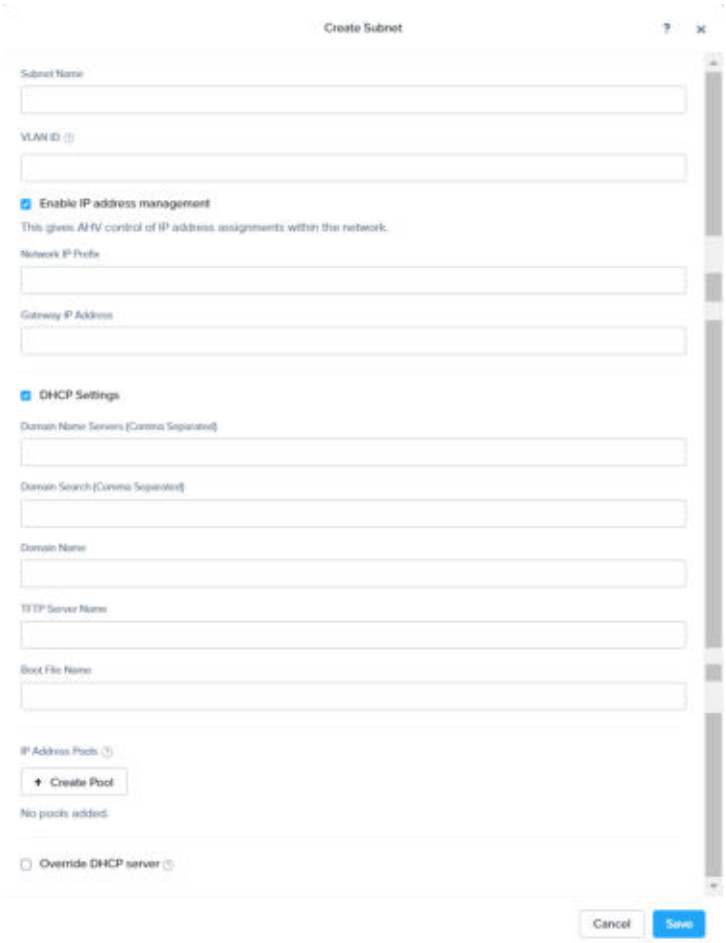

Figure 369: Create Subnet Page (expanded)

3. Do the following in the indicated fields:

   a. **Subnet Name**: Enter a name for the network.

   b. **Virtual Switch**: Select the virtual switch for the network from the drop-down list.

      You can associate the default virtual switch vs0 or create a virtual switch. To create a virtual switch, see Creating or Updating Virtual Switch on page 728.

      For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

   c. **VLAN ID**: Enter the number of the VLAN.

      Enter just the number in this field, for example 1 or 27. Enter 0 for the native VLAN. The value appears as **`vlan.1`** or **`vlan.27`** in displays.

   d. **Enable IP Address Management**: Check the box to have the cluster control IP addressing in the network.

      Checking this box display additional fields. If this box is not checked, no network management is attempted. In this case it is assumed management for this virtual LAN is handled outside the cluster.

   e. **Network IP Prefix**: Enter the IP address of the gateway for the network and prefix with the network prefix (CIDR notation, for example, 10.1.1.0/24).

   f. **Gateway IP Address**: Enter the VLAN default gateway IP address.

   g. **DHCP Settings**: Check this box to display fields for defining a domain.

      Checking this box displays fields to specify DNS servers and domains. Unchecking this box hides those fields.

   h. **Domain Name Servers (Comma Separated)**: Enter a comma-delimited list of DNS servers.

   i. **Domain Search (Comma Separated)**: Enter a comma-delimited list of domains.

   j. **Domain Name**: Enter the VLAN domain name.

   k. **TFTP Server Name**: Enter the host name or IP address of the TFTP server from which virtual machines can download a boot file. Required in a Pre-boot eXecution Environment (PXE).

   l. **Boot File Name**: Name of the boot file to download from the TFTP server.

4. To define a range of addresses for automatic assignment to virtual NICs, click the **+ Create Pool** button (under **IP Address Pools**) and enter the following in the Add IP Pool page:

   Note: If no pool is provided, the user must assign IP addresses to VMs manually.

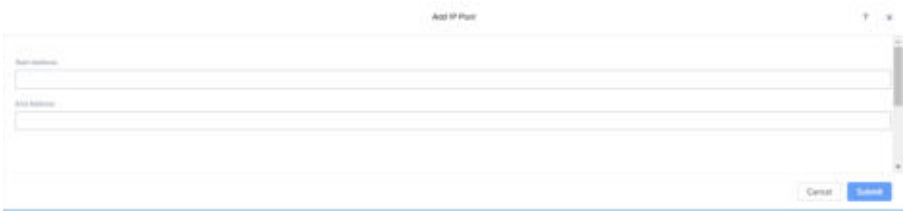

   Figure 370: Add IP Pool Page

   a. Enter the starting IP address of the range in the **Start Address** field.

   b. Enter the ending IP address of the range in the **End Address** field.

   c. Click the **Submit** button to close the page and return to the Create Network page.

5. To configure a DHCP server, check the **Override DHCP server** box and enter an IP address in the **DHCP Server IP Address** field.

   This address (reserved IP address for the Acropolis DHCP server) is visible only to VMs on this network and responds only to DHCP requests. If this box is not checked, the **DHCP Server IP Address** field is not displayed and the DHCP server IP address is generated automatically. The automatically generated address is `network_IP_address_subnet.254`, or if the default gateway is using that address, `network_IP_address_subnet.253`.

6. Click the **Save** button to configure the network connection, close the Create Subnet page, and return to the Network Configuration page.

7. Click the **Close** button to close the Create Subnet page.

Creating or Updating Virtual Switch

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can create or update a Virtual Switch (VS) using the Create Virtual Switch or Update Virtual Switch dialog box. The Create Virtual Switch and Update Virtual Switch dialog boxes have identical parameters and settings.

To open the Create Virtual Switch dialog box in Prism Central, do the following.

Procedure

1. Click **Subnets** > **Network Config**.
2. On the Network Configuration window, click the **Virtual Switch** > **Create VS** link.

About this task

To create a virtual switch using the Create Virtual Switch dialog box, do the following.

Procedure

1. In the Create Virtual Switch dialog box, provide the necessary information on the **General** tab.

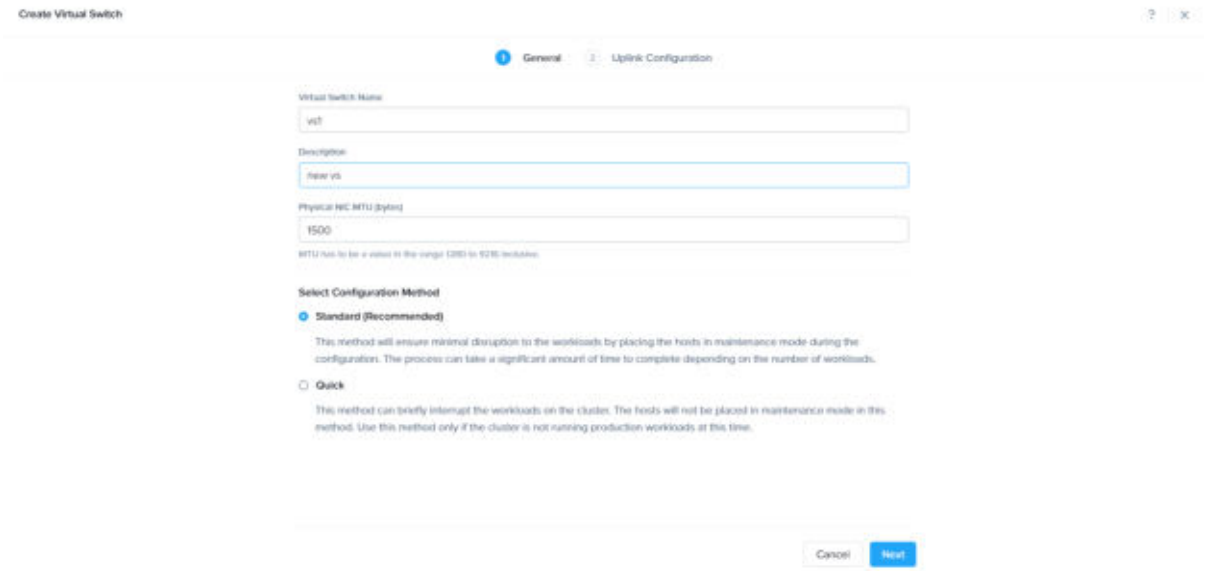

Figure 371: Create Virtual Switch - General tab

Field	Description
Virtual Switch Name	Enter a name for the virtual switch.
Description	Provide a description for the virtual switch that helps identify the virtual switch.
Physical NIC MTU (bytes)	MTU must be a value in the range 1280 to 9216 inclusive.

Field	Description
Select Configuration Method	Select one of the two methods that you can use to implement the VS configuration: • **Standard (Recommended)**: This method ensures no disruptions occur to the workloads by putting the hosts in maintenance mode and migrating the VMs out of the host before applying the configuration. This process requires a longer duration of time to complete. The time required depends on the number and configuration of VMs. In this method, the VS configuration is deployed in the rolling update process. • **Quick**: This method interrupts the workloads running on the hosts. Use this method only if you are not running production workloads because it may result in network interruptions. In this method, the VS configuration is deployed in a rolling update process but the nodes are not put in maintenance mode before modifying the VS configuration on the node.

2. Click **Next** to go to the Uplink Configuration tab.

   On the Uplink Configuration tab, provide the following details:

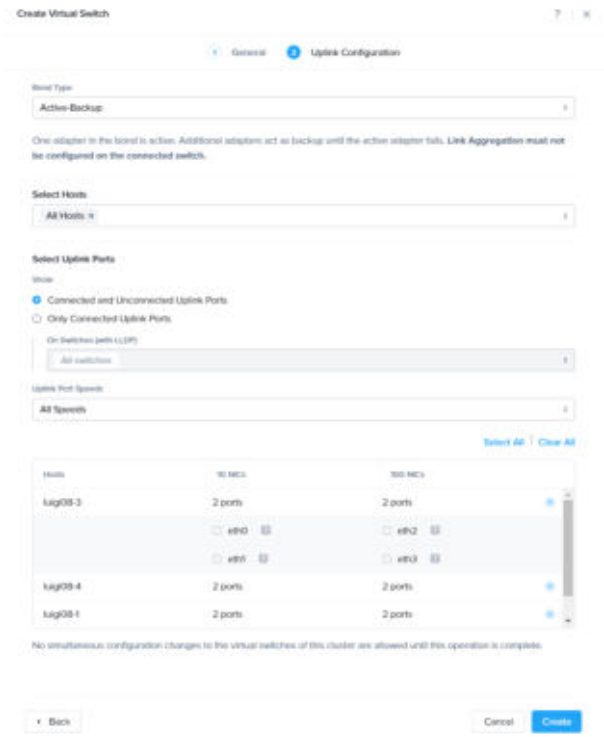

Figure 372: Create Virtual Switch - Uplink Configuration tab

Field	Description and Value
Bond Type	Select an appropriate bond type. See the *Bond Types* table for details about the bond types.
Select Hosts	Select the hosts that can host VMs.
Select Uplink Ports	Select the criteria that need to be satisfied for the uplink ports. The available uplink ports that satisfy the criteria are displayed in the (Host port) table at the bottom of this tab.
Show	(Port Type) Select one of the following: **Connected and Unconnected Uplink Ports**: Select this if you want to use ports that are not currently connected but may be connected later. **Only Connected Uplink Ports**: Select this option if you want to use only the connected ports. You must also select in the **On Switches (with LLDP)** drop down box the switches with LLDP.
Uplink Port Speeds	Select a speed to display the ports that have the selected speed. You can select speeds such as **1G**, **10G** or both (**All Speeds**). The speeds displayed depend on the NIC type that is installed on the host. Based on your selection the columns in the (Host Port) table change dynamically to display the ports with the speeds you selected.
(Host Port) table	Based on the Selections you made in this Select Uplink Ports section, a table displays the hosts that have the uplink ports that satisfy the selected criteria. Select the ports you need for this configuration from the list. Click the down arrow on the right side of the table to display the ports listed for each host. Note: A port listing is greyed out if it is unavailable because it is already associated with another virtual switch. Click **Select All** to select all the ports available and listed. Click **Clear All** to unselect all the ports available and listed.

Table 716: Bond Types

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Backup	Recommended. Default configuration, which transmits all traffic over a single active adapter.	10 Gb	10 Gb

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Active with MAC pinning Also known as balance-slb	Works with caveats for multicast traffic. Increases host bandwidth utilization beyond a single 10 Gb adapter. Places each VM NIC on a single adapter at a time. Do not use this bond type with link aggregation protocols such as LACP.	10 Gb	20 Gb
Active-Active Also known as LACP with balance-tcp	LACP and link aggregation required. Increases host and VM bandwidth utilization beyond a single 10 Gb adapter by balancing VM NIC TCP and UDP sessions among adapters. Also used when network switches require LACP negotiation. The default LACP settings are: • Speed—Fast (1s) • Mode—Active fallback-active-backup • Priority—Default. This is not configurable.	20 Gb	20 Gb
No Uplink Bond	No uplink or a single uplink on each host. Virtual switch configured with the No uplink bond uplink bond type has 0 or 1 uplinks. When you configure a virtual switch with any other bond type, you must select at least two uplink ports on every node.	-	-

Note: The Maximum VM NIC Throughput and Maximum Host Throughput values are not restricted to the value provided in this table. The values in the table are indicated for an assumption of 2 x 10 Gb adapters.

For more information about uplink configuration, see the *Uplink Configuration (AHV Only)* topic in the *AHV Administration Guide*.

3. Click **Create** to create the virtual switch.

   Click **Cancel** to exit without creating the virtual switch.

   Click **Back** to go back to the General tab.

What to do next

> Important:
>
> You can migrate or convert the bridges other than br0 in the cluster to virtual switches after you upgraded the minimum or compatible version of AOS and AHV. You can convert the other bridges only on Prism Web Console or using aCLI. You can convert only one bridge at a time. You need to repeat the workflow for every bridge that you want to convert to a virtual switch.
>
> For more information about converting the bridges other than br0, see *Migrating Bridges after Upgrade*in *Prism Web Console Guide*.

- If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

  For information about enabling LAG and LACP, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 733

Enabling LAG and LACP on the ToR Switch (AHV Only)

Active-Active with MAC pinning does not require link aggregation since each source MAC address is balanced to only one adapter at a time. If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

About this task

In an **Active-Active** NIC-teaming configuration, network traffic is balanced among the members of the team based on source and destination IP addresses and TCP and UDP ports. With link aggregation negotiated by LACP, the uplinks appear as a single layer-2 link so a VM MAC address can appear on multiple links and use the bandwidth of all uplinks. If you do not enable LAG and LACP in this **Active-Active** configuration, network traffic might be dropped by the switch.

Perform the following procedure to enable LAG and LACP on a TOR switch after enabling the Active-Active policy by using the Prism Element web console.

Procedure

1. Put the node in maintenance mode. This is in addition to the previous maintenance mode that enabled Active-Active on the node.
2. Enable LAG and LACP on the ToR switch connected to that node.
3. Exit maintenance mode after LAG and LACP is successfully enabled.
4. Repeat steps 1 to 3 for every node in the cluster.

Modifying Network Connections

About this task

To modify or delete a network configuration (defined on an Acropolis managed cluster), do the following:

Procedure

1. Do one of the following:

   » Go to the List tab of the VMs dashboard (see VMs Summary View on page 86) and click the **Network Config** button.

   » Go to the Networks dashboard (see Subnets View and click the **Network Config** button.

2. To modify a network configuration, in the **Subnets** tab, select the target network, and click **Edit** under the **Actions** column.

   The Update Subnet window appears, which contains the same fields as the Create Subnet page (see Configuring Network Connections on page 725).

   Do the following:

   a. Update the field values as desired.

   b. When the configuration is correct, click the **Save** button to close the page, update the network configuration, and return to the Network Configuration page.

3. To delete a network configuration, in the Subnets tab, select the target network, and click **Delete** under the **Actions** column.

   A window prompt appears to verify the action; click the **OK** button. The network configuration is removed from the list.

   > Note:
   >
   > You can also select the target network on the Networks dashboard and select **Delete** in the **Actions** drop-down list.

4. Click the **Close** button to close the Network Configuration page.

11

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 373: Dept Details View (example)

Creating a Category

About this task

To create a category, do the following:

Figure 374: Creating a Category

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275774441001 on Brightcove.

Procedure

1. Go to the categories dashboard (see Categories Summary View on page 157) and click the **Create Category** button.
   The Create Category page appears.

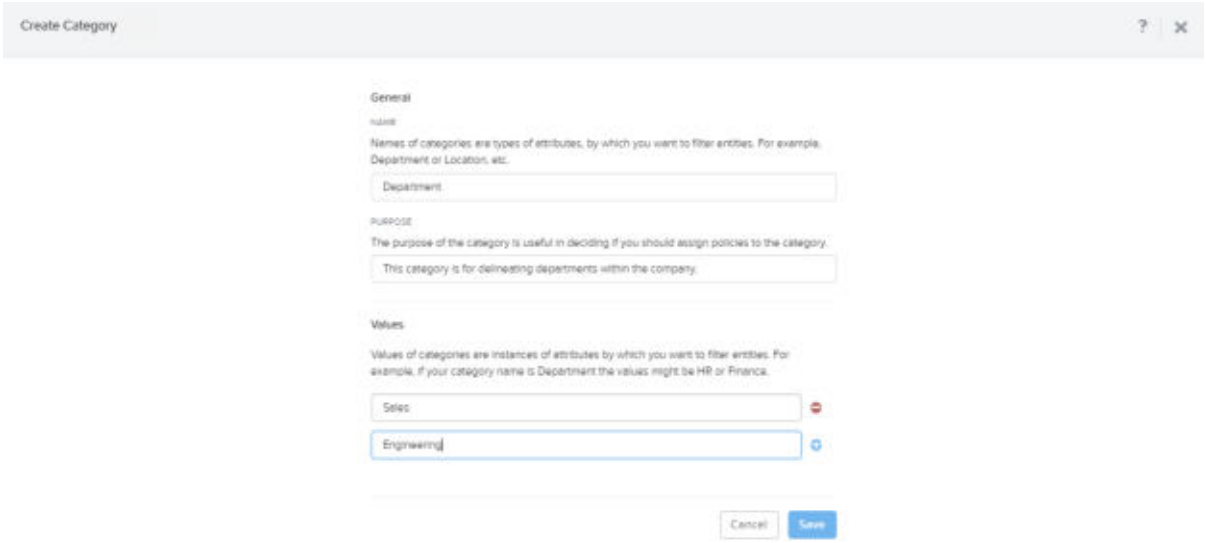

Figure 375: Create Category Page

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 157) 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 735). 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 197) 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 86) and select the target VMs.
   - Image: Go to the Images dashboard (see Images Summary View on page 117) and select the target images.
   - Subnet: Go to the Subnets dashboard (see Subnets View) 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.

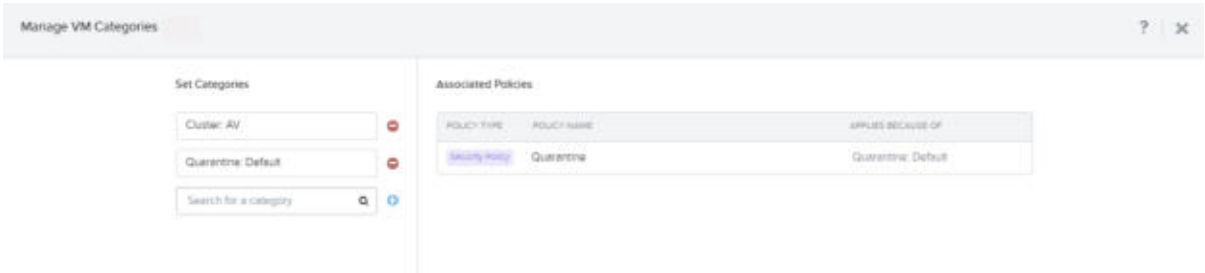

Figure 376: Manage VM Categories Page

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.

12

OPERATIONS MANAGEMENT

- Analysis: Analyze cluster activity (see Analysis Dashboard on page 739)
- App Discovery: Discover applications running in a cluster (see Application Discovery on page 771)
- Cost Management: Analyze and manage costs (see Cost Management (Xi Beam) on page 795)
- Integrations: Monitor select applications in a cluster (see Application Monitoring on page 796)
- Ops Policies: Manage operations policies (see Operations Policy Management on page 812)
- Planning: Plan for resource capacity and usage (see Resource Planning on page 816)
- Playbooks: Create playbooks for task automation (see Task Automation on page 841)
- Reports: Generate resource and activity reports (see Reports View on page 245)

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 739).
- You can create custom entity and metric performance charts (see Chart Management on page 749).

> Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

Analysis Dashboard

The Analysis dashboard allows you to create sessions with charts that can monitor dynamically a variety of performance measures. To view the Analysis dashboard at any time, select **Operations** > **Analysis** from the entities menu (see Entities Menu on page 13).

When you open the Analysis dashboard, it displays the session that you were last working in as the default session. A session helps you correlate the metrics with the alerts and events for troubleshooting. You can create new sessions to troubleshoot new situations and preserve the old sessions created for old issues.

When you open the Analysis dashboard for the first time after installing or upgrading Prism Central, the Analysis dashboard displays a session named **Analysis Session**. Prism displays the label **System** with the name of this session in the Sessions dashboard to indicate that this session is a system-generated session. Prism displays the label next to the **Analysis Session** session name to indicate that this session is system-generated. You can set the default session

for display in the Analysis dashboard in the Sessions dashboard. For more information about the Sessions dashboard, see Sessions Dashboard on page 765.

> Note:
>
> The updated Analysis dashboard and Sessions dashboard are available only with the Prism Pro and Prism Ultimate licenses. With the Prism Starter license, only the system-generated session is available.

Prism Central saves all session changes and new sessions automatically including **Session Name** changes, except changes made using specific options that open respective dialog boxes to facilitate the changes. For example, if you change the name of the session in the Analysis dashboard, then the change is automatically saved. However, when you edit a chart using the **Edit** option and the **Edit Chart** dialog box on the Analysis dashboard, you need to click the **Update** button to save the changes made in the dialog box. When you modify the session name or description in the Sessions dashboard using the **Edit Session Details** action, you need to click the **Update** button to save the changes made in the dialog box.

The following figures are sample views of the Analysis dashboard. Some fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated.

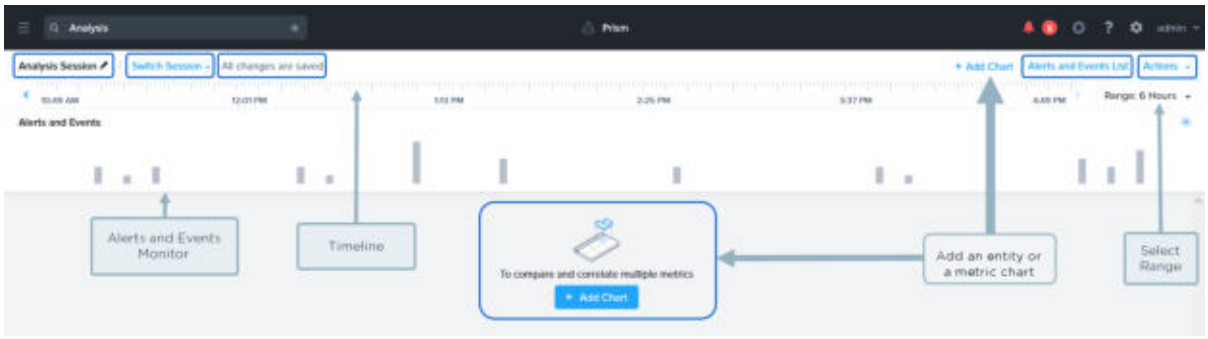

Figure 377: Analysis Dashboard - Prism-generated Session

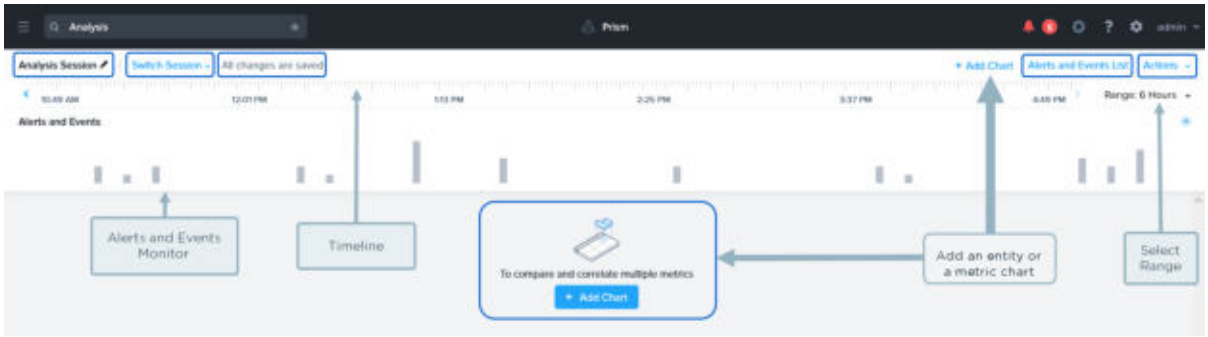

Figure 378: Analysis Dashboard - User-defined Session

The Analysis dashboard elements are described in the table starting with the elements in the top left corner to the right side of the dashboard.

Table 717: Analysis Dashboard Elements

Field	Description
Session	The name of the session that is displayed in the Analysis dashboard. The Analysis dashboard displays the session you viewed last by default. You can change the default session in **Settings** in the Sessions dashboard.
Switch Session	A dropdown menu of all available sessions created and saved. Select the session you want to view and modify the details for. Select **Switch Sessions** > **View All Sessions** to view the Sessions dashboard that lists all the sessions created and saved.
Changes saved status	The status whether the changes made to the session are saved. • Displays **Saving** when Prism is saving the changes. • Displays **All changes are saved** when Prism has saved all the changes.
Add Chart	You can add a chart by clicking: • The **Add Chart**option at the top left of the dashboard. • The **Add Chart** button at the bottom center of the dashboard. You can define a chart to monitor many entities with the same single metric. You can define a chart to monitor several metrics for a single entity. Note: You cannot define a single chart with many entities and many metrics.

Field	Description
Alerts and Events monitor	Displays the alerts and events that occurred. The alerts and events occurring at any specific point of time are displayed as a stacked bar with three colored segments. The hover image of the stacked bar displays the following: • Time range for the stacked bar. • Red segment depicting **Critical Alert** with the number of alerts. • Yellow segment depicting **Warning Alert** with the number of alerts. • Grey segment depicting **Events** with the number of events. Click the stacked bar to open the Alerts and Events List panel on the right. The Alerts and Events List panel displays the details of the alerts and events represented by the stacked bar.
Alerts and Events List	Provides more information about the alerts and events. When you click: • The Alerts and Events List option at the top right of the Analysis dashboard, the Alerts and Events List panel provides information about all the alerts and events that have occurred during the selected time range. • A stacked bar in the **Alerts and Events** monitor, the Alerts and Events List panel provides information about all the alerts and events that occurred in the time range indicated by the stacked bar that you clicked on the timeline. See the *Alerts and Events List Panel Elements* table for descriptions of the tabs and fields in the **Alerts and Events List** panel.

Field	Description
Actions	Provides the following three options for the current session: 1. **Edit Session Details** 2. **Delete** (session) This action is not available and appears dimmed if the current session is the system-defined session, because the system-defined session cannot be deleted. 3. **Close** (session) When you close the current session, the Sessions dashboard is displayed.
(Defined chart monitors)	Displays the monitors for the charts that you defined for the session.

Table 718: Alerts and Events List Panel Elements

Field	Description
Alerts tab	Displays the list of alerts. The tab name includes colored icons with inscribed numbers. The color and number provide the following information: • Color depicts the severity type of the alert. Red indicates critical and yellow indicates warning. • Number indicates the number of respective alerts. In the example depicted in the figure, the red icon has 1 inscribed indicating that there is one critical alert in the stacked bar instance.
Group by field in Alerts tab	Provides the grouping criteria that you can use to group the alerts. You can group the alerts by: • **Entity Type** • **Impact** • **Severity**

Field	Description
Events tab	Displays the list of events. The tab name includes a grey-colored icon with an inscribed number. The inscribed number indicates the number of events. You can see the event details page by clicking the event in the list. The events in the list are grouped into two default groups: • **System Action** • **User Action**
(Time range for the stacked bar)	The time range considered for the stacked bar plotted on the **Alerts and Events** monitor. The time range is specifically displayed only when you open the Alerts and Events List panel by clicking a stacked bar in the **Alerts and Events** monitor.
Hide/Show toggle	Displayed for each group of alerts or events in the list. This toggle allows you to hide or show the alerts or events in a group.

Modifying the Time Range

You can select a time range for the session and the charts you choose in the Analysis dashboard.

About this task

The time range you select applies to the session as a whole including the charts you add to the session. The default time range is **1 Day**.

> Note: The time range selection affects the whole session and all the charts added to the session.

You can select a predefined time range or customize a range. You can customize a range by moving the start and end time selectors on the time-line or you can set a range in the Select a Time Range dialog box.

To select a predefined time range for the session, do the following:

Procedure

1. Click the Range drop-down list.
2. Select one of the following:
   - » A predefined range from the list.
   - » **Custom Range**.

3. If you selected **Custom Range**, then in the Select a Time Range dialog box:
   a. Select the date and time in the **From** section to define the start of the range.
   b. Select the date and time in the **To** section to define the end of the range.
   c. Click **Select**.

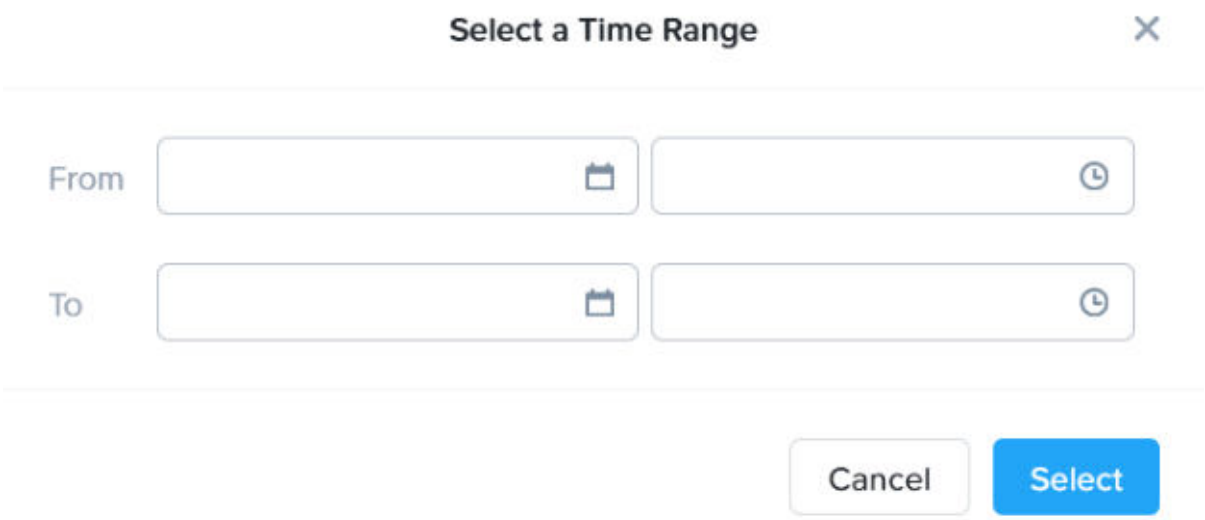

Figure 379:

If you selected **Custom Range**, you can also select the range dynamically on the time-line by sliding the range start and end sliders.

There is no minimum time range for **Custom Range**, and the maximum range is one year.

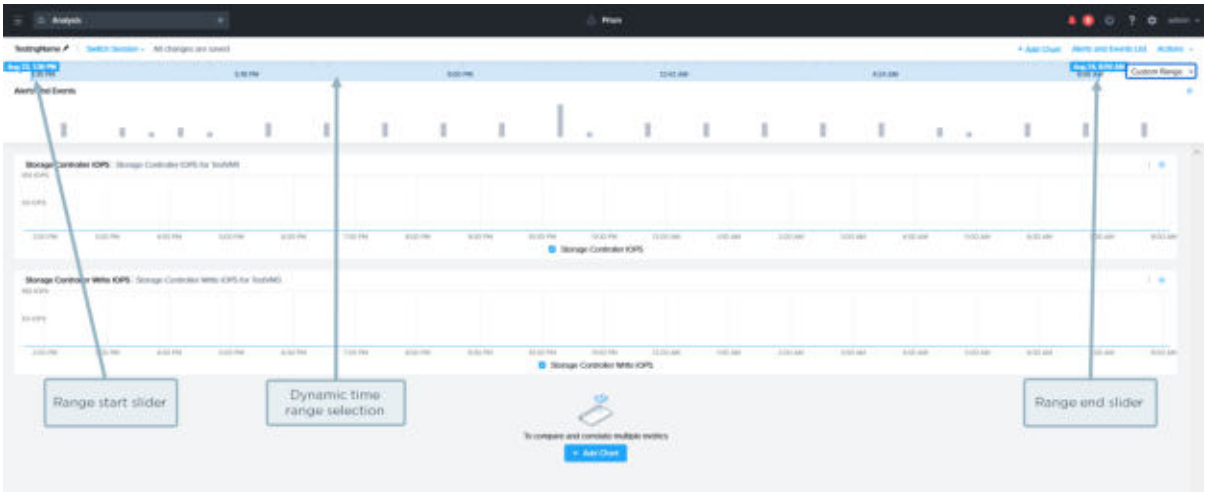

Figure 380: Time Range

About this task

You can view the alerts and events occurring during the selected time range in the following ways:

- Displayed on the **Alerts and Events** monitor available below the time-line.
- By clicking the **Alerts and Events List** option at the top right of the Analysis dashboard

Procedure

1. Click the stacked bar to see the details of the alerts and events in the Alerts and Events List panel.

   Point to a stacked bar on the monitor to see the hover image. The hover image provides more information about the alerts and events that occurred in the time range that the stacked bar depicts.

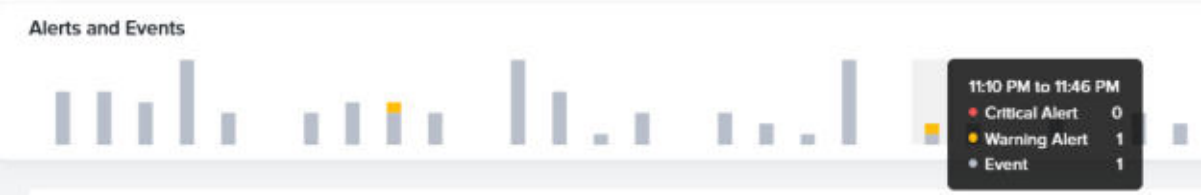

Figure 381: Stacked bar - Alerts and Events Information

2. Click the **Alerts and Events List** option at the top right of the Analysis dashboard.

   The Alerts and Events List panel displays the details of the alerts and events that occurred during the time range selected.

   See the Table 718: Alerts and Events List Panel Elements on page 743 table for more information.

   - The Alerts tab provides the details of the alerts. The alerts can be filtered.

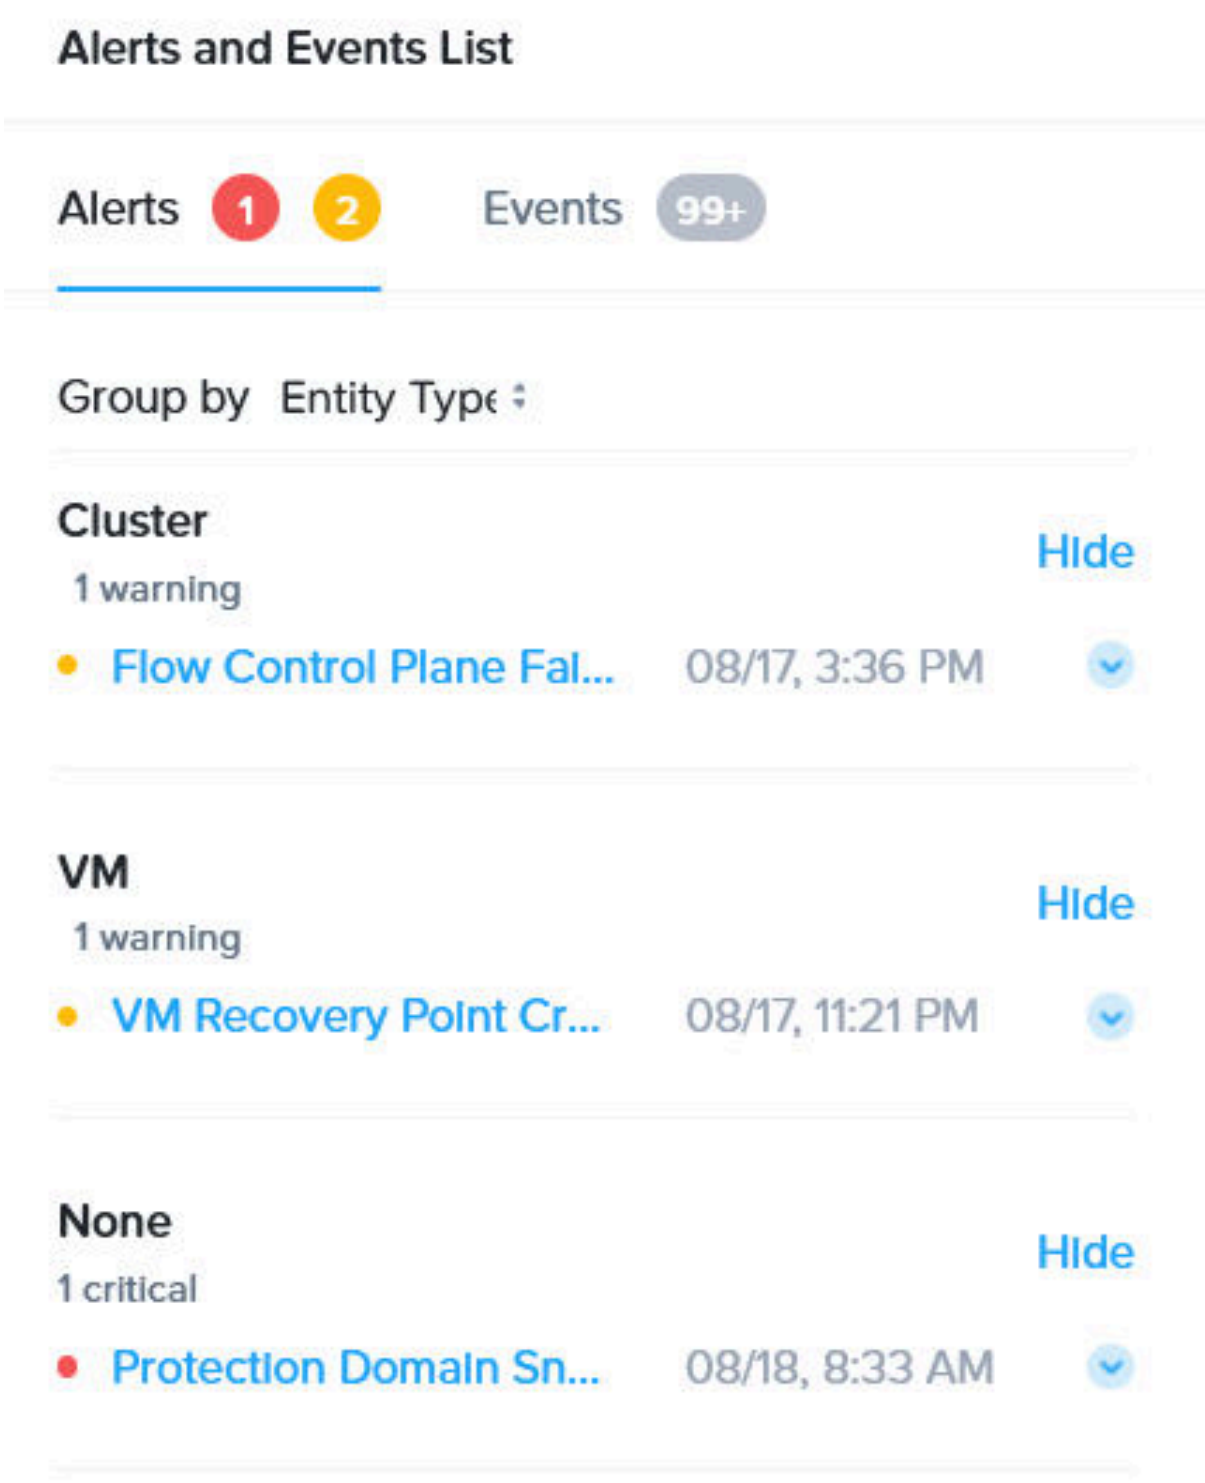

Figure 382: Alerts tab

Use the filters provided in the **Group by** drop down list to filter the category of alerts.

- The Events tab provides the details of the events. The events cannot be filtered.

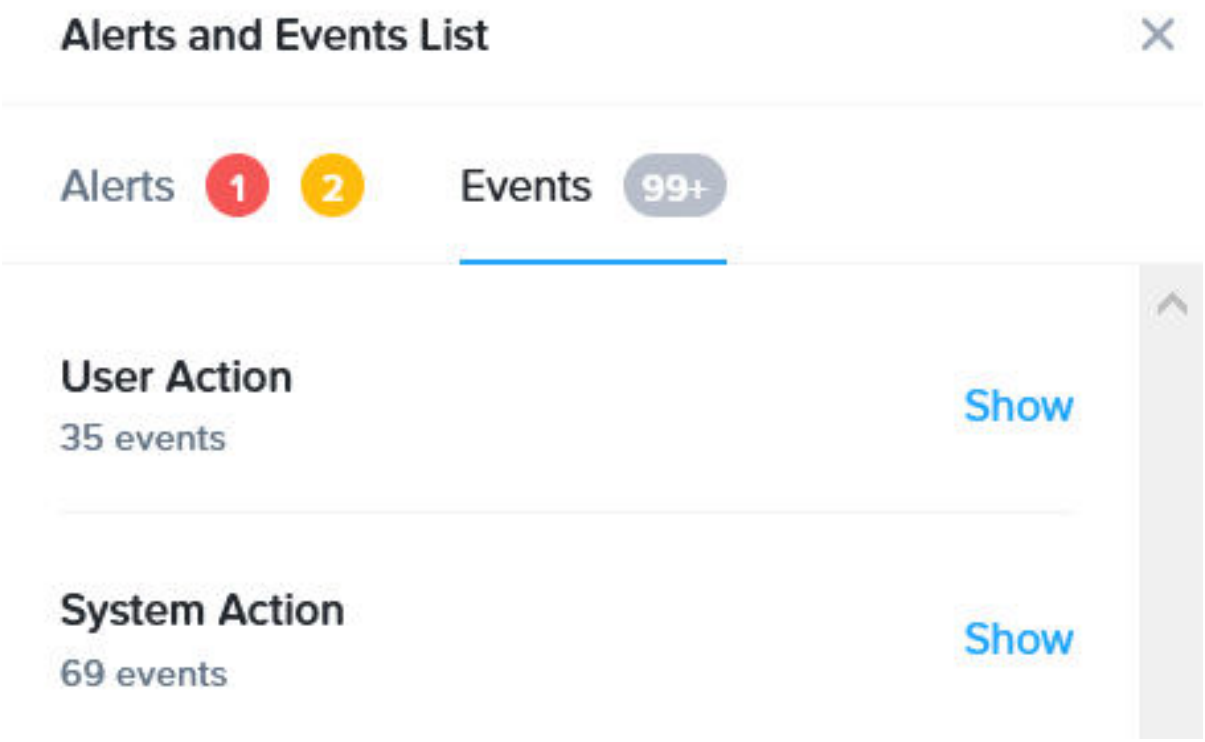

Figure 383: Events tab

Click **Show** to expand the category of alerts or events.

3. Click an alert or event to open the alert or event details page.

Chart Management

On the Analysis dashboard, you can:

Add charts	Adding Charts on page 750
Edit a chart	Editing a Chart on page 752
Delete a chart	Deleting a Chart on page 753
Clone a chart	Cloning a Chart on page 754
Export a chart as a JSON or CSV file	Exporting a Chart on page 754

Click the options button (three vertical dots) on the top left corner of a chart to see the aforesaid options, except **Add Chart**. The following sample figure shows the options.

Figure 384: Sample Chart - Chart Options

Adding Charts

This is a Nutanix task.

About this task

You can add charts in a session by clicking:

- The **Add Chart** option at the top left of the dashboard.
- The **Add Chart** button at the bottom center of the dashboard.

The **Add Chart** option and button opens the Add Chart dialog box. You can add multiple charts by opening the Add Chart dialog box once.

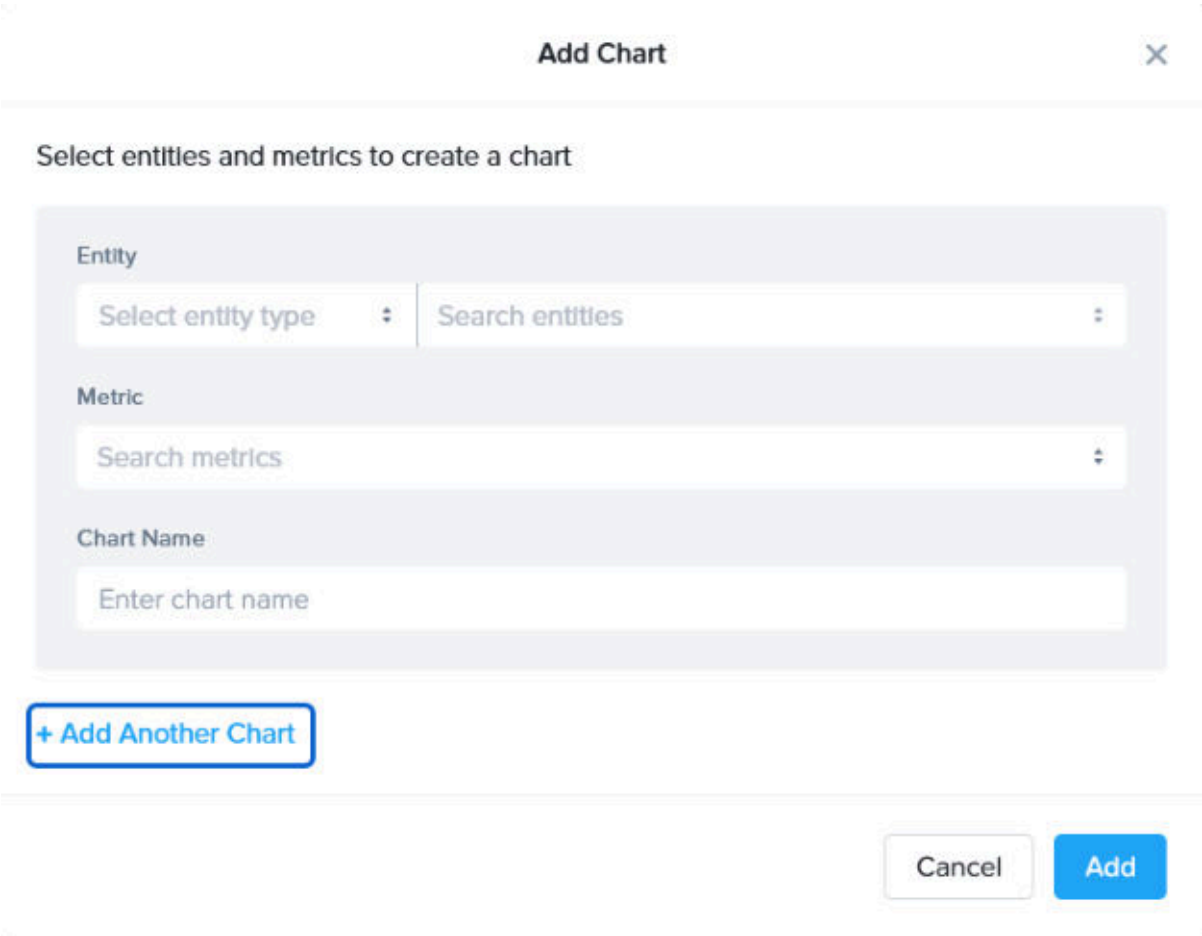

Figure 385: Add Chart dialog box

You can define a single chart to monitor a single metric for multiple entities of the same entity type. You can define a chart to monitor multiple metrics for a single entity.

> Note:
>
> You cannot define a single chart to monitor multiple metrics for multiple entities.

Procedure

1. Click one of the following:
   - » The **Add Chart** option at the top left of the dashboard.
   - » The **Add Chart** button at the bottom center of the dashboard.
2. In the Add Chart dialog box, do the following:
   a. Under **Entities**, select an entity type.
   b. Select the entities by selecting the check boxes for the entities in the list of entities.
   c. Under **Metric**, select the metrics you want by selecting the check boxes for the metrics in the list of metrics.

      See Chart Metrics on page 755 for the list of metrics you can select.

      > Note:
      >
      > In one chart, select metrics that use the same unit of measurement. Add a metric with a different unit of measurement to a separate chart.
      >
      > See the step to add another chart below.
      >
      > When adding a metric for multiple entities hosted on different hypervisors, select a metric only if it is supported on the hypervisors that are hosting the selected entities. Prism validates the list of available metrics provided for the **Metric** field. Prism removes from the list any metric that is not supported by all the hosting hypervisors for the selected entities. For example, for two entities hosted on ESXi and AHV hypervisors respectively, the list includes only the metrics that are supported on both ESXi and AHV.

   d. Provide an appropriate **Chart Name** with appropriate keywords.

      Prism provides a default name in the **Chart Name** field based on the entities and metrics you select in the preceding steps.

   e. To add another chart, click **Add Another Chart**.

      Follow the preceding steps a to d, to define the chart.

      Click the **Add Another Chart** option multiple times in the Add Chart dialog box. Each time you click the **Add Another Chart** option, provide the necessary details (steps a to d).

   f. Click **Add** after adding all the charts.

3. You can also add any chart or metric from anywhere in Prism Central. For example, you can add a chart from the Metrics tab of a Disk details page. On the chart or metric widget, do one of the following:

   » Click **Actions** > **Add to Analysis** to add the metric to the current session in the Analysis dashboard.

   » Click **Actions** > **Select Analysis Session** to add the metric to any available session or a new session.

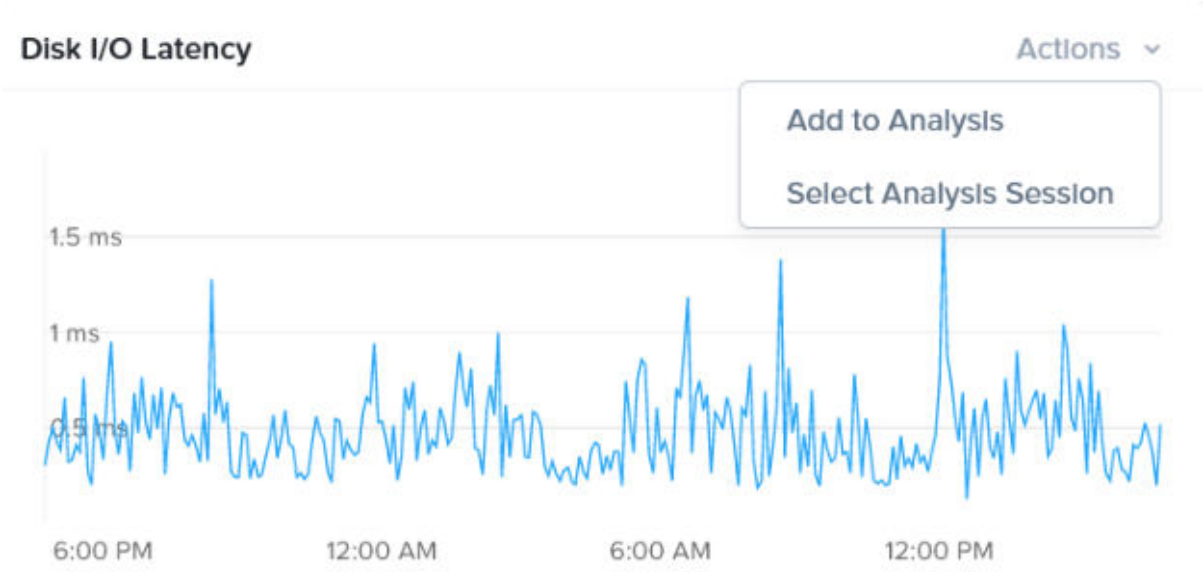

Figure 386: Example of chart or metric widget

Editing a Chart

About this task

You can edit a chart in the Analysis dashboard to modify the entity or the metric you are monitoring with the chart. You can also modify the name of the chart.

To edit a chart, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.

2. Click **Edit**.

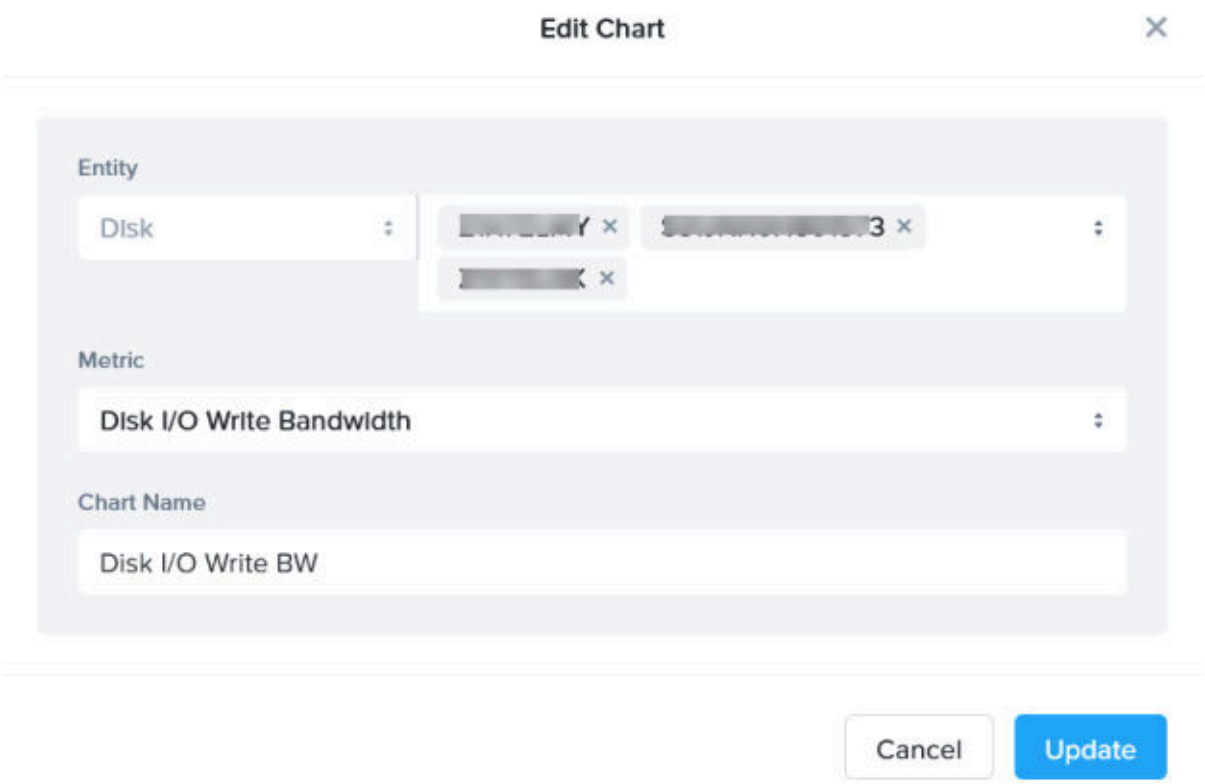

Figure 387: Edit Chart dialog box

3. In the **Edit Chart** dialog box, make the necessary changes in the fields.
4. Click **Update**.
   The changes are displayed in the chart on the Analysis dashboard.

Deleting a Chart

About this task

To delete a chart, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.

2. Click **Delete**.

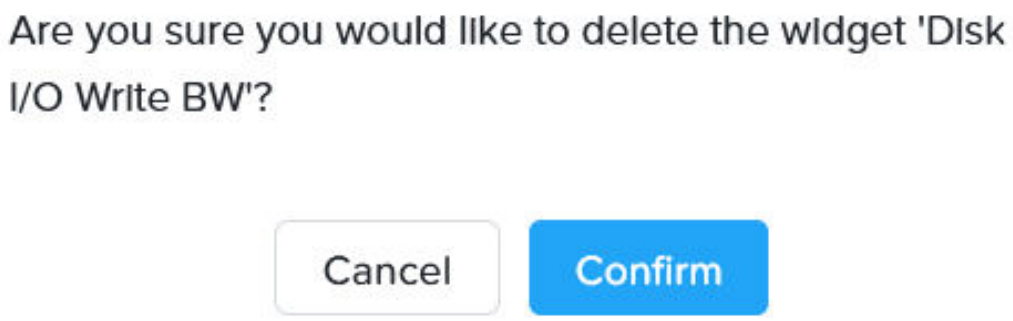

Figure 388: Delete confirmation dialog box

3. Do one of the following:
   - » Click **Confirm** to delete the chart widget.
   - » Click **Cancel** to return to the Analysis dashboard without deleting the chart.

Cloning a Chart

About this task

To clone a chart, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.
2. Click **Clone**.
   The cloned chart is the last chart in the row of charts in the Analysis dashboard.

Exporting a Chart

About this task

You can export a chart as a CSV or a JSON file.

To export a chart as CSV or JSON file, do the following:

Procedure

1. Click the options button (three vertical dots) on the top left corner of the chart that you want to modify.

2. Click one of the following:
   - » **Export CSV** to save as a CSV file.
   - » **Export JSON** to save as a JSON file.
3. Select **Save File**.
4. Browse to the location you want to save the file in the local folder and click **Save**.

Chart Metrics

These metrics can be added to charts.

> Note: The mapping between a metric and an entity type is hypervisor dependent.

Metric	Entity Type(s)	Description
Content Cache Hit Rate (%)	• Host • Cluster	Content cache hits over all lookups. ID: `CONTENT_CACHE_HIT_PPM`
Content Cache Hits	• Host • Cluster	Number of hits on the content cache. ID: `CONTENT_CACHE_NUM_HITS`
Content Cache Logical Memory Usage	• Host • Cluster	Logical memory (in bytes) used to cache data without deduplication. ID: `CONTENT_CACHE_LOGICAL_MEMORY_USAGE_BYTES`
Content Cache Logical SSD Usage	• Host • Cluster	Logical SSD memory (in bytes) used to cache data without deduplication. ID: `CONTENT_CACHE_LOGICAL_SSD_USAGE_BYTES`
Content Cache Lookups	• Host • Cluster	Number of lookups on the content cache. ID: `CONTENT_CACHE_NUM_LOOKUPS`
Content Cache Memory Saved	• Host • Cluster	Memory (in bytes) saved due to content cache deduplication. ID: `CONTENT_CACHE_SAVED_MEMORY_USAGE_BYTES`
Content Cache Physical Memory Usage	• Host • Cluster	Real memory (in bytes) used to cache data by the content cache. ID: `CONTENT_CACHE_PHYSICAL_MEMORY_USAGE_BYTES`
Content Cache Reference Count	• Host • Cluster	Average number of content cache references. ID: `CONTENT_CACHE_NUM_DEDUP_REF_COUNT_PPH`

Metric	Entity Type(s)	Description
Content Cache SSD Usage	• Host • Cluster	Real SSD usage (in bytes) used to cache data by the content cache. ID: `CONTENT_CACHE_PHYSICAL_SSD_USAGE_BYTES`
Content Cache SSD Usage Saved	• Host • Cluster	SSD usage (in bytes) saved due to content cache deduplication. ID: `CONTENT_CACHE_SAVED_SSD_USAGE_BYTES`
Deduplication Fingerprints Cleared	• Host • Cluster	Number of written bytes for which fingerprints have been cleared. ID: `DEDUP_FINGERPRINT_CLEARED_BYTES`
Deduplication Fingerprints Written	• Host • Cluster	Number of written bytes for which fingerprints have been added. ID: `DEDUP_FINGERPRINT_ADDED_BYTES`
Disk I/O Bandwidth	• Host • Cluster • Disk • Storage Pool	Data transferred per second in KB/second from disk. ID: `STATS_BANDWIDTH`
Disk I/O Bandwidth - Read	• Host • Cluster • Disk • Storage Pool	Read data transferred per second in KB/second from disk. ID: `STATS_READ_BANDWIDTH`
Disk I/O Bandwidth - Write	• Host • Cluster • Disk • Storage Pool	Write data transferred per second in KB/second from disk. ID: `STATS_WRITE_BANDWIDTH`
Disk I/O Latency	• Host • Cluster • Disk • Storage Pool	I/O latency in milliseconds from disk. ID: `STATS_AVG_IO_LATENCY`

Metric	Entity Type(s)	Description
Disk IOPS	• Host • Cluster • Disk • Storage Pool	Input/Output operations per second from disk. ID: `STATS_NUM_IOPS`
Disk IOPS - Read	• Host • Cluster • Disk • Storage Pool	Input/Output read operations per second from disk. ID: `STATS_NUM_READ_IOPS`
Disk IOPS - Write	• Host • Cluster • Disk • Storage Pool	Input/Output write operations per second from disk. ID: `STATS_NUM_WRITE_IOPS`
GPU Framebuffer Usage	Virtual Machine	Framebuffer usage in percentage. ID: `FRAMEBUFFER_USAGE_PPM` Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU Usage	Virtual Machine	GPU compute usage in percentage. ID: `GPU_USAGE_PPM` Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU video decoder Usage	Virtual Machine	GPU video decoder usage in percentage. ID: `DECODER_USAGE_PPM` Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU video encoder usage	Virtual Machine	GPU video encoder usage in percentage ID: `ENCODER_USAGE_PPM` Note: The Virtual Machine entity is applicable to AHV hypervisor.
Hypervisor CPU Ready Time (%)	Virtual Machine	Percentage of time that the virtual machine was ready, but could not get scheduled to run. ID: `STATS_HYP_CPU_READY_TIME`

Metric	Entity Type(s)	Description
Hypervisor CPU Usage (%)	• Host • Cluster • Virtual Machine	Percent of CPU used by the hypervisor. ID: `STATS_HYP_CPU_USAGE`
Hypervisor I/O Bandwidth	• Host • Cluster • Virtual Machine	Data transferred per second in KB/second from Hypervisor. ID: `STATS_HYP_BANDWIDTH` Note: Virtual Machine entity is applicable for ESXi. hypervisor
Hypervisor I/O Bandwidth - Read	• Host • Cluster • Virtual Machine	Read data transferred per second in KB/second from Hypervisor. ID: `STATS_HYP_READ_BANDWIDTH` Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Bandwidth - Write	• Host • Cluster • Virtual Machine	Write data transferred per second in KB/second from Hypervisor. ID: `STATS_HYP_WRITE_BANDWIDTH` Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency	• Host • Cluster • Virtual Machine	I/O latency in milliseconds from Hypervisor. ID: `STATS_HYP_AVG_IO_LATENCY` Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency - Read	• Host • Cluster	I/O read latency in milliseconds from Hypervisor. ID: `STATS_HYP_AVG_READ_IO_LATENCY` Note: Cluster entity is applicable only for AHV hypervisor.
Hypervisor I/O Latency - Write	• Host • Cluster	I/O write latency in milliseconds from Hypervisor. ID: `STATS_HYP_AVG_WRITE_IO_LATENCY` Note: Cluster entity is applicable only for AHV hypervisor.

Metric	Entity Type(s)	Description
Hypervisor IOPS	• Host • Cluster • Virtual Machine	Input/Output operations per second from Hypervisor. ID: `STATS_HYP_NUM_IOPS` Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor IOPS - Read	• Host • Cluster • Virtual Machine	Input/Output read operations per second from Hypervisor. ID: `STATS_HYP_NUM_READ_IOPS` Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor IOPS - Write	• Host • Cluster • Virtual Machine	Input/Output write operations per second from Hypervisor. ID: `STATS_HYP_NUM_WRITE_IOPS` Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor Memory Usage (%)	• Host • Cluster • Virtual Machine	Percent of memory used by the hypervisor. ID: `STATS_HYP_MEMORY_USAGE`
Logical Usage	Storage Container	Logical usage of storage (physical usage divided by replication factor). ID: `STATS_UNTRANSFORMED_USAGE`
Memory Usage(%)	• Virtual Machine	Percent of memory used by the hypervisor. Consolidated guest memory usage in percentage. ID: `MEMORY_USAGE_PPM`

Metric	Entity Type(s)	Description
Network Rx Bytes	Virtual Machine	Write data transferred in KB/second from the Storage Controller. ID: `HYPERVISOR_NUM_RECEIVED_BYTES` Note: Virtual Machine entity is applicable for AHV hypervisor.
Network Tx Bytes	Virtual Machine	Write data transferred per second in KB/second. ID: `HYPERVISOR_NUM_TRANSMITTED_BYTES` Note: Virtual Machine entity is applicable for AHV hypervisor.
Physical Usage	• Host • Cluster • Disk • Storage Pool • Storage Container	Actual usage of storage. ID: `STATS_TRANSFORMED_USAGE`
Read IOPS (%)	• Host • Cluster • Disk • Storage Pool	Percent of IOPS that are reads. ID: `STATS_READ_IO_PPM`
Replication Bandwidth - Received	• Cluster • Remote Site • Protection Domain • Replication Link	Replication data received per second in KB/second ID: `STATS_REP_BW_RECEIVED`
Replication Bandwidth - Transmitted	• Cluster • Remote Site • Protection Domain • Replication Link	Replication data transferred per second in KB/second ID: `STATS_REP_BW_TRANSFERRED`

Metric	Entity Type(s)	Description
Replication Bytes - Received	• Cluster • Remote Site • Protection Domain • Replication Link	Number of bytes received. ID: `STATS_REP_NUM_RECEIVED_BYTES` Note: Cluster entity is applicable for AHV hypervisor.
Replication Bytes - Total Received	Replication Link	Total number of bytes received. ID: `STATS_REP_TOT_RECEIVED_BYTES`
Replication Bytes - Total Transmitted	Replication Link	Total number of bytes transmitted. ID: `STATS_REP_TOT_TRANSMITTED_BYTES`
Replication Bytes - Transmitted	• Cluster • Remote Site • Protection Domain • Replication Link	Number of bytes transmitted. ID: `STATS_REP_NUM_TRANSMITTED_BYTES` Note: Cluster entity is applicable for AHV hypervisor.
Storage Controller Bandwidth	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Data transferred in KB/second from the Storage Controller. ID: `STATS_CONTROLLER_BANDWIDTH`
Storage Controller Bandwidth - Read	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Read data transferred in KB/second from the Storage Controller. ID: `STATS_CONTROLLER_READ_BANDWIDTH`

Metric	Entity Type(s)	Description
Storage Controller Bandwidth - Write	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Write data transferred in KB/second from the Storage Controller. ID: `STATS_CONTROLLER_WRITE_BANDWIDTH`
Storage Controller IOPS	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output operations per second from the Storage Controller ID: `STATS_CONTROLLER_NUM_IOPS` Note: The Host entity is applicable for AHV hypervisor only.
Storage Controller IOPS - Read	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output read operations per second from the Storage Controller ID: `STATS_CONTROLLER_NUM_READ_IOPS` Note: The Host entity is applicable for AHV hypervisor.
Storage Controller IOPS - Read (%)	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Percent of Storage Controller IOPS that are reads. ID: `STATS_CONTROLLER_READ_IO_PPM` Note: The Host entity is applicable for AHV hypervisor.

Metric	Entity Type(s)	Description
Storage Controller IOPS - Write	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output write operations per second from the Storage Controller ID: `STATS_CONTROLLER_NUM_WRITE_IOPS`
Storage Controller IOPS - Write (%)	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Percent of Storage Controller IOPS that are writes. ID: `STATS_CONTROLLER_WRITE_IO_PPM`
Storage Controller Latency	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	I/O latency in milliseconds from the Storage Controller. ID: `STATS_CONTROLLER_AVG_IO_LATENCY` Note: • Custer entity is applicable only for AHV hypervisor. • Virtual Machine entity is applicable for AHV and Hyper-V hypervisors.
Storage Controller Latency - Read	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Storage Controller read latency in milliseconds. ID: `STATS_CONTROLLER_AVG_READ_IO_LATENCY` Note: Custer entity is applicable for AHV hypervisor.

Metric	Entity Type(s)	Description
Storage Controller Latency - Write	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Storage Controller write latency in milliseconds. ID: `STATS_CONTROLLER_AVG_WRITE_IO_LATENCY` Note: Custer entity is applicable for AHV hypervisor.
Storage container own usage	Storage Container	Storage container's own usage + Reserved (not used). ID: `NEW_CONTAINER_OWN_USAGE_LOGICAL`
Swap In Rate	Virtual Machine	Rate of data being swapped in. ID: `STATS_HYP_SWAP_IN_RATE` Note: Virtual Machine entity is applicable for ESXi and Hyper-V hypervisors.
Swap Out Rate	Virtual Machine	Rate of data being swapped out. ID: `STATS_HYP_SWAP_OUT_RATE` Note: Virtual Machine entity is applicable for ESXi and Hyper-V hypervisors.
Virtual NIC bytes received packets with error.	Virtual Machine	Virtual NIC bytes received packets with error. `STATS_NETWORK_ERROR_RECEIVED_PACKETS` Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC bytes received rate.	Virtual Machine	Virtual NIC bytes received rate in kbps. `STATS_NETWORK_RECEIVED_RATE` Note: Virtual Machine entity is applicable only for ESXi hypervisor.
Virtual NIC bytes transmitted rate.	Virtual Machine	Virtual NIC bytes transmitted rate in kbps. `STATS_NETWORK_TRANSMITTED_RATE` Note: Virtual Machine entity is applicable for ESXi hypervisor.

Metric	Entity Type(s)	Description
Virtual NIC dropped transmitted packets.	Virtual Machine	Number of dropped transmitted packets by the Virtual NIC. `STATS_NETWORK_DROPPED_TRANSMITTED_PACKETS` Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC receive packets dropped.	Virtual Machine	Number of receive packets dropped by the Virtual NIC. `STATS_NETWORK_DROPPED_RECEIVED_PACKETS` Note: Virtual Machine entity is applicable for ESXi hypervisor.
Write IOPS (%)	• Host • Cluster • Disk • Storage Pool	Percent of IOPS that are writes. ID: `STATS_WRITE_IO_PPM`

Sessions Dashboard

You can view the Sessions dashboard by selecting **Switch Sessions** > **View All Sessions** in the Analysis dashboard. You can also view the Sessions dashboard by closing the current session. The Sessionsdashboard provides a list of sessions including the system-generated session. After an upgrade, the old sessions are also available in the Sessions dashboard.

> Note:
>
> Prism saves all session changes and new sessions automatically including **Session Name** changes, except changes made using specific options that open respective dialog boxes to facilitate the changes.

For example, when you edit a chart using the **Edit** option and the **Edit Chart** dialog box on the Analysisdashboard, you need to click the **Update** button to save the changes made in the dialog box. When you modify the session name or description in the Sessionsdashboard using the **Edit Session Details** action, you need to click the **Update** button to save the changes made in the dialog box. However, if you change the name of the session in the Analysis dashboard, then the change is automatically saved.

> Note:
>
> The updated Analysis dashboard and Sessions dashboard are available only with the Prism Pro and Prism Ultimate licenses. With the Prism Starter license, only the system-generated session is available.

The following figure is a sample view of the Sessions dashboard.

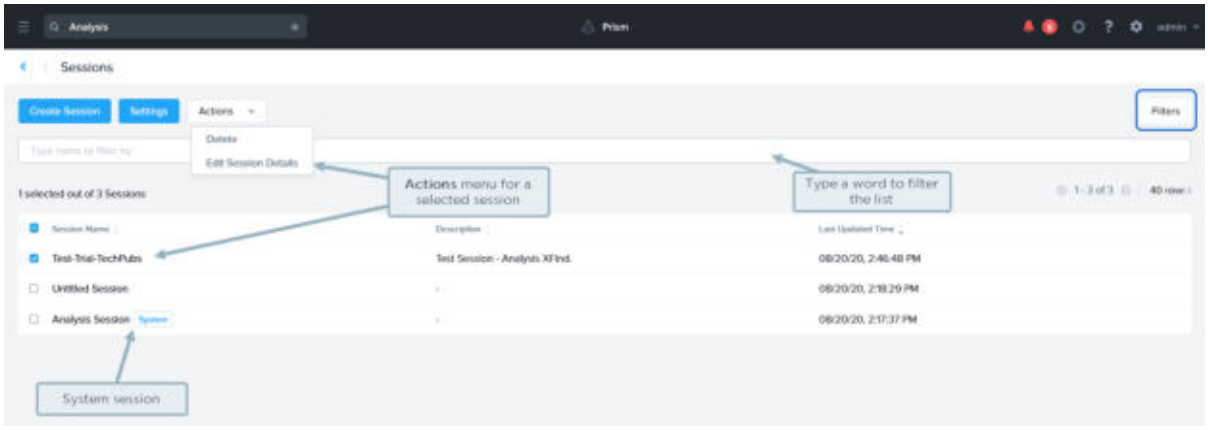

Figure 389: Sessions dashboard

Table 719: Sessions Dashboard Elements

Field	Allows you to...
Create Session	Create a new session.
Settings	Set the default session for the Analysis dashboard.
Actions	• **Delete** a session you select in the list. • **Edit Session Details** of a session you select in the list. The Actions menu is displayed only when you select a session in the list.

Field	Allows you to...
Filters	Filter the list of sessions. You can filter the list by: 1. **Session Name**: Filters on the session name. Select a condition from the pull-down list and enter a string in the field. It returns a list of sessions that satisfy the session name condition/string. Note: Adding a name for a session is mandatory. The default name is Untitled Session. Change the default name to a suitable name with appropriate keywords to facilitate filtering the sessions by session name. 2. **Description**: Filters on the description of the session. Select a condition from the pull-down list and enter a string in the field. It returns a list of sessions that satisfy the description condition/string. Note: Adding a description for a session is not mandatory. However, to facilitate filtering the sessions by description, enter a description with appropriate keywords for the sessions you create.
Type name to filter by	Allows you to filter the sessions by name. Type a word used in the session names that you want to filter to in this search box. Prism Central filters the list of all the sessions to provide a list of the sessions which have the word in their names. The filter is not case sensitive.
Session Name	Enter a name for the session you create or update.
Description	Provide a description for the session.
Last Updated Time	View the date and time when the session was last updated.

Viewing Sessions Dashboard

About this task

You can view the Sessions dashboard and manage the sessions.

To view the Sessions dashboard, do one of the following:

Procedure

1. Click **Switch Session** in the Analysis dashboard and select **View All Sessions**.

   This switches the view to the Sessions dashboard.

2. Click **Actions** in the Analysis dashboard and select **Close**.

   This closes the active session and switches to the Sessions dashboard.

Creating a New Session

About this task

You can create a new session from the Analysis dashboard or the Sessions dashboard.

To create a new session, do the following:

Procedure

1. In the Analysis dashboard, click **Switch Session** > **Create New Session**.

   Note: To create a new session from the Sessions dashboard, click **Create Session** in the Sessions dashboard.

2. In the new session displayed in the Analysis dashboard, change the **Session Name** and add a **Description** help quickly identify the session.

   To change the **Session Name** and add a **Description**, see Editing Session Details on page 770.

What to do next

You can add charts to the new session.

Setting the Default Session View

About this task

You can set the default session view for the Analysis dashboard in the Sessions dashboard.

To set the default session view, do the following:

Procedure

1. Click **Settings** in the Sessions dashboard.

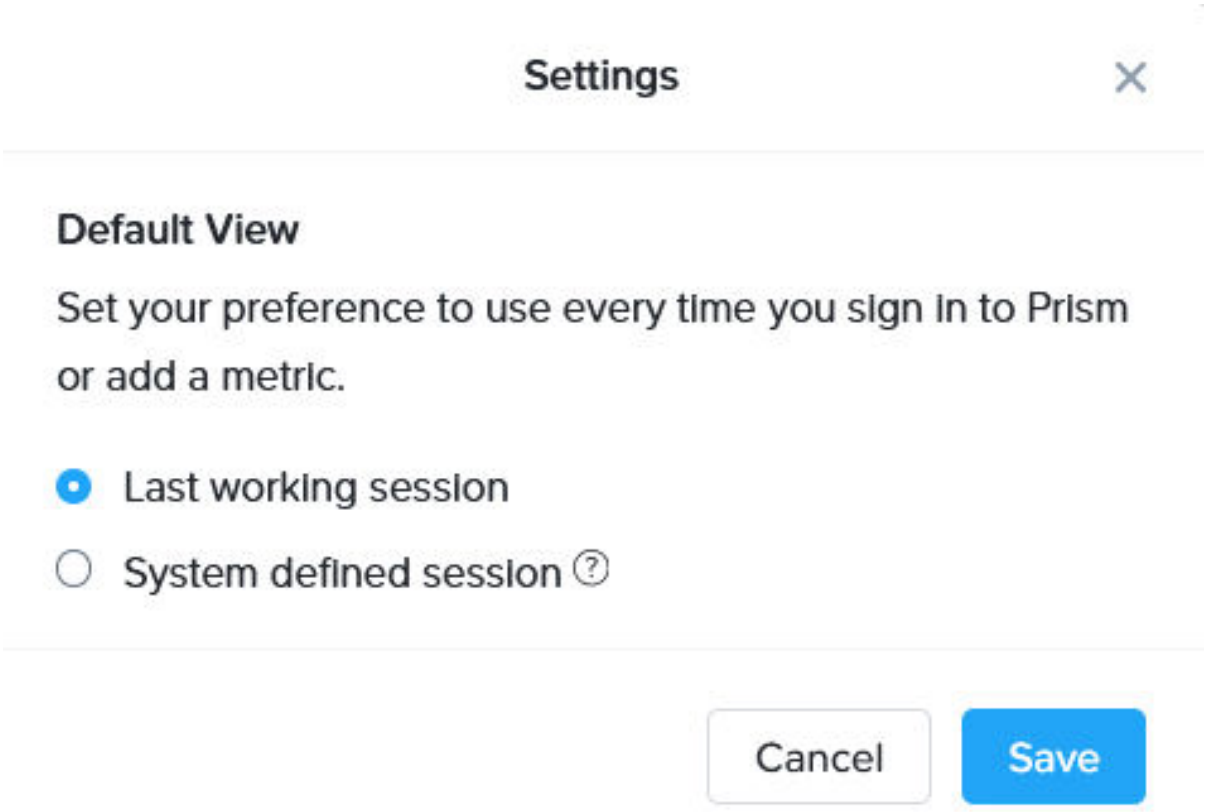

Figure 390: Settings dialog box

2. Select one of the following options:

   » **Last working session**

   This session is displayed by default in the Analysis dashboard. This is the last session that you worked on before signing out of Prism. This option is pre-selected.

   » **System defined session**

   This session is named **Analysis Session**. Prism creates and opens this session after installation of or upgrade to Prism Central 5.18 or later. Prism migrates and persists all the charts that you created before upgrading, to **Analysis Session** in Prism Central 5.18.

   Prism also opens this session in the Analysis dashboard after Prism services are restarted.

3. Click **Save**.

Deleting a Session

About this task

To delete a session in the Sessions dashboard, do the following:

Procedure

1. Select the session you want to delete.
2. Click **Actions** > **Delete**.
3. In the confirmation dialog box, do one of the following:
   - » Click **Confirm** to delete the session.
   - » Click **Cancel** to return to the Sessions dashboard without deleting the session.

Editing Session Details

About this task

To edit the session details like **Session Name** and **Description** of a session, do the following:

Procedure

1. Click **Actions** > **Edit Session Details**.

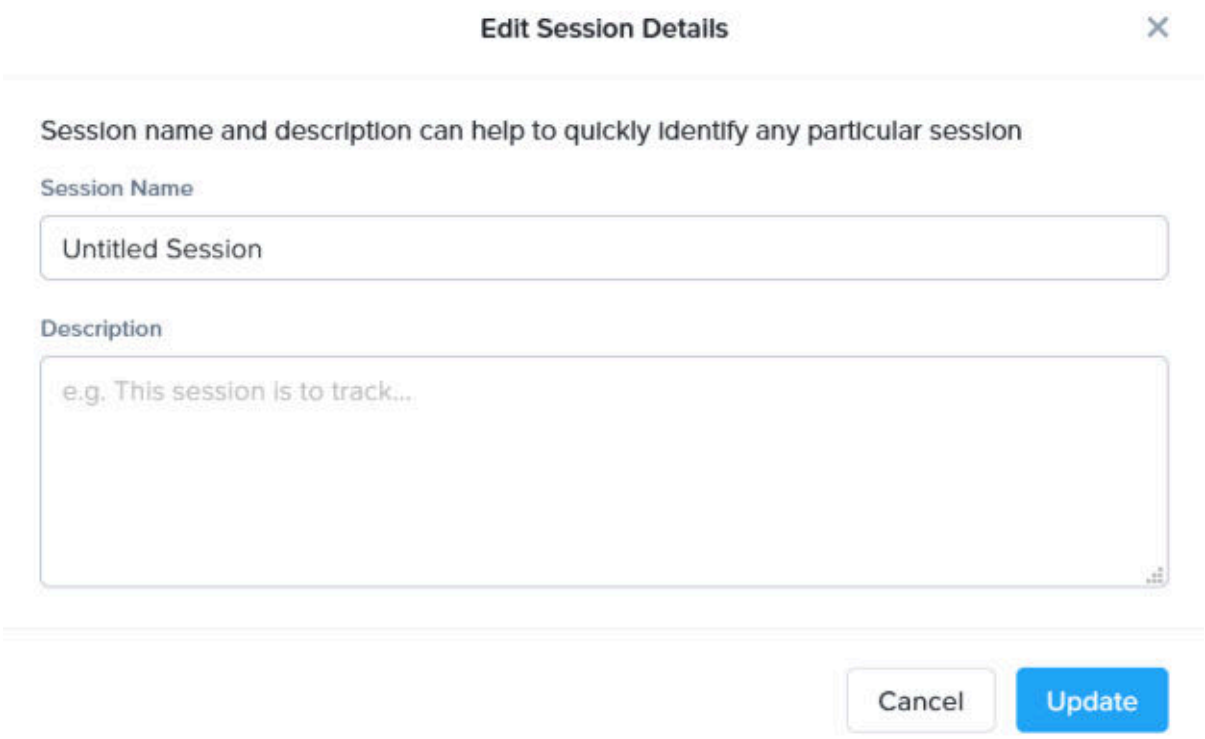

Figure 391: Edit Session Details dialog box

2. Provide appropriate **Session Name** and **Description** to help quickly identify the session.
3. Click **Update**.

Closing the Session

About this task

To close the current session in the Analysis dashboard, do the following:

Procedure

Click **Actions** > **Close**.

Switching Between Sessions

About this task

To switch between sessions, do the following:

Procedure

1. Click **Switch Session** in the Analysis dashboard.
2. From the drop-down list, select the session you want to view.
3. If you do not find the session you want to switch to in the list, do the following:
   a. Click **Switch Session** > **View All Sessions** to go to the Sessions dashboard.
   b. Select the session you want to open in the Analysis dashboard.

Application Discovery

Application discovery provides the capability to discover applications in a specified set of clusters, identify in which VM each application runs and what ports it uses, and publish that data to an internal database which is accessible through an external API call. You can monitor applications in Nutanix, non-Nutanix (vCenter managed), or mixed (cluster with mixed hypervisors) environment.

> Note: When you add a mixed cluster for application discovery, only ESXi hypervisor nodes are enabled for discovery.

Application discovery sends data about the VM network traffic stats to a SaaS-like service in the cloud for analysis. The information sent to the cloud includes network flow data (source IP, destination IP, destination port, protocol, packet count, byte count) and VM data (cluster UUID, source VM name, UUID, VLAN ID, and IP addresses).

- To configure application discovery:
  - enable application discovery (see Enabling Application Discovery on page 781)
  - connect to the cloud-based discovery service (see Connecting to the Cloud on page 787)
  - add clusters to monitor (see Adding (Removing) Clusters to Monitor on page 788)
  - authenticate with vCenter for ESXi cluster access (see Configuring vCenter Authentication on page 790)

- To generate current findings at any time:
  - start an application discovery (see Discovering Applications on page 791)
  - publish data for identified applications (see Publishing Applications on page 792)
  - create policies for unidentified applications (see Creating Discovery Policies on page 792)
- To view discovery findings, see the application discovery dashboard (see Application Discovery View on page 772).

Application Discovery View

The application discovery dashboard allows you to view information about discovered applications and associated policies.

To access the application discovery dashboard, select **Operations** > **App Discovery** from the entities menu (see Entities Menu on page 13). The dashboard includes seven tabs at the top (**Summary**, **Discovered Apps**, **Published Apps**, **Policies**, **App Relationships**, **Monitored Clusters**, and **Alerts**) with a display area below for the selected tab.

Note:

- This section describes the information and options that appear in the application discovery dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in various ways.
- If application discovery is not enabled, an enablement page appears instead of the dashboard (see Enabling Application Discovery on page 781).

Summary Tab

Clicking the **Summary** tab displays the following four widgets:

- **Discovered Highlights**: Displays the total number of application instances, the number of identified (known application to port mapping) and unidentified instances, the number of application types, and the number of ports.
- **Status & Settings**: Displays the number of monitored clusters, the number of discovery policies, and the status (active or inactive) of the discovery process, vCenter authentication, vCenter monitoring, and cloud connection. Click either **View Details** to see the vCenter authentication and monitoring details (see Configuring vCenter Authentication on page 790) or click**Manage** to see the cloud connection details (see Connecting to the Cloud on page 787). If vCenter monitoring is not configured, an **Add vCenter** link appears (instead of the **View Details** link); click the link to add clusters from any configured vCenter instances (see Adding (Removing) Clusters to Monitor on page 788).
- **Top Discovered Apps**: Lists the applications with the most instances.
- **Top Discovered Ports**: Lists the ports with the most instances.

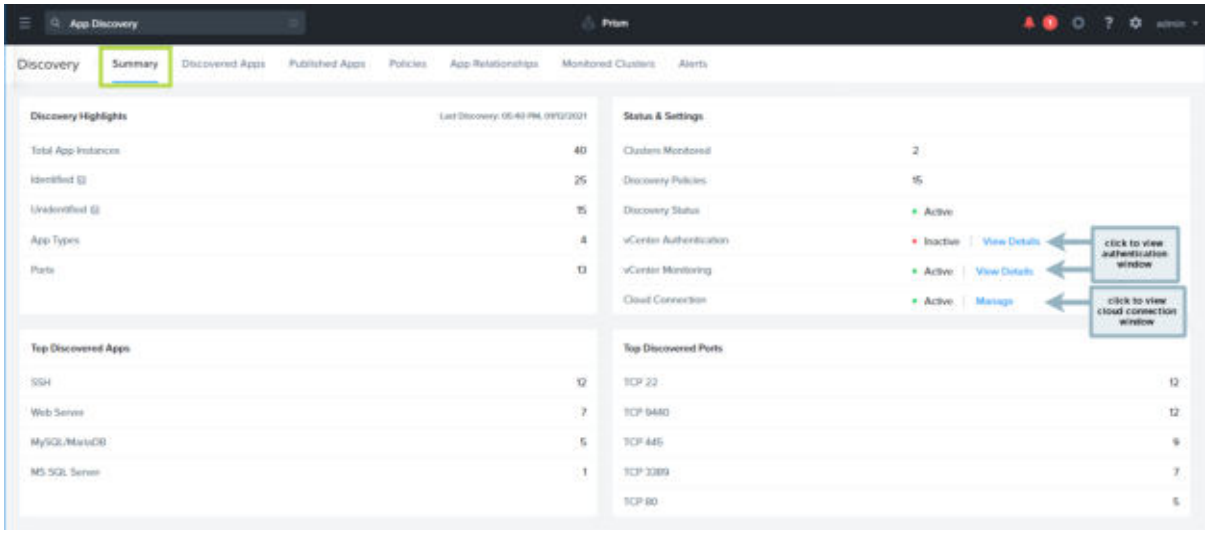

Figure 392: App Discovery Summary Tab

Discovered Apps Tab

Clicking the **Discovered Apps** tab, which appears by default when you first open the page, displays a list of the discovered applications.

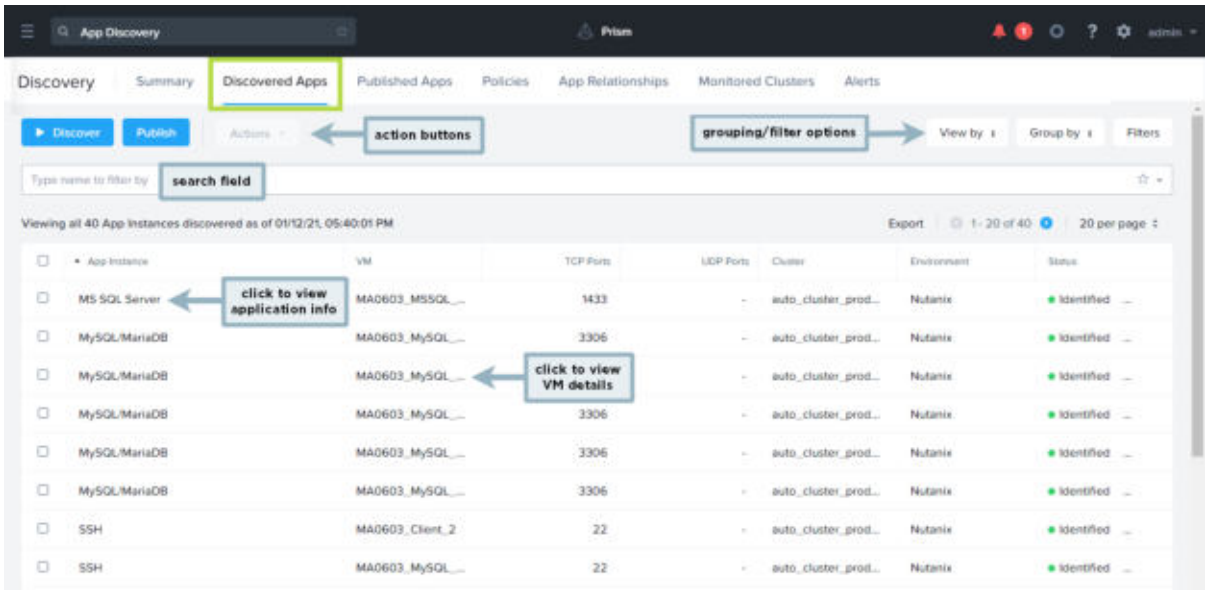

Figure 393: Discovered Apps Tab

The following table describes the fields that appear in the discovered applications list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 720: Discovered Apps Fields

Parameter	Description	Values
App Instance	Displays the application name for the instance. Click the name to display the Application Info window. The window displays the application name, VM name, TCP ports, and UDP ports.	(app name)
VM	Displays the VM name. Click the name to display the details page for the VM (see VM Details View on page 96).	(VM name)
TCP Ports	Displays one or more TCP ports used by the application.	(port numbers)
UPD Ports	Displays one or more UDP ports used by the application.	(port numbers)
Cluster	Displays the name of the cluster in which the application runs. Click the name to view the details for that cluster (see Cluster Details View on page 205).	(cluster name)
Environment	Displays whether the cluster resides in a Nutanix or non-Nutanix (vCenter managed) environment.	Nutanix, vCenter
Status	Indicates whether the instance is an identified application type and whether it is newly discovered ("new app" meaning discovered during the most recent pass) or discovered previously ("existing app" meaning discovered in an earlier pass). An "identified" application is one that is a known type. The discovery service can classify many applications, but for those it cannot identify, you can assign them to a type through a custom discovery policy (see Creating Discovery Policies on page 792).	Identified, Unidentified

You can filter the list either by entering a string in the search field or by opening the **Filter** pane (open by default) to select a filter option. The following table describes the filter options. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 721: Filter Pane Fields

Parameter	Description	Values
Status	Filters by whether the instance has been identified as a known application type.	Identified, Unidentified
App State	Filters by whether the instance is new.	New, Existing

Parameter	Description	Values
Ignored	Filters for ignored applications.	Show Ignored Apps (check box)
App Instance	Filters on the application instance name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.	(app instance name string)
Cluster	Filters on the cluster name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.	(cluster name string)
Environment	Filters based on whether the applications are in a Nutanix or non-Nutanix (vCenter) managed environment.	Nutanix, vCenter
VM	Filters on the VM name (same options as App Instance).	(VM name string)
TCP Ports	Filters on the TCP port numbers. Enter a range in the "**From <low> to <high>**" field.	([xx] to [yy] range)
UDP Ports	Filters on the UDP port numbers. Enter a range in the "**From <low> to <high>**" field.	([xx] to [yy] range)

To download the table contents in CSV format, click the **Export** link.

You can group the list in the following ways:

- The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.
- To group the list by application name, select **App Instance** from the **Group By** pull-down menu.

The discovered apps tab includes the following action buttons:

- Click the **Discover** button to start a discovery search (see Discovering Applications on page 791).
- Click the **Publish** button to publish application information (see Publishing Applications on page 792).
- An **Actions** menu appears when one or more instances are selected. The menu allows you to set the selected instances to [**Identify**|**Unidentify**] or [**Ignore**|**Unignore**] status (as appropriate). The available actions appear in bold; other actions are grayed out.

Published Apps Tab

Clicking the **Published Apps** tab displays a list of the published applications. A "published" application is one that is made available to the API. Publishing involves copying the data for the identified (and not ignored) applications into an internal database that you can access through an external API call. An "identified" application is one where the VM in which it runs is known and the used ports are mapped.

- The listed information for published applications is the same as for discovered applications, except there is no status field.

- The **Focus** , **Group**, and **Export** options are the same, but the **Filter** pane includes just **App Instance** and **VM** options.

The published apps tab includes the following action buttons:

- Click the **Get API Link** button to copy the API link to your clipboard. The link is the access point to consume the published applications through the API.
- An **Actions** menu appears when one or more instances are selected; select **Delete** from the pull-down list to delete the selected instances from the published list.

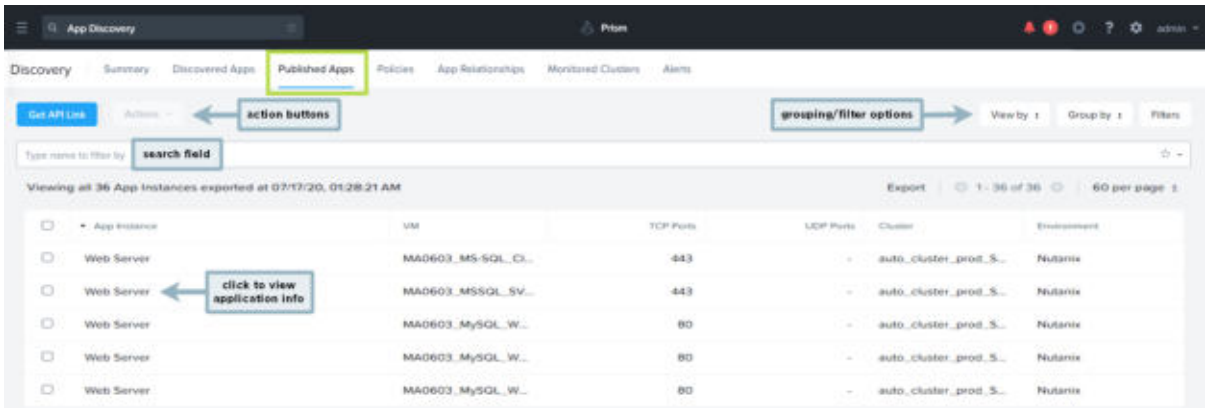

Figure 394: Published Apps Tab

Policies Tab

Clicking the **Policies** tab displays a list of application discovery policies. The policies define the application types.

> Note: Existing discovered applications are remapped whenever there is a policy action change, based on the existing VM traffic statistics.

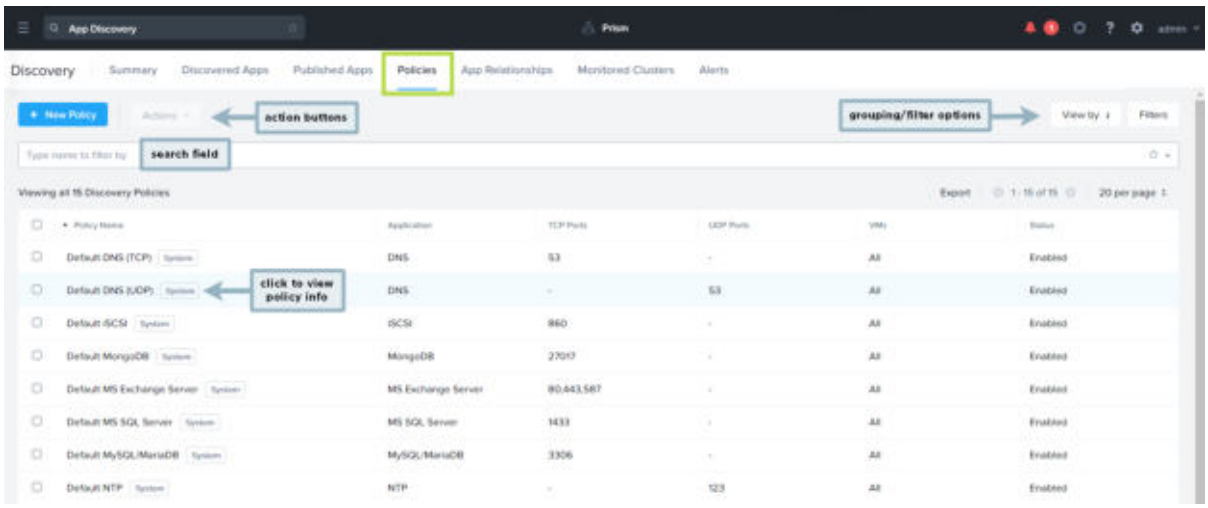

Figure 395: Policies Tab

The following table describes the fields that appear in the policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 722: Policies Fields

Parameter	Description	Values
Policy Name	Displays the name of the policy. Click the name to display the Edit Discovery Policy window for user-defined policies or the View Discovery Policy window for system-defined policies (see Modifying (Viewing) Discovery Policies on page 794).	(policy name)
Application	Displays the application type. Application discovery can identify various common applications by default (referred to as system-defined applications) such as DNS, NTP, SSH, Web Server, MS Exchange Server, MS SQL Server, Oracle DB, SAP HANA. For user-defined applications, the user provides the name (see Creating Discovery Policies on page 792).	(application type)
TCP Ports	Displays one or more TCP ports used by the application.	(port numbers)
UPD Ports	Displays one or more UDP ports used by the application.	(port numbers)
VMs	Displays the scope of the policy, that is to which VMs the policy applies. A value of "All" indicates it applies to all VMs. One or more names indicates that the policy applies to the named VMs only.	All, (VM names)
Status	Displays the current application operational status.	Enabled, Disabled

To download the table contents in CSV format, click the **Export** link.

You can filter the list either by entering a string in the search field or by opening the **Filter** pane to select one or both of the filter options.

- **Policy Name**: Check the box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field to filter the list for policies with that string in the name.
- **Status**: Check the **Enabled** or **Disabled** box to filter the list for enabled or disabled policies.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

The policies tab includes the following action buttons:

- Click the **New Policy** button to create a custom discovery policy (see Creating Discovery Policies on page 792).
- An **Actions** menu appears when one or more policies are selected. The menu allows you to **Enable**, **Disable**, or **Delete** the selected policies (as appropriate). You can only delete user-defined policies (not system policies).

App Relationships Tab

Clicking the **App Relationships** tab displays the top 20 applications based on the number of entities that are communicating with these applications. The displayed information includes the VMs where these applications are running, along with the number of entities that are communicating with these VMs.

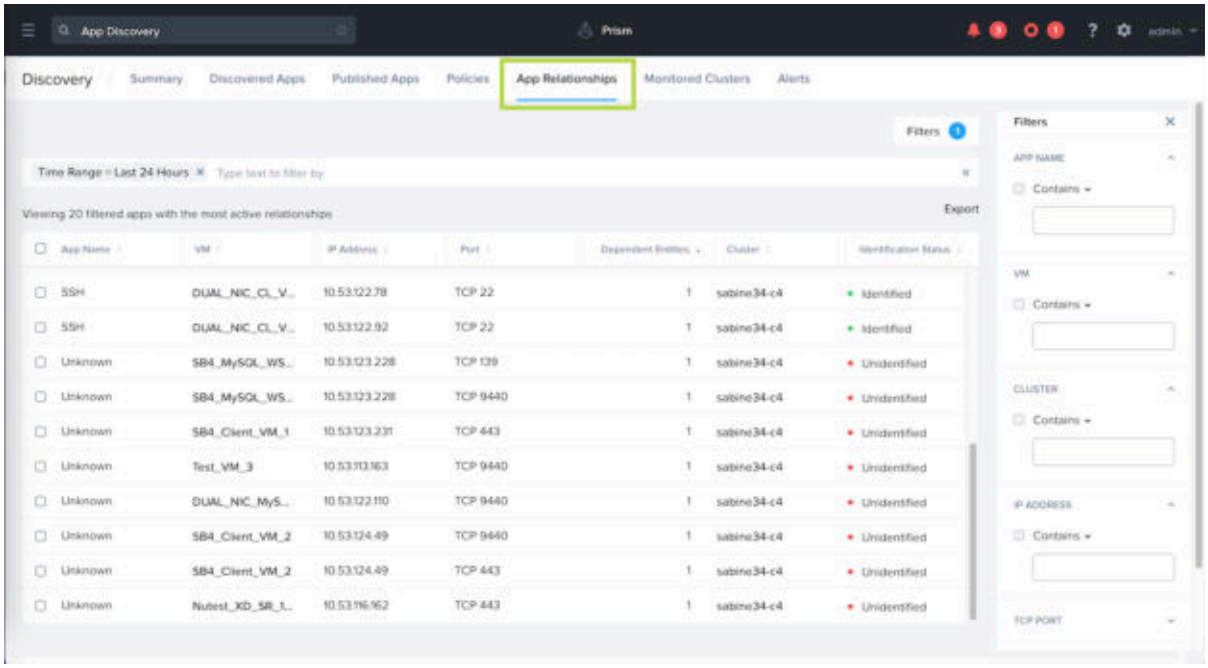

Figure 396: App Relationships Tab

The following table describes the fields that appear in the app relationships list.

Table 723: App Relationships Fields

Parameter	Description	Values
App Name	Displays the application name.	(app name)
VM	Displays the VM name. Click the name to display the details page for the VM (see VM Details View on page 96).	(VM name)
IP Address	Displays the IP Address of the VM.	(IP Address)
Ports	Displays the port used by the application.	(port numbers)
Dependent Entities	Displays the number of entities that are communicating with the application.	(number of dependent entities)
Cluster	Displays the name of the cluster in which the application runs. Click the name to view the details for that cluster (see Cluster Details View on page 205).	(cluster name)

Parameter	Description	Values
Identification Status	Indicates whether the application is an identified application type. An identified application is one that is a known type per application discovery service. You can use custom discovery policy to assign a type to the unidentified applications listed under App Relationships tab. If an application is displayed as unidentified regardless of it belonging to the known types, try refreshing the list of discovered applications. You can refresh the list by running the Discover operation from Discovered Apps tab.	Identified, Unidentified

You can filter the list by opening the **Filter** pane to select a filter option. The following table describes the filter options. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

> Note: When you use filters for information displayed on the App Relationships tab, the filters fetch the latest discovery and communication data and accordingly display the result set.

Table 724: Filter Pane Fields

Parameter	Description	Values
App Name	Filters on the application name. Select a condition from the pull-down list (**Contains**, **Equal to**, **Not equal to**, **Doesn't contain**, **Starts with**, **Ends with**) and enter a string in the field.	(app name string)
VM	Filters on the VM name. (same options as App Name).	(VM name string)
Cluster	Filters on the cluster name (same options as App Name).	(cluster name string)
IP Address	Filters on the IP address (same options as App Name).	(IP address string)
TCP Port	Filters on the TCP port number.	(TCP port number)
UDP Port	Filters on the UDP port number.	(UDP port number)
Time Range	Filters on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

To download the table contents in CSV format, click the **Export** link.

Monitored Clusters Tab

Clicking the **Monitored Clusters** tab displays the list of clusters that application discovery monitors when it does a discovery search.

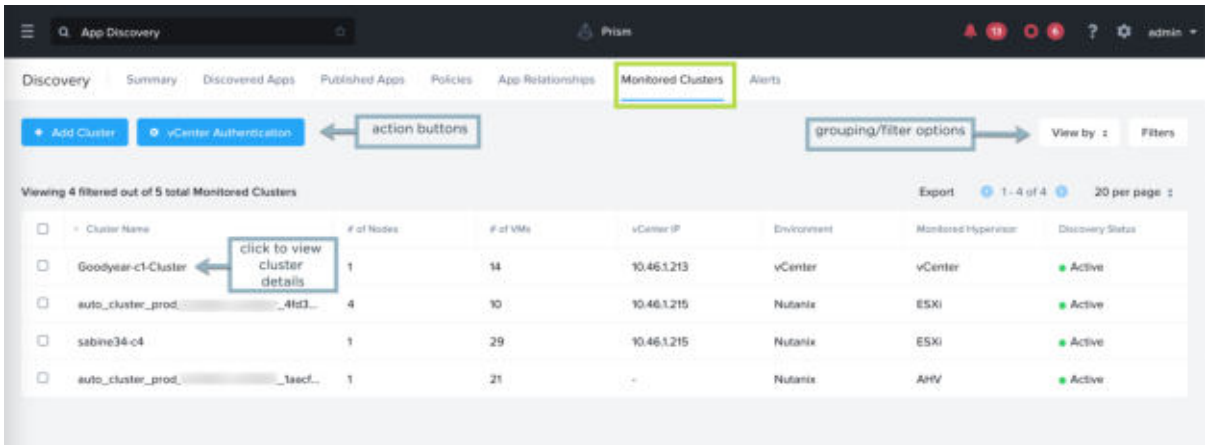

Figure 397: Monitored Clusters Tab

The following table describes the fields that appear in the monitored clusters list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 725: Monitored Clusters Fields

Parameter	Description	Values
Cluster Name	Displays the cluster name. Click the name to displays the details page for the cluster (see Cluster Details View on page 205).	(cluster name)
# of Nodes	Displays the number of nodes in the cluster.	(integer)
# of VMs	Displays the number of VMs in the cluster.	(integer)
vCenter IP	Displays the IP address of the vCenter for this cluster.	(IP address)
Environment	Displays whether the cluster resides in a Nutanix or non-Nutanix (vCenter managed) environment.	Nutanix, vCenter
Monitored Hypervisor	Displays the monitored hypervisor in the cluster.	ESXi, AHV, vCenter
Health Status	Displays the current VM network traffic monitoring status.	Active, Inactive

To download the table contents in CSV format, click the **Export** link.

You can filter the list either by entering a string in the search field or by opening the **Filter** pane to select a filter option. The following table describes the filter options. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 726: Filter Pane Fields

Parameter	Description	Values
vCenter IP	Filters on the vCenter IP address (same options as Cluster Name).	(IP address string)
Health Status	Filters on the selected health status. Check the box for each status to include.	Deploying, Deployed, Running, Stopped, Failed

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

The monitored clusters tab includes the following action buttons:

- Click the **Add Cluster** button to add a cluster for monitoring (see Adding (Removing) Clusters to Monitor on page 788).
- Click the **vCenter Authentication** button to authenticate application discovery on the monitored clusters with vCenter (see Configuring vCenter Authentication on page 790).
- A **Remove** button appears when one or more clusters are selected; click the button to remove the selected clusters from application discovery monitoring.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard (List tab), except it is filtered to display just application discovery-related alerts (see Alerts Summary View on page 261).

Enabling Application Discovery

Enabling application discovery (disabled by default) involves connecting to the discovery service in the cloud, specifying clusters to monitor, and authenticating to the vCenter instances for the specified clusters.

About this task

When you select application discovery for the first time, a wizard appears to guide you through the setup process. To configure application discovery initially, do the following:

Before you begin

Complete the Prerequisites on page 786.

Procedure

1. Select **Operations** > **App Discovery** from the entities menu (see Entities Menu on page 13).

   If application discovery is not enabled, a "Welcome to App Discovery" screen appears.

2. In the "Welcome to App Discovery" screen, check the "I have increased the PC VM's memory and vCPUs" box (as specified in the prerequisites) and then click the **Enable App Discovery** button.

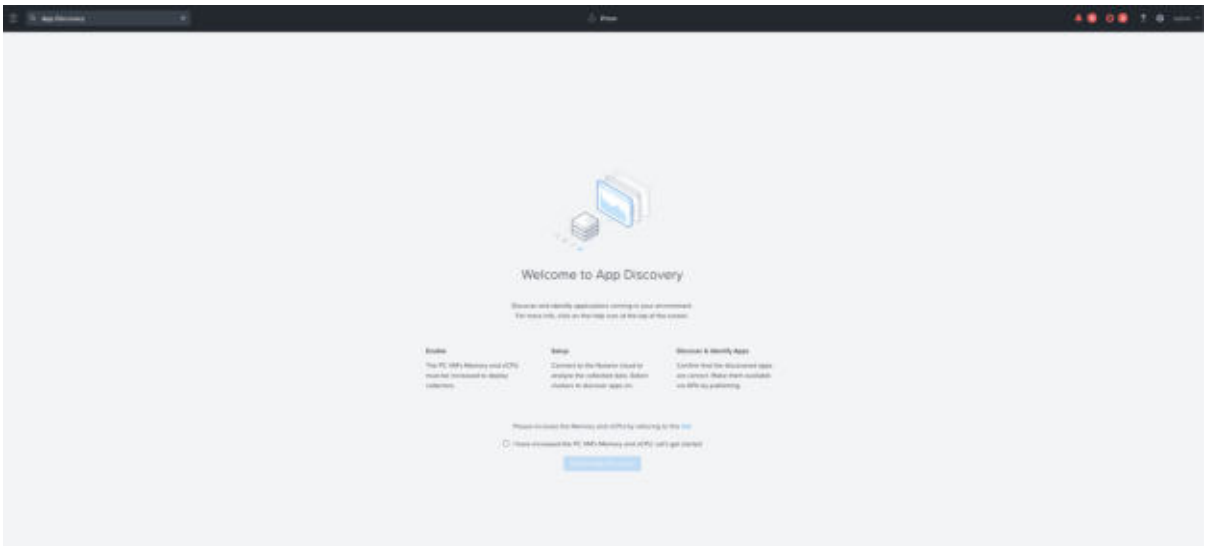

Figure 398: Welcome to App Discovery screen

3. In the "App Discovery Setup: Cloud Connection" screen, enter your API key and key authentication values in the indicated fields and then click the **Next** button.

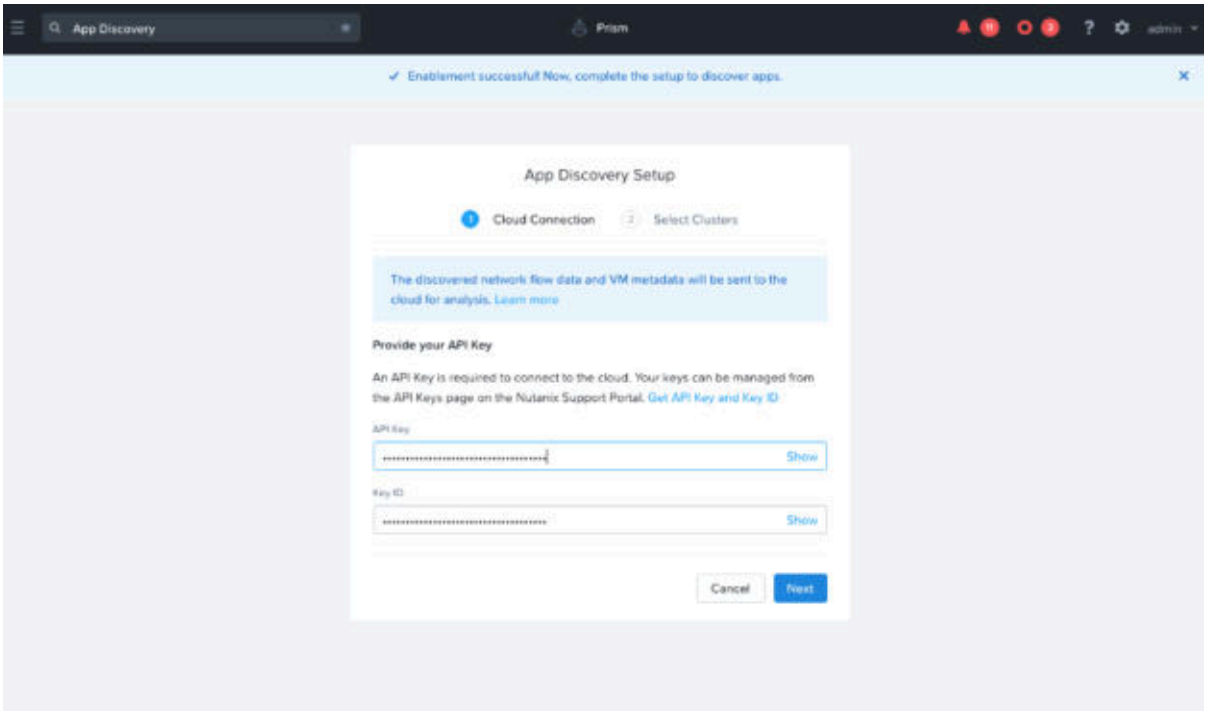

Figure 399: App Discovery Setup: Cloud Connection screen

4. In the "App Discovery Setup: Select Clusters" screen, do the following:

   » If the clusters to select are in a Nutanix environment, continue with this step.

   » If the clusters to select are in a non-Nutanix (vCenter managed) or mixed environment, click the **Configure vCenter Monitoring** link and then skip to step 5.

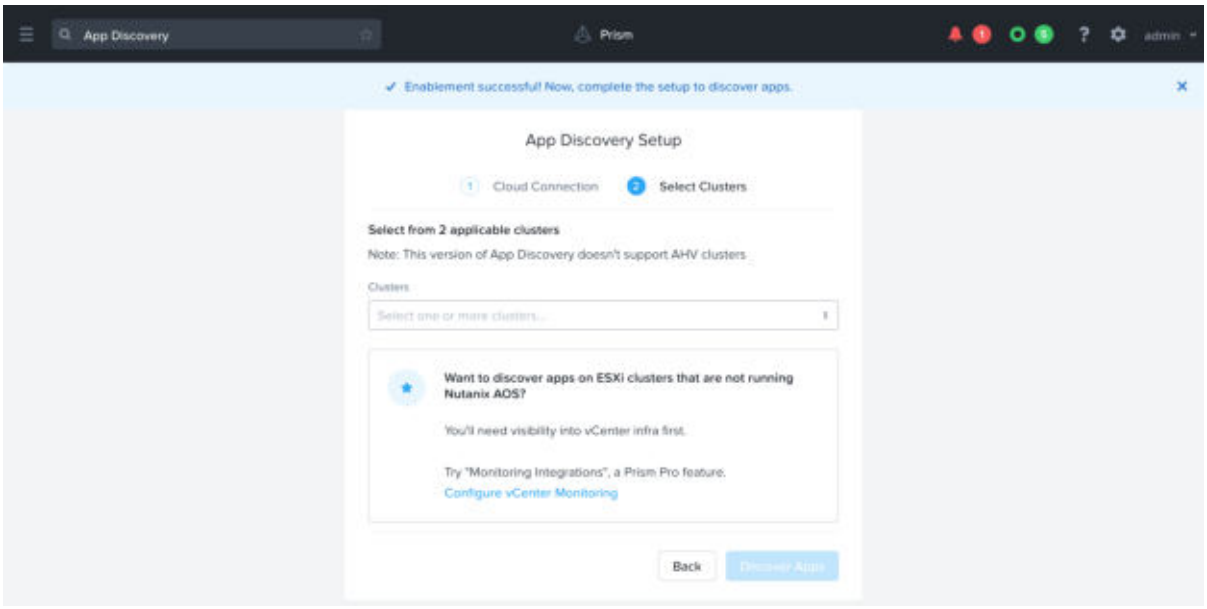

Figure 400: App Discovery Setup: Select Clusters (initial)

   a. Check the boxes for the target clusters from the pull-down list.

   b. Enter the vCenter administrator credentials (user name and password) in the indicated fields. The administrator must have write permission.

      An entry for each registered vCenter appears here (see Registering vCenter Server (Prism Central) on page 709).

   c. Click the **Discover Apps** button.

      This starts a discovery search. When it completes, the App Discovery window and dashboard appear (see Application Discovery View on page 772). This completes the setup for a Nutanix environment.

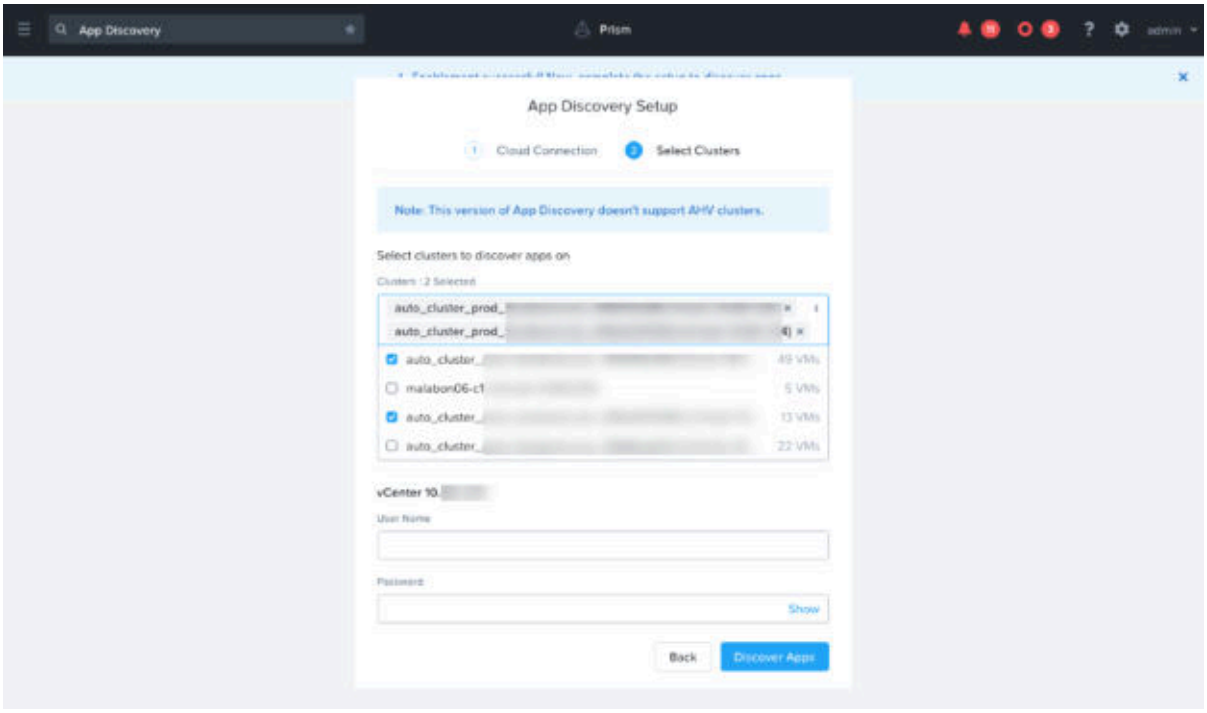

Figure 401: App Discovery Setup: Select Clusters

5. [vCenter only] select **vCenter** in the **External Entity Type** field and then click the **Enable** button.

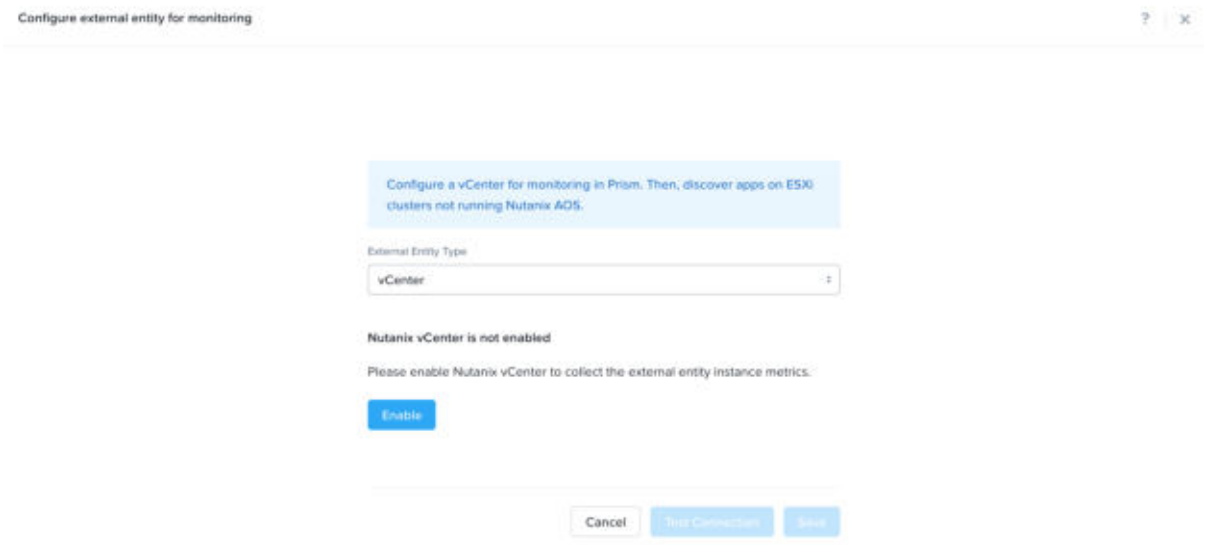

Figure 402: Enable vCenter Discovery

6. [vCenter only] Do the following in the indicated fields to configure a vCenter instance:

   a. **External Entity Type**: Select **vCenter** from the pull-down menu.

   b. **vCenter Host**: Enter the IP address of the vCenter host VM.

      If the IP address does not resolve automatically to the proper port number, enter the address in the form `ip_address:port#`. The standard vCenter Server port is 443.

   c. **Username**: Enter the vCenter administrator user name.

   d. **Password**: Enter the vCenter administrator password.

   e. **SSL Verify**: Select the SSL mode to use (**Verify** or **Don't Verify**) when connecting to the target.

      This setting determines whether to verify (or not verify) the target's SSL certificate as part of the connection process.

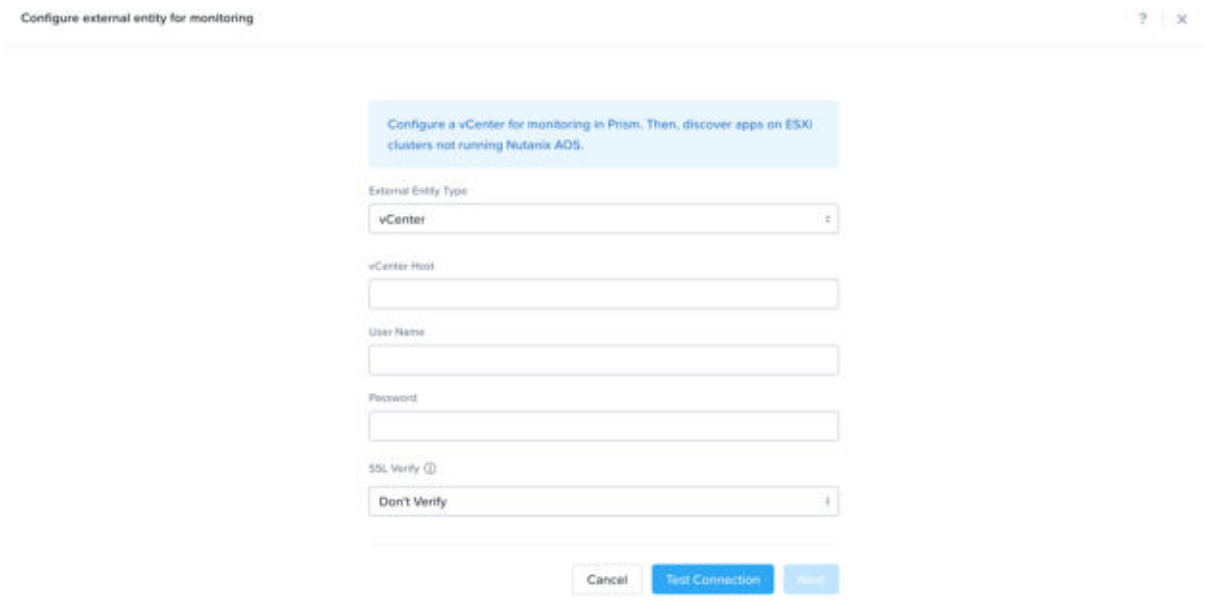

Figure 403: Configure vCenter Instance

7. [vCenter only] Click the **Test Connection** button to verify (test) the connection.

   A "Test connection success" message appears after a successful test. After a successful test, click the **Next** button.

8. [vCenter only] select the clusters to monitor.

   A list of discovered clusters for the vCenter instance appears. Select (check the boxes for) the clusters you want to monitor from the list and then click the **Save** button. The instance

configuration and data collection begins. When it completes (which can take a few minutes), a "vCenter Configuration successful" message appears.

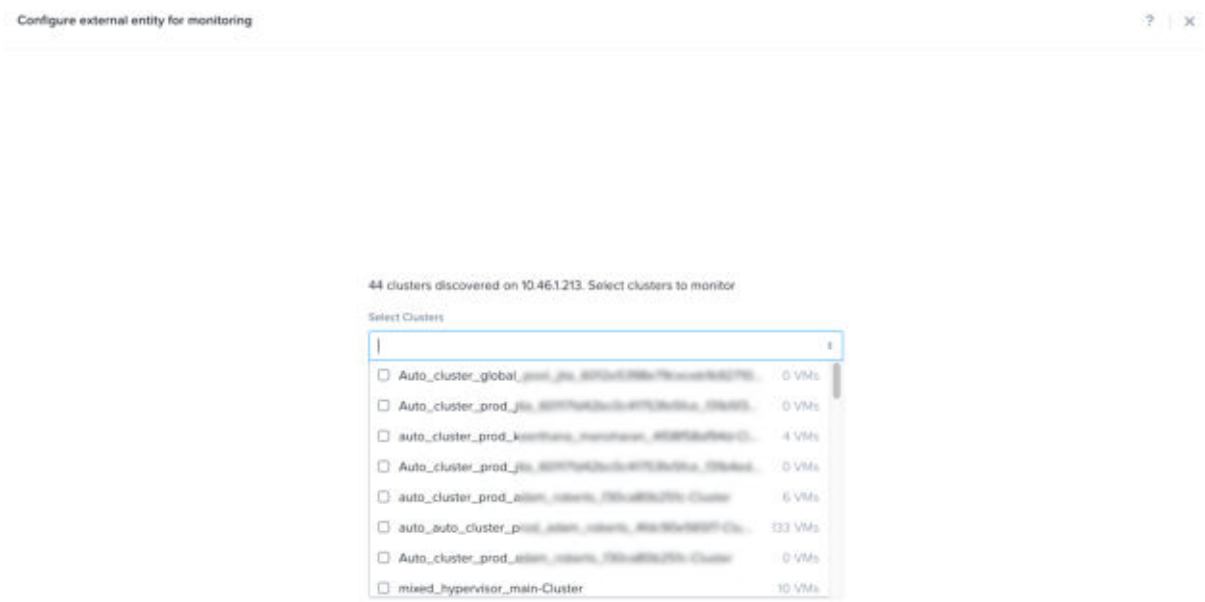

Figure 404: Select Clusters (vCenter)

9. [vCenter only] After the vCenter configuration completes, click the **Back to App Discovery** button.

   Clicking this button takes you to the monitored clusters page and opens a cluster selection window. The vCenter clusters specified in step 8 appear in the list for you to select.

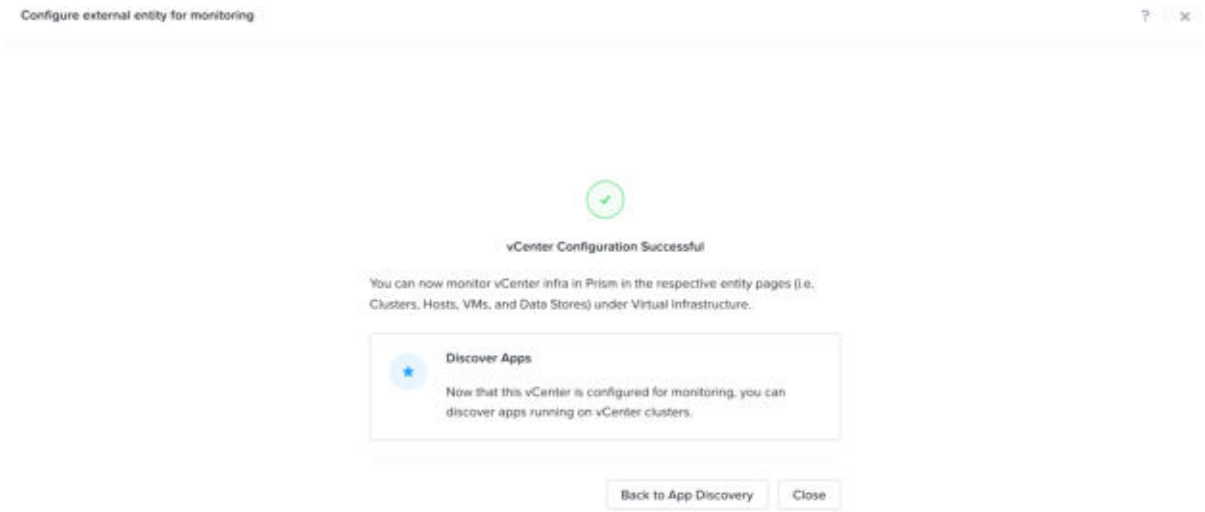

Figure 405: vCenter Configuration Complete

Prerequisites

Address the following before enabling application discovery:

- Sufficient Prism Central VM resources. Application discovery is one of the Prism Ultimate license features that requires additional Prism Central VM resources (more memory and vCPUs). The requirements are listed in KB 8932. If your Prism Central VM resources are below the recommended amount, manually increase the memory and vCPUs to the recommended amount.
- Internet connection. An Internet connection is required to connect to the Nutanix cloud.
- Up and running firewall in Prism Central.
- VDS enabled on vCenter. This requires a vSphere version (Enterprise Plus license) that supports vSphere Distributed Switch (VDS) and netflow version 10.
- My Nutanix account that is either admin role or account admin.
- API key and key ID. Log on to the Nutanix support portal and get your API key information. See the *Nutanix Licensing Guide* for instructions on how to generate a new key or view information about an existing key.

  Each Prism Central instance requires a unique API key/ID pair. If you have multiple Prism Central instances (or a key was revoked for any reason), create additional API key/ID pairs as needed so you have a unique and valid pair for each Prism Central instance.

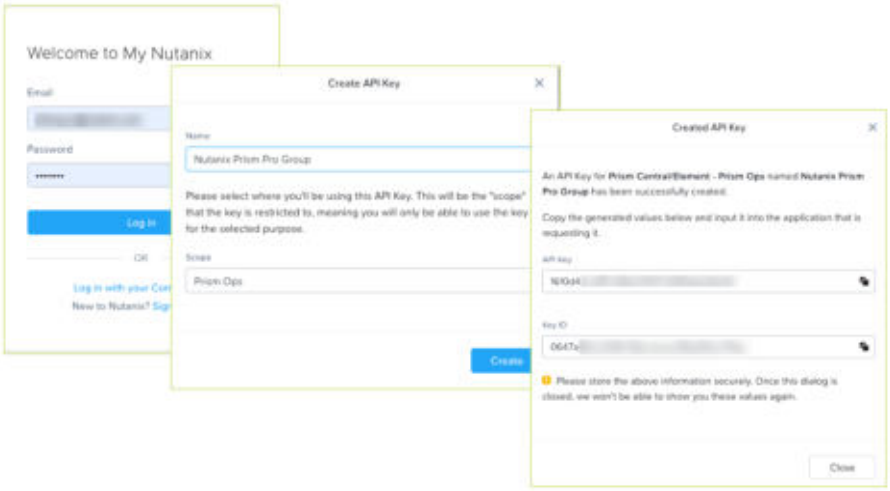

Figure 406: Create API Key Flow (example)

Connecting to the Cloud

Using application discovery requires connecting to a cloud-based discovery service.

About this task

Nutanix provides a SaaS-like service in the cloud that analyzes application data from the monitored clusters. The connection is established during initial setup (see Enabling Application Discovery on page 781), but to update the connection after setup, do the following:

Before you begin

Log on to the Nutanix support portal and get your API key information. See the *Nutanix Licensing Guide* or *License Manager Guide* for instructions on how to generate a new key or view information about an existing key.

Procedure

1. Go to the **Summary** tab of the application discovery dashboard and click **Manage** in the Cloud Connection line of the Status & Settings widget (see Application Discovery View on page 772).
2. In the Cloud Connection window, do the following in the indicated fields:
   a. **API Key**: Enter the API key value.
   b. **Key ID**: Enter the key identification number.
   c. Click the **Save** button.

   Only one API key/ID pair is allowed per Prism Central instance. If you have multiple Prism Central instances (or your initial key was revoked for any reason), create additional API key/ID pairs as needed (through the Nutanix support portal) so you have a unique and valid pair for each Prism Central instance.

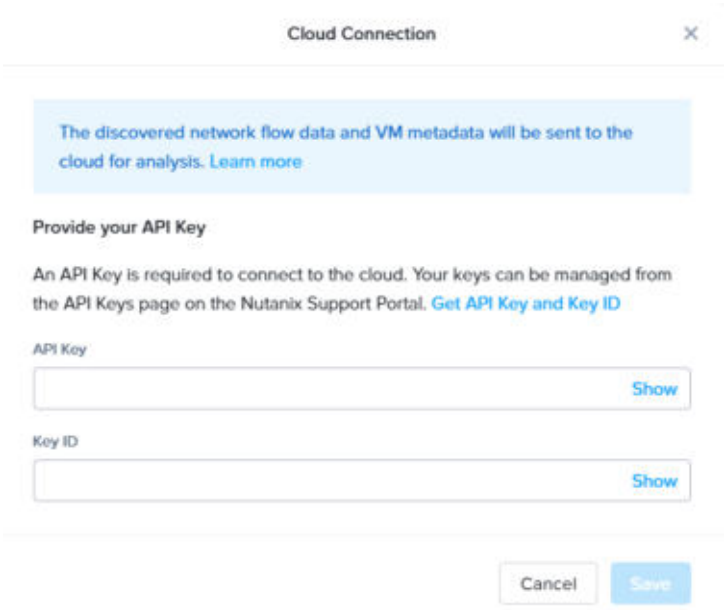

Figure 407: Cloud Connection Window

Adding (Removing) Clusters to Monitor

Using application discovery requires first specifying the clusters to monitor.

About this task

Specifying clusters to monitor is done during initial setup (see Enabling Application Discovery on page 781), but you can add or remove clusters after setup at any time.

Procedure

- To add a cluster to the monitored list, do the following:

  a. Go to the **Monitored Clusters** tab of the application discovery dashboard and click the **Add Cluster** button (see Application Discovery View on page 772).

  b. In the Add Cluster to App Discovery window, check the boxes for the target clusters from the pull-down list.

     If you add a cluster from a vCenter server, you must enter the vCenter administrator credentials (user name and password) in the indicated fields. The administrator must have write permission.

     Click the **Add** button. Network data collection begins immediate and may take some time to complete when multiple clusters are added.

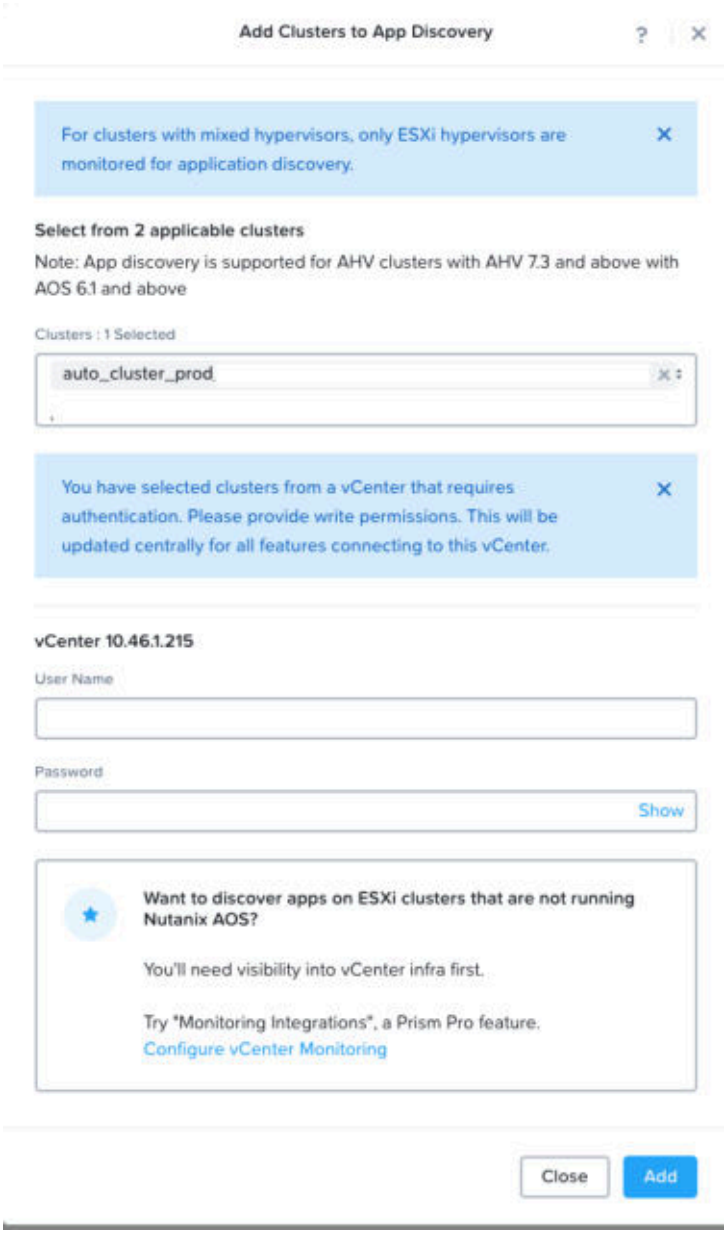

Figure 408: Add Clusters to App Discovery Window

c. Run discovery to identify applications on the new clusters (see Discovering Applications on page 791).

- To remove a cluster from the monitored list, select the target cluster(s) in the list and then click the **Remove** button.

  Note: Removing a cluster stops application monitoring for that cluster. You can effectively disable application monitoring by removing all clusters, which means no data is sent between Prism Central and the Nutanix cloud.

Configuring vCenter Authentication

Using application discovery with ESXi clusters requires first authenticating with the associated vCenter.

About this task

Authenticating with vCenter is done during initial setup (see Enabling Application Discovery on page 781), but to update authentication after setup, do the following:

Procedure

1. Go to the **Monitored Clusters** tab of the application discovery dashboard and click the **vCenter Authentication** button (see Application Discovery View on page 772).
2. In the vCenter Authentication window, do the following in the indicated fields:
   a. **User Name**: Enter the user name for a vCenter administrator or any user with write permission to the clusters.
   b. **Password**: Enter the user password.
   c. Click the **Save** button.

   An entry for each registered vCenter appears here (see Registering vCenter Server (Prism Central) on page 709). The authentication and monitoring status for each vCenter is

displayed. If a vCenter is not monitored currently, you can enable monitoring by clicking the **Configure** button (see Configuring Application Monitoring on page 809).

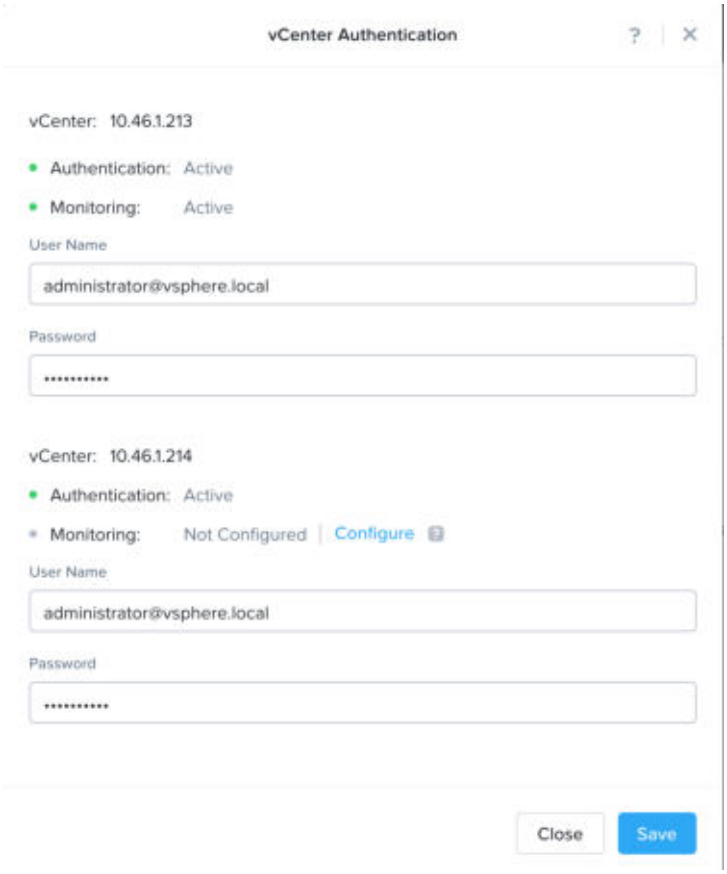

Figure 409: vCenter Authentication Window

Discovering Applications

You can check for the current applications on the monitored clusters at any time.

About this task

To start an application discovery search, do the following:

Procedure

1. Go to the **Discovered Apps** tab of the application discovery dashboard and click the **Discover** button (see Application Discovery View on page 772).
   The search begins immediate and may take a few minutes to complete depending on the number and size of the monitored clusters. When the search completes, the **App Discovery** window appears. The window displays the search status (successful, unsuccessful); number

of discovered, identified, and unidentified applications; number of application types; and number of new application instances.

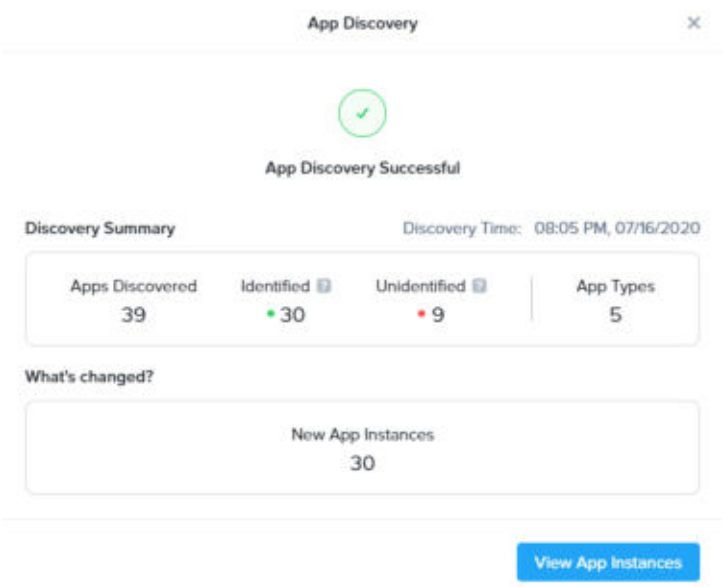

Figure 410: App Discovery Window

2. Click the **View App Instances** button to redisplay the tab.

   If some applications are misclassified, you can modify the results by manually identifying, unidentifying, or ignoring (hiding) selected applications (see the **Actions** menu options in Application Discovery View on page 772).

Publishing Applications

You can publish the application discovery findings to a database.

About this task

Publishing means copying the VM and port mapping data for the identified applications to an internal database, which is accessible through an external API call. To publish data for the current identified applications, do the following:

Procedure

1. Go to the **Discovered Apps** tab of the application discovery dashboard and click the **Publish** button (see Application Discovery View on page 772).
   Information for the identified applications is published immediately.

2. Go to the **Published Apps** tab to view the list of published applications.

   You can delete applications from the published list that you do not want in the database (see Application Discovery View on page 772).

Creating Discovery Policies

You can define custom discovery policies for unidentified applications.

About this task

Application discovery can identify a variety of common (system defined) application types such as DNS, NTP, SSH, Web Server, MS Exchange Server, MS SQL Server, Oracle DB, and SAP

HANA. However, it cannot determine the type for some applications. To create a custom (user defined) discovery policy for an unidentified application, do the following:

Procedure

1. Go to the **Policies** tab of the application discovery dashboard and click the **New Policy** button (see Application Discovery View on page 772).

2. In the New Discovery Policy window, do the following in the indicated fields:

   a. **Application**: Enter the application name.

   b. **Policy Name**: Enter a name for the policy.

   c. **Description**: Enter a description for the policy (optional).

   d. **VMs**: Do one of the following:

      » If you want the policy to apply to all VMs, click the **All VMs** radio button.

      » If you want the policy to apply to a limited set of VMs, click the **Specific VMs** radio button and then select the target VMs from the pull-down list.

   e. **TCP Ports**: Enter the TCP port number(s) used by the application. Enter a comma separated list for multiple port numbers.

   f. **UDP Ports**: Enter the UDP port number(s) used by the application. Enter a comma separated list for multiple port numbers.

   g. Click the **Save and Apply** button.

   Information about existing discovered applications is refreshed, reflecting any changes from the new policy.

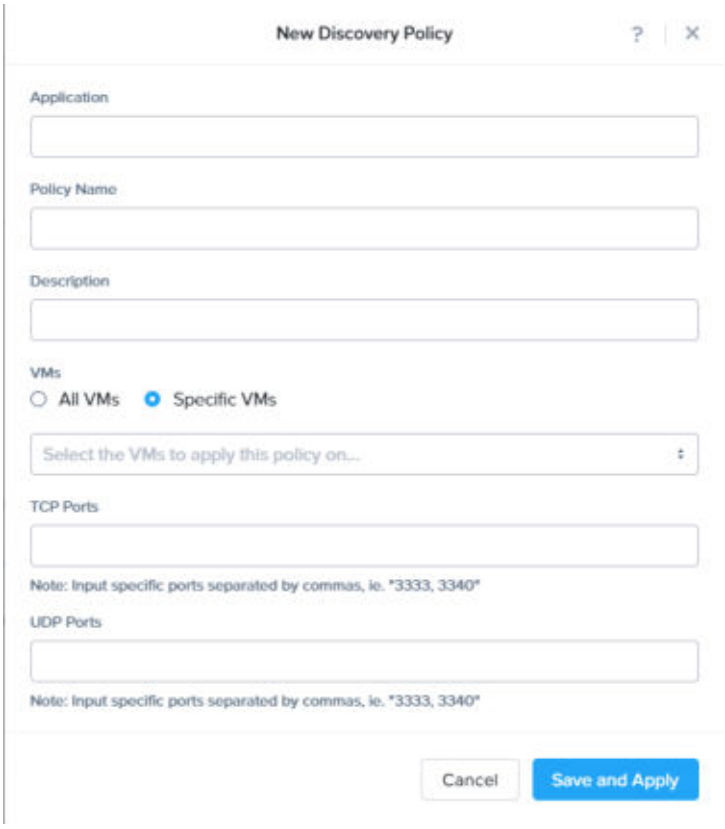

Figure 411: New Discovery Policy Window

Modifying (Viewing) Discovery Policies

You can view any discovery policy, enable or disable any policy, or modify a user-defined policy at any time.

About this task

To view or modify a policy, do the following:

Procedure

1. Go to the **Policies** tab of the application discovery dashboard, find the target policy in the list, and click the name (see Application Discovery View on page 772).

   If it is a system policy, the View Discovery Policy window appears. If it is a user-defined policy, the **Edit Discovery Policy** window appears. The displayed information is the same, but the view window is read-only.

2. To modify one or more parameters in a user-defined policy (see Creating Discovery Policies on page 792), change the entry in the appropriate field(s) and then click the **Save and Apply** button.

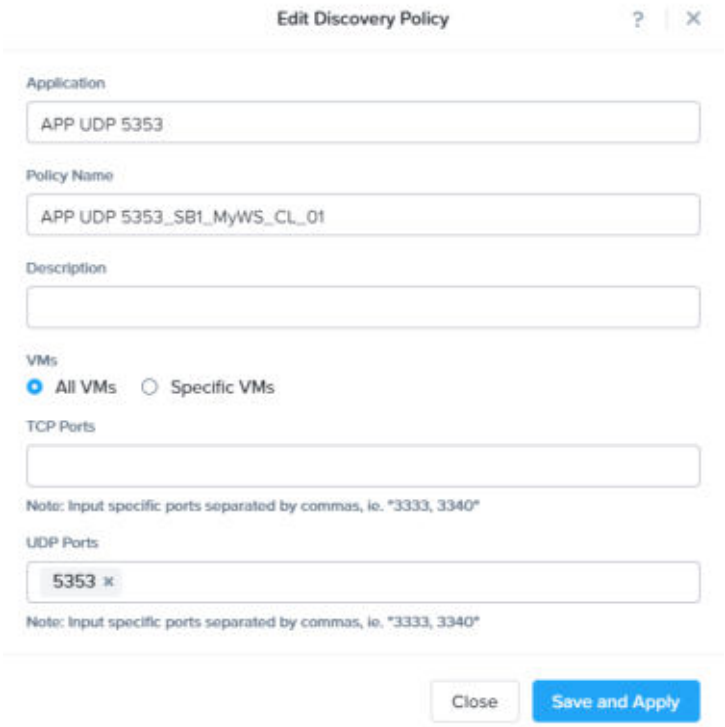

Figure 412: Edit Discovery Policy Window

3. To enable, disable, or delete a policy, select the target policy in the list and then select the desired option (**Enable**, **Disable**, or **Delete**) from the **Actions** pull-down menu. You can only delete user-defined policies (not system policies).

Cost Management (Xi Beam)

About this task

Xi Beam is a cost governance and security compliance SaaS product offering by Nutanix. Xi Beam helps cloud-focused organizations to gain visibility into cloud spend and security compliance status across multiple cloud environments.

To enable Xi Beam, do the following:

Procedure

1. Enable Pulse if it is not already enabled (see Configuring Pulse on page 933).
2. Set up a Xi Beam (SaaS) account.
3. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Operations** > **Cost Management** from the entities menu (see Entities Menu on page 13).
4. In the Welcome to Cost Management page, click the **Manage Cost** button.

   The Xi Beam site appears. See the *Xi Beam User Guide* for information about setting up a Xi Beam account and using Xi Beam to configure and monitor cost governance and security compliance in your environment.

Application Monitoring

Application monitoring provides visibility into integrated applications by collecting application metrics using Nutanix and third-party collectors, providing a single pane of glass for both application and infrastructure data, correlating application instances with virtual infrastructure, and providing deep insights into applications performance metrics.

Application monitoring provides visibility into the following application:

- Microsoft SQL Server
- VMware vCenter Server

Application Monitoring Summary View

The monitoring integrations dashboard allows you to view information about select applications, such as SQL Server or vCenter instances, running in the cluster.

To access the monitoring integrations dashboard, select **Operations** > **Integrations** from the entities menu (see Entities Menu on page 13). The monitoring integrations dashboard allows you to view summary information about application instances and access detailed information about each instance.

Note:

- This section describes the information and options that appear in the application discovery dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.
- If monitoring integrations is not enabled, a "Welcome to Monitoring Integrations" page appears instead of the dashboard. Click the **Get Started** button. This enables monitoring integrations, after which the dashboard appears. However, the dashboard is blank until you configure monitoring (see Configuring Application Monitoring on page 809).

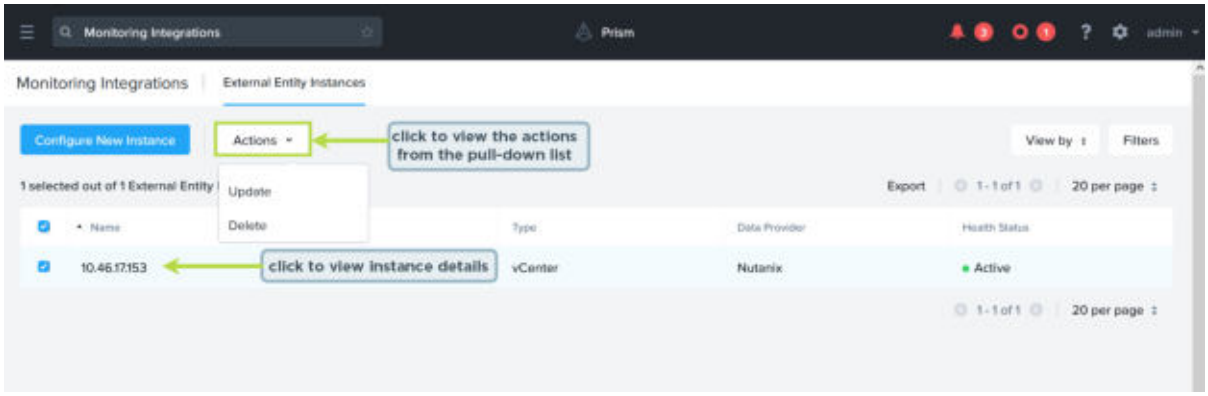

Figure 413: Monitoring Integrations Dashboard

The following table describes the fields that appear in the list of application instances. A dash (-) is displayed in a field when a value is not available or applicable.

Table 727: Application Instance Fields

Parameter	Description	Values
Name	Displays the name of the application instance. Click the name to display the details for that instance (see Application Instance Details View (SQL Server) on page 798 or Application Instance Details View (vCenter) on page 807).	(instance name)
Type	Displays what type of application the instance is.	Microsoft SQL Server, vCenter
Data Provider	Application metrics are supplied through an agentless collector, and this field displays who provided that collector.	Nutanix, (third-party name)
Health Status	Displays the current health status of the instance.	Active, Failed

To download the table contents in CSV format, click the **Export** link.

You can filter the list by opening the **Filter** pane to select one or more of the filter options.

- **Name**: Check the box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field to filter the list for instances with that string in the name.
- **Type**: Check the **Microsoft SQL Server** or **vCenter** box to filter the list for that application type.
- **Health Status**: Check the **Active** or **Failed** box to filter the list for active or failed application instances.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

The dashboard includes the following action buttons:

- Click the **Configure Instance** button to add an application instance for monitoring (see Configuring Application Monitoring on page 809).
- An **Actions** menu appears when an instance is selected:
  - Select **Update** from the pull-down list to update the instance configuration (Modifying Application Monitoring on page 812).
  - Select **Delete** from the pull-down list to stop monitoring the instance. The instance entry disappears from the list.

Application Instance Details View (SQL Server)

The application instance details page provides a range of data and performance metrics for an SQL Server instance.

To access the details page for an SQL Server instance, go to the application integrations dashboard and click the application instance name (see Application Monitoring Summary View on page 796). The instance name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Databases**, **Queries**, and **Metrics**. Click a tab to display that information on the right. (Click the **Back to External Entity Instances** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays the following:

- A **Properties** widget that displays summary information about the cluster (see following table).
- A **Metrics** widget that displays capacity and performance metrics (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, disk or other anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
- Click the **Back to External Entity Instances** link to return to the monitoring applications dashboard.

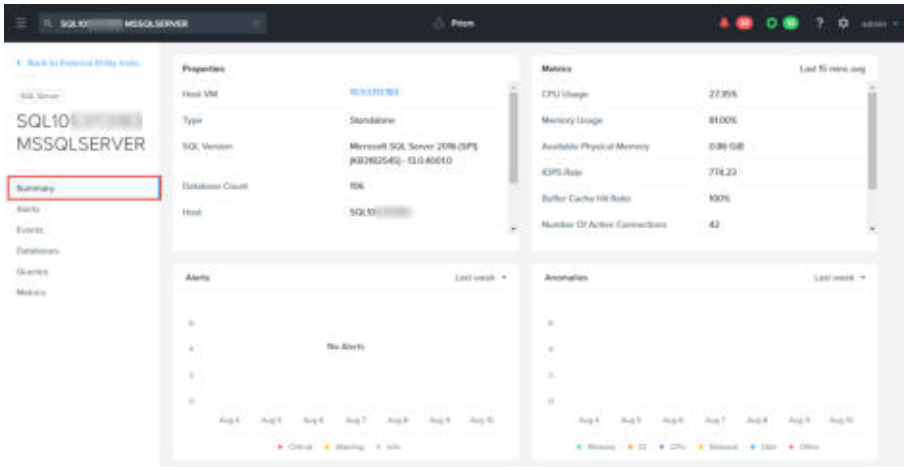

Figure 414: Application Instance Summary Tab

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

Table 728: Instance Properties Fields

Parameter	Description	Values
Properties		
Host VM	Displays the host VM IP address. Click the IP address to displays the details page for that VM (see VM Details View on page 96).	(IP address)
Type	Displays the type of SQL server.	Standalone
SQL Version	Displays the SQL server version number for the instance.	(version number)
Database Count	Displays the number of databases managed by the SQL server instance.	(integer)
Host	Displays the SQL server host name.	(host name)
Instance	Displays the instance type.	
VM Name	Displays the host VM name.	(name)
VM Alert Count	Displays the number of alerts assciated with the instance.	(integer)
Metrics		
CPU Usage	Displays the percentage of CPU capacity currently being used by the instance.	0 - 100%
Memory Usage	Displays the percentage of memory capacity currently being used by the instance.	0 - 100%
Available Physical Memory	Displays the amount of available memory.	xx [GiB]

Parameter	Description	Values
IOPS Rate	Displays the average number of combined read and write operations per second for the past 15 minutes.	(number)
Buffer Cache Hit Ratio	Displays the buffer cache hit rate for the past 15 minutes.	0-100%
Number of Active Connections	Displays the number of active connections.	(integer)
Read Ops Rate	Displays the average number of read operations per second for the past 15 minutes.	(number)
Write Ops Rate	Displays the average number of write operations per second for the past 15 minutes.	(number)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard (List tab), except it is filtered to display just SQL Server-related alerts (see Alerts Summary View on page 261).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just SQL Server-related events (see Events Summary View on page 272).

Databases Tab

The **Databases** tab displays a list of databases supported by this SQL Server instance. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

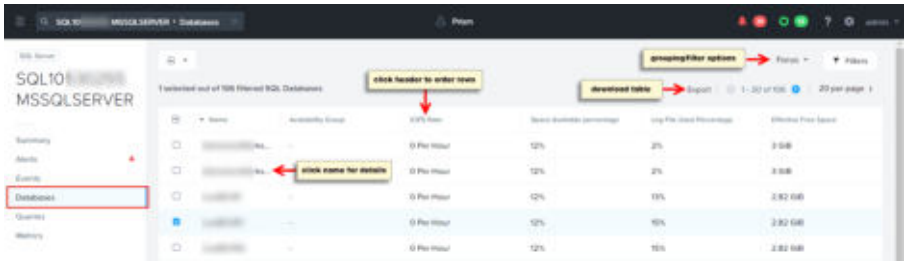

Figure 415: Application Instance Databases Tab

- To order the rows, click a column header, which toggles the list alphabetically or numerically (as appropriate) based on that parameter.
- To download the table in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 81).
- To filter the list by database name, open the **Filter** pane, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field.

- To select (or unselect) databases in the list:
  - Click the check box on the far left for each database to select. (Click the check box again to unselect that database.)
  - Click the check box in the header row to select (unselect) all the databases in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of databases.

Table 729: Databases Fields (SQL Server)

Parameter	Description	Values
Name	Displays the database name. Click on the name to display the details for that database (described below).	(database name)
Availability Group	Displays the name of the availability group to which the database belongs.	(availability group name)
IOPS Rate	Displays the average number of combined disk read and write read operations per hour.	<number> Per Hour
Space Available Percentage	Displays the percentage of space available on disk used by the database.	0-100%
Log File Used Percentage	Displays the percentage of the log file used for the database.	0-100%
Effective Free Space	Displays the amount of space on disk that is usable by the database. This includes the free space of the entire disk and any space being used by the database.	xxx [MiB|GiB]

Clicking a database name displays the details page for that database. The following figure is a sample view. The page displays four graphs. Each graph is a rolling time interval performance or usage monitor for a metric. The metrics are IOPS rate, space available, log file utilization, and effective free space.

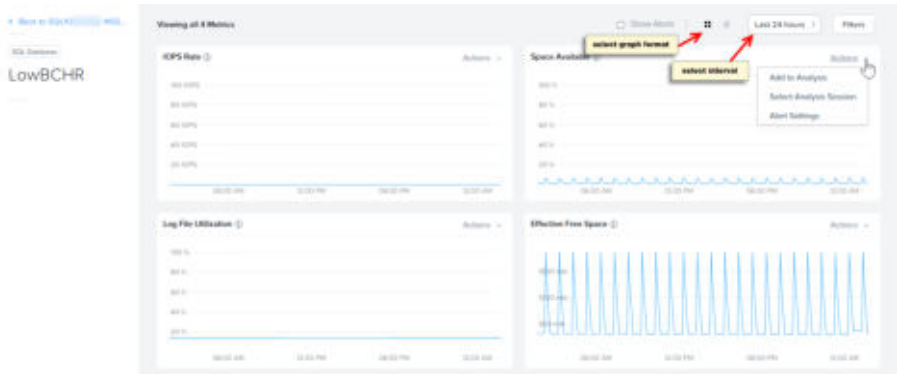

Figure 416: Database Instance Page

- Each graph includes a pull-down actions menu with the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard on page 739).
  - Select **Select Analysis Session**, which opens the **Select A Session** window, to add the metric to an existing or new analysis session. Choose an option from the pull-down list and then click the Add to Session button.

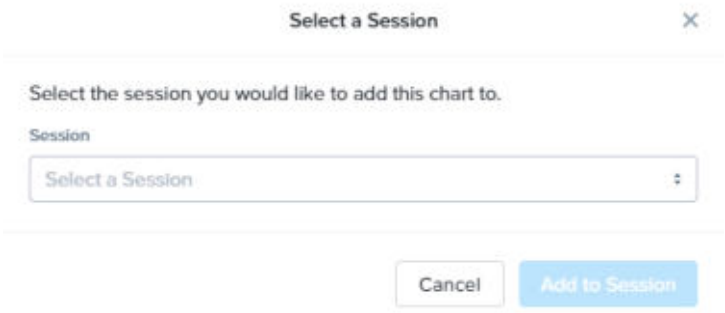

Figure 417: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 282).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired graphs.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to <instance_name>** link to return to the application instance details page.

Queries Tab

The **Queries** tab displays a list of SQL query statements. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

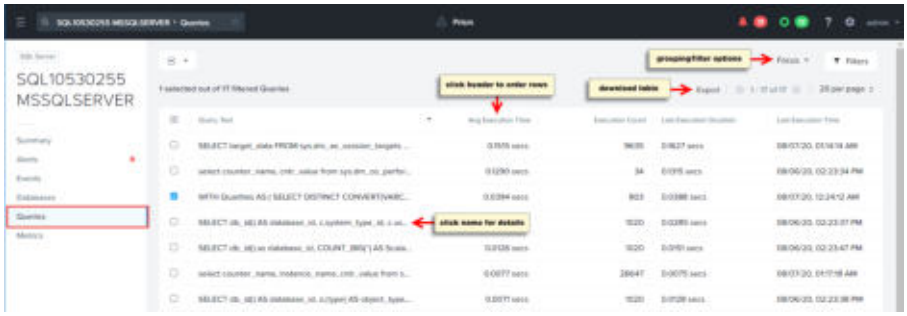

Figure 418: Application Instance Queries Tab

- To order the rows, click a column header, which toggles the list alphabetically, numerically, or chronologically (as appropriate) based on that parameter.

- To download the table in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 81).
- To filter the list for a text string in the query, open the **Filter** pane, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field.
- To select (or unselect) queries in the list:
  - Click the check box on the far left for each query to select. (Click the check box again to unselect that query.)
  - Click the check box in the header row to select (unselect) all the queries in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of queries.

Table 730: Query Fields

Parameter	Description	Values
Query Text	Displays the query statement text. Click on the text to display details about that query statement (described below).	(query text)
Avg Execution Time	Displays the average amount of time it took to run the query.	xxx secs
Execution Count	Displays the number of times the query statement has been executed.	(integer)
Last Execution Duration	Displays the amount of time it took to run the last query.	xxx secs
Last Execution Time	Displays the date and time when the query was last executed.	(date and time)

Clicking a query text displays details about that query, including a Properties widget that includes the Execution Count, Last Execution Duration, and Query text fields.

- To see the full query statement, click the **View** link.

Figure 419: Query Statement Window

- To copy the query statement, click the **Copy** link.
- Click the **Back to <instance_name>** link to return to the application instance details page.

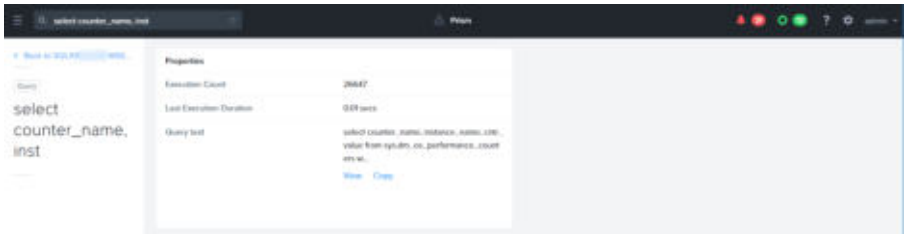

Figure 420: Query Details Page

Metrics Tab

The **Metrics** tab displays a set of graphs for various metrics. Each graph is a rolling time interval performance or usage monitor for that metric.

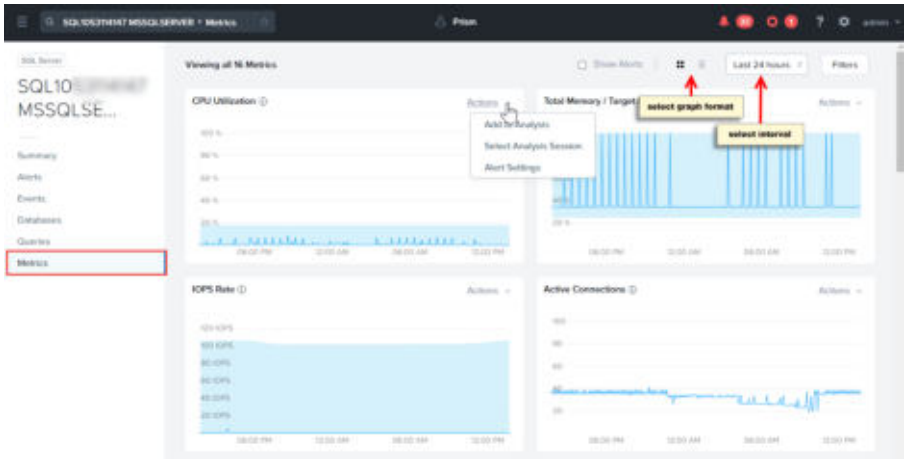

Figure 421: Application Instance Metrics Tab

- Each graph includes a pull-down actions menu with the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard on page 739).
  - Select **Select Analysis Session**, which opens the **Select A Session** window, to add the metric to an existing or new analysis session. Choose an option from the pull-down list and then click the Add to Session button.

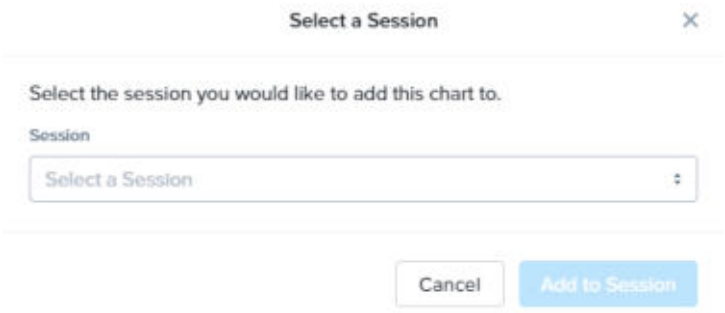

Figure 422: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 282).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired graphs.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to External Entity Instances** link to return to the monitoring applications dashboard.

The following table describes the metric graphs.

Table 731: Metric Performance Graphs

Metric	Description
CPU Utilization	Displays the CPU bandwidth usage by all requests in all workload groups belonging to all related resource pools.
Total Memory / Target Memory	Displays the percentage of total server memory divided by target server memory. When the total server memory and target server memory values are similar (close to 100%), there is no memory pressure on the server. Total server memory significantly lower than target server memory during normal SQL Server operations indicates there is memory pressure on the server, so SQL Server cannot get as much memory as needed.

Metric	Description
IOPS Rate	Displays the total amount of disk write and read operations. This is a calculated metric displaying the IOPS that occurred per second since the previous collection.
Active Connections	Displays the number of active connections to this instance.
Buffer Cache Hit Ratio	Displays the percentage of pages found in the buffer cache without having to read from disk. The ratio is the total number of cache hits divided by the total number of cache lookups over the last few thousand page accesses. Because reading from the cache is much less expensive than reading from disk, you want this ratio to be high. Generally, you can increase the buffer cache hit ratio by increasing the amount of memory available to SQL Server or by using the buffer pool extension feature.
Batch Request Rate	Displays the rate of transact-SQL command batches received. This statistic is affected by all constraints (such as I/O, number of users, cache size, complexity of requests, and so on). High batch requests mean good throughput. This is a calculated metric displaying the batch requests that occurred per second since the previous collection.
SQL Compilations Rate	Displays the rate of SQL compilations of the instance. This is a calculated metric displaying the SQL compilations that occurred per second since the previous collection.
SQL Recompilations Rate	Displays the rate of statement recompiles. Counts the number of times statement recompiles are triggered. Generally, you want the recompiles to be low. This is a calculated metric displaying the SQL Recompilations that occurred per second since the previous collection.
Buffer Page Life Expectancy	Displays the number of seconds a page will stay in the buffer pool without references.
Checkpoint Pages Rate	Displays the number of pages flushed to disk by a checkpoint or other operation that requires all dirty pages to be flushed. This is a calculated metric displaying the checkpoint pages that occurred per second since the previous collection.
Available Physical Memory	Displays the amount of physical memory available.
Buffer Pool Size	Displays the total size of the buffer pool.
Lock Wait Time	Displays the total wait time for locks in the last second.
User Connections	Displays the number of connections that are currently being used for user sessions.
Read Operations Rate	Displays the total amount of disk read operations. This is a calculated metric displaying the write operations that occurred per second since the previous collection.
Write Operations Rate	Displays the total amount of disk write operations. This is a calculated metric displaying the write operations that occurred per second since the previous collection.

Application Instance Details View (vCenter)

The application instance details page provides a range of data and performance metrics for a vCenter instance.

To access the details page for a vCenter instance, go to the application integrations dashboard and click the application instance name (see Application Monitoring Summary View on page 796). The instance name and tabs for **Summary** and **Alerts** appear on the left. Click a tab to display that information on the right. (Click the **Back to External Entity Instances** link to return to the summary view.)

Summary Tab

The **Summary** tab, which appears by default when you first open the page, displays a **Properties** widget that shows summary information about the instance. Click the **Back to External Entity Instances** link to return to the monitoring applications dashboard.

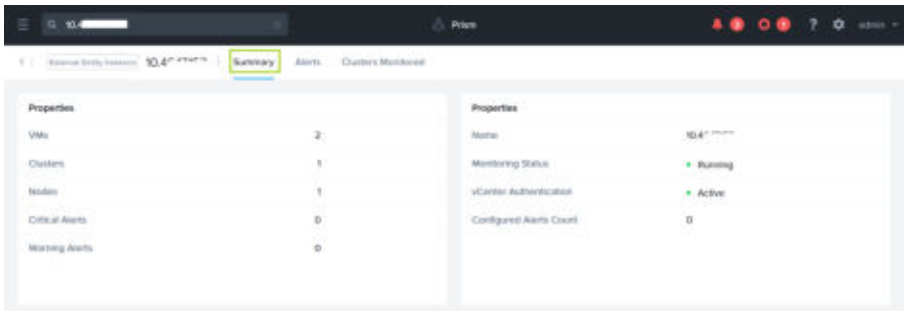

Figure 423: Application Instance Summary Tab

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

Table 732: Instance Properties Fields

Parameter	Description	Values
VMs	Displays the number of VMs managed by the vCenter instance.	(integer)
Clusters	Displays the number of clusters managed by the vCenter instance.	(integer)
Nodes	Displays the number of clusters managed by the vCenter instance.	(integer)
Critical Alerts	Displays the current number of critical alerts associated with the vCenter instance.	(integer)
Warning Alerts	Displays the current number of warning alerts associated with the vCenter instance.	(integer)
Name	Displays the name of the vCenter instance.	(name)
Configured Alerts Count	Displays the number of alerts configured to monitor the vCenter instance.	(integer)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard (List tab), except it is filtered to display just vCenter-related alerts (see Alerts Summary View on page 261).

Infrastructure Tabs (VMs, Clusters, and Hosts)

In addition to the application details view, tabs for non-Nutanix entities appear in the VMs, clusters, and hosts dashboards after enabling vCenter application monitoring.

- To view information about VMs managed by vCenter, select **Compute & Storage** > **VMs** from the entities menu and then click the **non-Nutanix** tab (see VMs Summary View on page 86). The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed VMs. The **Metrics** tab displays a subset of the full list (10 of the 12 metrics). The **List** tab displays fields for name, node name, hypervisor, memory, IP addresses, power state, and cluster name.

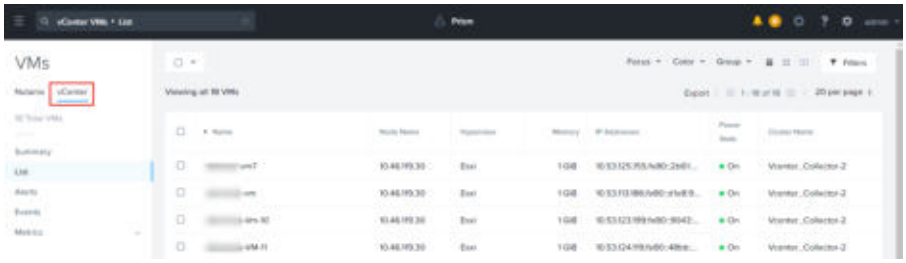

Figure 424: VMs dashboard (vCenter tab)

- To view information about clusters managed by vCenter, select **Hardware** > **Clusters** from the entities menu and then click the **non-Nutanix** tab (see Clusters Summary View on page 197). The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed clusters. The **Metrics** tab displays CPU usage, memory usage, and VM operations. The **List** tab displays fields for name, host count, VM count, storage usage, and storage capacity.

Figure 425: Clusters dashboard (vCenter tab)

- To view information about hosts managed by vCenter, select **Hardware** > **Hosts** from the entities menu and then click the **non-Nutanix** tab (see Hosts Summary View on page 217). The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed

hosts. The **Metrics** tab displays a subset of the full list (5 of the 7 metrics). The **List** tab displays fields for name, host IP, VM count, memory capacity, and cluster name.

Figure 426: Hosts dashboard (vCenter tab)

Configuring Application Monitoring

Monitoring an application instance requires configuring an instance profile.

Before you begin

Complete the Prerequisites on page 811.

About this task

To configure an instance profile, do the following:

Procedure

1. Go to the monitoring integrations dashboard and click the **Configure Instance** button (see Application Monitoring Summary View on page 796).

2. In the Configure external entity for monitoring page, do the following in the indicated fields:

   a. **External Entity Type**: Select the target type from the pull-down menu.

      The options are **Microsoft SQL Server** or **vCenter**. If this is the first time configuring an instance, an **Enable** button appears. Click the **Enable** button to provision instance monitoring.

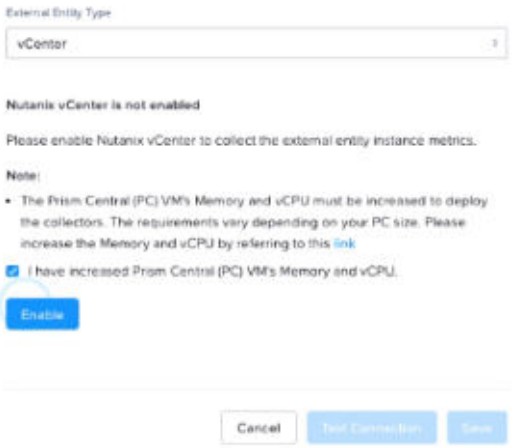

Figure 427: Configuration Page (enable)

When the provisioning is complete (which can take a few minutes), the following fields appear.

Figure 428: Configuration Page (SQL server example)

   b. [**Microsoft SQL Server** | **vCenter**] **Host**: Enter the IP address of the SQL Server or vCenter host VM. You can monitor SQL Server instances in both Nutanix and non-Nutanix (vCenter managed) clusters.

   c. [**Microsoft SQL Server** | **vCenter**] **Port**: Enter the target port number.

      The port number defaults to the standard port number for SQL Server (1433) or vCenter (443).

   d. **Username**: Enter the SQL Server or vCenter administrator user name.

   e. **Password**: Enter the SQL Server or vCenter administrator password.

   f. [vCenter only] **SSL Verify**: Select the SSL mode to use (**Verify** or **Don't Verify**) when connecting to the target.

      This setting determines whether to verify (or not verify) the target's SSL certificate as part of the connection process.

g. Click the **Test Connection** button to verify (test) the connection.

   A "Test connection success" message appears after a successful test.

h. After a successful test, do one of the following:

   » SQL Server: Click the **Save** button. The instance configuration and data collection begins. When it completes (which can take a few minutes), click the **Close** button to close the page. The monitoring integrations dashboard reappears with the new instance in the list.

   » vCenter: Click the **Next** button.

i. [vCenter only] Select the clusters to monitor.

   A list of discovered clusters for the vCenter instance appears. Select (check the boxes for) the clusters you want to monitor from the list and then click the **Save** button.

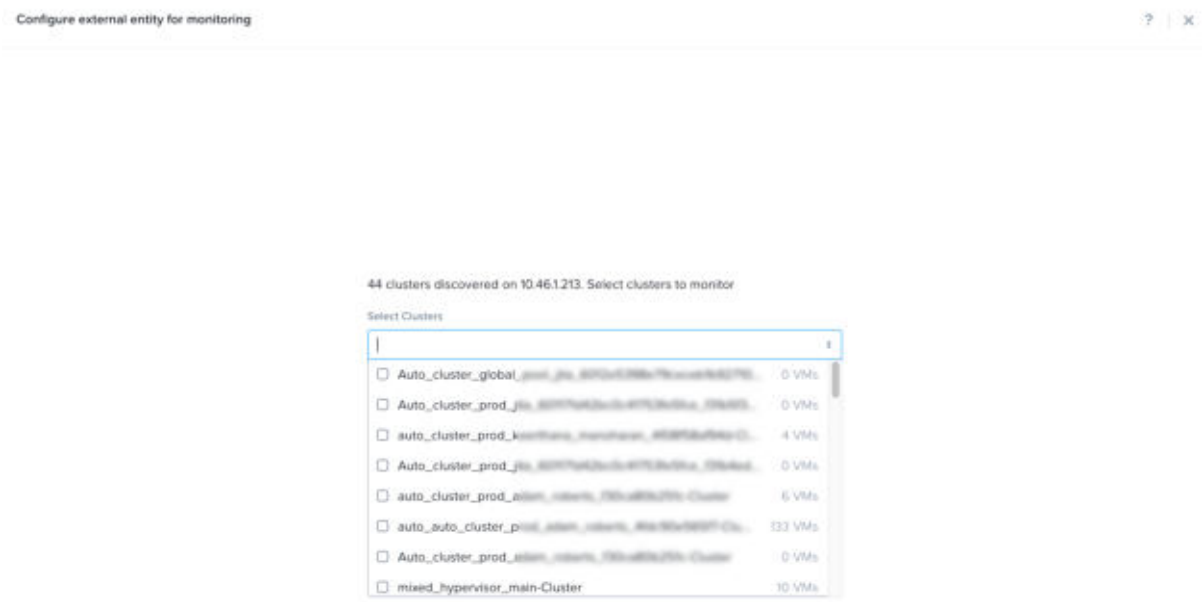

Figure 429: Select Clusters

Note: Nutanix clusters monitored by vCenter that are also registered with this Prism Central instance do not appear in the discovered clusters list. Conversely, non-registered Nutanix clusters selected here but later registered with Prism Central appear in both the Nutanix and non-Nutanix tabs of the clusters dashboard (see Application Instance Details View (vCenter) on page 807).

The instance configuration and data collection begins. When it completes (which can take a few minutes), a "vCenter Configuration Successful" screen appears. Either click the **Close** button to close the page or click the **App Discovery** link (if present) to enable application discovery for the selected clusters (see Enabling Application Discovery on page 781).

Prerequisites

Address the following before enabling application monitoring:

- Sufficient Prism Central VM resources. Application monitoring is a feature that requires additional Prism Central VM resources (more memory and vCPUs). The requirements are listed in KB 8932 (see instructions for Prism Ultimate features). If your Prism Central VM resources are below the recommended amount, manually increase the memory and vCPUs to the recommended amount.
- SQL Server-specific requirements:
  - Each SQL Server instance port must be accessible through your network, but no other special network configuration is required.
  - Create a read-only login to connect to your server:

    ```
    CREATE LOGIN <username> WITH PASSWORD = '<PASSWORD>';
    CREATE USER <username> FOR LOGIN <username>;
    GRANT SELECT on sys.dm_os_performance_counters to <username>;
    GRANT VIEW SERVER STATE to <username>;
    ```

    After creating the login make sure that your SQL Server instance supports SQL Server authentication by enabling "SQL Server and Windows Authentication mode" in the server properties (**Server Properties** > **Security** > **SQL Server and Windows Authentication mode**).
  - Active Directory users and Windows authentication are not supported for SQL Server monitoring.

Modifying Application Monitoring

You can view, update, or delete the application configuration for an instance at any time.

About this task

To modify application monitoring for an instance, do the following:

Procedure

1. Go to the monitoring integrations dashboard, select the target instance in the list, and then select **Update** from the **Actions** menu pull-down list (see Application Monitoring Summary View on page 796).

   The Configure external entity for monitoring page for that instance appears (see Configuring Application Monitoring on page 809).
2. To modify one or more parameters, change the entry in the appropriate field(s), click the **Test Connection** button to verify the changes, and then click the **Save** button.

   Note: [vCenter] Modifying the cluster selection occurs after testing the connection.
3. To remove application monitoring for an instance, select the target instance in the list, and then select **Delete** from the **Actions** menu pull-down list.

   The entry for that instance disappears from the list.

Operations Policy Management

Prism Central implements a variety of system services to analyze cluster integrity and performance. You have the option to apply policies to customize some services in your environment. You can apply the following policies:

- **Anomaly Detection Exclusion**. Prism Central includes machine-learning capabilities that analyze resource usage and identify potential problems (see Behavioral Learning Tools on

page 838). One is anomaly detection that records when performance or resource usage is outside an expected range. The anomaly detection exclusion policy allows you to specify VMs to ignore when looking for anomalies. For example, you might have some VMs with expected peak load on certain days, and you do not want that usage surge counted as an anomaly. In this case you can either define a category for those VMs and enable anomaly detection exclusion for that category or enable the exclusion for each VM individually.

- **Inefficiency Measurement Exclusion**. Another machine-learning capability is VM "right sizing" that identifies VMs which exhibit inefficient profiles. Again, you have the option to exclude selected VMs from the inefficiency analysis.

Operations Policies View

The operations policies dashboard allows you to view and modify operations policies that can be applied to VMs or other entities.

To access the operations policies dashboard, select **Operations** > **Ops Policies** from the entities menu (see Entities Menu on page 13). The dashboard displays a list of the operations policies.

> Note: This section describes the information and options that appear in the operations policies dashboard. See Entity Exploring on page 81 for instructions on how to view and organize that information in a variety of ways.

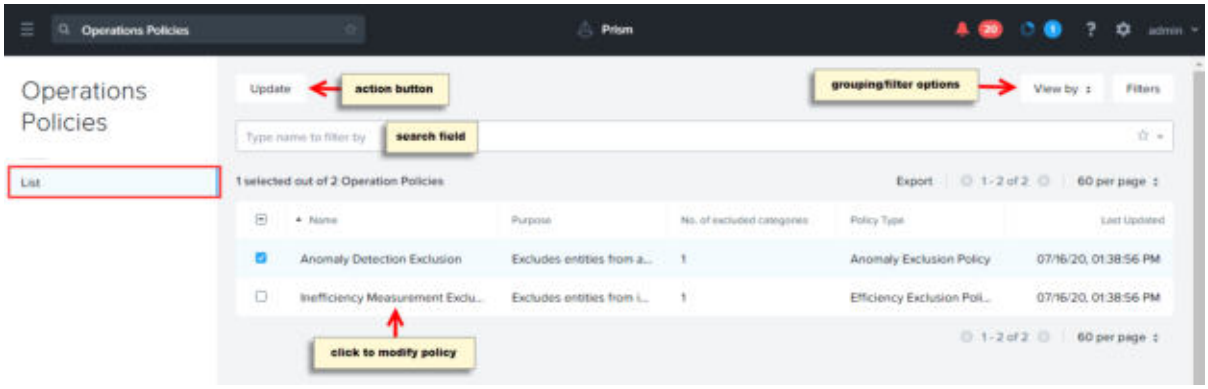

Figure 430: Operations Policies Dashboard

The following table describes the fields that appear in the operations policies list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 733: Operations Policy Fields

Parameter	Description	Values
Name	Displays the application name for the instance. Click the name to display the Application Info window. The window displays the application name, VM name, TCP ports, and UDP ports.	(app name)
Purpose	Describes the purpose of the policy.	(VM name)
No. of Excluded Categories	Displays the number of categories the policy applies to (see Category Management on page 735).	(integer)

Parameter	Description	Values
Policy Type	Describes what type of policy this is.	(type name)
Last Updated	Displays the date and time when the policy was last updated.	(date and time)

You can filter the list either by entering a string in the search field or by clicking the **Filters** button, checking the **Name** box, selecting a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and entering a string in the field.

The **Focus** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.

To modify a policy, select that policy and then select **Update** from the **Actions** pull-down menu (see Modifying an Operations Policy on page 814).

Modifying an Operations Policy

You can customize the list of categories to which an operations policy applies.

About this task

No categories are assigned to the Anomaly Detection Exclusion and Inefficiency Measurement Exclusion polices by default. To add (or delete) one or more categories to an operations policy, do the following:

Procedure

1. Go to the operations policy dashboard, select the desired policy, and then select **Update** from the **Actions** pull-down menu (see Operations Policies View on page 813).
2. In the policy page, click the **Add** link.

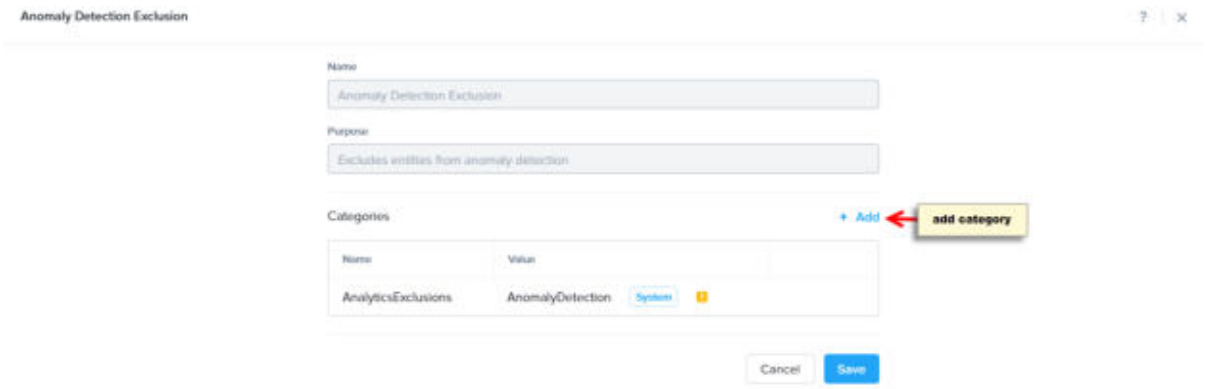

Figure 431: Anomaly Detection Exclusion Page (add category)

3. In the Add window, enter a category name and then click the plus icon. Repeat this step for all categories you want to add.

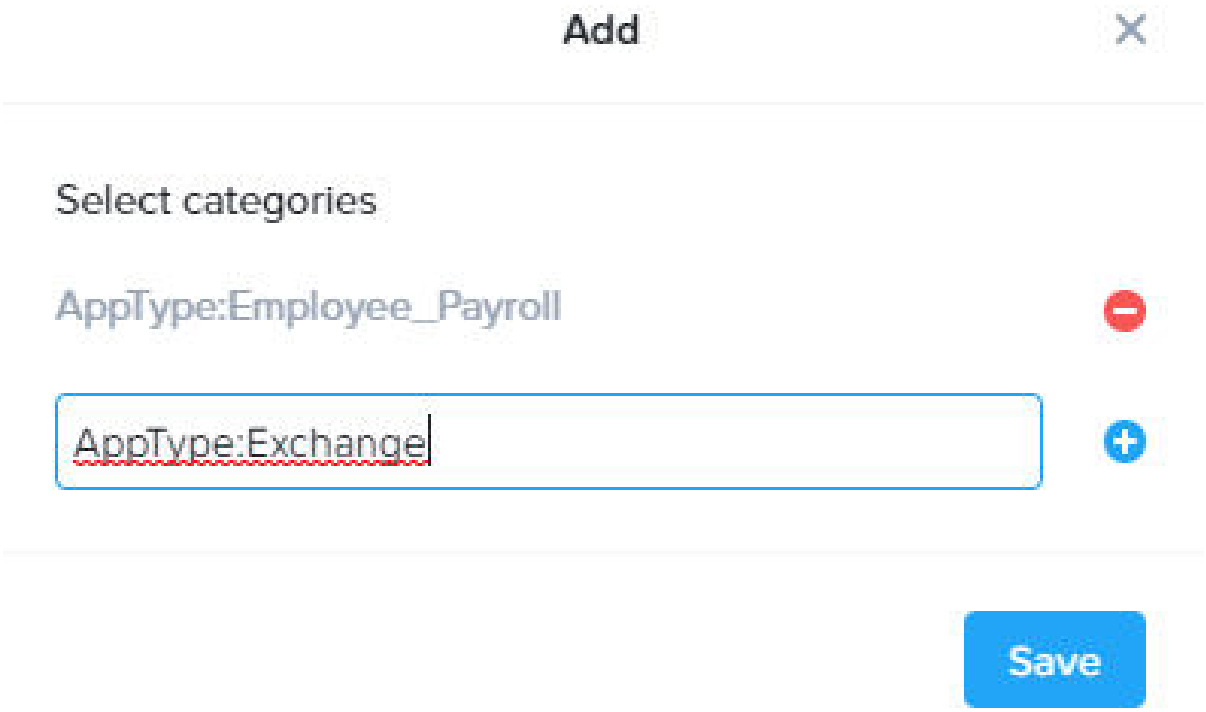

Figure 432: Add Window

4. Click the **Save** button.

   The policy is now enabled for entities in the selected categories.

5. To delete a category from the list, click the garbage pail icon for that category and then click the **Save** button.

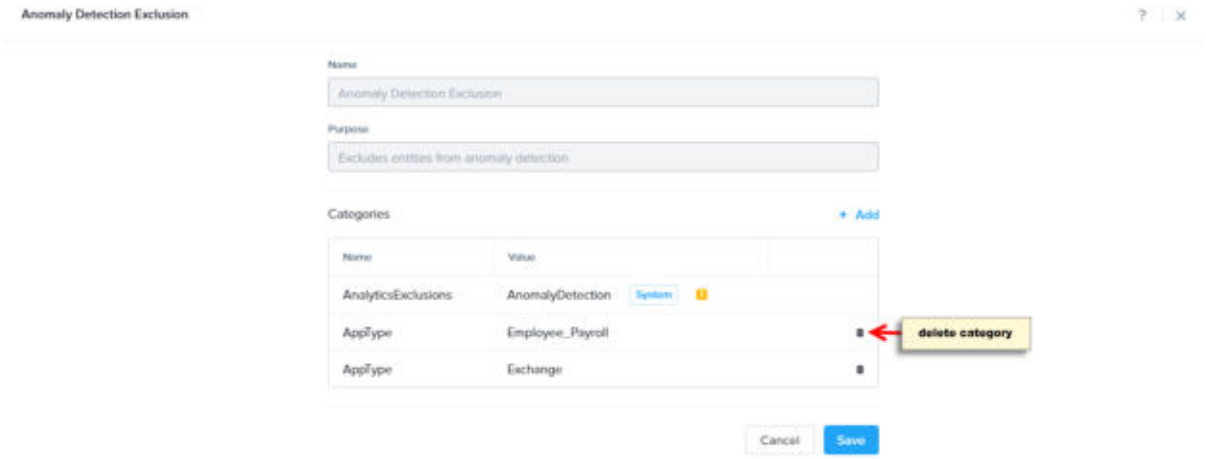

Figure 433: Anomaly Detection Exclusion Page (delete category)

Enabling an Operations Policy

You can enable (or disable) an operations policy for selected VMs.

About this task

To disable anomaly detection or inefficiency measurement for selected VMs or for all the VMs in selected clusters, do the following:

Procedure

1. Do one of the following:

   » Go to the List tab of the VMs dashboard, select the target VMs, and then select **Disable Efficiency Measurement** or **Disable Anomaly Detection** from the **Actions** pull-down menu (see VMs Summary View on page 86).

   » Go to the List tab of the clusters dashboard, select the target clusters, and then select **Disable Efficiency Measurement** or **Disable Anomaly Detection** from the **Actions** pull-down menu (see Clusters Summary View on page 197).

   This immediately disables the specified operation (enables the exclusion policy) for the target VMs.

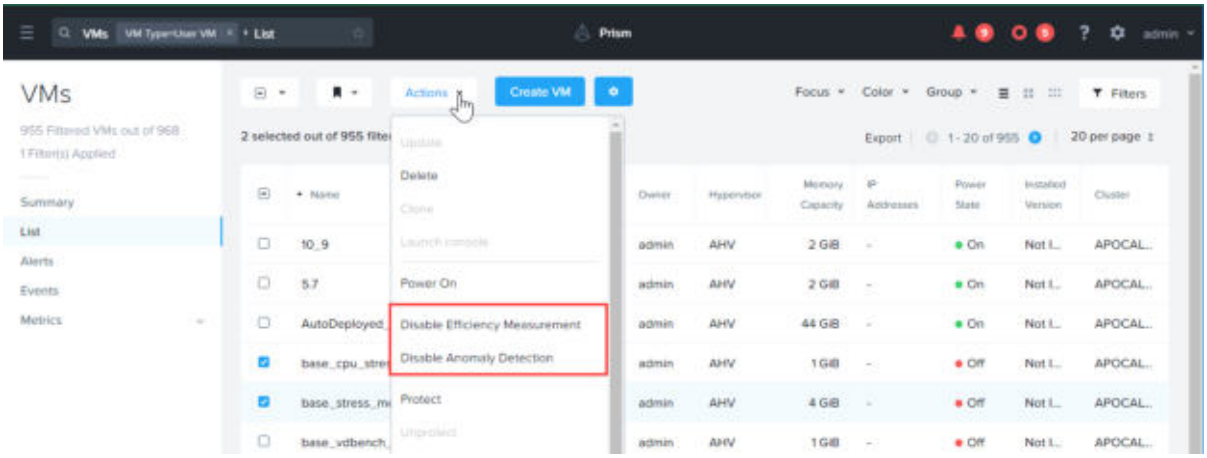

Figure 434: VM Actions Menu

2. To re-enable the operation, select the target VMs or clusters, and then select **Enable Efficiency Measurement** or **Enable Anomaly Detection** from the **Actions** pull-down menu.

   This immediately enables the specified operation (disables the exclusion policy) for the target VMs.

Resource Planning

You can review and analyze current and potential resource needs through the planning dashboard.

To access the planning dashboard, perform the following.

1. Click the hamburger icon.
2. Select **Operations** from the entities menu. (See Entities Menu on page 13).
3. Select **Planning** from the **Operations** menu.

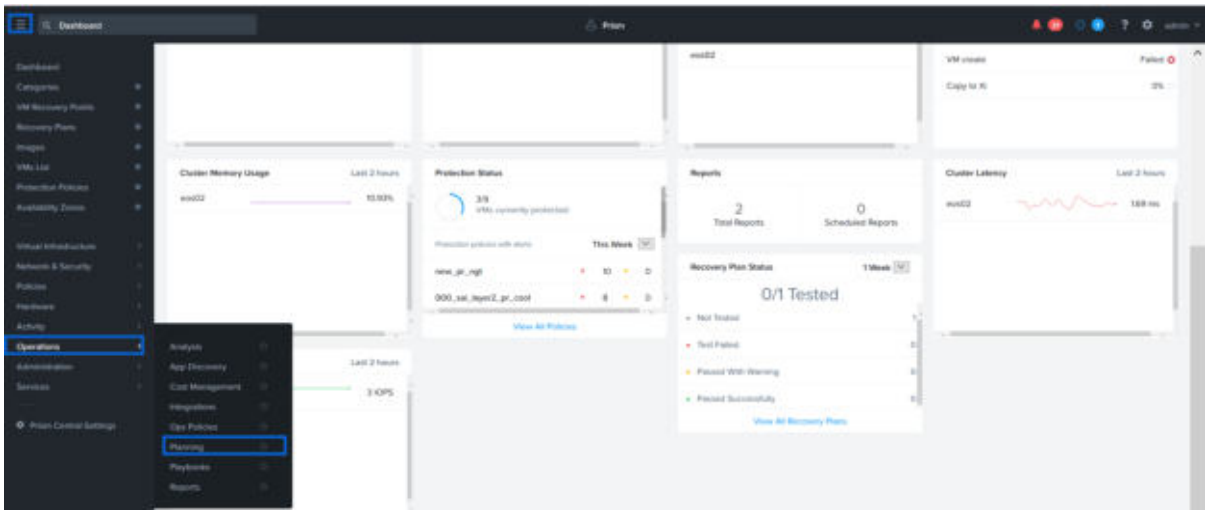

Figure 435: Planning

The planning dashboard includes two tabs:

> Note: You must have a Prism Pro license to use the resource planning tools.

- The **Capacity Runway** tab allows you to view current resource runway information across the registered Nutanix and Non-Nutanix clusters (see Capacity Runway View on page 817).
- 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 819).
- 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 837.

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 838).

Capacity Runway View

Click the **Capacity Runway** tab in the planning dashboard (see Resource Planning on page 816) to view summary resource runway information for the registered Nutanix and Non-Nutanix clusters and to access detailed runway information about each cluster.

> Note:
>
> - Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
> - Capacity planning checks the past 90 days of data for monthly seasonality, it also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
> - For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 796 for more information.

- Click a column header to order the rows by that value (alphabetically or numerically as appropriate).
- Click the cluster name to display the **Capacity** tab of the details page for that cluster (see Cluster Details View on page 205).

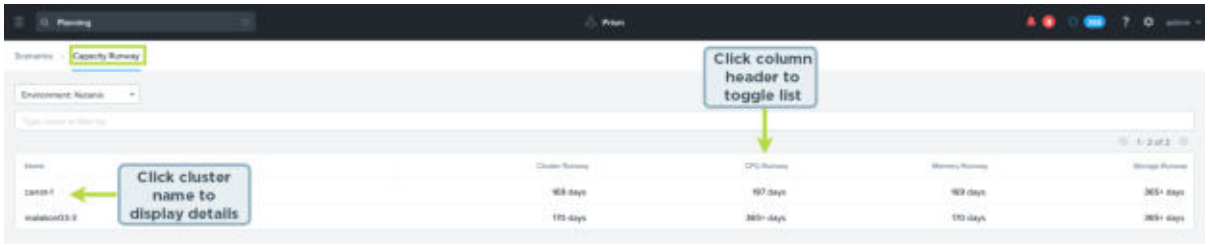

Figure 436: Capacity Runway View (Nutanix)

Figure 437: Capacity Runway View (Non-Nutanix)

The following table describes the fields that appear in the Capacity Runway view.

Table 734: Capacity Runway Tab Fields

Parameter	Description	Values
Name	Displays the cluster name. Clicking the name displays the **Capacity** tab of the details page for that cluster. This tab displays current, historical, and projected usage information. (see Cluster Details View on page 205).	(cluster name)
Cluster Runway	Displays the overall runway for the cluster. The "cluster" runway is the shortest CPU, memory, or storage runway; that is, how long will it be before one of these resources runs out of capacity? Note: The "runway" is the number of days the existing CPU, memory, or storage capacity can handle the workload based on the current workload consumption model.	xxx [days]
CPU Runway	Displays the cluster runway for CPU usage.	xxx [days]
Memory Runway	Displays the cluster runway for memory usage.	xxx [days]
Storage Runway	Displays the cluster runway for storage usage.	xxx [days]

Scenarios View

Click the **Scenarios** tab in the planning dashboard (see Resource Planning on page 816) 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 820).
- 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.

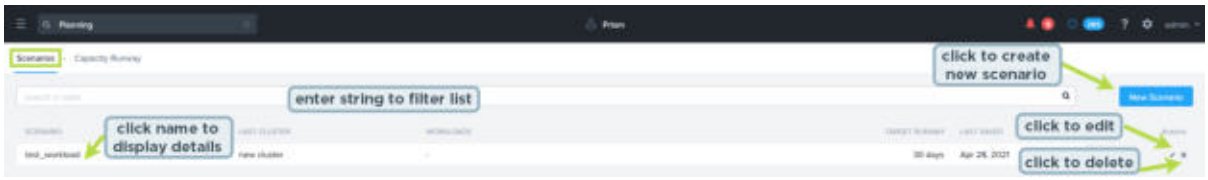

Figure 438: Scenarios View

The following table describes the fields that appear in the Scenarios view.

Table 735: 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 820).	(scenario name)
Last Cluster	Displays the name of the cluster to which the scenario applies. (The column title is "Last Cluster" because you can change the target cluster when modifying a scenario.)	(cluster name)
Workloads	Displays the workloads included in the scenario. This field lists workloads added to the scenario. In the figure above, the "tom" scenario lists "sql" in this field indicating an SQL workload was added to the scenario. A blank field indicates either there is no workload (for a scenario based on a new cluster) or only the existing (current) workload is applied (for a scenario based on an existing cluster).	(blank),

Parameter	Description	Values
Target Runway	Displays the target runway duration for this scenario, that is how many days am I planning for. This is specified when the scenario is created or modified ("Target" parameter).	xxx [days]
Last Saved	Displays the time and date for the last (most recent) time the scenario was saved.	(time and date)
Actions	• To edit the scenario, click the pencil icon (see Modifying a Scenario on page 836). • 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 819) and click the **New Scenario** button.
   The Scenario configuration page appears.

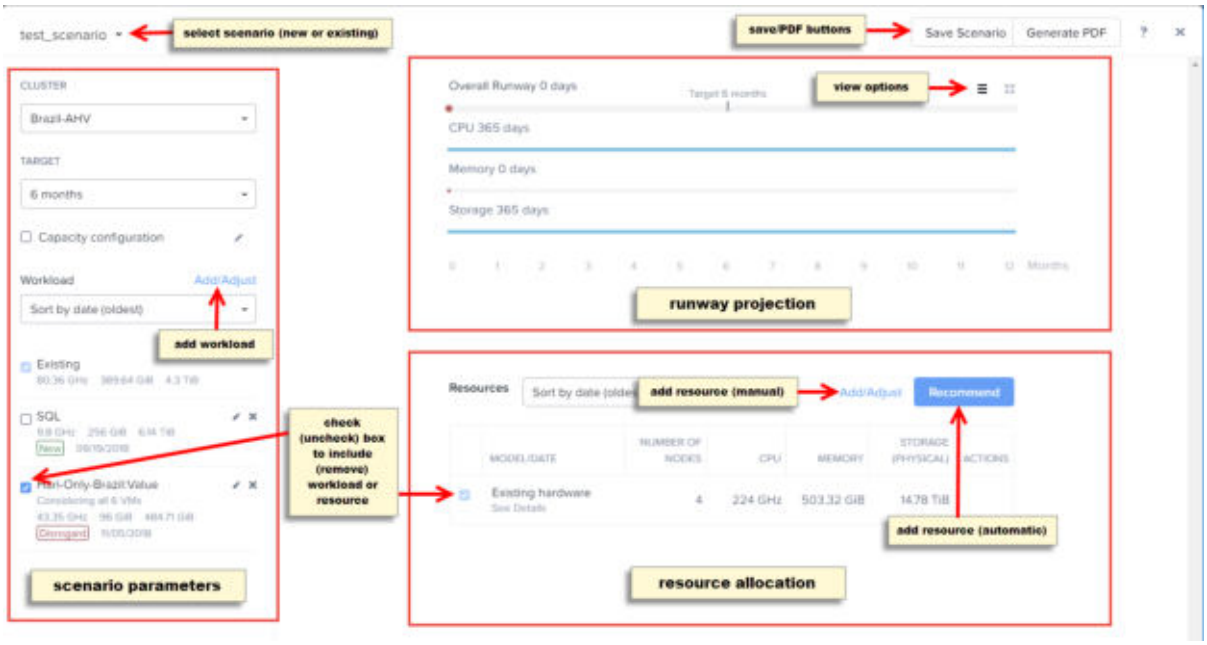

Figure 439: Scenario Configuration Page (example)

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 819), 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 440: 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.

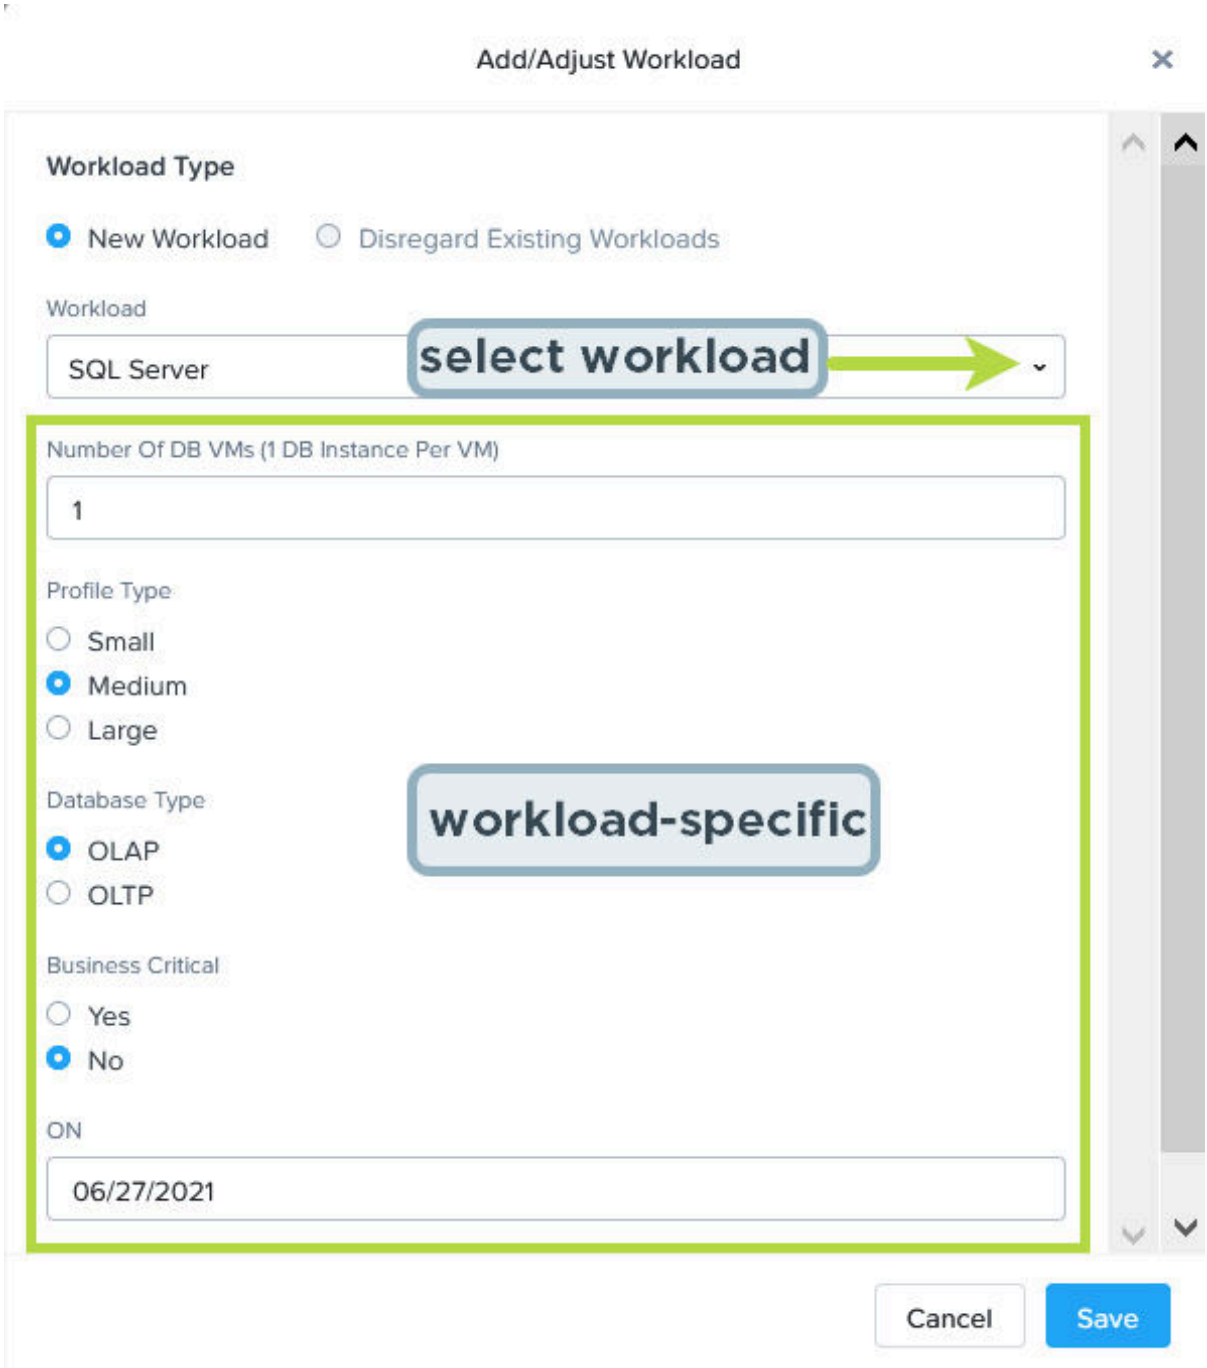

Figure 441: New Workload Fields (SQL Server example)

Table 736: 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 833 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 735).

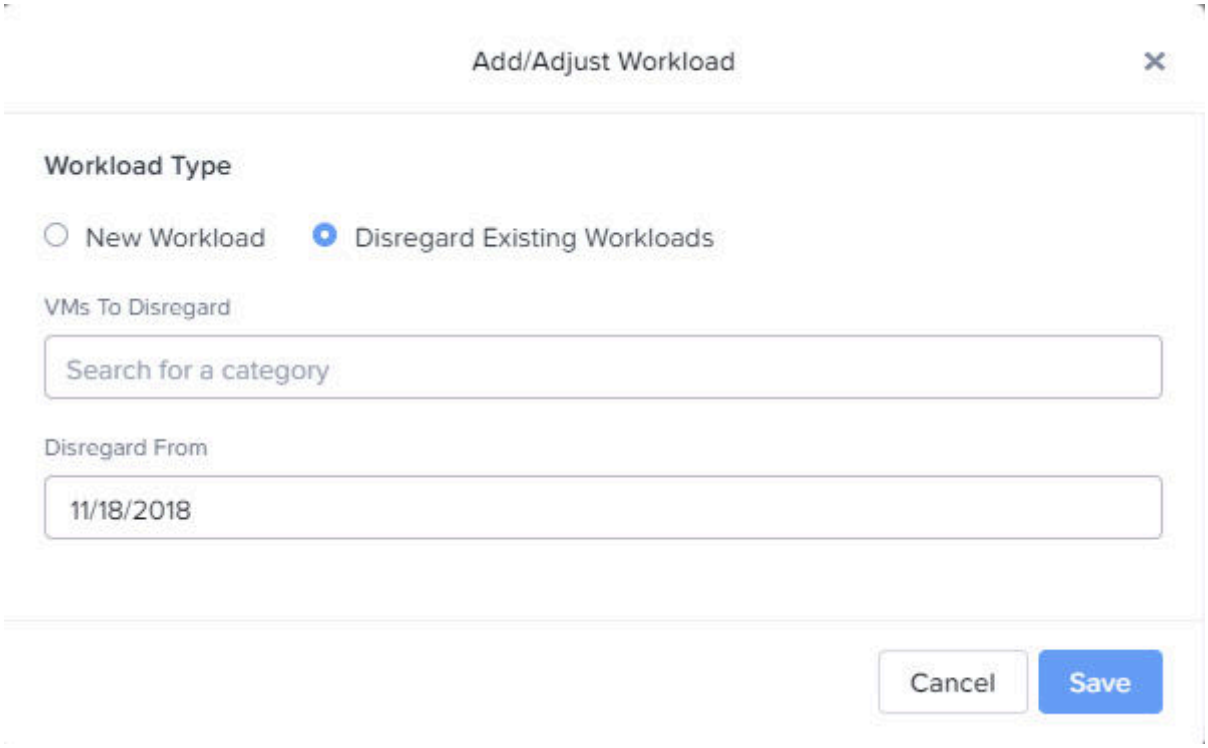

Figure 442: Disregard Existing Workloads Fields

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.

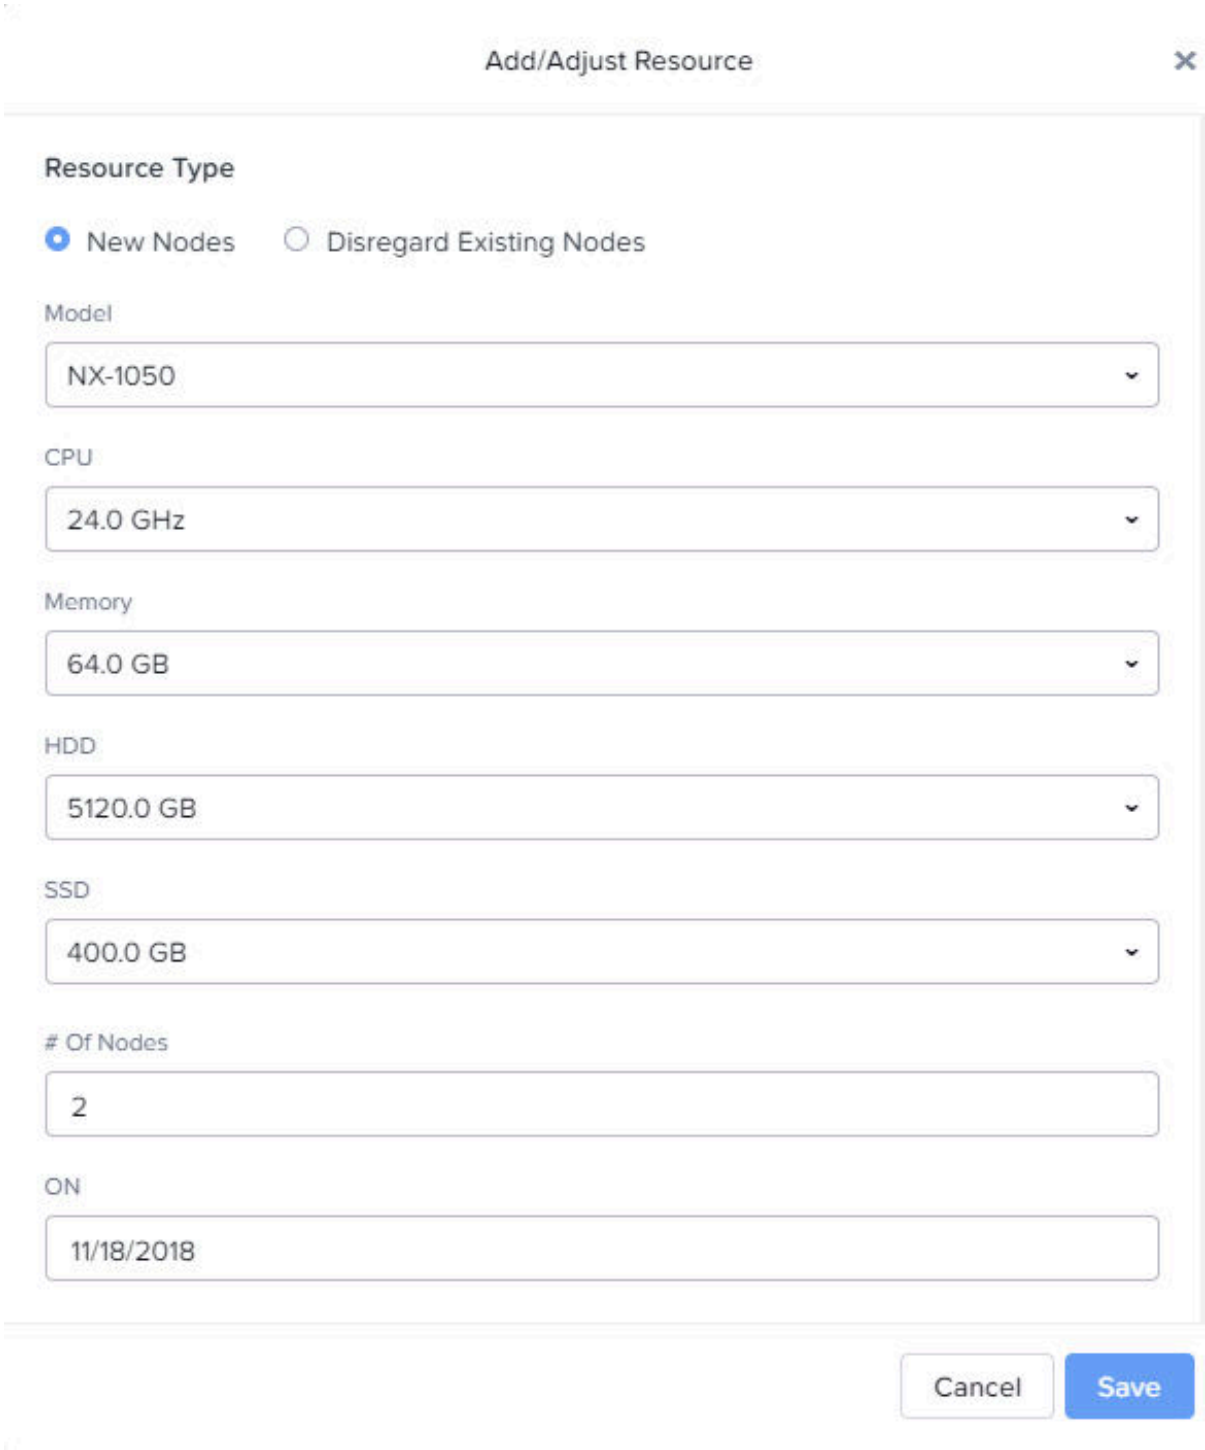

Figure 443: New Nodes Fields

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.

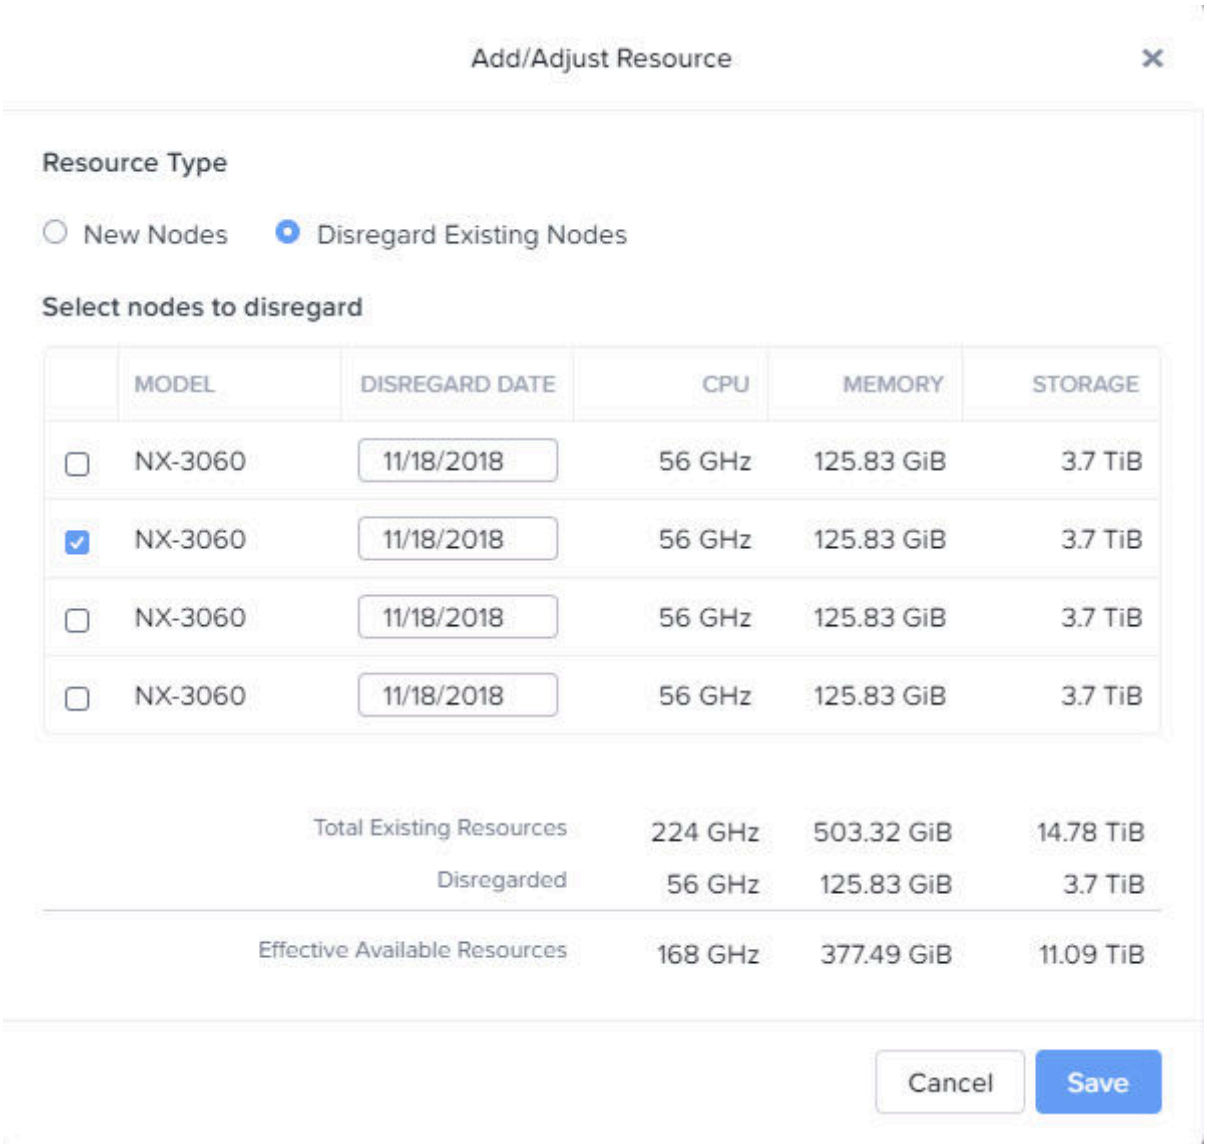

Figure 444: Disregard Existing Nodes Fields

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 819).

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.

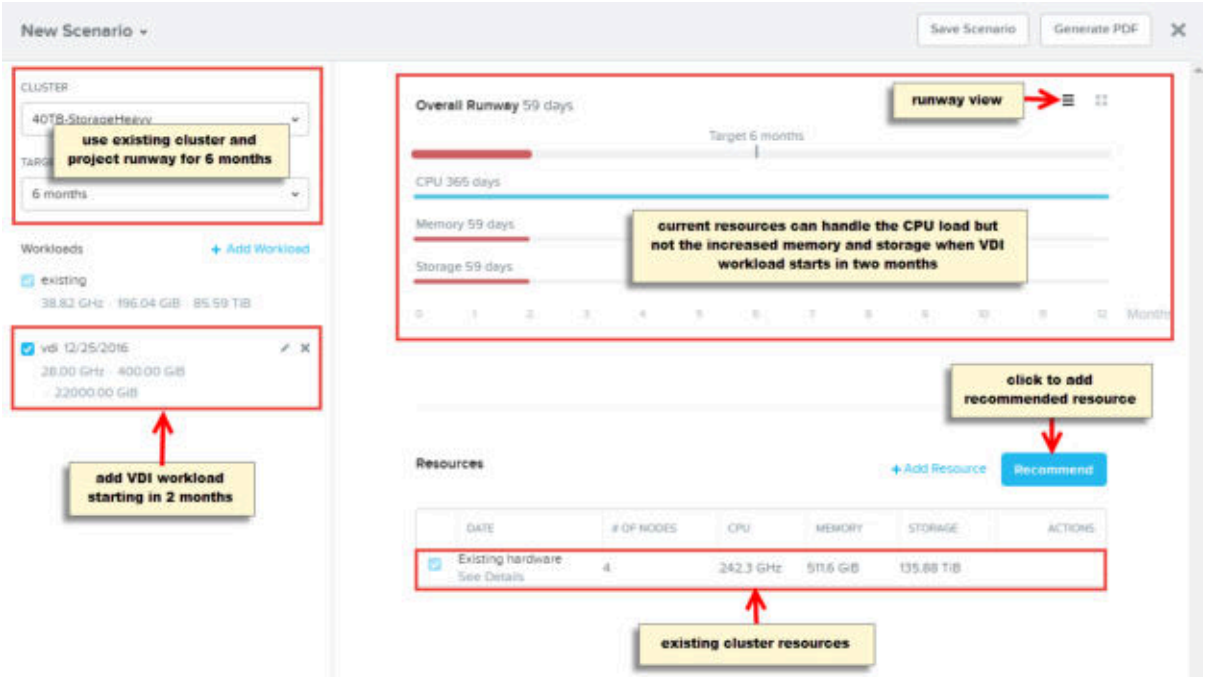

Figure 445: Scenario Screen (with added workload)

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.

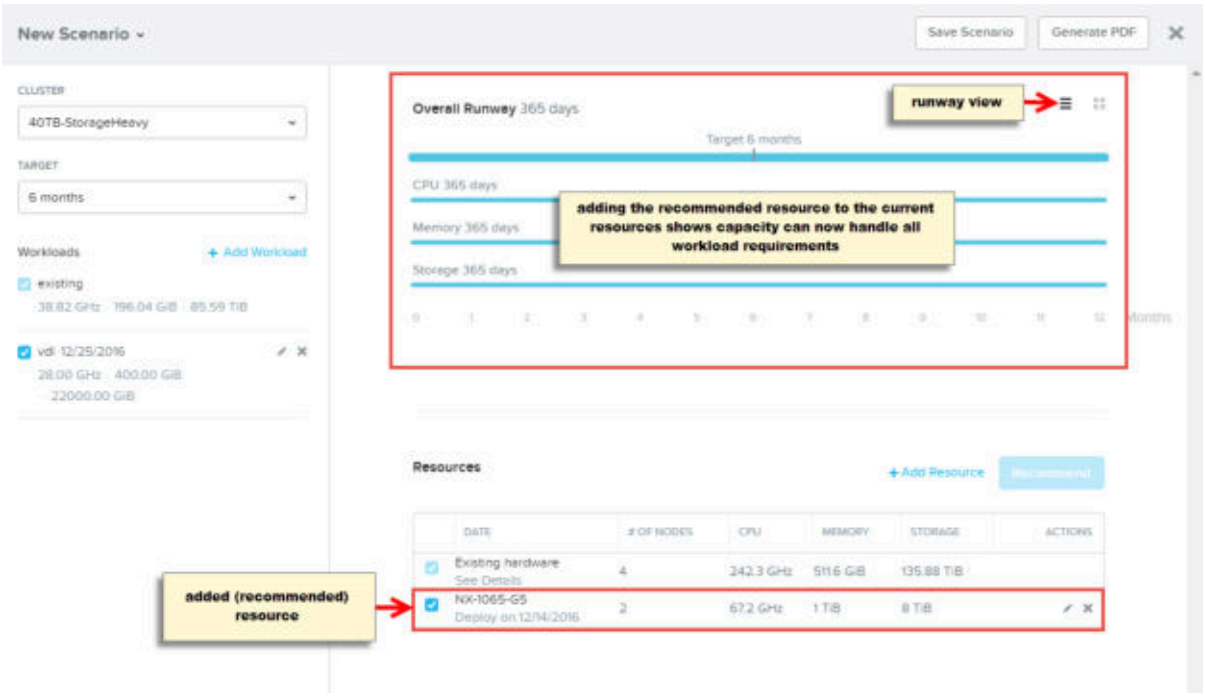

Figure 446: Scenario Screen (with added workload and recommended resource)

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.

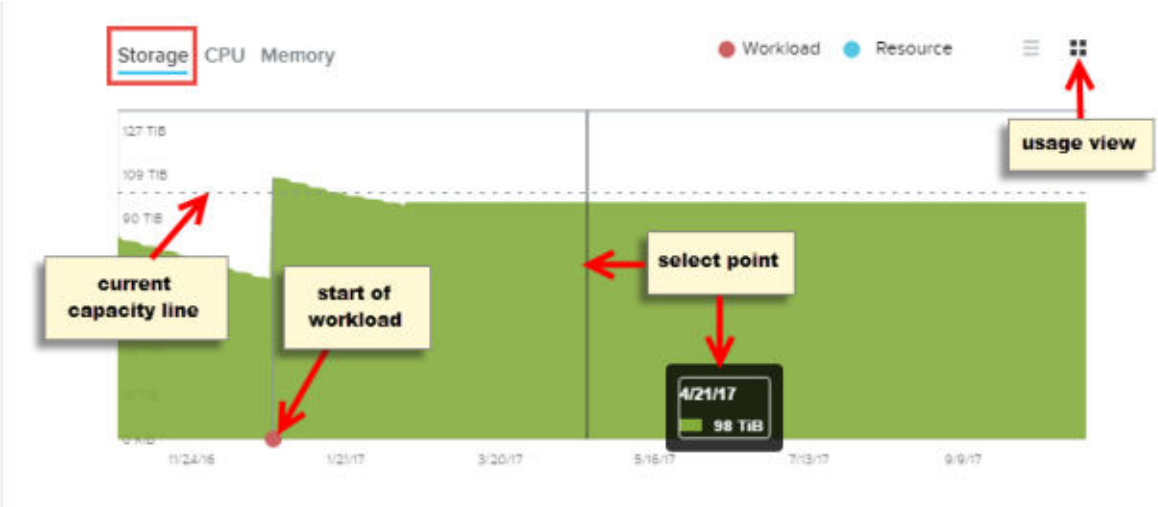

Figure 447: Storage Usage View

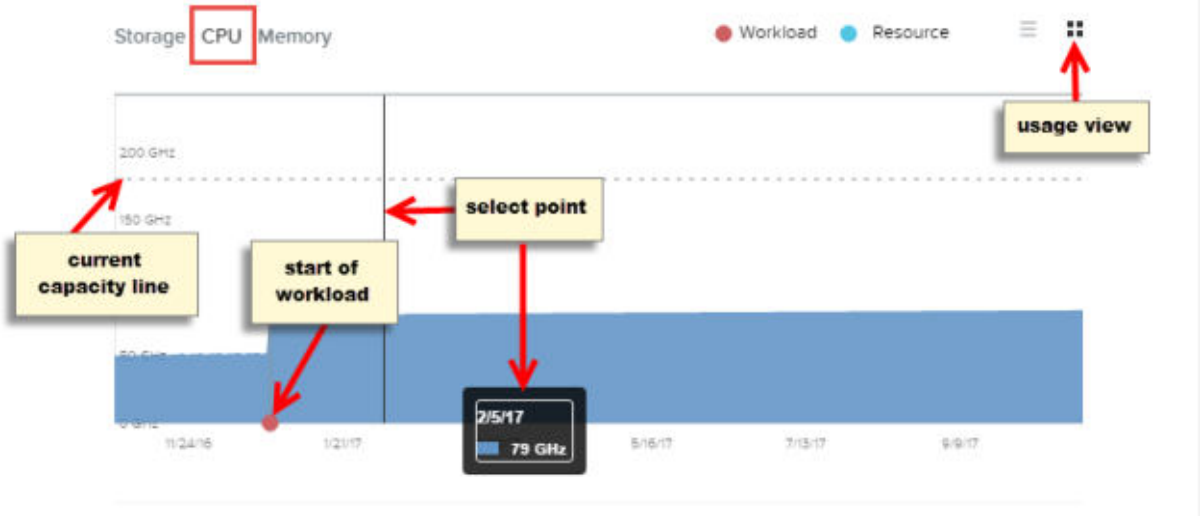

Figure 448: CPU Usage View

Figure 449: 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 819), 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 820 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 816) 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.

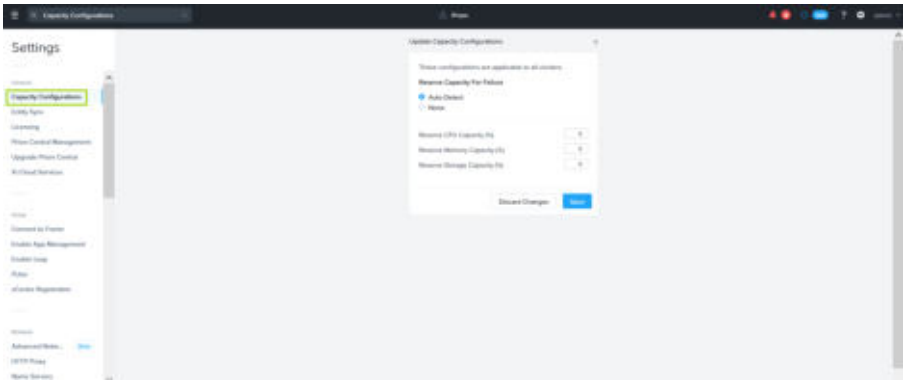

Figure 450: Update Capacity Configurations Window

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 737: Inefficient VM Profiles

Type	Description
Bully VM	A "bully" VM is one that consumes too many resources and causes other VMs to starve. A VM is considered a bully when it exhibits one or more of the following conditions for over an hour: • CPU ready time > 5% • Memory swap rate > 0 Kbps • Host I/O Stargate CPU usage > 85% One or more bully VMs might cause cluster performance to degrade. Identifying bully VMs can help in analyzing whether one or more of those VMs are misbehaving or need additional resources.
Constrained VM	A "constrained" VM is one that does not have enough resources for the demand and can lead to performance bottlenecks. A VM is considered constrained when it exhibits one or more of the following baseline values, based on the past 21 days: • CPU usage > 90% (moderate), 95% (severe) • CPU ready time > 5% , 10% • Memory usage > 90%, 95% • Memory swap rate > 0 Kbps (no moderate value) One or more constrained VMs might cause a performance bottleneck. To provide adequate host resources, resize (increase) the constrained VMs.
Over-provisioned VM	An "over-provisioned" VM is the opposite of a constrained VM, meaning it is a VM that is over-sized and wasting resources which are not needed. A VM is considered over-provisioned when it exhibits one or more of the following baseline values, based on the past 21 days: • CPU usage < 50% (moderate) or < 20% (severe) and CPU ready time < 5% • Memory usage < 50% (moderate) or < 20% (severe) and memory swap rate = 0 Kbps To prevent host resource wastage, resize (decrease) the over-provisioned VMs.
Inactive VM	A VM is "inactive" in either of the following states: • "Dead" VM: A VM is considered dead when it has been powered off for at least 21 days. • "Zombie" VM: A VM is considered a zombie when it is powered on but does fewer than 30 read or write I/Os (total) and receives or transfers fewer than 1000 bytes per day for the past 21 days. To prevent host resource wastage, delete dead and zombie VMs.

You can view information about inefficient VMs from the **VM Efficiency** and **Impacted Cluster** widgets on the home or a custom dashboard (see Main Dashboard on page 64 and Adding Dashboard Widgets on page 74) and from the **Capacity** tab of a cluster details page (see Cluster Details View on page 205).

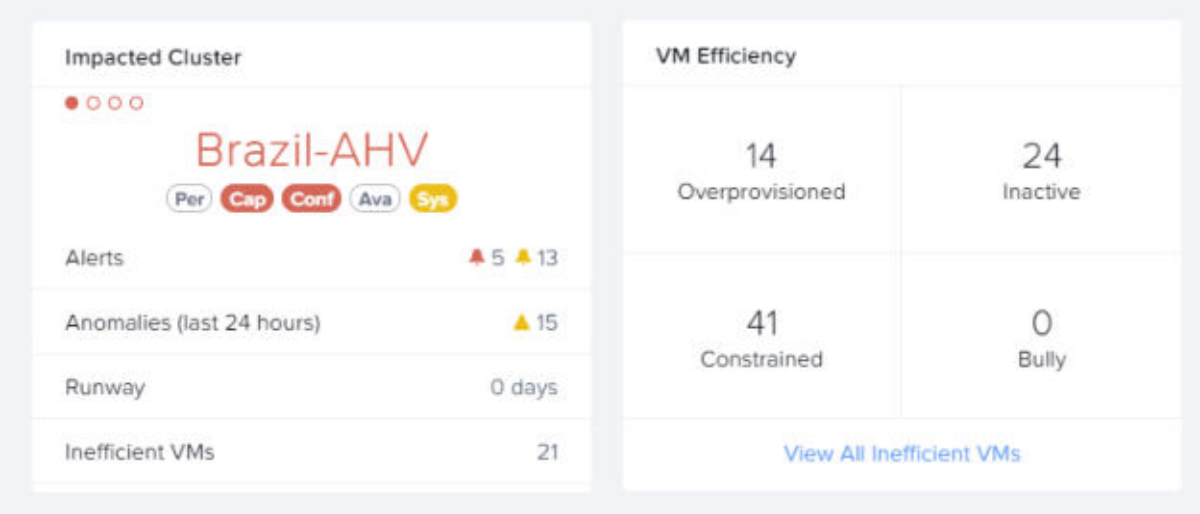

Figure 451: Impacted Cluster and VM Efficiency Widgets

You can also use a search filter from the VMs Summary View on page 86 to identify inefficient VMs.

Figure 452: Filtering VMs

Anomaly Detection

The system predicts a normal behavior band for various metrics based on historical data. The anomaly detection module monitors a predefined set of metrics on a daily basis and publishes baseline values for each of the metrics.

- Twenty-seven metrics are monitored for VMs, hosts, and clusters.

- Data for each metric from the past 21 days is recorded and analyzed, a normal behavior band is established, and predictions for the next 2 days are formulated.
- The behavior bands and predictions are adjusted accordingly when time period or trend patterns are observed, for example low CPU on weekends or increasing CPU usage.

The anomaly detection module measures usage every five minutes and compares that usage with the predicted values. If the observed value is outside the band, it flags that value as an anomaly. Each anomaly is recorded as an event (see Event Details on page 274). Anomalies appears as outliers in the behavioral anomaly event details screen and the VM details Metrics tab.

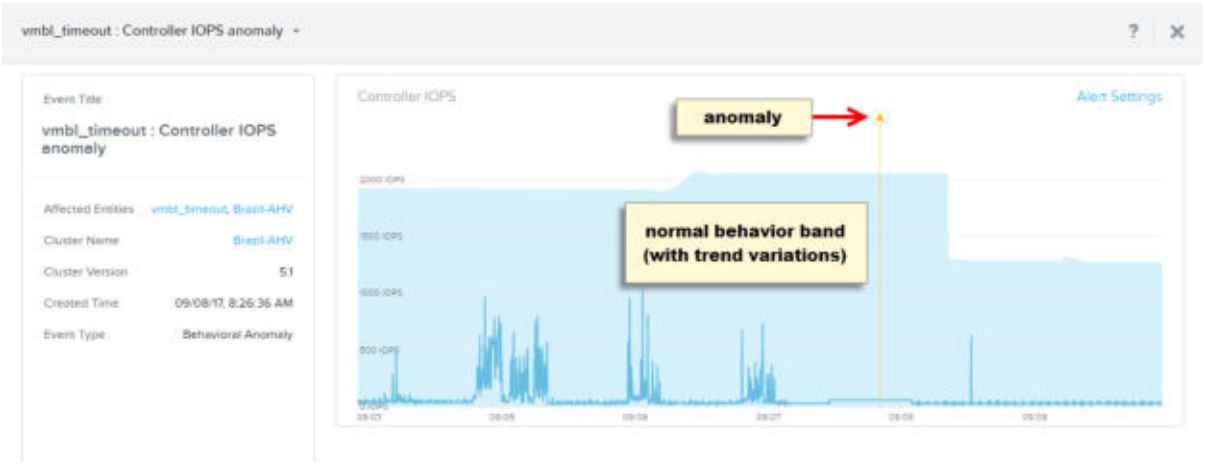

Figure 453: Anomaly Event Details Screen

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 282 for instructions on how to create a custom anomaly-based alert.

Cluster Efficiency Reports

One of the default reports (see Reports View on page 245) is the **Cluster Efficiency Summary** report. This report provides cluster statistics (host and VM counts), performance and usage statistics (average and peak CPU and memory usage), runway metrics, and information about inefficient VMs (bully, constrained, over-provisioned, and inactive). You can run this report to get a quick view of how efficiently the cluster is performing.

Task Automation

You can automate routine administrative tasks through Prism Central by using the X-Play feature.

Overview

The X-Play feature allows you to automate certain administrative tasks.

> Note: X-Play is a Prism Pro feature. If you disable the Prism Pro features, you can still create playbooks and actions, but you can't enable or run the playbooks..

X-Play is an easy to use automation tool that helps you to automate routine administrative tasks, and auto-remediate issues that may occur in your system. You can achieve this automation by creating **Playbooks**.

Playbook allows you to define a trigger that results in the execution of an action or a series of actions. A trigger may be an event that occurs in the system, such as an alert or a request made by you. The resultant actions that you configure can be VM actions, communication actions, alert actions, or report actions.

To get started, you can clone one of the predefined playbooks or create a playbook based on your requirement. You can create playbooks based on predefined alert triggers or you can create a manual trigger (see Creating Playbooks using Triggers on page 865).

You can view the instances of playbook execution using the Plays on page 842 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 843 table for details on the available system actions.

To view the action gallery, go to **Operations** > **Playbooks** > **Action Gallery**.

You can clone system actions from the action gallery, to create custom actions to use in your Playbooks.

> Note: You can not clone actions with any password filled in the action parameter.

You can update or delete any cloned actions later.

Plays

Plays are the execution instances of playbooks that you can use to monitor the status of all the playbooks that you have executed.

To view **Plays**, go to **Operations** > **Plays**. Alternatively, you can view plays from the **Plays** tab in the playbook details page.

You can view the following information within plays.

- Execution start time and execution end time of each action
- Result status
- Any inputs/outputs provided to the action
- Error messages
- Execution Flow (for branch actions)
- Autopilot-enabled playbook statistics and attempt details (see Configuring Autopilot for Playbook on page 868).

You can also select one of the following actions for the paused plays from the **Actions** menu.

- **Resume** - allows you to resume any play that is paused.
- **Abort** - allows you to prevent a paused play from resuming.

Playbook Actions

This section provides information on the playbook actions.

You can switch between the **Tiles** or **List** views by clicking the respective icons.

> Note: The VM actions configured in playbooks have certain prerequisites based on the hypervisor (AHV or ESXi) and the AOS version, see Prerequisites for VM Actions in Playbooks on page 852 for details.

You can integrate the following applications for Playbook Actions:

- Slack - see **Slack** under Communication Actions in the table.

Utility actions include:

- Lookup Cluster Details
- Lookup Host Details
- Lookup VM Details

The Utility Actions provide the following parameters as output when you run a playbook:

Table 738: Utility Actions and the Output Parameters

Utility Action	Parameters
Lookup Cluster Details	• Storage Capacity • Cluster VIP • Cluster Name • CVM IP List • Memory Size Cluster UUID • CPU Count

Utility Action	Parameters
Lookup Host Details	• Disk Capacity • Hosting Cluster VIP • CPU Count • Host IP • Hosting CVM IP • Host Name • Memory Size • Host UUID • VM Count
Lookup VM Details	• VM Name • VM Capacity • Hosting Cluster VIP • VM IP Address • VM Hypervisor Type • Host IP • vCPU Count • Hosting CVM IP • VM UUID • Power State • Memory Size • Core per vCPU

Table 739: Playbook Actions

Action	Action Description	Configurable attributes
Conditional Actions		
Branch Action	Switch execution based on a matching criteria.	See Using Branch Action (Conditional Execution) on page 873 for details.
Utility Actions		

Action	Action Description	Configurable attributes

String Parser	Parse data from a string. You can use this action to parse response from any other actions like **REST API** or **SSH Action**.	• **String to Parse** - Enter the string.
		• **Format** - Select one of the following options.
		  • **JSON** - Enter the field to retrieve as JSONPath.
		  • **XML** - Enter the node/tag/text to retrieve as XPath in the **XPath** field. Optionally, enter the comma separated list of namespaces as key=value in the XML Namespaces field.
		  • **Regular Expression** - Enter the string to parse as regular expression

Lookup Cluster Details	Captures cluster details	Select the cluster from the Target Cluster drop-down menu.
Lookup Host Details	Captures host details	Select the host from the **Target Host** drop-down menu.
Lookup VM Details	Captures VM details	Select the VM from the **Target VM** drop-down menu.
VM Actions		
Power Off VM	Power off a VM.	• Select the VM from the **Target VM** drop-down menu. • Select the **Type of Power Off Action** (Guest Shutdown, ACPI, or Power Off). Note: Guest shutdown requires NGT to be installed on the VM.
Power On VM	Power on a VM.	Select the VM from the **Target VM** drop-down menu.

Action	Action Description	Configurable attributes
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 852 for details.
VM Reduce Memory	Reduce memory for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the value for **Reduce Memory by** and **Minimum Limit**.
VM Reduce CPU	Reduce CPU for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **vCPUs to Remove** and **Minimum Number of vCPUs**.
VM Snapshot	Captures the state of a VM and saves it.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Time To Live** in number of days.
VM Add Disk	Add vDisk to a VM.	• Select the VM from the **Target VM** drop-down menu. • Select the **Bus Type** and enter the **Disk Size** in GiB.

Action	Action Description	Configurable attributes
VM Expand Disk	Expand vDisk for a VM.	• Select the VM from the **Target VM** drop-down menu. • Enter the **Expand Disk By** in GiB. • Enter the **Maximum Disk Size** in GiB. • Choose from: • **Specific Disk** - Select the **Bus Type** and enter the **Index** value for the selected disk type. • **Most Used Disk**
Communication Actions		
Email	Send an email.	Enter the **Recipient**, **Subject**, **Message**, and **Attachment**. Note: The **Attachment** field is applicable only if you have selected **Generate Forecast Report** as a previous action. Note: This action requires the SMTP server to be configured. See *Configuring an SMTP Server* in the *Prism Central Web Console Guide*.
Slack	Send a Slack message.	Enter the **Slack Token**, **Channel ID**, and **Message**. Note: You must integrate Slack to automate receiving messages in Slack, see *Integrating Slack* for details.

Action	Action Description	Configurable attributes
Microsoft Teams	Send a message to Microsoft Teams	Enter the **Webhook URL** that you configured in Microsoft Teams and the **Message**. See Microsoft Teams documentation for information about creating the Webhook URL for a Teams user or channel. For more information, see Integrating MS Teams Action in Playbooks on page 855.
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 one of the following ServiceNow target. • ServiceNow Instance	Enter the **ServiceNow Instance Name**, **Username**, and **Password** for the ServiceNow instance.
	• Mid Server	Enter the **Mid Server Address**, **Mid Server Port**, **Username**, and **Password**.
Send Alert to Webhook	Send a JSON Alert Object to a REST API endpoint.	Enter the **Alert**, **Method**, and **URL** for the outgoing webhook. Optionally, select the **Username**, **Password**, and **Request Headers**.
Report Actions		
Generate Forecast Report	Generate a just in time forecast report.	Select the cluster for which you want to generate the report and enter the days for the runway period.

Action	Action Description	Configurable attributes
Generate Report	Generate a custom report.	• Select the Report Type as **New Report**. You can also select any of previously created reports to be generated. • Enter the **Report Instance Name**. • Optionally, enter a **Description**. • Select the **Time Period for Report**.
Wait Actions		
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.

Action	Action Description	Configurable attributes
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.
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		

Action	Action Description	Configurable attributes
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.
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**.
Execution actions		

Action	Action Description	Configurable attributes
Add to Category	Add an entity like a VM, Host or Cluster to one or more Categories.	Select an **Entity Type** such as VM, Host or Cluster, select the target entity in the **Target <entity>** field and select the **Category** or categories you want to add the entity to.
Remove from Category	Remove an entity like a VM, Host or Cluster from one or more Categories.	Select an **Entity Type** such as VM, Host or Cluster, select the target entity in the **Target <entity>** field and select the **Category** or categories you want to remove the entity from.
Ansible	Run an Ansible **Workflow** or **Job** template.	Select **Template Type**. Enter **Workflow Template Name** or **Job Template Name** (as per template type selected), the Ansible **Tower/ AWS Hostname**, the OAuth authentication **Token** to log in to the Tower or AWS host, and optionally, any other parameters that you want to use in JSON format in the **Request Body JSON (Optional)** field.

> Note:
>
> The String Parser action is categorized under the Utility Actions.

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 740: 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.

Action	Hypervisor	Prerequisites	Notes
	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	
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)	

Action	Hypervisor	Prerequisites	Notes
VM Snapshot	AHV	None	• The VM snapshot action creates snapshots for AHV version lower than 5.10 and recovery points for AHV version 5.10 and above. • Maximum Time to Live value allowed is 365 days. • Snapshot can not be created for VMs which are protected by protection domain of legacy DR.
	ESXi	None	• The VM snapshot action creates recovery points for all ESXi VMs. • Maximum Time to Live value allowed is 365 days. • Snapshot can not be created for VMs which are protected by protection domain of legacy DR.
VM Add Disk	AHV	AHV does not support hot add of SATA disks.	
	ESXi	PCI disk type is not supported.	
VM Expand Disk	AHV		
	ESXi	• PCI disk type is not supported. • AOS 5.15 or above	

X-Play Integrations

This section provides information about integrating the various tools such as MS Teams, Ansible and Slack.

Integrating Slack

If you do not have a Slack account, create your account on the Slack website and set up a new workspace.

About this task

To integrate Slack for playbook actions, do the following.

> Note: You can follow these steps in an already existing workspace, however you will require permission from the workspace administrator to complete the installation steps.

Procedure

1. Go to https://api.slack.com/apps.
2. To create a Slack application, click **Create New App**.
3. In the **Create a Slack App** dialog box, do the following.
    1. Enter a name for the application.
    2. Select a development Slack Workspace. You must select the workspace on which you have the authority of accepting messages.
4. Set the app permissions - In the **Basic Information** menu, do the following.
    1. Click **Add features and functionality**.
    2. Go to **Permissions** > **Scopes** and Add **chat:write** or **chat:write.public** permission according to where you expect the application to post messages.
5. To install the application in your workspace, do the following.
    1. In **Basic Information** menu, click **Install your app to your workspace**.
    2. Click **Install App** (or **Reinstall App** if you have configured it previously). This redirects you to the **My Workspace** authorization page.
    3. Click **Authorize**.
6. To retrieve the Slack authorization token, do the following.
    1. Go to the **Install App** menu.
    2. Copy the **OAuth Access Token**. You must paste the copied token each time you use the **Slack** action in a playbook.

Integrating MS Teams Action in Playbooks

This action is used to send a message (without attachments) to a Microsoft Teams channel using a configured webhook URL.

Perform the following steps to configure a webhook for your Microsoft Teams channel:

1. In Microsoft Teams, go to the channel where you want to add a webhook and select **More options (...)** next to the channel name.
2. Select **Connectors**.
3. Search for **Incoming Webhook** in the list of connectors and select **Add**.
4. Enter a webhook name.
5. Click **Upload Image** to associate with data from this incoming webhook.
6. Click **Create**.
7. Copy the webhook URL to a clipboard and save the URL.

    > Note: This copied webhook URL is used to send information to Microsoft Teams.

8. Click **Done**.

Ansible Integration in Playbooks

The Ansible action in X-play can be used to trigger job or workflow templates in your Ansible Tower or AWX instance. Refer to the following sections to understand the prerequisites and other dependencies for Ansible actions defined in playbooks.

- More information about job templates and how to create them, see Ansible documentation.
- For information about workflow templates and how to create them, see Ansible documentation.
- More information about how to create application in Ansible Tower, see Configuring Application (Ansible Tower) on page 856
- More information about how to create application token in Ansible Tower, see Configuring Application Token (Ansible Tower) on page 858

  Note:

  - The X-Play Ansible action internally uses the following Tower APIs to launch the job and workflow templates respectively.
  - If you need to pass extra variables to your job or workflow template, ensure that you select **PROMPT ON LAUNCH** for **EXTRA VARIABLES** while creating your template from the Tower UI.
  - /api/v2/job_templates/{id}/launch/
  - /api/v2/workflow_job_templates/{id}/launch/

Configuring Application (Ansible Tower)

About this task

The X-Play Ansible action uses a token-based authentication to integrate with Ansible Tower. In Tower, you create an application that is representative of X-Play and then use it to create tokens for X-Play to use on behalf of a Tower user. Token-based authentication for users can be configured in the Applications window.

Procedure

1. In the Ansible Tower User Interface, click the Applications icon from the left navigation bar. The Applications window opens.

2. Click the (+) button located in the upper right corner of the Applications window.
   The New Application window opens.

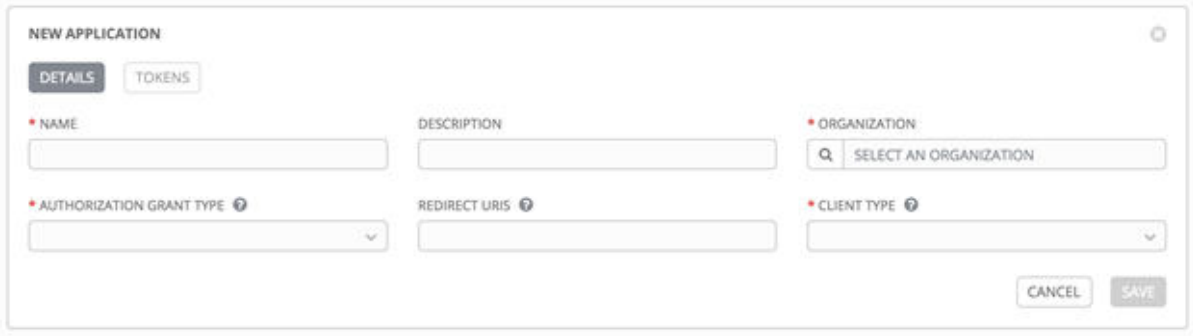

Figure 454: New Application Window

3. Enter the following details in the Create New Application window:
   a. **Name** (required): Enter a name for the application you want to create.
   b. **Description** (optional): Provide a short description for your application.
   c. **Organization** (required): Provide an organization for which this application is associated.
   d. **Authorization Grant Type** (required): Select **resource owner password-based** as the grant type.
   e. **Client Type** (required): Select the level of security of the client device as 'Confidential'.
4. Click **Save** to complete.
   Click **Cancel** to abandon your changes.

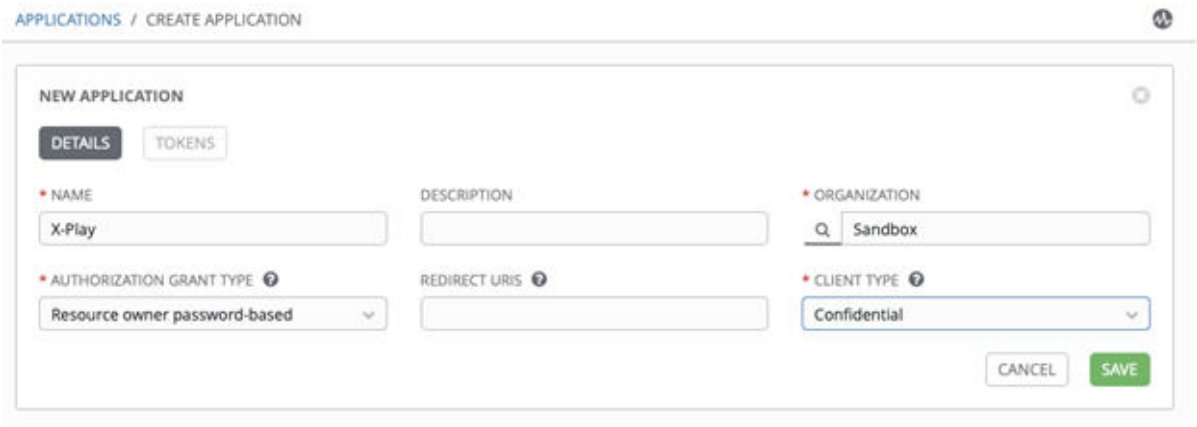

Figure 455: Create New Application Window

What to do next

For information on adding a token, see Configuring Application Token (Ansible Tower) on page 858

Configuring Application Token (Ansible Tower)

About this task

The X-Play Ansible action uses a token-based authentication to integrate with Ansible Tower. In Tower, you create an application that is representative of X-Play and then use it to create tokens for X-Play to use on behalf of a Tower user. Token-based authentication for users can be configured in the Applications window. Tokens are added through the Users screen and can be associated with an application at that time. Specifying an application can be performed directly in the User's token settings. You can create a token for your user in the Tokens configuration tab, meaning only you can create and see your tokens in your own user screen. Tokens can only access resources that its associated user can access, and can be limited further by specifying the scope of the token.

Perform the following steps to add a token.

Procedure

1. In the Ansible Tower User Interface, access the Users list view by clicking the Users icon from the left navigation bar.
2. Click on your user to configure your OAuth 2 tokens.
3. Click the **Tokens** tab from your user's profile.

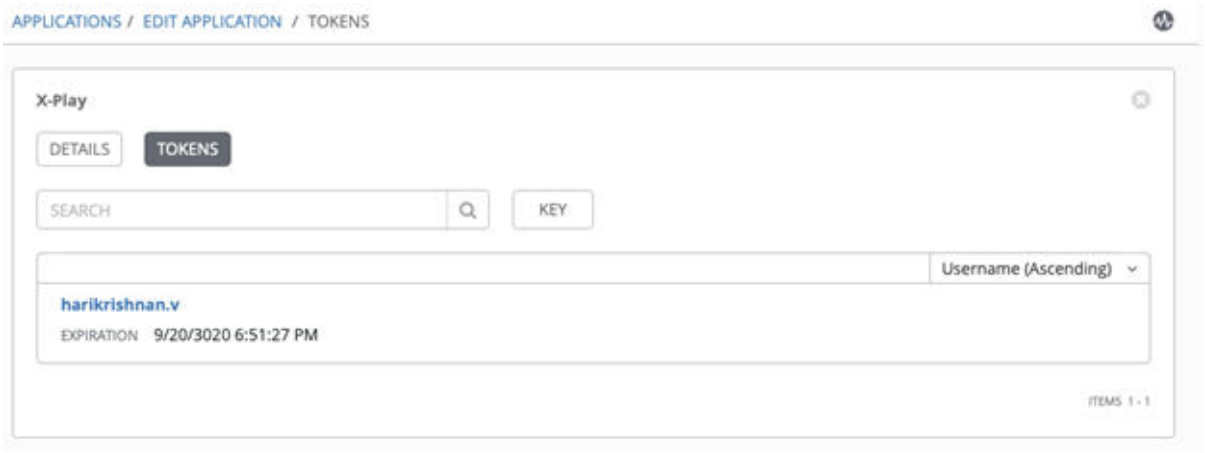

Figure 456: Tokens

4. Click the (**+**) button, which opens the Create Token window.
5. Enter the following details in the Create Token window:
   a. **Application**: Enter the name of the application with which you want to associate your token. Alternatively, you can search the application name by clicking the search button. This opens a separate window that allows you to choose from the available options. Use the Search bar to filter by name if the list is extensive.
   b. **Description** (optional): Provide a short description for your token.
   c. **Scope** (required): Select **write** so that the token can be used to launch job or workflow templates.

6. Click **Save** to complete.
   Click **Cancel** to abandon your changes.
   After the token is saved, the newly created token for the user displays with the token information and when it expires.

What to do next

- Copy the **TOKEN** field and save it so that you can provide it while creating a playbook in X-Play that uses the Ansible action.

Figure 457: Token Information

- To verify the application in the example above now shows the user with the appropriate token, go to the **Tokens** tab of the Applications window.

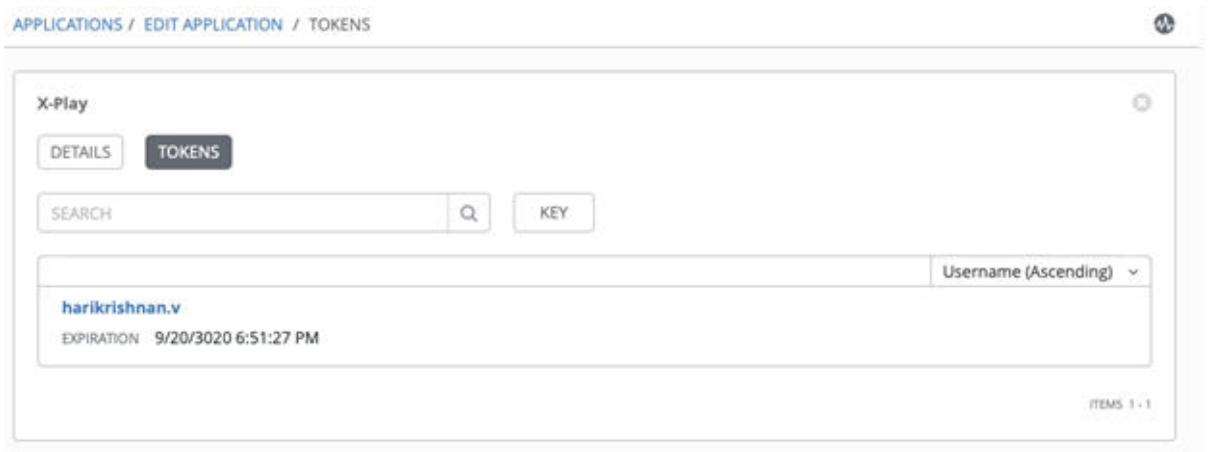

Figure 458: X-Play Tokens

ServiceNow Integration with Prism Central (X-Play Support)

ServiceNow Integration Overview

ServiceNow provides IT services management (ITSM), and IT operations management (ITOM) as a software service.

ServiceNow integration with Prism Central allows seamless incident management for Prism Central alerts using ServiceNow. Export of alerts from Prism Central is achieved by creating an X-Play playbook in Prism Central. The X-Play Action (Send to ServiceNow) pushes the Prism Central alerts to ServiceNow to be processed by the ServiceNow event management pipeline.

> Note: To integrate Prism Central with ServiceNow, enable the **Event Management** plug-in on your ServiceNow instance .

Enabling Event Management Plug-in on ServiceNow (Production Instance)

The ServiceNow **Event Management** plugin requires a separate subscription and can be activated by installing it from the ServiceNow store. . See ServiceNow documentation at https://docs.servicenow.com to enable the plugin for your production instance.

Enabling Event Management Plug-in on ServiceNow (Developer Instance)

About this task

To enable X-Play support on the ServiceNow developer instance, do the following.

Procedure

1. Go to the ServiceNow Developer Portal and create an account.
2. Log on to your account and go to **Manage** > **Instance**.
3. On the instance page, request a ServiceNow instance.
4. Once instance is provisioned, select **Action** > **Activate Plugin** for your instance.
5. On the dialog box that appears, enable the **Event Management** plug-in.

Configuring Alert Playbook for ServiceNow Alert Management

About this task

About this task

You can create a **Playbook** to send Prism Central alerts to ServiceNow.

Procedure

1. In Prism Central, click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Create Playbook**, then click **Select** in the **Alerts Matching Criteria** (option) card.
   - **Alert**: Allows you to select an alert policy as a trigger. When an alert is generated from this policy, the action configured in the playbook is triggered. For more information, see Creating Playbooks Using Alert on page 866.
   - **Alerts Matching Criteria**: Allows you to trigger playbook on any alert matching specified criteria such as severity levels, impact type, and so on. For more information, see Creating Playbooks Using Alerts Matching Criteria on page 868.
3. Click **Add Action**. Click **Select** in the **Send to ServiceNow** in the actions deck.

4. Configure the ServiceNow action.

   Note: This action requires the *Event Management Plug-in* to be enabled on the ServiceNow instance.

   a. Click **Parameters** and select `Alert Matching Criteria: Alert` in the **Alert** dropdown menu.

   b. Enter any additional information in the **Additional Info (Optional)**

   c. Select the appropriate ServiceNow instance based on the ServiceNow deployment for the **Send to** parameter.

   Select **Cloud Instance**, **Private Instance** or **MID Server** .

   If you select **MID Server**, you need to provide the MID Server IP address and Port details in next sections

   Note: The option to send the alert to ServiceNow is a default action in X-Play that calls the ServiceNow Event Collection API when a Nutanix alert matches the criteria specified.

   d. Enter the **ServiceNow Instance Name** in case of **Cloud Instance** or **Private Instance**. Enter the **MID Server IP Address** and **MID Server Port** details in case you selected **MID Server** in previous step.

   Note: Do not enter the complete URL of your ServiceNow instance in the **ServiceNow Instance Name** field. If your instance URL is *https://foo.servicenow.com*, the **ServiceNow Instance Name** you must enter is *foo*.

   e. Enter the **Username** and **Password** for the ServiceNow instance.

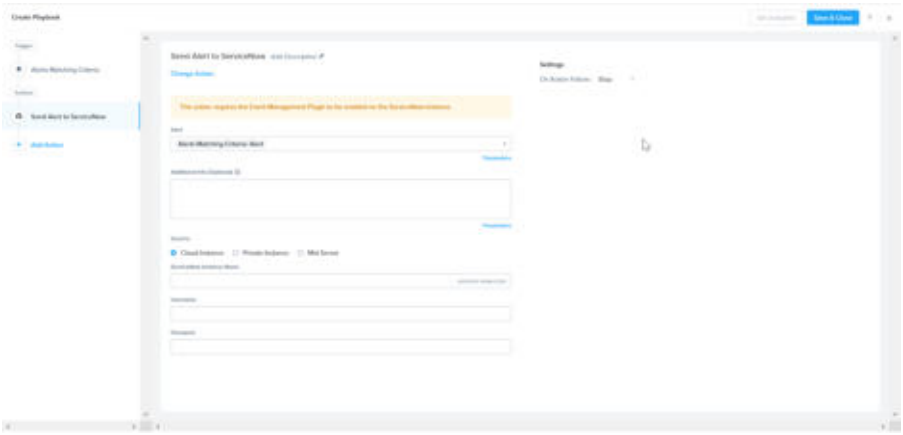

Figure 459: ServiceNow Action

5. To save the playbook, do the following.

   a. Click **Save & Close**.

   b. Specify the **Name** and provide a **Description**.

   c. Click the slider to enable the playbook. Prism Central executes only the playbooks that are enabled.

   d. Click **Save**.

   All Prism Central alerts of the specified type are forwarded to ServiceNow with the execution of the playbook created above.

What to do next

Log on to ServiceNow to view Nutanix alerts in the **Alert Intelligence All Alerts** menu after completing Importing Nutanix Event Rules to ServiceNow on page 862 and Activating the Create Incident Rule on page 862.

You can also create playbooks for actions that you want ServiceNow to initiate on Prism Central for the alerts sent to Service now. See .

Importing Nutanix Event Rules to ServiceNow

Ensure that you have imported Nutanix event rules to ServiceNow after you setup the alert playbook.

Before you begin

In ServiceNow, go to **Event Management** > **Event Rules** and check if the Nutanix event rules are already available.

The Nutanix event rules are provided in the `Nutanix Event Rule Creation` xml file. Download the *Nutanix Event Rule Creation* xml file from `download.nutanix.com` or contact Nutanix Support.

About this task

To import event rules using the *Nutanix Event Rule Creation* xml file, do the following.

Procedure

1. In ServiceNow, go to **System Update Sets** > **Retrieved Update Sets**.
2. On the right pane, click **Import Update Set from XML**.
   Import Update Set from XML is a clickable link on the right pane.
3. Browse and select the `Nutanix Event Rule Creation` xml file that you downloaded. Click **Upload**.
4. On the Update Sources page, click the uploaded Update Set to open update set details page.
5. Click **Preview Update Set**.
   If you find errors after the preview run is complete, click **Accept Remote Update** for each error. This resolves the error. Resolve all the errors.
6. After resolving all the errors, click **Commit Update Set**.

What to do next

Go to **Event Management** > **Event Rules**, Select **Source** in the search dropdown and type *Nutanix* in the **Search** box. The Nutanix Event Rules are now displayed.

Activating the Create Incident Rule

You need to activate the **Create Incident** rule in ServiceNow.

About this task

The incidence creation does not occur in ServiceNow until you set the active status of the **Create Incident** rule in Alert Management Rules in Event Management of the ServiceNow instance to `true` (Active status being `false` by default).

To change the Active status of the **Create Incident** to `true`, do the following:

Procedure

1. Go to **Event Management** > **Alert Management**. A list of Alert Management Rules is displayed.
2. Click the rule named **Create Incident**. Click (to check) the **Active** status checkbox below the rule name to set the **Create Incident** rule status to active.

Figure 460: Activate the Create Incident Rule

Configuring a Webhook Playbook for ServiceNow Action communications to Prism

X-Play now supports a webhook feature to invoke a playbook through REST calls. You can create webhook based playbooks for all the actions that you want ServiceNow to initiate on Prism Central for the alerts sent to ServiceNow.

About this task

Create webhook based playbooks using the **Webhook** trigger.

> Note: You can add one or more actions that you want SNOW to perform on Prism in one playbook.

To create a webhook based playbook, do the following.

Procedure

1. In Prism Central, click the hamburger icon from the main menu and go to **Operations** > **Playbook**.

2. Click **Create Playbook**, then click **Select** in the **Webhook** (option) card.

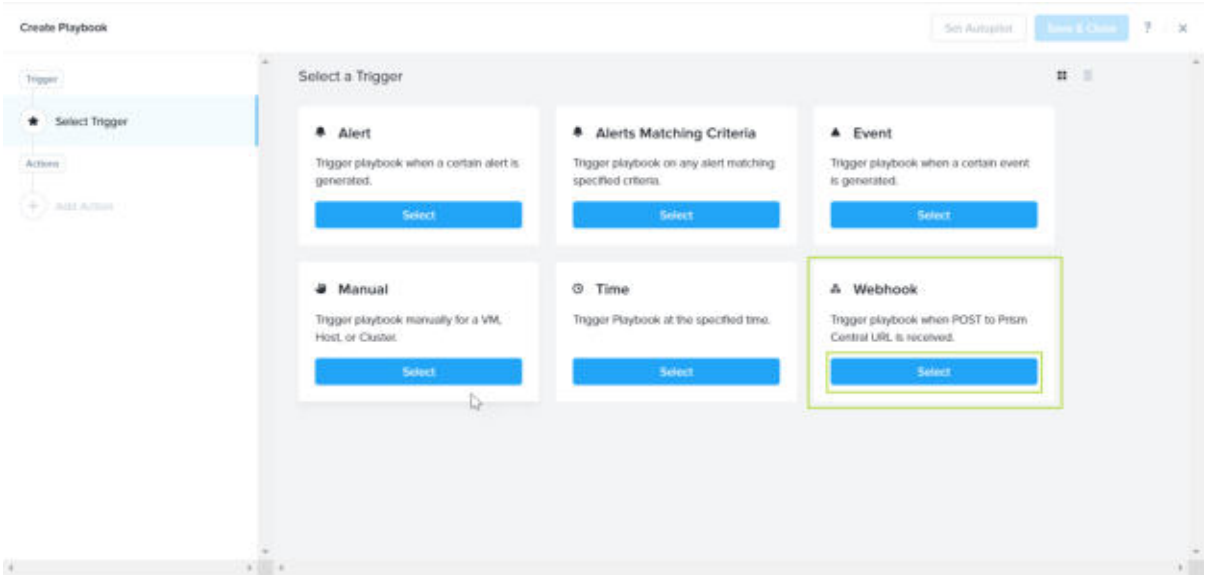

Figure 461: Select a Trigger

The Webhook page displays a message informing you to return to this playbook and review the instructions configured for this trigger.

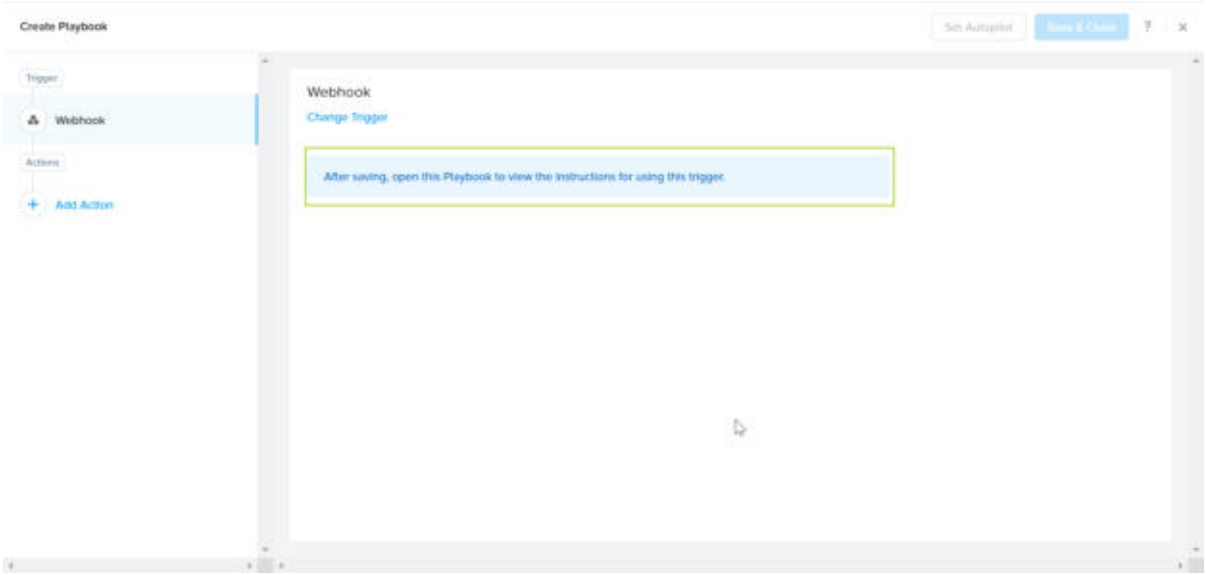

Figure 462: Webhook page

3. Click **Add Action**. Click **Select** in the appropriate action card(s) in the X-Play actions deck.

   You can select one or more actions for the playbook. After you complete configuring an X-Play action, click **Add Action** to add another action to the playbook.

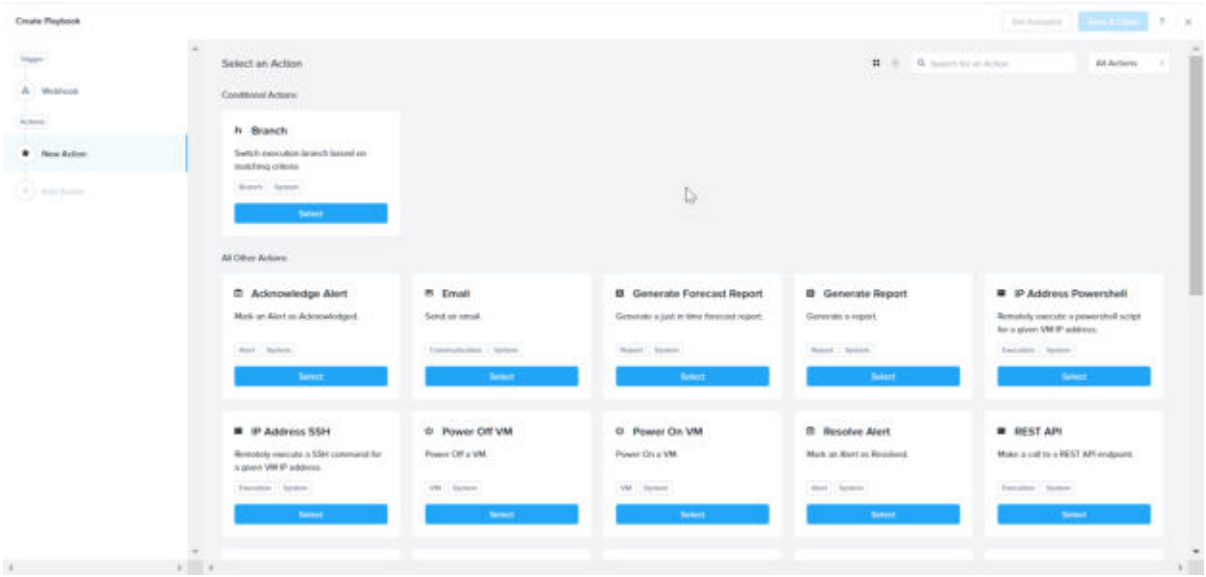

Figure 463: Select an Action page

   When you click **Select** on a X-Play action card, the action page opens. Provide all the necessary values for the parameters on the action page.

4. To save the playbook, do the following.

   a. Click **Save & Close**.

   b. Specify the **Name** and provide a **Description**.

   c. Click the slider to enable the playbook. Prism Central executes only the playbooks that are enabled.

   d. Click **Save**.

What to do next

Open the playbook you just created. Notice that the Webhook section is populated with the URL details and the body of the webhook. The URL of the webhook points to Prism Central; each webhook has a unique ID. Call this URL with the requisite parameters to trigger the playbook.

Creating Playbooks using Triggers

You can create playbooks based on the following triggers.

- Alert Trigger - The **Alert** trigger allows you to select an alert policy as a trigger. When an alert is generated from this policy, the action configured in the playbook is triggered.
- Alerts Matching Criteria - The **Alerts Matching Criteria** trigger allows you to trigger playbook on any alert matching specified criteria.
- Event - The **Event** trigger allows you to trigger playbook when an event of a certain type is generated.

- Manual Trigger - The **Manual** trigger allows you to select an entity type (VM, cluster, or host) to manually trigger the playbook from the **Explore** view for the selected entity type.
- Time - The **Time** trigger allows you to trigger playbook at the specified time.
- Webhook - The **Webhook** trigger allows triggering the playbook from a POST to the webhook URL.

The configured playbooks are available in the **Entities** menu (see Entities Menu on page 13).

Creating Playbooks Using Alert

About this task

To create a playbook using a predefined alert, do the following.

> Note: Alert as trigger allows you to select only a specific alert that can be used in a playbook. To select multiple alerts based on a matching criteria, see Creating Playbooks Using Alerts Matching Criteria on page 868.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Alert** option.
4. In the search text field, enter the name of an existing alert policy.
5. Configure the filter criteria to target the most relevant alerts and the entities associated with the alert.
6. Select the target entity type.
   - » All Applicable Clusters - The alerts from all the applicable clusters AND matching the alert criteria are used as trigger.
   - » Clusters in Categories - The alerts from only the clusters (or entity type) within the specified categories AND matching the alert criteria are used as trigger.
   - » Specific Clusters - The alerts from only the specified clusters AND matching the alert criteria are used as trigger.

7. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.

   - Branch
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   The **Parameters** option allows you to select certain inputs based on the **Trigger** that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.

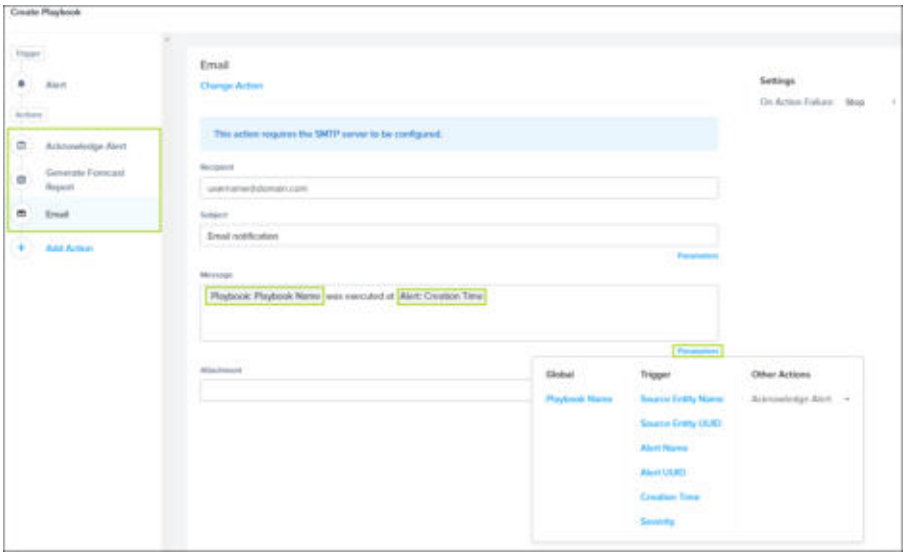

Figure 464: Parameters

To configure an action, set the action attributes based on your requirement. Refer to the Playbook Actions on page 843 table for action attributes and instructions on how to set the attributes. For each action, you must define the one of the following setting that you want to do in the event of action failure.

- Stop - Stop any further execution of the playbook when the action fails.
- Continue - Continue the execution of playbook when the action fails.

> Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action. You can also add actions before or after a particular action.

8. (Optional) Configure Autopilot for playbook.

   The Autopilot feature allows you to control the execution playbooks and automate the optimization of plays by adding specific metrics like key performance indicator (KPI) and optimal threshold for your automation tasks, see Configuring Autopilot for Playbook on page 868.

9. To save the playbook, do the following.
   a. Click **Save and Close**.
   b. Specify the **Name** and **Description**.
   c. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   d. Click **Save**.

Configuring Autopilot for Playbook

The Autopilot feature allows you to control the execution playbooks and automate the optimization of plays by adding specific metrics like key performance indicator (KPI) and optimal threshold for your automation tasks. You can use the autopilot to runs playbooks in the autopilot mode until the specified KPI is m new change

About this task

To configure autopilot for new or existing playbooks, do the following.

Before you begin

> Note: The autopilot feature is only applicable for playbooks that have metric-based alerts as triggers.

Procedure

1. Click **Set Autopilot**.
2. Select the **KPI Metric** from the drop-down menu. For example, for a **VM Memory Constrained** alert, select **Memory Usage (%)** as the KPI metric.
3. In **KPI Target Range**, select the Range option and specify the KPI target range value. Or, you can select **<** (less than) or **>** (greater than) options and specify a fixed value for the KPI target.
4. Enter the **Monitoring Duration** in minutes.
5. For **Max # Attempts**, enter the number of attempts that the playbook must run in the autopilot mode to achieve the KPI target.
6. (Optional) Check **Wait for approval for every attempt** to get an alert (INFO level) for every execution attempt of the playbook.
   If you select this option, you will have to resume the playbook after each run from the **Plays** page. The playbook remains in a suspended state until it is manually resumed after each run.
7. Click the **Autopilot Status** toggle-button to enable or disable the Autopilot. Click Save.

What to do next

You can view the details of the autopilot enabled playbooks from the Plays on page 842 page.

Creating Playbooks Using Alerts Matching Criteria

About this task

To create a playbook using alerts matching criteria as trigger, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Alerts Matching Criteria** option.
4. Select the following for either **All Alert Policies** or **Specific Alert Policies** from the following options.

   - **Impact Type** - Choose **Any** to select all impact types; or choose a single value or a combination of values from **Capacity**, **Performance**, **Configuration**, **Availability**, and **System Indicator** impact types.
   - **Severity** - Choose **Any** to select alerts having any severity levels; or choose a single value or combination of values from **Critical**, **Warning**, and **Info** severity levels.
   - **Cluster** - Choose **All Clusters** for alerts from from all clusters or **Specific Clusters** for alerts from only the specified clusters.

   For **Specific Alert Policies**, the available options are **Policies**, **Severity**, and **Cluster**.

   The actions defined in this playbook will be available for the selected entity type in the entities view.
5. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.

   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

6. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

Creating Playbooks Using Event

About this task

To create a playbook using event, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Event** option.

   > Note: VM related event triggers are applicable only to VMs running on AHV cluster. Event triggers for VM events for ESXi or HyperV VMs are not supported.

4. Select an event type from the events drop-down menu.
5. Select the target entity type by choosing one of the following options.
   - All Applicable VMs - The selected event type for all the applicable VMs is used as the trigger.
   - VMs in Categories - The selected event type for the VMs in the specified categories is used as the trigger.
   - Specific VMs - The selected event type for the specified VMs is used as the trigger.
6. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

7. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

Creating Playbooks Using Time

About this task

To create a playbook using a time trigger, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Time** option.
4. Select the **Trigger Action(s)** from the following occurrence type.
   - **Once** - Specify the day, time, and timezone.
   - **At a recurring interval** - Specify the **Recurrence** interval from daily, weekly, or monthly along with the required interval details.
5. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

6. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

What to do next

Once you create and enable a playbook with time trigger, you can go to the list view and click on the playbook to see the next time it will be executed.

Creating Playbooks Using Manual Triggers

About this task

To create a playbook using manual trigger for a VM, host, or cluster, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.

2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Manual** option.
4. Select the **Entity Type** from the following options.
   - VM
   - Host
   - Cluster

   The actions defined in this Playbook will be available for the selected entity type in the entities view.
5. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.
   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

6. To save the playbook, do the following.
   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

Creating Playbooks Using Webhook

About this task

To create a playbook using a webhook as trigger do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Click **Get Started** and then click **Create Playbook**.
   The **Get Started** option appears only when you are creating a playbook for the first time.
3. Click **Select** within the **Webhook** option.

4. Click **Add Action** to view the list of all the available actions on the deck. Click the action type filter to filter from a list of following action categories.

   - VM Actions
   - Communication Actions
   - Alert Actions
   - Report Actions
   - Wait Actions
   - Execution Actions

   To configure an action, set the action attributes based on your requirement. Refer to the *Playbook Actions* table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue)

   > Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

5. To save the playbook, do the following.

   1. Click **Save and Close**.
   2. Specify the **Name** and **Description**.
   3. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   4. Click **Save**.

   After saving, open this playbook to view the instructions for using this trigger.

6. Once saved, open the playbook to view and copy the webhook parameters.

   Using the parameters (**API Method**, **URL**, **ID**, and **Body**) and values shown to create an incoming webhook for the external application to trigger playbook.

Using Branch Action (Conditional Execution)

You can create an X-Play playbook with conditional branch action for a trigger. Branch action allows you to switch execution based on a matching criteria.

About this task

To use a branch action, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Select and configure one of the following trigger.

   - Alert Trigger
   - Alerts Matching Criteria
   - Event
   - Manual Trigger
   - Time
   - Webhook

3. Click **Add Action** and select **Branch**.
4. Define condition for the branch.
   - Branch - Click the pencil icon to add a description for the conditional branch.
   - Condition - Select either **If**, **Else**, or **Else If** depending on the condition context that you want to set.

     Note:
     - You can create multiple branch actions for a trigger.
     - Once the playbook is run, you can view the **Execution Flow** of your playbook. To view the execution flow of any previously run playbook, go to **Operations** > **Plays** and click on the playbook name. To view the complete execution flow, click **View Full Playbook**.
   - Operand - Click **Parameters** to select a **Global** or **Trigger** entity.
   - Operator - Select the required operator from the drop-down menu. You can choose **regexp** for pattern matching.
   - Value - Enter the value for the operator, or click **Parameters** to select a **Global** or **Trigger** entity.

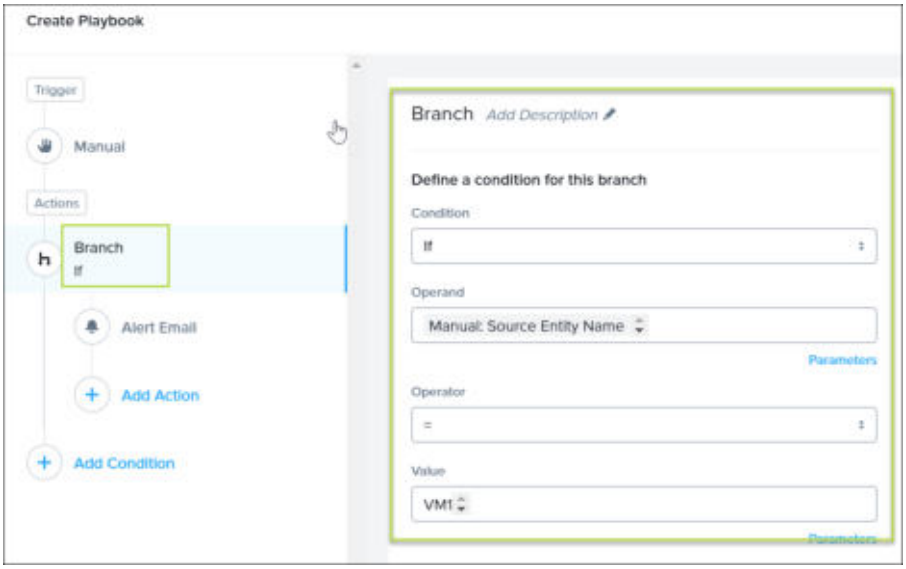

Figure 465: Branch Action

5. Click **Add Action** to view the list of all the available actions on the deck.

   To configure an action, set the action attributes based on your requirement. Refer to the Playbook Actions on page 843 table for action attributes and instructions on how to set the attributes.

   Note: For each action branch condition, you must add at least one action.

6. Create the execution flow using more branch actions (**Else-If** or **Else** options) and adding required actions for each condition.

   You can also configure nested branch actions as required. To delete any added branch action, click the three-dots icon select **Delete**.

7. To save the playbook, do the following.

   a. Click **Save and Close**.

   b. Specify the **Name** and **Description**.

   c. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.

   d. Click **Save**.

Running a Playbook (Manual Trigger)

About this task

You can run a playbook that has a manual trigger associated with the playbook.

Procedure

1. Click the hamburger icon from the main menu and go to **Compute & Storage** > **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 878 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.

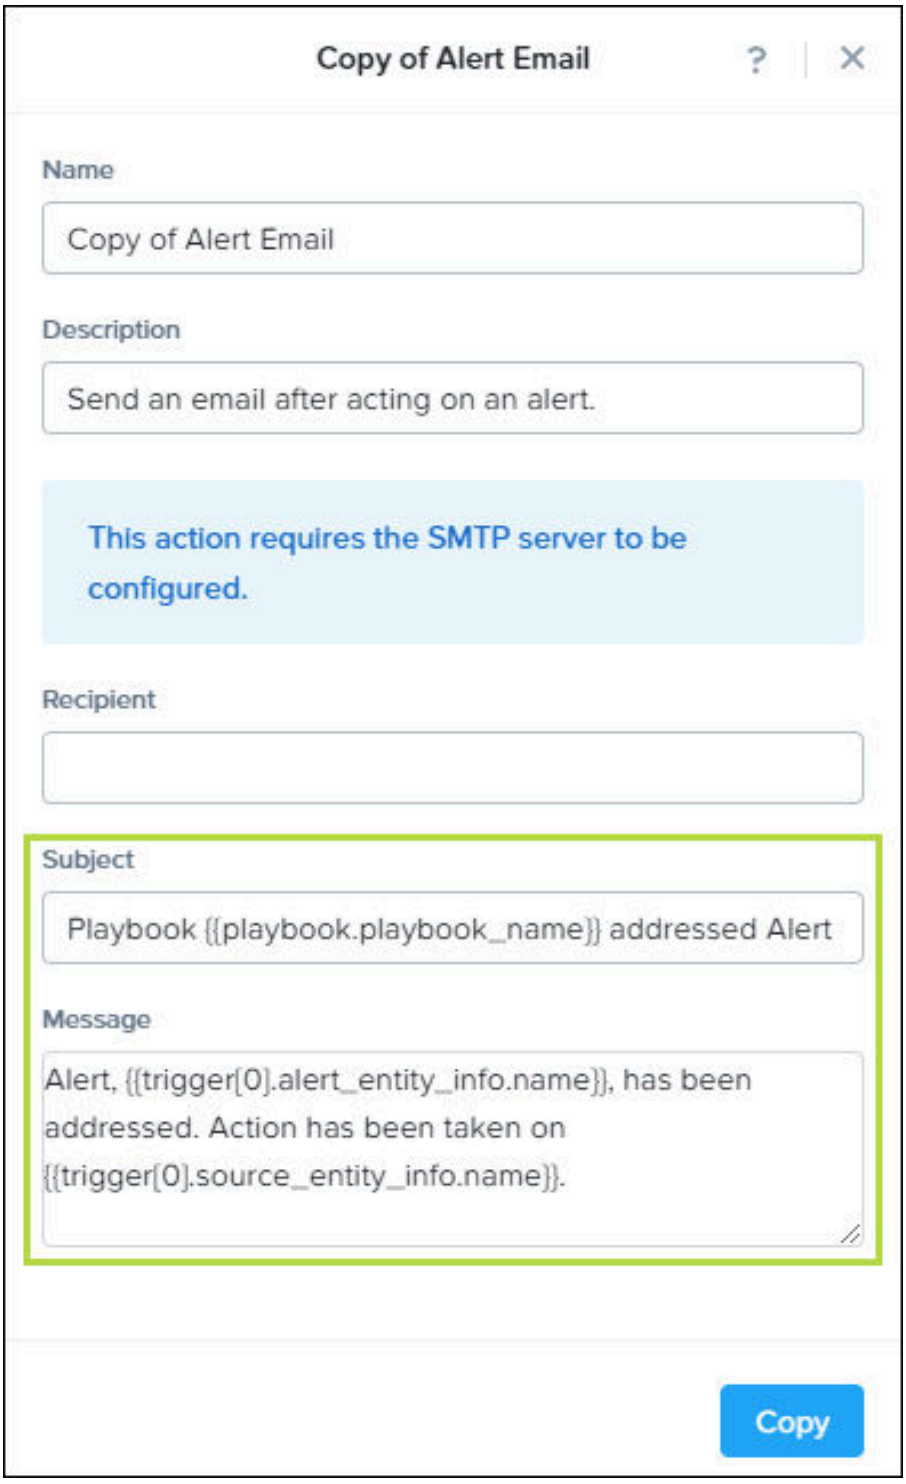

Figure 466: Configuring parameters manually during Clone operation

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 878 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 741: Parameters Table

Entity	Parameter	
Playbook		
Playbook Name (string)	{{playbook.playbook_name}}	
Alert Trigger		
Source Entity Name (string)	{{trigger[0].source_entity_info.name}}	
Source Entity UUID (string)	{{trigger[0].source_entity_info.uuid}}	
Source Entity (object)	{{trigger[0].source_entity_info}}	
Alert Name (string)	{{trigger[0].alert_entity_info.name}}	
Alert UUID (string)	{{trigger[0].alert_entity_info.uuid}}	
Alert (object)	{{trigger[0].alert_entity_info}}	
Creation Time (long)	{{trigger[0].creation_time}}	
Severity (string)	{{trigger[0].severity}}	
Manual Trigger		
Source Entity Name (string)	{{trigger[0].source_entity_info.name}}	
Source Entity UUID (string)	{{trigger[0].source_entity_info.uuid}}	
Source Entity (object)	{{trigger[0].source_entity_info}}	
Playbook UUID (string)	{{trigger[0].action_rule_uuid}}	
Action		
Generate Forecast Report	Report PDF (file)	{{action[index].report}
	Report Name (String)	{{action[index].report.fileName}}
	Report Data (bytes)	{{action[index].report.fileData}}
IP Address Powershell	Return Code (long)	{{action[index].ret_code}}
	Standard Output (string)	{{action[index].stdout}}

Entity	Parameter	
	Standard Error (string)	{{action[index].stderr}}
IP Address SSH	Return Code (long)	{{action[index].status_code}}
	Standard Output (string)	{{action[index].response_headers}}
	Standard Error (string)	{{action[index].response_body}}
Slack	Channel Id (string)	{{action[index].channel_id}}
	Message Timestamp (string)	{action[index].timestamp}}
VM Add CPU	Resultant vCPU Count (long)	{{action[index].result_vcpu_count}}
VM Reduce Memory	Resultant Memory (double)	{{action[index].total_memory}}
VM Snapshot	Recovery Point UUID (string)	{{action[index].vm_recovery_point_entity.
	Recovery Point Name (string)	{{action[index].vm_recovery_point_entity.
	Recovery Point (object)	{{action[index].vm_recovery_point_entity}
VM SSH	Return Code (long)	{{action[index].ret_code}}
	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
Wait for Some Time	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}
Wait Until Day of Month	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (String)	{{action[index].planned_resume_time}}
Wait Until Day of Week	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}

Exporting or Importing Playbooks

You can export a playbook for backup reasons or import any previously exported playbook. If you export a playbook using a Prism Central instance, you can import that playbook to another Prism Central instance that is running the same or newer version only.

About this task

To export a playbook, do the following.

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook**.
2. Select one or more playbooks by clicking the check-box against the playbook name
3. Click **Actions** > **Export**.
4. Optionally edit the default **Filename** and click **Export**.
   The playbook backup is downloaded in the .pbk format.

About this task

To import a playbook, do the following

Procedure

1. Click the hamburger icon from the main menu and go to **Operations** > **Playbook** and click **Import**.
2. Click **Browse** to select the previously exported playbook backup file and click **Import**.

PagerDuty Integration with Prism

PagerDuty is an incident management platform that provides reliable notifications, automatic escalations, on-call scheduling, and other functionality to help teams detect and fix infrastructure problems quickly.

The PagerDuty integration with Prism Central allows receiving Prism alerts as incidents in PagerDuty interface. You can create a Playbook using the new alert action called **Send Alert to PagerDuty**. This alert action sends Prism alerts (from Prism Central and Prism Element) to PagerDuty. The alerts appear as PagerDuty incidents in the PagerDuty **Incidents** page.

Contact Nutanix Support for any issues with PagerDuty integration with Prism.

Setting up the PagerDuty Integration

This procedure explains how to set up the PagerDuty integration key. This key is used as the routing key to integrate PagerDuty with Playbooks.

About this task

To integrate Prism Central with PagerDuty service, do the following in the PagerDuty interface:

Procedure

1. Log in to PagerDuty
2. From the **Services** menu, click **New Service**.

   Create a Service page appears.

   Creating a new service includes these steps:

   - Step 1: Name
   - Step 2: Assign
   - Step 3: Alert Grouping
   - Step 4: Integrations

   Follow the instructions outlined in Create a New Service section of PagerDuty documentation until the Step 3: Alert Grouping.
3. In the **Integrations** step, select **Events API V2** as the integration type, and click **Create Service**.

   > Note: The integration type 'Nutanix Prism' is not yet available in the PagerDuty instance. Therefore, you must select Events API V2 as the integration type until the Nutanix Prism name is added in the PagerDuty.
4. You will be redirected to the **Integrations** tab of your newly created service.

   An **Integration Key** will be available on this screen. Keep this key saved in a safe place, as it will be used when you configure the integration with Prism.

Configuring a Playbook with the PagerDuty Integration

You can create a Playbook to send Prism alerts to PagerDuty. The Playbook sends the alerts from Prism Central and Prism Element to the PagerDuty instance.

About this task

This procedure shows how to create a Playbook to send Prism alerts to PagerDuty.

Procedure

1. In Prism Central, click the hamburger icon and go to **Operations > Playbook.**
2. Click **Create Playbook**.
   The **Select a Trigger** page appears.
3. Select **Alerts Matching Criteria** as a trigger.

   The **Alerts Matching Criteria** allows you to choose all or specific alert policies.
4. In the **Alerts Matching Criteria** page, do the following:

   » Select **All Alerts Policies** to send alerts coming from all existing alert policies, or select **Specific Alert Policies** to send alerts coming from one or more specific alert policies. You can type in the alert policy name and search, or scroll the list of existing alert policies and select multiple alert policies.

   » Specify the **Impact Type**.

   You can specify more than one Impact Type.

   Impact type is the category in which the alerts are classified such as Availability, Configuration, Capacity, Performance, and System Indicator alerts.

   » Specify the **Severity** of the alerts.

   » Specify the cluster, either choose **All Clusters** or **Specific Clusters**.

   If you choose a specific cluster, then only alerts coming from that cluster trigger the action specified in Step 6.

   Every time there is an alert raised by the selected alert policy, the system will trigger actions defined in this Playbook.
5. Configure the filter criteria to target the most relevant alerts and the entities associated with the alert. The filter criteria options that appear depend on the type of alert policy that you have selected in the Step 4.

   The filter options can be any of these:

   - Severity
   - Target VM
   - Target Cluster
   - Target Host

6. Click **Add Actions**.
   A list of all the available actions on the deck appears. Select **Send Alert to PagerDuty** option.

   You can search by typing the action name in the search bar, or by using the **Alert Actions** filter option.

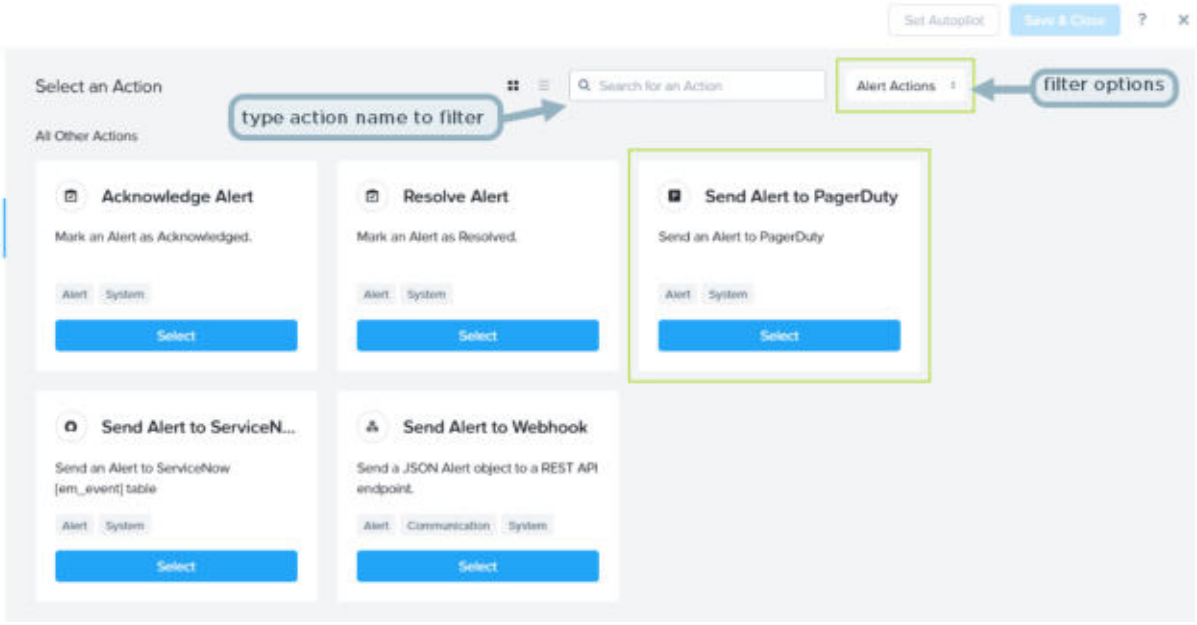

Figure 467: Send Alert to PagerDuty

7. Enter the **Routing Key**.
   The Routing Key is an integration key that you can find in your PagerDuty instance at this location: **Service Directory> Event API Integration> Integrations** tab.

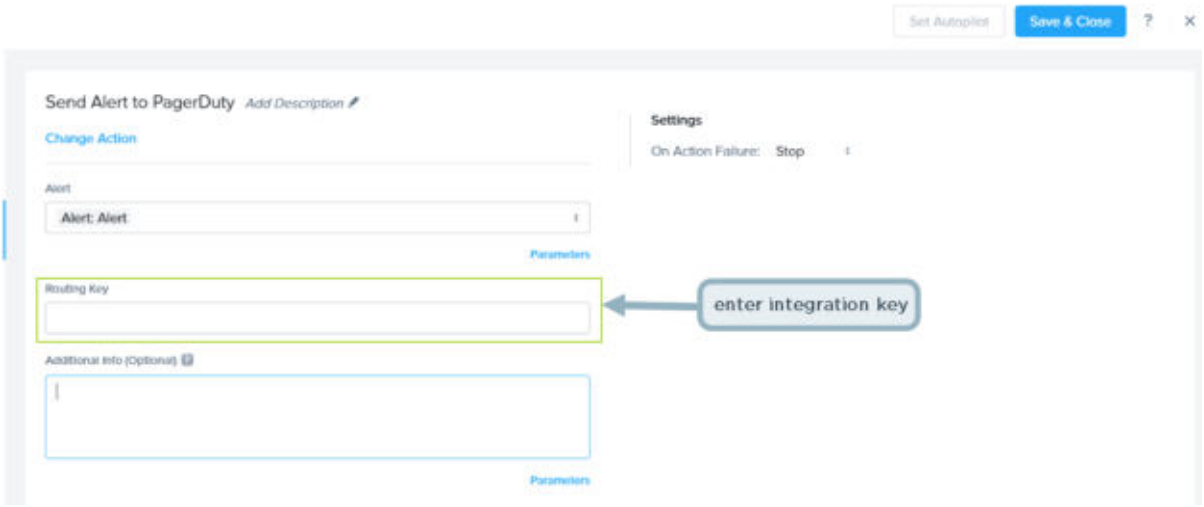

Figure 468: Routing Key

8. Select **Parameters**.
   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.

9. To save the playbook, do the following.
   a. Click **Save and Close**.
   b. Specify the **Name** and **Description**.
   c. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
   d. Click **Save**.

Reports Management

The report management feature provides you with an ability to configure and deliver the historical reports containing information about the infrastructure resources across Nutanix managed and Non-Nutanix managed environments. This feature provides the operational insights of your infrastructure into your mailbox according to the schedule that you have configured.

A generated report can be divided into three main components.

- Style: Defines the way to represent the report, for example background color, logos, and so on.
- Representation of data: Defines the way that you want to represent the data. You can accomplish this by including different views (pre-defined and customizable).
- Data: Defines the actual data depending on your selection, for example information about different metrics; CPU usage, memory usage, IOPS bandwidth, VM count, host count, cluster count, license summary, and so on.

By default the **Reports** dashboard displays the **Cluster Efficiency Summary** and **Environment Summary** reports. The cluster efficiency summary report provides a detailed information about the predictive utilization and runway information of all the resources. The environment summary report provides information about the summary configuration of each cluster (host count, VM count, license information, CPU usage information, IOPS bandwidth, Controller VM I/O bandwidth) along with the licensing information that is registered to Prism Central.

> Note: You cannot modify or delete the pre-defined or factory shipped reports such as the 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 884.

> Note: All the built-in roles except the **Prism Viewer** role can manage the reports (create, update, delete). Users with **Prism Viewer** privileges can only view the reports.

Internationalization in simplified Chinese and Japanese is supported. The generated report can have English, Chinese, and Japanese characters.

Some of the salient features of report management are as follows.

- Ability to generate and customize end-to-end reports.
- Ability to add different views to customize what data is displayed and how that data is represented.
- Ability to download the report in PDF and CSV format.
- Create a report definition from another report definition. You can use this feature to save time for a report that has a minor change from an existing one and copy from a shared read-only report from the other users.

- Customize the report definition. Ability to add a logo and make changes to the appearance of the report generated.
- Ability to send the reports through Email. Generated report in the PDF format or CSV format, or both, are sent as an attachment to the Email.
- Ability to schedule the report. Note that you can define only one schedule for a report definition.
- Ability to retain the reports for the specified period of time.
- Ability to check the report log details for the status of the report and any error message (if the report generation fails).
- Ability to create Role Based Access Control (RBAC) to define different roles and assign permissions accordingly.

Requirements

- License requirements: Prism Central should be running Prism Pro license.

Limitations

The report management feature in Prism Central has the following limitations.

- The report management feature currently only supports generation of reports in PDF format for 2000 entities. Nutanix recommends generating reports in CSV format for such scenarios.
- Generation of a report in CSV format supports a maximum of 25000 entities.
- Generation of a report in PDF format can only consist of a maximum 10 columns for a data table.
- The entire process of report generation fails if either of the report formats (PDF or CSV) fails to generate. This behavior is observed when both the formats are selected during report generation.

Alerts Generation

Alerts are generated during following failure scenarios.

- Generating a report
- Sending of a scheduled Email

Creating a New Report

Perform the following to create a new report.

Procedure

1. In the **Reports** dashboard (see Reports View on page 245), 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 900.
3. Configure a schedule for the report. For more information about scheduling, see Scheduling a Report on page 920.
4. Configure the settings specific for this report. For more information about configuring report settings, Configuring Report Settings on page 923.

5. To save the report, do one of the following.

   - Click **Save**.

     This option saves the report and does not generate the report.

   - Click **Save and Run Now**.

     This option saves and generates the report. A **Run Report** window appears. For more information, see Generating a Report on page 885.

   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).

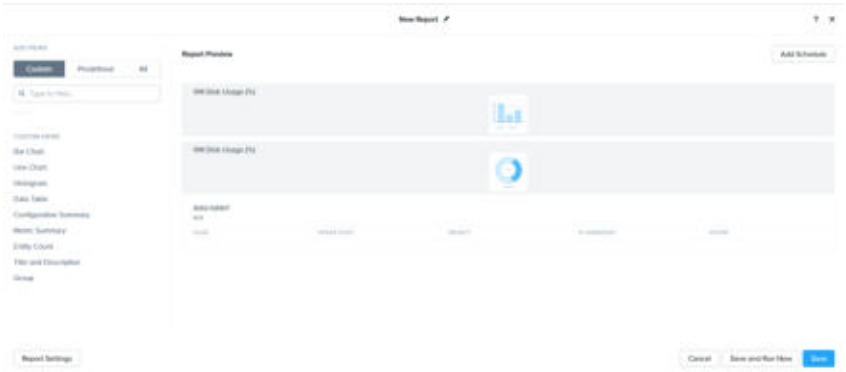

Figure 469: Create Report

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 885.
- To edit a report, see Editing a Custom Report on page 886.
- To clone a report, see Cloning a Report on page 887.
- To delete a report, see Deleting a Custom Report on page 887.

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 245), 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 923.

7. Click **Run**.
   A report is generated.

Editing a Custom Report

You can edit an existing custom report. You can modify views and change other report settings in the Edit mode.

About this task

To edit a custom report, do the following.

Procedure

1. In the **Report** dashboard (see Reports View on page 245), 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 884 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 913.

     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 custom 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 245), 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 884 to add views and create a new instance of the report.
   The report is cloned.

Deleting a Custom Report

Deleting a custom report removes all the generated report instances along with the selected report configuration.

About this task

To delete a custom report, do the following.

Procedure

1. In the **Report** dashboard (see Reports View on page 245), click the check box against the report that is already generated.
2. From the **Actions** drop-down menu, select **Delete**.
   The custom 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 245), 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 470: Viewing Instance Information (Nutanix)

Figure 471: Viewing Instance Information (vCenter)

> Note: If the report generation fails then the Status column displays the status as `Failed` and if you hover on the field, you can view the reason for the failure.

a. **Resend Report**:

- Select this option to send the report again to all the Email recipients that you have configured.
- Select the format of the report (**PDF** or **CSV** or both) before you send an email, under **Report Format**.

You also have an option to send the report to additional recipients by clicking **+ Add Additional Recipients** link. If you do not add any additional recipients, the report is sent to the recipients that are already configured in the **Report Settings** flow.

b. **Delete**: Select this option to delete the selected report instance.

Example Report - Environment Summary

A generated environment summary report has the following information.

Sample Environment Summary Report

- Stylistic information: The first page of the report is used to represent this kind of information. The stylistic information includes the name of the report, time it was generated, any logo

that you have configured, header and footer information (for example, report name and page number).

Figure 472: 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 environment summary information for Nutanix and vCenter clusters.

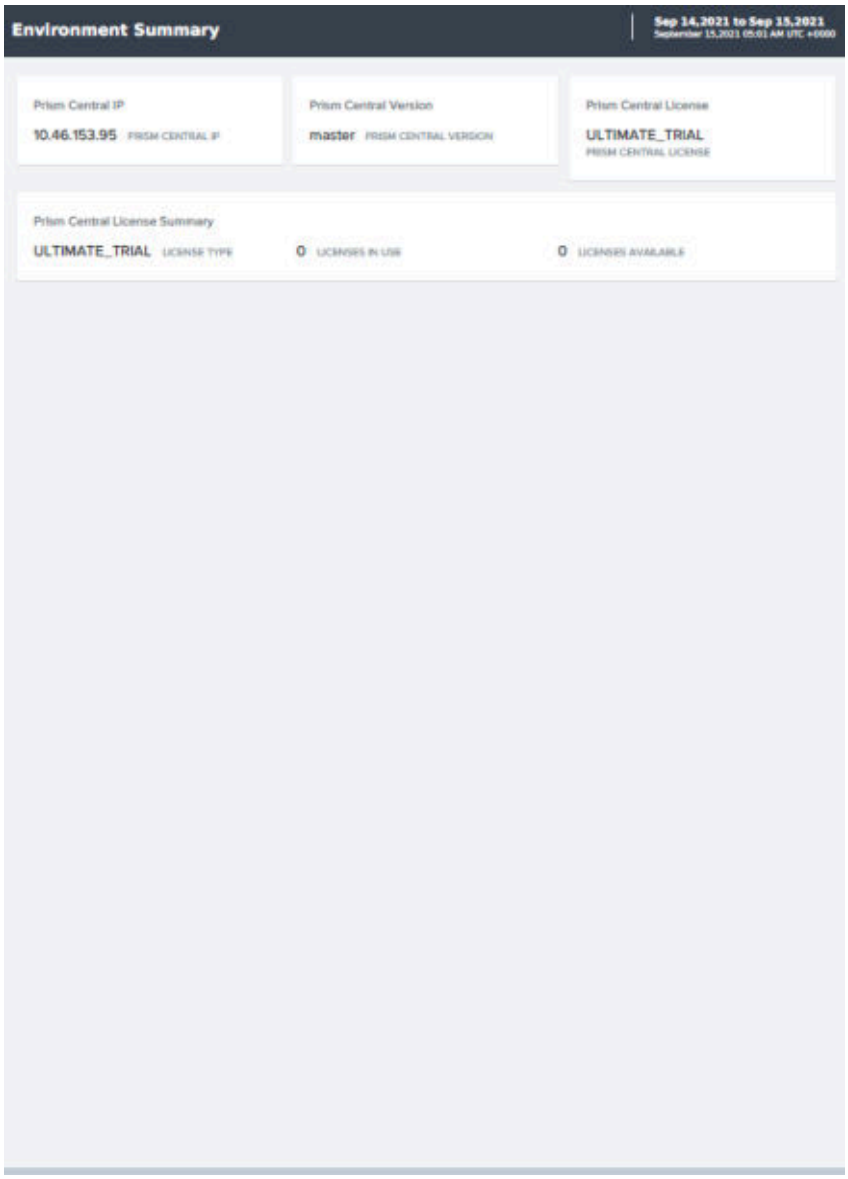

Figure 473: Detailed Metrics Information (Nutanix: Image 1)

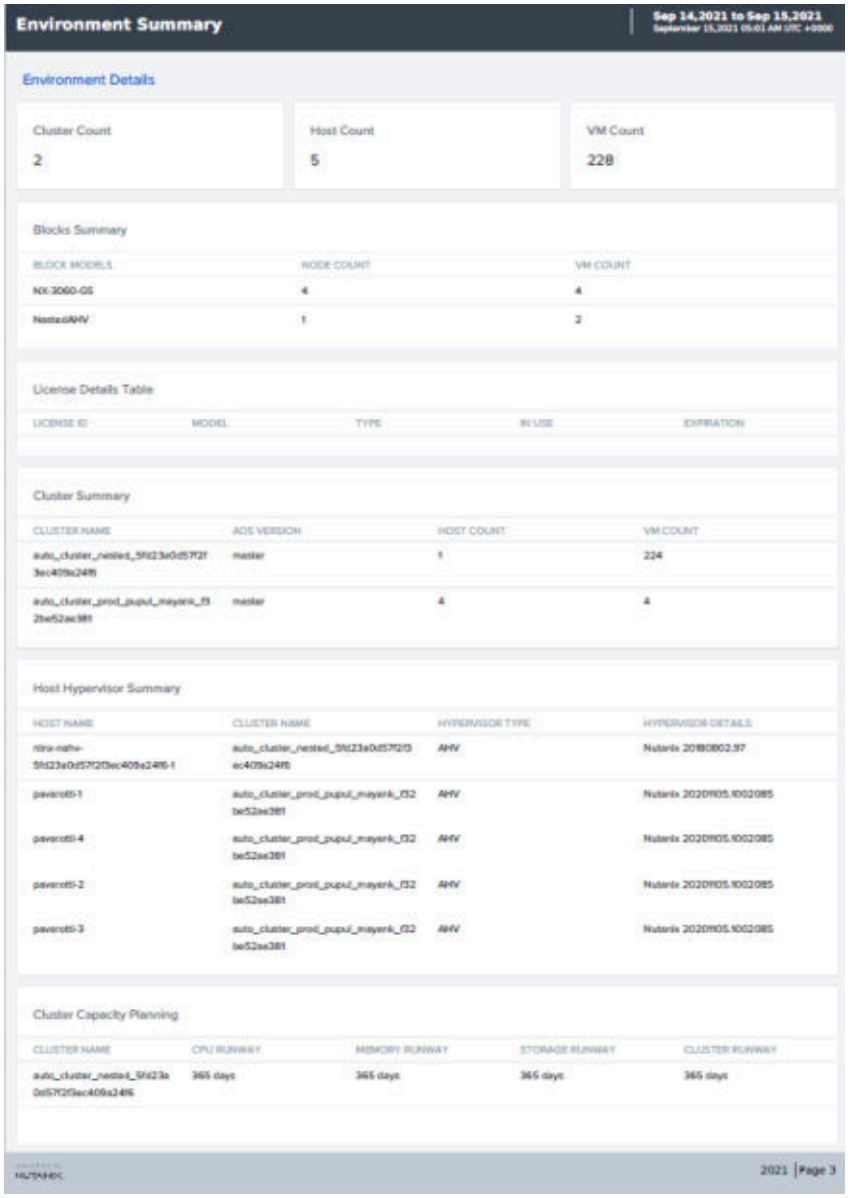

Figure 474: Detailed Metrics Information (Nutanix: Image 2)

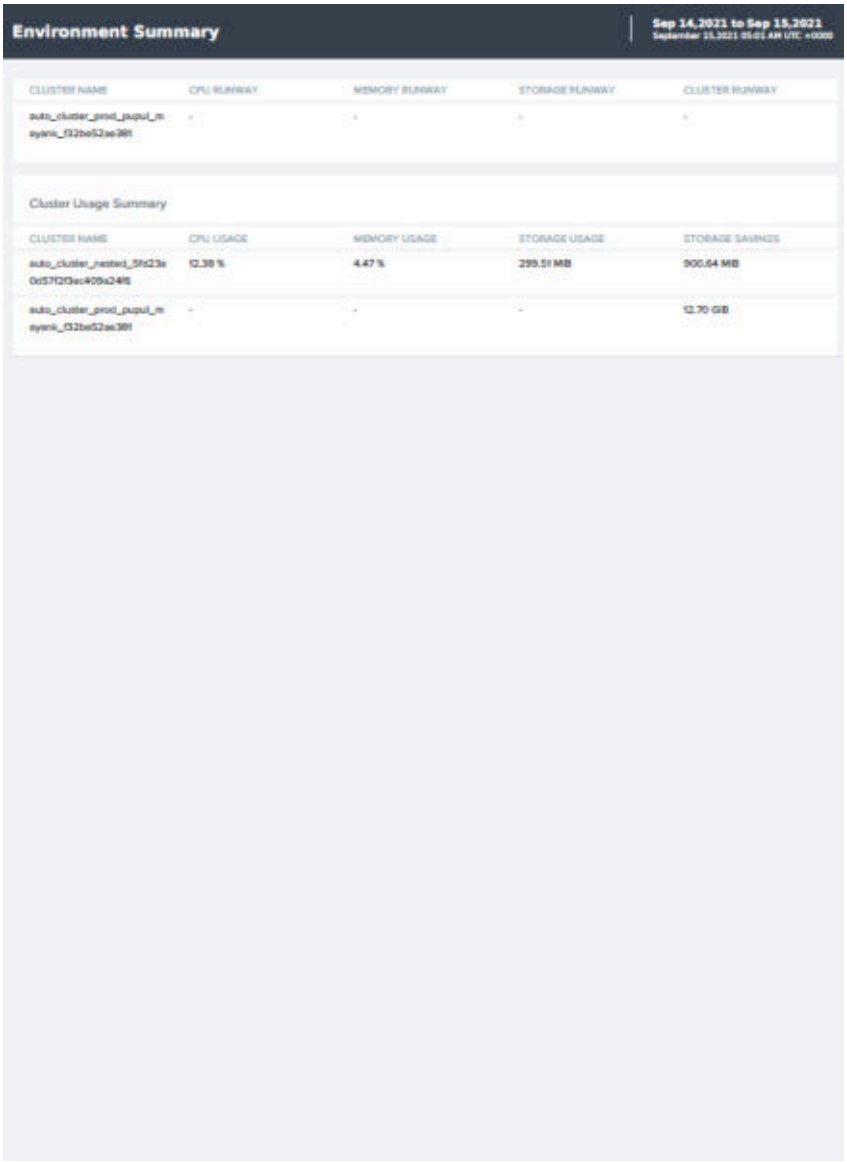

Figure 475: Detailed Metrics Information (Nutanix: Image 3)

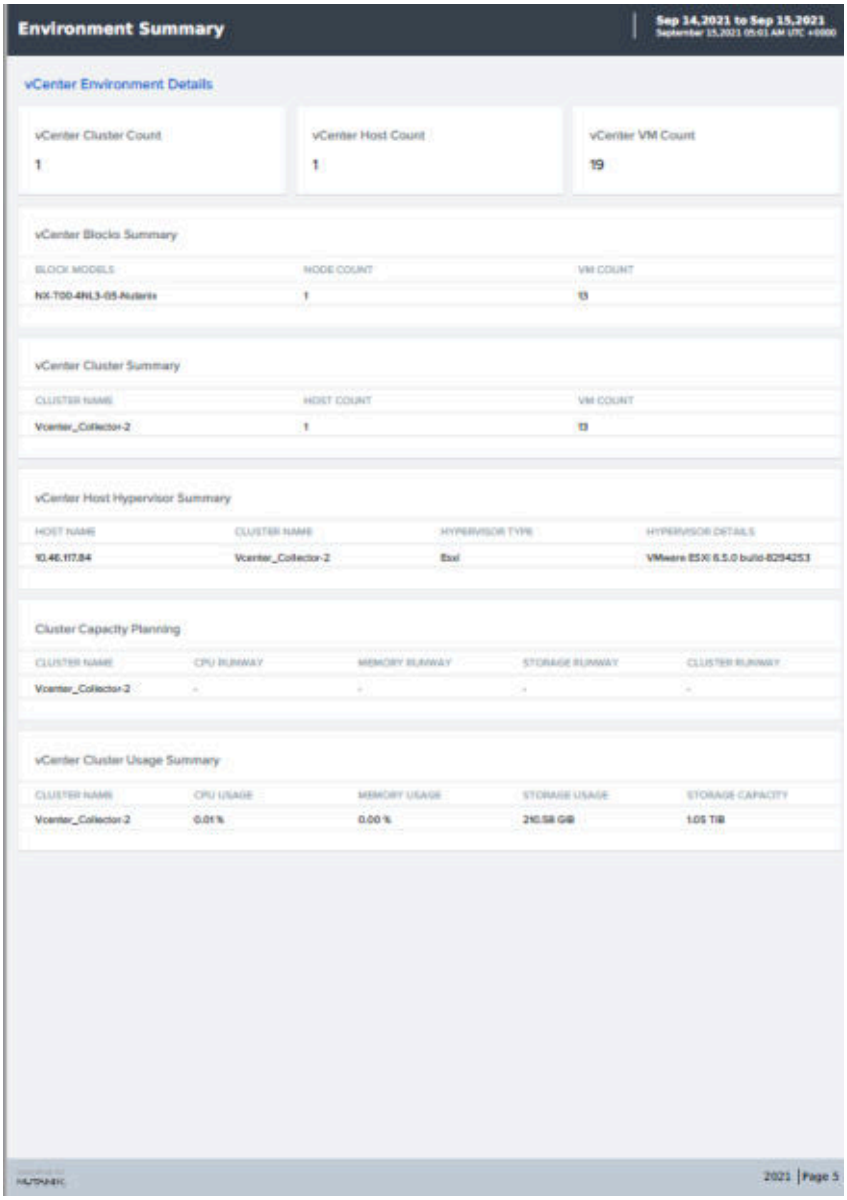

Figure 476: Detailed Metrics Information (vCenter: Image 1)

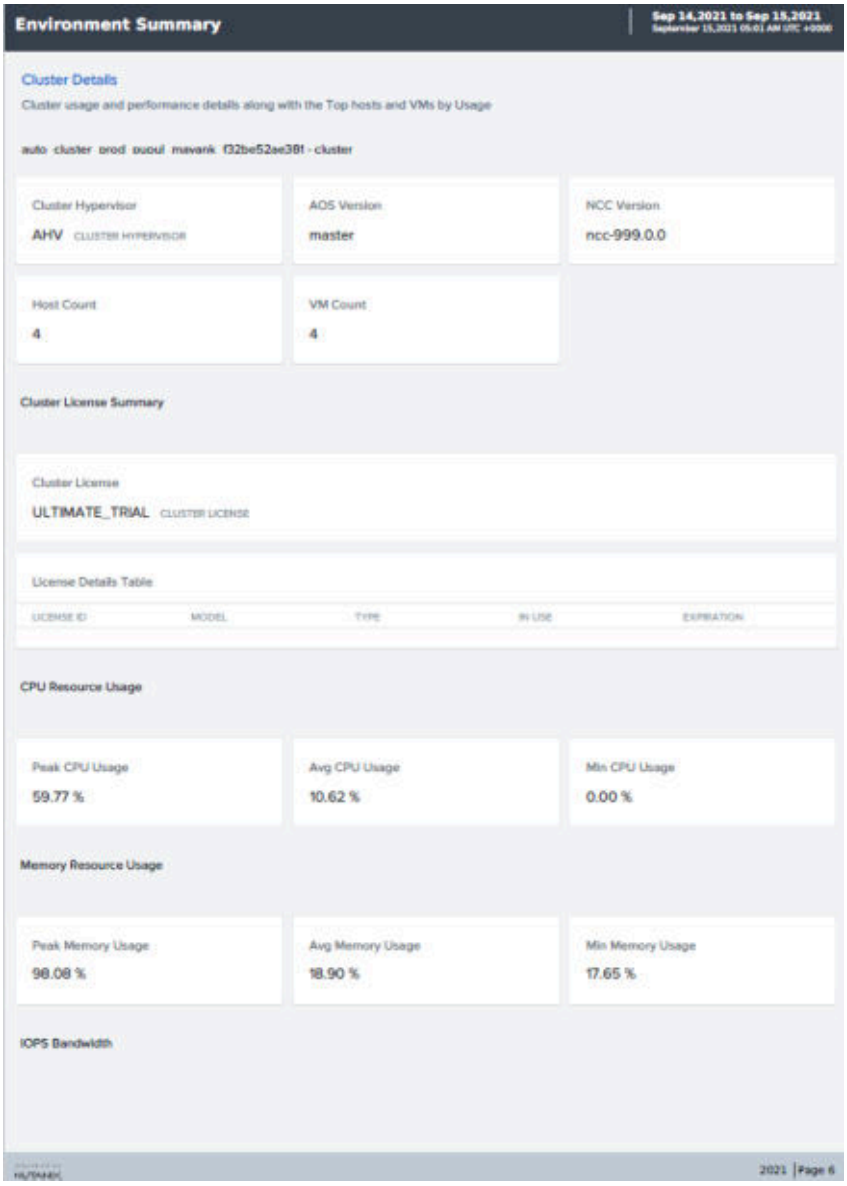

Figure 477: Detailed Metrics Information (vCenter: Image 2)

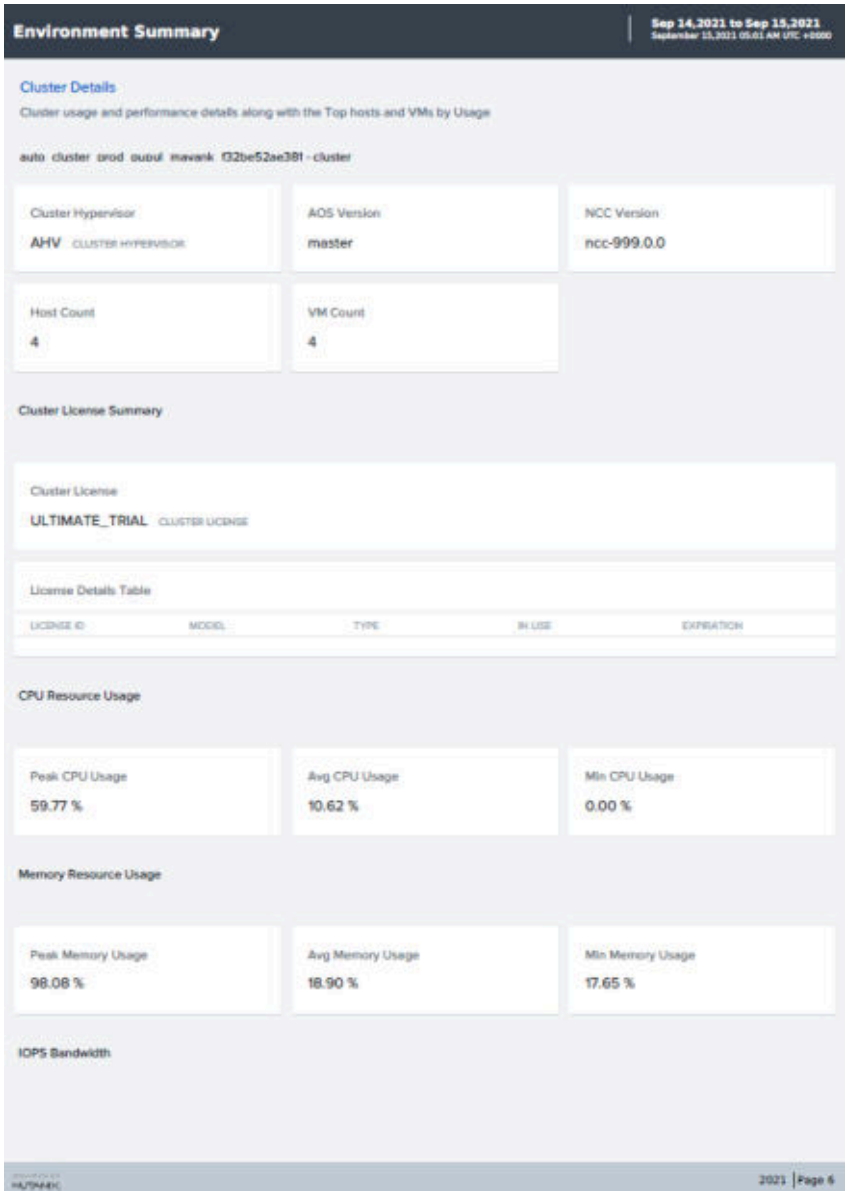

Figure 478: Detailed Metrics Information (vCenter: Image 3)

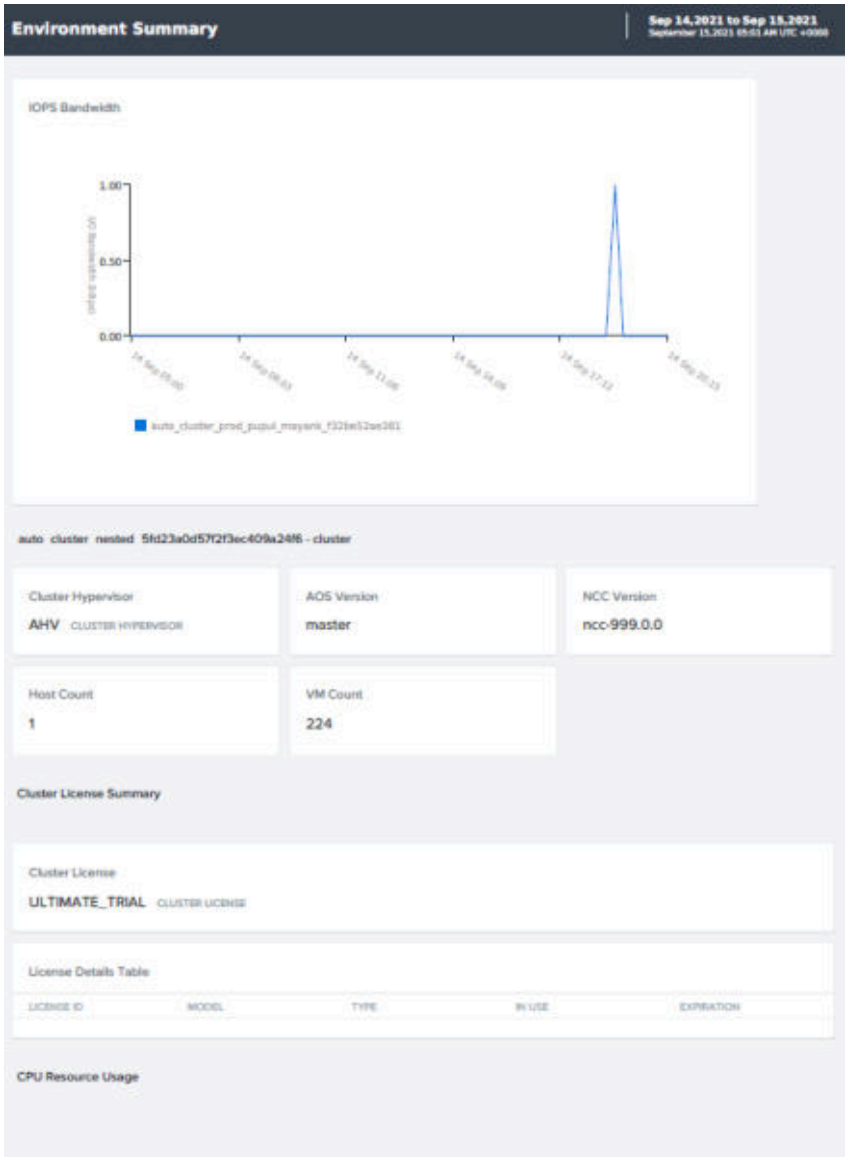

Figure 479: Detailed Metrics Information (vCenter: Image 4)

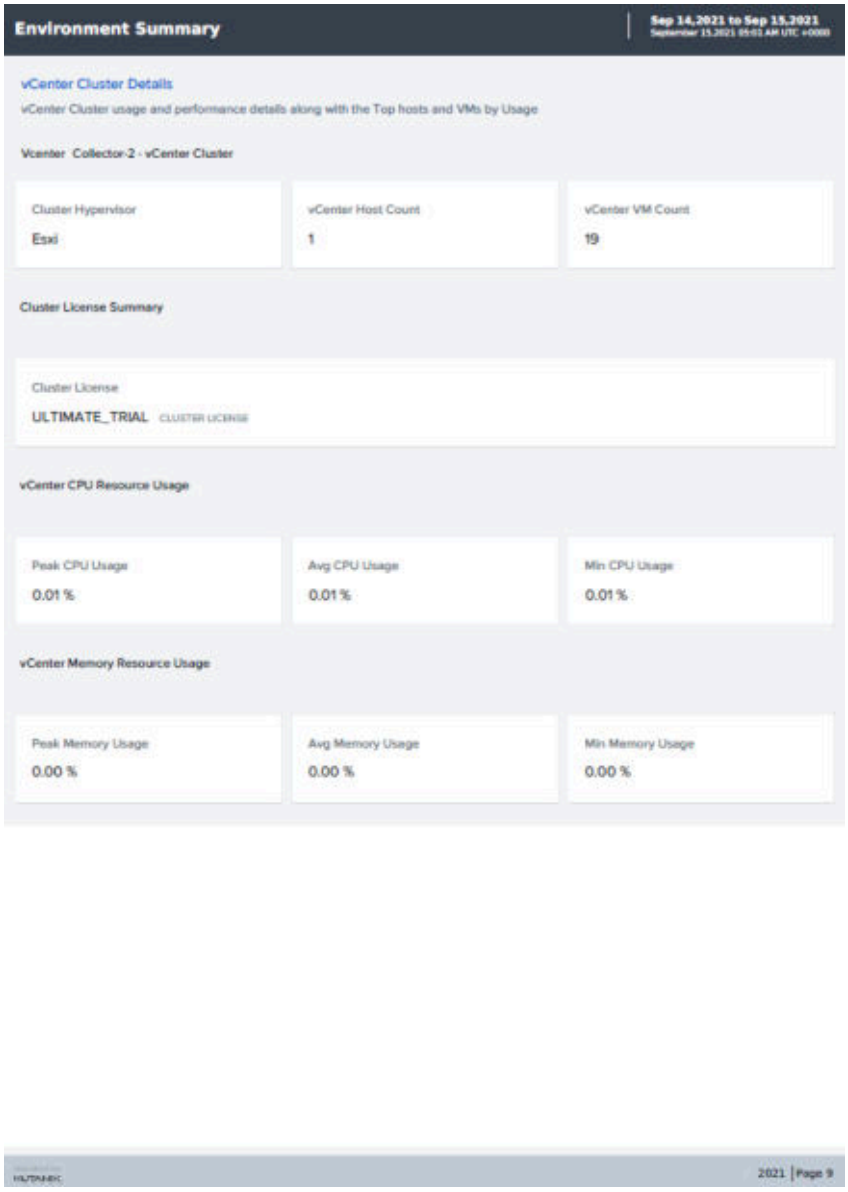

Figure 480: Detailed Metrics Information (vCenter: Image 5)

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 **Predefined Views** pane.

Figure 481: Custom View

Figure 482: Predefined View

Figure 483: All View (Custom and Predefined Combined)

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 (Nutanix) on page 901 and Adding a Bar Chart or a Line Chart or a Histogram View to a Report (vCenter).
- To add a data table view, see Adding a Data Table View to a Report (Nutanix) on page 904 and Adding a Data Table View to a Report (vCenter).
- To add a configuration summary view, see Adding Configuration Summary View on page 907.
- To add a metric summary view, see Adding Metric Summary View (Nutanix) on page 909 and Adding Metric Summary View (vCenter).
- To add an entity count view, see Adding Entity Count View (Nutanix) on page 911 and Adding Entity Count View (vCenter).
- To add a title and description view, see Adding Title and Description View on page 912.
- To add a group view, see Adding a Group View on page 913.
- To add a predefined view, see Adding a Predefined View on page 917.

Note:

- After you add the views to a report, the report displays the dummy data for the view that you have added. Only when you generate the report, the data for the view is updated and you can download the report to view its information.
- Collecting some VM data requires that NGT is installed. As a result reports include such data only from NGT-enabled VMs.

Adding a Bar Chart or a Line Chart or a Histogram View to a Report (Nutanix)

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.

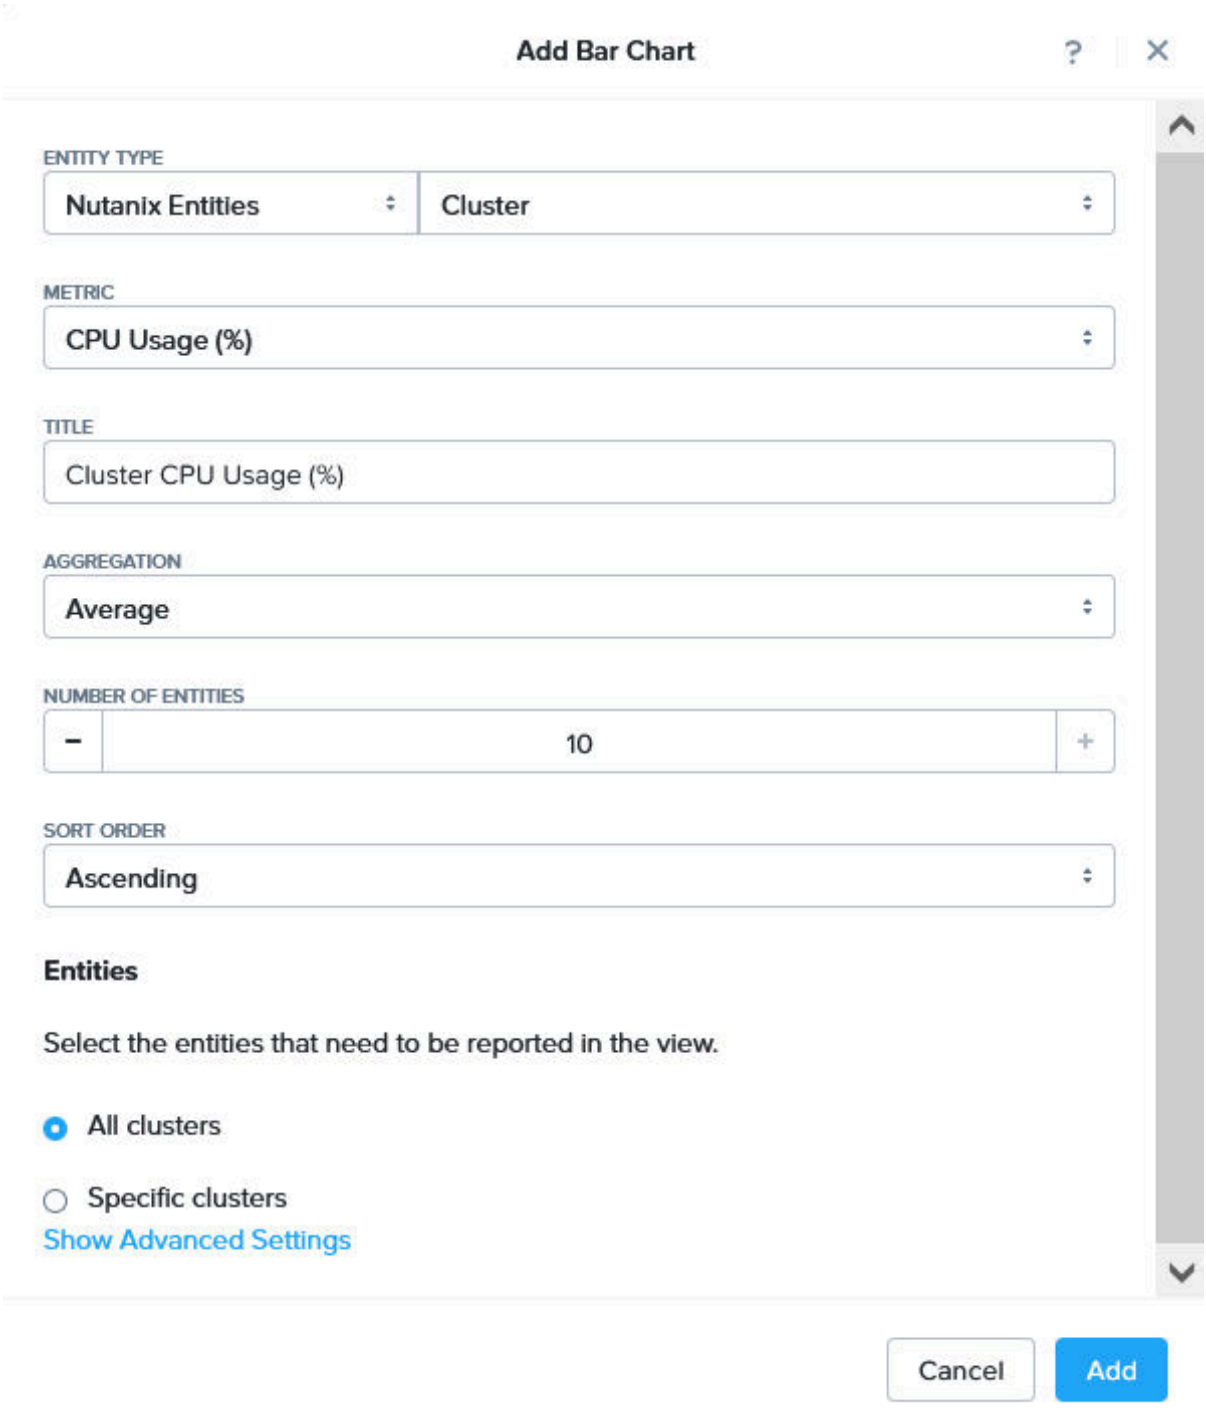

Figure 484: Bar Chart View

   a. Select **Nutanix Entities** and the entity type (**VM**, **Disk**, **Storage Container**, **Host**, **Cluster**, **Virtual Disk**, or **Virtual Group**) from the **Entity Type** drop-down menu for which you want to configure the view.

Note: If a vCenter is enabled, you can select **vCenter Entities** in the entity type. This option does not appear in the drop-down menu if a vCenter is not configured. To configure the view for vCenter entities, see Adding a Bar Chart or a Line Chart or a Histogram View to a Report (vCenter)

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, **Free Physical Storage**, **Controller Read I/O Bandwidth**, **Controller AVG Write I/O Bandwidth**, **Storage Savings Ratio**, **Controller I/O Bandwidth**, **Controller AVG Write I/O Latency**, **Controller IOPS**, **Controller Read IOPS**, **Controller Write I/O Bandwidth**, **Memory Usage (%)**, **Physical Usage**, **Memory Usage Total**, **Storage Logical Usage**, **Controller Write IOPS**, **Controller AVG Read I/O Latency**, **Storage Savings**

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 how you intend to represent the entities (**Ascending** or **Descending**) from the **Sort Order** drop-down menu.

g. (Only for Histogram view) Enter number of buckets in the **Number of Objects** field.

   You can define the frequency distribution for all the entities in the entity type for the selected metric by using this field. For example, for a selected metric if the minimum value is 10 and maximum value is 100 and you have entered number of buckets as 5 then you will have 18 buckets configured (100-10/5). Depending on the value of the metric, different metrics will fall in the appropriate buckets.

h. Select the entities that you want to add.

   - **All Entities**: Select this option if you want to include all the entities that you have selected in the view
   - **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

i. To configure advanced setting, click the **Show Advanced Settings** link. 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.

3. Click **Add**.
   The configured view gets added to the report.

Adding a Data Table View to a Report (Nutanix)

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.

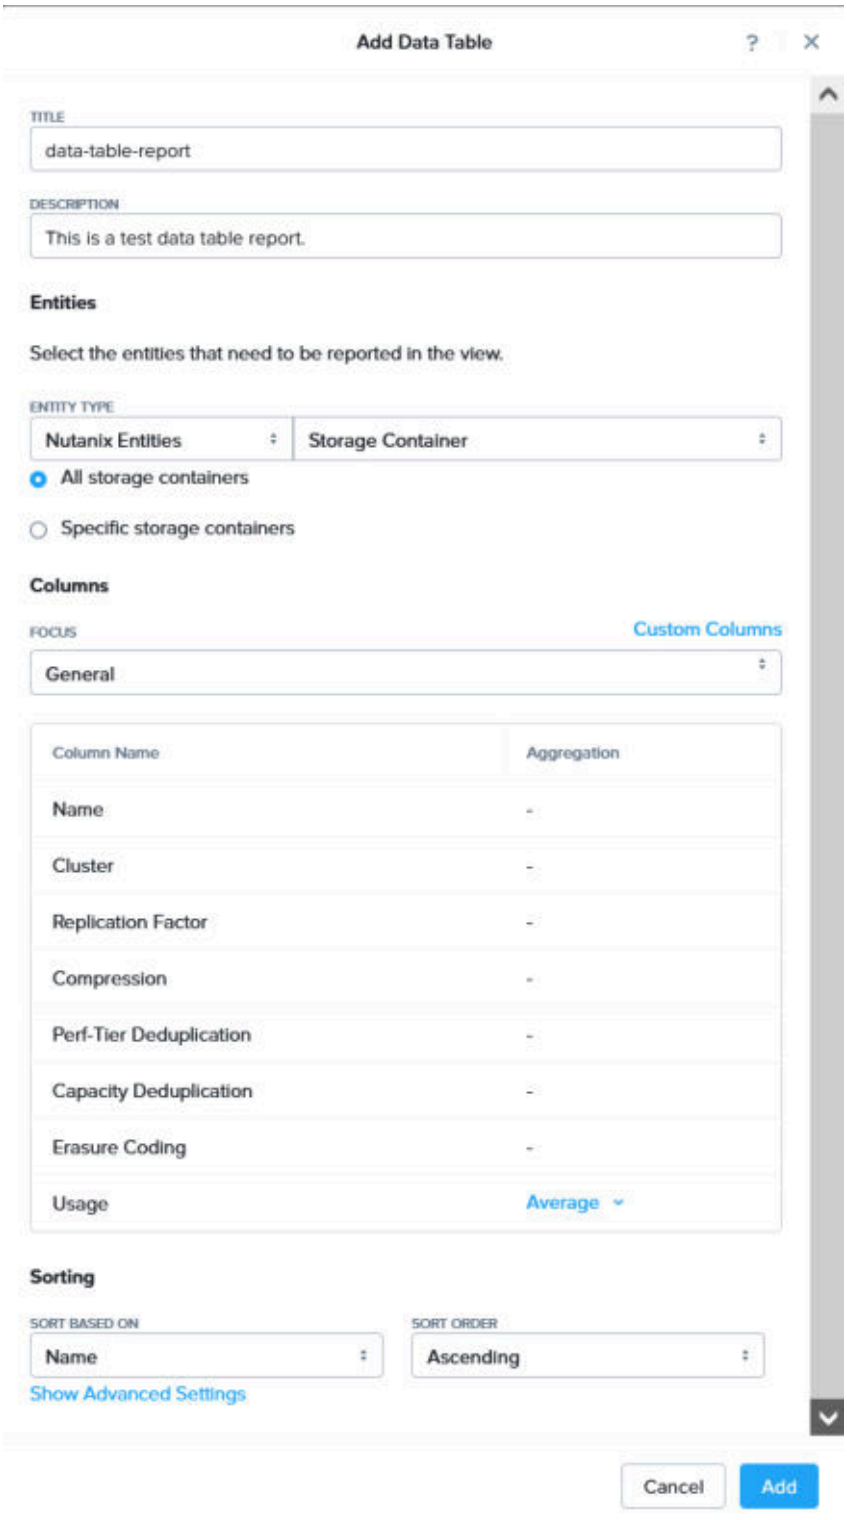

Figure 485: Data Table View

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 **Nutanix Entities** and 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.

   > Note: You can define sorting only on selected entities from the entity list.

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.

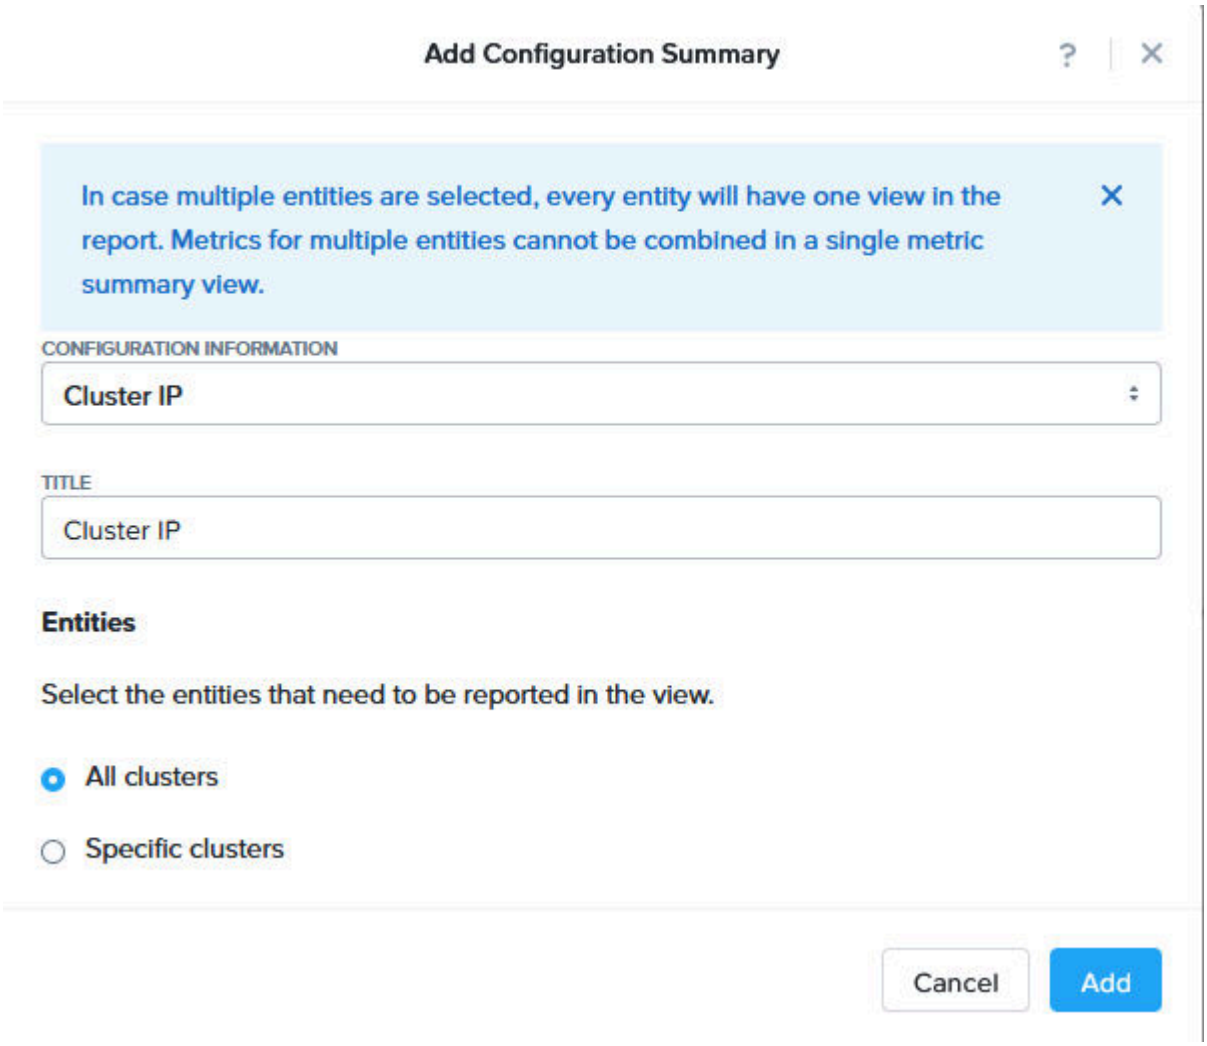

Figure 486: Configuration Summary View

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 (Nutanix)

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.

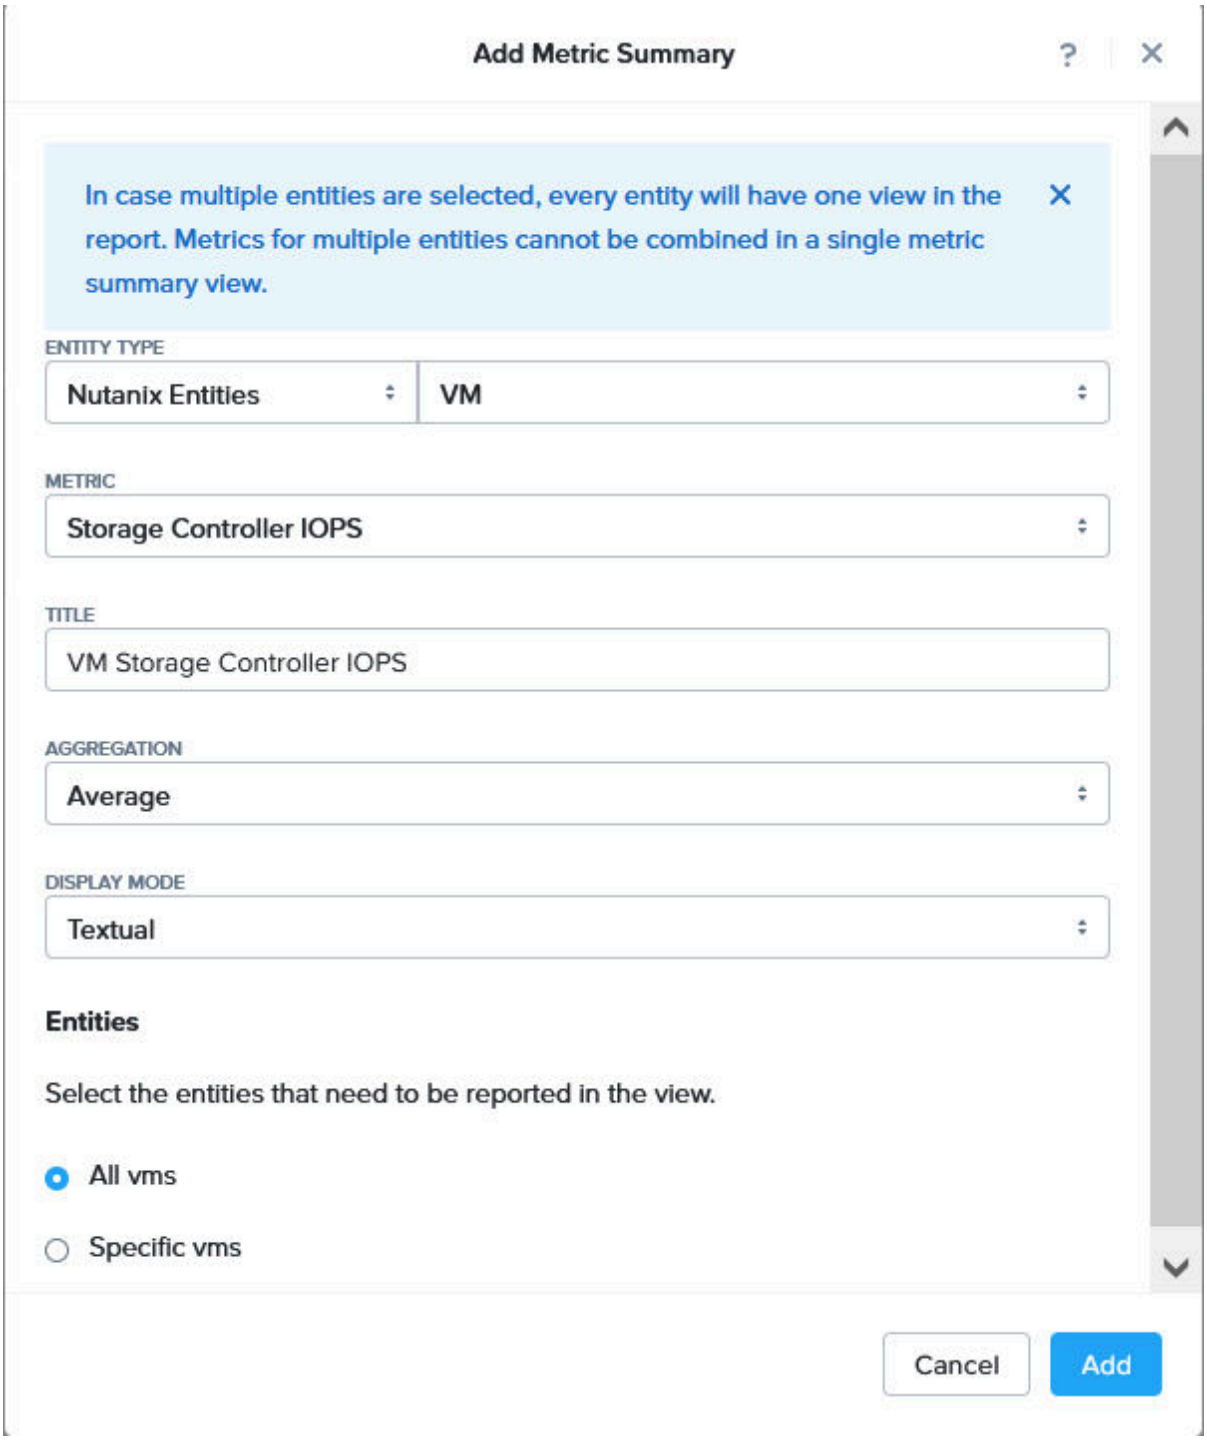

Figure 487: Metric Summary View

a. Select **Nutanix Entities** and 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 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 Entity Count View (Nutanix)

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.

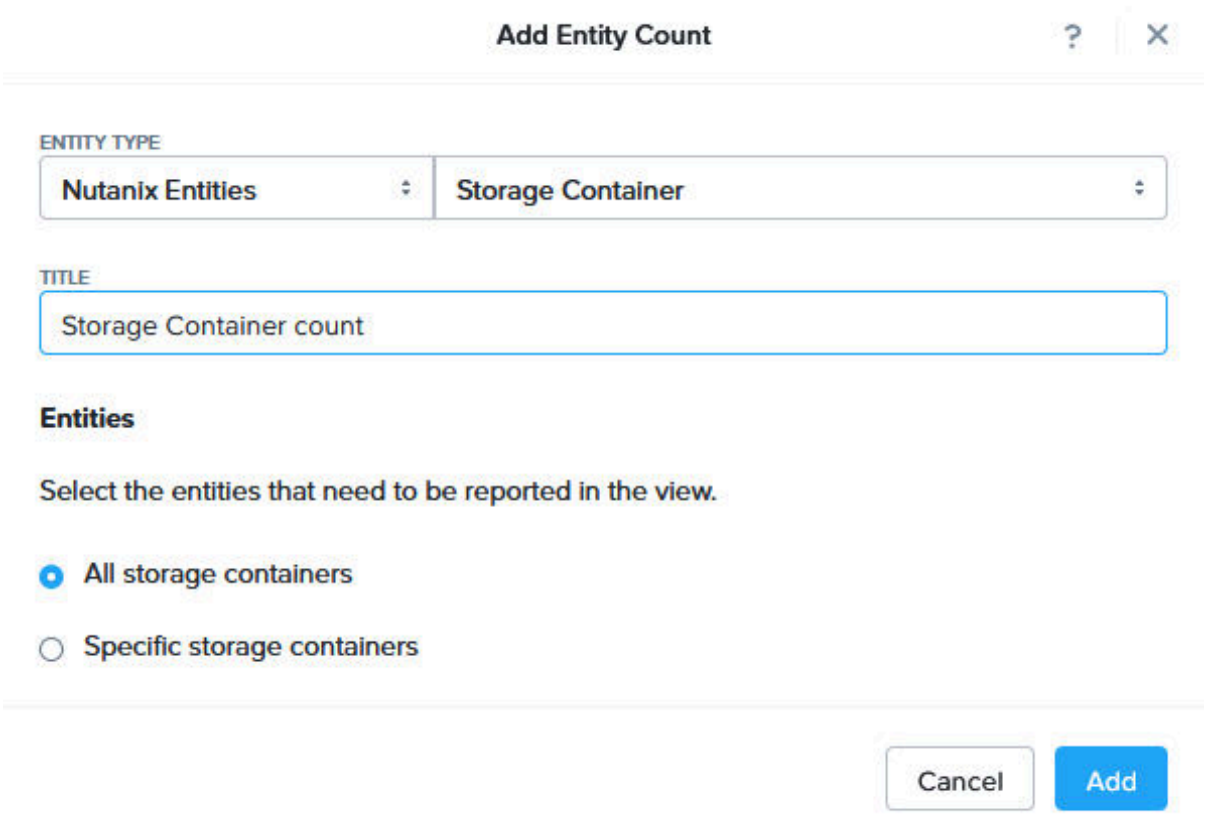

Figure 488: Entity Count View

a. Select **Nutanix Entities** and 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 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 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. You cannot group Nutanix and vCenter entities in a single view. For example, a Nutanix group for VM entity cannot be added to a vCenter group for VM entity.

> 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.

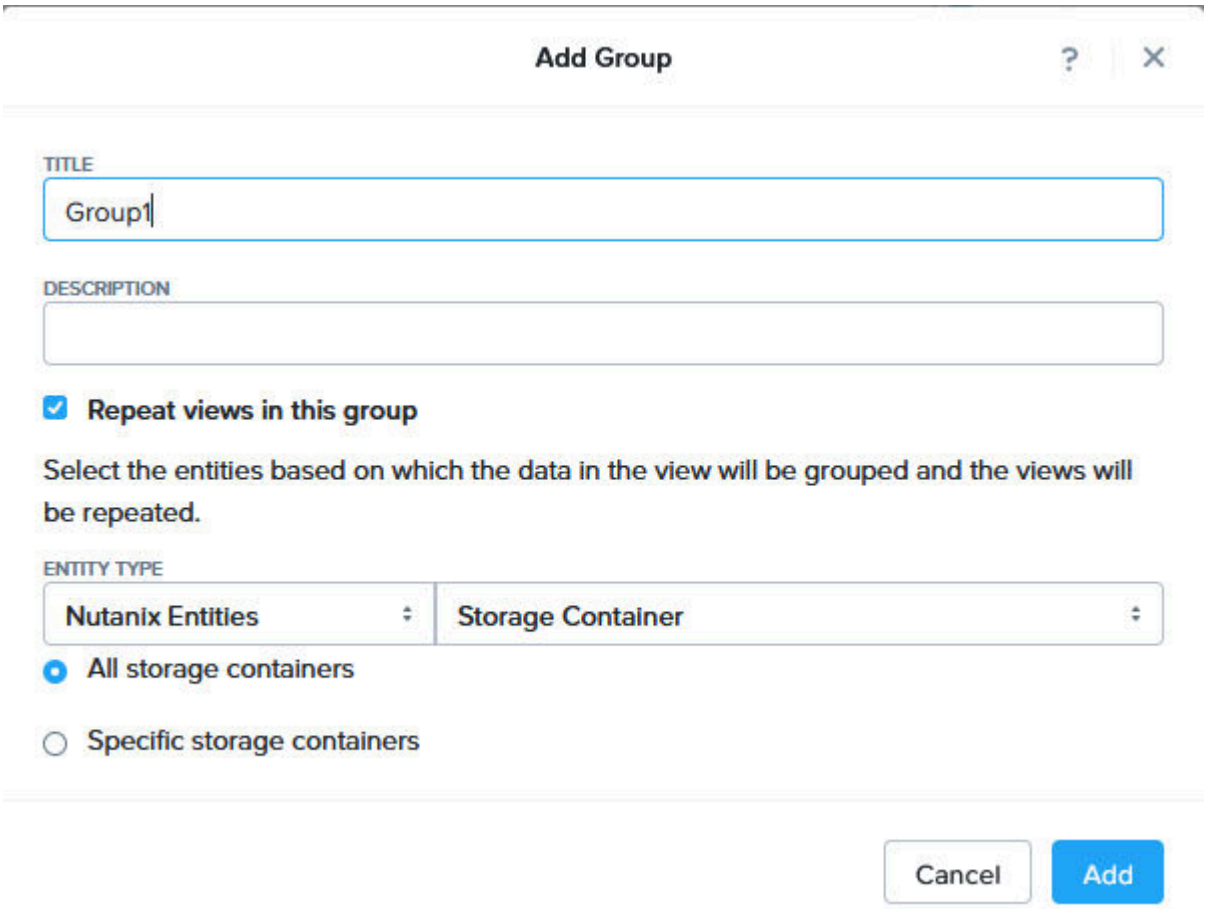

Figure 489: Group View (Nutanix)

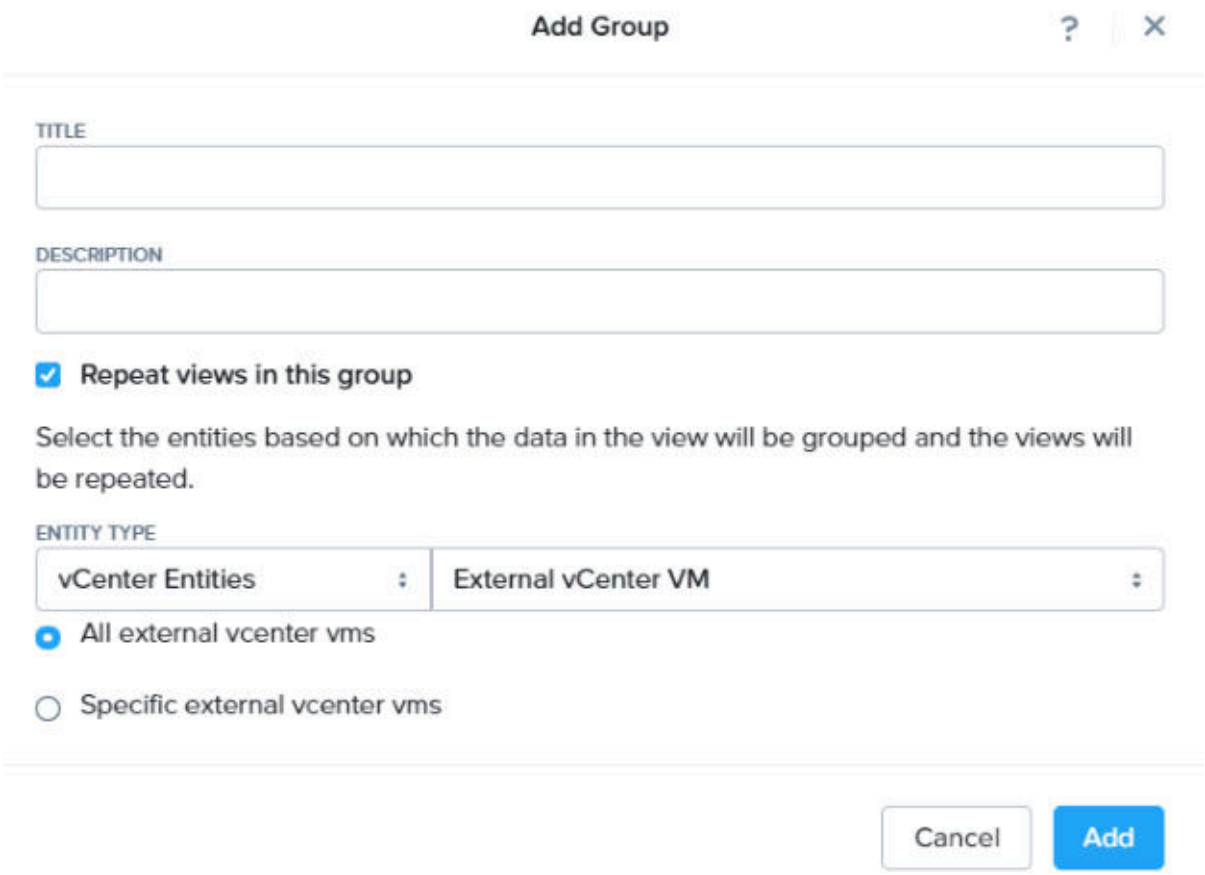

Figure 490: Group View (vCenter)

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, for Nutanix entities, 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. Similarly, for vCenter entities, if you select entity type as External vCenter VM and desire to limit the view for the 50 powered on VMs, you can select **Specific external vcenter 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 916

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 913

   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**.

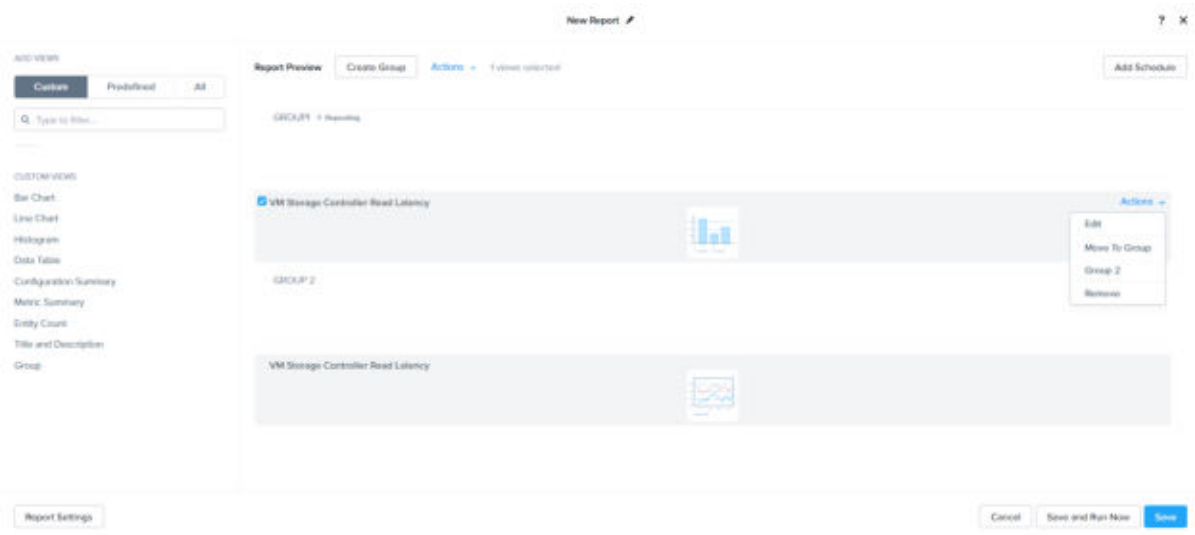

Figure 491: Adding Views to Group view

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 913.
   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 Predefined 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 **Predefined** views, you can add the following views.

> Note:
>
> - You can only include the **License Details Table** view as part of the group view.
> - You can only view DR specific widgets when a DR report is edited.

Procedure

Select the view that you want to add from the **Predefined 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.
- **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.
- **Blocks Summary:** Provides summarized information about the different block models, number of nodes, and number of VMs running in the Prism Central.
- **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.
- **Inactive VMs List:** 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.
- **Constrained VMs List:** 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.
- **Overprovisioned VMs List:** 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.
- **Bully VMs List:** 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 Reclaim:** Provides summarized information about the potential CPU resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Memory Reclaim:** Provides summarized information about the potential memory resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Storage Reclaim:** Provides summarized information about the potential storage resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **vCenter Blocks Summary:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter CPU Runway:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Memory Runway:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Storage Runway:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Inactive VMs List:** Provides detailed information about the inactive VMs on the vCenter 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.
- **vCenter Constrained VMs List:** Provides detailed information about the constrained VMs on the vCenter cluster (that require more resources to function properly); including the VM name, efficiency status and the reasons for being marked as constrained.
- **vCenter Overprovisioned VMs List:** Provides detailed information about the over-provisioned vCenter registered VMs on the cluster (that have more resources than they deserve); including the VM name, efficiency status, reasons for being marked as over-provisioned and potential resource gains that can be accomplished by deleting the inactive VMs or resizing the over-provisioned VMs.
- **vCenter Bully VMs List:** Provides detailed information about the vCenter registered VMs whose activities are affecting other VMs that is not directly related to their sizes; including VM name, efficiency status and the reasons for being marked as bully.

- **vCenter Overprovisioned VMs List:** Provides detailed information about the over-provisioned vCenter registered VMs on the cluster (that have more resources than they deserve); including the VM name, efficiency status, reasons for being marked as over-provisioned and potential resource gains that can be accomplished by deleting the inactive VMs or resizing the over-provisioned VMs.
- **vCenter Inactive VMs List:** Provides summarized information about the potential storage resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **vCenter Constrained VMs List:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Bully VMs List:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Potential CPU Reclaim:** Provides summarized information about the potential storage resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **vCenter Potential Memory Reclaim:** Provides summarized information about the potential storage resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **vCenter Potential Storage Reclaim:** Provides summarized information about the potential storage resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **Configuration Alerts:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **DR Alerts Summary Table:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Recovery Plan Events:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Primary:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Recovery:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Entity Protection Summary:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **List of Entities with RPO Not Met:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **List of Entities Meeting RPO:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Recovery Plan Events Details:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.

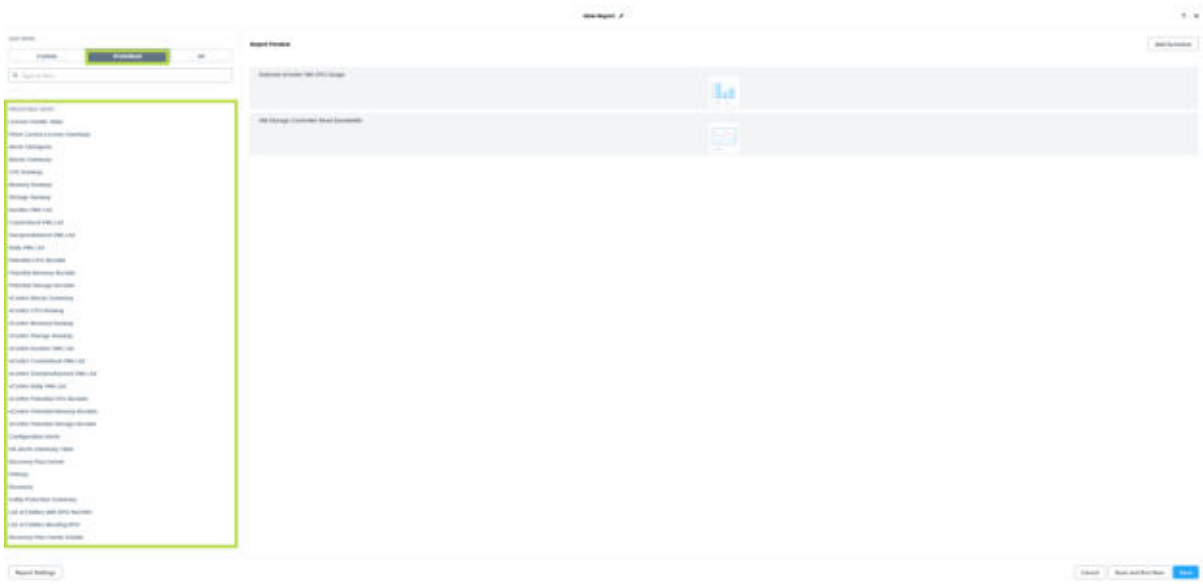

Figure 492: Pre-Defined Views

Scheduling a Report

After you add all the views to the reports, you have an option to run the report at a custom generated time. This scheduling is applicable only for the specific report for which you are configuring the schedule.

Procedure

1. In the **New Report** wizard, click **Add Schedule**.

2. Select when you want to run the report from the **Report Runs** drop-down menu.
   You an run the report on daily, weekly, monthly, or yearly basis.

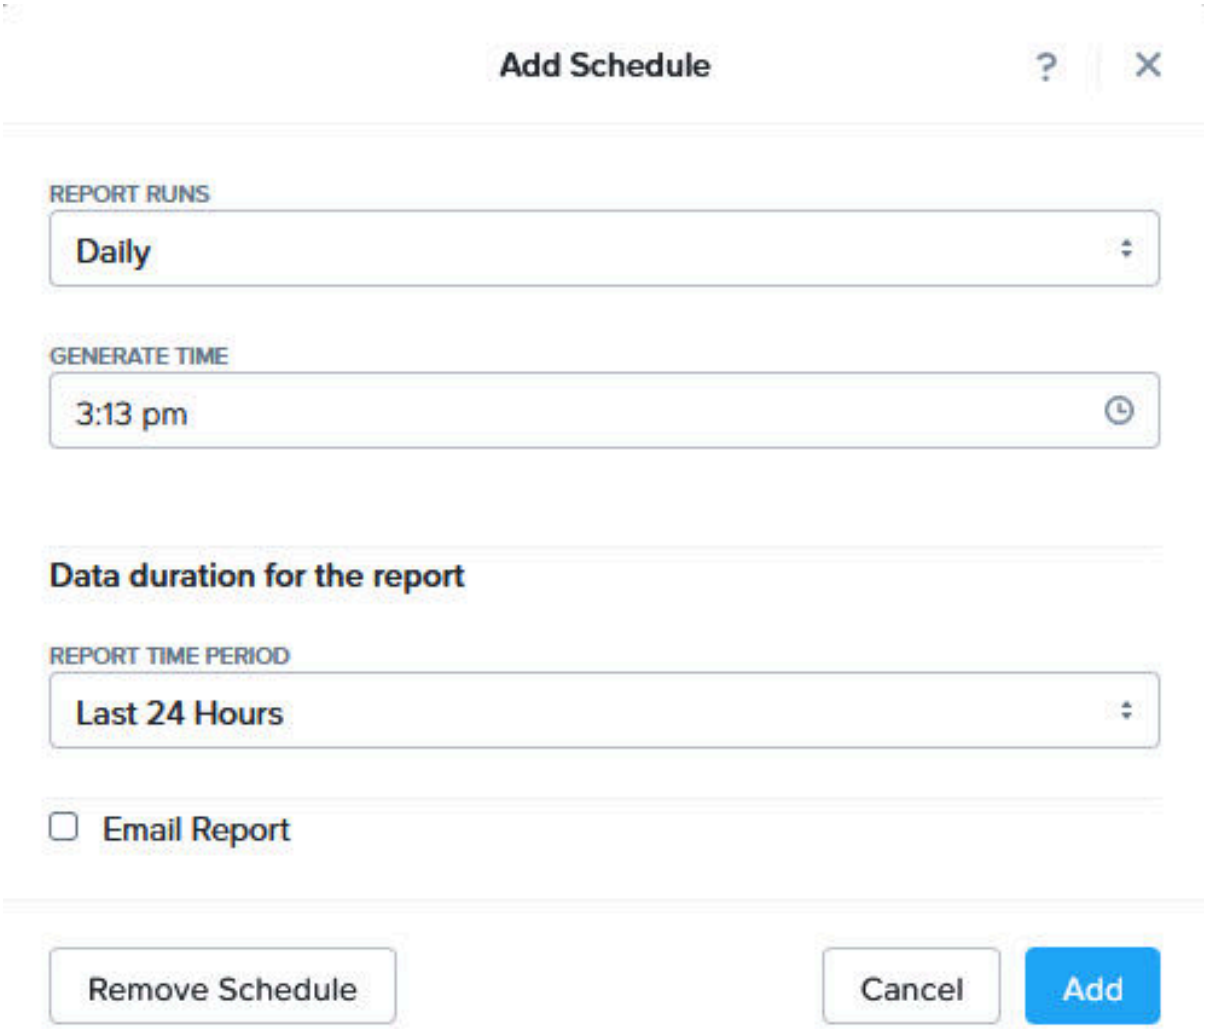

Figure 493: Report Scheduling

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 245), 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 884.

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 245).
   - » To configure the report settings at the report level, in the **New Report** wizard, click **Report Settings**.

2. Configure the appearance of the report.

   The appearance of the report is divided into two parts; cover page settings and content page settings. Configuration of logo and the background color is included in the cover page settings. Configuration of color of the header is included in the content page settings.

   Note: The copyright information will not be displayed in the footer section, therefore the ability to customize copyright information is not allowed.

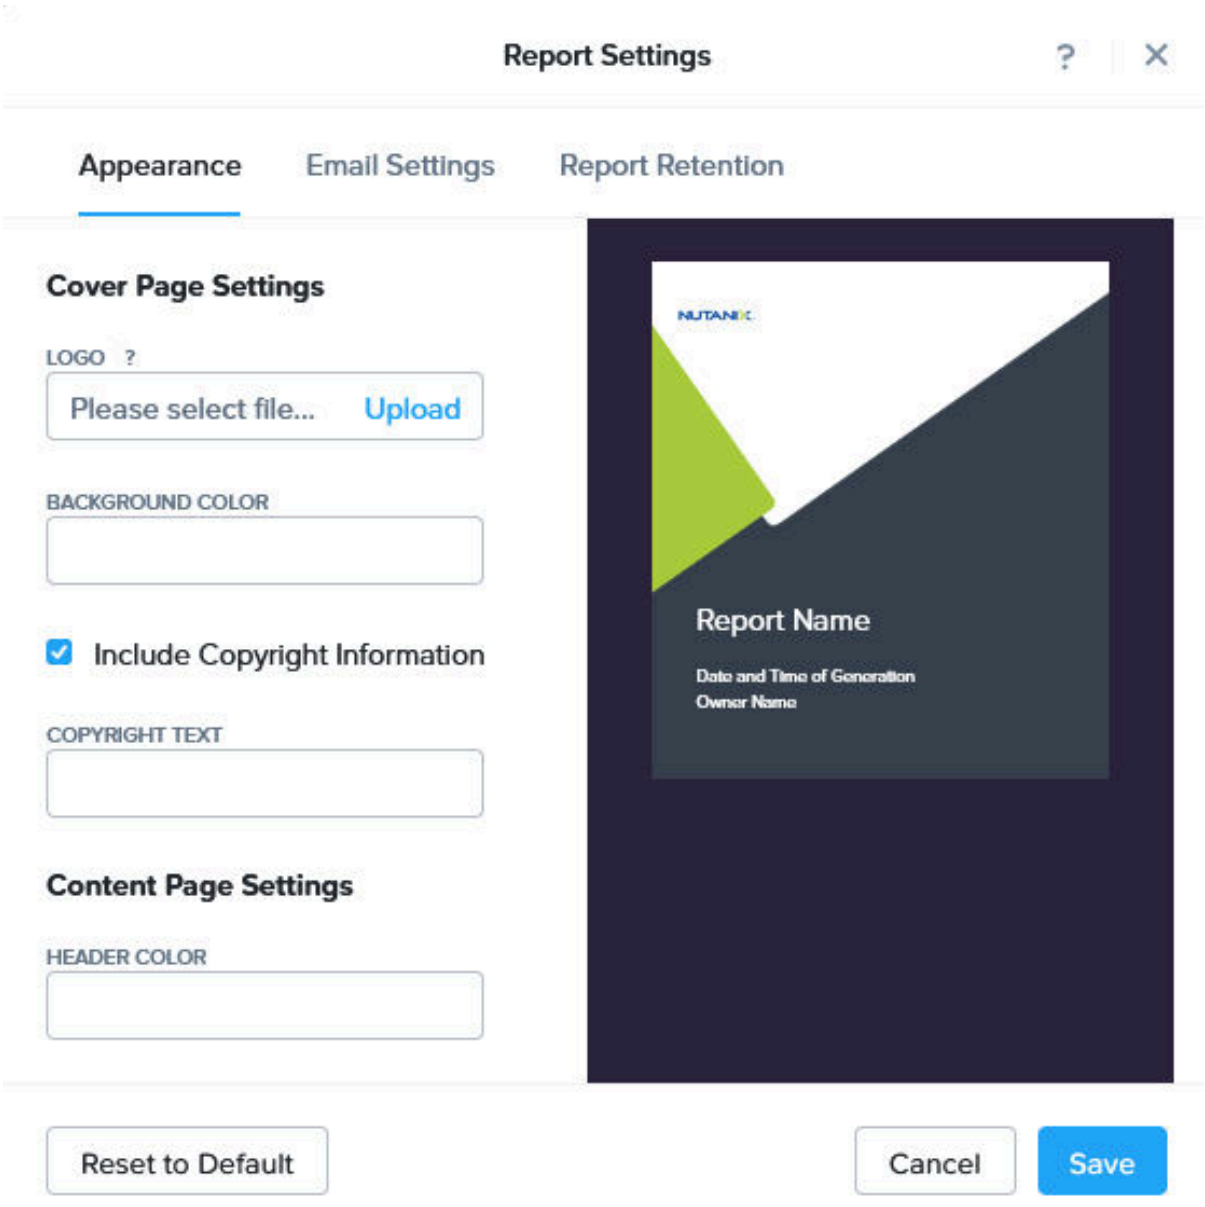

Figure 494: Report Settings Configuration

   a. To upload the logo for the report, click the **Upload** button.

      Note: The logo size cannot be more than 1 MB.

b. To change the background color of the report, select the background color from the **Background Color** field.

c. To change the header color of the report, select the header color from the **Header Color** field.

3. Configure the Email settings.

a. Enter the subject of the Email in the **Prepend Content to the Email Subject** text box.

You can view the preview of the subject that you are adding in the Preview pane and is prepended to the subject *`subject_email`* Prism Report *`report_name`*.

b. Enter the body of the Email in the **Append Content to the Email Body** text box.

You can view the preview of the body that you are adding in the Preview pane and is appended to the body `Please find attached the generated report` *`email_body`*.

c. Enter the recipients (comma-separated) of the report in the **Email Recipients** text box.

The recipients that you have configured will get the report with the subject and body of the Email that you have entered.

4. Configure the retention policy of the report.

a. Click the **Define Report Retention** check box.

You can define the retention policy by entering the number of instances of the report that you want to retain or the duration of time that you want to retain the report.

b. To configure the retention policy in instances form, click the **Number of Instances** option and enter the desired number of instances.

> Note: You cannot retain more than 25 instances of a report. If you have 25 instances of a report and you generate a new instance of the same report, the first instance that was generated gets deleted.

c. To configure the retention policy in duration form, click the **Time Duration** option and enter the desired number and duration in terms of **Days**, **Weeks**, or **Months** from the drop-down menu.

For the time-based retention, you can retain a report for maximum of 3 months.

> Note: If you do not define any retention policy, by default 10 instances of a report are retained.

5. Configure the Report Format of the report.

a. Click the **PDF** check box or **CSV** check box or both to set a report format.

A PDF generated report displays a graphical representation of infrastructure resources. A CSV generated report contains information separated by commas. You can use this CSV file to load CSV data into other systems.

b. If you select the PDF format and CSV format in **Report Format**, you can select the **Recipient Format**. Select the format of report in which you want to send an e-mail.

> Note: You cannot set the report format at the global level. You can only set the report format at the report level while creating a new report or editing an existing report.

6. Click **Save**.
The report configuration is saved and next time you generate the report, the report generation follows the configured setting. If you apply the report settings both the global

level (for all the reports) or for report level (when you are creating a new report), the setting applied at the report level takes precedence.

7. (Optional) Click **Reset to Default** to discard all the changes that you have made in the Report Settings page.

Role Based Access Control

Role-Based Access Control (RBAC) restricts Reports Management access to authorized users. For any role you are assigned, privileges are given to you to modify certain configurations in Reports Management. The following table shows the roles and certain privileges given to that role.

Table 742: 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				#		#	#	#

Access/ Permissior	Consumer	Developer	Operator	Prism Admin	Prism Viewer	Project Admin	Self-Service Admin	Super Admin
View Common Report Config	#	#	#	#	#	#	#	#
View Report Config	#	#	#	#		#	#	#
View Report Instance	#	#	#	#	#	#	#	#

13

SERVICES ENABLEMENT

You can enable selected services through Prism Central.

- Enabling Calm on page 928
- Enabling Files on page 929
- Enabling Foundation Central on page 929
- Enabling Karbon on page 929
- Enabling Objects on page 930

Enabling Calm

About this task

You can select, provision, and manage your business applications across all your infrastructure for both private and public clouds through the Nutanix Calm feature. Nutanix Calm provides automated application life cycle management, custom blueprints for the setup and management of enterprise applications, a marketplace to publish the blueprints to end users, and automated hybrid cloud management to provision your hybrid cloud architecture. See the *Nutanix Calm Administration and Operations Guide* for information about configuring and using Calm.

To enable Calm, do the following:

Procedure

1. Do one of the following:
   - » Click the collapse menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Calm** from the entities menu (see Entities Menu on page 13).
   - » Select **Enable App Management** from the Settings menu (see Settings Menu on page 16).
2. In the **Enable App Management** page, do the following:
   a. Check the **Enable App Management** box to enable application management.
   b. Check the **Enable Nutanix seeded blueprints** box to load your application store with a set of pre-defined blueprints.

      While loading the pre-defined blueprints is optional, it is recommended as it provides blueprints for many common use cases.
   c. Click the **Save** button.

Enabling Files

About this task

Nutanix Files provides file services to clients for file sharing across user work stations from a centralized and protected location.

To enable Files, do the following:

Procedure

1. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Files** from the entities menu (see Entities Menu on page 13).
2. In the Nutanix Files page, click the **Enable Files** button.

   After Files is enabled, the Files dashboard appears. See the *Nutanix Files Guide* for information about configuring and using Files.

Enabling Foundation Central

About this task

Foundation Central can manage several Foundation instances from a single pane of glass, allowing you to create clusters of remote nodes without needing to configure each of them individually.

To enable Foundation Central, do the following:

Procedure

1. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Foundation Central** from the entities menu (see Entities Menu on page 13).
2. In the Nutanix Foundation Central page, click the **Enable Foundation Central** button.

   After Foundation Central is enabled, the Foundation Central dashboard appears. See the *Foundation Central Guide* for information about configuring and using Foundation Central.

Enabling Karbon

About this task

Nutanix Karbon is a curated turnkey offering that provides simplified provisioning and operations of Kubernetes clusters. Kubernetes is an open-source container orchestration system for deploying and managing container-based applications.

To enable Karbon, do the following:

Procedure

1. Click the collapse (hamburger) menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Karbon** from the entities menu (see Entities Menu on page 13).
2. In the Karbon page, click the **Enable Karbon** button.

   After Karbon is enabled, you can click the **here** link to go to the Karbon console. See the *Nutanix Karbon Guide* for information about configuring and using Karbon.

Enabling Objects

Nutanix Objects is a feature that allows you to create and manage object stores in a cluster.

About this task

To enable Objects, do the following:

Procedure

1. Click the collapse menu button in the main menu (see Main Menu on page 10) and then select **Services** > **Objects** from the entities menu (see Entities Menu on page 13).
2. In the Objects page, click the **Enable** button to enable Objects.

   After Objects is enabled, the Object Stores dashboard appears. See the *Nutanix Objects Guide* for information about configuring and using object stores.

14

CUSTOMER SUPPORT SERVICES

Nutanix provides customer support services in several ways.

- Nutanix customer support can monitor your clusters and provide assistance when problems occur through the Pulse mechanism (see Pulse Health Monitoring on page 931 and Configuring Pulse on page 933).
- 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 943).
- If you need help, you can create a support ticket directly from Prism Central (see Creating a Support Case on page 936).

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 515). For the complete list of required ports, see Port Reference.

> Note: When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse.
>
> - To enable Pulse, click the **Continue** button in the message and follow the prompts.
> - To continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

Pulse Transport Methods

Configure one of the following Pulse transport methods (in order of preference):

- [first choice] Enable Pulse and use Prism Central as a proxy for the Pulse data transmitted by each node (for clusters registered with Prism Central). Advantages: The configuration is

automatic (as described in Prism Central Proxy for Pulse Data on page 936), 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 508). 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 931 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 515). For the complete list of required ports, see Port Reference.
- When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse. To enable Pulse, click the **Continue** button in the message and follow the prompts; to continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

You can enable (or disable) Pulse at any time. To configure Pulse, do the following:

Procedure

1. Go to the Settings menu (see Settings Menu on page 16) and select **Pulse**.

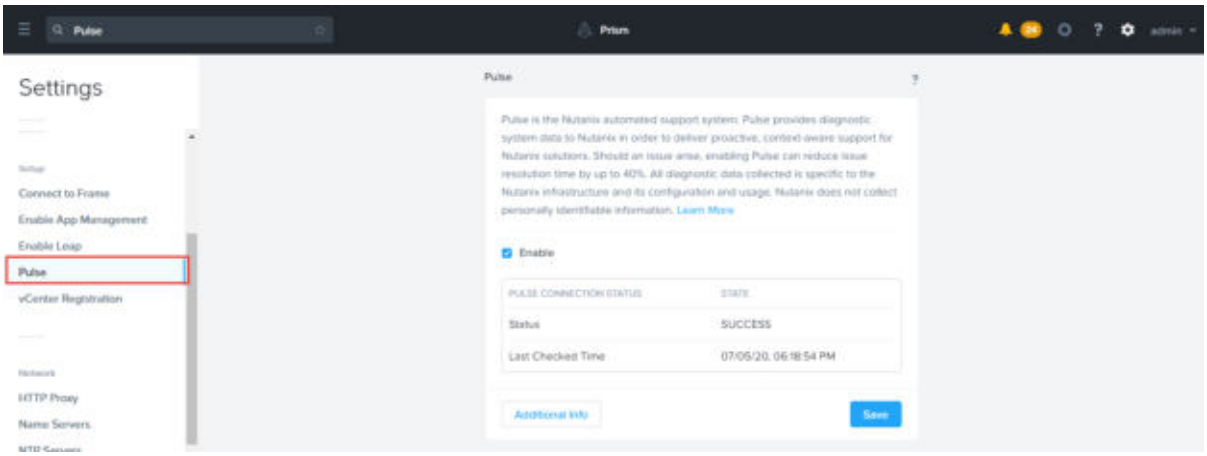

Figure 495: Pulse Window

2. To enable (disable) this feature, check (uncheck) the **Enable** box. Nutanix recommends enabling Pulse to allow Nutanix Support to receive system data and deliver proactive and context-aware support.

   When you enable Pulse, Pulse periodically sends data for troubleshooting purposes to Nutanix Insights tool and support team. Pulse is the underlying technology that securely transmits system-level diagnostic data to the Insights platform, enabling predictive health and context-aware support automation workflows. Nutanix Insights is an integrated service

that utilizes this data to augment product support, reducing customer case volume and expediting issue resolution time.

Nutanix does not share any data that Pulse sends, with any third parties unless permitted by your agreement with Nutanix or by the Nutanix Privacy Statement. Certain Nutanix products require Pulse enablement for functionality and/or features. See the Nutanix Privacy Statement and applicable product documentation for more details.

You may elect to obfuscate (mask) entity names and IP addresses which are not obfuscated or masked by default. Entity names and IP addresses are not linked to individuals. These relate to Nutanix entities such as cluster names. (See KB 9326 for more information.)

- To check if partial scrubbing is enabled, use the following command:

  ```
  nutanix@cvm$ curl -H "Content-Type: application/json" -X GET -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
  ```

  If the output contains the string `"identificationInfoScrubbingLevel":"PARTIAL"` or `"identificationInfoScrubbingLevel":"AUTO"`, that means partial scrubbing is enabled. If the output contains the string `"identificationInfoScrubbingLevel":"ALL"`, that means `identificationInfoScrubbingLevel` is still set to `ALL` and the **Open In Prism** button on the Insights Portal remains deactivated.

- To obfuscate (mask) entity names and IP addresses, and update the PII scrub level for specific Prism Element clusters or Prism Central cluster, use the following command:

  ```
  nutanix@cvm$ curl -H "Content-Type: application/json" -X PUT -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse
  ```

- To obfuscate (mask) entity names and IP addresses, and update the PII scrub level for all registered Prism Element clusters from Prism Central, use the following command on a Prism Central cluster:

  ```
  nutanix@cvm$ curl -H "Content-Type: application/json" -X PUT -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
  ```

3. Click the **Additional info** button to display additional email recipient fields.

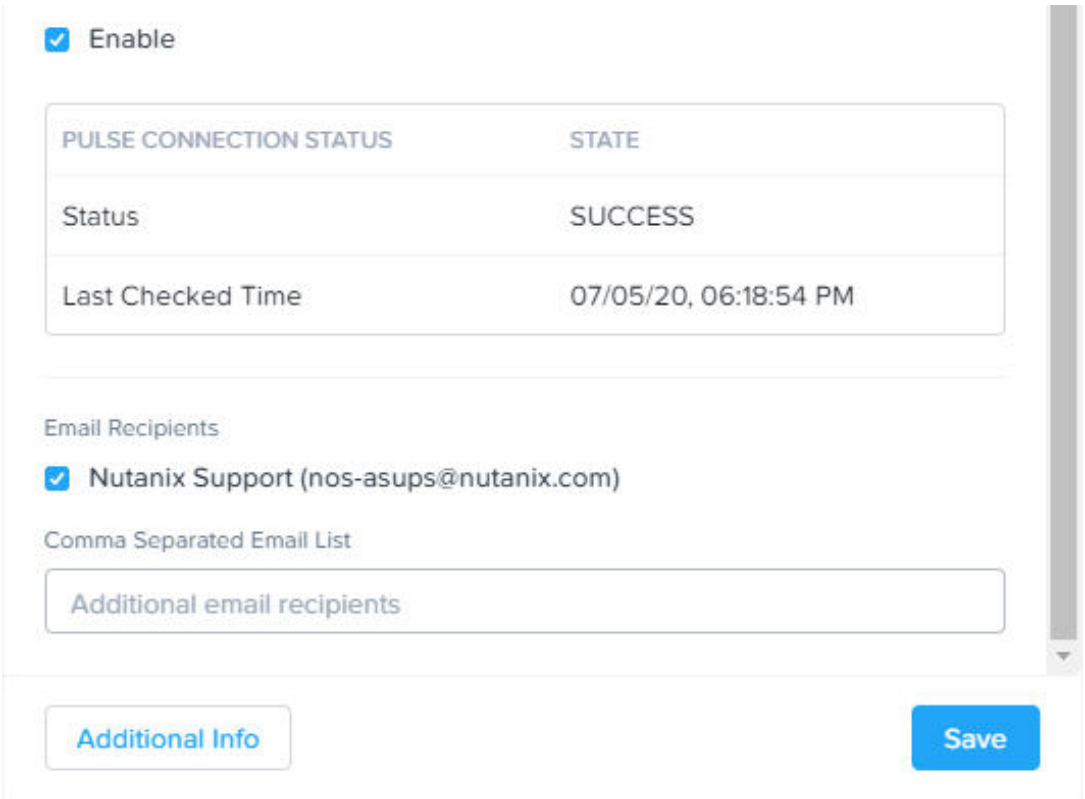

Figure 496: Pulse Additional Info

4. Check the **Pulse Connection Status** field for connection information.

   - **Status**: Displays the transport mechanism status, which is **Success** when the HTTP client can access the Pulse server successfully or **Failure** when it cannot (or **Unknown** when the status is not known).
   - **Last Checked Time**: The last time Pulse sent data and checked the connectivity with the Pulse server.

5. To add (remove) Nutanix customer support as a recipient of the cluster information, check (uncheck) the box next to **Nutanix Support (nos-asups@nutanix.com)** in the **Email Recipients** field.

   Cluster information is e-mailed automatically to Nutanix customer support at **nos-asups@nutanix.com** when this feature is enabled. If you do not want Nutanix customer support to receive the cluster information, you can uncheck this box at any time (or restart the e-mails by checking it again).

6. To send the cluster information to additional recipients, enter one or more e-mail addresses in the **Comma Separated Email List** field.

   In addition to (or instead of) sending the cluster information to Nutanix customer support, you can e-mail the information to yourself or others by entering recipient e-mail addresses in this field. Use a comma to separate multiple addresses.

7. Click the **Save** button to save the new setting and close the window.

Prism Central Proxy for Pulse Data

Prism Central can automatically act as a proxy for Pulse data transmitted by each node in a Prism Element cluster registered to that Prism Central instance.

How Do I Enable Prism Central Proxy for Pulse Data?

You do not have to explicitly enable this feature. It depends on your Prism Central and Prism Element configuration.

Pulse data from Prism Element nodes is automatically routed through Prism Central and then sent to Nutanix Support if you satisfy these requirements:

- You enable Pulse on Prism Central and each registered Prism Element cluster.
- Prism Central and each Prism Element cluster node are running NCC 3.5.2 or later.
- You have not configured an HTTP proxy on Prism Central or, if you have configured an HTTP proxy on Prism Central, it does not require basic authentication (a user name and password).
- You have not configured an HTTP proxy on Prism Element clusters registered to this Prism Central instance. If your Prism Element clusters are configured to use an HTTP proxy, the cluster nodes bypass the Prism Central Pulse proxy and transmit Pulse data to Nutanix Support.

If your Prism Central deployment is not available, the cluster nodes bypass the Prism Central Pulse proxy and transmit Pulse data to Nutanix Support.

For a Prism Central scale out deployment, each Prism Element node selects a Prism Central VM at random to act as its proxy.

Can I Use This Feature If I Have Configured an HTTP Proxy on Prism Central and Prism Element?

If you have configured an HTTP proxy on Prism Central, you can use this feature automatically if you satisfy these requirements:

- You have configured an HTTP proxy on Prism Central but it does not require basic authentication (a user name and password). Configuring an HTTP Proxy on page 508 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 943). 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 30). 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 942). 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.

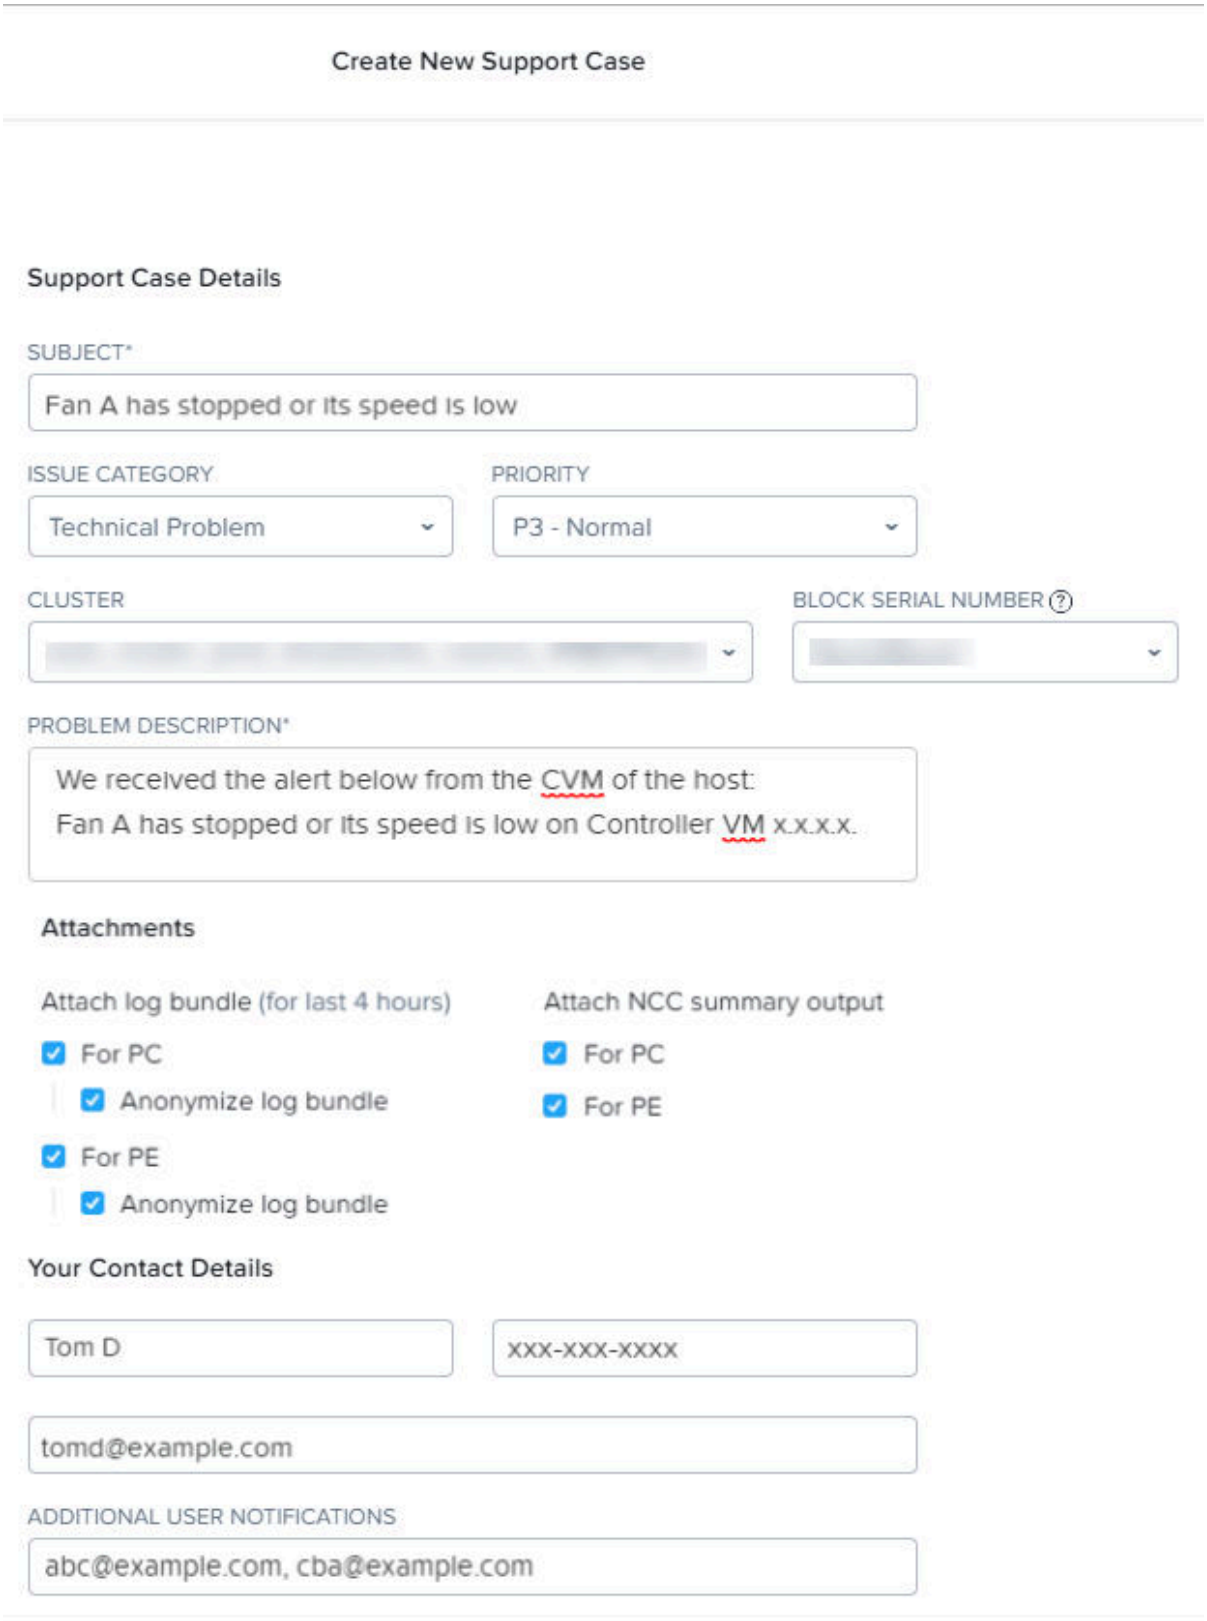

Figure 497: Create New Support Case Page

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 241). 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 942).

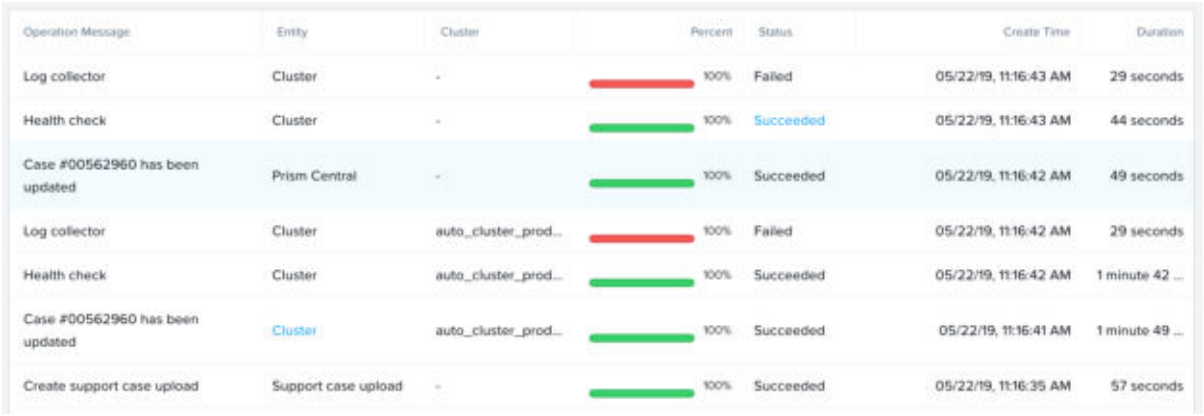

Figure 498: Tasks Dashboard for Prism Central (create case tasks)

Figure 499: 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.

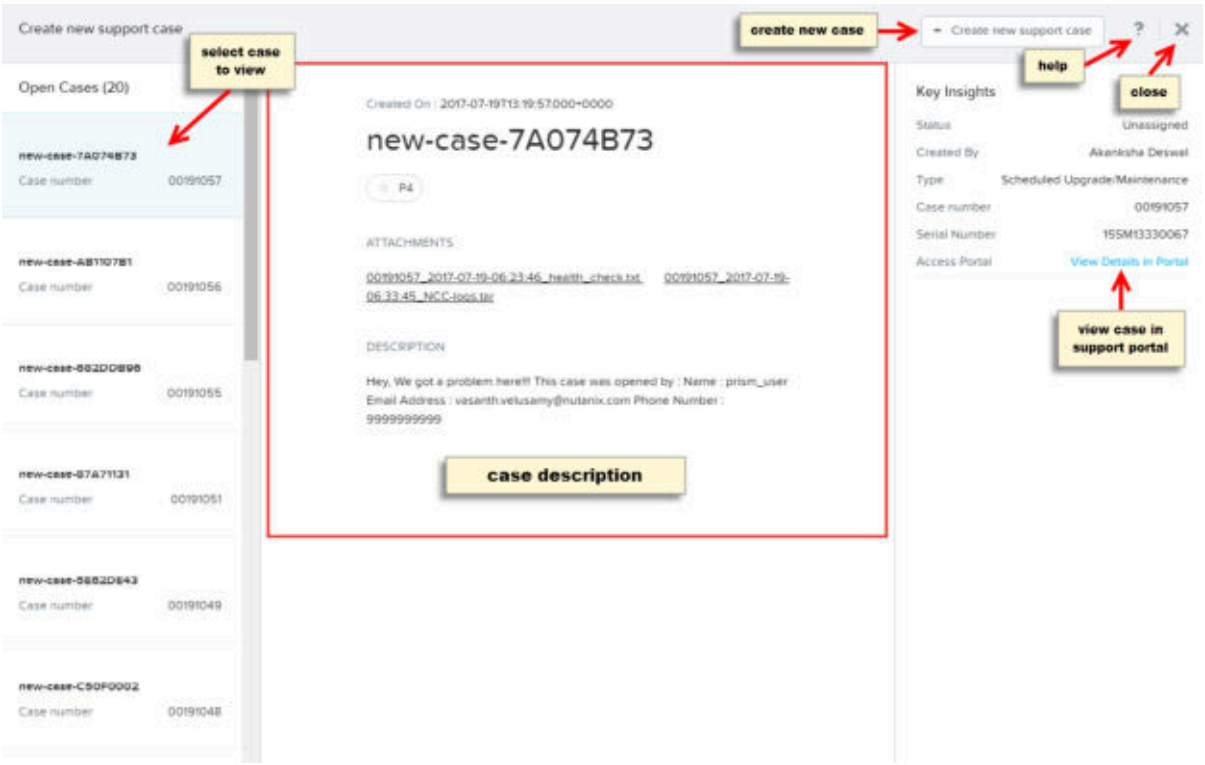

Figure 500: Create New Support Case Page (open cases)

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 743: Main Menu Options

Category	Option	Description
Documentation	Software Documentation	Displays a page from which you can view the Nutanix software manuals.
	Hardware Replacement Documentation	Displays a page from which you can view the Nutanix hardware replacement manuals.
	Knowledge Base	Displays a page from which you can view the knowledge base (KB) articles.
	Solutions Documentation	Displays a page from which you can view documents that describe how to implement the Nutanix platform to solve a variety of business applications.
	EOL Information	Displays a page from which you can view the end of life policy and bulletins.
	Field Advisories	Displays a page from which you can view field advisories.
	Training	Provides a link to the separate Nutanix training portal.
	Security Advisories	Displays a page from which you can view security advisories.
	AOS Upgrade Paths	Displays a page where you can see the supported AOS release upgrade paths.
	Compatibility Matrix	Displays a page from which you can view a compatibility matrix broken down (filtered) by hardware model, AOS version, hypervisor type and version, and feature version (NCC, Foundation, BMC/BIOS).
	Webinar Recordings	Displays a page with links to a selection of Nutanix training webinars.
Support	Open Case	Displays a form to create a support case.
	View Cases	Displays a page from which you can view your current support cases.
	.NEXT Forums	Provides a link to the (separate) Nutanix Next Community forum.
	Terms & Conditions	Displays a page from which you can view various warranty and terms and conditions documents.
Downloads	AOS (NOS)	Displays a page from which you can download AOS releases.

Category	Option	Description
	Acropolis File Services (AFS)	Displays a page from which you can download the Acropolis File Services.
	Acropolis Container Services (ACS)	Displays a page from which you can download the Acropolis Container Services.
	Hypervisor Details	Displays a page from which you can download Acropolis hypervisor versions. You can also download supporting files used when manually upgrading a hypervisor version (AHV, ESXi, or Hyper-V).
	Prism Central	Displays a page from which you can download the Prism Central installation bundle. There are separate bundles for installing on AHV, ESXi, or Hyper-V.
	Phoenix	Displays a page from which you can download Phoenix ISO files.
	Foundation	Displays a page from which you can download Foundation releases.
My Products	Installed Base	Displays a table of your installed Nutanix appliances, including the model type and serial number, location, and support coverage.

Category	Option	Description
	Licenses	Displays a table of your product licenses along with buttons to add or upgrade licenses for your clusters.

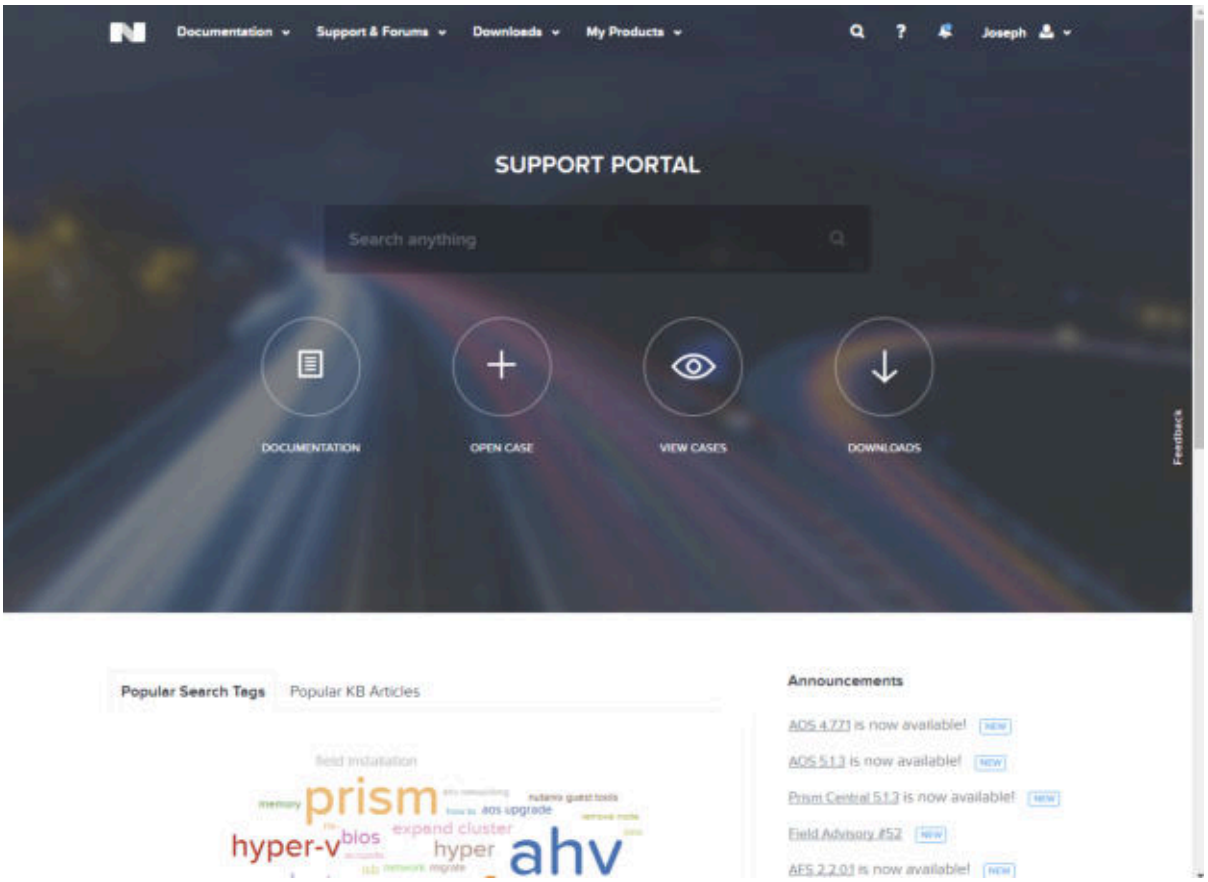

Figure 501: Nutanix Support Portal

Accessing the REST API Explorer

About this task

Nutanix provides a utility with the web console to help you get started with the REST API. The Explorer displays the parameters and format for the API calls that can be included in scripts. Sample API calls can be made to show the type of output you should expect to receive.

The v3 API can be viewed in the REST API Explorer.

Procedure

1. Log on to the web console, click the user icon in the upper-right corner of the web console, and click **REST API Explorer**.

   The REST API Explorer displays a list of the objects that can be managed by the API. Each line has four options:

   - **Show/Hide**: Expand or reduce the detail shown for the object
   - **List Operations**: Show all operations that can be run on this object
   - **Expand Operations**: Show the detailed view of the operations that can be run on this object

2. Find the line for the object you want to explore and click **Expand Operations**.
   For example, click **alerts**.
3. Click **GET** to show the details for this API call.
4. Click **Try it out!** to test the API call when used with your cluster.

15

HELP RESOURCES

There are several information sources that you can access at any time when you need help:

- Prism Central help documentation (see Accessing Online Help on page 948).
- Nutanix customer support portal (see Accessing the Nutanix Support Portal on page 943).
- Nutanix community forum (see Accessing the Nutanix Next Community on page 950).
- Glossary of terms (see Glossary on page 951).

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.

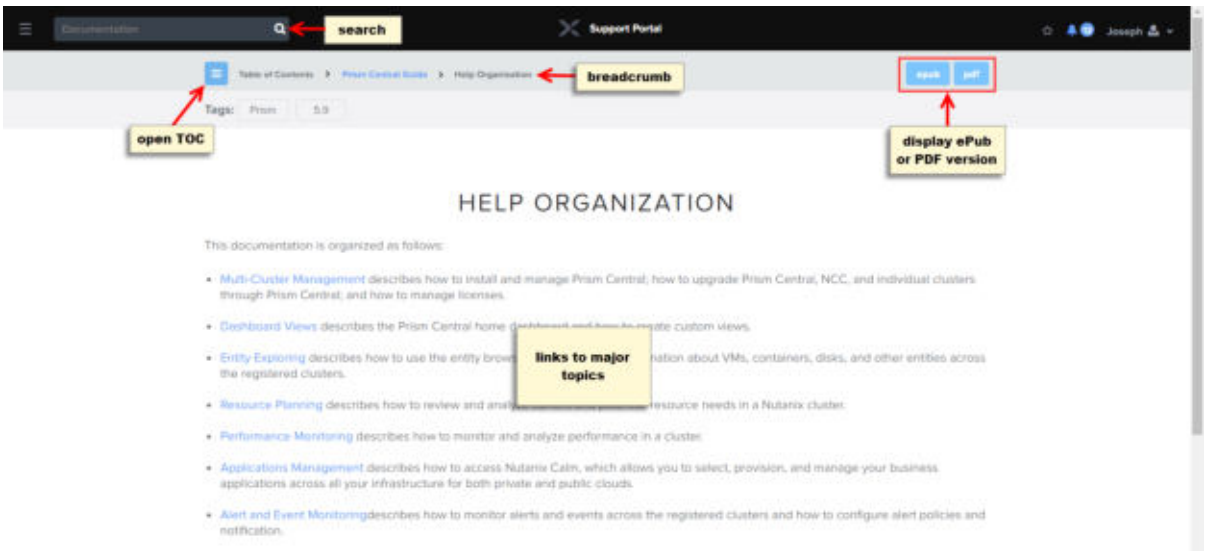

Figure 502: Help Organization Page

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.

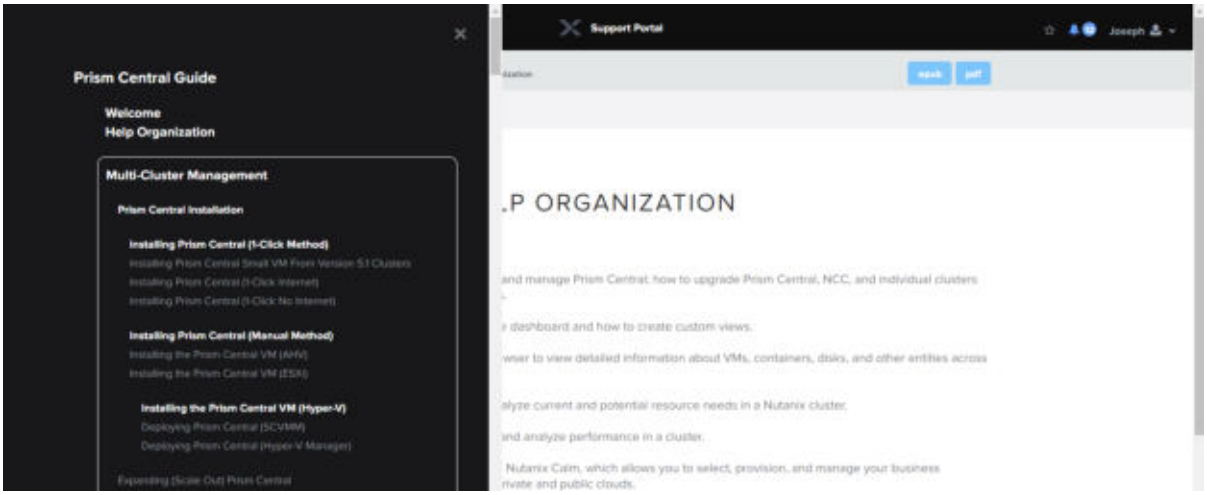

Figure 503: Table of Contents Pane

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.

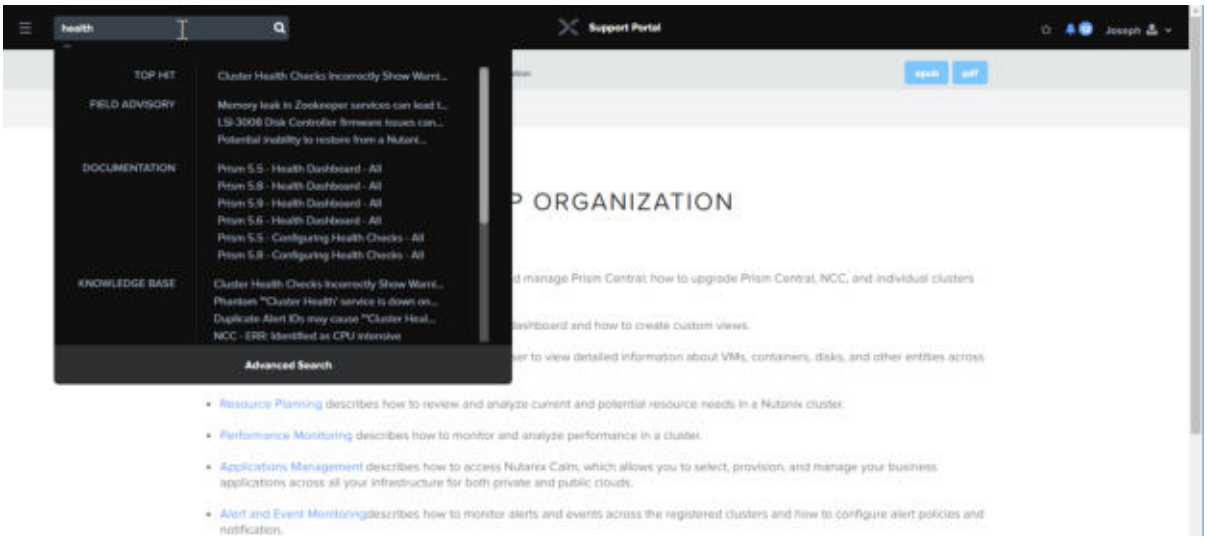

Figure 504: Search Results Example

Accessing the Nutanix Next Community

Nutanix maintains a community forum for customers and partners to facilitate a peer-to-peer exchange of ideas, tips, and information about Nutanix technologies and the rapidly changing landscape of data center IT.

Procedure

- To access the Nutanix next community forum from Prism Central, select **Nutanix Next Community** from the question mark icon pull-down list of the main menu (seeMain Menu on page 10.

  The Nutanix Next Community main page appears in a new tab or window. From this page you can search existing posts, ask questions, and provide comments.

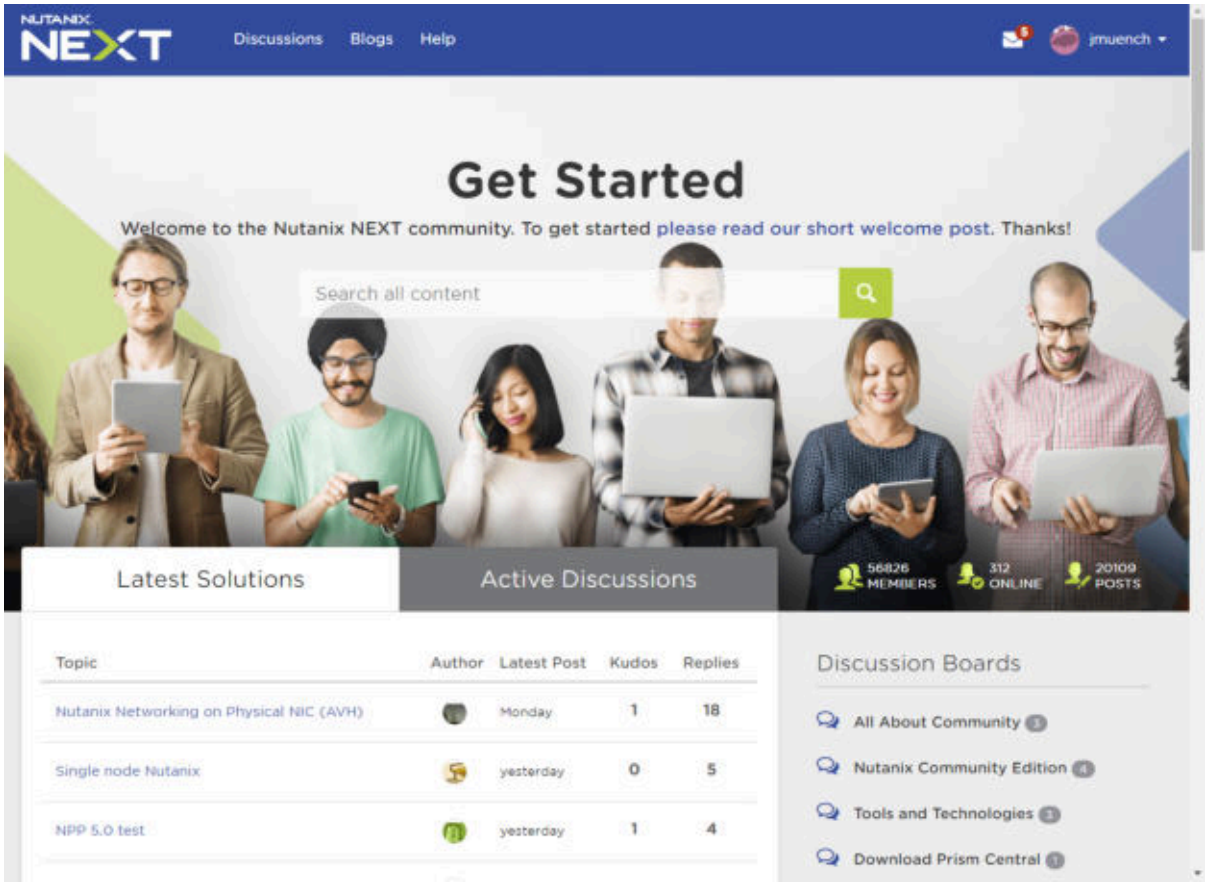

Figure 505: Next Community Screen

Glossary

aCLI

Acropolis command-line interface.

Acropolis

The Nutanix converged software fabric for virtualization and storage management. It consists of the Acropolis base software, Acropolis Distributed Storage Fabric, AHV, App Mobility Fabric, Prism, and Acropolis APIs.

Acropolis App Mobility Fabric

Provides virtualization management, volume management, and other distributed management functions for Acropolis.

AHV

Nutanix-specific hypervisor host providing core server virtualization and optimized for Nutanix cluster and guest VM use.

block

A set of Nutanix nodes contained in a single enclosure.

block fault tolerance

When certain conditions are met, redundant copies of any data required to serve I/O are placed on nodes that are not in the same block. In the case where multiple nodes in a block fail, the cluster can continue to run because cluster configuration data exists on other blocks. Guest VMs can continue to run because redundant copies of guest VM data and metadata exist on other blocks.

clone

A writeable copy of a vDisk.

Cloud Connect

A feature that enables you to back up and restore copies of virtual machines and files to and from an on-premise cluster and a Nutanix Controller VM located on the Amazon Web Service (AWS) cloud. The cloud-based cluster is managed as a remote site from the Prism Web Console or nCLI.

cluster

A group of nodes contained in one or more Nutanix blocks.

cold data

Data that did not have many rewrites or overwrites for a long time. For example, snapshots, file servers, archives, or backups.

compression

An optional storage container setting that enables one of two types of compression.

compression, inline

Data is compressed as it is written. This type of compression is recommended for workloads that perform batch processing.

compression, post-process

Data is compressed after it is written. The delay time between write and compression is configurable. Nutanix has recommends a delay value of 60. However, every workload has a different I/O profile, so you might require a different delay value otherwise. This type of compression is recommended for most workloads.

Consistency Group

A subset of VMs in a Protection Domain. All VMs within a Consistency Group for that Protection Domain will be snapshotted in a crash-consistent manner. For all VMs in a Consistency Group, a snapshot creates one snapshot for all VMs in the group.

Controller VM

A Nutanix VM that manages storage and other cluster functions on a node.

data resiliency

A measure of the types of failures a cluster can withstand; affected by block awareness and redundancy factor.

datastore

A logical storage container for files necessary for VM operations.

deduplication

The sharing of identical guest VM data on premium tiers (RAM and Flash) for improved performance or on capacity tiers (HDD) for storage space savings. Enabled by properties of a storage container or vDisk.

Distributed Storage Fabric

All storage functionality, including snapshots and clones, data protection, disaster recovery, data path redundancy, replication factors, deduplication, compression, erasure coding, and so on.

Enterprise Cloud Platform

The Nutanix solution that natively converges compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence

erasure coding

Optional algorithm included in the Acropolis base software to help reduce the storage used for fault tolerance. It helps to increase the effective or usable capacity on a cluster, depending on cluster size and data coldness.

failback

Planned failover initiated from recovery site.

failover

Moving VMs from a primary site to a recovery site.

failover, disaster

Failover when the primary site is down.

failover, planned

Failover when both sites are up.

guest VM

A VM running on a Nutanix cluster that executes a workload, such as VDI or Exchange, as opposed to a VM that is involved in cluster operations, such as a Controller VM.

host

An instance of the hypervisor that runs on a Nutanix node.

image service

A workflow in the Prism web console that enables a user to upload ISO or disk images (in ESXi or Hyper-V format) to a Nutanix AHV cluster by specifying a remote repository URL or by uploading a file from a local machine.

local replication

Multiple copies of data within a storage container. These copies exist for fault tolerance: if a physical disk fails, the cluster can recover data from another copy. The cluster manages the replicated data, and the copies are not visible to the user.

local snapshot

Snapshots stored on the same cluster where they originated.

nCLI

Nutanix command-line interface.

node

A physical server contained in a Nutanix block; runs a hypervisor host.

oplog

A write cache on a faster, more expensive storage tier.

Prism

Web-based management interface for managing Nutanix clusters.

Prism Central

Centralized management tool that runs as a separate VM configured as a single-node cluster to monitor and manage multiple clusters through a single web console.

Prism Element

A single cluster being managed by and available through the Prism Central web console.

Protection Domain

A group of VMs to be backed up locally on a cluster or replicated on the same schedule to one or more clusters. Protection Domains may be associated with remote sites. Protection Domain names must be unique across sites. A VM can be in at most one Protection Domain.

Protection Domain, active

A Protection Domain that manages live VMs and makes, replicates, and expires snapshots.

Protection Domain, inactive

A Protection Domain that receives snapshots from a remote cluster.

redundancy factor

The number of nodes plus 1 that the cluster can tolerate being down at one time. By default, Nutanix clusters have a redundancy factor of 2, which means that they can tolerate 1 node being down. They are configurable to redundancy factor 3 to enable tolerating 2 nodes being down.

remote replication, one-to-one

Replicates a production cluster with one idle cluster as a DR target.

remote replication, reciprocal

Cross replication within running (non-idle) production clusters.

remote site

A pair of clusters that can replicate data to each other. A configured remote site can also be located in the cloud (based on Amazon AWS, for example).

remote snapshot

A snapshot copied asynchronously from one cluster to another.

replication factor

The number of copies of data kept by a storage container. The default is 2. Storage Containers on clusters with redundancy factor of 3 can have replication factor of 3.

reserved capacity

A property of a storage container or vDisk that guarantees that a certain amount of storage space is available.

schedule

A property of a Protection Domain that specifies the intervals to take snapshots and how long the snapshots should be retained. A schedule optionally specifies which remote site to replicate to.

Self-Service Restore

Allows a user to restore a file within a virtual machine from the Nutanix protected snapshot with minimal Nutanix administrator intervention.

shadow clone

A cache of a vDisk on all the nodes in the cluster. When a vDisk is read by multiple VMs (such as the base image for a VDI clone pool), the cluster creates shadow clones of the vDisk.

snapshot

A read-only copy of the state and data of a VM at a point in time. Snapshots for a VM are crash consistent, which means that the VMDK on-disk images are consistent with a single point in time. That is, the snapshot represents the on-disk data as if the VM crashed. The snapshots are not, however, application consistent, meaning that application data is not quiesced at the time of snapshot, unless the Protection Domain is configured to use application-consistent snapshots.

storage container

A subset of available storage within a storage pool.

storage pool

A group of physical disks from one or more tiers.

storage replication adapter

A Nutanix-provided module that allows VMware Site Replication Manager (SRM) to use native remote replication.

tier

A type of physical storage in a Nutanix node.

vDisk

Data associated with a VM represented as a set of files on a datastore.

VM high availability

In virtualization management VM high availability, when a node becomes unavailable, VMs that are running on that node are restarted on another node in the same cluster.

VM mobility

The ability to export your existing VMs from one non-AHV cluster to an AHV cluster. This option requires that you install the Nutanix VM Mobility installer on all the VMs. The Nutanix VM Mobility installer deploys the drivers that are required at the destination AHV cluster. After you prepare the source VMs, they can be exported to the AHV cluster.

vStore

A separate mount point within a storage container which has its own NFS namespace. This namespace maps to a Protection Domain. Each vStore is exported as a device through the Nutanix SRA.

COPYRIGHT

Copyright 2022 Nutanix, Inc.

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

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

License

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

Conventions

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

Version

Last modified: March 22, 2022 (2022-03-22T16:59:34+05:30)

BOOKLIST

Related Guides

Refer to the following guides for more information:

- *Acropolis Upgrade Guide*
- *Security Guide (AOS Security 5.18)*
- *Data Protection and Recovery with Prism Element*

PD-Based DR 6.1

Data Protection and Recovery with Prism Element

March 23, 2022

Contents

1. Data Protection and Recovery with Prism Element 6
   - Terminology 7
   - Protection Strategies 9
   - Nutanix DR Solutions Workflow 13
   - Snapshots 15
     - Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro) 16
     - Conditions for Application-consistent Snapshots 18
   - VStore Data Protection 25

2. Data Protection Dashboard 27
   - Data Protection Overview View 28
   - Data Protection Table View 29

3. Data Protection with Async DR and NearSync Replication 42
   - Data Protection with Asynchronous Replication (One-hour or Greater RPO) 42
     - Conditions for Implementing Asynchronous Replication 42
     - Limitations of Data Protection with Asynchronous Replication 46
     - Configuring Data Protection with Asynchronous Replication 49
   - Data Protection with NearSync Replication (1–15 minutes RPO) 57
     - Requirements of Data Protection with NearSync Replication 60
     - Limitations of Data Protection with NearSync Replication 62
     - Handling Network Partition (NearSync) 63
     - Configuring Data Protection with NearSync Replication 64
   - Interoperability Between Async DR and NearSync in a Protection Domain 71
   - Modifying Asynchronous or NearSync Replication Configurations 73
     - Deleting a Protection Domain 76
   - Failover and Failback Operations for Asynchronous and NearSync Replication 77
     - Performing Failover 77
     - Failing Back a Protection Domain 80
     - Performing a Failback 81

4. Metro Availability (ESXi and Hyper-V 2016) 82
   - Conditions for Implementing Data Protection (Metro Availability) 83
   - Limitations of Data Protection with Metro Availability 86
   - vSphere Cluster Services on Nutanix Clusters 87
   - vSphere Metro Storage Cluster for Nutanix Metro Availability 88
   - Upgrade Best Practices and Requirements with Metro Availability Enabled 89
   - Metro Availability on ESXi 90
     - Configuring a Protection Domain (Metro Availability) 90
     - Protection Domain Failover and Failback (ESXi) 103
     - Reconfiguring Data Protection in a Metro Availability and Asynchronous Replication Configuration 107
     - Metro Availability Protection Domain and Container Management (ESXi Clusters) 107
   - Metro Availability on Hyper-V 2016 110
     - Summary: Preparing to Configure Metro Availability for Hyper-V 111

Configuring a Metro Availability Relationship Between the Nutanix Clusters 112
Performing a Planned Failover (Manual Live Migration) 113
Performing an Unplanned Failover (Disaster) 113
Re-enabling Metro Availability (Failback) 114

5. Witness Option 116
Metro Availability Witness Option 116
Witness VM Dashboard 121
Installing a Witness VM 122
Upgrading a Witness VM 125
Registering a Witness VM 126
Changing the Witness VM for a Protection Domain 128
Configuring Witness Mode 128
Changing the IP Address of a Witness VM 129
Recovery Procedures (Witness VM installed) 131
Configuring a Witness (Two-node Cluster) 138
Witness VM Dashboard (Two-node Cluster) 140

6. Synchronous Replication (ESXi and Hyper-V 2012) 142
Failover Scenarios for Synchronous Replication in Hyper-V 143
VMs are Originally Registered and Running on the Primary Site and Failure Occurred on the Primary Site 143
VMs are Originally Registered and Running on the Secondary Site and Failure Occurred on the Primary Site 143
VMs are Originally Registered and Running on the Secondary Site and Failure Occurred on the Secondary Site 144
Failback Scenarios for Synchronous Replication in Hyper-V 144
Failback VMs that are Currently Running on the Secondary Site but were Originally Registered on the Primary Site 144
Failback VMs that are Currently Running on the Secondary Site and were Originally Registered on the Secondary Site 145
Failing Over a Protection Domain Manually in Synchronous Replication (Planned Failover) 146
VM Registration and De-registration 146
Registering a VM on a Hyper-V Cluster 146
Unregistering a VM on a Hyper-V Cluster 147
Deployment Scenario: Transferring a Synchronous Replication Setup to a Standby Site 148

7. Remote Site Configuration 152
Configuring a Remote Site (Physical Cluster) 152
Modifying a Remote Site (Physical Cluster or Cloud) 160
Network Mapping 162
Single-Node Replication Target Clusters 163
Single-Node Replication Target Requirements and Limitations 165
Imaging Single-Node Replication Target Nodes 167
Setting Up Single-Node Replication Target Clusters 168
Recovery Procedures by Using Single-Node Replication Target Clusters 168

8. Self-Service Restore 169
Requirements and Limitations of Self-Service Restore 169
Enabling Self-service Restore 171
Self-Service Restore for Windows VMs 171

Restoring a File by Using Web Interface for Windows VMs....171
Restoring a File by using ngtcli for Windows VMs....173
Self-Service Restore for Linux VMs....174
Restoring a File by Using Web Interface for Linux VMs....174
Restoring a File by using ngtcli for Linux VMs....176

9. Restoration of Protected Entities....178
Restoring an Entity from a Protection Domain....179
Recovering from the Remote Snapshots on a Backup or DR Site....182
Recovery of vGPU-enabled VMs....184

10. Shutting Down and Starting Clusters in Both Primary and Remote Sites....186
Shutting Down the Clusters....186
Starting the Clusters....187

11. Nutanix Cross Hypervisor Disaster Recovery....189
UEFI and Secure Boot Support for CHDR....190
Recovering Virtual Machines in the Remote Cluster....192
Bringing Multiple SCSI Disks Online....193
Blacklisting Devices for Multipathing....194

12. Asynchronous Replication Using Cloud Connect (On-Prem to Cloud)....196
Implementation Considerations for Cloud Connect....196
AWS Account Requirements....198
Instance Types for AWS....201
Configuring a Remote Site in AWS....201
Recovery Procedures by using Cloud Connect....205
Upgrading a Nutanix Controller VM in the Cloud....205

13. Support Services....207
Pulse Health Monitoring....207
Configuring Pulse....209
Pulse Access Requirements....211
Controlling Remote Connections....211
Configuring HTTP Proxy....212
Accessing the Nutanix Support Portal....215
Nutanix REST API....218
Accessing the REST API Explorer....218

14. Help Resources....222
Accessing Online Help....222
Accessing the Nutanix Next Community....223
Glossary....223

Copyright..... 228
License..... 228
Conventions..... 228
Version..... 228

1

DATA PROTECTION AND RECOVERY WITH PRISM ELEMENT

This guide describes the concepts and procedures for configuring data protection using protection domains. You can configure data protection using protection domains in Prism Element (in other words, by signing in to an individual cluster through its web console).

For information about configuring data protection by using the newer, entity-centric approach, see the *Leap Administration Guide*. Entity-centric data protection focuses its workflows on the entity that you want to protect. Entity-centric data protection is configured by using Prism Central.

> Note: You cannot protect a VM simultaneously by a protection domain of protection domain-based DR and a protection policy of Leap. If you want to use a protection policy to protect a VM that is part of a protection domain, first remove the VM from the protection domain, and then include it in the protection policy. For information, see *Migrating VMs from a Protection Domain to a Protection Policy* in the *Leap Administration Guide*.

You can enable network segmentation for disaster recovery on a stretched Layer 2 network that does not have a gateway. For more information, see *Segmenting a Stretched L2 Network for Disaster Recovery* in the *Securing traffic through Network Segmentation* section of the *Security Guide*.

This guide describes the concepts and procedures for configuring data protection using protection domains. You can configure data protection using protection domains in Prism Element (in other words, by signing in to an individual cluster through its web console).

For information about configuring data protection by using the newer, entity-centric approach, see the *Leap Administration Guide*. Entity-centric data protection focuses its workflows on the entity that you want to protect. Entity-centric data protection is configured by using Prism Central.

Considerations that Affect Data Protection and Recovery

- You cannot protect a VM simultaneously by a protection domain of protection domain-based DR and a protection policy of Leap. If you want to use a protection policy to protect a VM that is part of a protection domain, first remove the VM from the protection domain, and then include it in the protection policy. For information, see *Migrating VMs from a Protection Domain to a Protection Policy* in the *Leap Administration Guide*.
- Ensure that specific ports are open in your firewall to allow communications between the primary site and the replication site when you configure a protection domain or remote site. See *Firewall Requirements* in the *Security Guide*.
- You cannot recover vGPU console-enabled guest VMs efficiently.

  When you perform disaster recovery of vGPU console-enabled guest VMs, the VMs recover with the default VGA console (without any alert) instead of vGPU console. The guest VMs fail to recover when you perform cross-hypervisor disaster recovery (CHDR).

Terminology

The data protection features for a Nutanix cluster employ the following components and capabilities.

This section provides information about the terminology used and concepts discussed in this guide.

Category

A category is a key-value pair that groups similar VMs. Associating a policy with a category ensures that the policy applies to all the VMs in the group regardless of how the group scales with time. For example, you can associate a group of VMs with the `Department: Marketing` category, where `Department` is a category that includes a value `Marketing` along with other values such as `Engineering` and `Sales`.

Categories operate in the same way on on-prem Prism Central instances and in Xi Cloud Services. For more information about categories, see the *Prism Central Guide*.

Consistency Group

A consistency group is a subset of the entities in a protection domain. Consistency groups become operational when you create a protection domain. The system captures snapshots for all the entities within a consistency group for that protection domain in a crash-consistent manner. The system generates a snapshot for each VM in the consistency group simultaneously at the same time.

Primary Site

Primary site is the local cluster or clusters that host the entities that require protection and whose snapshots must be replicated on another cluster or a remote site.

Prism Element

The interface that provides you the ability to configure, manage, and monitor a single Nutanix cluster running on an AHV or ESXi hypervisor. It is a service built into the platform for every Nutanix cluster deployed.

Protection Domain

There are two types of protection domains.

- Protection Domain (Async DR). A standard (Async DR) protection domain is a defined group of entities (virtual machines, volume groups, or storage containers) that must be protected. These entities are backed up locally on a cluster and optionally replicated to one or more remote sites (see Configuring Data Protection with Asynchronous Replication on page 49). This type of protection domain uses asynchronous data replication to create snapshots. protection domain names must be unique across sites. Configure an entity in only one protection domain.

  A protection domain on a cluster is in one of two modes:

  Active

  Manages volume groups and live VMs; makes, replicates, and expires snapshots.

  Inactive

  Receives snapshots from a remote cluster.

  > Note: The protection domains at a remote site or cluster are inactive by default. Activate them from the Prism Element Web Console on the remote cluster.

- Protection Domain (Metro Availability). A metro availability protection domain consists of a specified (active) storage container in the local cluster. This active storage container is linked to a (standby) container on a remote site. The standby container must have the same name. Synchronous data replication occurs to the linked standby container when metro availability is enabled (see Configuring a Protection Domain (Metro Availability) on page 90).

Recovery Point

A recovery point is a copy of the state of a system at a particular point in time. A recovery point refers to a snapshot.

Recoverable Entity

A VM that you can recover from a recovery point.

Recovery Plan

A configurable policy that orchestrates the recovery of VMs at the recovery site.

Recovery Point Objective (RPO)

A time interval that refers to the acceptable data loss if a failure occurs. For example, if the RPO is 1 hour, the system creates a recovery point every 1 hour. In the event of a recovery, you can recover VMs with data as of up to 1 hour ago.

Recovery Time Objective (RTO)

The time period from failure event to the restored service. For example, an RTO of 30 minutes enables you to back up and run VMs in 30 minutes in the event failure.

Related Entities

Related entities are the VMs and volume groups that are associated with the entities that you want to protect. The association can be direct or transitive. Direct means direct hypervisor attachment or in-guest iSCSI. Transitive means if a volume group is attached to two VMs, then the VMs are considered to be related to each other.

For example, if a volume group named VG1 is directly attached to a VM named VM1 as a SCSI disk, VG1 and VM1 are related entities. If VG1 is connected to a VM named VM2, VG1 and VM2 are related entities. In this configuration, VM1 and VM2 are considered to be related entities too because they are associated with the same volume group.

For the requirements to include related entities and limitations, see Conditions for Implementing Asynchronous Replication on page 42.

Remote Site

A remote site is a separate cluster used as a target location to replicate and store data asynchronously or a standby storage container for metro availability. A remote site maybe another cluster in the same physical site as well.

You can configure one or more remote sites for a cluster (see Configuring a Remote Site (Physical Cluster) on page 152 orAsynchronous Replication Using Cloud Connect (On-Prem to Cloud) on page 196).

Replication

Replication is the process of asynchronously copying snapshots from one cluster to one or more remote sites. Several replication scenarios are supported (see Protection Strategies on page 9).

Snapshot

A snapshot is read-only copy of the state and data of a VM or volume group at a point in time.

A local snapshot is a snapshot taken and stored on the local cluster or site based on the protection domain configuration.

A remote snapshot is a snapshot replicated on remote sites from the protected entities based on the replication scenarios.

Crash-consistent Snapshots

A snapshot is crash-consistent if it captures all of the data components (write order consistent) at the instant of the crash. VM snapshots are crash-consistent (by default), which means that the vDisks that the snapshot captures are consistent with a single point in time. Crash-consistent snapshots are more suited for non-database operating systems and applications which may not support quiescence (freezing) and un-quiescence (thawing) and such as file servers, DHCP servers, print servers.

Application-consistent Snapshots

A snapshot is application-consistent if, in addition to capturing all of the data components (write order consistent) at the instant of the crash, the running applications have completed all their operations and flushed their buffers to disk (in other words, the application is quiesced). Application-consistent snapshots capture the same data as crash-consistent snapshots, with the addition of all data in memory and all transactions in process. Therefore, application-consistent snapshots may take longer to complete.

Application-consistent snapshots are more suited for systems and applications that can be quiesced and un-quiesced or thawed, such as database operating systems and applications such as SQL, Oracle, and Exchange.

Schedule

Schedule is a property of a protection domain that specifies the time intervals used to take snapshots and time durations that the snapshots must be retained for. Schedules are set up when configuring a protection domain. A schedule optionally specifies the remote site or sites that the data is replicated to.

Time Stream

A time stream is a set of snapshots that are stored on the same cluster as the source VM or volume group.

Ports and Protocols

Nutanix deploys a number of ports and protocols in its software. To see the ports and protocols used for Disaster Recovery, see Ports Reference.

Protection Strategies

Replication is a fundamental component of any enterprise data protection solution, ensuring that critical data and applications can be reliably and efficiently replicated to a different site or a separate infrastructure. While enterprise IT architects have many technology options, replication capabilities are important requisites for any successful enterprise data protection initiative.

Nutanix provides data protection functions at the virtual machine, file, and volume group levels to ensure that VMs and data remain safe.

You can implement a data protection strategy by configuring protection domains and remote sites through the web console. You can restore a protected VM or volume group at any time through the Prism Element web console.

The web console also allows you to monitor data protection implementations and status across the cluster.

Nutanix supports several types of protection strategies including one-to-one or one-to-many replication. These strategies are as follows:

- Per-VM Backup. The ability to designate certain VMs for backup to a different site is particularly useful in branch office environments. Typically, only a subset of VMs running in a branch location require regular back up to a central site. Such per-VM level of granularity, however, is not possible when replication is built on traditional storage arrays. In these legacy environments replication is performed at a coarse grain level, entire LUNs or volumes, making it difficult to manage replication across multiple sites.
- Selective Bi-directional Replication. In addition to replicating selected VMs, a flexible replication solution must also accommodate a variety of enterprise topologies. It is no longer sufficient to simply replicate VMs from one active site to a designated passive site, which can be "lit up" in event of a disaster. Supporting different topologies demands that data and VMs can be replicated bi-directionally.
- Synchronous Datastore Replication (metro availability). Datastores can be spanned across two sites to provide seamless protection in the event of a site disaster (see Conditions for Implementing Data Protection (Metro Availability) on page 83).

The Nutanix native replication infrastructure and management supports a wide variety of enterprise topologies to meet real-world requirements. The following are four replication options to consider.

1. Two-Way Mirroring. The ability to mirror VM replication between multiple sites is necessary in environments where all sites must support active traffic. Consider a two-site example. Site 2 is used as the target for selected workloads running on Site 1. At the same time, Site 1 serves as the data protection target for designated workloads running at Site 2. In this scenario there are active workloads running on both sites simultaneously, such that there are no idle resources in either location. Utilizing storage, compute and networking resources at both locations has a significant advantage over traditional data protection strategies where servers sit idle in anticipation of a future data disaster event.

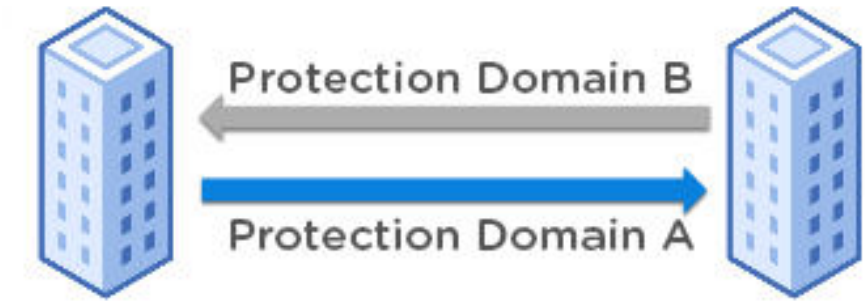

Figure 1: Two-Way Mirroring Topology

2. One-to-Many. In a different scenario, there may be one central site with multiple remote locations. Consider an example where tier-one workloads run at site 1, and sites 2 and 3 serve as remote back-up locations. Site 1 workloads can then be replicated to both 2 and 3

locations. In the event of a data disaster event, the protected workloads can be started on either the desired replication sites for greater overall VM availability.

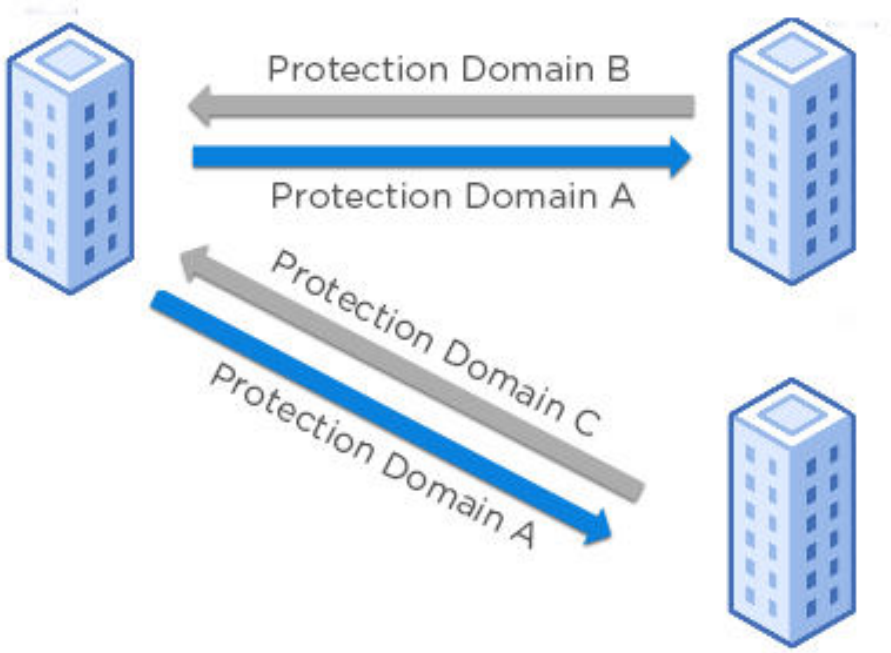

Figure 2: One-to-Many Topology

A one-to-many topology can also be designed to optimize bandwidth between sites. For example, assume the available wide area network (WAN) bandwidth between sites 1 and 3 is greater than that between sites 1 and 2. (Sites 1 and 2 could be in the same city, whereas site 3 may be across the country). In this case, the replication schedule can be set such that larger size VMs running at site 1 are replicated to site 2 in order to conserve bandwidth and improve performance. Similarly, smaller VMs is backed up to site 3, to make better use of lower bandwidth resources.

3. Many-to-One. In a hub and spoke architecture workloads running on site 1 and 2, for example, can be replicated to a central site 3. Centralizing replication to a single site may

improve operational efficiency for geographically disperse environments. Remote and branch offices (ROBO) are a classical use case of a many-to-one topology.

> Note: Many-to-one replication with the same protection domain name is not supported.

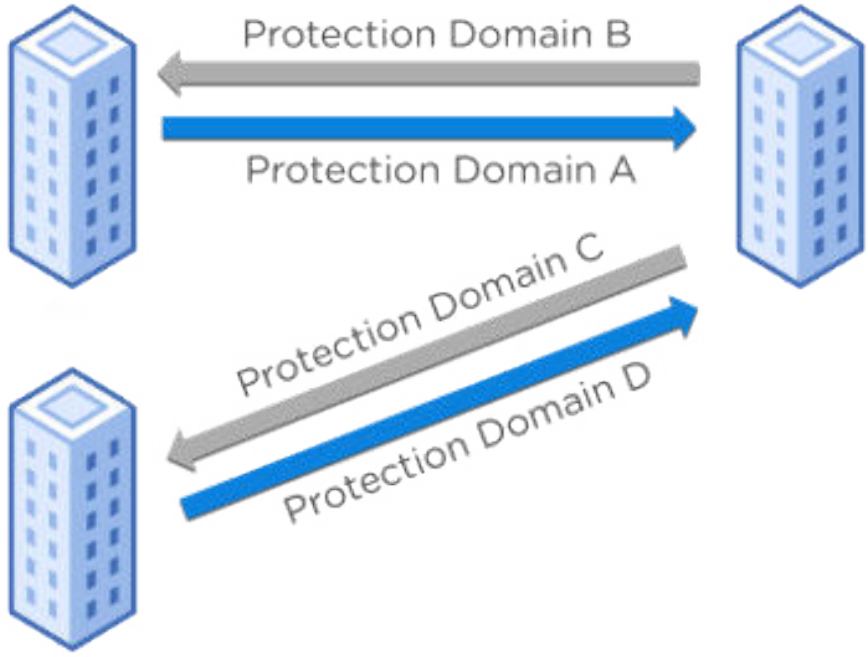

Figure 3: Many-to-One Topology

4. Many-to-Many. This topology allows for the most flexible setup. Here, IT departments have a maximum amount of control and flexibility to ensure application and service level continuity.

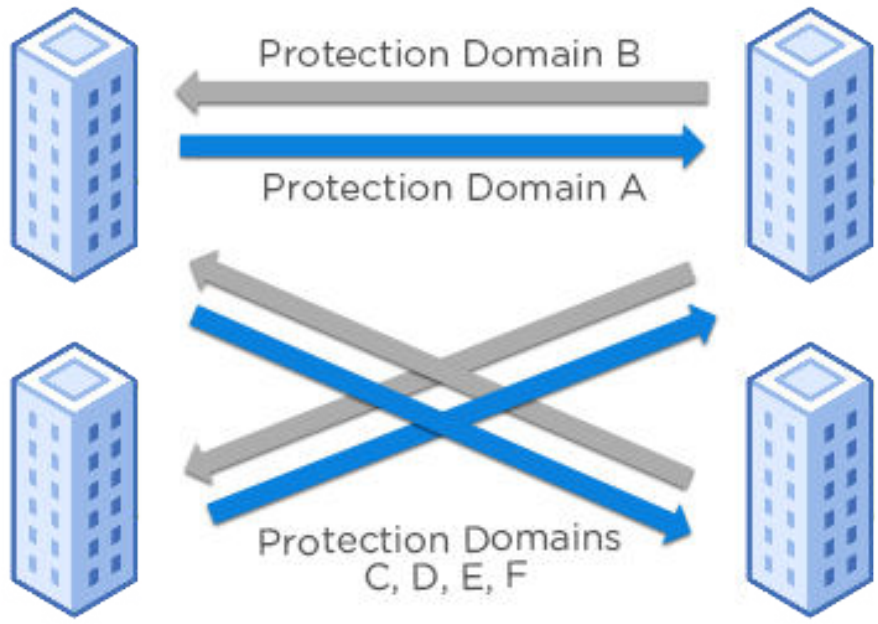

Figure 4: Many-to-Many Topology

Nutanix DR Solutions Workflow

The following flowchart provides you with the detailed representation of the Nutanix DR solutions. This workflow helps you to make quick decision on which DR strategy is best suited for your environment. This workflow covers all the three; protection-domain-based DR, Leap (on-prem to on-prem), and Xi Leap (On-prem to Xi Cloud Services) DR solutions of Nutanix.

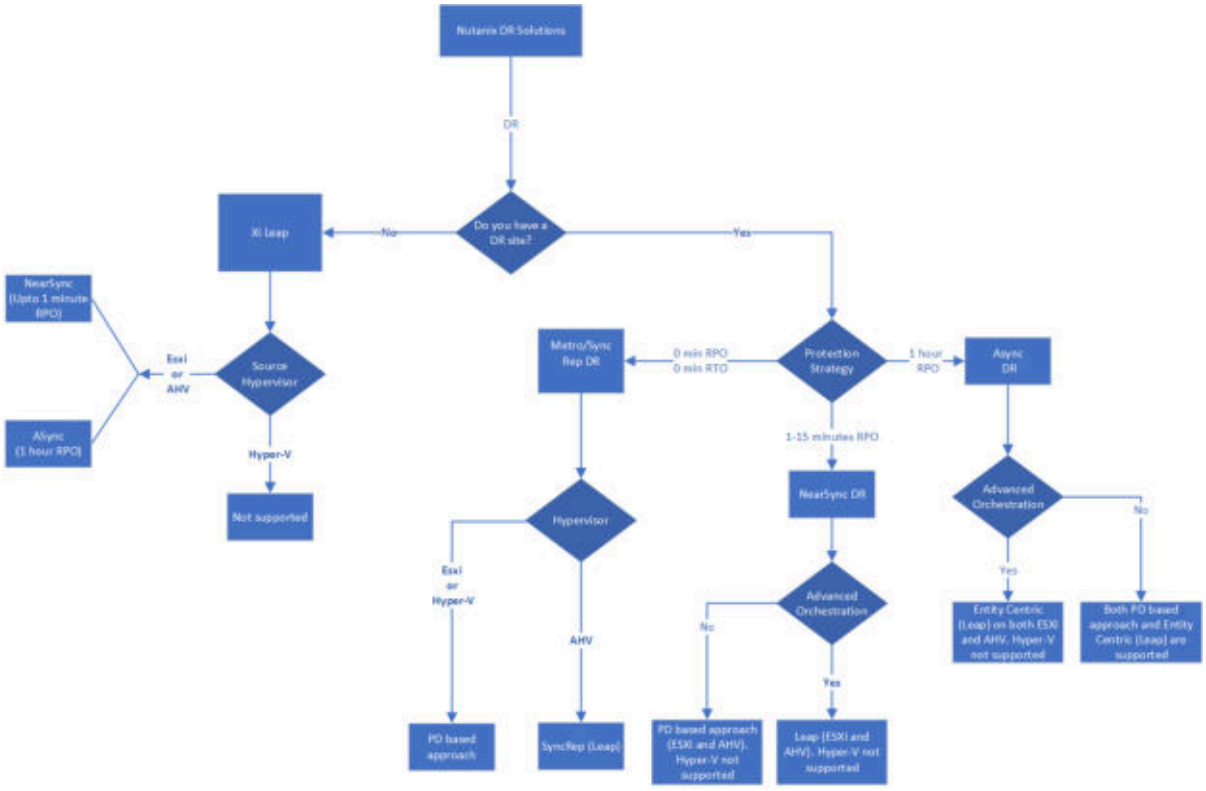

Figure 5: Nutanix DR Workfow

Nutanix provides disaster recovery in Prism Element Web Console with the following Recovery Point Objective (RPO) ranges:

Async DR: Sixty minutes or greater

Schedules with an RPO greater than or equal to 60 minutes are referred to only as asynchronous schedules and use full snapshots. Schedules can be configured in terms of hours, days, weeks, and months.

NearSync DR: Between 1 and 15 minutes

Schedules with an RPO between 1 and 15 minutes are called NearSync schedules. In reality, these schedules are also asynchronous schedules, but they have 15-minute or lower RPOs. Schedules can be configured in terms of minutes.

No disaster Recovery with RPO Between 15 and 59 minutes

Data Protection or disaster recovery schedules between 15 and 59 minutes are not possible.

For detailed information about

- Asynchronous disaster recovery or Async DR, see Data Protection with Asynchronous Replication (One-hour or Greater RPO) on page 42
- NearSync disaster recovery, see Data Protection with NearSync Replication (1–15 minutes RPO) on page 57
- Metro Availability, see Metro Availability (ESXi and Hyper-V 2016) on page 82
- Leap (on-prem to on-prem) or Xi Leap (on-prem to Xi), see the *Leap Administration Guide* which provides detailed information about various protection strategies using protection policy-based data protection and disaster recovery.

Snapshots

The hypervisor creates snapshots of the entities that are protected using any of the disaster recovery solutions that you may choose to use for your deployment.

Snapshot is a point-in-time state of entities such as VM and Volume Groups, and used for restoration and replication of data.. Nutanix DR solutions create snapshots for virtual machine, containers (for metro availability) or for the hypervisor itself. Using the read-on-write (ROW) implementation for its snapshots, Nutanix reduces any performance impacts associated with other implementations.

Snapshots are created locally and stored locally or remotely. Snapshots are created based on time schedules that you construct. Each Nutanix DR solution permits time schedules with certain frequency for creation of snapshots. The frequency of creating snapshots also determines the resources like memory that are required to run the DR solution successfully. See Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro) on page 16 for information about the supported frequency for the snapshots across various hardware configurations.

> Note: You can view the expired snapshots, both local and remote, by selecting the **Include Expired** option to include the expired snapshots in the table in the Local Snapshots and Remote Snapshots tabs in the Data Protection table view.

Nutanix data protection snapshots: Nutanix uses the self-service restore feature to provide file level restoration. This feature allows virtual machine administrators to do a self-service recovery using the Nutanix data protection snapshots with minimal administrator intervention. See Self-Service Restore on page 169 for more information.

Snapshots are of two types in terms of consistency of metadata and data captured - Crash-consistent snapshots and Application-consistent snapshots.

Crash-consistent Snapshots

VM snapshots are by default crash-consistent, which means that the vDisks captured are consistent with a single point in time. The snapshot represents the on-disk data as if the VM crashed or the power cord was pulled from the server—it does not include anything that was in memory when the snapshot was taken.

Application-consistent Snapshots

Application-consistent snapshots capture the same data as crash-consistent snapshots, with the addition of all data in memory and all transactions in process. Because of their extra content, application-consistent snapshots take longer time to complete.

Nutanix supports Application-consistent snapshots. The Nutanix application-consistent snapshot uses Nutanix Volume Shadow Copy Service (VSS) to quiesce the file system for ESXi and AHV prior to taking the snapshot. You can configure which type of snapshot each protection domain should maintain. Nutanix also supports third-party scripts for pre-freeze and post-thaw quiesce and unquiesce operations.

See Conditions for Application-consistent Snapshots on page 18 section for more information.

> Note:
>
> The VM snapshot names have the following restrictions:
>
> - The maximum length of the name is 80 characters.

- You can include only the following characters:
  - Uppercase and lowercase standard Latin letters (A-Z and a-z)
  - Decimal digits (0-9)
  - Dots (.)
  - Hyphens (-)
  - Underscores (_)

Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro)

For DR solutions with Asynchronous, NearSync, and Synchronous (using metro availability) replication schedules to succeed, the nodes must have certain resources. This section provides information about the node configurations necessary to support the asynchronous, NearSync, and synchronous replication schedules.

- The conditions and configurations provided in this section apply to Local and Remote snapshots.
- Any node configuration with two or more SSDs, each of 1.2 TB or greater capacity, supports NearSync.
- Any node configuration that supports snapshots at a frequency of 6 hours also supports metro availability.
- Both the primary cluster and the replication target cluster must fulfill the same minimum resource requirements for node configurations for the required RPO.
- Ensure that any new node or disk addition meets the minimum requirements provided in this section.
- Features such as Deduplication and RF3 may require additional memory depending on the DR schedules and other workload run on the cluster.

Note:

Deduplication is supported without additional memory if CVM is already configured with 32GB memory. However, RF3 would require additional 8GB memory

Note: The default minimum snapshot or recovery point frequency with the default Foundation configuration is 6 hours. To increase the snapshot frequency, you must also modify the Foundation configuration (SSD and CVM) accordingly. For example, an all-flash setup with a capacity between 48 TB to 92 TB has the default snapshot interval is 6 hours. If you want to increase the snapshot interval to 1 hour, you must modify the default Foundation configuration to:

- 14 vCPUs for CVM
- 40 GB for CVM

The table lists the supported frequency for the snapshots across various hardware configurations.

Table 1: Snapshot Frequency

Type of disk	Capacity per node	Minimum snapshot frequency	Foundation Configuration - SSD and CVM requirements
Hybrid	Total HDD capacity per node: 32 TB or lower. And Total Capacity Per Node: (HDD + SSD) of 40 TB or lower.	• NearSync • Async (Hourly)	No change required —existing Foundation defaults. • 2 x SSDs • Each SSD must be minimum 1.2 TB or more for NearSync.
	Total HDD capacity per node: Between 32-64 TB. And Total Capacity Per Node: (HDD + SSD) Between 40-92 TB. Up to 32 TB SSD (4 x 7.68 TB SSDs)	• NearSync • Async (Hourly)	Modify Foundation configurations to minimum: • 4 x SSDs • Each SSD must be minimum 1.2 TB or more for NearSync. • 14 vCPU for CVM • 40 GB for CVM
	Total HDD capacity per node: Between 32-64 TB. And Total Capacity Per Node: (HDD + SSD) Between 40-92 TB Up to 32 TB SSD (4 x 7.68 TB SSDs)	Async (Every 6 hours)	No change required —Default Foundation configuration.
	Total HDD capacity per node: Between 64-80 TB. And Total Capacity Per Node: (HDD + SSD) 96 TB or lower.	Async (Every 6 hours)	No change required —Default Foundation configuration.

Type of disk	Capacity per node	Minimum snapshot frequency	Foundation Configuration - SSD and CVM requirements
	Total HDD capacity per node: Greater than 80 TB. And Total Capacity Per Node: (HDD + SSD) Between 96-136 TB. Maximum 136 TB.	Async (Every 6 hours)	Modify Foundation configurations to minimum: • 12 vCPU for CVM • 36 GB for CVM
All Flash	Total Capacity Per Node: 48 TB or lower	• NearSync • Async (Hourly)	No change required —Default Foundation configuration. Each SSD must be minimum 1.2 TB or more for NearSync.
	Total Capacity Per Node: Between 48-92 TB	• NearSync • Async (Hourly)	Modify Foundation configurations to minimum: • 14 vCPU for CVM • Each SSD must be minimum 1.2 TB or more for NearSync. • 40 GB for CVM
	Total Capacity Per Node: Between 48-92 TB	Async (Every 6 hours)	No change required —Default Foundation configuration.
	Total Capacity Per Node: Greater than 92 TB	Async (Every 6 hours)	Modify Foundation configurations to minimum: • 12 vCPU for CVM • 36 GB for CVM

Conditions for Application-consistent Snapshots

This section provides the conditions and limitations for application-consistent snapshots that you generate during the replication.

Requirements for Application-consistent Snapshots

- Before taking an application-consistent snapshots, consider the workload type of your guest VM.

  Consider the difference between application-consistent snapshots and crash-consistent snapshots, and the workloads that support each type of snapshot. Ascertain the type of snapshot that your workload supports.

  Crash-consistent Snapshots: A snapshot is crash-consistent if it captures all of the data components (write order consistent) at the instant of the crash. VM snapshots are crash-consistent (by default), which means that the vDisks that the snapshot captures are consistent with a single point in time. Crash-consistent snapshots are more suited for non-database operating systems and applications which may not support quiescence (freezing) and un-quiescence (thawing) and such as file servers, DHCP servers, print servers.

  Application-consistent Snapshots: A snapshot is application-consistent if, in addition to capturing all of the data components (write order consistent) at the instant of the crash, the running applications have completed all their operations and flushed their buffers to disk (in other words, the application is quiesced). Application-consistent snapshots capture the same data as crash-consistent snapshots, with the addition of all data in memory and all transactions in process. Therefore, application-consistent snapshots may take longer to complete. Application-consistent snapshots are more suited for systems and applications that can be quiesced and un-quiesced, such as database operating systems and applications such as SQL, Oracle, and Exchange.
- Applications running in your guest VM must be able to quiesce I/O operations. For example, you can quiesce I/O operations for database applications and similar workload types.
- Ensure that Windows VMs do not have in-guest mounted VHDX disks. Nutanix does not support application-consistent snapshots on Windows VMs that have in-guest mounted VHDX disks.
- Configure one consistency group for each VM. That is, each consistency group must consist of one VM. If the consistency group includes more than one VM, the system ignores the application-consistent snapshot feature.
- Do not delete the .snapshot folder in the vCenter.
- For protection of VAAI clones with application consistency, place each VAAI clone in its own protection domain.
- The snapshot consistency (whether a snapshot is application-consistent or crash-consistent) depends whether you have selected **Create application consistent snapshot** check box while creation of the consistency group and creation of the schedules. The rules are as provided in the following tables.

Table 2: Snapshot Consistency

Consistency group setting	Schedule setting	Result
Checked	Unchecked	Crash-consistent snapshot
Checked	Checked	Application-consistent snapshot
Unchecked	Unchecked	Crash-consistent snapshot
Unchecked	Checked	Crash-consistent snapshot

Application-consistent Snapshots with Microsoft Volume Shadow Copy Service (VSS)

- To take application-consistent snapshots on Windows guest VMs, enable Microsoft VSS services.

  When you configure a protection domain and select **Use application consistent snapshots**, the Nutanix cluster transparently invokes the VSS (also known as Shadow copy or volume snapshot service). This option creates an application-consistent snapshot of a VM limited to consistency groups consisting of a single VM.

  > Note: This option is available for ESXi and AHV only; it is not available for Hyper-V. However, you can use a third-party product to invoke VSS for Hyper-V.

- To take application-consistent snapshots on guest VMs that use VSS, systems invoke Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent. VSS framework takes application-consistent snapshots without causing VM stuns (temporary unresponsive VMs).
- VSS framework enables third-party backup providers like Commvault and Rubrik to take application-consistent snapshots on Nutanix platform in a hypervisor-agnostic manner.
- The default and only backup type for VSS snapshots is VSS_BT_COPY (copy backup).

  Third party Backup products can choose between VSS_BT_FULL (full backup)and VSS_BT_COPY (copy backup) backup types.
- Guest VMs with any iSCSI attachments (LUNs) do not support Nutanix VSS snapshots. Nutanix VSS snapshots fail for such guest VMs.
- Do not take Nutanix enabled application-consistent snapshots while using any third-party backup provider enabled VSS snapshots (for example, Veeam).

NGT Support and Role in Application-consistent Snapshots

- Before taking an application-consistent snapshots, install and enable Nutanix Guest Tools (NGT) on your guest VM.

  For installing and enabling NGT, see *Nutanix Guest Tools* in the *Prism Web Console Guide*.

  For guest VMs running on ESXi, consider these points.

  - Install and enable NGT on guest VMs running on ESXi also. Application-consistent recovery points fail for EFI boot-enabled Windows 2019 VMs running on ESXi without installing NGT.
  - (vSphere) If you do not enable NGT and then try to take an application-consistent snapshot, the system creates a Nutanix native snapshot with a single vSphere host-based snapshot. The system deletes the vSphere host-based snapshot. If you enable NGT and then try to take application-consistent snapshot, the system directly captures a Nutanix native snapshot.

- The following table provides detailed information on whether a snapshot is application-consistent or not depending on the operating systems and hypervisors running in your environment.

  Note:

  - Installed and active means that the VM has the following:
    - NGT installed.
    - VSS capability enabled.
    - Powered on.
    - Actively communicating with the Controller VM.
  - The system captures hypervisor-based snapshots only when you have VMware Tools running on the VM and the VM does not have any independent disks attached to it. If these requirements are not met, the system captures crash-consistent snapshots.
  - The consistency group and schedule settings determine which type of snapshot the system captures, whether application-consistent or crash-consistent. See the snapshot consistency table to get more information on snapshot consistency.
  - The pre_freeze and post_thaw scripts can be Python or shell scripts or any executable files for Linux and batch files for Windows. These scripts contain commands specific to the particular applications that are running on the Linux or Windows VMs. Backup vendors like Commvault can provide these scripts. You can also write your own scripts. See the Pre_Freeze and Post_Thaw Scripts on page 23 for more information.
  - In AHV: If you receive a non-zero return code from the pre_freeze script, the system captures a non application-consistent snapshot and raises an alert on the Prism web console. If you receive a non-zero return code from the post_thaw script, the system attempts to capture an application-consistent snapshot once again. If that attempt fails, the system captures a non application-consistent snapshot, and raises an alert on the Prism web console.

Table 3: Supported Operating Systems (NGT Installed)

Operating system	Version
Windows	• Windows 2008 R2 through Windows 2019

Operating system	Version
Linux	• CentOS 6.5 through 6.9 and 7.0 through 7.3 • Red Hat Enterprise Linux (RHEL) 6.5 through 6.9 and 7.0 through 7.3. • Oracle Linux 6.5 and 7.0 • SUSE Linux Enterprise Server (SLES) 11 SP1 through 11 SP4 and 12 SP1 through 12 SP3 • Ubuntu 14.04

Table 4: Application-consistent Snapshots

Server	ESXi		AHV	
	NGT status	Result	NGT status	Result
Microsoft Windows Server edition	Installed and active. Also pre_freeze and post_thaw scripts are present.	Nutanix script-based VSS snapshots	Installed and active. Also pre_freeze and post_thaw scripts are present.	Nutanix script-based VSS snapshots
	Installed and active	Nutanix VSS-enabled snapshots.	Installed and active	Nutanix VSS-enabled snapshots
	Not enabled	Hypervisor-based application-consistent or crash-consistent snapshots.	Not enabled	Crash-consistent snapshots
Microsoft Windows Client edition	Installed and active. Also pre_freeze and post_thaw scripts are present.	Nutanix script-based VSS snapshots	Installed and active. Also pre_freeze and post_thaw scripts are present.	Nutanix script-based VSS snapshots
	Not enabled	Hypervisor-based snapshots or crash-consistent snapshots.	Not enabled	Crash-consistent snapshots
Linux VMs	Installed and active. Also pre_freeze and post_thaw scripts are present.	Nutanix script-based VSS snapshots	Installed and active. Also pre_freeze and post_thaw scripts are present.	Nutanix script-based VSS snapshots

Server	ESXi		AHV	
	NGT status	Result	NGT status	Result
	Not enabled	Hypervisor-based snapshots or crash-consistent snapshots.	Not enabled	Crash-consistent snapshots

Note: For information about the operating systems that support the AOS version you have deployed, see the Compatibility Matrix.

Pre_Freeze and Post_Thaw Scripts

- You can use Nutanix Guest Tools (NGT) to perform specific operations before and after a replication (snapshot) occurs. For more information about general NGT requirements, see *Nutanix Guest Tools Requirements and Limitations* in the *Prism Web Console Guide*.
- Any operation that the system must perform on any VM before replication or a snapshot capture, is a pre-freeze operation. Any operation that the system must perform on any VM after replication or a snapshot capture, is a post-thaw operation. You need scripts to perform pre-freeze and post-thaw operations using NGT tools.
- On a virtual machine, the pre-freeze script runs before the software snapshot creation. The post-thaw script runs after the software snapshot creation. For example, if a virtual machine hosts a database, you can enable a hot backup of the database before replication using a pre-freeze script. NGT calls this script prior to the snapshot freeze and quiesce operation.
- You can take application-consistent snapshots on NGT and VSS enabled VMs. However, some applications require more steps before or after the VSS operations to fully quiesce the VMs to an appropriate restore point or state in which the system can capture a snapshot. Such applications need pre-freeze and post-thaw scripts to run the necessary extra steps.
- The pre_freeze and post_thaw scripts could be Python or shell scripts or any executable files. These scripts contain commands and routines necessary to run specific operations on one or more applications that are running on the Linux or Windows VMs.
- Vendors such as *CommVault* provide such scripts. You can also write your own scripts.

Only the following applications support application-consistent snapshots without pre-freeze and post_thaw scripts:

- Microsoft SQL Server 2008, 2012, 2016, and 2019
- Microsoft Exchange 2010
- Microsoft Exchange 2013
- Microsoft Exchange 2016

The following sections cover the important considerations and samples for pre-freeze and post-thaw scripts.

Location

- For Windows VMs, create pre_freeze.bat and post_thaw.bat scripts under `system_drive:\Program Files\Nutanix\scripts\pre_freeze.bat` and `system_drive:\Program Files\Nutanix\scripts\post_thaw.bat` respectively. For example, if your system_drive is C, create these scripts at `C:`

`\Program Files\Nutanix\scripts\pre_freeze.bat` and `C:\Program Files\Nutanix\scripts\post_thaw.bat` respectively.

- For the Linux VMs on AOS 5.1 or later releases running NGT 1.1 version, you must create the pre_freeze `/usr/local/sbin/pre_freeze` and post_thaw `/usr/local/sbin/post_thaw` scripts.

  The scripts must have root ownership and root access with 700 permissions.

Requirements

- For Windows VMs, as an administrator you must have read and write/run permissions on the scripts.
- For Windows Server operating systems if scripts are present, and you run the pre_freeze script, then the system performs the VSS quiesce, and then the system runs the post_thaw script.
- For Linux VMs, the scripts must have root ownership and root access with 700 permissions.
- For completion of any operation before or after replication, you must have both the pre_freeze and post_thaw scripts for the operation.
- Timeout for both the scripts is 60 seconds.
- A script must return 0 to indicate a successful run. Any other value that the script returns, implies that the script execution failed. The necessary log entries are available in the NGT logs.
- Irrespective of whether the pre-freeze script execution is successful, the corresponding post-thaw script runs.

Sample Pre_Freeze and Post_Thaw Scripts

Sample `pre_freeze` and `post_thaw` scripts for MySQL on Linux are as follows.

> Note:
>
> The scripts provided in this section are samples only. These scripts must be modified to fit your deployment. These scripts require other scripts like quiesce.py and unquiesce.py that are provided by vendors such as Commvault (https://documentation.commvault.com). Use these scripts after you have modified them to work in your deployment.

```
#!/bin/sh
#pre_freeze-script
date >> '/scripts/pre_root.log'
echo -e "\n attempting to run pre_freeze script for MySQL as root user\n" >> /scripts/
pre_root.log
if [ "$(id -u)" -eq "0" ]; then
python '/scripts/quiesce.py' &
echo -e "\n executing query flush tables with read lock to quiesce the database\n" >> /scripts/
pre_freeze.log
echo -e "\n Database is in quiesce mode now\n" >> /scripts/pre_freeze.log
else
date >> '/scripts/pre_root.log'
echo -e "not root useri\n" >> '/scripts/pre_root.log'
fi
```

```
#!/bin/sh
#post_thaw-script
date >> '/scripts/post_root.log'
```

```
echo -e "\n attempting to run post_thaw script for MySQL as root user\n" >> /scripts/
post_root.log
if [ "$(id -u)" -eq "0" ]; then
python '/scripts/unquiesce.py'
else
date >> '/scripts/post_root.log'
echo -e "not root useri\n" >> '/scripts/post_root.log'
fi
```

Sample **pre_freeze** and **post_thaw** for Windows are as follows.

```
@echo off
echo Running pre_freeze script >C:\Progra~1\Nutanix\script\pre_freeze_log.txt
```

```
@echo off
echo Running post_thaw script >C:\Progra~1\Nutanix\script\post_thaw_log.txt
```

> Note:
>
> If any of these scripts prints excessive output to the console session, the script could freeze and not exit. Therefore, consider the following best practices when you create or update these scripts:
>
> - Always add **@echo off** to your scripts.
> - Redirect all the script output to a log file.

VStore Data Protection

VStore refers to a separate mount point that is created with its own NFS namespace in a storage container.

VStore Protection Domain

The namespace of a VStore is mapped to a protection domain when you protect the VStore. Such a protection domain is termed as VStore protection domain.

A VStore protection domain allows you to protect the VStore. You can also unprotect the protected VStore. Files in a protected VStore are replicated to a remote site at a defined frequency. You can recover these files from the remote site in the event of a disaster.

You can configure a NearSync replication schedule of as less as 20 seconds to protect a VStore. See Configuring Data Protection with NearSync Replication (20-59 seconds RPO) for VStore Protection Domains on page 70 for more information.

VStore Mapping for Data Protection

Nutanix recommends mapping the VStore container in the primary site to a container in the remote site. If a container containing a protected VStore in the primary site is not mapped to a specific container in the remote site, the container in the primary site is automatically mapped to the default (Self-service) container in the remote site.

If you need to reconfigure a VStore mapping for a protected VStore (for example, you add, remove, or modify the sVStore mappings), then:

1. Unprotect the VStore first.
2. Reconfigure the VStore mapping.
3. Protect the VStore.

You can configure VStore mapping while configuring or modifying the Remote Site. See Remote Site Configuration on page 152 for more information.

Note: Changing a VStore mapping causes associated protection domains to initiate full replication of protected entities to the newly specified target container as if it were initial replication. A mapping change therefore results in the overconsumption of storage resources at the remote site. Contact Nutanix Support for help with cleaning up snapshots in the previously specified container.

Considerations for VStore Data Protection

Consider the following for VStore data protection:

- Use nCLI to list the VStore entities on the cluster and to protect or unprotect a VStore. For more information about VStore protection, see the **`vstore`** commands in *Nutanix Command - line Interface Reference* in the *Command Reference* for AOS.
- Add the VMs or entities you need to protect to the VStore protection domain that is created after you protect the VStore. You must use nCLI only to add or remove entities to or from a VStore protection domain.
- Unprotect the VStore before you add, remove or modify the VStore mappings applied to the VStore. Protect the VStore again after completing the VStore mapping changes.
- See the *Maximums for Protected VStore* table below. You can protect these maximum number of files using one VStore. NCC health check verifies if these thresholds are met. If exceeded, NCC raises a Warning alert indicating a possibility of performance degradation. A number of factors including VMware log file creation during processes can increase the number of files in the VStore.
- Self-service restore is not supported for vStore protection domains.
- To use VStore with VMware SRM, you need the Nutanix SRA version corresponding to the VMware SRM version.

Table 5: Maximums for Protected VStore

Maximum number of files in a protected VStore.	3600 Note: The system continues to create and replicate snapshots even after the number of files exceeds this limit. Warnings are triggered when this limit is exceeded.
Maximum number of files in a consistency group in a protected VStore.	1200 Note: The system cannot create snapshots if the number of files exceeds this limit.

2

DATA PROTECTION DASHBOARD

The Data Protection dashboard displays dynamically updated information about the data protection configuration in a cluster. To view the Data Protection dashboard, select **Data Protection** from the pull-down list on the far left of the main menu.

Menu Options

In addition to the main menu (see Main Menu Options in Cluster Management in the *Prism Web Console Guide*), the Data Protection screen includes a menu bar with the following options:

- The Data Protection dashboard allows you to select from two viewing modes:
  - Click the **Overview** button on the left to display data protection and recovery information in a summary view (see Data Protection Overview View on page 28).
  - Click the **Table** button to display data protection information in a tabular form. The table screen is further divided into protection domain and remote site views; click the **Async DR** or **Metro Availability** tabs to view protection domain information or the **Remote Site** tab to view remote site information (see Data Protection Table View on page 29).
- The Data Protection dashboard allows you to perform actions using the action buttons on the dashboard:
  - Click the **Remote Site** button on the right to configure a remote site (see Configuring a Remote Site (Physical Cluster) on page 152).
  - Click the **Protection Domain** button on the right and then select **Async DR** from the pull-down list to configure a protection domain for asynchronous data replications used to create backup snapshots (see Configuring Data Protection with Asynchronous Replication on page 49) or **Metro Availability** to configure a protection domain for a metro availability configuration (see Configuring a Protection Domain (Metro Availability) on page 90).
- The Data Protection dashboard allows you to go to the **Protection Domains** and the **Remote Sites** pages in the Table view. These pages list the protections Domains and the remote sites configured. Each page lists a maximum of 10 protection domains and remote sites respectively. If there are more than 10 items in the list, the left and right paging arrows appear on the right, along with the total count and the count for the current page.
- The Data Protection dashboard allows you to export the table content as JSON or CSV. In the Table view, you can export the table information to a file in either CSV or JSON format by clicking the gear icon on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.
- You can search the table content. In the Table view, you can search for specific content in the table by entering a string in the search field.

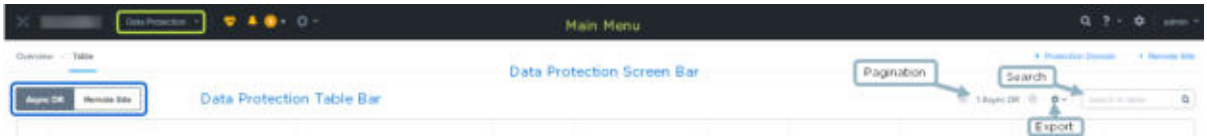

Figure 6: Data Protection Dashboard Menu

Data Protection Overview View

The Data Protection Overview view displays data protection-specific performance and usage statistics on the left plus the most recent data protection-specific alert and event messages on the right. 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.

The following figure is a sample view, and the table describes each field in this view.

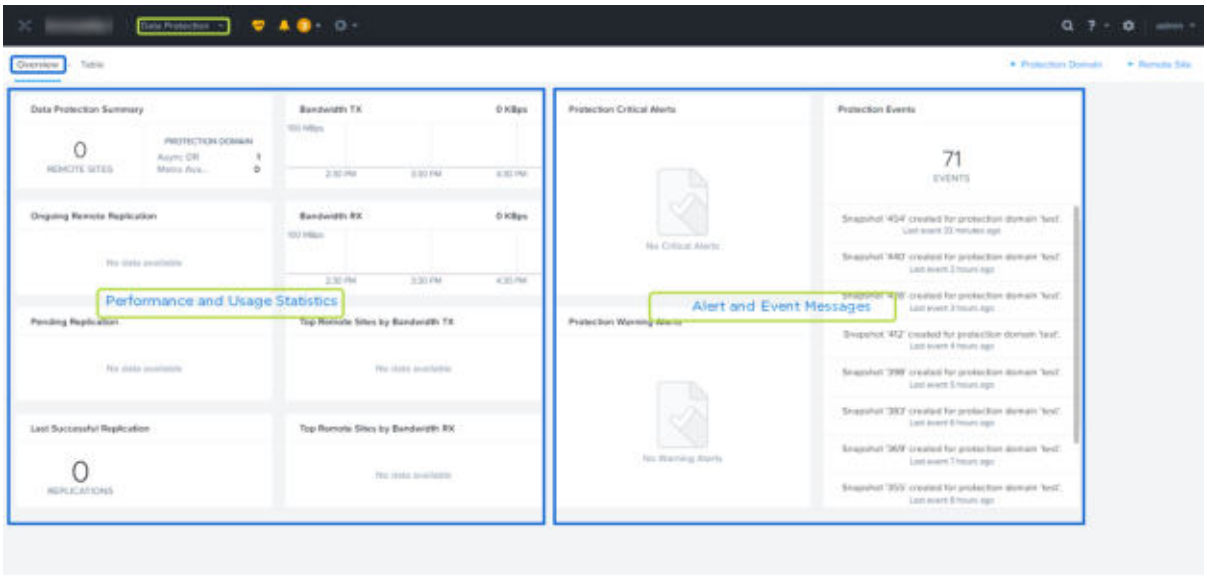

Figure 7: Data Protection Overview View

Table 6: Data Protection Overview View Fields

Name	Description
Data Protection Summary	Displays the current number of remote sites, active protection domains, and inactive protection domains.
Ongoing Remote Replication	Displays the number of ongoing outbound and inbound replications.
Pending Replication	Displays the number of pending replications.
Last Successful Replication	Displays the date and time of the last successful replication (and the remote site name) for each protection domain.

Name	Description
Bandwidth TX	Displays the amount of replication data (in GB or TB) transmitted from protection domains to remote sites in the cluster. 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 in Performance Monitoring in the *Prism Web Console Guide*).
Bandwidth RX	Displays the amount of recovery data (in GB or TB) received by protection domains from remote sites in the cluster.
Top Remote Sites by Bandwidth TX	Displays the names of the protection domains using the most replicated data storage. Clicking on a protection domain name displays the Table detail page for that protection domain (see Data Protection Table View on page 29).
Top Remote Sites by Bandwidth RX	Displays the names and amount of storage (in GB or TB) at the remote sites storing the largest amounts of replicated data.
Protection Critical Alerts	Displays the five most recent unresolved data protection-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the **view all alerts** button at the bottom of the list (see Alerts Dashboard in the *Prism Web Console Guide*).
Protection Warning Alerts	Displays the five most recent unresolved data protection-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the **view all alerts** button at the bottom of the list.
Protection Events	Displays the 10 most recent data protection-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the **view all events** button at the bottom of the list.

Data Protection Table View

The Data Protection table view displays information about remote sites and protection domains in a tabular form. Click the **Async DR** or **Metro Availability** tab in the screen menu bar to display protection domain information; click the **Remote Site** tab to display remote site information. The displayed information is dynamically updated to remain current.

> Note: The figures in this section are sample views for illustrative purposes only.

The Data Protection table view is divided into two sections:

- The top section is a table. Each row represents a single protection domain (configured for asynchronous data replication or metro availability) or remote site and includes basic information about that protection domain or remote site. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details content and set of tabs varies depending on what has been selected.

The following figure is a sample view of the Data Protection Table view.

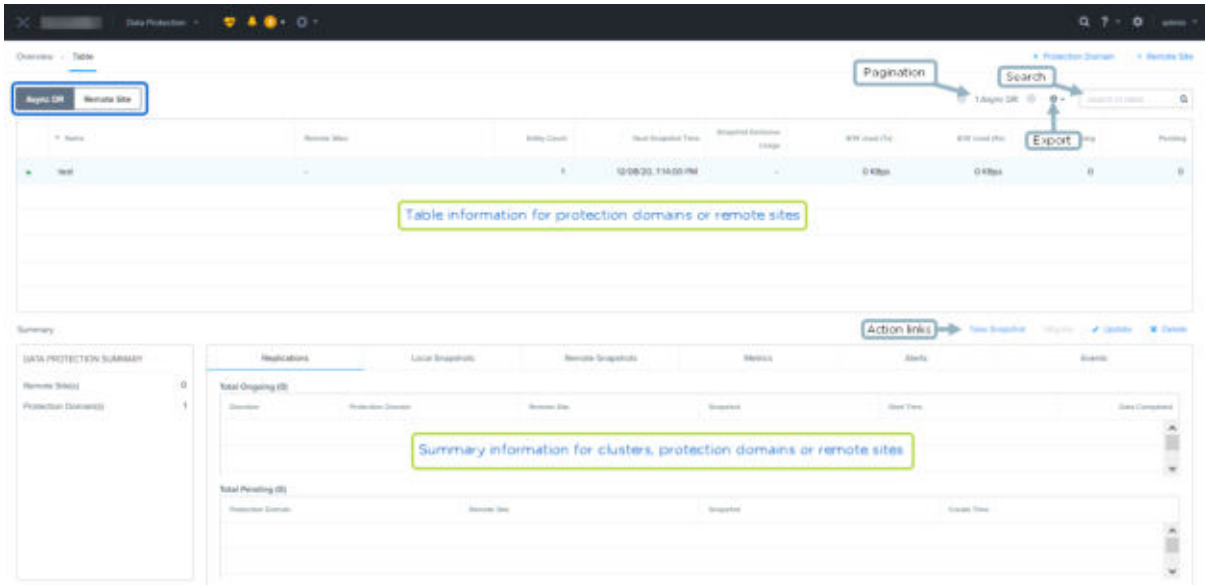

Figure 8: Data Protection Table View

Async DR Tab

Clicking the **Async DR** tab displays information about protection domains configured for asynchronous data replication in the cluster (see Configuring Data Protection with Asynchronous Replication on page 49). This type of protection domain consists of a defined group of virtual machines to be backed up (snapshots) locally on a cluster and optionally replicated to one or more remote sites.

- The table at the top of the screen displays information about all the configured protection domains, and the details column (lower left) displays additional information when a protection domain is selected in the table. The following table describes the fields in the protection domain table and detail column.
- When a protection domain is selected, **Summary:** ***protection_domain_name*** appears below the table, and action links relevant to that protection domain appear on the right of this line. The actions vary depending on the state of the protection domain and can include one or more of the following (see Modifying Asynchronous or NearSync Replication Configurations on page 73 for more information about these actions):
  - Click the **Take Snapshot** link to create a snapshot (point-in-time backup) of this protection domain.
  - Click the **Migrate** link to migrate this protection domain.
  - Click the **Update** link to update the settings for this protection domain.
  - Click the **Delete** link to delete this protection domain configuration.
- Eight tabs appear that display information about the selected protection domain (see following sections for details about each tab): **Replications**, **Entities**, **Schedules**, **Local Snapshots**, **Remote Snapshots**, **Metrics**, **Alerts**, **Events**.

Table 7: Async DR Protection Domain Table and Detail Fields

Parameter	Description	Values
Protection Domain Table Fields (upper screen)		
(color coded circle)	Displays a color coded circle that indicates whether protection is active or inactive for the protection domain. There are separate colors for regular protection domains (green=active, gray=inactive) and protection domains created by vStore protect (blue=active vStore, light blue=inactive vStore). The protection domains created by vStore protect are special ones with limited administrative operations (update only).	(green, gray, blue, or light blue circle)
Name	Displays the name of the protection domain.	(name)
Remote Sites	Displays the number of remote sites configured for this protection domain.	[0-unlimited]
Entity Count	Displays the number of entities in the protection domain.	[0-unlimited]
Next Snapshot Time	Displays the date and time of the next scheduled snapshot.	[date and time|No Schedule]
Snapshot Exclusive Usage	The amount of space that is reclaimed after the snapshot is deleted.	xxx [GB|TB]
B/W Used (Tx)	Displays the amount of replication data transmitted from the protection domain to a remote site.	xxx [GB|TB]
B/W Used (Rx)	Displays the amount of recovery data received by the protection domain from a remote site.	xxx [GB|TB]
Ongoing	Displays the number of ongoing replications to this remote site.	[0-unlimited]
Pending	Displays the number of pending replications scheduled for this remote site.	[0-unlimited]
Protection Domain Detail Fields (lower screen)		
Name	Displays the name of the protection domain.	(name)
Mode	Displays whether protection is active or inactive for the protection domain.	[Active|Inactive with color coded circle]
Next Snapshot Time	Displays the date and time of the next scheduled snapshot.	[date and time|No Schedule]
VM Count	Displays the number of protected VMs in this domain.	(quantity or ID number)
Volume Group Count	Displays the number of protected volume groups in this domain.	(quantity or ID number)

Parameter	Description	Values
File Count	Displays the number of protected files (outside of a VM) in this domain. One reason to save individual files outside of a VM is to store templates or other common files that might be used when configuring VMs.	(quantity or name)
Remote Site(s)	Displays the name of the remote sites configured for the domain.	(remote site names)

Metro Availability Tab

Clicking the **Metro Availability** tab displays information about protection domains configured for metro availability in the cluster (see Configuring a Protection Domain (Metro Availability) on page 90). This type of protection domain consists of a primary (active) storage container in the local cluster linked to a standby storage container of the same name on a remote site in which synchronous data replication to the remote site occurs when metro availability is enabled.

- The table at the top of the screen displays information about all the configured protection domains, and the details column (lower left) displays additional information when a protection domain is selected in the table. The following table describes the fields in the protection domain table and detail column.
- When a protection domain is selected, **Summary:** ***protection_domain_name*** appears below the table, and action links relevant to that protection domain appear on the right of this line. The actions vary depending on the state of the protection domain and can include one of the following:
  - Click the **Take Snapshot** link to take a snapshot of the protection domain locally and to one or more remote sites.
  - Click the **Update** link to modify the protection domain.
  - Click the **Activate** link to set a disabled protection domain to active status.
  - Click the **Disable** link to disable an active protection domain.
  - Click the **Promote** link to promote this protection domain from standby to active status. This option is available only when metro functionality is disabled on the active site. See the "Performing Site Maintenance" section in Conditions for Implementing Data Protection (Metro Availability) on page 83 for instructions on how to safely migrate VMs and promote a site from standby to active.
- Four tabs appear that display information about the selected protection domain (see following sections for details about each tab): **Local Snapshots**, **Metrics**, **Alerts**, **Events**.

Table 8: Metro Availability Protection Domain Table and Detail Fields

Parameter	Description	Values
Protection Domain Table Fields (upper screen)		
(color coded circle)	Displays a color coded circle that indicates whether the protection domain is active (dark green), standby (light green), or inactive (gray) in the metro availability configuration.	(dark green, light green, or gray circle)

Parameter	Description	Values
Name	Displays the name of the protection domain.	(name)
Role	Displays the role of this protection domain in the metro availability configuration: • **Active** means this is the active (primary) protection domain. • **Standby** means this is the standby protection domain.	[Active|Standby]
Metro Remote Site	Displays the name of the remote site containing the paired metro availability protection domain.	(remote site name)
Storage Container	Displays the name of the storage container configured for this protection domain.	(storage container name)
Status	Displays the status of this protection domain in the metro availability configuration: • **Enabled** means metro availability is enabled currently for this protection domain. • **Disabled** means metro availability is not enabled currently for this protection domain.	[Enabled|Disabled]
Failure Handling	Failure handling timeout range in seconds.	10 to 30 seconds
B/W Used (Tx)	Displays the amount of replication data transmitted from the protection domain to the remote site.	xxx [GB|TB]
B/W Used (Rx)	Displays the amount of replication data received by the protection domain from the remote site.	xxx [GB|TB]
Protection Domain Detail Fields (lower screen)		
Name	Displays the name of the protection domain.	(name)
Role	Displays the role of this protection domain in the metro availability configuration: • **Active** means this is the active (primary) protection domain. • **Standby** means this is the standby protection domain.	[Active|Standby with color coded circle]
Storage Container Name	Displays the name of the storage container configured for this protection domain.	(storage container name)
Remote Site	Displays the name of the remote site containing the paired metro availability storage container.	(remote site name)

Parameter	Description	Values
Status	Displays the status of this protection domain in the metro availability configuration: • `Enabled` means metro availability is enabled currently for this protection domain. • `Disabled` means metro availability is not enabled currently for this protection domain.	[Enabled|Disabled]
Failure Handling	Failure handling timeout range in seconds.	10 to 30 seconds

Remote Site Tab

Clicking the **Remote Site** tab displays information about remote sites available to the cluster (see Configuring a Remote Site (Physical Cluster) on page 152).

- The table at the top of the screen displays information about all the configured remote sites, and the details column (lower left) displays additional information when a remote site is selected in the table. The following table describes the fields in the remote site table and detail column.
- When a remote site is selected, **Summary:** ***`remote_site_name`*** appears below the table, and action links appear on the right of this line (see Modifying a Remote Site (Physical Cluster or Cloud) on page 160) for more information about these actions):
  - Click the **Test Connection** link to check whether the remote site is alive and accessible.
  - Click the **Update** link to update the settings for this remote site.
  - Click the **Delete** link to delete this remote site configuration.
- Five tabs appear that display information about the selected remote site (see following sections for details about each tab): **Replications**, **Remote Snapshots**, **Metrics**, **Alerts**, **Events**.

Table 9: Remote Site Table and Detail Fields

Parameter	Description	Values
Remote Site Table Fields (upper screen)		
Name	Displays the name of the remote site.	(name)
Remote Site Type	Displays the type of remote site.	[physical or cloud]
Remote Addresses	Displays the IP addresses of the remote site.	(IP addresses)
vStore Mappings	Displays the vStore map entries (local source storage container name and remote target storage container name) for the remote site.	(***`local_name:remote_name`***)

Parameter	Description	Values
Capabilities	Displays how the remote site can be used. A value of Backup means the remote site is a backup (replication) target only. A value of Disaster Recovery means the remote site is both a backup target and a source for dynamic recovery (see Restoring an Entity from a Protection Domain on page 179).	[Backup|Disaster Recovery]
Compress on Wire	Displays whether compression is enabled for transmission to this site.	[On|Off]
Use SSH Tunnel	Displays whether an SSH tunnel is enabled for secure transmission between the local cluster and the remote site.	[Yes|No]
B/W Limit (Tx)	Displays the bandwidth limit set for transmissions to the remote site.	xxx [KBps]
B/W Used (Tx)	Displays the amount of replication data transmitted from protection domains to the remote site.	xxx [GB|TB]
B/W Used (Rx)	Displays the amount of recovery data received by protection domains from the remote site.	xxx [GB|TB]
Remote Site Details Fields (lower screen)		
Name	Displays the name of the remote site.	(name)
Capabilities	Displays how the remote site can be used. A value of Backup means the remote site is a backup (replication) target only. A value of Disaster Recovery means the remote site is both a backup target and a source for dynamic recovery (see Restoring an Entity from a Protection Domain on page 179).	[Backup|Disaster Recovery]
Cluster	Displays the cluster ID number.	(ID number)
Remote Address	Displays the IP addresses of the remote site.	(IP addresses)
vStore Map	Displays the configured association between the local source storage container and the remote target storage container. The entry syntax is source_storage_container_name : target_storage_container_name.	(source_container : target_container)
Max Bandwidth	Displays the bandwidth limit set for transmissions to the remote site.	xxx [KBps]
Compression on Wire	Displays whether compression is enabled for transmission to this site.	[On|Off]
Proxy Enabled	Displays whether a proxy is enabled, which allows specified addresses to be used as a proxy to communicate with other Nutanix components on the remote site.	[Yes|No]

Parameter	Description	Values
Use SSH Tunnel	Displays whether an SSH tunnel is enabled for secure transmission between the local cluster and the remote site.	[Yes|No]

Cluster Summary Information

When a protection domain or remote site is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Data Protection Summary** column (on the left) includes two fields:
  - **Remote Site(s)**. Displays the number of remote sites configured in the cluster.
  - **protection domain(s)**. Displays the number of protection domains configured in the cluster.
- Six tabs appear that display cluster-wide information (see following sections for details about each tab): **Replications**, **Local Snapshots**, **Remote Snapshots**, **Metrics**, **Alerts**, **Events**.

Replications Tab

The Replications tab displays information in tabular form about ongoing and pending replications in the cluster. When a protection domain is selected, it also displays tabular information about successful replications for that protection domain. The following information is displayed:

- **Direction** (ongoing and successful tables only). Displays the direction (incoming or outgoing) of the replication, this is whether the replication is coming in from a remote site to the cluster or going out from the cluster to a remote site.
- **Protection Domain** (ongoing and pending tables only). Displays the name of the protection domain being replicated.
- **Remote Site**. Displays the names of the target remote sites for this replication.
- **Snapshot** [**ID**]. Displays the snapshot ID number assigned to this replication.
- **Start Time** (ongoing and successful tables only). Displays the time when the replication started.
- **Create Time** (pending table only). Displays the time when the pending replication is scheduled to start.
- **Time Remaining**. Displays the approximate time remaining to complete the replication.
- **End Time** (successful table only). Displays the time when the replication completed successfully.
- **Data Completed** (ongoing table only). Displays the amount of data replicated successfully as replication progresses (*xxx* KB|MB|TB).

Figure 9: Data Protection Table View: Replications Tab

Entities Tab

The entities tab, which appears only when a protection domain is selected, displays the following information about entities (VMs and volume groups) in the selected protection domain:

- **Name**. Displays the name of the entity.
- **Type**. Displays the type of entity.
- **Consistency Group**. Displays the name of the consistency group in which the entity is assigned.
- **Power State on Recovery**. Displays the power setting that will be started after recovering the entity.

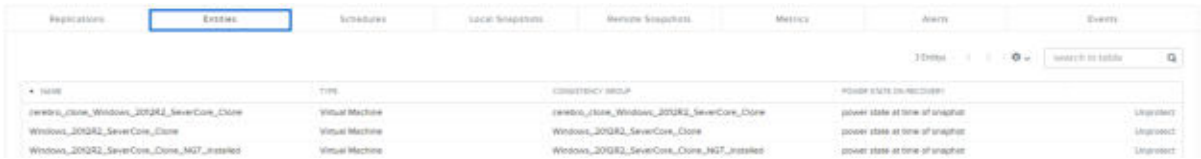

Figure 10: Data Protection Table View: VMs Tab

Schedules Tab

The Schedules tab, which appears only when a protection domain is selected, displays backup schedule information in tabular form. The following information is displayed:

- **Type**. Displays the schedule type, which indicates the interval metric (minutes, hours, days) used to set the schedule.
- **Repeat On**. Displays when to start the next backup run, such as every 60 minutes (when the type is set to minutes).
- **Start Date**. Displays the time when the first run is to start (or did start).
- **End Date**. Displays when the schedule is disabled, that is the latest date a backup run should start because of this schedule. A dash (-) indicates the backup schedule should continue indefinitely (no end date).
- **Retention Policy**. Displays how many snapshots from this schedule to keep locally and on remote sites (if configured).
- Click the pencil icon (far right) to update a schedule; click the X icon to delete a schedule. See the Modifying Asynchronous or NearSync Replication Configurations on page 73 for more information.

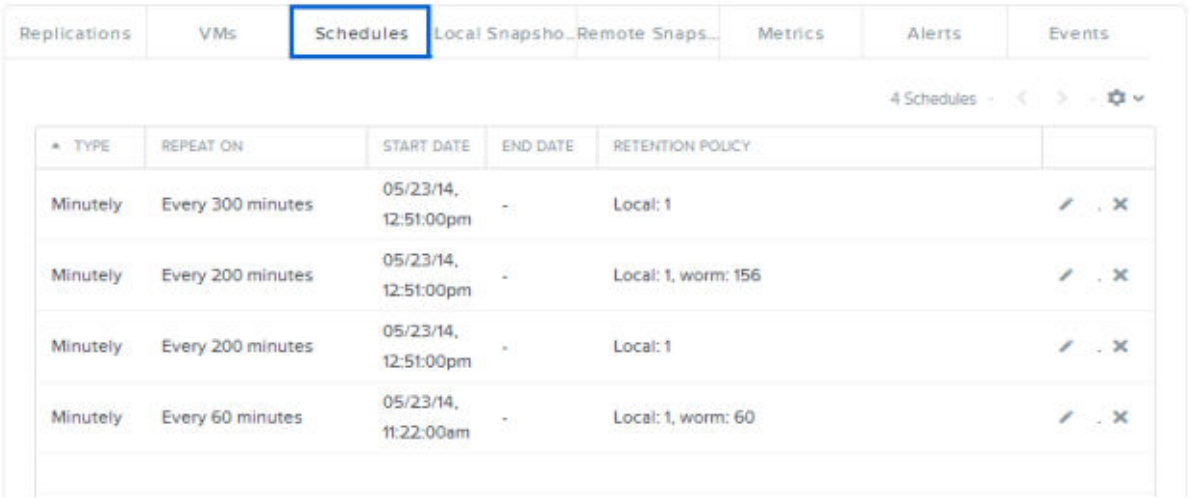

Figure 11: Data Protection Table View: Schedules Tab

Local Snapshots Tab

The Local Snapshots tab displays information in tabular form about snapshots stored locally. (This tab does not appear when a remote site is selected.) The following information is displayed:

- **ID**. Displays the snapshot identification number.
- **Create Time**. Displays the time when the snapshot was created.
- **Reclaimable Space**. Displays the total space that can be recovered after the snapshot is deleted.
- **Expiry Time**. Displays the date and time when the snapshots are about to expire.
- **VM Recovery**. Displays whether the snapshot that is present on the local cluster is recoverable or not. Click **Recovery Details** link for more information.
- Use the links in the far right column to do the following:
  - Click the **Details** link to display information about the VMs backed up in the snapshot.
  - Click the **Restore** link to restore one or more VMs from this snapshot (see Restoring an Entity from a Protection Domain on page 179).
  - Click the X icon to delete the snapshot.

    Note: You must delete the protection domain to delete all the snapshots in the protection domain. Deleting the schedules alone is not sufficient. If you do not delete the protection domain, but delete all the snapshots, a reference snapshot is kept to allow for delta replication in future.

Note: Select the **Include Expired** option to include the expired snapshots in the table.

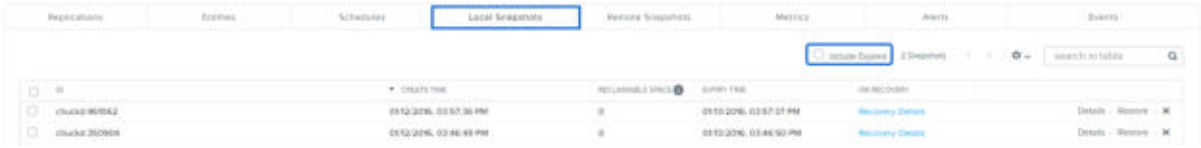

Figure 12: Data Protection Table View: Local Snapshots Tab

Remote Snapshots Tab

The Remote Snapshots tab displays information in tabular form about snapshots stored on remote sites. The following information is displayed:

- **ID**. Displays the snapshot identification number.
- **Create Time**. Displays the time when the snapshot was created.
- **Location**. Displays the name of the remote site on which the snapshot is stored.
- **Expiry Time**. Displays the date and time when the snapshots are about to expire.
- **VM Recovery**. Displays whether the snapshot that is present on the remote cluster is recoverable on the local cluster or not. Click **Recovery Details** link for more information.
- Use the links in the far right column to do the following:
  - Click the **Details** link to display information about the VMs backed up in the snapshot.
  - Click the **Retrieve** link to copy (replicate) the snapshot from the remote site to the local cluster.
  - Click the X icon to delete the snapshot.

    Note: You must delete the protection domain to delete all the snapshots in the protection domain. Deleting the schedules alone is not sufficient. If you do not delete the protection domain, but delete all the snapshots, a reference snapshot is kept to allow for delta replication in future.

Note:

- Select the **Include Expired** option to include the expired snapshots in the table.
- To find the reclaimable space of the remote snapshot in a protection domain, sign in to Prism web console for the remote site, navigate to the **Data Protection** dashboard, select the relevant protection domain and go to **Local Snapshot** tab. The **Reclaimable Space** column displays the total space used by the snapshot on the remote site.

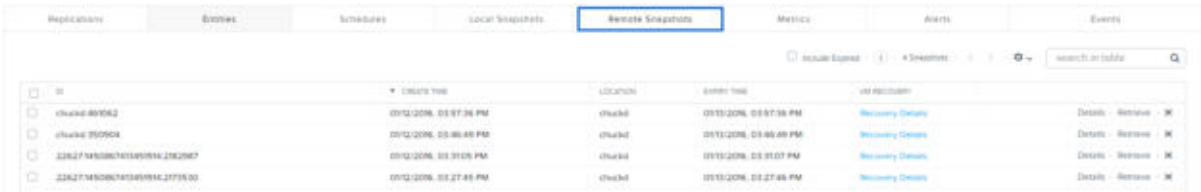

Figure 13: Data Protection Table View: Remote Snapshots Tab

Metrics Tab

The Metrics tab displays two graphs of replication data transmissions during a rolling time interval:

- **Bandwidth Tx**. Displays the amount of replication data transmitted from protection domains in the cluster to remote sites.
- **Bandwidth Rx**. Displays the amount of recovery data received by protection domains in the cluster from remote sites.
- For more in depth analysis, you can add either of these charts to the analysis page by clicking the blue link in the upper right of the chart (see Analysis Dashboard in Performance Monitoring in the *Prism Web Console Guide*).

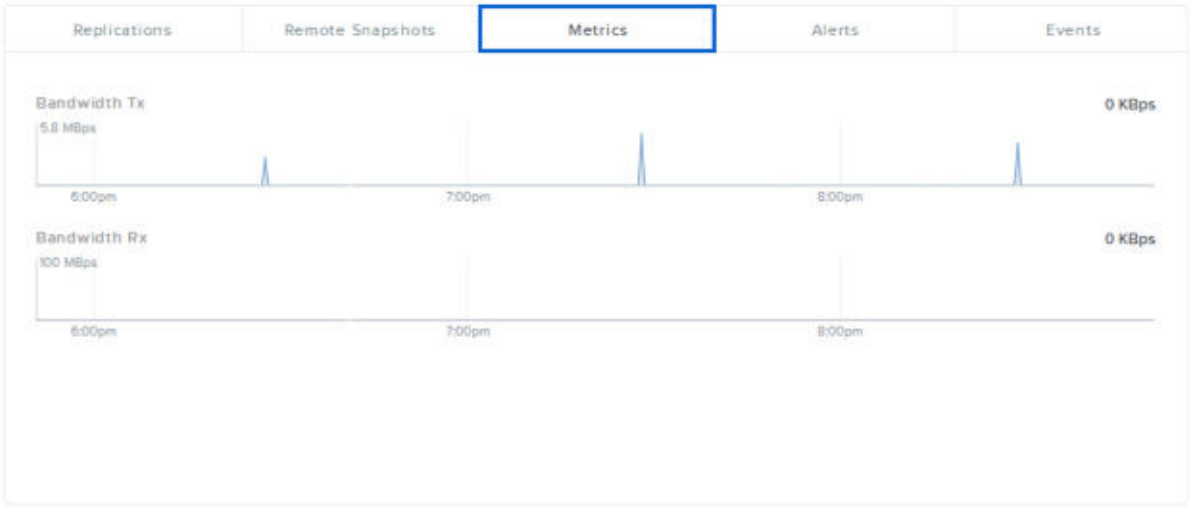

Figure 14: Data Protection Table View: Metrics Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about protection domains or remote sites in the same form as the Alerts page (see Alert Messages View in the *Prism Web Console Guide*). You can click the **drop-down** button to retrieve the alerts according to severity, resolution, and duration.

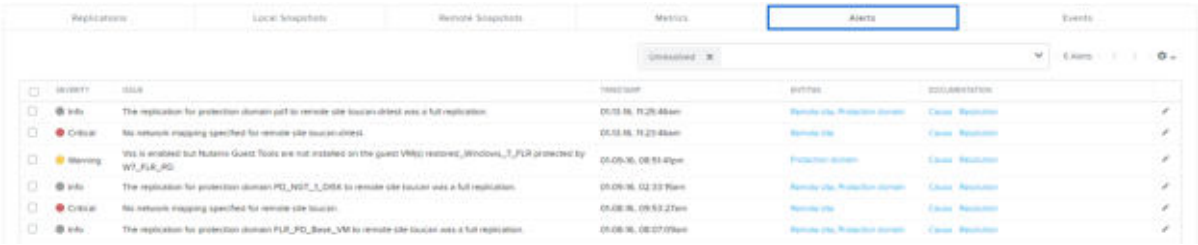

Figure 15: Data Protection Table View: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about protection domains or remote sites in the same form as the Events page (see Event Messages View in the *Prism Web Console Guide*). Click the **Include Acknowledged** button to also display acknowledged events.

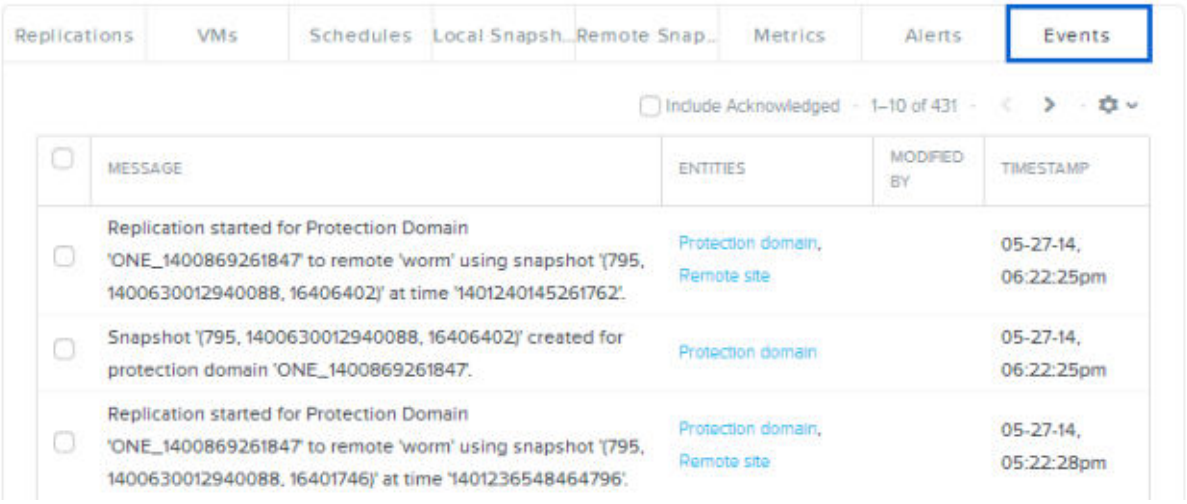

Figure 16: Data Protection Table View: Events Tab

3

DATA PROTECTION WITH ASYNC DR AND NEARSYNC REPLICATION

Data Protection with Asynchronous Replication (One-hour or Greater RPO)

Configure asynchronous replication schedules for backing up a group of entities (virtual machines, volume groups, or storage containers) locally to the Nutanix cluster. You can also optionally configure the replication to one or more remote sites. The replication occurs between two or more on-prem datacenters or clusters. The source cluster at the primary site and the target cluster at a remote site connected over the Internet.

> Note: Nutanix provides multiple disaster recovery solutions to protect your environment. See Nutanix DR Solutions Workflow on page 13 for the detailed representation of the workflow.

Aysnc DR: RPO of 60 minutes or greater

Asynchronous replication in Prism Element Web Console has recovery schedules with an RPO greater than or equal to 60 minutes are referred to only as asynchronous schedules and use full snapshots. Configure schedules in terms of hours, days, weeks, and months.

NearSync: RPO Between 1-15 minutes

See the Data Protection with NearSync Replication (1–15 minutes RPO) on page 57 for more information.

RPO Between 15-59 minutes

Nutanix does not provide any data protection method with this RPO interval.

To configure a protection domain for:

- Async DR replication, see Configuring Data Protection with Asynchronous Replication on page 49.
- NearSync replication, see Configuring Data Protection with NearSync Replication on page 64.

Conditions for Implementing Asynchronous Replication

This section provides the conditions necessary for the implementation of protection domain-based asynchronous replication.

General Conditions

- Nutanix supports replications between the all the latest supported LTS and STS released AOS versions. To check the list of the latest supported AOS versions, see KB-5505. Check

the EOL document to determine if the AOS versions currently running on the clusters are EOL.

Upgrade the version to the next available supported LTS/STS release. Before you upgrade the AOS, check the Upgrade Paths page to determine if an upgrade path is supported.

> Note: If both clusters have different AOS versions that are EOL, upgrade the cluster with lower AOS version to match the cluster with higher AOS version and then perform the upgrade to the next supported LTS version as advised above.

For example, the clusters in a metro configuration are running AOS versions 5.5.x and 5.10.x respectively. Upgrade the cluster on 5.5.x to 5.10.x. After both the clusters are on 5.10.x, proceed to upgrade each cluster to 5.15.x (supported LTS). Once both clusters are on 5.15.x you can upgrade the clusters to 5.20.x or later.

Nutanix recommends that both the primary and the replication clusters or sites run the same AOS version.

- All local Controller VMs must be able to communicate with all remote Controller VMs for successful replication.

  You can enable network segmentation for disaster recovery on a stretched Layer 2 network that does not have a gateway. For more information, see *Segmenting a Stretched L2 Network for Disaster Recovery* in the *Securing traffic through Network Segmentation* section of the *Security Guide*.

- If you need encryption of the data replication, set up an encrypted site-to-site tunnel. Specify the tunnel IP address when you create the remote site (by specifying the tunnel IP address in the Addresses parameter).
- If bandwidth between sites is limited, set a limit on the bandwidth that replication uses (by specifying the Maximum Bandwidth parameter).
- A consistency group is a subset of the entities in a protection domain. Consistency group typically does not exceed more than 20 entities.
- Multiple protection domains cannot protect the same VM. Configure one protection domain to protect a VM.
- One-time snapshots have no expiry time and, therefore, Nutanix recommends you specify the retention time when you create one-time snapshots.
- Do not include the source and the destination cluster in the same datacenter because the VMs could get deleted from both the source and destination clusters post the migration process.
- Do not make any changes in the VM or its configuration on the source cluster after the VM is powered off during the migration of the protection domain. Such changes are lost.
- Do not have VMs with the same name on the primary and the secondary clusters. Otherwise, it may affect the recovery procedures.
- If a VM has ongoing tasks, the snapshot operation fails after six retries of the protection domain that contains that VM. The snapshot operation succeeds only if within these six retries the ongoing tasks on the VM are completed.

- VM snapshots of VMs that also have VMware snapshots could be corrupted under certain unknown circumstances, if you use Prism to take storage snapshots of the protect domain, which includes such VMs. To avoid this issue, apply the following recommended practices:

  - Ensure that a VMware snapshot (through vCenter) does not exist when a storage snapshot is taken (by Prism).
  - Perform storage snapshots during times when VM snapshots are less likely to be taken.

    Schedule Nutanix storage snapshots to have limited or no overlap with manual or backup initiated VM snapshots.
  - Restore VMs only from a storage snapshot that was taken at a time when no VMware snapshots were taken or when the VMs were powered off.
- If you are deploying an intrusion prevention system (IPS) appliances or software, consider whether any configured filters or other network monitoring aides could block packet transfer during replication operations.

  Add the IP address of any appliances or systems running such appliances or software to the whitelist as described in *Configuring a Filesystem Whitelist* in the *Prism Web Console Guide*.
- If a VM running NGT is restored from a hypervisor-based snapshot that was created before NGT was installed on the VM, the VM is restored without NGT. Therefore, any native snapshots of the VM created after restoration is based on stale NGT information.

  To avoid this issue, after restoration, either reinstall or disable NGT on the VM.

Conditions for ESXi and VMware Components

The following conditions apply to clusters with ESXi and other VMware components:

- Nutanix cluster nodes to vSphere cluster mapping: The Nutanix data distribution architecture does not support the mapping of the nodes of a Nutanix cluster to multiple vSphere clusters.

  > Note: If the nodes of a Nutanix cluster are not in the same vSphere cluster and are split into multiple vSphere clusters, disaster recovery operations like migration and recovery fail.

  Ensure that all the nodes in the Nutanix cluster are in the same vSphere (ESXi) host cluster and that the network is available on all the nodes in the Nutanix cluster at the primary or remote site. Any configuration using Nutanix DR where there is more than one ESXi cluster on top of a single Nutanix cluster is not supported.
- Nutanix recommends that any ESXi clusters used in a disaster recovery configuration be registered with vCenter.

  This condition also applies to a cross-hypervisor disaster recovery (CHDR) configuration.
- Nutanix recommends adding files associated with the VM gold image to a protection domain, to protect VMs created by VMware View Composer or Citrix XenDesktop.

  Use the nCLI command to create the protection domain. For example, to protect the `replica-ABC.vmdk` file in a protection domain named `vmware1-pd` and a consistency group named `vmware1-cg`:

  ```
  ncli> protection-domain protect \
  ```

```
files=/container1/view-gold-image/replica-ABC.vmdk name=vmware1-pd cg-name=vmware1-cg
```

> Note:
>
> Disable the VMware View pool from View Composer.
>
> For more information about disabling VMware View pool from View Composer, see *VMware View documentation*.

- Asynchronous replication supports multiple ESXi VMs with multi-writer sharing enabled only if these VMs are protected in a single consistency group within the protection domain.
- An ESXi VM with GPU attached cannot be added to any protection domain.
- Nutanix recommends that you do not configure the storage containers with the same name on both the sites, if you are using the same vCenter Server to manage both the primary and remote sites.
- Nutanix recommends that you do not use the VMware level encryption that is used to encrypt the existing virtual machine or virtual disk along with Async DR configuration. The use of such encryption with Async DR is not supported.
- Ensure that you configure forward (DNS A) and reverse (DNS PTR) DNS entries for VM migration (as part of data replication) to succeed in an ESXi hypervisor environment.

  Configure these entries for each ESXi management host, on the DNS servers that are used by the Nutanix cluster.

  The **Hardware** page on the web console shows the host name and hypervisor IP address for each management host. The nCLI command `ncli host ls` also lists the IP address for each hypervisor.
-
- Replacing a vDisk in a VM that is currently in a protection domain with some existing snapshots does not cause problems when taking subsequent snapshots. Therefore, you do not need to change the protection or migration plan after changing a vDisk.
- If the local and remote sites are running ESXi hypervisor and are registered with the vCenter Server, the MAC address retention after a failover works as follows:
  - If you configure a static MAC address for a virtual network adapter on the local site, the adapter retains its MAC address after the protection domain is activated at the remote site and the VM is powered on.
  - If you configure automatic MAC address assignment for the virtual network adapter on the local site, and both the sites (local and remote) are registered with the same vCenter Server, the MAC address changes when the VM is registered at the remote site.
  - If you configure automatic MAC address assignment for the virtual network card on the local site, and the local and remote sites are registered with separate vCenter Servers, the MAC address changes only after the VM is powered on at the remote site.

Conditions for Hyper-V

The following conditions apply to Hyper-V:

- Nutanix recommends creating a VM in their unique folders instead of using a default folder. If you use a default folder to create the VMs, you cannot protect these VMs.
- Do not reuse the (same) path-prefix for the same VMs in a protection domain.

- VMs configured with Hyper-V host snapshots do not support Nutanix native snapshots and replication.

  Configuring such VMs with Nutanix native snapshot or replication can result in unexpected snapshot behavior.

VM Registration Behavior for DR Operations Between Windows Server 2012 R2 and Windows Server 2016

The following matrix describes how VMs are registered on the remote site if your data protection configuration involves replicating and migrating VMs between Windows Server 2012 R2 and Windows Server 2016:

Table 10: VM Registration Behavior

	Replication	Migration
2012 R2 to 2016	• Replication works as expected. • VMs created from the replicated snapshots are registered with the 2012 R2 VM version, which is 5.0.	• Migration works as expected. • Migrated VMs are registered with the 2012 R2 VM version, which is 5.0.
2016 to 2012 R2	• Replication works as expected. • VM registration fails because a 2012 R2 host cannot register VMs with the 2016 VM version, which is 8.0.	A warning indicating version incompatibility is displayed, but you can complete the migration process. If you continue to migrate the VMs, the protection domain is made active on the remote site. However, VM registrations fail.

Limitations of Data Protection with Asynchronous Replication

General Limitations

Nutanix asynchronous replication:

- Nutanix cluster nodes to vSphere cluster mapping: The Nutanix data distribution architecture does not support the mapping of the nodes of a Nutanix cluster to multiple vSphere clusters.

  > Note: If the nodes of a Nutanix cluster are not in the same vSphere cluster and are split into multiple vSphere clusters, disaster recovery operations like migration and recovery fail.

  Ensure that all the nodes in the Nutanix cluster are in the same vSphere (ESXi) host cluster and that the network is available on all the nodes in the Nutanix cluster at the primary or remote site. Any configuration using Nutanix DR where there is more than one ESXi cluster on top of a single Nutanix cluster is not supported.

- Supports configuration maximums as provided on the Nutanix Configuration Maximums page.

- Recommends that each application that constitutes a set of entities is protected by a unique protection domain.
- Supports one snapshot per hour as the shortest possible snapshot frequency for asynchronous replication.
- Does not support snapshots of the entire file systems or storage containers.
- Supports a VM to be a part of a protection domain only if the VM is entirely on a Nutanix datastore and not on an external storage.
- Does not support running multiple third-party backups in Parallel on the same set of VMs. The system becomes unstable.
- Supports encryption for the network connection over which data replication between sites takes place.
- Does not support VMX editing when you restore a VM. As a result, the characteristics of the restored VM such as MAC addresses may be in conflict with other VMs in the cluster.

Other limitations of asynchronous data replication are as follows.

- The replication speed reduces when deduplication is enabled on storage containers containing protected VMs.
- Consistency groups cannot define boot ordering.
- Do not delete an inactive protection domain that contains entities such as VMs.

  > CAUTION: Do not deactivate and then delete a protection domain that contains VMs. If you do so, the VMs in the protection domain are deleted.

  Before deleting a protection domain:

  - Activate the protection domain.
  - Unprotect all the entities.
  - Delete all the replication schedules and all the snapshots of the protection domain.

  For information about deleting a protection domain, see Deleting a Protection Domain on page 76.
- Some protection domain operations like snapshots may fail:
  - When some VMs do not appear in the web console.
  - In nCLI command results during rolling upgrades, planned Controller VM maintenance.
  - When hosts are down or unavailable.
- Minimize the number of snapshots retained (snapshot retention) on both the primary site to avoid metadata bloat.
- The following limitations apply to the inclusion of related entities in a protection domain.
  - If the number of entities in a consistency group exceeds the number provided in the Nutanix Configuration Maximums page, the protection of related entities fails.

    If you want to include more than the number provided in the Nutanix Configuration Maximums page in a protection domain, protect the entities in separate consistency groups within the same protection domain. If two related entities are in separate

consistency groups but in the same protection domain, their attachment configuration is included in the snapshots and restored during recovery.

- Snapshot creation and recovery of attachments are not supported if you have configured volume groups with the following.
  - Challenge-Handshake Authentication Protocol (CHAP). The iSCSI target secret is cleared after a volume group is restored from a snapshot.
  - IP addresses. If you attach volume groups to VMs by allowlisting the IP addresses of the VMs, entities are recovered, but their attachment configuration is not recovered. You must manually reattach the entities after recovery.

- For information about snapshot frequencies supported for node capacities, see Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro) on page 16.

Limitations Specific to vSphere Environment

- If a protection domain fails over or back from one ESXi cluster to another ESXi cluster, a thick-provisioned disk changes to a thin-provisioned disk.
- Nutanix native snapshots cannot be used to protect VMs on which VMware fault tolerance is enabled.
- Virtual machines connected to a vSphere Distributed Switch (dvSwitch) are not automatically connected to their port group after failover. Manually connect such VMs to their port group after failover or restore operation.

Limitations Specific to Hyper-V Environment

- A snapshot fails and raises an alert if:
  - Any VM files (for example, configuration, snapshots, virtual disks, or ISOs) are residing in non-Nutanix storage containers.
  - All virtual disks associated with a VM are located in different directories or folders. For the replication to succeed, all virtual disks associated with a VM must be located in a single directory or folder.
  - The main VM folder and snapshot folder for a VM are located in different directory/folder paths. That is, a snapshot folder is typically located under the folder path of the VM. The snapshot folder must be under the VM folder or the replication fails.
- The run-as account must be a domain account and must have local administrator privileges for the Nutanix hosts.

  The local administrator privilege may be a domain administrator account. When the Nutanix hosts are connected to the domain, the domain administrator accounts automatically takes administrator privileges on the host.

  If the run-as account (domain account) in System Center Virtual Machine Manager (SCVMM) is not a domain administrator account, you must manually add it to the list of local administrators on each host by running the server configuration tool (Sconfig).
- Nutanix does not support Hyper-V replica VM in the Async DR protection domain.
- The name of the Hyper-V virtual switches between primary and remote site must be the same, otherwise restore operation fails.

- If the base VM and the Hyper-V differencing disk are in the same protection domain, after migrating the protection domain to the secondary site, the in-place restore fails.

  See Restoration of Protected Entities on page 178 for information about in-place and out-of-place restoration.

- Any VMs created during hypervisor upgrade, including VMs created through disaster recovery operations, experience downtime if they are not configured for high availability and the corresponding node is yet to be upgraded. Nutanix recommends that you do not create any VMs, either manually or through disaster recovery operations, during hypervisor upgrade.
- After a VM is migrated, its entry is not removed from failover cluster manager on the source cluster.
- If a highly available VM is protected, the clones of the VM do not inherit the high availability configuration.
- A highly available VM is no longer highly available after it is migrated to a remote site unless an entry with the same VMID exists in failover cluster manager.
- During an in-place restore or failback operation, the high availability (HA) property of a VM is honored only if the entry is not cleaned up in failover cluster manager.
- The state of the high availability property of a VM is not captured when a snapshot is created. If the VMID entry for a VM is present in failover cluster manager when it is being restored from its snapshot, then the VM is HA protected.
- The owner node of a VM can change when a disaster recovery operation, such as in-place restoration or migration, is performed on the VM.

Limitations Specific to Volume Groups

- You cannot include volume groups in a protection domain that is configured for metro availability.
- You cannot include volume groups in a protected vStore.
- You cannot include volume groups in a consistency group that is configured for application-consistent snapshots.
- Replication and protection for a volume group fail if a remote site is running an AOS release that does not support volume groups or disaster recovery for volume groups.
- Specify the volume group UUIDs when configuring volume group protection in the nCLI. You cannot specify volume group names.

Configuring Data Protection with Asynchronous Replication

Before you begin

- Ensure that you have met the protection domain requirements for configuring Async DR before proceeding, see Conditions for Implementing Asynchronous Replication on page 42.
- See the *DR Best Practices Guide* at the following URL for information to set up disaster recovery in your environment: https://portal.nutanix.com/#/page/solutions .

About this task

Perform this procedure to configure a protection domain that supports backup snapshots for entities (selected VMs and volume groups) using asynchronous data replication .

Note:

- To create a protection domain that supports a metro availability configuration through synchronous data replication, see Configuring a Protection Domain (Metro Availability) on page 90.
- For Nutanix Files (Files), protection domain operations are not available during file server operations.
- Asynchronous data replication (Async DR) is supported between Nutanix NX/SX series and Dell XC series clusters and between Nutanix NX/SX series and Lenovo HX series clusters.
- Asynchronous and NearSync schedules can coexist in a protection domain. See Interoperability Between Async DR and NearSync in a Protection Domain on page 71.

To configure a standard (Async DR) protection domain, do the following:

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Protection Domain** button and select **Async DR** from the drop-down list.
   The Protection Domain window opens.

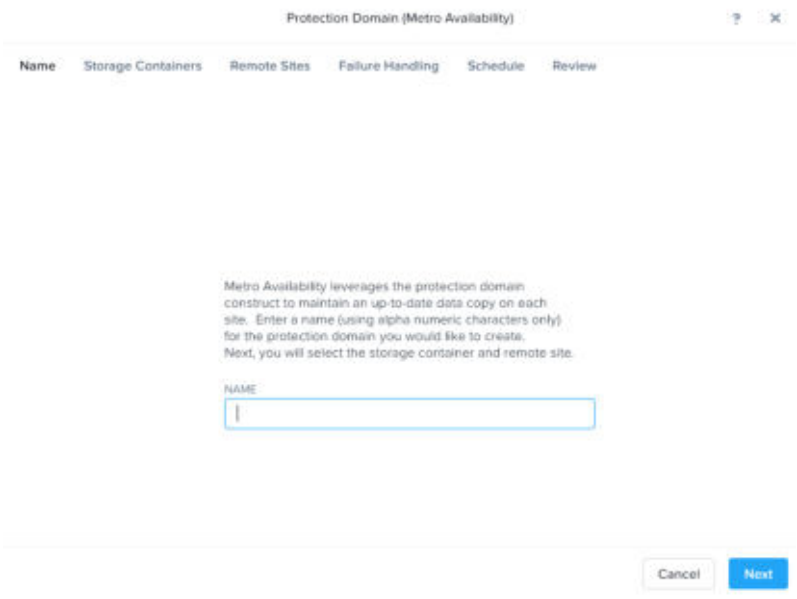

Figure 17: Protection Domain (Async DR): Create Screen

2. Enter a name for the protection domain on the first screen and then click **Create**.

   > Note: In your environment if you have more than one primary and secondary clusters, ensure that the names of the protection domain are unique for each pair of clusters.

   Clicking the **Create** button immediately creates the protection domain (even if you do not continue to the next step).

   > Note: This entity has the following naming restrictions:
   >
   > - The maximum length is 75 characters.
   > - Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

   The Entities display appears. This display is divided into two sections for unprotected entities (left side) and protected entities (right side).

3. Do the following in the indicated fields:

   a. Check the boxes next to the entities you want in a consistency group from the list of unprotected entities on the left. You can filter the list by entering an entity or host name in the **Filter By:** fields above the list.

      > Attention: The list includes the entities that are actively stored in a Nutanix container. If entities are stored on third-party storage or storage local to the host, the entities may not show in the list.

      A consistency group is a subset of the entities in a protection domain that are treated as a single group and backed up collectively in a snapshot. Protection domains can have multiple consistency groups. You can protect a maximum of 200 entities (VMs or volume groups) within any number of consistency groups (up to 200) in a single protection domain.

      Different protection domains cannot protect the same VM. A VM cannot be protected by different protection domains.

      > Note: If a VM is deployed through a specialized workflow like Calm, it may be assigned automatically to an Entity-centric protection domain. VMs in such a protection domain do not appear here, so it is impossible to add any of those VMs to an Async DR (or any) protection domain. To make such VMs available here, sign in to Prism Central, go to the VMs dashboard, select the VMs in question, and click the **Unprotect** action.

   b. If there are more entities you want in this consistency group that are already protected, check the boxes next to those entities in the list of protected entities on the right, and then click the **Unprotect Selected Entities** button at the bottom of the column to move them into the unprotected entities list.

   c. Select a consistency group (in the left column) for the checked entities. Click the **Use Entity Name** button to create a consistency group for each checked entity with the same name as that entity, click the **Use an existing CG** button and select a consistency group from the pull-down list to add the checked entities to that consistency group, or click the

Create a new CG button and enter a name in the field to create a consistency group with that name for all the checked entities.

> Note: Only one entity is allowed in a consistency group if the hypervisor is Hyper-V, so always click **Use Entity Name** for Hyper-V. (multiple entities are allowed in a consistency group for ESXi or AHV.)
>
> The **Use Entity Name** option is not available on the ESXi and AHV clusters.

d. [ESXi and AHV only] Check the **Use application consistent snapshots** box to ensure that application consistency is maintained in the snapshots.

   This option invokes VSS to create an application-consistent snapshot for a VM and is limited to consistency groups with just a single VM.

   > Note:
   >
   > - This option is available for ESXi and AHV only; it does not appear when the hypervisor is Hyper-V.
   > - For AHV, you must install NGT on the VM for which you are planning to take Nutanix VSS-based application-consistent snapshots. For more information, see Nutanix Guest Tools in the *Prism Web Console Guide*.
   > - For ESXi, you can take application-consistent snapshots without installing NGT on the VMs, but these snapshots are hypervisor-based and lead to VM stuns (temporarily unresponsive VMs). See Conditions for Application-consistent Snapshots on page 18.

   See Snapshots on page 15 section for information about snapshots in general and application-consistent snapshots.

   This box is ignored if there are multiple VMs in the consistency group. Therefore, check this box only if you have selected just one VM (and **Use VM Name** or **Create a new CG**) or selected multiple VMs and the **Use VM Name** button to create a separate consistency group for each of those VMs.

   You cannot create application-consistent snapshots of volume groups.

e. Click **Auto protect related entities** if you want to include all related entities in the consistency group.

   By default, the check box is selected. If you choose to leave the check box selected when you add entities to the protected entities list, all related entities are also added. All protected entities are included in the same consistency group and snapshots of the related entities are created for subsequent recovery. Also, the attachment configuration of all related entities is included in the snapshot so that you do not have to reattach entities again after using the snapshot for recovery.

   If you choose to clear the check box when adding entities for protection, only the entities that you select are included in the protection domain.

   If you choose to protect only a subset of the related entities, the selected entities are protected, but their attachment configuration is not captured, and you must manually

reattach the entities after recovery. If the entities you excluded, must be included later, update the protection domain and add those entities.

For a description of the term, related entities, see Terminology on page 7. For requirements and limitations to the inclusion of related entities and supporting operating systems, see Conditions for Implementing Asynchronous Replication on page 42.

f. Click the **Protect Selected Entities** button at the bottom of the column to move the selected entities into the **Protected** entities column.

g. Repeat this step as desired to create more consistency groups for the remaining unchecked entities (or as many of them as you want in the protection domain).

h. When you have created all the required consistency groups, click the **Next** button (lower right).

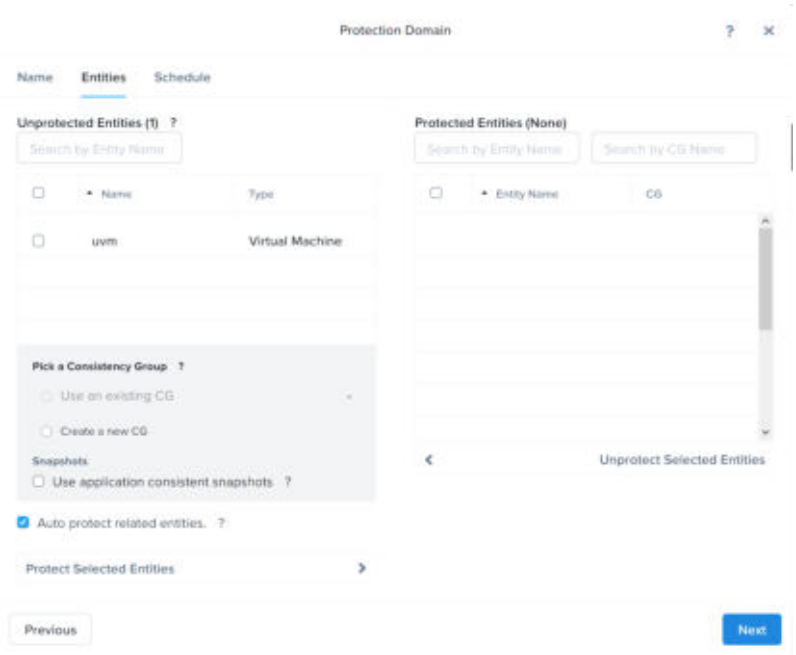

Figure 18: Protection Domain (Async DR): Virtual Machines Screen

The Schedule screen appears. There are no default backup schedules for a protection domain. Define all schedules required.

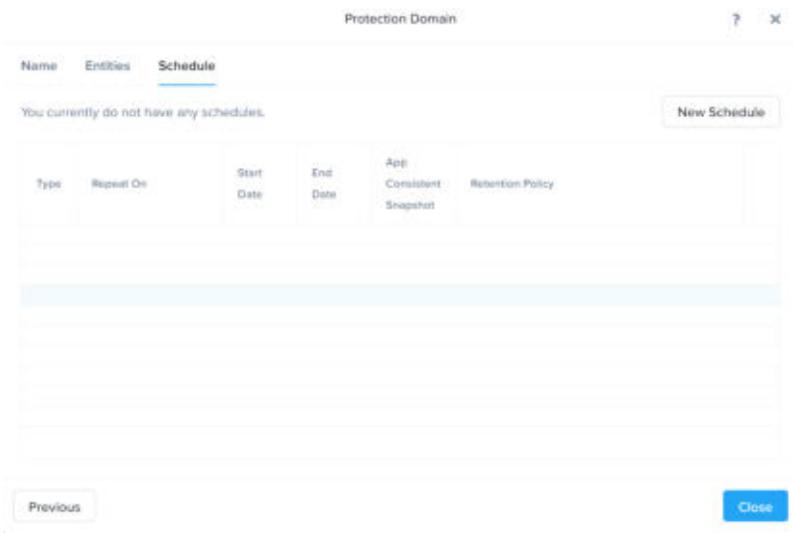

Figure 19: Protection Domain (Async DR): Schedule Screen (initial)

4. To create a schedule, click the **New Schedule** button to display the create schedule screen and then do the following in the indicated fields:

   a. **Repeat every ## [minutes|hours|days]**: Click the appropriate circle for minutes, hours, or days and then enter the desired number in the box for the scheduled time interval.

      The interval cannot be less than an hour, so the minutes value must be at least 60.

   b. **Occur [weekly|monthly]**: Select which days the schedule must run.

      - If you select weekly, check the boxes for the days of the week the schedule must run.
      - If you select monthly, enter one or more integers (in a comma-separated list) to indicate which days in the month to run the schedule. For example, to run the schedule on the 1st, 10th, and 20th days, enter "1,10,20".

   c. **Start on**: Enter the start date and time in the indicated fields.

      The default value is the current date and time. Enter a new date if you want to delay the schedule from starting immediately.

      > Note: A storage container-level protection domain requires a system metadata scan (Curator process) to populate the file list. Sometimes, this scan could take a few hours. Any snapshots started before the scan completes do not contain any data.

   d. **End on**: To specify an end date, check the box and then enter the end date and time in the indicated fields.

      The schedule does not have an end date by default, and it runs indefinitely unless you enter an end date here.

   e. **Retention Policy**: Enter how many intervals of snapshots that you want to retain.

      For example, if you enter 1 in the retention policy check box and select one of the following:

      - For **Repeat every minutes/hours/days** option: Total of only one snapshot is retained.
      - For **Repeat weekly/monthly** option: Total of 1 for each period/occurrence is retained. Suppose if you select occur weekly with Monday, Tuesday, Friday, a total of three snapshots are retained.

      - Enter the number to save locally in the **Local** line "**keep the last *xx* snapshots**" field. The default is 1.
      - A separate line appears for each configured remote site. To replicate to a remote site, check the remote site box and then enter the number of snapshots to save on that site in the appropriate field. Only previously configured remote sites appear in this list (see Configuring a Remote Site (Physical Cluster) on page 152).
      - The snapshots that are saved are equal to the value that you have entered in the **keep the last ## snapshots** field + 1. For example, if you have entered the value **keep the last ## snapshots** field as 20, a total of 21 snapshots are saved. When the next (22nd) snapshot is taken, the oldest snapshot is deleted and is replaced by the new snapshot.

   f. Select the **Create application consistent snapshot** check box to make the snapshots that are created as part of the new schedule application consistent.

      You can use this option to control the time when an application-consistent snapshot must be created. If you already have an application-consistent snapshot attached to a protection domain from a previous version, after upgrading to the new version the same schedules are marked as application consistent. For more information about application-consistent snapshot rules, best practices, and snapshot consistency (whether

a snapshot is application-consistent or crash-consistent), see Conditions for Implementing Asynchronous Replication on page 42.

g. When all the field entries are correct, click the **Create Schedule** button.

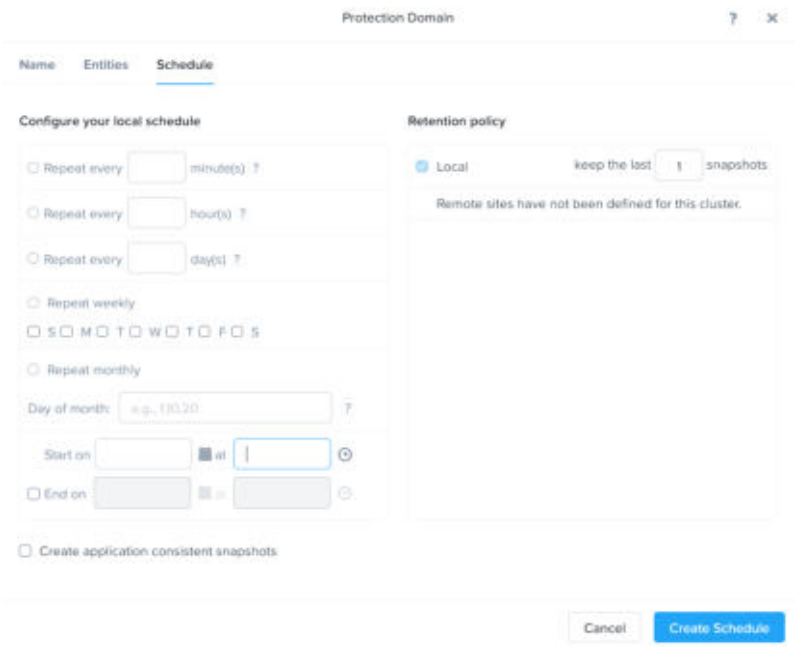

Figure 20: Protection Domain (Async DR): Schedule Screen (configure)

The Schedules screen returns.

5. Click the **New Schedule** button to create another schedule, update, or delete the schedule you created, or click the **Close** button to close the window.

Note:

- Replications can fall behind if too many schedules have the same start time. Therefore, it is recommended that start times be staggered across schedules.
- If you delete a schedule, the snapshots that were taken when the schedule was active does not get deleted.

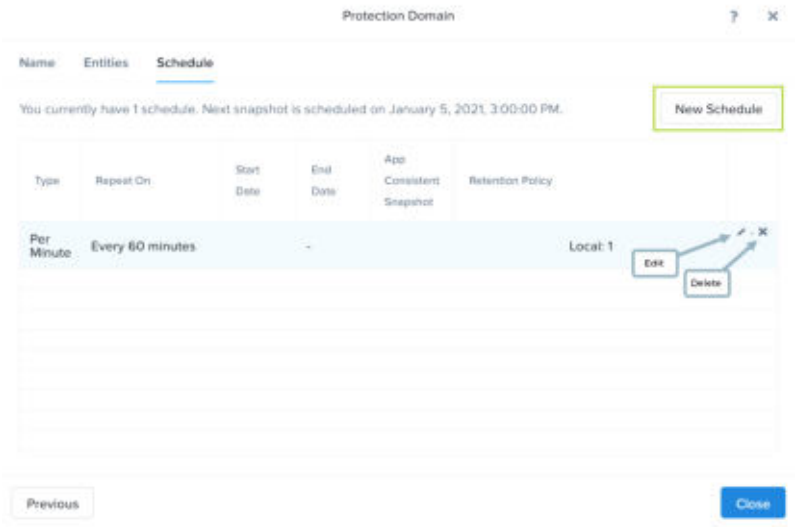

Figure 21: Protection Domain (Async DR): Schedule Screen (populated)

Results

Information about the new protection domain appears in the Data Protection dashboard (see Data Protection Table View on page 29).

See Remote Site Configuration on page 152 for information about configuring a remote site to store the snapshots.

Data Protection with NearSync Replication (1–15 minutes RPO)

RPO Between 1–15 minutes

NearSync replication, in Prism Element Web Console, has recovery schedules with an RPO between 1–15 minutes. Called NearSync schedules, they are also practically asynchronous schedules with 15–minute or lower RPOs. You can configure the schedules for even 1 minute.

Note: You cannot have a schedule with an interval of between 16-59 minutes. Nutanix provides multiple DR solutions to protect your environment. See Nutanix DR Solutions Workflow on page 13 for the detailed representation of the workflow.

The NearSync replication allows you to protect your data with an RPO of as low as 1 minute.

Note: You can configure NearSync replication with an RPO of as low as 20 seconds for VStore protection domains on an ESXi cluster. The schedule creates snapshots in every 20 seconds and replicates to another cluster that is registered to the same or different vCenter Server.

Some of the advantages of NearSync replication are as follows.

- Protecting mission-critical applications.
- Securing data with minimal loss in a disaster.
- Providing more granular control during the restore process.
- No latency or distance requirements for fully synchronous replication.
- Allowing resolution of a disaster event in minutes.

To implement the NearSync replication schedules, Nutanix introduces a technology called lightweight snapshots (LWS) to generate snapshots that continuously replicate incoming data generated by workloads running on the active cluster. The LWS snapshots are generated at the metadata level only. These snapshots are stored in the LWS store, which is allocated on the SSD tier. The LWS store is automatically allocated when you configure NearSync replication schedule for a protection domain.

> Note:
>
> - The maximum LWS store allocation for each node is 360 GB. For the hybrid systems, it is 7% of the SSD capacity on the node.

Transitioning in and out of NearSync

When you configure a NearSync schedule on a protection domain, the protection domain remains in an hourly schedule until the transition to NearSync is completed.

To transition a cluster into NearSync replication (or NearSync), perform initial seeding of the remote site with the data. The snapshots for the data seeding process are generated on an hourly basis. These snapshots are replicated to the remote site. Once the system determines that the snapshots containing the seeding data have replicated within a specified amount of time, the cluster is automatically transitioned into NearSync depending on the bandwidth and the change rate. After the cluster transitions into NearSync, you are able to see the configured NearSync snapshots in the web interface.

Following are the major characteristics of the process.

- Until the cluster transitions into NearSync, you only see the hourly snapshots generated in the web interface.
- If the cluster transitions out of this NearSync (for whatever reason), data seeding of the delta happens when you are again transitioned back into NearSync.
- Appropriate alerts are raised in the **Alerts** dashboard and the protection domain is transitioned back to the hourly schedule if the cluster transitions out of NearSync. The system continuously tries to get to the NearSync schedule that you have configured. If the transition is successful, the protection domain is automatically transitioned back into NearSync and alerts specific to this condition is also raised in the **Alerts** dashboard.

To transition the cluster out of NearSync, delete the NearSync schedule that you have configured.

If you are manually deleting the protection domain, perform the following actions depending on where you are deleting the protection domain .

On the Primary Site

1. Delete the NearSync schedule.
2. Delete all the snapshots.

3. Delete the protection domain .

On the Remote Site

1. Delete the NearSync schedule from the primary site.
2. Stop the replication by using the `pd pause-replication pd_name replication_ids` nCLI command on the primary site.

   Replace `pd_name` with the name of the protection domain. Replace `replication_ids` with a list of comma-separated identifier of replications.
3. Delete all the snapshots.
4. Delete the protection domain .

Repeated transitions in and out of NearSync may occur because of the following reasons:

- LWS store usage may be high.
- Change rate of data may be high for the available bandwidth between the primary and the remote sites.
- Internal processing of Lightweight Snapshots may take more time because the system may be overloaded.

Retention Policy

Depending on the scheduling configuration (1–15 minutes), snapshots are retained for a specific amount of time. With NearSync, you can configure the retention policy for days, weeks, or months on both the primary and remote sites instead of defining the number of snapshots you want to retain.

For example, if you desire RPO of 1 minute and want to retain the snapshots for 5 days, the following retention policy is applied:

- For every 1 minute, a snapshot is created and retained for the maximum of 15 minutes.

  Note: The recent 15 snapshots are only visible in the web console and are available for the restore operation.
- For every hour, a snapshot is created and retained for 6 hours.
- One daily snapshot is created and retained for 5 days.

You can also define snapshot retention in weeks or months. For example, if you configure a 3-month schedule, following retention policy is applied.

- For every 1 minute, a snapshot is created and retained for 15 minutes.
- For every hour, a snapshot is created and retained for 6 hours.
- One daily snapshot is created and retained for 7 days.
- One weekly snapshot is created and retained for 4 weeks.
- One monthly snapshot is created and retained for 3 months.

Note:

- You can define different retention policies on the primary and remote sites.
- Subhourly and hourly snapshots are retained for 15 minutes and 6 hours respectively. Maximum retention time for days, weeks, and months is 7 days, 4 weeks, and 12 months respectively.

- If you change the protection domain configuration from Async DR to NearSync, the first snapshot is not created according to the new schedule (if you have configured the new schedule as a NearSyncreplication schedule). The snapshots are created according to the start time of the old schedule that you have configured in Async DR. If you want to get the maximum retention for the first snapshot after modifying the schedule, update the start time accordingly for NearSync.

Failover/Failback and Recovery Process

At any given point in time, depending on the retention policy that you have configured, you have recovery points available for the recovery purpose on the primary and the remote site. You can recover from any of these snapshots.

Note:

- Currently, retrieval of snapshots from the remote site to the primary site is not possible. If you want to restore from a remote snapshot, you must first restore the snapshot on the remote site and then migrate the protection domain to the primary site.
- In-place restore is not supported. If you want to restore an entity by using in-place restore, delete the schedule and then perform the operation.
- If you are configuring a remote site with backup only capability, you cannot retrieve NearSync snapshots from the remote site.

Planned Failover by Using NearSync

Before performing the **Migrate** operation, ensure that you fulfill the following conditions.

1. You can migrate protection domain configured with the NearSync schedule to the remote site that is configured in the schedule.
2. If the remote site is not configured in the schedule, then first remove the schedule and then attempt the migrate operation. This operation is similar to the Async DR migrate operation.

Note: The unplanned failover and failback operation is similar to the way that you perform in the Async DR replication. For more information, see Failover and Failback Operations for Asynchronous and NearSync Replication on page 77.

Third-Party Backups

Third-party backup support is as follows:

- (vSphere) Any third-party backup software that uses vSphere snapshots is supported.
- (AHV) All third-party backup software that uses Nutanix `v3/vm_snapshot` APIs are supported.

Requirements of Data Protection with NearSync Replication

This section provides the general conditions and requirements of using NearSync replication. It also provides information about general limitations of NearSync replication. Ensure that you implement these conditions and are aware of the limitations before configuring NearSync replication.

For the general data recovery conditions and limitations, see Conditions for Implementing Asynchronous Replication on page 42.

Requirements and Conditions

Nutanix NearSync replication:

- Nutanix supports replications between the all the latest supported LTS and STS released AOS versions. To check the list of the latest supported AOS versions, see KB-5505. Check the EOL document to determine if the AOS versions currently running on the clusters are EOL.

  Upgrade the version to the next available supported LTS/STS release. Before you upgrade the AOS, check the Upgrade Paths page to determine if an upgrade path is supported.

  > Note: If both clusters have different AOS versions that are EOL, upgrade the cluster with lower AOS version to match the cluster with higher AOS version and then perform the upgrade to the next supported LTS version as advised above.

  For example, the clusters in a metro configuration are running AOS versions 5.5.x and 5.10.x respectively. Upgrade the cluster on 5.5.x to 5.10.x. After both the clusters are on 5.10.x, proceed to upgrade each cluster to 5.15.x (supported LTS). Once both clusters are on 5.15.x you can upgrade the clusters to 5.20.x or later.

  Nutanix recommends that both the primary and the replication clusters or sites run the same AOS version.

- Supports multi-site replication (1:n sites). However, the system replicates a protection domain to only one remote site in NearSync and to multiple remote sites in Async DR.
- Supports only with a minimum of three or more nodes in a cluster on both primary and remote sites.
- Supports replication schedules of 1 minute up to 15 minutes. You cannot configure schedule for 16 minutes up to 59 minutes.

  > Note: Supports replication schedules of 20-59 seconds for VStore protection domains. For more information, see Configuring Data Protection with NearSync Replication (20-59 seconds RPO) for VStore Protection Domains on page 70.

- Supports the minimum capacity as detailed in Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro) on page 16.
- Supports configuration maximums as provided on the Nutanix Configuration Maximums page.
- Group the virtual machines that form an application in a single protection domain for failover granularity of a particular application.
- Supports ESXi and AHV.
- Supports multiple ESXi VMs with multi-writer sharing enabled only if these VMs are protected in a single consistency group within the protection domain.
- Supports only homogeneous clusters—all the nodes in a cluster must run the same type of hypervisor.

  > Note: NearSync treats all storage-heavy nodes within a cluster as homogenous.

- Supports only when both the primary and the recovery clusters register with a vCenter for ESXi hypervisors.
- Supports all platforms except IBM Power Systems.
- Supports linked clone VMs.

- Supports only when the remote site container space is equal to or exceeds the working size configured for the protected entities in the primary site.

  For example, if you are protecting a VM using 30 GB of space on the container of the primary site, the remote site container requires 30 GB or more space.

- Supports addition and deletion of the entities to a protection domain. If you are adding an entity, NearSync transitions out the protection domain. Depending on the entity that you add, it takes some time to transition back into NearSync. For more information about transitioning in and out of NearSync, see Data Protection with NearSync Replication (1–15 minutes RPO) on page 57

- Requires bandwidth that is approximately equal to or higher than the rate of change of the protected entities between the two sites.

- Supports network segmentation for disaster recovery on a stretched Layer 2 network that does not have a gateway. For more information, see *Segmenting a Stretched L2 Network for Disaster Recovery* in the *Securing traffic through Network Segmentation* section of the *Security Guide*.

- Requires that you configure the corresponding vStore mappings for a protection domain configured to replicate to the remote site.

  > Note: If you must reconfigure VStore mappings for the protection domain (for example, you add/remove/modify vStore mappings), then
  >
  > - First delete or unprotect the reconfigured vStore
  > - Then protect the vStore again.
  >
  > For more information, see Configuring Data Protection with NearSync Replication on page 64

- Nutanix cluster nodes to vSphere cluster mapping: The Nutanix data distribution architecture does not support the mapping of the nodes of a Nutanix cluster to multiple vSphere clusters.

  > Note: If the nodes of a Nutanix cluster are not in the same vSphere cluster and are split into multiple vSphere clusters, disaster recovery operations like migration and recovery fail.

  Ensure that all the nodes in the Nutanix cluster are in the same vSphere (ESXi) host cluster and that the network is available on all the nodes in the Nutanix cluster at the primary or remote site. Any configuration using Nutanix DR where there is more than one ESXi cluster on top of a single Nutanix cluster is not supported.

  > Note: For information about the operating systems that support the AOS version you have deployed, see the Compatibility Matrix.

Limitations of Data Protection with NearSync Replication

NearSync replication has the following limitations:

- Nutanix cluster nodes to vSphere cluster mapping: The Nutanix data distribution architecture does not support the mapping of the nodes of a Nutanix cluster to multiple vSphere clusters.

  > Note: If the nodes of a Nutanix cluster are not in the same vSphere cluster and are split into multiple vSphere clusters, disaster recovery operations like migration and recovery fail.

  Ensure that all the nodes in the Nutanix cluster are in the same vSphere (ESXi) host cluster and that the network is available on all the nodes in the Nutanix cluster at the primary or remote site. Any configuration using Nutanix DR where there is more than one ESXi cluster on top of a single Nutanix cluster is not supported.

- NearSync does not support application-consistent snapshots.

  > Note: To take application-consistent snapshots of a NearSync protection domain, add another hourly schedule on the same protection domain, and enable **Application-Consistent Snapshots**. This way, the minutely snapshots are not application-consistent but the hourly snapshots are.

  > Note: Configure only one VM in a consistency group to use application-consistent snapshots.

- Hyper-V does not support NearSync replication schedules.

- Deduplication enabled on storage containers having protected VMs lowers the replication speed.

- NearSync replication schedules of less than 15 minutes do not support self-service restore.

  For the NearSync workflow (up to 15 minutes), self-service restore uses the last hourly full snapshot taken. NearSync replication generates hourly recovery points to avoid too many delta recovery points.

- The following replication methods have minimum RPO as follows.

  - Cloud Connect—1 hour.
  - Single-node and two-node clusters—6 hours.
  - Cross-Hypervisor DR—1 hour.

- If a protection domain fails over from one ESXi cluster to another ESXi cluster, a thick-provisioned disk changes to thin-provisioned.

- Ensure that you locate all the files associated with the VMs running on ESXi in the same folder as the VMX configuration file. The files not located in the same folder as the VMX file could become irrecoverable on a remote cluster. When you migrate a protection domain that includes a VM running on ESXi to a remote cluster, the VM fails to power-on with the following error message `Operation failed: InternalTaskCreationFailure: Error creating host specific VM change power state task. Error: NoCompatibleHost: No host is compatible with the virtual machine`.

Handling Network Partition (NearSync)

Perform the following procedure, if the deployment requires network isolation between the primary and remote site.

About this task

Following scenarios may occur:

Procedure

1. Network between primary site (site A) and remote site (site B) is restored and both the sites are working.
   a. Primary site tries to transition into NearSync automatically between site A and site B. No manual intervention is required.
2. Site B is not working or destroyed (for whatever reason). If you create a new site (site C) and want to establish subhourly schedule from A to C.
   a. Configure subhourly schedule from A to C.
      The configuration between A to C succeeds. No other manual intervention is required.
3. Site A is not working or destroyed (for whatever reason). If you create a new site (site C) and try to configure subhourly schedule from B to C.
   a. Activate the protection domain on site B and set up the schedule between site B and site C.

Configuring Data Protection with NearSync Replication

This procedure creates a protection domain that supports snapshots for entities (selected VMs and volume groups) through asynchronous data replication with an RPO of as low as 1 minute (20 seconds for vStore protection domains).

Before you begin

- Ensure that you have met the protection domain requirements for configuring NearSync before proceeding. See Requirements of Data Protection with NearSync Replication on page 60.

About this task

To configure a NearSync, do the following.

> Note:
>
> - This feature is supported between Nutanix NX/SX Series and Dell XC Series clusters and between Nutanix NX/SX Series and Lenovo HX Series clusters.
> - Asynchronous and NearSync schedules can coexist in a protection domain. See Interoperability Between Async DR and NearSync in a Protection Domain on page 71.

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Protection Domain** button and select **Async DR** from the drop-down list.

2. Enter a name for the protection domain on the first screen and then click **Create**.

   Clicking the **Create** button immediately creates the protection domain (even if you do not continue to the next step).

   > Note: This entity has the following naming restrictions:
   >
   > - The maximum length is 75 characters.
   > - Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

   > Note: In your environment if you have more than one primary and secondary clusters, ensure that the names of each protection domain are unique for each pair of clusters.

   The Entities display appears. This display is divided into two sections for unprotected entities (left side) and protected entities (right side).

3. Do the following in the indicated fields:

   a. Check the boxes next to the entities you want in a consistency group from the list of unprotected entities on the left. You can filter the list by entering an entity or host name in the **Filter By:** fields above the list.

      > Attention: The list includes the entities that are actively stored in a Nutanix container. If entities are stored on third-party storage or storage local to the host, the entities may not show in the list.

      A consistency group is a subset of the entities in a protection domain that are treated as a single group and backed up collectively in a snapshot. Protection domains can have multiple consistency groups. You can protect a maximum of 10 entities (VMs or volume groups) within any number of consistency groups (up to 10) in a single protection domain.

   b. If there are more entities you want in this consistency group that are already protected, check the boxes next to those entities in the list of protected entities on the right, and then click the **Unprotect Selected Entities** button at the bottom of the column to move them into the unprotected entities list.

   c. Select a consistency group (in the left column) for the checked entities. Click the **Use Entity Name** button to create a new consistency group for each checked entity with the same name as that entity, click the **Use an existing CG** button and select a Consistency Group from the pull-down list to add the checked entities to that consistency group, or click the **Create a new CG** button and enter a name in the field to create a new consistency group with that name for all the checked entities.

   d. Click **Auto protect related entities** if you want to include all related entities in the consistency group.

      By default, the check box is selected. If you choose to leave the check box selected when you add entities to the **Protected Entities** list, all related entities are also added. All protected entities are included in the same consistency group and snapshots of the related entities are created for subsequent recovery. Also, the attachment configuration

of all related entities is included in the snapshot so that you do not have to reattach entities again after using the snapshot for recovery.

If you choose to clear the check box when adding entities for protection, only the entities that you select are included in the protection domain.

If you choose to protect only a subset of the related entities, the selected entities are protected, but their attachment configuration is not captured, and you must manually

reattach the entities after recovery. If the entities you excluded must be included at a later time, update the protection domain and add those entities.

e. Click the **Protect Selected Entities** button at the bottom of the column to move the selected entities into the Protected entities column.

f. To create more consistency groups for the remaining unchecked entities (or as many of them as you want in the protection domain), repeat this step.

g. When all desired consistency groups have been created, click the **Next** button (lower right).

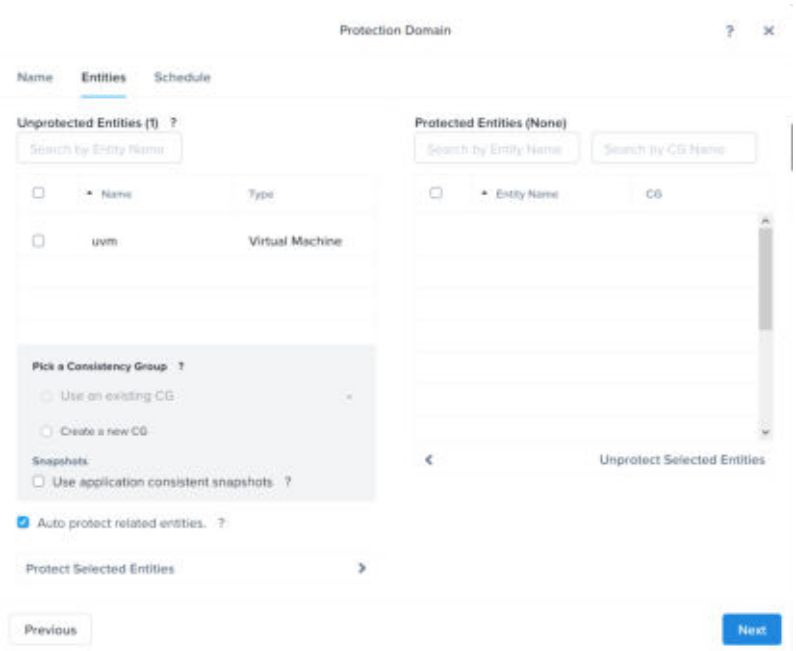

Figure 22: Virtual Machines Screen

The Schedule screen appears. There are no default backup schedules for a protection domain. Define all schedules you require.

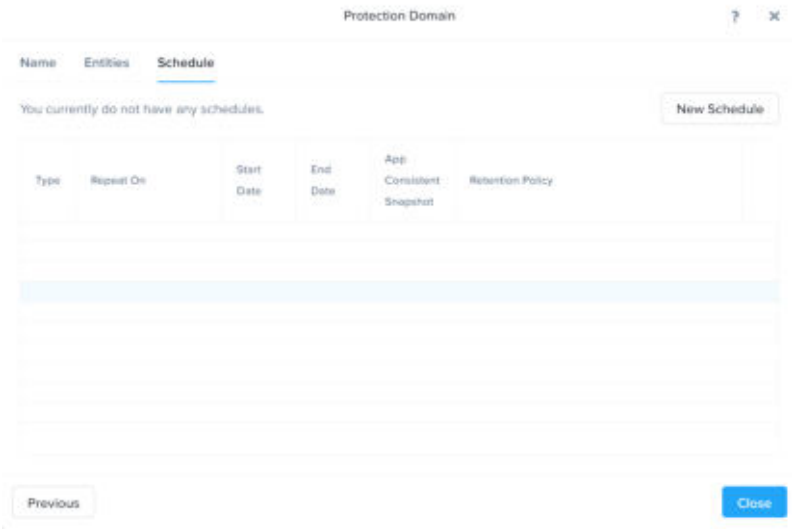

Figure 23: Schedule Screen (initial)

4. To create a schedule, click the **New Schedule** button to display the create schedule screen and then do the following in the indicated fields:

   a. To configure NearSync, in the **Repeat every ## minutes** enter a desired number between 1–15 (both included) as the scheduled time interval.

   a. **Start on**: Enter the start date and time in the indicated fields.

      The default value is the current date and time. Enter a new date if you want to delay the schedule from starting immediately.

   b. **End on**: To specify an end date, check the box and then enter the end date and time in the indicated fields.

      The schedule does not have an end date by default, and it runs indefinitely unless you enter an end date here.

   c. **Retention Policy**: Enter the snapshot retention time (days/weeks/months) that you want to retain on the local and remote sites.

      > Note: You can define different retention policies on the primary and remote sites.

      For example, if you desire RPO of 1 minute and you want to retain snapshots for 5 days, following retention policy is applied:

      - For every 1 minute, a snapshot is created and retained for 15 minutes.

        > Note: The recent 15 snapshots are only visible in the web console and are available for the restore operation.

      - For every hour, a snapshot is created and retained for 6 hours.
      - One daily snapshot is created and retained for 5 days.

      You can also define snapshot retention in weeks or months. The retention policy is updated accordingly.

      > Note: Subhourly and hourly snapshots are retained for 15 minutes and 6 hours respectively. Maximum retention time for days, weeks, and months is 7 days, 4 weeks, and 12 months respectively.

      To replicate to a remote site, select the remote site box and then enter the snapshot retention time (days/weeks/months). This retention time on the remote site can be different from what you have entered for the local site. The retention policy is similar to the primary site. For example, on the primary site, you can select to retain the snapshot

for 5 days, however, on the remote site, you can define the retention time as 4 weeks. In this case, the following retention policy is applied on the remote site.

- For every 1 minute, a snapshot is created and retained for 15 minutes.

  Note: The recent 15 snapshots are only visible in the web console and are available for the restore operation.

- For every hour, a snapshot is created and retained for 6 hours.
- One daily snapshot is created and retained for 7 days.
- One weekly snapshot is created and retained for 4 weeks.

d. When all the field entries are correct, click the **Create Schedule** button.

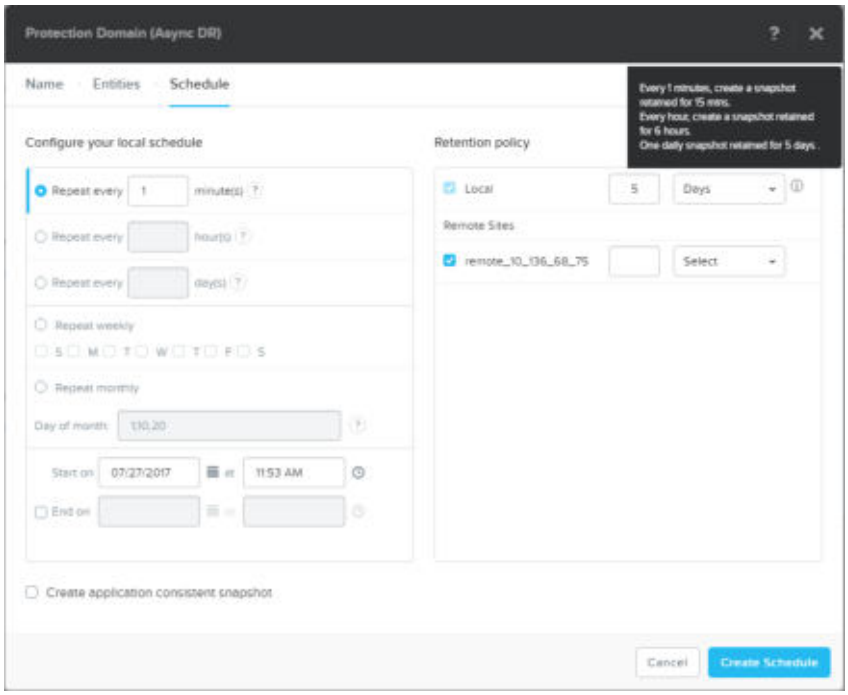

The Schedules screen returns.

5. Click the **New Schedule** button to create another schedule, update, or delete the schedule you created, or click the **Close** button to close the window.

   Note:

   - Replications can fall behind if too many schedules have the same start time. Therefore, Nutanix recommends that start times be staggered across schedules.
   - You can also define local NearSync schedule. If you define local near sync schedule, snapshots are created on the local cluster and you can view the snapshots in the **Local Snapshots** tab.

Results

Information about the new protection domain appears in the Data Protection dashboard (see Data Protection Table View on page 29).

See Remote Site Configuration on page 152 for information about configuring a remote site to store the snapshots.

Configuring Data Protection with NearSync Replication (20-59 seconds RPO) for VStore Protection Domains

A vStore is a separate mount point within a storage container which has its own NFS namespace mapped to a protection domain. You can configure NearSync replication with an RPO of as low as 20 seconds for VStore protection domains on an ESXi cluster. The schedule enables generating snapshots with a low RPO (20–59 seconds) and replicating to another cluster registered to the same or different vCenter Server. Based on the specified RPO, the schedule creates a stream of lightweight snapshots for the preceding 15 minutes. Using the slider, you can generate snapshots from the available stream of lightweight snapshots.

About this task

Use the following procedure to configure a NearSync replication schedule with an RPO between 20-59 seconds, generate a snapshot, and restore from the generated snapshot.

Before you begin

See Requirements of Data Protection with NearSync Replication on page 60 and Limitations of Data Protection with NearSync Replication on page 62 before you start.

Procedure

1. Protect the VStore protection domain using the `vstore protect` command.
   For more information about VStore protection, see *Nutanix Command -line Interface Reference* in the *Command Reference* for AOS.
   You can unprotect the VStore using the `vstore unprotect` command.
2. Logon to the Prism web console of the cluster hosting the VStore container.
3. Go to **Data Protection** > **Async DR** and select the VStore protection domain.
4. Click **Update**.
5. Go to the **Schedule** tab and do the following.
   Click **Next** in the **Entities** tab.

   > Note: Since the VStore protection domains are created using nCLI, you can add entities to protection domains or remove entities from protection domains using nCLI only.

   a. Under **Configure your local schedule**, enter the time in seconds (20–59 seconds). This time defines the interval between two lightweight snapshots.
   b. Enter the snapshot retention number for **Local** and **Remote sites**.
   c. Click **Create Schedule**.
   The VStore protection domain is updated with the schedule and the stream of lightweight snapshots starts generating.
6. Go to **Summary** > **Local Snapshots**.
7. Click **Snapshot stream** > **Actions** > **Generate Snapshot**.
8. Select a time for the snapshot by moving the slider within the progress bar.
   The progress bar shows the time for which the lightweight snapshots are available.

9. Click **Generate Snapshots**.
   The nearest available snapshot to the selected time is created and shown under **Summary** > **Local Snapshots**. You can restore using this snapshot. For information about restoration, see Restoring an Entity from a Protection Domain on page 179

Interoperability Between Async DR and NearSync in a Protection Domain

Asynchronous and NearSync schedules can coexist in a protection domain. Coexisting NearSync and asynchronous schedules must be aligned if the asynchronous schedule has a 1-hour RPO. Alignment is not required if an asynchronous schedule with a 1-hour RPO does not exist in the protection domain. This section applies only if you create NearSync schedules in a protection domain with a 1-hour asynchronous schedule.

How Alignment Is Performed

An asynchronous schedule with a 1-hour RPO creates full snapshots every hour. Every NearSync schedule also includes an internally generated 1-hour schedule to roll up the NearSync snapshots. Therefore, a protection domain that has both asynchronous and NearSync schedules effectively includes two 1-hour schedules that must be aligned to reduce the number of full hourly snapshots.

Alignment of Co-Existing Asynchronous and NearSync Schedules

The alignment of the schedules of coexisting NearSync and asynchronous schedules depends on the order in which you create the schedules. It is as follows.

- If the protection domain includes a NearSync schedule when you create an asynchronous schedule, you are not allowed to select the start time for the asynchronous schedule. The start time of the asynchronous schedule is aligned with the next NearSync roll-up time.

  For example, if the hourly NearSync schedule takes snapshots at 10:00 am, 11:00 am, and so on, and you create an asynchronous schedule at 2:30 pm., the start time of the asynchronous schedule is set to 3:00 pm to align it with the hourly NearSync schedule.

- If a protection domain includes an asynchronous schedule when you create a NearSync schedule, you are allowed to specify the start time of the NearSync schedule. However, the 1-hour schedule that is generated internally for NearSync is aligned with the next asynchronous snapshot while the other internally generated lower-frequency schedules are assigned the start time that you specify.

  For example, if the 1-hour asynchronous schedule has a start time of 10:00 am, and you create a NearSync schedule to begin at 2:30 pm, only the internally generated hourly NearSync schedule is aligned with the asynchronous schedule (the hourly schedule is set to start at 3:00 pm) while the minutely, daily, and higher schedules (if configured) are set to start at 2:30 pm.

The snapshot generation for the schedule you are creating, is not held up if the start time of the existing schedule is in the future. However, the start time of the hourly snapshots for the schedule you are creating (start time of the internally generated hourly schedule if you are creating a NearSync schedule) is adjusted in such a way that an alignment is achieved when the existing schedule begins.

For example, if the existing schedule is a NearSync schedule that is set to begin at 10 am 5 days from the current date, and you configure an asynchronous schedule to begin at 8:30 am on the current date, the start time of the asynchronous schedule is set to 9:00 am on the current date.

Conversely, if the existing schedule is a 1-hour asynchronous schedule that is set to begin at 10 am five days from the current date, and you configure a NearSync schedule to begin at 8:30

am on the current date, the start time of the internally generated hourly schedule is set to 9:00 am on the current date. The start time of the other NearSync schedules is set to 8:30 am on the current date.

Start Time Updatability

After you have configured a NearSync and an asynchronous schedule (in any order) for a protection domain, you cannot update the start time of the hourly schedule. However, you can update the start time of the NearSync schedule. If you do so, the start times of the internally generated schedules other than the hourly schedule are updated. The hourly NearSync schedule remains aligned with the asynchronous schedule.

Behavior When a Co-Existing Schedule Is Deleted

When the NearSync Schedule Is Deleted

If you delete the NearSync schedule in a protection domain, you can update the start time of the asynchronous schedule.

When the 1-hour Asynchronous schedule is deleted

If you delete the 1-hour asynchronous schedule, the hourly schedule that was generated internally for the NearSync schedule is aligned with the start time that was specified for the NearSync schedule.

For example, consider that a 1-hour asynchronous schedule exists when you configure a NearSync schedule, and that the start time of the asynchronous schedule was 10:00 am. Also consider that you create a NearSync schedule to begin at 2:30 pm on the current date. The start time for the schedules other than the internally generated hourly schedule is set to the specified time of 2:30 pm on the current date. The start time for the internally generated hourly schedule is set to 3:00 pm on the current date for alignment with the asynchronous schedule. If you delete the 1-hour asynchronous schedule, alignment is no longer required, and the internally generated hourly schedule for NearSync is realigned with the start time of the other NearSync schedules (that is, 2:30 pm).

Requirements for Interoperability

The primary and remote cluster must run AOS 5.9 or later versions.

Limitations with Interoperability

You can create any number of schedules for a protection domain, but with the following interoperability limitations:

- There can only be a maximum of two remote sites across all schedules in a protection domain.
- A protection domain can have any number of asynchronous schedules but only one NearSync schedule. Among the asynchronous schedules, only one schedule can have an interval of 1 hour.
- The schedules in the protection domain (asynchronous and NearSync) do not support self-service restore.

  For more information about self-service restore, see Self-Service Restore on page 169.

Modifying Asynchronous or NearSync Replication Configurations

About this task

You can modify an existing protection domain by changing the domain mode (active/inactive), replicating the domain, migrating the domain, updating the domain settings, or deleting the domain. To modify a protection domain in one or more of these ways, do the following:

> Note: This procedure is applicable to both Async DR and NearSync replications.

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Table** view.
2. Click the **Async DR** tab and select the target domain line in the table (top section of screen).
   The `Summary` line (middle of screen) displays the domain name with a set of relevant action links on the right. The possible actions are **Take Snapshot**, **Migrate**, **Update**, and **Delete**. Further steps describe how to perform each action.
3. To change the protection domain mode from inactive to active, click the **Activate** action link.
   This action is typically taken to recover from a site disaster (see Performing Failover on page 77). An active protection domain cannot be deactivated through the web console, but you can deactivate it through the nCLI. When a protection domain is inactive, backups (scheduled snapshots) in that protection domain are disabled.
4. To create a snapshot (point-in-time backup) of the protection domain locally and to one or more remote sites, click the **Take Snapshot** action link.
   The Replicate protection domain dialog box appears. Do the following in the indicated fields:
   a. **Local**: Do nothing in this field.
      The box is checked by default to indicate that a local snapshot is always created. If no remote site is checked, the snapshot is saved locally. If one or more remote sites are checked, the snapshot is saved locally and at the specified remote sites.
   b. **Remote sites**: Check one or more boxes of the target remote sites.
      > Note: If you do not check a remote site, the snapshot is saved locally.
   c. **Replication start time**: Select the start time from the pull-down list.
      Choices range from now (current time) to 48 hours from now (current time).
   d. **Retention time**: Select the retention interval (how long the snapshot is saved on the remote sites) from the pull-down list.
      Choices range from one day to `indefinitely`.
      Snapshots are retained based on the number of days, weeks or months you select. If you select **No Expiration**, the snapshots are retained indefinitely. For example, if you have set a retention time of two days at 00:00:00 hrs on day1, the snapshots generated at 00:00:10 hrs are retained until 00:00:09 hrs on day3. If you modify the retention time to five days at 00:00:20 on day2, all snapshots taken up to 00:00:20 on day 2 are retained for two days (until 00:00:20 hrs on day4) and snapshots taken from 00:00:20 hrs on day2 are retained for five days (until 00:00:20 hrs on day7).
      The snapshots created with the retention of two days expire in two days even if the schedule is changed from two days to five days. The existing snapshots honor the old

retention period. The snapshot created after the schedule is modified to retain it for five days expire in five days. See the following example.

Sample Schedule: Daily snapshot at 1 am retention of two days changed to five days*:

Date	Snapshot Create Time Stamp	Retention	Snapshot Expiration Time Stamp
Mar 1	Mar 1, 1am	2 days	Mar 3, 1am
Mar 2	Mar 2, 1 am	2 days	Mar 4, 1 am
Mar 3	Mar 3, 1am	5 days*	Mar 8, 1 am

Mar 4	Mar 4, 1 am	5 days	Mar 9, 1 am

e. Select the **Create application consistent snapshot** check box to make the snapshots application consistent.

   For more information about application-consistent snapshot, see Conditions for Application-consistent Snapshots on page 18.

f. When all the field entries are correct, click **Save**.

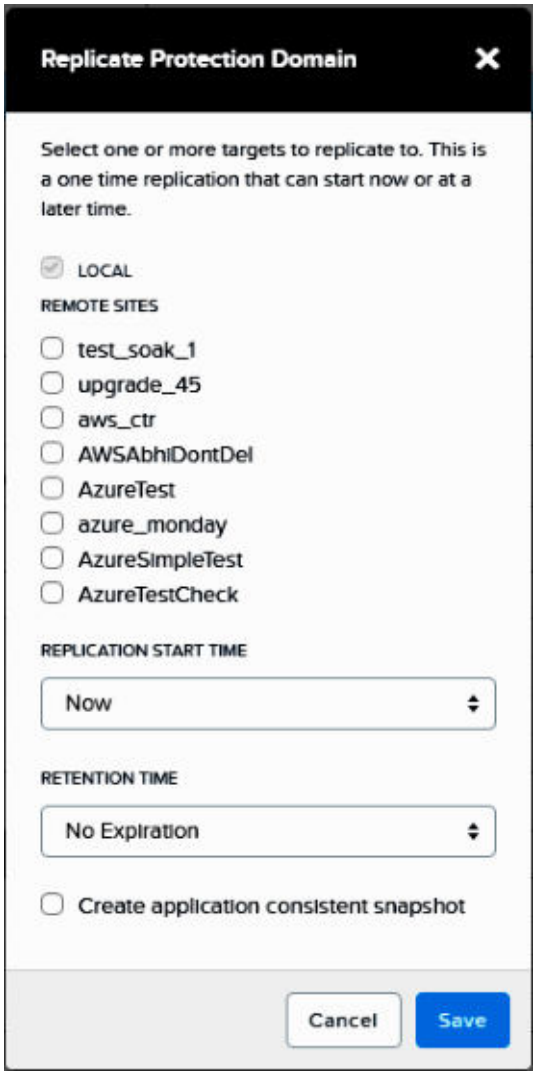

Figure 24: Replicate Protection Domain Dialog Box

5. To migrate the protection domain to a remote site, click the **Migrate** action link.

   See the *Migration (Planned) Failover* section of Performing Failover on page 77 and follow the instructions to migrate a protection domain.

6. To update the protection domain settings, click the **Update** action link. Update the displayed screens as required and click **Save**.

   The Update Protection Domain dialog box appears (see Configuring Data Protection with Asynchronous Replication on page 49).

   Note:

   - If you are adding entities to an existing consistency group, click the **Use an existing CG** option and select the consistency group from the list.
   - For NearSync, addition or deletion of entities to a protection domain is disabled if you have configured this feature to ensure that the recovery points of other entities are not affected. If you want to edit the protection domain that is configured by using this feature, first remove the configured NearSync schedule and edit the protection domain (add or delete the entities) and then reconfigure the new schedule.

7. To delete the protection domain, see Deleting a Protection Domain on page 76.

Deleting a Protection Domain

Before you begin

Before deleting a protection domain, unprotect all the entities, delete all the replication schedules from the protection domain, and delete all the snapshots of the protection domain.

About this task

To delete a protection domain, do the following.

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Table** view.
2. Click the **Async DR** tab and select the target domain in the table (top section of the screen).
3. Click the **Update** action link and go to the **Entities** tab.
4. Select all the protected entities in the protection domain.

   Tip: You can search for entities by entering their name or the name of their consistency group in the **Search by Entity Name** and **Search by CG Name** respectively.

5. Click **Unprotect Selected Entities** to unprotect all the selected entities and then click **Next**.
6. In the **Schedule** tab, delete all the replication schedules one by one and then click **Close**.
7. Select the target protection domain in the table and go to the **Local Snapshots** tab at the bottom.
8. Select all the snapshots and click **Delete** to delete all the snapshots of the protection domain.
9. Select the target protection domain in the table and click **Delete**.

10. Click **OK** to delete the protection domain.
    The deleted protection domain disappears from the list of domains in the table.

> CAUTION: Do not delete a protection domain that contains entities but is inactive. Deleting such protection domains also deletes their entities. Either activate the protection domain before deleting it or remove the entities from the protection domain.

Failover and Failback Operations for Asynchronous and NearSync Replication

After a protection domain is replicated to at least one remote site, you can carry out a planned migration of the contained entities by failing over the protection domain. You can also trigger failover in the event of a site disaster.

Failover and failback events re-create the VMs and volume groups at the other site, but the volume groups are detached from the iSCSI initiators to which they were attached before the event. After the failover or failback event, you must manually reattach the volume groups to iSCSI initiators and rediscover the iSCSI targets from the VMs.

Performing Failover

Migration (Planned) Failover

About this task

System maintenance or expansion might dictate moving a protection domain to another site as a planned event. To migrate a protection domain to a remote site, do the following:

Procedure

1. Log into the web console for the primary site (see *Logging In to the Web Console* in the *Prism Web Console Guide*).
2. Go to the **Async DR** tab of the **Data Protection** table view (see Data Protection Table View on page 29).

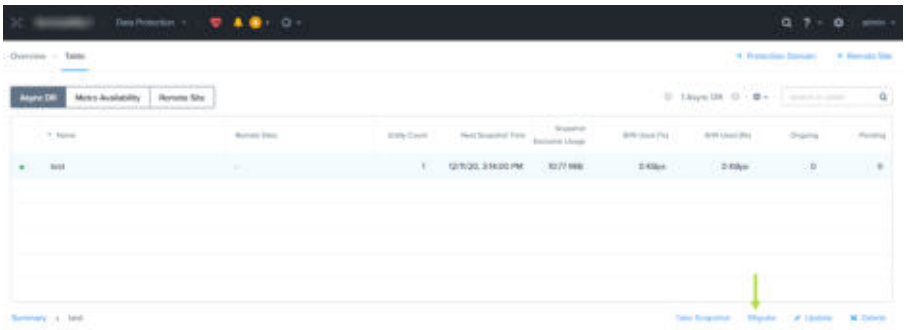

Figure 25: Migrating a Protection Domain

3. Select the target protection domain and click the **Migrate** action link.
   The Migrate Protection Domain dialog box appears. Select the site where you want to migrate the protection domain. Type `MIGRATE` in the text box.

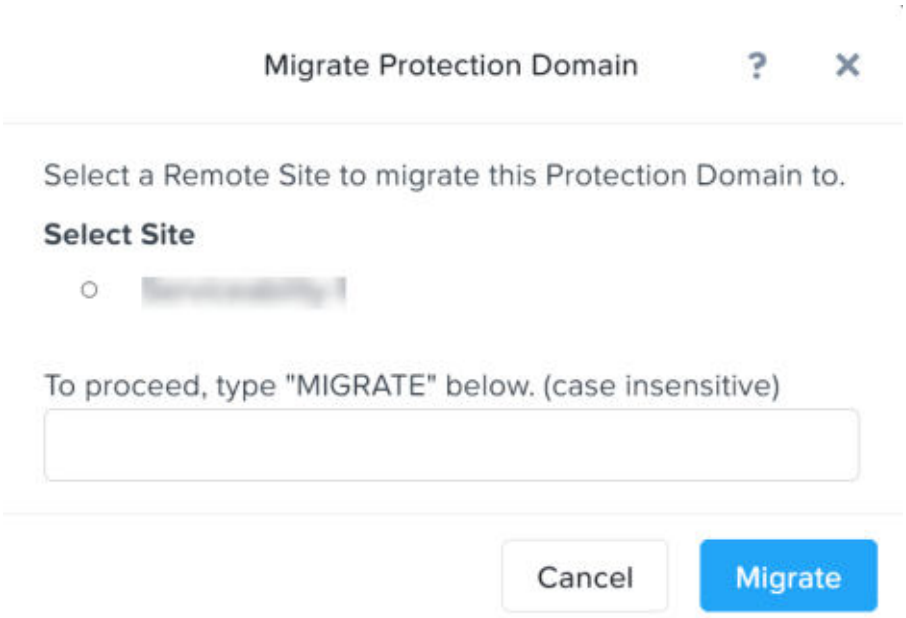

Figure 26: Migrate Protection Domain Dialog Box

4. Click **Migrate**.

   Migrating a protection domain does the following:

   - Creates and replicates a snapshot of the protection domain.
   - Powers off the VMs on the local site.

     Note: The data protection service waits for 5 minutes for the VM to shutdown. If the VM does not get shutdown within 5 minutes, it is automatically powered off.

   - Creates and replicates another snapshot of the protection domain.
   - Unregisters all VMs and volume groups and removes their associated files.
   - Marks the local site protection domain as inactive.
   - Restores all VM and volume group files from the last snapshot and registers them with the same UUIDs at the remote site.

     Note: If you are migrating a protection domain to a different hypervisor (for example, ESXi to AHV), all VM and volume group files register with a new UUID at the recovery site.

   - Marks the remote site protection domain as active.

     Note: Perform the following additional steps for the cluster with Hyper-V nodes.

     - If VMs are registered with Microsoft Failover Cluster, you need to manually re-register these VMs with the Microsoft Failover Cluster after the Async DR is completed.
     - After the migration of the protection domain is completed, you need to manually remove stale entries belonging to the migrated VMs from the Microsoft Failover Cluster manager on the source cluster. If you do not perform this operation, then data corruption might occur during the failback operation.

   The VMs on the remote site are not powered on automatically. This allows you to resolve any potential network configuration issues, such as IP address conflicts, before powering on the VMs. Additionally, you must manually reattach the volume groups that were affected by the migration or restore operation, perform in-guest iSCSI attachment. If you use the nCLI for attaching volume groups, note that the UUIDs of volume groups change when they are restored at the remote site, and so do their iSCSI target names, which contain the volume group UUID. See Volume Group Configuration in the *Prism Web Console Guide*.

Disaster (Unplanned) Failover

About this task

When a site disaster occurs, do the following to fail over a protection domain to a remote site:

Procedure

1. Log into the web console for the target remote site (see *Logging In to the Web Console* in the *Prism Web Console Guide*).
2. Go to the **Async DR** tab of the **Data Protection** table view (see Data Protection Table View on page 29).

3. Select the target protection domain and click the **Activate** button. A window prompt appears; click the **Yes** button.

   This operation does the following:

   - Restores all VM and volume group files from the last fully-replicated snapshot.

     The process first detaches the volume groups that are included in the protection domain or attached to the VMs in the protection domain.
   - Registers the VMs and volume groups on the recovery site.
   - Marks the failover site protection domain as active.

     Note: Perform the following additional steps for the cluster with Hyper-V nodes.

     - If VMs are registered with Microsoft Failover Cluster, manually re-register these VMs with the Microsoft Failover Cluster after the Async DR is completed.
     - After the migration of the protection domain is completed, manually remove stale entries belonging to the migrated VMs from the Microsoft Failover Cluster manager on the source cluster. If you do not perform this operation, then data corruption might occur during the failback operation.

   The VMs are not powered on automatically. This allows you to resolve any potential network configuration issues, such as IP address conflicts, before powering on the VMs. Also, you must manually reattach the volume groups that were affected by the migration or restore operation, perform in-guest discovery of the volume groups as iSCSI targets, and log in to the targets. If you use the nCLI for attaching volume groups, note that the UUIDs of volume groups change when they are restored at the remote site, and so do their iSCSI target names, which contain the volume group UUID. See Volume Group Configuration in the *Prism Web Console Guide*.

Failing Back a Protection Domain

If an unplanned failure occurs on the primary site, all the entities failover to the remote site. After the primary site is restored, failback the entities to the primary site.

Procedure

1. Sign in to the primary site.

   All the hosts are down.
2. Start all the hosts of the primary site.

   - Controller VMs are automatically restarted and the cluster configuration is established again.
   - All the protection domains that were active before the disaster occurred gets recreated in an active state. However, you cannot replicate the protection domains to the remote site since the protection domains are still active at the remote site.
3. Log on to one of the Controller VMs at the primary site and enter the hidden nCLI command mode by using the following nCLI command.

   ```
   nutanix@cvm:~$ ncli -h true
   ```

4. Deactivate and destroy the VMs by using the following hidden nCLI command. Run this command only when the protection domain is active on the secondary site because this command deletes the VMs from the primary site.

   ```
   ncli> pd deactivate-and-destroy-vms name=protection_domain_name
   ```

   Replace *protection_domain_name* with the name of the protection domain that you want to deactivate and destroy.

   > Note: Shut down all the VMs on the primary site before running the command.

   > CAUTION: Do not use this command for any other workflow. Otherwise, it deletes the VMs and data loss occurs.

   VMs are deleted from the primary site and the protection domain is no longer active on the primary site. Remove the orphaned VMs from the inventory of the primary site.

5. (Optional) If you want to schedule frequent replications, log into the remote site and schedule replication to the primary site.

6. Log into the remote site and click **Migrate** to migrate the protection domain to the primary site.

   The VMs are unregistered from the secondary site (not removed) and data is replicated to the primary site. Then the VMs are registered on the primary site. The protection domain starts replicating from the primary site to the remote site based on the schedule that was originally configured on the primary site.

7. Start the VMs in the correct sequence and configure the IP address again (if necessary). Also, manually reattach any volume groups that were either included in the protection domain or attached to the VMs in the protection domain. You can start using the VMs from the primary site.

8. Remove the orphaned VMs from the inventory of the secondary site.

Performing a Failback

Perform the following steps to failback a protection domain from remote site to primary site.

Procedure

1. Login to the Web console of the remote site - the site where the protection domain is currently active.

2. From the **Async DR** tab under Data Protection, select the protection domain that you want to failback.

3. Click **Migrate**.

   The Migrate Protection Domain dialog box appears. Select the site where you want to migrate the protection domain.

4. When the field entries are correct, click the **Save** button.

What to do next

Manually reattach the volume groups that were affected by the migration or restore operation. Note that the UUIDs of volume groups change when they are restored at the remote site, and so do their iSCSI target names, which contain the volume group UUID.

4

METRO AVAILABILITY (ESXI AND HYPER-V 2016)

You configure a metro availability protection domain by specifying an active storage container in the local cluster (primary site) and linking it to a standby storage container with the same name on a remote site in which synchronous data replication occurs when metro availability is enabled.

> Note: Nutanix provides multiple DR solutions to protect your environment. See Nutanix DR Solutions Workflow on page 13 for the detailed representation of the workflow.

Nutanix native replication infrastructure provides flexible options to backup and restore VMs from snapshots (see Protection Strategies on page 9). In addition, Nutanix offers a metro availability option that is a policy applied on a datastore, which effectively spans the datastore across two sites. This is accomplished by pairing a storage container on the local cluster with one on a remote site and then synchronously replicating data between the local (active) and remote (standby) storage containers. When metro availability is enabled, everything in the active storage container is replicated synchronously to the remote storage container. Metro availability configurations can include VMs, but they cannot include volume groups.

> Note: After you enable metro availability, the cluster no longer automatically takes a backup snapshot of the storage container at the secondary site, as this snapshot can consume a large amount of storage. This snapshot was previously intended to protect pre-existing data on that storage container.

Metro availability policies apply per storage container (not cluster), so a cluster can be active for one datastore and standby for another. For example, consider a cluster configuration with an Oracle datastore (Datastore A) in storage container 1 and an Exchange datastore (Datastore B) in storage container 2. Cluster A can be active for Datatstore A with Cluster B as the standby, while Cluster B can be active for the Datastore B with Cluster A as the standby. In addition, metro availability storage containers can co-exist with regular storage containers in the same cluster.

Metro availability minimizes application downtime due to unplanned failure events like site disasters and planned events such as site maintenance. A metro availability policy for a storage container interoperates seamlessly with the compression data management policy.

> Note: Metro availability is supported on clusters running ESXi or Hyper-V only. Metro is supported on ESXi clusters that include storage-only nodes (which run AHV) but not Hyper-V clusters with storage-only nodes.

You have the option of adding a Witness to a metro availability configuration (see Conditions for Implementing Data Protection (Metro Availability) on page 83). For information about Witness VMs, see Witness Option on page 116.

> Note: You can deploy a Witness VM only on AHV and ESXi hypervisors. For information about Witness options with AHV Metro Availability, see the *AHV Metro (Witness Options)* in the *Leap Administration Guide*.

For information about Witness VMs, see Witness Option on page 116.

Conditions for Implementing Data Protection (Metro Availability)

General Conditions for Applying a Metro Availability Policy

A protection domain is created by specifying a primary storage container on the local cluster and a standby storage container of the same name on a remote cluster. Following are settings and conditions for applying metro availability.

> Note: Do not create a metro availability protection domain when there are active replications running in the cluster. When there are any ongoing replications in the cluster, you must wait until the replications are completed before creating a metro availability protection domain.

When several overlapping schedules are running, it may not be possible to get any time period during which replications are not running. In such a case, you need to suspend the schedules, and allow ongoing replications to complete before you create a metro availability protection domain.

To suspend replication schedules, use the `ncli pd suspend-schedules` command. To resume replication schedules, use the `ncli pd resume-schedules` command. For more information on these schedules, see the *Command Reference* .

Container Conditions and Settings

- Nutanix recommends that the VMs run locally with the active container on which the data of the VMs is residing to reduce network traffic and associated network latency.

  Protected VMs can run on the Nutanix cluster either associated with the active storage container or associated with the standby storage container. However, when VMs are running on the cluster associated with the standby storage container, reads and writes are sent over a network connection to the active storage container.

- Ensure that the datastore name must be the same as the storage container name when a storage container is used with metro availability.

- Ensuring datastore availability across two sites requires data to be present at both sites in real time. Therefore, it is required that the round trip latency between clusters be less than five (5) ms. Maintain adequate bandwidth to accommodate peak writes. Nutanix recommends that you have a redundant physical network between the sites.

- Minimize the number of snapshots retained (**Retention policy**) on both the primary and remote sites to avoid metadata bloat. If you need a large snapshot retention, configure a third site where you can set an appropriate retention policy to do so without impacting the metro protection domains.

  Metro availability uses snapshots taken for the entire container. As the snapshot retention increases, the number of `vmware.log` files increases, thus causing a metadata bloat.

- Ensure that the gold image must reside in the same metro availability storage container where linked clones are.

- Container Configuration Settings
  - Compression settings for a container can be different across clusters.
  - Redundancy factor settings for a container can be different across clusters.
  - Deduplication enabled on storage containers having protected VMs lowers the replication speed.

VM Related Conditions and Settings

- Linked cloned VMs must have host affinities to either primary site or secondary site. The linked cloned VMs must always be registered on the same site as their root VM.
- Ensure that all virtual disks associated with a given VM reside on a container enabled for metro availability.
- Ensure that the following ports remain open for all the Controller VM IP addresses in Cluster A and Cluster B:
  - TCP 2009
  - TCP 2020

  Communication between the two metro clusters happens over these ports and, therefore, requires that these ports must remain open.
- Metro availability does not support replication of multiple ESXi VMs with multi-writer sharing enabled because if multiple ESXI VMs have multi-writer sharing enabled for disks, then they can be protected only if they are in same CG within the protection domain. Metro availability only supports replication of VMs if they are placed in different consistency groups within vStore protection.
- Cluster Configuration Settings
  - Cluster models can be different for primary and secondary sites.
  - vMotion may require the Enhance vMotion Compatibility (EVC) feature.
  - Cluster resource sizes can be different.
  - Redundancy factor can be different across clusters.
- Remote Site Configuration Conditions
  - Ensure that you do not enable a proxy on remote sites that you use with a metro availability protection domain.
  - Ensure that you set up the remote site clusters using the virtual IP address of the remote site.
  - Ensure that you have redundant remote replication links between the clusters.
  - Ensure that you enable remote site compression in bandwidth-limited environments.

  See Remote Site Configuration on page 152 for information about configuring a remote site to store the snapshots.

VMware and vSphere Configuration Conditions

- You cannot add an ESXi vGPU-enabled VM to a protection domain for Nutanix DR solutions like asynchronous, NearSync or Metro Availability replication methods. Metro availability does not support snapshot replications of vGPU-enabled VMs.
- Ensure that VMware network port group names are identical between the VMware hosts in each Nutanix cluster.
- Nutanix cluster nodes to vSphere cluster mapping: The Nutanix data distribution architecture does not support the mapping of the nodes of a Nutanix cluster to multiple vSphere clusters.

  > Note: If the nodes of a Nutanix cluster are not in the same vSphere cluster and are split into multiple vSphere clusters, disaster recovery operations like migration and recovery fail.

  Ensure that all the nodes in the Nutanix cluster are in the same vSphere (ESXi) host cluster and that the network is available on all the nodes in the Nutanix cluster at the primary or remote site. Any configuration using Nutanix DR where there is more than one ESXi cluster on top of a single Nutanix cluster is not supported.
- Nutanix recommends using an empty container for metro availability configuration.

  Disable SIOC (which is enabled by default on a container) and delete all the files from the container that are related to SIOC.

  For more information about disabling SIOC on a container, see *vSphere Administration Guide*.
- Nutanix strongly recommends enabling VMware vSphere HA and DRS.

  While these features are not required, they improve performance and availability. If you choose to enable these features, make sure that the required vSphere licenses are installed.
- VMware HA Settings
  - Configure the hosts in the two Nutanix clusters as a single VMware HA cluster in the vCenter Server. This provides high availability protection for VMs across the Nutanix clusters.

    > Note: Do not configure multiple VMware HA clusters. Multiple VMware HA clusters are not a supported configuration and prevent the proper usage of VMware HA, which Nutanix Metro Availability relies on.
  - Configure the VM restart priority of all Controller VMs to **Disabled**.
  - Configure the host isolation response of all Controller VMs to **Disabled**.

    > Note: For guest VMs, it is recommended to change the VM restart priority and host isolation response settings to **Disabled**.
  - Configure the VM monitoring setting for all the Controller VMs to **Disabled**.
  - Configure VMware HA admission control based on the cluster configuration. For example, assuming you have a balanced configuration, you can set **Percentage of cluster resources reserved as failover spare capacity** setting to 50 percent or lower. This setting ensures that enough resources exist in a single Nutanix cluster to support both sites.

- vSphere DRS Affinity Settings
  - Configure DRS affinity rules such that VMs are not migrated to hosts on a cluster that owns the standby container.
  - Configure **Should** affinity rules to allow automated VM restart on the nodes owning the standby datastore.
  - Configure **Must** affinity rules if you do not want automated VM restart against the standby datastore.
  - Configure or modify manually the **Must** rules to allow VMs to restart against blocked hosts.

    Configuring or modifying rules assumes the availability of the vCenter Server.
- vCenter Server Configuration and Conditions
  - Ensure that a single vCenter Server instance manages the VMware clusters between the sites.

    Nutanix recommends maintaining the vCenter server in a third site to allow cluster management regardless of the site failure.
  - Nutanix recommends that you configure Metro availability to protect and replicate the vCenter Server if the third site is not available.
- Datastore Monitoring Settings
  - During the configuration of datastore, select **Select only from my preferred datastores** option and select the datastores that are used for metro availability.
  - If the VMware cluster has only one datastore, add the advanced option *das.ignoreInsufficientHbDatastore=true*

System Maximums

Supports configuration maximums as provided on the Nutanix Configuration Maximums page.

Shut Down and Start a Cluster

If you are planning to stop your cluster for any maintenance-related activities, perform the following before you run the `cluster stop` command.

- For the protections domains that are active on this cluster, perform a planned failover. For more information, see Failing Over a Protection Domain (Planned Failover) on page 103.
- For the protection domains that are standby on the container of cluster, perform a disable operation on the remote site that is configured in the metro availability configuration.

For information about shutting down and starting clusters at both the primary and remote sites, see Shutting Down and Starting Clusters in Both Primary and Remote Sites on page 186.

Limitations of Data Protection with Metro Availability

Following are the limitations of data protection with metro availability:

- Platforms up to 120 TB HDD tier support metro availability.

  Note: Metro Availability takes snapshots of the protected entities every 6 hours.

- Metro Availability does not support inclusion of volume groups that are configured for metro availability in a protection domain.
- Metro availability requires deleting the Async DR protection domains if you want to enable metro availability on a container with VMs configured in an Async DR protection domain.
- Metro availability may increase latency on vDisks during the synchronization between two clusters.
- Metro availability does not allow forwarding of I/O from the standby site to the primary site if metro availability is disabled.

  Commands run on the hypervisors on the standby site take time to execute as the hypervisors cannot access the underlying storage protected by that protection domain.
- Metro availability does not support restoring snapshots of the protection domains with overwite access in metro availability configuration.
- Metro availability does not support symbolic links and hard links.
- Metro availability does not support hosting of VMs on the secondary cluster during the metro enable operation for the same container.
- Metro availability does not support hosting of VMs on the primary cluster during promote operation of the secondary site for the same container.
- Re-enabling metro availability triggers an asynchronous replication in the Prism web console. This replication completes when the metro availability synchronization is complete.

  After configuring metro availability and completing the synchronization, disable metro availability and then re-enable it. Navigate to the **Async DR** tab in Prism and you see an outgoing replication in progress even if you have not configured any asynchronous disaster recovery replications.

See the *Acropolis Release Notes* for more limitations and workarounds related to metro availability.

vSphere Cluster Services on Nutanix Clusters

vSphere Cluster Services (vCLS) ensures availability of cluster services such as vSphere DRS and vSphere HA to maintain the resources and health of the workloads running in the clusters. It ensures that the cluster services run even if vCenter Server instance becomes unavailable.

In this section about vCLS, supported vCenter release means appropriate release of vCenter which supports vCLS.

vCLS is deployed on each cluster when vCenter Server is upgraded to supported vCenter release or deploy a new cluster with the appropriate vCenter release. You can only upgrade the cluster to supported vCenter releases if it is already running a minimum vSphere version of 6.5 .

vCLS uses agent virtual machines (VMs) to maintain cluster services health in each cluster.

> Note: vCLS creates a maximum of three vCLS VMs in a cluster of three or more hosts.

> Note: See VMware vSphere Cluster Services documentation for more information about how the number of vCLS VMs increase based on the number of hosts in the cluster and the vSphere releases that support vCLS.

Datastore Selection and Requirements for Metro Availability

- In a fresh deployment of Nutanix cluster on supported vCenter release where there is no shared storage: vCenter creates the vCLS VMs on the local storage.
- When you upgrade the vCenter deployment to the vCLS supported vCenter release: vCenter creates the vCLS VMs on the shared storage. Shared storage includes metro containers. Therefore, the vCLS VMs may reside in metro containers also.

  When vCLS VMs reside in the metro containers:

  - Each vCLS VM contributes 13 files to the metro container capacity usage.
  - Before planned failovers, the vCLS VMs need to be storage vMotioned. vCLS VMs are like system VMs and VMware recommends that you do not perform any operations on these VMs since such operations could disrupt the DRS functionality. Migrate the vCLS VMs to a non-metro shared storage container.

Nutanix recommends that you migrate and maintain the vCLS VMs in a separate shared container on each metro site. Storage migrate at least one vCLS VM to each metro site cluster after vCenter is upgraded to supported vCenter release.

Example:

- Create containers on SiteA - SiteA_vCLS_ctr - and mount them on all hosts on SiteA.
- Create containers on SiteB - SiteB_vCLS_ctr - and mount them on all hosts on SiteB.
- Storage migrate one vCLS VM to SiteA_vCLS_ctr and one vCLS VM to SiteB_vCLS_ctr.

vSphere Metro Storage Cluster for Nutanix Metro Availability

This is a Nutanix reference.

VMware vSphere Metro Storage Cluster (vMSC) is a certified configuration for stretched storage cluster architectures. A vMSC configuration is designed to maintain data availability beyond a single physical or logical site. A storage device configured in the vMSC configuration is supported after successful vMSC certification. All supported storage devices are listed on the VMware Storage Compatibility Guide.

Requirements

The following requirements must be satisfied to support this configuration:

- Ensure that the vCenter Server is able to connect to the ESXi hosts in both storage clusters or datacenters that vMSC stretches across.

  Nutanix recommends that you do not configure the vCenter Server as a part of the metro container to ensure that you have the management plane separate from data plane.

  If you choose to configure the vCenter Server in the metro container, then, for smooth functionality during unplanned disaster events, ensure that you set up a redundancy option for the vCenter such as vCenter HA.
- Ensure that you configure the hosts in the two Nutanix clusters as a single VMware HA cluster in the vCenter Server to provide high availability protection for VMs across the Nutanix clusters.

  > Note: Do not configure multiple VMware HA clusters. Multiple VMware HA clusters configuration is not supported because it prevents the proper usage of VMware HA, which Nutanix metro availability relies on.

- You need a minimum of two Nutanix clusters. The cluster sizes may be different, for example, the primary cluster may be a 3-node cluster and replication (remote) cluster may be a 5-node cluster.
- Ensure that you configure **Remote site** on both the primary and replication clusters.
- Ensure that the maximum Round Trip Latency (RTT) between clusters on two different sites must be less than 5 ms, to maintain a synchronous replication.
- Ensure that the IP subnets used to connect the VMs on ESXi hosts are accessible on both the primary and replication clusters for smooth functionality during a VMware HA failover event.
- Ensure that the datastore names on the primary and the remote clusters are the same.
- Nutanix recommends that for best I/O performance the VMs are run on the cluster where the container is active. Configure the DRS affinity rules in the ESXi cluster to achieve this. Create one group for the hosts on each cluster and another group for the VMs in the metro container. Then, use these groups in a rule to run the VM group on hosts in the active cluster. This rule ensures that the VMs are running in the active cluster unless all the hosts are unavailable.

Upgrade Best Practices and Requirements with Metro Availability Enabled

Following are the best practices and requirements if you are planning to upgrade your cluster and have metro availability enabled.

- Nutanix supports replications between the all the latest supported LTS and STS released AOS versions. To check the list of the latest supported AOS versions, see KB-5505. Check the EOL document to determine if the AOS versions currently running on the clusters are EOL.

  Upgrade the version to the next available supported LTS/STS release. Before you upgrade the AOS, check the Upgrade Paths page to determine if an upgrade path is supported.

  > Note: If both clusters have different AOS versions that are EOL, upgrade the cluster with lower AOS version to match the cluster with higher AOS version and then perform the upgrade to the next supported LTS version as advised above.

  For example, the clusters in a metro configuration are running AOS versions 5.5.x and 5.10.x respectively. Upgrade the cluster on 5.5.x to 5.10.x. After both the clusters are on 5.10.x, proceed to upgrade each cluster to 5.15.x (supported LTS). Once both clusters are on 5.15.x you can upgrade the clusters to 5.20.x or later.

  Nutanix recommends that both the primary and the replication clusters or sites run the same AOS version.
- Nutanix strongly recommends using the 1-click upgrade available through Prism to upgrade your cluster if you have metro availability enabled.
- Ensure that you successfully upgrade one site before upgrading the other site. You can upgrade either the primary or secondary site first. Do not upgrade both the sites simultaneously, as race conditions are observed in this case.
- Cluster can host both active and standby protection domains during an upgrade.

Metro Availability on ESXi

Configuring a Protection Domain (Metro Availability)

This procedure describes how to create a Protection Domain that supports synchronous data protection for a metro availability configuration.

Before you begin

Do the following:

- Ensure that you have met the protection domain requirements for configuring metro availability before proceeding. See Conditions for Implementing Data Protection (Metro Availability) on page 83.

  > Note: Metro Availability is now supported across different hardware vendors (NX/SX, Dell, or Lenovo). You can also establish a Metro Availability relationship between Nutanix NX/SX Series and other non-NX clusters (Dell or Lenovo). However, mixing of NX and non-NX nodes in the same cluster is not supported.

- Create containers with the same name on the clusters. See *Creating a Storage Container* in the *Prism Web Console Guide*.

  > Note: Do not make this container the default container.

About this task

To configure a protection domain for metro availability, do the following:

> Note: The following user interface–based procedure creates a protection domain in which each VM is included in its own consistency group (per-VM consistency) by default, and the user interface does not allow you to alter this behavior. If you want to include all the VMs in the container in a single consistency group, before you perform this user interface–based procedure, log on to the Controller VM by using SSH, create the protection domain, and then enable Metro Availability on the container. When enabling Metro Availability, disable per-VM consistency.
>
> ```
> nutanix@cvm$ ncli pd add name=pd-name
> nutanix@cvm$ ncli pd metro-avail-enable name=pd-name per-vm-consistency=false \
>  remote-site=remote-cluster-name ctr-name=container-name \
>  failure-handling=failure-handling
> ```
>
> In the commands, replace ***pd-name*** with a name for the protection domain, ***remote-cluster-name*** with the name of the remote cluster, ***container-name*** with the name of the container, and ***failure-handling*** with the failure handling mode (one of **Automatic**, **Manual**, and **WitnessVM**).
>
> After you enable Metro Availability at the CLI, log on to the Prism web console and update the protection domain by adding a schedule on the Schedule tab.

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Protection Domain** button and select **Metro Availability** from the drop-down list.
   The Protection Domain (Metro Availability) window appears.

2. In the Name screen, enter a name for the protection domain and then click the **Next** button (lower right).

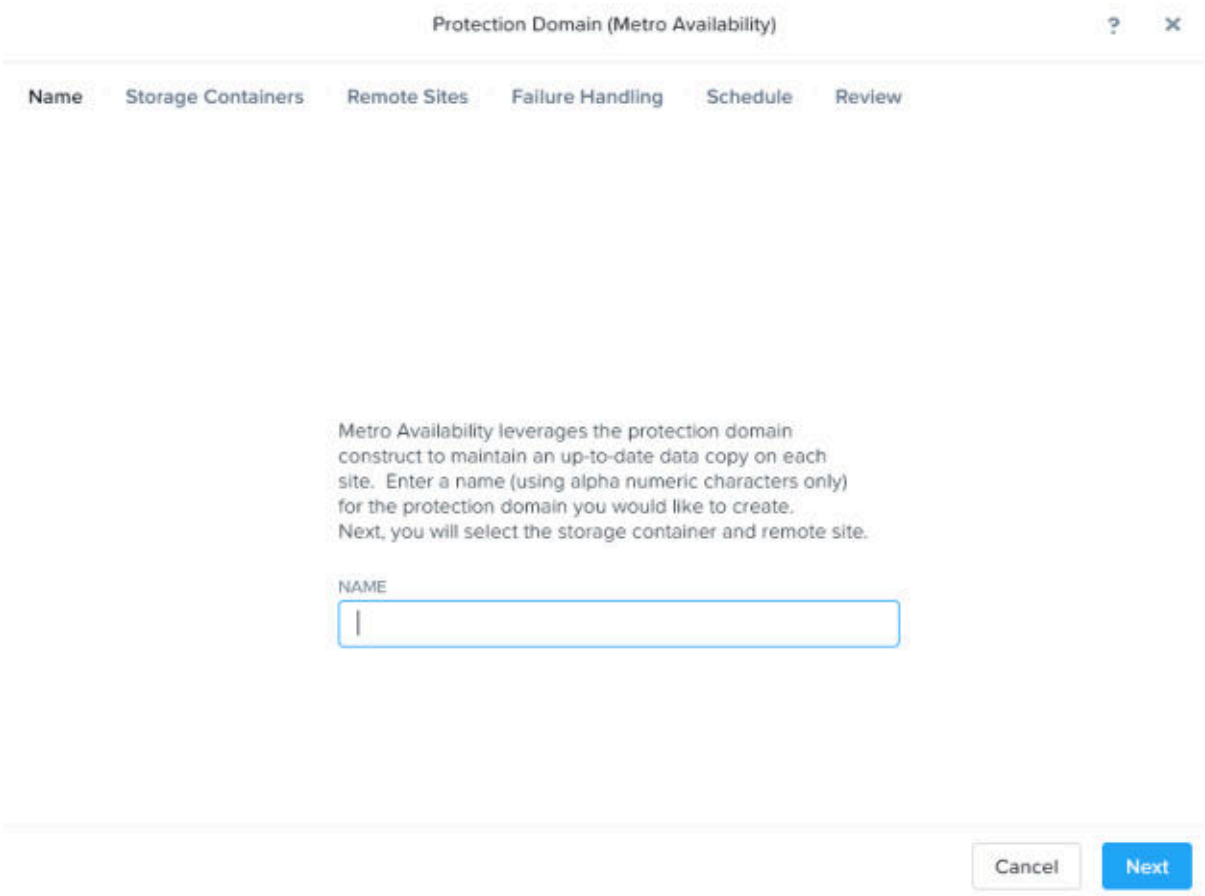

Figure 27:

3. In the Storage Containers screen, select (click the option for) the target storage container and then click the **Next** button.

   The screen displays a list of existing storage containers in the cluster. The selected storage container will be the primary storage container in a metro availability pair.

   Note:

   - You cannot select multiple storage containers. Create a separate protection domain for each storage container you want to protect. If the desired storage container does not appear in the list, click the **Cancel** button, create the storage container (see Creating a Storage Container in the *Prism Web Console Guide*), and then restart this procedure.

   - The Nutanix native backup snapshots are supported for VMs in a metro availability storage container. These snapshots are crash consistent. Restoration of such snapshots results in restoring the files that constitutes the VMs. However,

unless SRM is used (in a metro availability deployment with NearSync schedules), you need to manually register these VMs on the hosts.

- See the *Maximum Limits for Metro Availability* table in Conditions for Implementing Data Protection (Metro Availability) on page 83 topic to get more information on limits before proceeding.

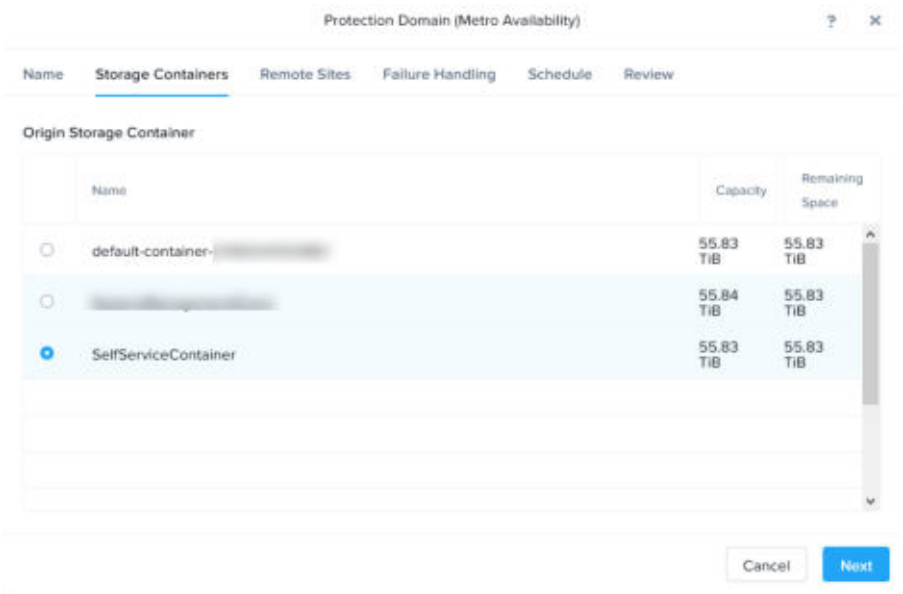

Figure 28: Protection Domain (Metro Availability): Storage Containers Screen

4. In the Remote Sites screen, select the target storage container in the Compatible Remote Sites section and then click the **Next** button.

   See Remote Site Configuration on page 152 for information about configuring a remote site to store the snapshots.

   The screen is broken into the following two sections:

   - Compatible Remote Sites (upper screen). This section lists all storage containers on remote sites that could be used as the standby storage container in a metro availability pair. The storage container must have the same name as the primary storage container selected in the previous step. In this example, the **second_stretch_container** was selected as the primary storage container, and a storage container with the same name appears in the **italy** remote site.

     Note: Make sure the target standby storage container does not contain any data. A storage container must be empty before you can select it as the standby in a metro availability pair.

   - Incompatible Remote Sites (lower screen). This section lists remote sites that are incompatible as a standby site. For a remote site to be compatible, it must first be set up to support metro availability, have a storage container with the same name, and have

a transmission latency between the clusters of less than 5 ms. Remote sites that do not satisfy all three requirements cannot be used as the standby site.

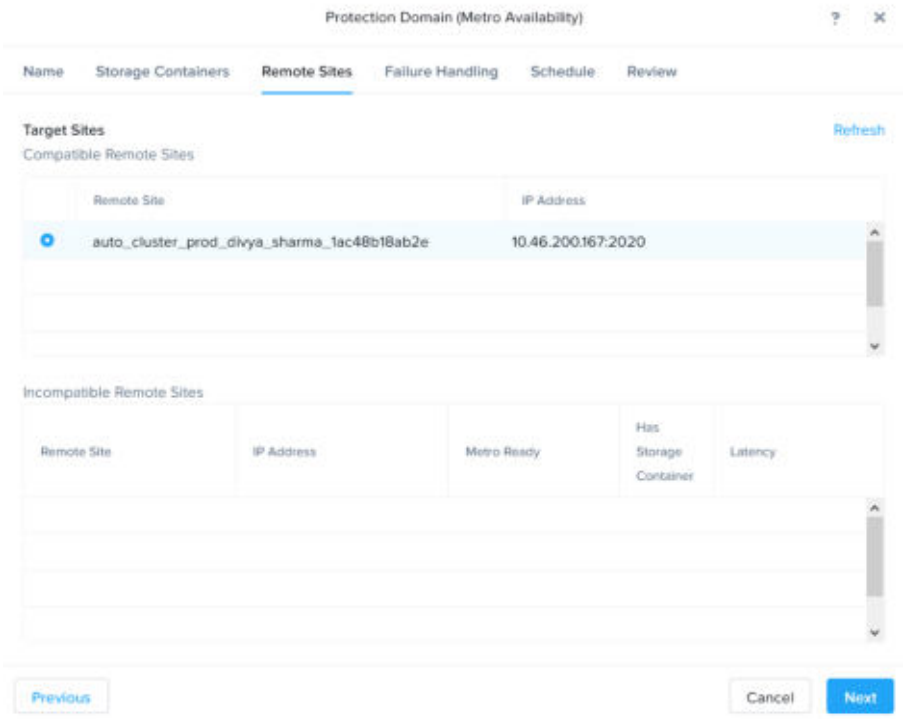

Figure 29: Protection Domain (Metro Availability): Remote Sites Screen

5. In the Failure Handling screen, select the failure handling mode (manually or automatically) when a network or other problem interrupts the connection between clusters.

   - (Recommended) Selecting the **Witness** option means the Witness VM can automatically distinguish a site failure from a network interruption between the metro availability sites and decide whether to failover in case of a site failure or network interruption. For details

on the recovery scenario, see "Recovering from a Storage-only Failure on the Primary Site" in Recovery Procedures (Witness VM installed) on page 131.

> Note: You must configure a Witness VM to use this option (see Metro Availability Witness Option on page 116).

- Selecting the **Automatic Resume** option means VM writes resume automatically (after 10 seconds by default). This option prefers VM availability. VMs will continue to write locally to the Active site only until replication is resumed in Prism.
- Selecting the **Manual** option means VM writes does not resume until a user disables metro availability manually or the problem is resolved. This option supports storage container consistency between the clusters at the cost of VM availability.

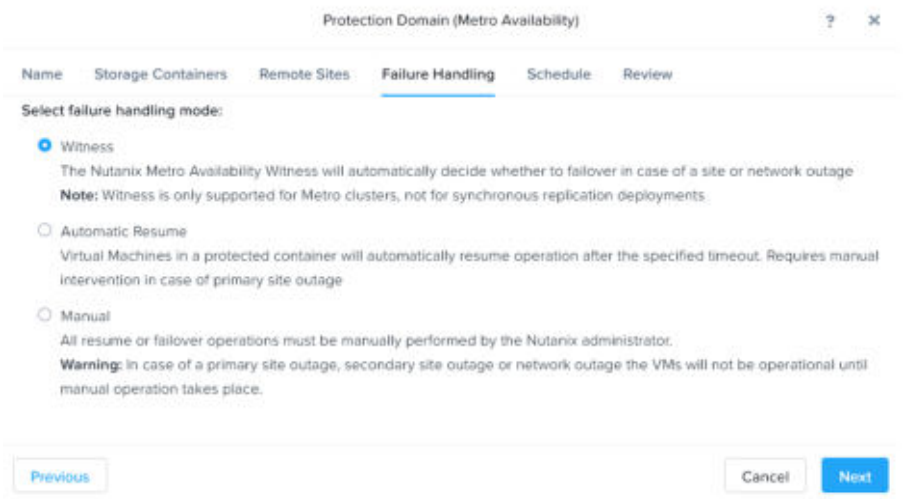

Figure 30: Protection Domain (Metro Availability): Failure Handling Screen

Metro uses VMCP APD implementation to handle storage-only failures.

> Note: You must manually enable APD on all the ESXi hosts in the cluster for automatic VM failover to work.
>
> - For each host, go to **Advanced System Settings** of the host on vSphere client and set **Misc.APDHandlingEnable** to **1**. (default setting is 0).
> - Enable APD and set the VM handling at the cluster level.
>
> Go to **Cluster Configure** > **vSphere Availability** > **Turn on vSphere HA** and set the following failure conditions and responses.
>
> - Set **All Paths Down (APD)** > **Failure and responses** to **Power off and restart VMs (aggressive)**.
> - Set **Response recovery** in **All Paths Down (APD)** > **Failure and responses** to **Disabled**.

> Note: If you do not enable VMCP APD, you must trigger a manual HA to recover from the storage-only outage, see Triggering a Manual HA Event for Storage-Only Outage (VMCP APD DISABLED) on page 138.

6. Optionally, on the **Schedule** tab, create a schedule.

   Apart from replicating data synchronously to the remote site, the active Nutanix cluster creates and replicates snapshots at a built-in schedule of 6 hours. These recovery points minimize the time spent in re-enabling Metro Availability. If you want a different snapshot frequency for the metro availability configuration, configure your schedule on the Schedule tab. A user-configured schedule with a replication interval of 6 hours or lower suppresses the built-in schedule. The built-in schedule continues to run if the user-configured schedule is

greater than 6 hours. Only metadata is replicated, so a more aggressive replication schedule does not incur a higher replication cost.

A schedule also applies to any additional remote sites that you select in the Remote Sites list. To these remote sites, data is replicated asynchronously based on the defined schedule.

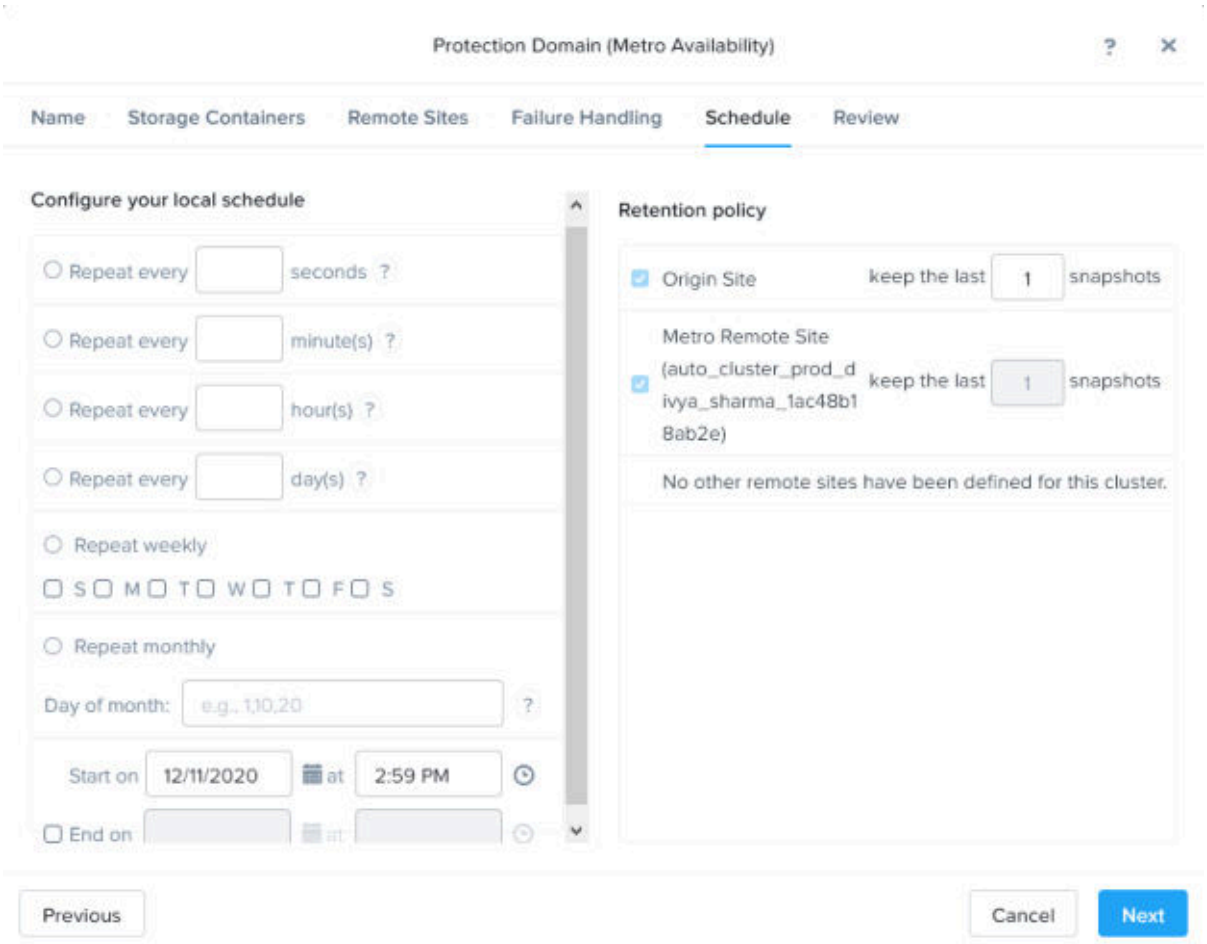

Figure 31: Protection Domain (Metro Availability): Schedule Screen

a. **Repeat every ## [minutes|hours|days]**: Click the appropriate circle for minutes, hours, or days and then enter the desired number in the box for the scheduled time interval.
   The interval cannot be less than an hour, so the minutes value must be at least 60.

b. **Repeat [weekly|monthly]**: Select which days the schedule should run.

   - If you select weekly, select the boxes for the days of the week the schedule should run.
   - If you select monthly, enter one or more integers (in a comma separated list) to indicate which days in the month to run the schedule. For example, to run the schedule on the first, tenth, and twentieth days, enter "1,10,20".

c. **Start on**: Enter the start date and time in the indicated fields.
   The default value is the current date and time. Enter a new date if you want to delay the schedule from starting immediately.

> Note: A storage container-level protection domain requires a system metadata scan to populate the file list. Any snapshots that started before the scan completes do not contain any data.

d. **End on**: To specify an end date, check the box and then enter the end date and time in the indicated fields.

   The schedule does not have an end date by default, and the schedule runs indefinitely unless you enter an end date.

e. **Retention Policy**: Enter the number of snapshots to save locally and at the remote sites.

   - Enter a number in the **Origin Site** line "**keep the last ## snapshots**" field. The default is 1. The number that you enter in the **keep the last ## snapshots** field is automatically copied in the **Target Site** line.
   - Enter a number of snapshots to save on the remote site in the "**keep the last ## snapshots**" field. This number can be different from the number that you have entered in the **Origin Site** line. This replication is an async replication and after the replication is completed, the protection domain is going to be available in the **Async DR** tab of the remote site.
   - The snapshots that are saved is equal to the value that you have entered in the **keep the last ## snapshots** field. For example, if you have entered the value **keep the last ## snapshots** field as 20, a total of 20 snapshots are saved. When the next (21st) snapshot is taken, the oldest snapshot is deleted and replaced by the new snapshot.

     The exception to this retention policy is on a metro PD when you have a schedule that is less frequent than the default automatic checkpoint frequency (6 hours). In such a case, the number of snapshots saved is equal to the value that you have entered in the **keep the last ## snapshots** field + 1. in such a case, for example, if you have entered the value **keep the last ## snapshots** field as 20, a total of 21 snapshots are saved. When the next (22nd) snapshot is taken, the oldest snapshot is deleted and replaced by the new snapshot.

     > Note: Nutanix recommends that start times be staggered across protection domains and across schedules.

7. In the Review screen, verify that the configuration is correct and then click the **Create** button.

   This configures the metro cluster but does not activate it.

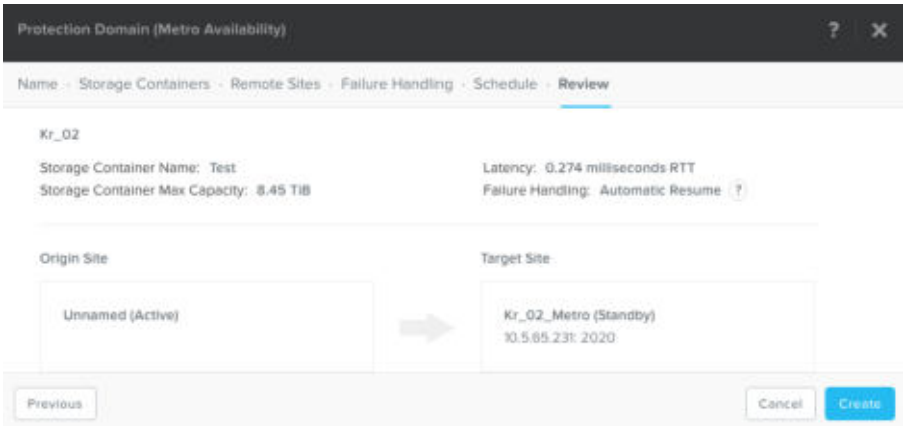

Figure 32: Protection Domain (Metro Availability): Review Screen

8. Click **Yes** in the confirmation dialog box.

   Metro availability is enabled on the storage container across two sites. You can modify an existing protection domain to create the backup schedules and also replicate the domain to local site and remote site. For more information, see Modifying Protection Domain (Metro Availability) on page 98.

Modifying Protection Domain (Metro Availability)

You can modify an existing protection domain by replicating the domain, updating the domain settings, or disabling the domain.

About this task

- **Taking snapshot of the Protection Domain**: You can take a snapshot of the protection domain locally and to one or more remote sites.
- **Disabling the Protection Domain**: You can disable the metro availability configuration on the selected storage container.
- **Updating the Protection Domain**: You can create the backup schedules for the metro availability configuration. You can also replicate a protection domain to one or more remote sites (sites that are not part of the metro availability configuration). This replication to the remote sites is an async replication and after the replication is completed, the protection domain is going to be available in the **Async DR** tab of the third site.

  Note: The backup schedule information in only available in **Schedules** tab of the primary site. If you perform a failover to the remote site, only then you can add a new schedule in the remote site.

Procedure

1. In the Data Protection dashboard, click the **Table** view.
2. Click the **Metro Availability** tab.
3. Select the protection domain that you want to modify.

4. To replicate a protection domain, click **Take Snapshot**. Select one or more targets where you want to replicate the protection domain.

   a. **Local**: Do nothing in this field.

      The box is checked by default to indicate a local snapshot is always created. If no remote site is checked, the snapshot is saved locally only. If one or more remote sites are checked, the snapshot is saved locally and at the specified remote sites.

   b. **Remote sites**: Check the boxes of the target remote sites.

      Note: If no remote site is checked, the snapshot is saved locally only.

   c. **Replication start time**: Select the start time from the pull-down list.

      Choices range from now to 48 hours from now.

   d. **Retention time**: Select the retention interval (how long the snapshot is saved on the remote sites) from the pull-down list.

      Choices range from one day to 3 months.

   e. When all the field entries are correct, click the **Save** button.

5. To disable a protection domain, click **Disable**.

   See the guidances for disabling metro availability before you perform disable operation. For more information, see Enabling and Disabling Metro Availability on page 101.

6. To update a protection domain, click **Update**.

   You can create the backup schedules for the metro availability configuration from the **Schedule** tab. You can also select the method to disable metro availability (manually or automatically) when a network or other problem interrupts the connection between clusters from the **Failure Handling** tab.

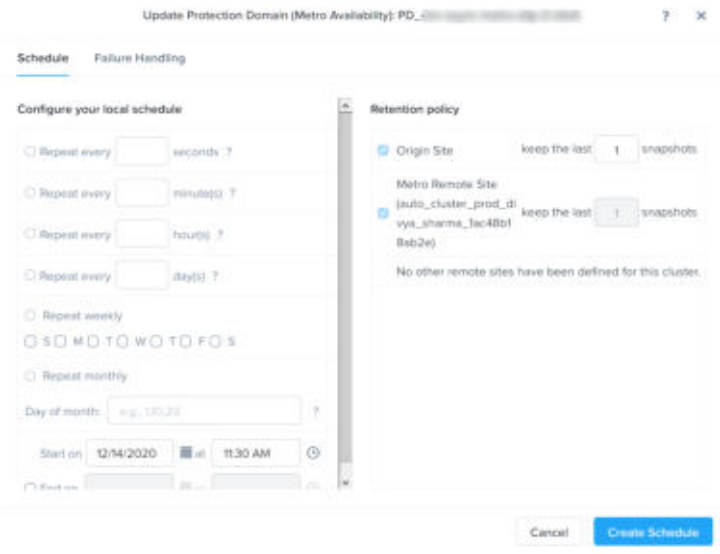

Figure 33: Updating the Protection Domain

7. To create a schedule, click the **New Schedule** to display the create schedule screen and then do the following in the indicated fields:

   a. **Repeat every ## [minutes|hours|days]**: Click the appropriate circle for minutes, hours, or days and then enter the desired number in the box for the scheduled time interval.

      The interval cannot be less than an hour, so the minutes value must be at least 60.

   b. **Occur [weekly|monthly]**: Select which days the schedule should run.

      - If you select weekly, select the boxes for the days of the week the schedule should run.
      - If you select monthly, enter one or more integers (in a comma separated list) to indicate which days in the month to run the schedule. For example, to run the schedule on the first, tenth, and twentieth days, enter "1,10,20".

   c. **Start on**: Enter the start date and time in the indicated fields.

      The default value is the current date and time. Enter a new date if you want to delay the schedule from starting immediately.

      > Note: A storage container-level protection domain requires a system metadata scan to populate the file list. Any snapshots that started before the scan completes do not contain any data.

   d. **End on**: To specify an end date, check the box and then enter the end date and time in the indicated fields.

      The schedule does not have an end date by default, and the schedule runs indefinitely unless you enter an end date.

   e. **Retention Policy**: Enter the number of snapshots to save locally and at the remote sites.

      - Enter a number in the **Origin Site** line "**keep the last ## snapshots**" field. The default is 1. The number that you enter in the **keep the last ## snapshots** field is automatically copied in the **Target Site** line.
      - Enter a number of snapshots to save on the remote site in the "**keep the last ## snapshots**" field. This number can be different from the number that you have entered in the **Origin Site** line. This replication is an async replication and after the replication is completed, the protection domain is going to be available in the **Async DR** tab of the remote site.
      - The snapshots that are saved is equal to the value that you have entered in the **keep the last ## snapshots** field. For example, if you have entered the value **keep the last ## snapshots** field as 20, a total of 20 snapshots are saved. When the next (21st) snapshot is taken, the oldest snapshot is deleted and replaced by the new snapshot.

        The exception to this retention policy is on a metro PD when you have a schedule that is less frequent than the default automatic checkpoint frequency (6 hours). In such a case, the number of snapshots saved is equal to the value that you have entered in the **keep the last ## snapshots** field + 1. in such a case, for example, if you have entered the value **keep the last ## snapshots** field as 20, a total of 21 snapshots

are saved. When the next (22nd) snapshot is taken, the oldest snapshot is deleted and replaced by the new snapshot.

f. When all the field entries are correct, click the **Create Schedule**.

You can create multiple schedules, update or delete the schedule you have created.

> Note: Nutanix recommends that start times be staggered across schedules.

8. To disable metro availability (manually or automatically) when a network or other problem interrupts the connection between clusters, click Failure Handling tab.

   - Selecting the **Manual** option means VM writes does not resume until a user disables metro availability manually or the problem is resolved. This option supports storage container consistency between the clusters at the cost of VM availability.

     See the guidances for disabling metro availability before you perform disable operation. For more information, see Enabling and Disabling Metro Availability on page 101.

   - Selecting the **Automatic Resume** option means VM writes resume automatically after 30 seconds. This option supports VM availability at the cost of storage container consistency between the clusters.

   - Selecting the **Witness** option means the Witness VM can automatically distinguish a site failure from a network interruption between the metro availability sites and decide whether to failover in case of a site failure or network interruption.

     > Note: The **Promote** operation is not run if only the clusters lose connectivity between them as long as the Witness is reachable.

9. To delete the protection domain, click **Delete** button.

10. In the confirmation box, enter DELETE word in the text box and then click the **Delete** button.

    The selected protection domain is deleted.

Enabling and Disabling Metro Availability

Following are the general guidances before you enable or disable metro availability on a storage container.

Enabling Metro Availability

1. All the VMs in a metro availability protection domain should be on the primary site.
2. Set the DRS rules to ensure that the VMs have affinity to the hosts of the primary site. The DRS rules should be **Fully Automated**.
3. Ensure that the default schedule on the metro availability protection domain is available and is not paused.

Disabling Metro Availability

1. Disabling metro availability is not allowed on the decoupled protection domain on the active site.
2. All the VMs in a metro availability protection domain should be on the primary site.
3. Set the DRS rules to ensure that the VMs have affinity to the hosts of the primary site. The DRS rules should be **Fully Automated**.
4. Datastore on the secondary site will be active, but it is in the read-only state. Do not promote the secondary site at this point as this can lead to the datastore being readable and

writeable on both the sites, which may lead to VM inconsistency or even corruption when a VM is live-migrated from the primary to the secondary site.

5. When Witness promotes the standby site, it causes the protection domain on the active site to go into decoupled state. Disable operation will not be allowed on a decoupled witnessed PD.

Restoring a VM from a Metro Availability Snapshot

Since metro availability uses storage container-level protection, you can only restore the entire contents of the snapshot. The restoration happens by using a separate path to prevent overwriting the protected entities.

About this task

CAUTION: You can restore a VM from a snapshot by using the following procedure only. Do not perform any manual operations with the files located in .snapshot directory to avoid cluster instability.

Note: You can only use nCLI to perform the restore operation.

If you perform the snapshot-restore operation when metro availability is enabled, the restored file path is also replicated to the secondary site. The restored file paths are like user-created files and they are synchronously restored in both the active and the standby sites.

You can only restore from the snapshots that were created after the most recent time you enabled the metro availability protection domain.

You can restore older snapshots when the metro availability is **Disabled**. The restoration can be done either on the active side or after the standby side is **Promoted** (in both these cases, the status of metro availability is **Disabled**).

Tip: Place the VMs that are frequently snapshotted and restored together in the same protection domain to avoid unnecessary storage space utilization. This also avoids the restoration and subsequent deletion of unnecessary VMs.

Procedure

1. Perform the restore operation on the metro protection domain only when the replication schedules are disabled. To suspend the schedules for a protection domain, do the following.

   a. Log into the Controller VM to know the replication status of the protection domain.

   ```
   ncli> pd suspend-schedules name=pd_name
   ```

   Replace *pd_name* with the name of the protection domain. This command returns a status of either true or false. True means all schedules for that protection domain have been suspended. If the status is false, repeat the command until the status is true.

   b. Verify the status of the protection domains.

   ```
   ncli> pd status
   ```

2. Log into the Controller VM and run the following nCLI command.

   ```
   ncli> pd restore-snapshot name=protection_domain_name snap-id=snapshot_id \
   ```

```
path-prefix=/path_restore files
```

Replace *protection_domain_name* with the name of the protection domain.

Replace *snapshot_id* with the ID of the snapshot.

Replace *path_restore files* with the path where you want to restore the files.

For example, for the protection domain *test_pd* with metro availability enabled on the storage container *test_storage_container* and the snapshot ID of *16388*, to restore the files to the path *test_restore* on the *test_storage_container*, use the following command.

```
ncli pd restore-snapshot name=test_pd snap-id=16388 path-prefix=/test_restore
```

The / in the restored path creates a folder test_restore in the test_storage_container. All the files that were in that snapshot are restored to the /test_restore location. You can then go to the folder to retrieve the necessary files by copying them to a different location.

If you omit the path-prefix parameter, a folder with a starting name of /Nutanix-Clone is created automatically.

3. Delete the unnecessary files from the restore folder after the restore operation is complete and all the desired files are restored, before you proceed with the next snapshot operation.
4. Resume the suspended schedules.

   ```
   pd resume-schedules name=pd_name
   ```

   Replace *pd_name* with the name of the suspended protection domain.

Protection Domain Failover and Failback (ESXi)

The failover and failback procedures are dependent on the nature of the failure and the site that has failed. You can perform the planned failover by using vMotion. If a disaster occurs, the recovery procedure is dependent on whether you have witness VM installed and enabled in your cluster or not.

> Note: If you have witness VM installed and enabled, see Recovery Procedures (Witness VM installed) on page 131 for recovery workflows.

Failing Over a Protection Domain (Planned Failover)

You can use vMotion to move virtual machines to the standby site. Migration with vMotion provides the highest uptime during planned failovers. Use vMotion to migrate only VMs. Do not use Storage vMotion.

Procedure

1. Configure the DRS setting to **Manual** to prevent live migration of VMs between sites.

   Instead of changing the DRS setting to **Manual** you could also modify the existing DRS rules to allow vMotion.
2. Perform vMotion of VMs from site A (active) to B (standby).

   > Note: If any VMs remain on the hosts of site A, they are disrupted during failover.
3. Change the DRS rules to have VM affinities to the hosts of the site B (standby).
4. After configuring VM affinities, change the DRS setting to **Fully Automated**.

5. Login to the web console of site B (standby), and select the metro availability protection domain and click **Promote**.

   If you log in to the web console of site A, the metro availability protection domain is in the **Decoupled** state. The datastore on the hosts of the site A should remain active, but it is in read-only state. It is not recommended to disable the protection domain on the site A as it can lead to the datastore being readable and writeable on both the sites, which can lead to VM inconsistency.

   > Note: You cannot disable any decoupled protection domains that have Witness enabled.

Failing Over a Protection Domain (Disaster or Site Failure)

Perform the following steps to failover a protection domain in case of disaster (unplanned failover) to remote site (B) from primary site (A). If the site A has failed, all the writes are stopped on the active storage containers of site B.

About this task

If the primary site has failed, you can promote storage containers on site B. After you promote the storage containers, the storage containers become active again to the VMware HA cluster. VMware HA updates its compatibility list and virtual machines that reside in that storage container automatically gets powered on. VMware DRS “should run on” affinity rules is overridden and those virtual machines are allowed to restart in the surviving site. “Must run on” affinity rules are enforced and has to be updated to allow virtual machines to restart on site B.

Procedure

1. Login to the Web console of site B.
2. Click **Promote**.

   HA restart process continues for DRS “should” rules.

   This operation does the following.

   - Restores all VM files from the last acknowledge write before site A failed.
   - Registers the VMs on the recovery site (site B) and power on the VMs.
   - Marks the failover site protection domain as active.

Re-enabling Metro Availability Configuration After Failover

After the original primary site (A) recovers, it comes back in a decoupled state. The datastores on the hosts of site A are active, but they are in a read-only state. Perform the following steps to re-enable metro availability after either a planned or disaster failover.

Before you begin

Ensure all the VMs are running on the promoted site (site B).

Procedure

1. Disable DRS to ensure VMs are not moved to site A.

2. Sign in to the web console of site B (original standby site but now the active site), go to the **Metro Availability** tab, select the target protection domain, and click **Re-enable**.

   A window for re-enabling metro availability appears. Read the instructions and verify that the specified action is what you want before clicking **Yes**. This step restarts the replication process (but now from site B to site A); allow the replication to finish.

   CAUTION:

   - For manual and automatic failure handling methods, perform the re-enable operation immediately after clicking **Disable** on the decoupled site A, because during this phase the containers are mutable and can host the VMs. If an unexpected host failure occurs, it can start an HA event. At that point, both sites could be writing to the VMs, which can lead to integrity issues (data inconsistency, duplicated IP addresses, and other issues). For protection domains with Witness-based failure handling, re-enable the metro availability configuration from the promoted site after the site is recovered.
   - Do not re-enable in the original metro availability direction, because it can lead to inconsistency due to stale data on the original primary site.

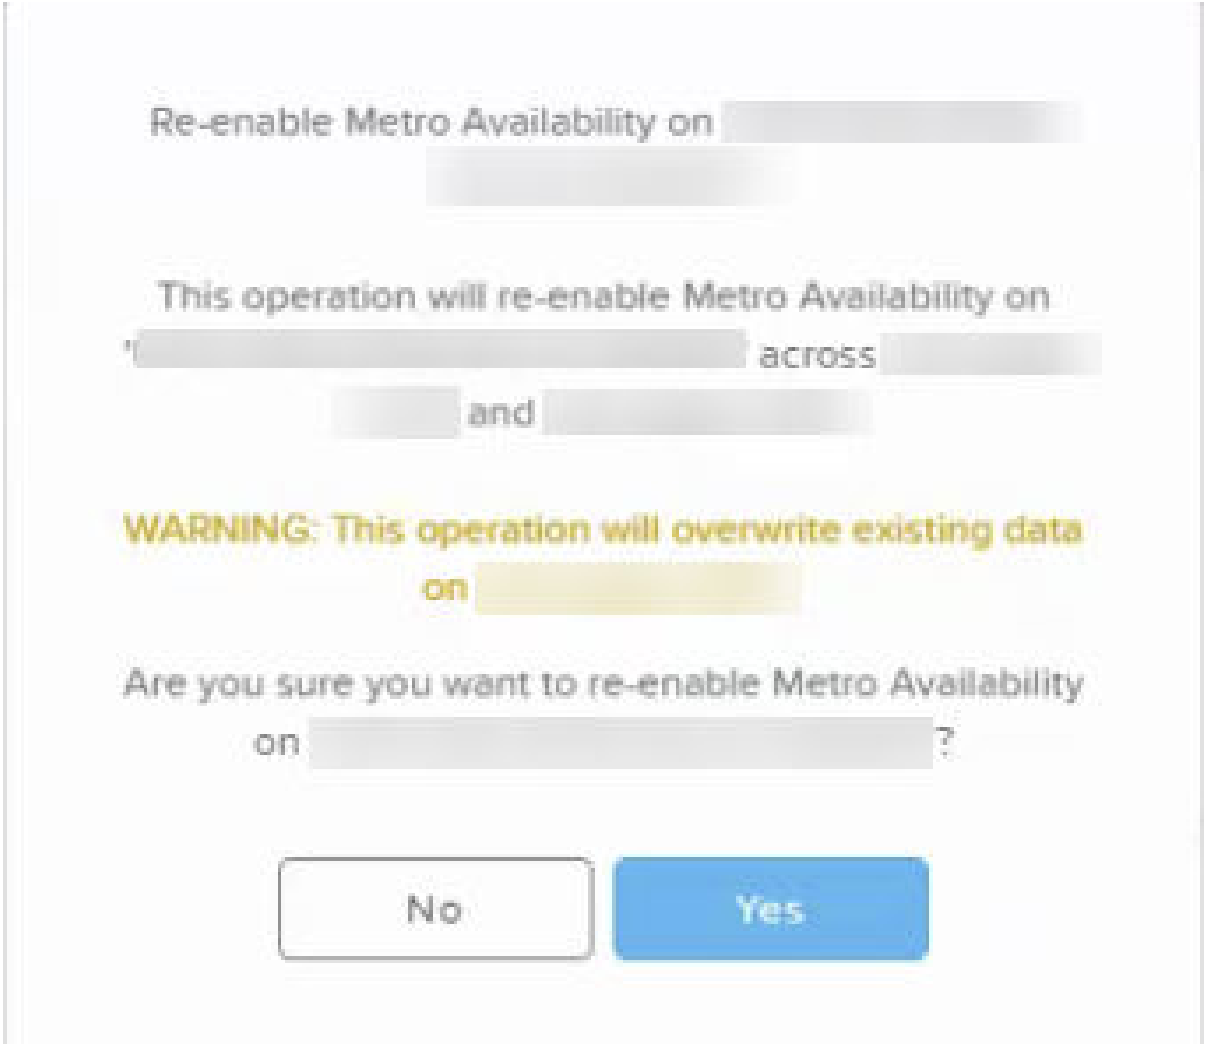

Figure 34: Re-enable Metro Availability Window

What to do next

> Note:
>
> Re-enabling metro availability starts an Async DR replication in the Prism web console.
>
> Navigate to the **Async DR** tab in Prism and you see an outgoing replication in progress even if you have not configured any asynchronous disaster recovery replications. This replication completes when the metro availability synchronization is complete.

This procedure re-enables metro availability but in the reverse direction from the original configuration.

```
Original metro availability state  {Site X [Metro Role, Metro Status]} :
Site A [Active, Enabled (in sync)] ---> Site B [Standby, Enabled (in sync)]

After failover:
Site A [Active, Decoupled] <--> Site B [Active, Disabled]

After reenabling metro availability:
Site A [Standby, Enabled (in sync)] <--- Site B [Active, Enabled (in sync)]
```

Site B is now the primary site, which works fine because metro availability does not care which site is primary. However, you can reset site A as the primary site. To make site A primary again, initiate another failover (see Failing Over a Protection Domain (Planned Failover) on page 103) and then repeat this re-enable procedure.

Handling Secondary Site Failure

Perform the following procedure to handle secondary site failure. You must have set the DRS rules to **Fully Automated** and configured VM affinities to the hosts of the primary site.

Procedure

1. If you have configured metro availability with the manual setting, then login to the primary site and disable metro availability by clicking **Disable**.
2. If there are any VMs running on the secondary site (which is against the best practices), these should failover to the hosts of the primary site by using vSphere HA functionality.

   If there is network failure between the primary and secondary sites, there might be some VMs on the secondary site in the paused state.

Handling Network Isolation (no Witness VM Installed)

Perform the following procedures depending on where the network isolation has occurred.

About this task

Context for the current task

Procedure

1. Network Isolation of the Primary Cluster
   a. Login to the Web console of the secondary site, and click **Promote**. The vSphere HA will start the VMs on the secondary site.
   b. Set the DRS rules so that all the VMs have affinities set to the secondary site.
   c. After restoring the network on the primary you need to destroy all the stale VMs that are still residing on the old primary as follows.
      - Run the `esxcli vm process list` on every host of the old primary and note down the world ID of these stale VMs.
      - Run the `esxcli vm process kill -w world_id -t force` to destroy the stale VMs. Ensure that you only destroy the VMs on the datastore of metro availability protection.
   d. Manually change the protection domain state from **Decoupled** to **Disabled** on the original primary site.
   e. Re-enable from the new primary to the intended secondary (old primary).
2. Network Isolation of the Secondary Cluster
   a. If you have configured metro availability with the manual setting, then login to the primary site and disable metro availability by clicking **Disable**.
   b. After restoring the network on the secondary re-enable from the primary site to the secondary site.

Reconfiguring Data Protection in a Metro Availability and Asynchronous Replication Configuration

If you are having a metro availability configuration from site A to site B and at the same time performing asynchronous replication from site A to site C for the same protection domain, and if site A fails you can use the following procedure to reconfigure data protection to keep replication going to site C.

Procedure

1. Login to the Web console of site B, and select the protection domain and click **Promote**.
2. Establish Async replication from B to C as follows.
   a. Create a remote site on B that points to C.
   b. Create a remote site of C that points to B.
   c. Create the snapshot schedule on site B by updating the protection domain of site B.

   Note: This schedule should point to the remote site as site C.

Metro Availability Protection Domain and Container Management (ESXi Clusters)

A metro availability (MA) protection domain (PD) consists of a specified container in the primary cluster linked to a container of the same name on the replication (remote) cluster. When MA PD is enabled, synchronous data replication occurs from the primary PD to the remote (replication) PD.

If you need to stop using metro availability, clean up the configuration and data in the corresponding containers. There are two ways to remove MA PD configuration when it is not needed anymore and to clean up the data in the corresponding container(s).

Deleting Metro Availability PD and Container Using vSphere and Prism Element

This is a safe procedure to remove the metro availability Protection Domain and cleanup the configurations. This procedure requires actual data movement where VMs need to be migrated with Storage vMotion from the metro container in the primary and remote clusters.

About this task

Perform the following steps using vSphere and Prism Element:

Procedure

1. Migrate all the VMs from the metro container to a new container using VMware's Storage vMotion. The VM files on the metro container are deleted when Storage vMotion completes.

   See VMware documentation for steps to migrate VMs using Storage vMotion.

2. Unmount the metro container from all ESXi hosts, on both the primary and replication (remote) Nutanix clusters.

   Use Prism Element on both the Nutanix clusters involved in metro availability.

   > Note: If the metro availability container is used for datastore heartbeats in the VMware HA cluster settings, select another container for heartbeats to allow the metro availability container to be unmounted.

3. Disable metro availability PD.

   Use Prism Element on the primary Nutanix cluster involved in metro availability to disable the metro availability PD.

Use Prism Element on the primary Nutanix cluster involved in metro availability for following steps.

4. Delete the metro availability PD.
5. Delete the metro container from the primary cluster.

Use Prism Element on the replication (remote) Nutanix cluster involved in metro availability for following steps.

6. Delete the metro availability PD.
7. Delete the metro container from the replication cluster.

Deleting Metro Availability PD and Container Manually

Use this procedure when you do not want to migrate VMs with Storage vMotion to other containers, and you need to retain the VMs in the current container (on either the primary or replication cluster) after disabling the metro availability Protection Domain.

About this task

Ascertain which cluster (primary or replication cluster) you want to retain the VM in the current container. Perform the following steps.

Procedure

1. To retain the VMs in the current container in the replication or standby cluster, perform planned failover of the primary cluster to make the metro PD in the replication cluster as the Active PD.

   > Note: Ensure that the cluster (whether primary or replication) where you want to retain the VMs in the current container becomes or remains as the Active cluster.

2. To ensure that the VMs on the ESXi hosts on the non-active (standby) cluster are not running, run the following NCC check command:

   ```
   nutanix@cvm$ ncc health_checks metro_availability_checks data_locality_check
   ```

   > Note: After planned failover in step 1, replication cluster becomes Active and primary cluster becomes non-active (standby). In such a case, ensure that no VMs on this standby cluster are not running.

   If the NCC check displays a warning and points to VMs running on the standby cluster, then fix the DRS rules. If DRS is disabled, then manually migrate the VMs to the ESXi hosts on the Active cluster.

   Verify that no VMs are running on the ESXi hosts on the standby cluster by running the NCC check again after fixing any warnings displayed in the previous run.

3. Disable metro availability PD on the Active cluster.

   Use Prism Element on the primary Nutanix cluster involved in metro availability to disable the metro availability PD.

Use Prism Element on the Active Nutanix cluster involved in metro availability for following steps.

4. Delete the metro availability PD.

Use Prism Element on the standby Nutanix cluster involved in metro availability for following steps.

5. Delete the metro availability PD.

Manually cleanup the metro container on the standby cluster.

Proceed with this step very cautiously since any mistake can lead to data loss.

If you opt to not to delete the container used for metro availability and instead only to delete the VM files and folders, ensure that this container is NOT used again for metro availability pairing with the previously configured site. Doing so leads to data inconsistency

> CAUTION:
>
> Do not attempt to delete the files from the NFS datastore from vCenter Server via browse datastore. In the case of Nutanix metro availability, the vCenter recognizes the containers with the same name in different AOS clusters as a single NFS datastore. See VMware KB1019286. When a datastore browser is used via vCenter, it selects one of the ESXi hosts with a datastore mounted to handle file operations. This means any attempt to delete files via vCenter datastore browser may be handled by ESXi host from the

side where you want to keep your VM/data, instead of the side where you want to clean up data. And this will lead to data inconsistency

Note: Ensure you are not logging on to the Active cluster where you need to retain the VMs in the current container.

6. To ensure that the VMs on standby site are not running and that ESXi can release it (not used for HA datastore heartbeats), unmount the container from all of the ESXi hosts on standby site using Prism Element.
7. Using SFTP protocol, port 2222 and Prism Element local admin user credentials in WinSCP, connect to a CVM IP on the standby cluster.

   Note: Ensure you are not logging on to the Active cluster where you need to retain the VMs in the current container.

8. In the WinSCP browser, open the container that needs manual cleanup and delete all the VM files and folders.

   Note: Deleting several files from the NFS datastore takes up to 30 minutes time to complete.

9. Delete the metro container on the standby cluster.

Metro Availability on Hyper-V 2016

A metro availability configuration on Hyper-V combines the existing metro availability concepts with the site-aware failover clustering feature introduced in Hyper-V 2016. A metro availability configuration provides for the following:

- Movement of running VMs between two Nutanix clusters hosting a stretched container.
- Promotion of a site hosting a container in standby mode to standalone mode without causing any disruption to the workloads running on the standby site.
- Hypervisor-provided HA and site-aware features to guarantee that, during normal operating conditions, the workloads stay on the site hosting the storage in the active mode.
- Workload failover to the promoted standalone (previously standby) site during disaster.

To meet these requirements, a metro availability configuration on Hyper-V includes the following configurations:

Failover Cluster

The failover cluster includes the hosts in the active and standby Nutanix clusters.

Sites

The failover cluster contains two fault domains of type site. One site contains the hosts from the active cluster and the other site contains the hosts in the standby cluster. A site enables you to group the hosts in a cluster into a single fault domain.

Metro Availability Configuration

Metro availability is configured between the active and standby clusters. To enable failover of VMs without interruption, the configuration uses a stretched storage container, which is achieved by creating a container with the same name on both clusters and configuring synchronous replication between the containers. The stretched storage container also makes storage migration unnecessary.

In addition to the container name being the same, the complete path to the container, as seen from SCVMM or Failover Cluster Manager, must be the same. An attribute named

metro SMB name is introduced to make the paths consistent and to uniquely identify a metro availability relationship between two clusters. For example, consider a stretched storage container named `ctr1` on two clusters named `cluster1` and `cluster2`. The UNC path to the containers on the two clusters would be `\\cluster1.domain.name\ctr1` and `\\cluster2.domain.name\ctr1`. To make both paths the same, you could specify a metro availability relationship name such as `cluster12`. With this relationship name, the UNC path on both clusters is `\\cluster12.domain.name\ctr1`.

Finally, you create a metro availability protection domain that creates and replicates snapshots on a defined schedule.

VM Affinity Rules

VMs using the stretched storage container must have site affinity, preferred owners, and possible owners configured. These settings specify which hosts must run the VMs in scenarios such as single node failure, complete site outage, and disaster drills or planned migration.

Summary: Preparing to Configure Metro Availability for Hyper-V

Before you begin

If the following tasks have not been completed, complete them before you configure metro availability. These procedures, unless mentioned otherwise, are described in *Hyper-V Administration for Acropolis*:

- Image the clusters with Hyper-V 2016 and a release of AOS that supports metro availability for Hyper-V. For information about imaging clusters, see the *Nutanix Field Installation Guide*.
- Join the clusters and hosts to the domain. See "Joining the Cluster and Hosts to a Domain".
- Create a failover cluster from one of the clusters that you want to add to the metro availability relationship. See "Creating a Failover Cluster for Hyper-V".
- In Failover Cluster Manager, add the hosts in the other Nutanix cluster to the failover cluster that you created in the previous step. For information about adding hosts, see the Microsoft documentation about adding nodes in Failover Cluster Manager.
- Optionally, register the Nutanix clusters with SCVMM. If you use SCVMM, whitelist the SCVMM on the clusters (see Configuring a Filesystem Whitelist in the *Prism Web Console Guide*).
- Create a container with the same name on each of the clusters.

About this task

To configure metro availability for Hyper-V, do the following:

Procedure

1. Configure a metro availability relationship between the Nutanix Clusters. See Configuring a Metro Availability Relationship Between the Nutanix Clusters on page 112.

2. If you use SCVMM, add the metro container you created in the previous step to the failover cluster in SCVMM.

   See "Adding Hosts and Storage to SCVMM Manually (SCVMM User Interface)" in *Hyper-V Administration for Acropolis*.

   You do not have to perform the entire procedure. Perform only the step that describes adding the file shares to the Nutanix cluster by navigating to **VMs and Services**. The file share you add must be of the form `\\new-metro-name\container-name`.

3. Create sites in the failover cluster and then map the Hyper-V hosts to the sites (use the `New-ClusterFaultDomain` and `Set-ClusterFaultDomain` PowerShell commands, respectively). For information about these commands, see the Microsoft documentation about fault domains.

   Verify the configuration by using the `Get-ClusterFaultDomain` PowerShell command.

4. Enable metro availability between the AOS clusters. For information about configuring metro availability, see Configuring a Protection Domain (Metro Availability) on page 90.

   When configuring the vStore name mapping, in both **Source VStore** and **Destination VStore**, select the container that you created on both the clusters.

5. Create VMs in the metro availability container on the active cluster.

   > Note: Make sure that the paths to the VMs in the metro container use the metro availability relationship name instead of the cluster FQDN.

6. Configure HA and affinity rules for each of the VMs in the metro availability container.

   a. In Failover Cluster Manager, configure the High Availability Role for all the VMs in the metro container.

   b. Configure the preferred owners. The preferred owners must be the hosts in the active cluster.

   c. Optionally, for the VMs in the metro availability container, configure the preferred site (site affinity) for each of the VMs.

      ```
      (get-clustergroup vm_name).PreferredSite="preferred_site"
      ```

      Replace *vm_name* with the name of the VM and *preferred_site* with the name of the preferred site.

   d. Configure the possible owners. The possible owners must be the hosts in the site in which the container is active and enabled.

7. Also configure possible owners for the VMs that are not in the metro availability container. The possible owners must be the hosts in the cluster on which the VMs reside.

Configuring a Metro Availability Relationship Between the Nutanix Clusters

Perform the following tasks on each of the clusters that you want to include in the metro availability configuration.

About this task

Procedure

Set up metro availability on the clusters.

a. Log on to any CVM in the cluster using SSH.

b. Add the metro availability relationship name to the Active Directory, update DNS and host file entries, and add an entry to the Kerberos keytab file for the metro availability relationship name.

```
nutanix@cvm$ setup_hyperv.py setup_metro
```

Follow the prompts to specify values for the following parameters:

- Domain administrator user name.
- Password of the domain administrator user name.
- Metro SMB name. Name for uniquely identifying the metro availability relationship between the two clusters for which you are configuring metro availability. All the VHDs in this metro availability relationship are registered with this name, and the path to the VHDs in the relationship is `\\new-metro-name\container-name\file-path`.

  Note: This name must be the same on both clusters.

- Metro SMB account password. Password for metro SMB object.

  Note: The password must be the same on both the clusters.

Performing a Planned Failover (Manual Live Migration)

A planned failover operation involves using Failover Cluster Manager to perform live migration of the VMs from the active site to the standby site and promoting the container on the standby site.

About this task

To live-migrate VMs from the active site to the standby site, do the following:

Procedure

1. In Failover Cluster Manager, under the failover cluster in the left pane, click **Roles**.
2. In the right pane, right-click the VM, go to **Move** > **Live Migration**, and then click **Select Node**.
3. In the **Move Virtual Machine** dialog box, do the following, select the host to which you want to migrate the VM, and then click **OK**.
   The VM is migrated to the selected host.
4. Login to the web console of the cluster to which you migrated the VM, select the metro availability protection domain, and then click **Promote**.
   The metro availability protection domain in the site from which you migrated the VMs is shown to be in the **Decoupled** state.

Performing an Unplanned Failover (Disaster)

About this task

To perform an unplanned failover, at the standby site, do the following:

Procedure

1. Sign in to the Nutanix web console.

2. In the **Data Protection** dashboard, click **Metro Availability**.
3. Click the protection domain, and then click **Promote**.
4. Read the message in the confirmation prompt, and then click **Promote**.
   The VMs fail over to the hosts in the site that has the container promoted. If the failure handling is set to manual, you must manually power on the VMs. If the failure handling is set to Automatic Resume, the VMs are powered on automatically.

Re-enabling Metro Availability (Failback)

When the primary site is recovered, the protection domain at the primary site is in the decoupled state. To reinstate the metro availability configuration to the state in which it was before the unplanned failover, re-enable metro availability from the site at which the VMs are running. This procedures uses the names Site A for the primary site, which experienced a disaster and has subsequently been recovered, and Site B, which was promoted during the duration of unavailability of Site A.

About this task

To re-enable metro availability, do the following:

Procedure

1. Disable the protection domain at Site A.
   a. Sign in to the Prism web console of Site A.
   b. On the **Data Protection** dashboard, click **Metro Availability**.
   c. Click the protection domain, and then click **Disable**.
   d. At the confirmation prompt, click **Yes**.
      The status of the container at Site A changes from `Decoupled` to `Disabled`.
2. Re-enable the protection domain at Site B.
   a. Sign in to the Prism web console of Site B.
   b. On the **Data Protection** dashboard, click **Metro Availability**.
   c. Click the protection domain, and then click **Re-enable**.
   d. At the confirmation prompt, click **Yes**.
      Data synchronization begins from Site B to Site A, and the status column displays the text `Synchronizing` until the process completes. Wait for the status column to display `Enabled (in sync)` before you proceed.
3. By using Failover Cluster Manager, live-migrate the VMs from Site B to Site A.
4. Promote the protection domain at Site A.
   a. Sign in to the Prism web console of Site A.
   b. On the **Data Protection** dashboard, click **Metro Availability**.
   c. Click the protection domain, and then click **Promote**.
   d. At the confirmation prompt, type `PROMOTE`, and then click **Yes**.
      The state of the container changes from `Enabled (in sync)` to `Disabled`.

5. Disable the protection domain at Site B.
   a. Sign in to the Prism web console of Site B.
   b. On the **Data Protection** dashboard, click **Metro Availability**.
   c. Click the protection domain, and then click **Disable**.
   d. At the confirmation prompt, click **Yes**.
      The state of the container at Site B changes from `Decoupled` to `Disabled`.
6. Re-enable the protection domain at Site A.
   a. Sign in to the Prism web console of Site A.
   b. On the **Data Protection** dashboard, click **Metro Availability**.
   c. Click the protection domain, and then click **Re-enable**.
   d. At the confirmation prompt, click **Yes**.
      Data synchronization begins from Site A to Site B, and the status column displays the text `Synchronizing` until the process completes. After the process completes, the status column displays `Enabled (in sync)`.

5

WITNESS OPTION

Witness option provides a separate VM called a Witness VM. Nutanix recommends the Witness option for metro availability deployments and two-node clusters. A Witness VM is a special VM that monitors the metro availability or the two-node cluster configuration health. The Witness VM resides in a separate failure domain to provide an outside view that can distinguish a site failure from a network interruption between the metro availability sites. The goal of the Witness option is to automate failovers in case of site failures or inter-site network failures.

For more information about adding a Witness to a two-node cluster, see Configuring a Witness (Two-node Cluster) on page 138.

For information about adding Witness to a metro availability configuration, see Conditions for Implementing Data Protection (Metro Availability) on page 83.

Metro Availability Witness Option

You have the option of adding a Witness to a metro availability configuration (see Conditions for Implementing Data Protection (Metro Availability) on page 83). For information about Witness VMs, see Witness Option on page 116.

> Note: You can deploy a Witness VM only on AHV and ESXi hypervisors. For information about Witness options with AHV Metro Availability, see the *AHV Metro (Witness Options)* in the *Leap Administration Guide*.

The main functions of a Witness include:

- Making a failover decision in the event of a site or inter-site network failure.
- Avoiding a split brain condition where the same storage container is active on both sites due to (for example) a WAN failure.
- Handling situations where a single storage or network domain fails.

In case of a failure of the primary site or the link between the sites, metro availability must be disabled and the target storage container (at the other site) must be promoted to active status.

> Note: When Witness is configured, it takes 120 seconds, from the time the communication with the primary site is lost, to promote the target storage container from standby to active status.

Metro Availability Failure Process (no Witness)

In the event of either a primary site failure (the site where the Metro storage container is currently active) or the link between the sites going offline, the Nutanix administrator is required to manually disable metro availability and promote the target storage container on the remote (or current) site to Active.

In case of a communication failure with the secondary site (either due to the site going down or the network link between the sites going down), metro availability does one of the following depending on the setting (automatic or manual):

- Automatic: The system automatically disables metro availability on the storage container on the primary site after a short pause if the secondary site connection does not recover within the specified time.
- Manual: The system waits for the administrator to manually take action.

Figure 35: Site Disaster (no witness)

Metro Availability Failure Process (with a Witness)

When a Witness is added, the process of disabling metro availability and promoting the storage container in case of a site outage or a network failure is fully automated. Thus, in the event of the Witness performs the functions of disabling metro and promoting the target storage container to active status.

The Witness functionality is only used in case of a failure. A failure of the Witness VM itself does not affect the VMs running on either site.

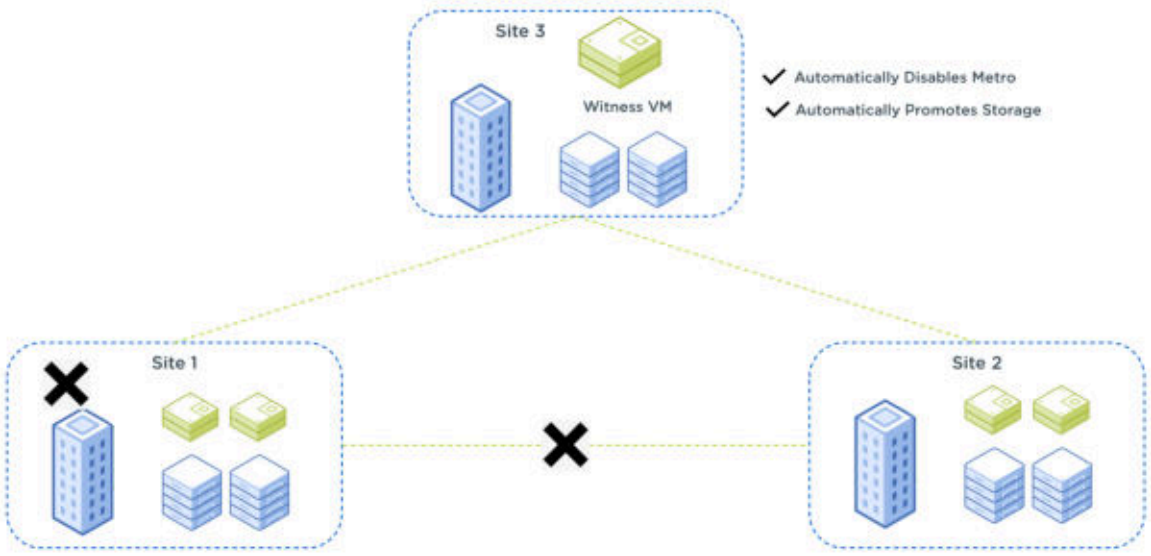

Figure 36: Site Disaster (with a witness)

Metro Availability Operational Modes

After adding a Witness, you can select from three metro availability operational modes: Witness Mode (new), Automatic Resume Mode, or Manual Mode. The metro availability response to a failure scenario varies depending on which operational mode is selected. The following table details the failure scenarios and the response behavior based on the operational mode.

Table 11: Operational Mode Response Behaviors

Failure Scenario	Witness Mode	Automatic Resume Mode	Manual Mode
Site 1 outage or complete network failure in Site 1	Automatically fails over to Site 2	Metro availability stops. You must manually activate the protection domain on Site 2 and wait for VMware HA to restart the VMs.	Metro availability stops. You must manually activate the protection domain on Site 2 and restart VMs.
Site 2 outage or complete network failure in Site 2	VMs continue to run on Site 1	VMs continue to run on Site 1	All VMs are paused. Administrative intervention required.
Connection loss between Sites 1 and 2	VMs continue to run on Site 1	VMs continue to run after timeout	All VMs are paused. Administrative intervention required.
Witness failure	No impact. However, ability to recover automatically in some scenarios is lost.	n/a	n/a
Connection loss between Witness and Site 1	VMs continue to run on Site 1	n/a	n/a
Connection loss between Witness and Site 2	VMs continue to run on Site 1	n/a	n/a
Connection loss between Witness and both Sites 1 and 2	VMs continue to run on Site 1	n/a	n/a
Failure at both Site 1 and Site 2	Metro availability stops. Administrative intervention required.	Metro availability stops. Administrative intervention required.	Metro availability stops. Administrative intervention required.

Failure Scenario	Witness Mode	Automatic Resume Mode	Manual Mode
Connection loss between Site 1 and Site 2 and between Witness and Site 1	VMs on Site 1 are paused. Site 2 will automatically be promoted and vSphere HA will automatically start the VMs. Administrative intervention required to hard-stop the VMs on site 1.	Metro availability is disabled, but VMs on Site 1 continue to run. Administrative intervention required.	All VMs are powered down. Administrative intervention required.
Storage-only outage on Site 1	The Witness promotes the storage on the Site 2. The storage on the Site 1 goes into an inaccessible state. If VMCP is enabled for APD, then VMs will failover to Site 2. Otherwise, refer to, see "Recovering from a Storage-only Failure on the Primary Site" in the Recovery Procedures (Witness VM installed) on page 131 topic.	All VMs remain unavailable until the storage outage is fixed or until the storage is manually promoted on site 2. After promoting storage on site 2, VMs have to be migrated manually by generating a VMware HA event.	All VMs will pause until, either the storage outage is fixed, or the storage is manually promoted on Site 2 using Prism. After promoting storage on site 2, VMs have to be migrated manually by generating a VMware HA event.
Storage-only outage on Site 2	Site 1 remains active and automatically breaks replication to Site 2 to be able to continue serving IO requests locally. The storage on Site 2 goes into a read-only state and any VMs that were running on hosts on Site 2 goes into a paused state. For details on the recovery scenario, see "Recovering from a Storage-only Failure on the Primary Site" in the Recovery Procedures (Witness VM installed) on page 131 topic.	Site 1 automatically breaks metro replication after a short pause.	All VMs will pause until, either the storage outage is fixed or the Metro availability is manually disabled from Site 1 in Prism.

Witness VM Requirements

There are several requirements when setting up a Witness VM:

- The Witness VM requires (at a minimum) the following:
  - 2 vCPUs
  - 6 GB memory
  - 25 GB storage
- The Witness VM must reside in a separate failure domain, which means independent power and network connections from each of the metro availability sites. It is recommended that the Witness VM be located in a third physical site. This site should have dedicated network connections to Site 1 and Site 2 to avoid a single point of failure.
- Communication with the Witness VM happens over port TCP 9440 and, therefore, requires that this port be open for the Controller VMs on any Metro cluster that uses the Witness VM.
- If an HTTP proxy is configured, you need to whitelist the Witness VM. See "Configuring HTTP Proxy" in the *Prism Web Console Guide* for instructions on how to whitelist a VM.
- Network latency betweens each metro availability site and the Witness VM must be less than 200 ms.
- The Witness VM may reside on any supported hypervisor, and it can run on either Nutanix or non-Nutanix hardware.

  Note: You can deploy a Witness VM only on AHV and ESXi hypervisors.

- You can register multiple (different) Metro cluster pairs to a single Witness VM. One Witness VM can support up to 100 instances, that is the combined total number of protection domains in the registered Metro cluster pairs and the number of registered two-node clusters. (A single Witness VM can support both metro availability and two-node clusters.)

  Warning: Ensure that there is no Unicode character (for example, &, !, #, $) in the name of the protection domains. If the name contains Unicode characters, both the primary and the standby clusters raise a false positive alert of Witness VM unreachability.

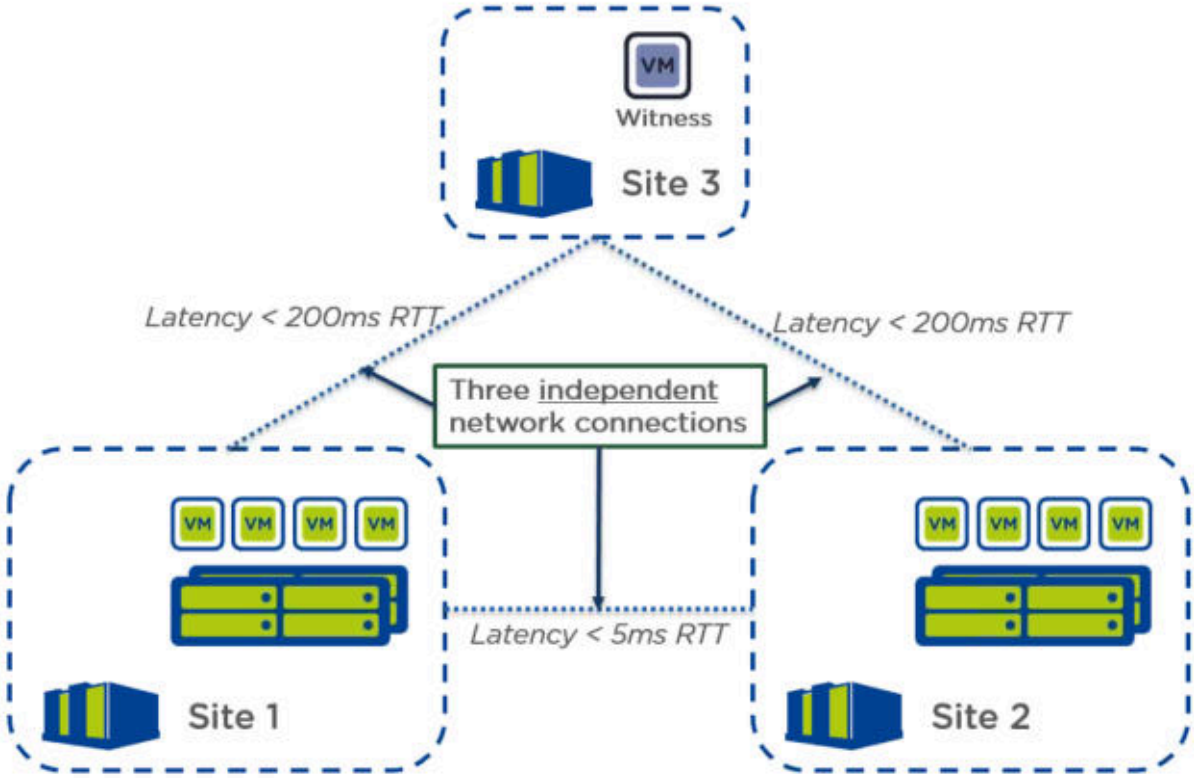

Figure 37: Witness Network Connections

Witness VM Dashboard

The Witness provides a dashboard of status information about all the registered storage containers. To view the Witness dashboard, do the following:

1. Open a browser to `https://witness_vm_ip_address`.
2. Log in to the Witness VM. (The user name is `admin` and the password `Nutanix/4u` by default.)

The Witness dashboard appears. A line appears for each storage container across all the registered protection domains. Each line includes the following fields:

> Note: The dashboard includes a pane on the left to select either the **2-Node Clusters** (see Witness VM Dashboard (two-node cluster) in the *Prism Web Console Guide*) or **Metro Clusters** view. Select **Metro Clusters** to view the storage container information in the right pane.

- **Storage Container Name**: Displays the name of the storage container.
- **Active Cluster**: Displays the name of the active (primary) cluster for this storage container.
- **Status**: Lists the current Metro Availability status of the storage container:
  - **In Sync**: Both sites (clusters) are in a steady state.
  - **Promoted**: One site got promoted.
  - **Disabled**: Metro Availability is broken between the primary and secondary sites.
- **Standby Cluster**: Displays the name of the standby cluster for this storage container.
- **PD Name**: Displays the name of the protection domain that includes the storage container.

To log out from the witness VM, go to the user icon menu (upper right of the screen) and select **Logout** from the pull-down list.

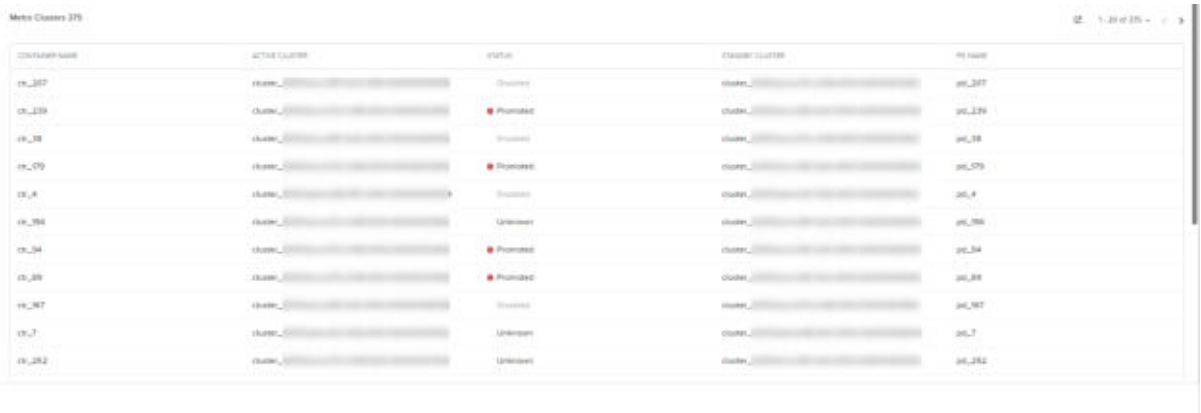

Figure 38: Witness Dashboard

Installing a Witness VM

Before you begin

Verify that the requirements are satisfied before attempting to install a witness VM (see Metro Availability Witness Option on page 116).

> Note: Metro availability and two-node clusters use the same witness VM. For more information about using a witness VM with a two-node cluster, see *Two-Node Clusters* in the *Prism Web Console Guide*. To upgrade a witness VM, seeUpgrading a Witness VM on page 125.

About this task

To install a witness VM, do the following:

Procedure

1. Download the Witness VM image file to your workstation as follows:
   a. Sign in to the Nutanix support portal (http://portal.nutanix.com) and go to the **Downloads** > **Witness VM** page.
   b. Click the download button for your hypervisor to download the file(s) for an initial deployment.
      » ESXi: `version#-witness_vm.ova`
      » AHV: `version#-witness_vm-boot.qcow2`, `version#-witness_vm-data.qcow2`, and `version#-witness_vm-home.qcow2`
2. Upload the VM image to the desired location.
   » If it is being uploaded to a Nutanix cluster, use the Prism image service to upload the VM image. For instructions on uploading an image file, see "Configuring Images" in the *Prism Web Console Guide*.
   » If it is being uploaded to non-Nutanix hardware, refer to the documentation from that vendor for instructions.

3. Create and power on the Witness VM.

   If the hypervisor is ESXi, see your ESXi documentation for instructions. If the hypervisor is AHV, do the following (see "Creating a VM (AHV)" in the *Prism Web Console Guide* for more information):

   > Note: The VM must have at least 2 vCPUs, 6 GBs of memory, and 25 GBs of storage.

   a. Logon to the web console and click the **Create VM** button in the VM dashboard.

   b. In the Create VM window, enter appropriate information in the name, compute details, and memory fields, and (if needed) click the **Add New NIC** button to create a network interface for the VM.

   c. Click the **Add New Disk** button and attach a boot disk.

      Enter the following information in the Add Disk window and then click the **Add** button:

      - **Type**: Select **Disk**.
      - **Operation**: Select **Clone from Image Service**.
      - **Bus Type**: Select **SCSI**.
      - **Storage Container**: Select the desired storage container from the list.
      - **Image**: Select `witness_vm_boot.img` (or whatever you named the boot image, originally downloaded as `version#-witness_vm-boot.qcow2`) from the list of images.
      - **Size (GiB)**: Enter the disk size (in GiB). This value is populated automatically from the selected image; do not change the value that appears in this field unless directed to do so by Nutanix customer support.

   d. Click the **Add New Disk** button again and attach the home disk.

      The steps are the same as for the boot disk except the image file name, which is `witness_vm_home.img` (or whatever you named the home image).

   e. Click the **Add New Disk** button again and attach the data disk.

      The steps are the same as for the boot disk except the image file name, which is `witness_vm_data.img` (or whatever you named the data image).

   f. When all the settings are correct, click the **Save** button to create the VM.

   g. Go to the VM dashboard table view, select the new VM (in the table), and then click the **Power on** action link (below the table) to start up the VM.

4. Open a console window and log in to the Witness VM.

   - The user name is `nutanix` and the password `nutanix/4u` by default when logging in through SSH.
   - The user name is `admin` and the password `Nutanix/4u` by default when logging in through Prism.

5. Assign a static IP address to the Witness VM as follows:

   a. Open the **ifcfg-eth0** file for editing.

      The following command opens the file using the **vi** editor:

      ```
      $ sudo vi /etc/sysconfig/network-scripts/ifcfg-eth0
      ```

   b. Update the `NETMASK`, `IPADDR`, `BOOTPROTO`, and `GATEWAY` entries as needed.

      ```
      NETMASK="xxx.xxx.xxx.xxx"
      IPADDR="xxx.xxx.xxx.xxx"
      BOOTPROTO="none"
      GATEWAY="xxx.xxx.xxx.xxx"
      ```

      - Enter the desired netmask value in the `NETMASK` field. (Replace *xxx.xxx.xxx.xxx* with the appropriate value.)
      - Enter the appropriate static IP address (usually assigned by your IT department) for the Witness VM in the `IPADDR` field.
      - Enter `none` as the value in the `BOOTPROTO` field. (You might need to change the value from `dhcp` to `none` if you employ DHCP. Only a static address is allowed; DHCP is not supported.)
      - Enter the IP address for your gateway in the `GATEWAY` field.

        Warning: Carefully check the file to ensure there are no syntax errors, whitespace at the end of lines, or blank lines in the file.

   c. Save the changes.

   d. Restart the Witness VM.

      ```
      $ sudo reboot
      ```

6. Enter the following command to create a Witness VM cluster:

   ```
   $ cluster -s vm_ip_address --cluster_function_list=witness_vm create
   ```

   The *vm_ip_address* is the Witness VM IP address.

   This completes the Witness VM installation.

7. (Optional) To change the Witness VM default password, enter the **passwd** command while logged in to the Witness VM and follow the prompts.

   Note: The following example is for changing the admin user password when logged in as the nutanix user. It is recommended that you change the password for both the admin and nutanix users.

   Witness VM registration requires the credentials of the admin user. If you change the password of the admin user after registering the Witness VM, you must register the Witness VM again.

   CAUTION: Password of the connected clusters must also be updated when you change the password of the admin user after registering the Witness VM.

   ```
   $ sudo passwd admin
   Changing password for user admin.
   Changing password for admin.
   (current) UNIX password: admin_password
   New password: new_admin_password
   ```

```
Retype new password: new_admin_password
passwd: all authentication tokens updated successfully.
```

This is sufficient if the password is changed just after creating the Witness VM. However, if the password is changed after the Witness VM is registered with the Metro Availability clusters, additional steps are required:

> Note: Registered clusters using the old password can lock the admin login access to the Witness VM for 15 minutes until the registered clusters get updated with the new password.

a. Log in to a Controller VM on the primary Metro Availability cluster using SSH and then use the nCLI to display information about the Witness VM registered with that cluster:

```
nutanix@cvm$ ncli
ncli> cluster list-witness

Id : id_number
Cluster Name : cluster_name
External IP Address : ip_address
Marked for Removal : value
```

b. Update the password for the registered Witness VM by entering the following command:

```
ncli> cluster update-witness id=id_number password=pw_string
```

The *id_number* is the ID number obtained from the **cluster list-witness** command in the previous step, and the *pw_string* is the same as what you entered to create the new password in the **passwd** (or **nuclei**) command.

c. Repeat these steps for the secondary Metro Availability cluster.

Upgrading a Witness VM

About this task

To upgrade a Witness VM to a newer version, do the following. You can upgrade the Witness VM at any time. It is not dependent on upgrade order when upgrading other cluster software components. The Witness VM version is independent of the cluster AOS and PC versions.

> Note:
>
> - A new Witness VM may not be generated for every new AOS release.
> - A Witness VM upgrade is blocked if you attempt to upgrade using an AOS or Prism Central package instead of the Witness VM upgrade package.
> - The Witness VM upgrade may not be successful if there is a time synchronization issue. The clock on VMs can be unreliable, so to avoid such issues it is recommended that you have three (or more) NTP servers configured for the cluster (see Cluster Time Synchronization) before attempting an upgrade. If the Witness VM is not located in a Nutanix cluster, see the appropriate NTP configuration documentation for your environment.

> CAUTION: Do not attempt to install a lower version such as installing 5.15.x upgrade package on 5.19.x. Such attempts are not blocked by the upgrade process. However, such attempts to install a lower version result in breaking the Witness VM.

Procedure

1. Log in to the Witness VM using SSH as the "nutanix" user and go to the `~/tmp` directory (`cd ~/tmp`).
2. Download the Metro Availability Witness Upgrade package from the Nutanix support portal using the `wget` command from the download directory (`~/tmp`):

   ```
   user@host$ wget https://nutanix_witness_package-release_version.tar.gz
   ```

   To determine the exact URL to use, log on to the Nutanix support portal (http://portal.nutanix.com), and go to the **Downloads** > **Witness VM** page, find the target Witness VM version line and copy the download link address into the `wget` command.

   > Note: If the `wget` command does not work, there might be a DNS issue. In this case you can download the installation bundle to your workstation and then copy it to the Witness VM's `/home/nutanix/tmp` folder using a program such as WinSCP, or you can configure a DNS server for the Witness VM by editing the `/etc/resolv.conf` file and adding a nameserver entry (`nameserver x.x.x.x`) as illustrated in the following example:
   >
   > ```
   > nutanix@cvm$ cat /etc/resolv.conf
   > # Auto generated by DnsConfig on Fri Apr  7 12:20:54 2017
   > nameserver 10.4.8.15
   > ```

3. Extract the bundle by entering the following command:

   ```
   user@host$ tar -xvzf nutanix_witness_package-release_version.tar.gz
   ```

   The installation bundle (`nutanix_witness_package-release_version.tar.gz`) is a large file, so delete it after the `tar` command completes successfully.
4. Upgrade the cluster (Witness VM) to the target AOS release by entering the following command:

   ```
   user@host$ ./install/bin/cluster -i install/ upgrade
   ```

   This command must be run from the same directory as the `tar` command in step 3 (`~/tmp`).

   The upgrade process typically takes a few minutes to complete. The Witness VM reboots after upgrade, and the services come up automatically.
5. If you upgraded a pre-AOS 5.1.1 Witness VM, the default `admin` user password is automatically reset to `Nutanix/4u`. When you log in through SSH as the `admin` user (with the `Nutanix/4u` password) the first time after the upgrade, you are prompted to set a custom password.

   > CAUTION: Once the custom password has been set, you need to update the Witness VM password on any Nutanix clusters using (registered with) that Witness VM. Clusters using the old password can lock the admin user access for 15 minutes until they get updated with the new password.

Registering a Witness VM

About this task

To register a Witness VM for a Metro Availability site, do the following:

Procedure

1. Open a browser and log in to Prism for the primary site metro availability cluster.

2. Click the gear icon in the main menu and then select **Configure Witness** in the **Settings** page.
   The Configure Witness window appears. Do the following in the indicated fields:

   Note: To unregister an existing Witness VM, click the **Unregister** button.

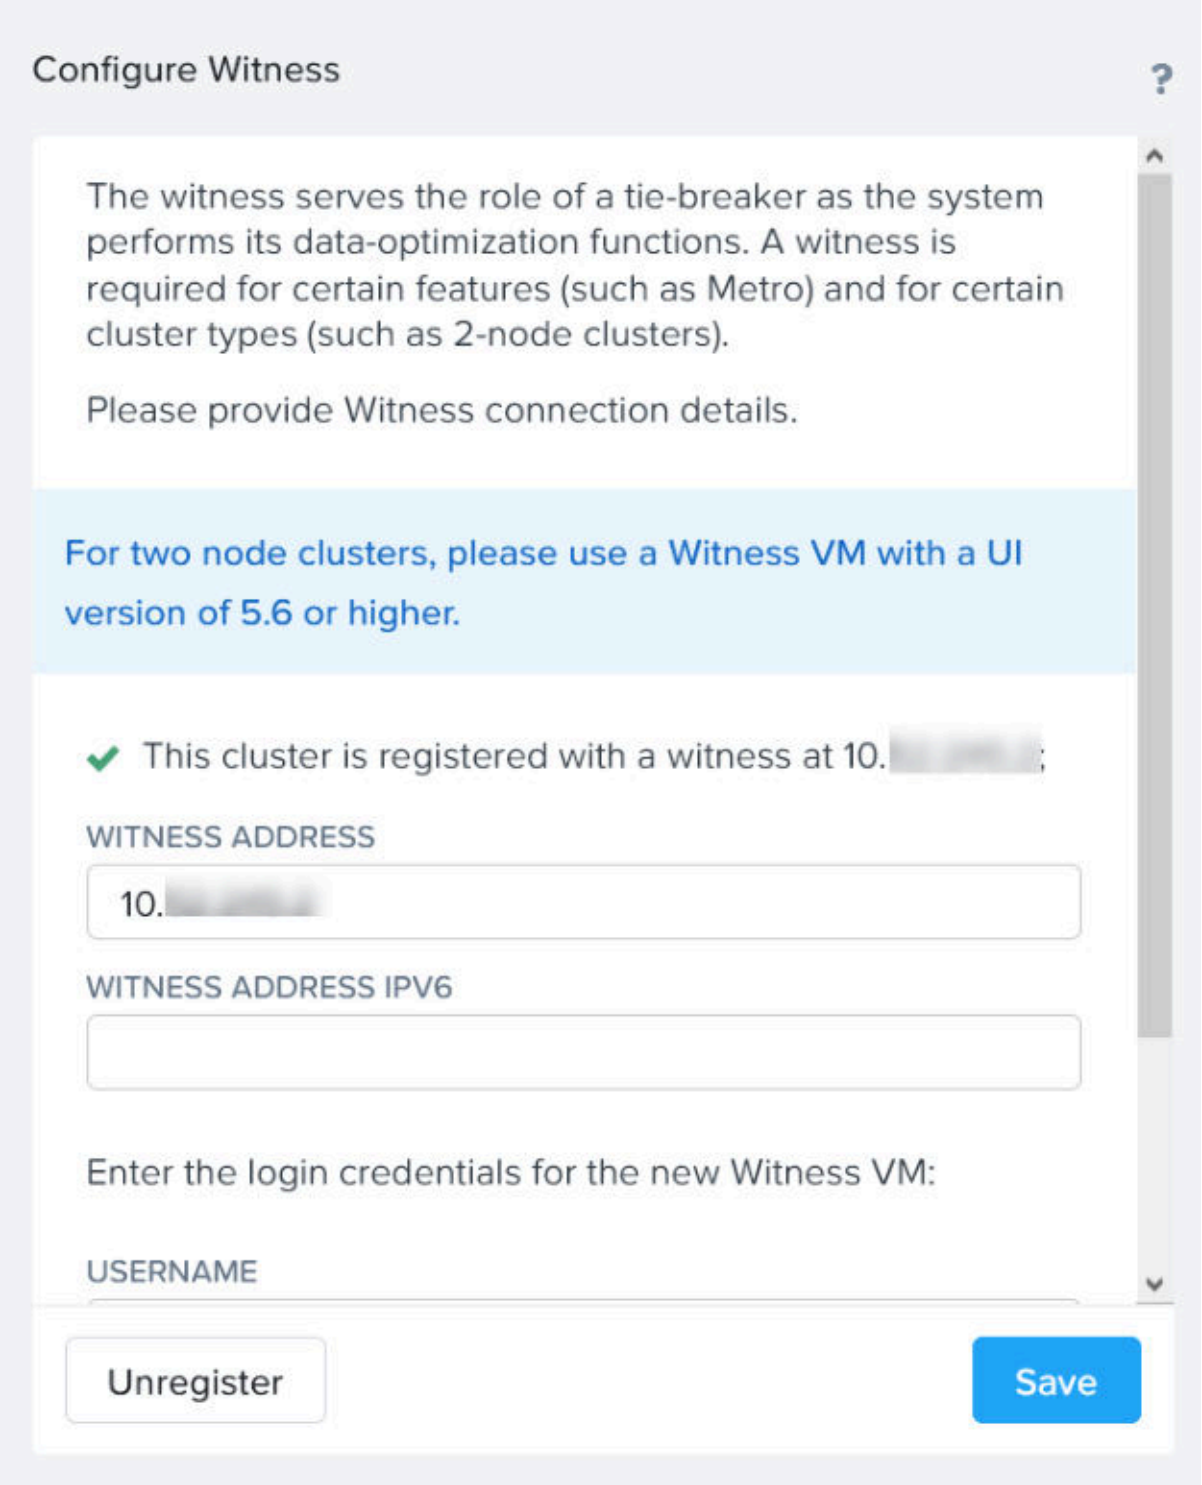

Figure 39: Configure Witness Window

a. **Witness Address**: Enter the Witness VM IP address.

b. **Username**: Enter the administrator user name (`admin`).

c. **Password**: Enter the user password (`Nutanix/4u` by default).

d. Click the **Save** button.

3. Log in to the standby metro availability cluster and repeat these steps.

> Note: You can register multiple metro availability pairs on the same Witness VM.

Changing the Witness VM for a Protection Domain

About this task

If it becomes necessary to change the Witness VM for a protection domain, for example if the Witness VM goes away permanently, do the following to associate a protection domain with a new Witness VM:

Procedure

1. Create the new Witness VM if it is not created already (see Installing a Witness VM on page 122).
2. On the primary metro availability cluster, change the failure handling setting from **Witness** to **Automatic** (or Manual) for each metro availability protection domain that uses the Witness VM and then disable these protection domains (see Modifying Protection Domain (Metro Availability) on page 98).
3. First unregister the current Witness VM and then register the new Witness VM on both the primary and secondary clusters (see Registering a Witness VM on page 126).
4. Change the failure handling settings back to **Witness** on all the protection domains changed in step 2 and then re-enable those protection domains.

Configuring Witness Mode

About this task

To configure a Witness for metro availability, do the following:

Procedure

1. Log in to Prism for the primary site metro availability cluster.
2. Select **Data Protection** from the view option pull-down list of the main menu and then do one of the following :

   » If you are creating a new metro availability configuration (see Configuring a Protection Domain (Metro Availability) on page 90), click the **Protection Domain** button and select **Metro Availability**.

   » If you are adding a Witness to an existing metro availability configuration (see Modifying Protection Domain (Metro Availability) on page 98), click the **Table** tab, click the **Metro**

Availability tab, select the primary site from the table list, and then click the **Update** button.

> Note: You need to first disable metro availability on the protection domain before changing the failure handling mode. Re-enable Metro Availability on the protection domain after changing the failure handling mode.

3. Complete the protection domain creation or update steps as normal. In the Failure Handling screen, select **Witness**.

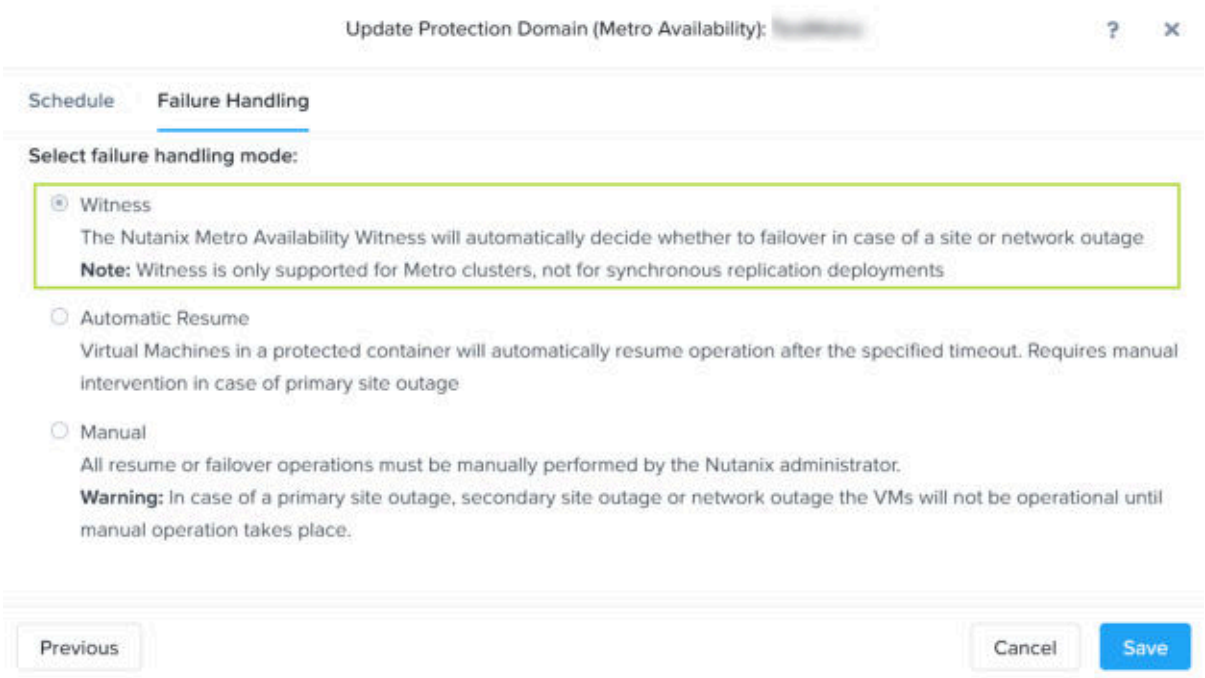

Figure 40: Failure Handling Screen (Protection Domain creation work flow)

Changing the IP Address of a Witness VM

About this task

To change the IP address of a Witness VM for metro availability, do the following.

Procedure

1. Log on to Prism on the primary site.
2. Change the recovery mode from Witness to Automatic (or Manual) for each metro availability protection domain that uses the Witness VM (see Modifying Protection Domain (Metro Availability) on page 98.
3. Unregister the Witness VM from the primary site (see Registering a Witness VM on page 126).
4. Log on to Prism on the secondary site and unregister the Witness VM.

5. Open a console window and log in to the witness VM.

   - The user name is **nutanix** and the password **nutanix/4u** by default when logging in through SSH.
   - The user name is **admin** and the password **Nutanix/4u** by default when logging in through Prism.

6. Run the following command to stop the witness VM services

```
nutanix@cvm$: cluster stop
```

7. Run the following command to destroy the Witness VM.

   Warning: This command removes all registered cluster from the Witness VM UI

```
nutanix@cvm$: cluster destroy
```

8. Assign a static IP address to the witness VM as follows.

   a. Open the **ifcfg-eth0** file for editing.

   The following command opens the file using the **vi** editor:

```
$ sudo vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

   b. Update the NETMASK, IPADDR, BOOTPROTO, and GATEWAY entries as needed.

```
NETMASK="xxx.xxx.xxx.xxx"
IPADDR="xxx.xxx.xxx.xxx"
BOOTPROTO="none"
GATEWAY="xxx.xxx.xxx.xxx"
```

   - Enter the desired netmask value in the NETMASK field. (Replace *xxx.xxx.xxx.xxx* with the appropriate value.)
   - Enter the appropriate static IP address (assigned by your IT department) for the witness VM in the IPADDR field.
   - Enter none as the value in the BOOTPROTO field. (You need to change the value from dhcp to none if you employ DHCP. Only a static address is allowed; DHCP is not supported.)
   - Enter the IP address for your gateway in the GATEWAY field.

     Warning: Carefully check the file to ensure there are no syntax errors, whitespace at the end of lines, or blank lines in the file.

   c. Save the changes.

   d. Reboot the witness VM.

```
$ sudo reboot
```

9. After the witness VM boots, logon to the console and re-create the witness again with the following command.

```
nutanix@cvm$: $ cluster -s NEW_vm_ip_address --cluster_function_list=witness_vm create
```

10. Reboot the witness VM.

```
$ sudo reboot
```

11. After the witness VM boots, logon to the Prism Element web console.
    The default user name is **admin** and the default password **Nutanix/4u** when logging on to Prism Element web console.
12. Register the Witness VM to the primary and the secondary site (see Registering a Witness VM on page 126).
13. Disable the target (changed in step 2) metro availability protection domains.
14. Change the recovery mode on the target metro availability protection domains back to Witness.
15. Re-enable the target metro availability protection domains.

Recovery Procedures (Witness VM installed)

The steps for recovering from a metro availability failure when a Witness is enabled depend on the nature of the failure. This section describes the steps needed (or not needed) when a failure occurs.

Recovering from a Primary Site Failure

When Site 1 (primary site) goes down, Site 2 (secondary site) detects the outage and acquires the lock from the Witness. The storage container automatically becomes active on Site 2. If the hypervisor has HA enabled, any protected VMs are failed over to Site 2. When Site 1 recovers, however, it cannot get the Witness lock and so cannot automatically make the local storage container active to re-enable metro availability. To re-establish replication (in this case from Site 2 to Site 1), do the following:

1. Log in to Prism on Site 2 and re-enable metro availability for the storage container.

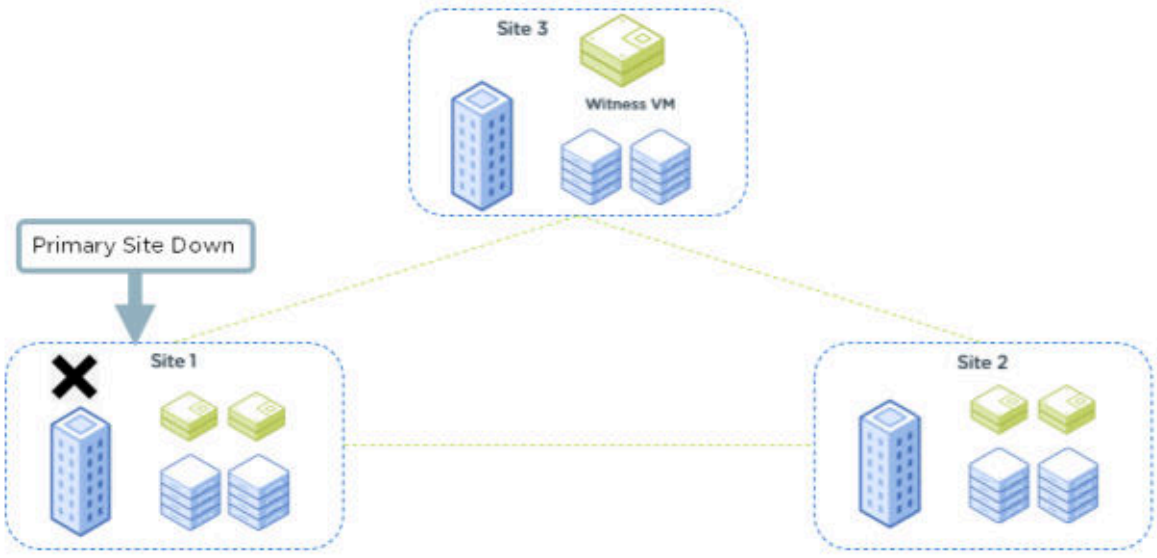

Figure 41: Metro Availability Primary Site Failure

Recovering from a Secondary Site Failure

When Site 2 (secondary site) goes down, metro availability on the protection domain is disabled automatically. When Site 2 is operational again, the administrator must re-enable the protection domain on Site 1 to resume replication to Site 2.

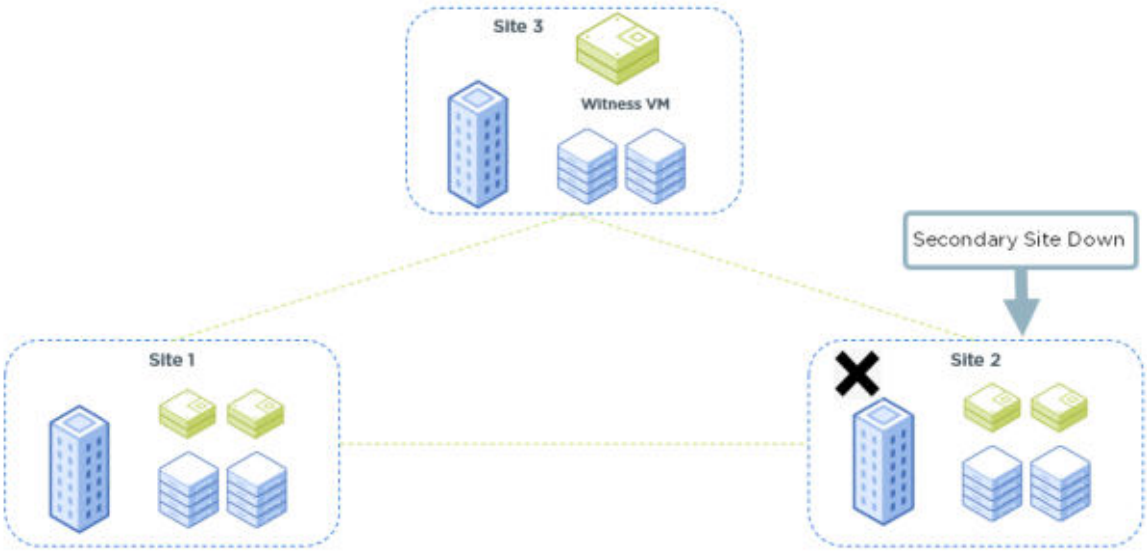

Figure 42: Metro Availability Secondary Site Failure

Recovering from a Storage-only Failure on the Primary Site

In the event of a storage-only outage on the active site, the Witness will promote the storage on the standby site. This causes the storage on the active site to go into a read-only state, as the protection domain is decoupled. A storage-only outage is defined as a condition where the ESXi hosts are up but have lost connectivity to the storage. Because the hosts are up, an HA event is not triggered.

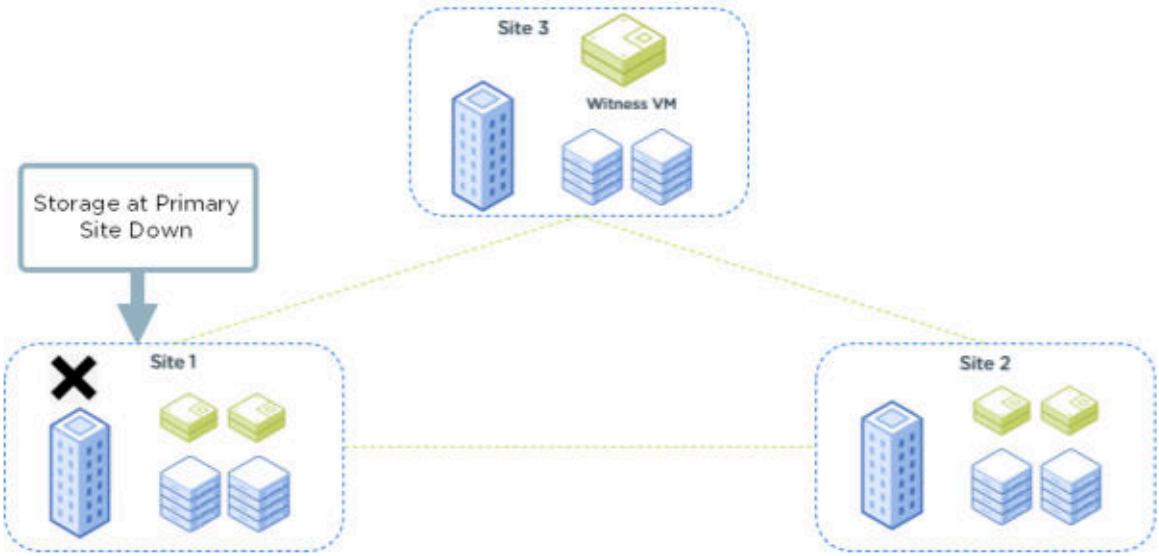

Figure 43: Metro Availability Primary Site Storage-only Failure

Metro (with a Witness) relies on a VMware HA event to power up the VMs on the standby site.

VM Component Protection helps protect vSphere environments from storage connectivity loss. When a host loses a storage device, VMCP marks it in one of the following states.

- PDL (Permanent Device Loss) - A device is marked as permanently lost if the storage array responds with a SCSI sense code marking the device as unavailable.
- APD (All Paths Down) - If the PDL SCSI code is not returned from a device, this is marked as All-Paths-Down and the ESXi host continues to send I/O requests until the host receives a response.

Metro uses VMCP APD implementation to handle storage-only failures. In case of storage-only failure on the primary site, metro detects the APD condition and automatically fails over the VMs on the affected site to the secondary site after the storage is promoted on the secondary site. This scenario makes the storage on the primary (affected) site unavailable for reads or writes.

Note: You must manually enable APD on all the ESXi hosts in the cluster for automatic VM failover to work.

- For each host, go to **Advanced System Settings** of the host on vSphere client and set **Misc.APDHandlingEnable** to **1**. (default setting is 0).
- Enable APD and set the VM handling at the cluster level.

  Go to **Cluster Configure** > **vSphere Availability** > **Turn on vSphere HA** and set the following failure conditions and responses.

  - Set **All Paths Down (APD)** > **Failure Response** to **Power off and restart VMs (aggressive)**.
  - Set **Response recovery** in the Failures and responses section of **All Paths Down (APD)** to: `Disabled`.

Note: If you do not enable VMCP APD, you must trigger a manual HA to recover from the storage-only outage, see Triggering a Manual HA Event for Storage-Only Outage (VMCP APD DISABLED) on page 138.

Recovering from a Witness Failure

When the Witness goes down, an alert is generated but metro availability is otherwise unaffected. When connection to the Witness is re-established, the Witness process resumes automatically. No administrator intervention is required.

If the Witness VM goes down permanently (unrecoverable), do the following to implement a new Witness VM:

1. Log in to Prism on the primary site.
2. Change the recovery mode from Witness to Automatic (or Manual) for each metro availability protection domain that uses the Witness VM (see Modifying Protection Domain (Metro Availability) on page 98).
3. Unregister the Witness VM from the primary site (see Registering a Witness VM on page 126).
4. Log in to Prism on the secondary site and unregister the Witness VM.
5. When a new Witness VM is installed and ready (see Installing a Witness VM on page 122), register that Witness VM on both the primary and secondary sites.
6. Disable the target (changed in step 2) metro availability protection domains.
7. Change the recovery mode on the target metro availability protection domains back to Witness.
8. Re-enable the target metro availability protection domains.

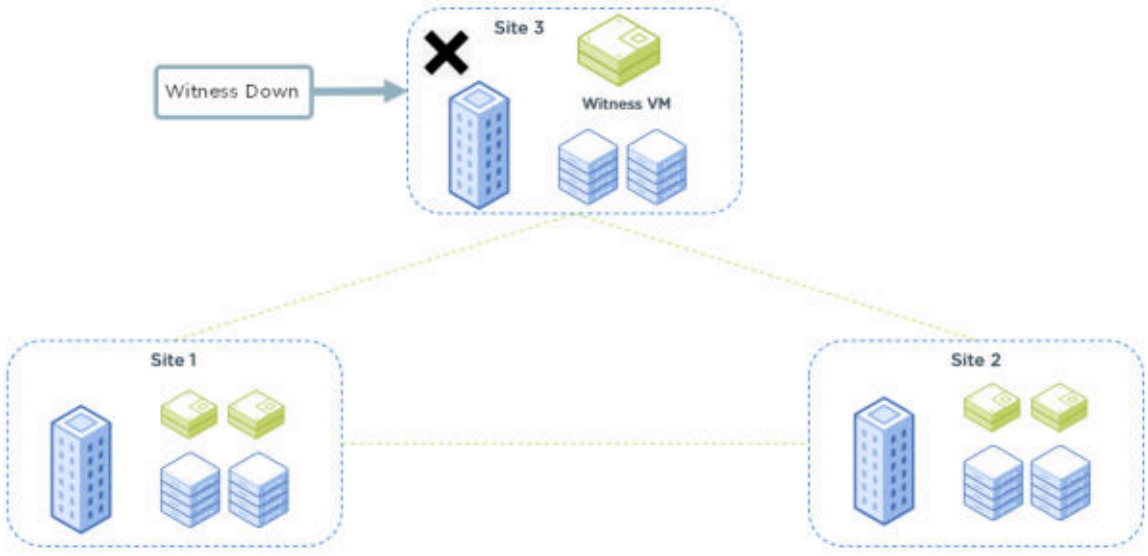

Figure 44: Witness Failure

Recovering from a Network Failure Between the Metro Sites

If there is a network interruption between Site 1 (primary) and Site 2 (secondary) but the network connections remain good between the sites and the Witness, both Site 1 and Site 2 attempt to acquire the Witness lock. There is a delay built into the Site 2 request, so Site

1 should get the lock first. When Site 1 obtains the lock, the protection domain on Site 1 is disabled automatically. The protection domain on Site 2 remains in "Standby" mode. When the network connection is re-established, the administrator must re-enable the protection domain on Site 1 to resume replication to Site 2.

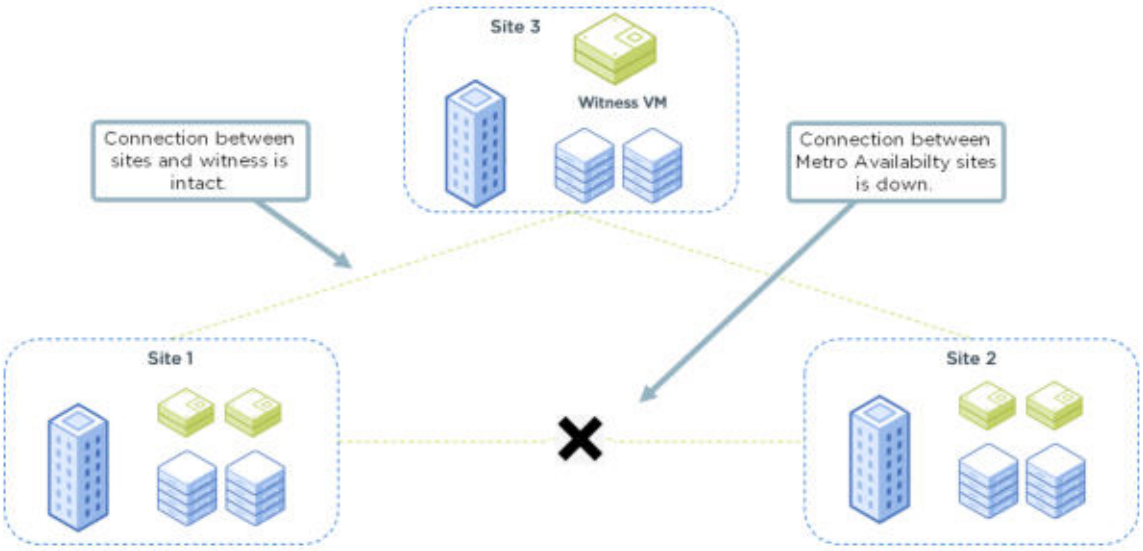

Figure 45: Network Failure between the Metro Sites

Recovering from a Double Network Failure (primary site isolated)

When there is a double network failure where the connection is broken both between the metro pair and between the primary site and the Witness, both Site 1 (primary) and Site 2 (secondary) attempt to acquire the Witness lock. Because Site 1 cannot connect to the Witness, Site 2 will acquire the lock after the built-in delay passes. Site 1 halts all I/O and the VMs are frozen. The storage container automatically becomes active on Site 2. If hypervisor HA (vCenter) is enabled and on the same network as the Witness VM, it restarts the Site 1 VMs on Site 2 after the secondary site is promoted. To re-establish replication (in this case from Site 2 to Site 1) when the connections are restored, do the following:

1. If the Site 1 VMs were restarted on Site 2, set DRS rules so that all VMs have affinities set to Site 2.
2. After restoring network connections to Site 1, kill all the stale VMs still residing on Site 1 as follows:
    1. Run the following command on every host in Site 1 and note down the world ID of the stale VMs:

        ```
        esxcli vm process list
        ```

    2. Run the following command to kill the stale VMs. Be careful to kill only the VMs identified in the previous step.

        ```
        esxcli vm process kill -w world_id -t force
        ```

3. Log in to Prism on Site 2 and re-enable metro availability for the storage container.

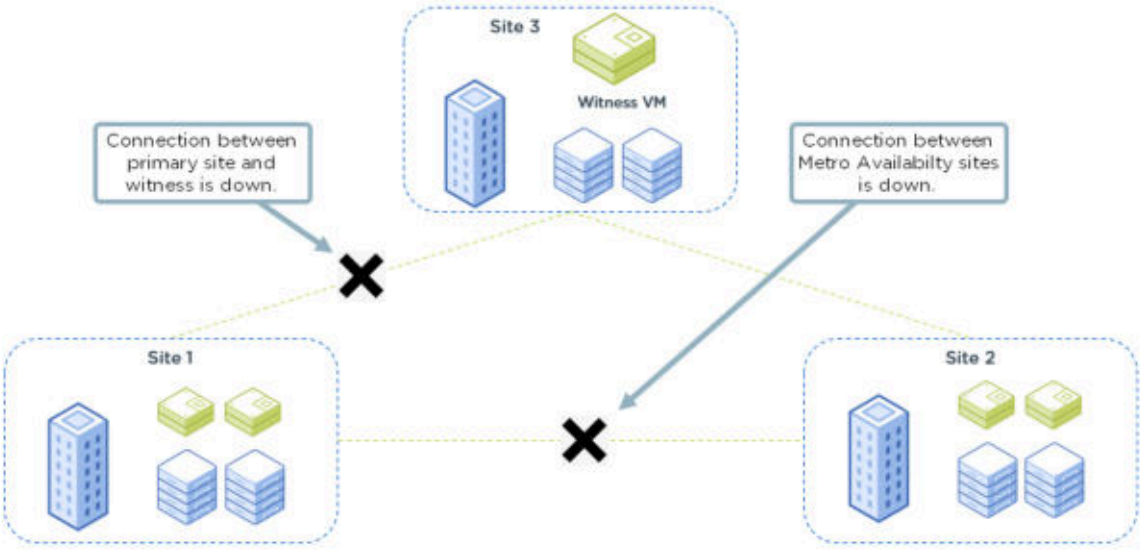

Figure 46: Network Double Failure (primary site isolated)

Recovering from a Double Network Failure (secondary site isolated)

When there is a double network failure where the connection is broken both between the metro pair and between the secondary site and the Witness, Site 1 gets the Witness lock and metro availability on the protection domain is disabled automatically. When the connections are restored, the administrator must re-enable the protection domain on Site 1 to resume replication to Site 2.

Figure 47: Network Double Failure (secondary site isolated)

Recovering from a Complete Network Failure (both sites isolated)

When there is a complete network failure where neither Site 1 nor Site 2 can connect to the other site or the Witness, neither site can get the Witness lock, so Site 1 halts all I/O and the VMs are frozen. When the connections are restored, the administrator must perform the same recovery steps described in the primary site double failure section.

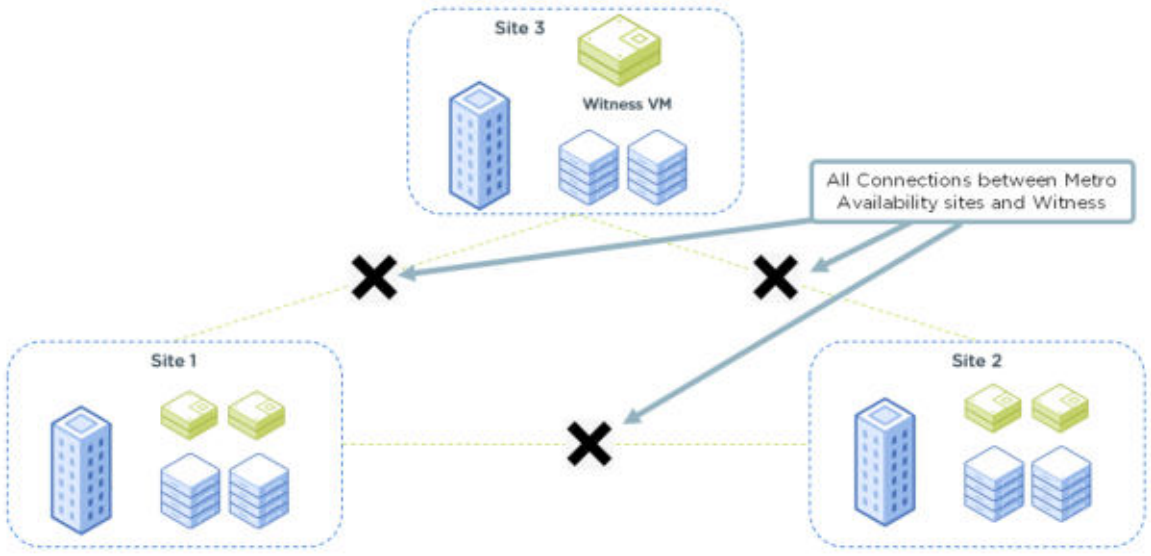

Figure 48: Network Complete Failure (both sites isolated)

Additional Recovery Considerations

In addition to the recovery procedures described for each failure scenario, the following recovery considerations might apply:

- If the Witness VM goes down and is no longer recognized when it is restarted or a new Witness VM is created, it might be necessary to re-register the Witness VM.
  - If there are any metro availability protection domains already in use, (1) change the failure handling to **Automation Resume** or **Manual**, (2) unregister the Witness VM, (3) register the Witness VM again with the default password on both clusters, and then (4) change the failure handling back to **Witness**.
  - If the password is changed after the Witness VM is registered, the Witness VM credentials on the cluster must be updated. To do this, enter the following nCLI command:

    ```
    nutanix@cvm$ ncli cluster update-witness id=id_number password=new_password
    ```

- After a site failure, consider the following:
  - Before promoting a remote cluster to be the primary site, first migrate any VMs on the primary cluster to the remote cluster.
  - It may take some time after the resynchronise of the protection domains before the migration can complete successfully.
  - Review the DRS rules to make sure they do not conflict with a recovery scenario. For example, do not create DRS rules for the VMs to run only on the nodes that belong to a

specific cluster (Site 1), because this would prevent the moved VMs from starting on the other cluster (Site 2) after an HA failover.

Triggering a Manual HA Event for Storage-Only Outage (VMCP APD DISABLED)

To recover from the storage-only outage situation when VMCP APD is not enabled, trigger a manual HA event as follows:

1. For the VMs protected by the protection domains promoted by the Witness, change the DRS rules on the vCenter from should run on site A (original site) to "must run" on site B (the promoted side).
2. If there are any protection domains that were not promoted automatically by the Witness, manually move the VMs in those protection domains (if possible) and then promote those protection domains on the standby side.

   > Note: Because of underlying storage connectivity issues, the ESXi host may be slow to respond, and vMotion may not work. If vMotion does not work, simply proceed to step 3.

3. Trigger an HA event for the impacted VMs hosted on the site A. You can do this in either of the following ways.
   1. Reboot the ESXi hosts on the site.
   2. Kill all the impacted VMs on hosts. First note the world ID of the impacted VMs by running the following command on every host in Site A.

      ```
      esxcli vm process list
      ```

      Then run the following command to kill the impacted VMs.

      ```
      esxcli vm process kill -w world_id -t force
      ```

      > Warning: This method may not work in all cases and is slower than rebooting the hosts mentioned in the step 3.a.

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 through the **Configure Witness** option, 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 in the Prism Web Console Guide*).

  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.

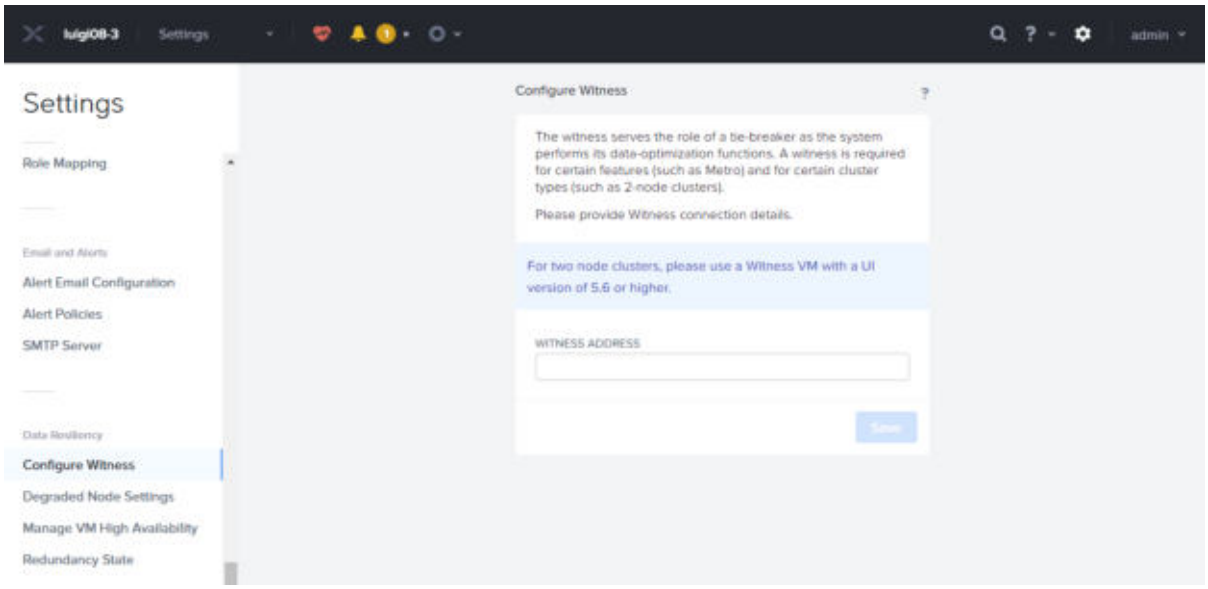

Figure 49: Configure Witness Window

- To configure a Witness when logging on to the Prism web console, do the following:

  a. Get the credentials of the target Witness VM (IP address, user name, and password).

  b. Log in to the Prism web console for the cluster as described in the *Logging into the Web Console* topic in the *Prism Web Console Guide*.

  c. When the Configure Witness screen appears, enter the Witness VM IP address, user name, and password in the indicated fields, and then click the **Apply & Continue** button.

     After clicking the **Apply & Continue** button, the cluster is registered with the Witness and the Prism web console home dashboard is displayed (see the *Home Dashboard* topic in the *Prism Web Console Guide*). If you check the **Skip for now (not recommended)** box,

you still go to the home dashboard, but the cluster is not registered with a Witness. This leaves the cluster vulnerable, so you should register it as soon as possible.

> Note: The Configure Witness screen appears only for clusters that require a Witness (such as a two-node cluster) and only when no Witness is configured.

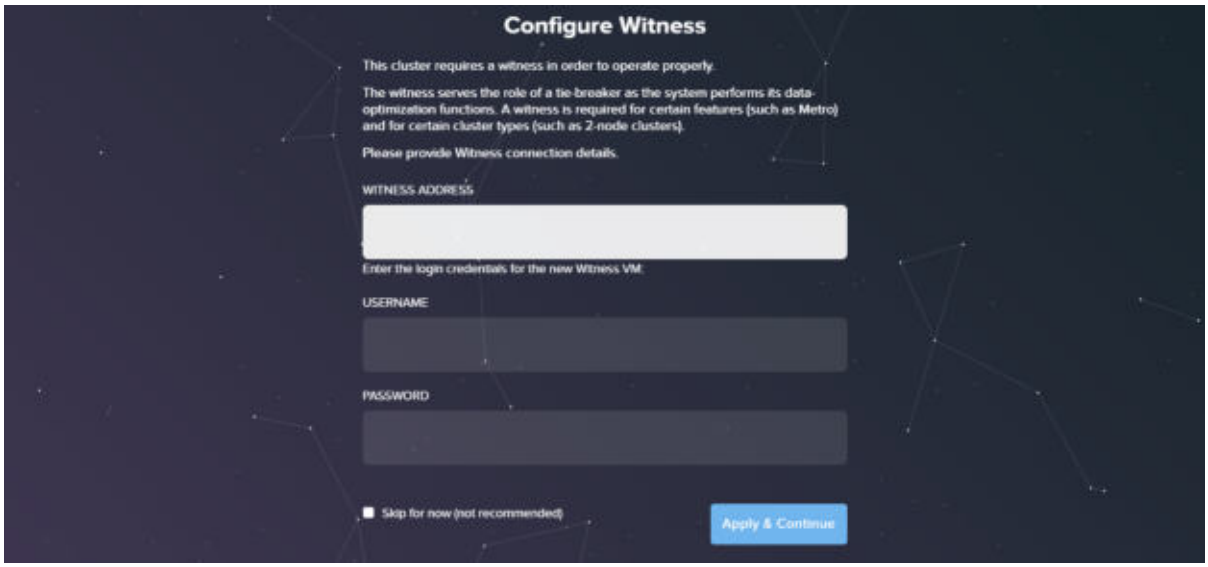

Figure 50: Configure Witness Screen

Witness VM Dashboard (Two-node Cluster)

The Witness provides a dashboard of status information about all the registered two-node clusters. (The two-node view appears only in a 5.6 or later version of the Witness.) To view the Witness dashboard, do the following:

1. Open a browser to `https://witness_vm_ip_address`.

2. Log in to the Witness VM. (The user name is **admin** and the password **Nutanix/4u** by default.)

Figure 51: Witness Login Screen

The Witness dashboard appears. In the **Witness Entities** pane on the left, click **2-Node Clusters** to display a table of registered two-node clusters in the right pane. A line appears for each cluster that includes the following fields:

- **Cluster Name**: Lists the name of the cluster.
- **Node IPs at Creation**: Displays the IP addresses of the two nodes in the cluster.
- **Cluster UUID**: Displays the cluster UUID value.
- **Status**: Indicates the current cluster status, either **OK** or a problem statement such as "one node down".
- **Last Keep Alive**: Lists the date and time when the current connection was established (continuous monitoring since that time).

To log out from the Witness VM, go to the user icon menu (upper right of the screen) and select **Logout** from the pull-down list.

Figure 52: Witness Dashboard (2-node cluster view)

6

SYNCHRONOUS REPLICATION (ESXI AND HYPER-V 2012)

Synchronous Replication replicates data synchronously between two sites in a metro availability configuration. In a disaster on any one site, real-time data is available on the other site.

The protection domain can be configured for metro availability or for synchronous replication depending on the following:

- Metro availability: The hypervisor is ESXi and the metro availability requirements for a single vSphere cluster across the sites and vSphere HA are satisfied. For more information on metro availability configuration, see Metro Availability tab section of Data Protection Table View on page 29 topic and to configure metro availability, see Configuring a Protection Domain (Metro Availability) on page 90.
- Synchronous replication: The hypervisor is either ESXi (not setup for metro availability) or Hyper-V.

> Note: Metro Availability and synchronous replication are supported across different hardware vendors (NX/SX, Dell, or Lenovo). You can also establish a Metro Availability or synchronous replication relationship between Nutanix NX/SX Series and other non-NX clusters (Dell or Lenovo). However, mixing of NX and non-NX nodes in the same cluster are not supported.

Configuring Synchronous Replication

1. To establish synchronous replication, configure the protection domain in metro availability configuration between primary and secondary site storage containers. For more information, see Configuring a Protection Domain (Metro Availability) on page 90.

Best Practices and Requirements

- Avoid whitelisting the primary and secondary clusters on each other.
- Ensure that no VMs are in **Critical State** on any storage container on any host. If a VM is in Critical State, the storage is blacklisted, thus impacting import of a VM on a host.
- Ensure that VMs are not deployed on the stretched storage container on the secondary site. If you want to deploy VMs on your secondary site, create a new storage container and then deploy the VMs. If you want to establish synchronous replication for these VMs, then this new storage container must be stretched to a remote site.
- Ensure that the VM's **AutomaticStopAction** property is set to `ShutDown` or `TurnOff` to protect a VM with Configuration Version 8.0 that is hosted on Hyper-V 2016.
- Ensure that, all the nodes in the Nutanix cluster must be part of the same ESXi host cluster and the network must be available on all the nodes in the Nutanix cluster at a given site (primary or remote).

Limitations

- All the limitations of metro availability configuration are applicable for synchronous replication. For more information, see Limitations section of Conditions for Implementing Data Protection (Metro Availability) on page 83 topic.

Failover Scenarios for Synchronous Replication in Hyper-V

VMs are Originally Registered and Running on the Primary Site and Failure Occurred on the Primary Site

About this task

If the VMs are originally registered and running on the primary site and a failure occurs on the primary site, then all the VMs can be started on the secondary site.

Procedure

1. Login to the Web console of secondary cluster.
2. Go to **Data Protection** > **Table** > **Metro Availability**. The status of protection domain is displayed as Remote Unreachable.
3. Select the protection domain and click **Promote** to make the protection domain on secondary independent of the primary.

   Promoting the secondary site unblocks all the VMs running on the secondary site for the stretched storage container.
4. Register the VMs on the secondary site. For more information about registering the VMs on the primary or secondary site, see Registering a VM on a Hyper-V Cluster on page 146.
5. Start the VMs.

VMs are Originally Registered and Running on the Secondary Site and Failure Occurred on the Primary Site

About this task

The secondary metro availability is in remote unreachable state and VMs are in hung state as storage is not accessible. All the VMs can be started on the secondary site.

Procedure

1. Login to the Web console of secondary cluster.
2. Go to **Data Protection** > **Table** > **Metro Availability**. The status of protection domain is displayed as Remote Unreachable.
3. Select the protection domain and click **Promote** to make the protection domain on secondary independent of the primary.

   All the VMs should start functioning on secondary site after restart.

VMs are Originally Registered and Running on the Secondary Site and Failure Occurred on the Secondary Site

About this task

All the VMs can be started on the primary site.

Procedure

1. Login to the Web console of primary cluster.
2. Register the VMs on the primary site. For more information about registering the VMs on the primary or secondary site, see Registering a VM on a Hyper-V Cluster on page 146.
3. Start the VMs.

Failback Scenarios for Synchronous Replication in Hyper-V

Failback VMs that are Currently Running on the Secondary Site but were Originally Registered on the Primary Site

About this task

When the primary site recovers, failback VMs back to the primary site that are currently running on the secondary site. These VMs were originally registered on the primary site when the failure on the primary site had occurred.

Procedure

1. Login to the Web console of recovered primary cluster and go to **Data Protection** > **Table** > **Metro Availability**.

   The status of the protection domain is in Decoupled state.
2. Select the protection domain and click **Disable**.

   > Note: You should not re-enable the protection domain from here because VMs might be running on secondary.
3. Unregister the VMs from the recovered primary cluster. For more information about unregistering the VMs on primary or secondary site, see Unregistering a VM on a Hyper-V Cluster on page 147.
4. Login to the Web console of secondary cluster and go to **Data Protection** > **Table** > **Metro Availability**.
5. Select the protection domain and click **Re-enable**. This ensures that the latest content of your VMs is synchronously replicated to the recovered primary site.
6. Click **Yes**. Wait for the operation to finish. On completion of the operation, the protection domain state changes to Enabled (In Sync). The resynchronization time varies on the data that needs to be replicated.
7. Shutdown the VMs from your current primary site (originally the secondary site).
8. Select the protection domain and then click **Disable** from your current primary site.
9. Login to the Web console of your current secondary site (recovered primary site) and go to **Data Protection** > **Table** > **Metro Availability**.

10. Select the protection domain and then click **Promote**.
11. Register the VMs on this site. For more information about registering the VMs on the primary or secondary site, see Registering a VM on a Hyper-V Cluster on page 146.
12. Start the VMs.
13. Unregister the VMs from where the VMs were originally replicated.
14. To establish the metro availability configuration again, login to the Web console of primary cluster and go to **Data Protection** > **Table** > **Metro Availability**.
15. Select the protection domain and click **Re-enable**.
16. Click **Yes**. Wait for synchronization to complete. After the synchronization is completed, the original metro availability configuration is established.

Failback VMs that are Currently Running on the Secondary Site and were Originally Registered on the Secondary Site

About this task

When the primary site recovers, failback VMs that were running on the secondary site when the failure on the primary site had occurred. This procedure establishes original synchronous replication configuration.

Procedure

1. Login to the Web console of the recovered primary cluster and go to **Data Protection** > **Table** > **Metro Availability**.
2. Select the protection domain and click **Disable**.
3. Login to the Web console of secondary cluster and go to **Data Protection** > **Table** > **Metro Availability**.
4. Select the protection domain and click **Re-enable**.
5. Click **Yes**. Wait for synchronization to finish.
6. Shut down the VMs from secondary site.
7. Select the protection domain and click **Disable**.
8. Login to the Web console of primary cluster and go to **Data Protection** > **Table** > **Metro Availability**.
9. Select the protection domain and click **Promote**.
10. Select the protection domain and click **Re-enable**.
    Now you can start the VMs on the secondary site.

Failback VMs that are Currently Running on the Primary Site but were Originally Registered on the Secondary Site

About this task

When the secondary site recovers, failback VMs back to the secondary site that are currently running on the primary site. These VMs were originally registered on the secondary site when failure on the secondary site had occurred.

Procedure

1. Unregister the VMs from the recovered secondary site.
2. Login to the Web console of primary cluster and go to **Data Protection** > **Table** > **Metro Availability**.
3. Select the protection domain and click **Re-enable**.
4. Shut down the VMs on the primary site.
5. Unregister the VMs from the primary site.
6. Register the VMs from the secondary site.

Failing Over a Protection Domain Manually in Synchronous Replication (Planned Failover)

Perform the following procedure to perform the manual failover to the secondary site if you have configured synchronous replication.

Procedure

1. Power off the VMs from the primary site.

   Note: Ensure that the container of the primary site does not have any VMs that are running on it.

2. Promote the secondary site.

   The secondary site becomes active. At the same time an alert is displayed on the primary site to disable the metro availability.

   CAUTION: If any schedules are enabled on the primary site, ensure that you first disable or remove the schedules.

3. Register the VMs on the secondary site.
4. Remove the VMs from the primary site.
5. Disable metro availability configuration from the primary site.
6. Power on the VMs on the secondary site.
7. Re-enable the metro availability configuration from secondary site to the primary site.

   This step resynchronizes the metro availability configuration and deletes the data of the primary site container. The roles are successfully reversed, the secondary site will become primary and primary site will become as secondary.

VM Registration and De-registration

Registering a VM on a Hyper-V Cluster

You can register the VMs on a Hyper-V cluster by performing following procedure.

Procedure

1. Generate a compatibility report for the VM that needs to be imported by using the `Compare-VM` cmdlet.

2. Fix incompatibilities that may occur because of incorrect configuration of Nutanix cluster name in the UNC paths (virtual hard disk path) by using the correct Nutanix cluster name.

```
> $report= Compare-VM -Path configuration_file_path -Register
```

Replace *configuration_file_path* with the path to the configuration file of the virtual machine to be compared.

To view the incompatibility report type, `> $report.Incompatibilities | Format-Table -AutoSize`

In the following example, the SMB path incompatibility should be fixed in the configuration XML file of VM to match the destination cluster.

```
> $report= Compare-VM -Path \\winne-smb\Mixed_vms\L4_IO\Virtual Machines'
\A7351476-4096-91E7-424A57B3FD08.xml -Register
> $report
```

```
VM                      :Microsoft.HyperV.PowerShell.VirtualMachine
OperationType           :ImportVirtualMachine
Destination             :WINNIE-1
Path                    :\\winne-smb\Mixed_vms\L4_IO\Virtual Machines
\A7351476-4096-91E7-424A57B3FD08.xml
SnapshotPath            :\\winne-smb\Mixed_vms\L4_IO\Snapshots
VhdDestinationPath      :
VhdSourcePath           :
Incompatibilities       :(40010)
```

To view the incompatibility report type, `> $report.Incompatibilities | fl *`

```
Message                 :Virtual Hard Disk file not found.
MessageId               :40010
Source                  :Microsoft.HyperV.PowerShell.HardDiskDrive
```

Incompatibility occurred while registering a VM on the secondary site (in the example, site WINNIE) and **Virtual Hard Disk file file not found - 40010 message** is displayed. This occurred because the VM was registered originally on the primary site and not on the secondary site. You must fix this issue before proceeding.

3. After all the incompatibilities are fixed, import the virtual machine by using the **Import-VM** cmdlet.

Unregistering a VM on a Hyper-V Cluster

You can unregister the VMs on a Hyper-V cluster by performing following procedure.

About this task

> Note: To unregister a VM you need to delete the VM. Hence, ensure that the VM has been registered successfully and all the VM files are present on the other site. Also, ensure that the protection domain state on the web console of the site is in the Disabled state.

Procedure

1. Stop the VM.
2. If the VM is part of the failover cluster, remove the VM from the cluster by using the **Remove-ClusterGroup** cmdlet.
3. Delete the VM by using the **Remove-VM** cmdlet.
4. Verify that the VM got unregistered from the host by using **Get-VM** cmdlet.

Deployment Scenario: Transferring a Synchronous Replication Setup to a Standby Site

If you have a synchronous replication configuration running at one site (the production site), you can use a combination of NearSync and asynchronous disaster recovery to keep the similar standby site up-to-date and then perform failover from the production site to the standby site when the primary cluster at the production site becomes unavailable due to a disastrous event or planned maintenance event.

The following diagram illustrates such a deployment.

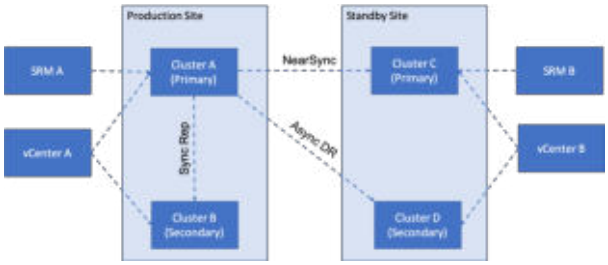

Figure 53: Production and Standby Sites for Transferring a Synchronous Replication Setup

In this setup, all the clusters run ESXi clusters at the production site are registered with a vCenter server. The primary cluster is registered with an SRM server. This setup is mirrored at the standby site.

Data from the primary cluster, Cluster A, is synchronously replicated to the secondary cluster, Cluster B. Concurrently, NearSync replication is used to replicate the same data to Cluster C (primary at the standby site) and asynchronous replication is used to replicate that data to Cluster D. The purpose of these additional replications is to reduce the time required for promoting the entire standby site to a production site that eventually performs synchronous replication between Cluster C and Cluster D.

Minimum Requirements

- SRA for SRM version 2.5

Referenced Documentation

All the documents that are referenced in this procedure are available on the Nutanix Support Portal.

Initial Configurations for Setting Up the Primary and Standby Sites

To set up the primary site, do the following:

1. Register the two clusters with vCenter (in the diagram, vCenter A). The clusters must belong to separate vCenter clusters. See *Creating a Nutanix Cluster in vCenter* in the *vSphere Administration Guide for Acropolis (using vSphere HTML 5 Client)*.
2. Configure the primary Nutanix cluster to work with the SRM server (in the diagram, SRM A). See the Nutanix Storage Replication Adapter (SRA) for Site Recovery Manager Administration Guide.
3. Ensure that the storage containers that you want to replicate have the same name on all clusters.
4. On all clusters, configure remote sites for replication. For SRM functionality, add vstore mappings to the recovery site configuration between cluster A and cluster C. Do not add vstore mappings for clusters B and D.

5. On Cluster A, configure a protection domain for synchronous replication between Cluster A and Cluster B. See Synchronous Replication (ESXi and Hyper-V 2012) on page 142.
6. On Cluster A, on the same protection domain that you created for synchronous replication, configure the following:

   - A NearSync schedule for replication between Clusters A and C. Specify an RPO in the range of 20 seconds through 15 minutes.
   - An asynchronous schedule between clusters A and D.

   See Data Protection with Asynchronous Replication (One-hour or Greater RPO) on page 42.

To set up the secondary site, do the following:

1. Register the two clusters with vCenter (in the diagram, vCenter B). See *Creating a Nutanix Cluster in vCenter* in the *vSphere Administration Guide for Acropolis (using vSphere HTML 5 Client)*.
2. Register the cluster that you want to use as the primary cluster with an SRM server (in the diagram, SRM B). See the Nutanix Storage Replication Adapter (SRA) for Site Recovery Manager Administration Guide.

After you configure the previously described replications, the protection domain is synchronised from Cluster A to clusters B, C, and D.

Snapshot Replication Behavior between the Clusters

The NearSync schedule creates 20-second RPO snapshots on Cluster A. These snapshots are not replicated to Cluster B. As a part of the roll-up schedule, the hourly and lower frequency snapshots are synchronized with Cluster B.

The 20-second RPO snapshots are replicated to Cluster C, and are rolled up into lower interval snapshots (for example, hourly snapshots, daily snapshots, and so on). The 20-second RPO ensures that initial replication from Cluster B to Cluster C is minimal if Cluster A becomes unavailable.

The asynchronous replication between clusters A and D creates hourly snapshots that are replicated to D on the basis of the configured schedule. These snapshots ensure that initial replication is minimal when it is time to promote the standby site to a production site that implements synchronous replication between Clusters C and D.

Promoting the Secondary Site To Production Status (Unplanned Failover If Only Cluster A Becomes Unavailable)

If Cluster A becomes unavailable, perform the following tasks to promote the standby site to production status, with synchronous replication between clusters C and D.

1. Promote the protection domain on Cluster B. See Performing an Unplanned Failover (Disaster) on page 113.

   Note: Only promote the protection domain. Do not start the VMs.

2. Sign in to the Controller VM in cluster C and do the following:

   - Determine whether any replication is in progress from Cluster A to Cluster C.

     ```
     nutanix@cvm$ ncli pd ls-repl-status name=pd-name
     ```

     Replace *pd-name* with the name of the protection domain.

   - If any replications are in progress, abort the replications.

     ```
     nutanix@cvm$ ncli pd abort-repl name=pd-name replication-ids=replication-id
     ```

     Replace *pd-name* with the name of the protection domain. Replace *replication-id* with the replication ID obtained from the **ncli pd ls-repl-status** command.

3. Log on to Cluster B, and then do the following:

   - Promote the protection domain on Cluster B, because B has the latest data to be replicated to Cluster C. See Failover Scenarios for Synchronous Replication in Hyper-V on page 143.

   - Create and replicate point-in-time backups (out-of-band snapshots) and replicate them to Cluster C. See Modifying Asynchronous or NearSync Replication Configurations on page 73.

     After this step, Cluster C has the latest data.

   - Deactivate the protection domain on Cluster B to prepare it to receive replication from Cluster C.

     ```
     nutanix@cvm$ ncli pd deactivate-and-destroy-vms name=pd-name
     ```

     Replace *pd-name* with the name of the protection domain.

4. Sign in to the SRM server and initiate unplanned failover from Cluster A to Cluster C. See the Nutanix Storage Replication Adapter (SRA) for Site Recovery Manager Administration Guide.
5. Configure synchronous replication from Cluster C to Cluster D.

   ```
   nutanix@cvm$ ncli pd metro-avail-enable name=pd-name ctr-name=container-name \
    remote-site=remote-site-name failure-handling=failure-handling
   ```

   This command converts an existing asynchronous protection domain into a protection domain for synchronous replication.

   Replace *pd-name* with the name of the protection domain, *container-name* with the name of the container, *remote-site-name* with the name of the remote site, and *failure-handling-mode* with one of **Automatic**, **Manual**, or **Witness**.

   At this stage, the former standby site is running as a fully operational production site. You can configure either NearSync or asynchronous disaster recovery between Cluster C and Cluster B.

6. After Cluster A becomes available, log on to Cluster A (it will be in a decoupled state) and disable metro availability and deactivate the protection domain.

   ```
   nutanix@cvm$ ncli pd deactivate-and-destroy-vms name=pd-name
   ```

   Replace *pd-name* with the name of the protection domain.

7. Configure a NearSync schedule on the protection domain from Cluster C to Cluster A. See Configuring Data Protection with NearSync Replication on page 64.

   At this point, the former production site is available as a standby site.

Promoting the Secondary Site To Production Status (Unplanned Failover If Both Clusters A and B Become Unavailable)

If both clusters A and B become unavailable, do the following:

- Sign in to the SRM server and perform unplanned failover on Cluster C. See the Nutanix Storage Replication Adapter (SRA) for Site Recovery Manager Administration Guide.
- Configure synchronous replication from Cluster C to Cluster D.

  ```
  nutanix@cvm$ ncli pd metro-avail-enable name=pd-name ctr-name=container-name \
   remote-site=remote-site-name failure-handling=failure-handling
  ```

  This command converts an existing asynchronous protection domain into a protection domain for synchronous replication.

  Replace *pd-name* with the name of the protection domain, *container-name* with the name of the container, *remote-site-name* with the name of the remote site, and *failure-handling-mode* with one of `Automatic`, `Manual`, or `Witness`.

Performing Planned Failover from Cluster A to Cluster C

1. Log on to Cluster A and power off the VMs in the storage container. See "Managing a VM (ESXi)" in VM Management in the Prism Web Console Guide.
2. Disable metro availability on the container. Enabling and Disabling Metro Availability on page 101.
3. Use SRM to perform planned failover from A to C. See the Nutanix Storage Replication Adapter (SRA) for Site Recovery Manager Administration Guide.

Performing Test Failover

Sign in to the SRM server managing Cluster A and perform a test failover to Cluster C. See the Nutanix Storage Replication Adapter (SRA) for Site Recovery Manager Administration Guide.

Restoring VMs from Snapshots

Do one of the following:

- If you restoring snapshots of hourly or lower frequency, see Restoring a VM from a Metro Availability Snapshot on page 102.
- If you are restoring LWS snapshots, first disable the metro availability protection domain and then perform the tasks described in Restoring a VM from a Metro Availability Snapshot on page 102.

7

REMOTE SITE CONFIGURATION

A remote site is the target location to store data replications for protection domains. All Nutanix DR solutions provide the means to use a remote site for disaster recovery. Thus, configuring and maintaining remote sites for disaster recovery are important activities in the disaster recovery operations.

The remote site can be one of the following:

- Another physical cluster (see Configuring a Remote Site (Physical Cluster) on page 152)
- A public cloud provider like AWS (see Configuring a Remote Site in AWS on page 201 in the Asynchronous Replication Using Cloud Connect (On-Prem to Cloud) on page 196 section.)

> Note: Ensure that, all the nodes in the Nutanix cluster must be part of the same ESXi host cluster and the network must be available on all the nodes in the Nutanix cluster at a given site (primary or remote).

Configuring a Remote Site (Physical Cluster)

Before you begin

- To use network mapping, create the network connections and VLANs on both source and destination cluster. For more information about configuring network connections, see *Network Configuration for VM Interfaces* in the *Prism Web Console Guide*.
- Set up the network mapping on both the source and the destination (remote) clusters. For more information on network mapping, see Network Mapping on page 162.

About this task

A remote site is the target location to store data replications for protected domains (see Configuring Data Protection with Asynchronous Replication on page 49). The remote site can be either another physical cluster or a cluster located in a public cloud. For information about a remote site on public cloud using Amazon Web Service (AWS), see Asynchronous Replication Using Cloud Connect (On-Prem to Cloud) on page 196

> Note: Do not create multiple remote sites pointing to the single destination cluster. An alert is generated if you create multiple remote sites.

To configure a remote physical cluster that can be used as a replication target, do the following:

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Remote Site** button and then select **Physical Cluster** from the pull-down list.
   The physical cluster selection refers to the physical cluster that you want to configure as **Remote Site**.
   The Remote Site dialog box appears.

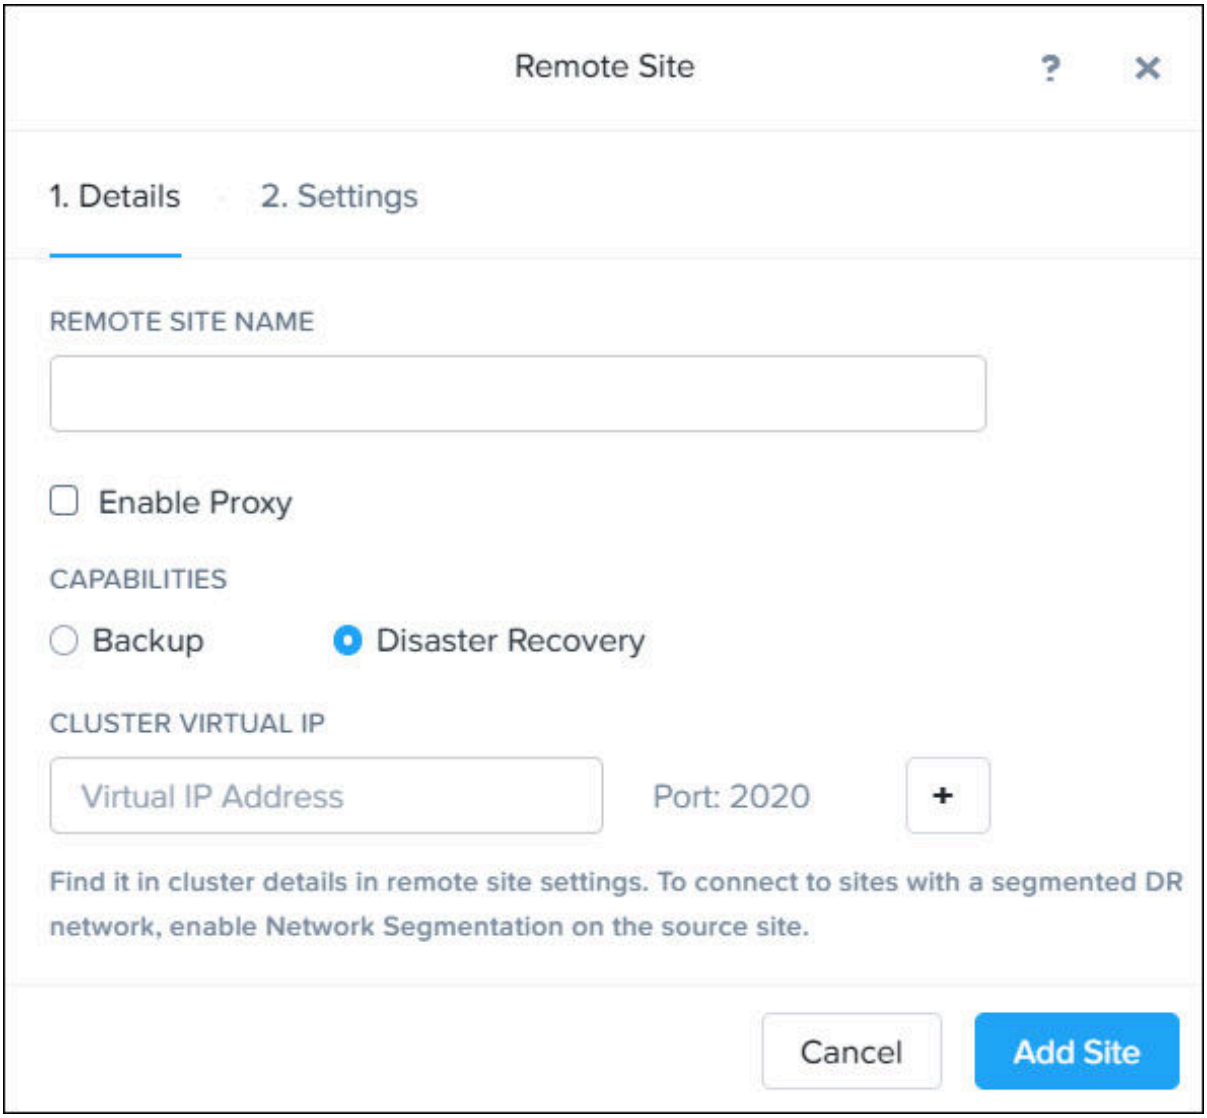

Figure 54: Remote Site Window

2. Do the following in the indicated fields for the remote site cluster:

   a. **Remote Site Name**: Enter the host name of the remote site cluster.

      Note: This entity has the following naming restrictions:

      - The maximum length is 75 characters.
      - Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

   b. **Enable Proxy**: Check the box to enable a proxy.

      No proxy is used by default.

      Enabling this field allows the specified IP addresses—remote CVM IP addresses or remote Virtual IP address—to be used as a proxy to communicate with a Nutanix cluster on the remote site. In this case, the source cluster communicates with only one of the remote proxy IP addresses (remote CVM IP addresses or remote Virtual IP address), which forwards the requests to the appropriate remote CVMs. The proxy setting on the remote site limits the replication traffic to the defined destination remote site IP address (many to one) rather than to each individual CVM IP address in the destination cluster (many to many). The many-to-one replication approach can be used to simplify firewall rules configuration.

      Note: It is recommended that you use a Virtual IP address for the remote site cluster proxy. Configuring a remote CVM for the proxy makes the CVM the single point of failure and break the entire communication if the CVM goes down.

      Network Address Translation (NAT) performed by any device between the two Nutanix clusters is not currently supported.

      CAUTION: Do not enable a proxy on remote site clusters that are used with a metro availability protection domain.

   c. **Backup**: Select this option to enable backup (only) to this remote site cluster.

      Backup allows the remote site cluster to be used as a backup (replication) target. This means data can be backed up to this remote site cluster and snapshots can be retrieved

from the site to restore locally, but failover protection (that is, running failover VMs directly from the remote site) is not enabled.

d. **Disaster Recovery**: Select this option to enable backup to and recovery from this site.

   Disaster recovery allows the remote site to be used both as a backup target and as a source for dynamic recovery. This means that failover VMs can be run directly from the remote site.

e. **Cluster Virtual IP**

   1. (Only on remote site clusters without network segmentation for disaster recovery) Do the following:

      In **Addresses** and **Port**, enter the virtual IP address of the remote site cluster and port number respectively, and then click the **Add** button.

      If you do not include a port number, a default port is used. The default port numbers are 2009 and 2020.

      > Note: Ensure that the virtual IP address and the nodes in the remote cluster are in the same subnet.

   2. (Only on remote site clusters with network segmentation for disaster recovery) Do the following:

      - In **Segmented Subnet (Gateway IP/Prefix Length)**, enter the address of the subnet in the remote site cluster that is configured for isolating disaster recovery traffic.
      - In **Port**, enter the port number of the remote site cluster.
      - In **Segmented Virtual IP**, enter the virtual IP address (IPv4 or IPv6) of the disaster recovery service configured in the remote site cluster.

      If you do not include a port number, a default port is used. The default port numbers are 2009 and 2020.

      > Note: Ensure that the virtual IP addresses and the nodes in the remote cluster are in the same subnet.

3. To configure more parameters for the remote site cluster, click the **Add Site** button and do the following in the indicated fields:

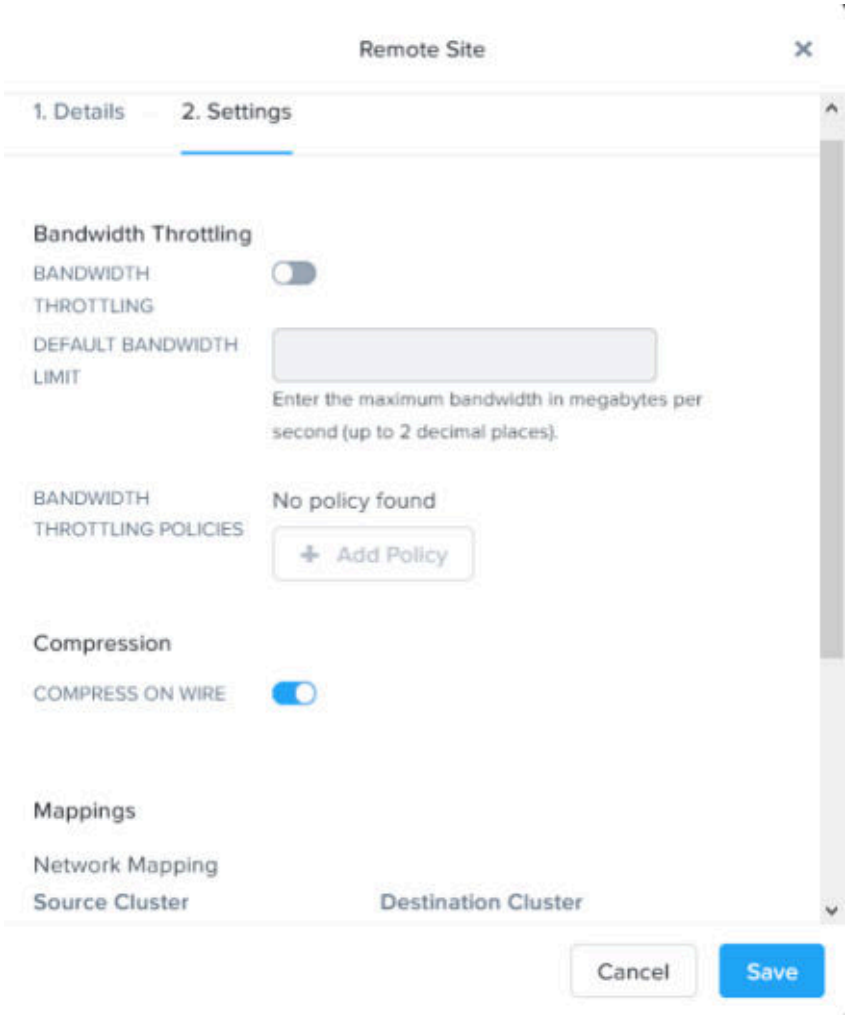

a. If you want to define a bandwidth policy for the network between the primary and remote site clusters, set the **Bandwidth Throttling** button to enabled.

   The bandwidth throttling policy provides you with an option to set the maximum limit of the network bandwidth. You can specify the policy depending on the usage of your network. For example, you can define a policy that a Nutanix cluster should replicate data

from site A to site B at less than 10 MBps between 9 a.m. to 5 p.m. on weekdays because there might be other critical traffic between the two sites then.

> Note: NearSync does not support bandwidth throttling.

b. **Default Bandwidth Limit**: Enter the maximum bandwidth allowed (in MBps) that should be used for replication purpose if no bandwidth policy is applied.

> Note: Maximum bandwidth of 2048 MBps is only supported for replication purpose.

Leaving the field blank means that there is no restriction on bandwidth use.

c. Click **Add Policy** to configure the bandwidth throttling policy.

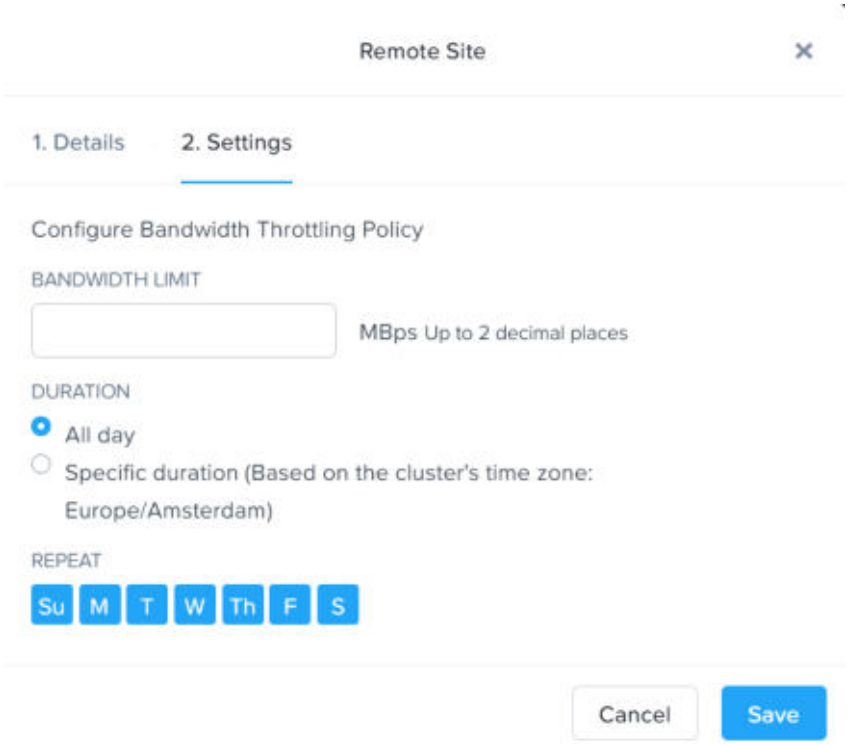

1. Enter the maximum bandwidth allowed (in MBps) for transmitting the replicated data to the remote site in the **Bandwidth Limit** text box.

> Note: Maximum bandwidth of 2048 MBps is only supported for replication purpose.

2. Select **All day** (default) to apply the bandwidth specified for the entire day and for the entire week. Otherwise, select the days that you want to apply the bandwidth

throttling. To specify the days that you want to apply bandwidth policy for the entire day, click and de-select the days that you do not want to apply the policy.

3. Select **Specific time** and specify the time and also the days that you want to use the bandwidth throttling. For example, you can specify the maximum bandwidth limit of 10 MBps between 9 a.m. to 5 p.m. from Monday to Friday.
4. Click **Save**.

d. **Compress on Wire**: Click the **Yes** button to enable compression on wire (also known as network compression).

   Data compression during transit (on wire) between primary and remote site clusters is not used by default. Enabling compression on wire consumes more CPU resources.

e. **Network Mapping**: Select the VLAN ID from the source and destination clusters drop-down menu.

   All the VLANs that are available on the source and destination clusters are mapped in the drop-down menu. For example, if VLAN01 is configured on the source cluster and VLAN02 is configured on the destination cluster. If you select VLAN01 on the source cluster and VLAN02 on the destination cluster, in failover scenario, a VM that is running in

VLAN01 on the source cluster is automatically brought up in VLAN02 on the destination cluster.

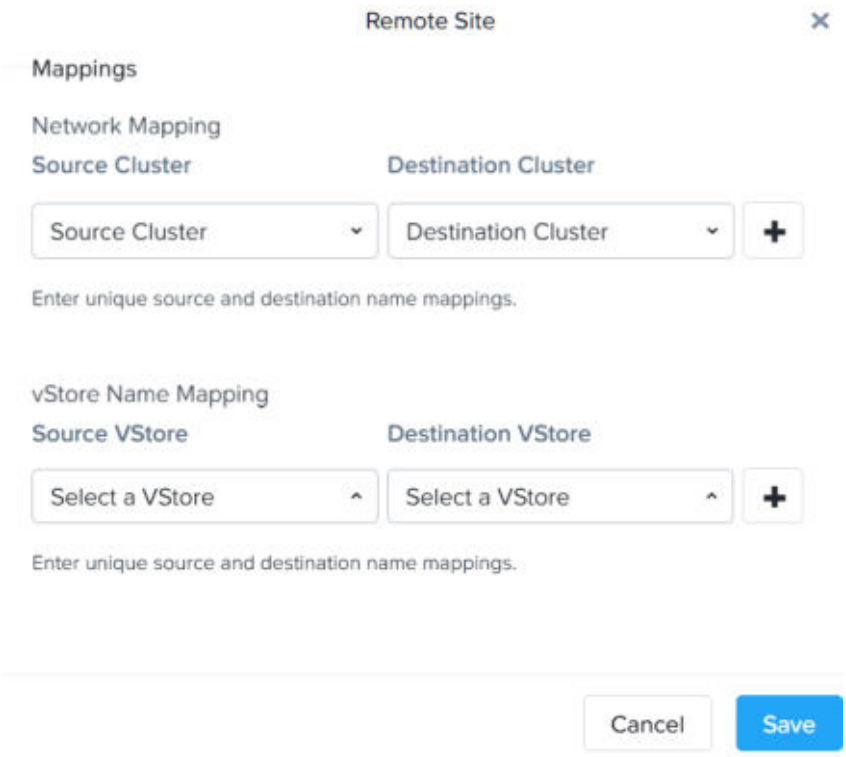

f. **vStores**: To specify a vStore, select the source storage container from the pull-down list that contains the VMs to protect, enter the destination storage container name on the remote site, and then click the **Add** button. Repeat for more vStore entries.

   Each **vStores** entry maps the association between a local source storage container and a remote target storage container at the remote site.

   Note:

   - If you perform restore of snapshots on vStore protection domains, you need to modify the VMDK paths in the VM files and point it to the restored path that is specified as part of the restore operation.
   - If a protection domain is configured with a NearSync schedule and is configured to be replicated to the remote site, you need to configure the corresponding vStore mappings.

4. When all the field entries are correct, click the **Save** button.

What to do next

Ensure that the local cluster is configured as a remote site on the remote site cluster.

Note: The local cluster must also be configured as a remote site on the remote cluster in order to use the remote site as a target location to store data replications. If this is not currently the case, log into the remote cluster and configure the local cluster as a remote site for that cluster.

Modifying a Remote Site (Physical Cluster or Cloud)

About this task

You can monitor or modify an existing remote site (see Configuring a Remote Site (Physical Cluster) on page 152) by testing the network connection, updating the remote site settings, or deleting the remote site configuration. To modify a remote site, do the following:

> Note: Changing a VStore mapping causes associated protection domains to initiate full replication of protected entities to the newly specified target container as if it were initial replication. A mapping change therefore results in the overconsumption of storage resources at the remote site. Contact Nutanix Support for help with cleaning up snapshots in the previously specified container.

> Note: Before you delete a remote cloud-based site, ensure you have saved any critical data and then delete the remote VM snapshots from the remote site. Otherwise, when you delete the remote site as described in this topic, the remote snapshots are automatically deleted.

Procedure

1. In the Data Protection dashboard, click the **Table** view.
2. Click the **Remote Site** tab and select the target remote site cluster in the table (top section of screen).

   The Summary line (middle of screen) displays the remote site cluster with a set of action links on the right. The actions are **Test Connection**, **Update**, and **Delete**. The following steps describe how to perform each action.

Figure 55: Remote Site Action Links

3. To verify a network connection, click the **Test Connection** action link.
   This action invokes an RPC call to the remote site, which checks whether the remote site is alive and accessible. A green check mark # appears if the connection test succeeds; a red # appears if the connection test fails. Check the alert log for potential problems if the test fails.
4. To update the remote site cluster settings (physical cluster), click the **Update** action link.
   Update the settings in the dialog box as desired and then click the **Save** button.

   The Remote Site window appears, which includes the same fields that appear when configuring a remote site cluster (see Configuring a Remote Site (Physical Cluster) on page 152.

5. To update the remote site settings (cloud), click the **Update** action link.

   The Remote Site window appears, which includes the same fields that appear when configuring a remote site cluster (see Implementation Considerations for Cloud Connect on page 196. Additionally, you can configure bandwidth throttling.

   > Note: You can only configure bandwidth throttling only while updating the remote site cluster. This option is not available during the configuration of the remote site cluster.

   a. If you want to define a bandwidth policy, set the **Bandwidth Throttling** button to enabled.

      The bandwidth throttling policy provides you with an option to set the maximum limit of the network bandwidth. You can specify the policy depending on the usage of your network. For example, you can define a policy that a Nutanix cluster replicate data from site A to site B at less than 10 MBps between 9 am to 5 am on weekdays because there might be other critical traffic between the two sites at that time.

   b. **Default Bandwidth Limit**: Enter the maximum bandwidth allowed (in MBps) that must be used for replication purpose if no bandwidth policy is applied.

      > Note: Maximum bandwidth of 2048 MBps is only supported for replication purpose.

      Leaving the field blank means that there is no restriction on bandwidth use.

   c. Click **Add Policy** to configure the bandwidth throttling policy.

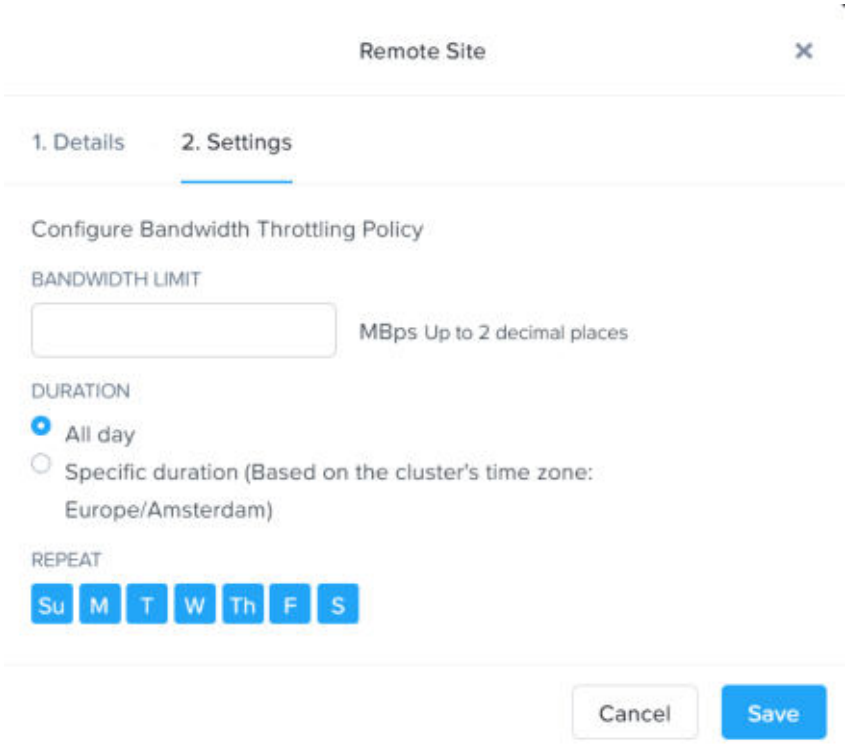

   1. Enter the maximum bandwidth allowed (in MBps) for transmitting the replicated data to the remote site in the **Bandwidth Limit** text box.

      > Note: Maximum bandwidth of 2048 MBps is only supported for replication purpose.

   2. Select **All day** (default) to apply the bandwidth specified for the entire day and for the entire week. Otherwise, select the days that you want to apply the bandwidth

throttling. To specify the days that you want to apply bandwidth policy for the entire day, click and de-select the days that you do not want to apply the policy.

3. Select **Specific time** and specify the time and also the days that you want to use the bandwidth throttling. For example, you can specify the maximum bandwidth limit of 10 MBps between 9 a.m. to 5 p.m. from Monday to Friday.
4. Click **Save**.

Network Mapping

Network mapping allows you to control network configuration for the VMs when they start on the remote site. By using the network mapping feature, you can specify network mapping between the source cluster and the destination cluster. The remote site wizard provides you with an option to create one or more network mappings and allows you to select source and destination network from the drop-down list. You can also modify or remove network mappings as part of modifying the remote sites.

> Note: Set up the network mapping on both the source and the destination clusters. If you do not set network mapping on both the clusters, the protection domain or the VM fails over without the VNIC. After the failover, you can see the guest OS but not the devices.

Periodic synchronization of network mapping between remote sites occurs and it ensures the sanity of the network mapping. Any modification on either site is reflected on another site.

> Note: Currently, only L2 network configuration mapping is supported.

> Note: Network mapping is supported for AHV and for cross-hypervisor VM mobility between ESXi and Acropolis. For more information about cross hypervisor VM mobility, see Nutanix Cross Hypervisor Disaster Recovery on page 189.

If you configure network mapping between source and remote sites, if there is failover the VMs register on the remote site with the network configuration that is specified in the network mapping. An example of one of the scenarios is as follows.

1. VMs are running on the source cluster on VLAN01.
2. On the remote site (destination cluster), VLAN02 network is present.
3. Network mapping is implemented between VLAN01 and VLAN02.
4. Snapshots are getting replicated from the source cluster to the destination cluster.
5. If there is failover of this protection domain to the remote site, all the VMs that were on VLAN01 on source cluster gets registered with VLAN02 on remote site automatically.

For more information about configuring networking mapping on remote sites, see Configuring a Remote Site (Physical Cluster) on page 152.

Network Mapping Behavior for AHV-Managed VMs After Failover/Disaster Recovery Migration

Sometimes, after failover or DR-related migration to a remote site, a VM is registered without any network configuration or virtual network card (VNIC). A VM can be registered without any network configuration if these conditions are met:

- You have specified a network mapping in the remote site.
- The VM is migrated to the remote site.
- The network where the migrated VM is mapped has changed or no longer exists.

AOS networking mapping helps ensure that migrated VMs are registered on the new network with an IP address in most cases.

Note: If a VM is configured with multiple VNIC and no mapping exists for one of the VNIC or AHV is unable to assign an IP to the VNIC on the destination site, the entire registration defaults to VM registered without a VNIC.

Table 12: VNIC and IP address Configuration Behaviour Scenarios

Source IP Address Management (IPAM)	Destination IPAM	Same Subnet?	Destination DHCP Pool?	VM IP Address Registration Result
No	No	Not applicable	Not applicable	VM NIC registered and original IP address configuration applied
No	Yes	Not applicable	Yes	VM registered with a randomly selected IP address from the DHCP pool
Yes	No	Not applicable	Not applicable	VM NIC registered and original IP address configuration applied
Yes	Yes	Yes	Not applicable	VM NIC registered and same IPAM DHCP address applied
Yes	Yes	No	No	VM registered without a VNIC
Yes	Yes	No	Yes	VM registered with a randomly selected IP address from the DHCP pool

Single-Node Replication Target Clusters

Single-node replication targets (SNRT) are single-node clusters that you can use as a recovery location. Small and medium businesses (SMBs) and remote office/branch office (ROBO) deployments can use a single cluster for the backup and recovery of infrastructure VMs such as the DHCP server, DNS server, SQL Server. If VMs on the primary site fail, you can recover them by using their replicated snapshots. See Single-Node Replication Target Requirements and Limitations on page 165 for requirements and limitations.

Specifications and Process

Platform

Nutanix provides different options of appliance that support the single-node replication Target feature. See Single-Node Replication Target Requirements and Limitations on page 165 for the options. For complete information about hardware configuration, see *Hardware Administration Guide*.

Read Operation

The read operation is same as that of a normal cluster (three nodes or more).

Write Operation

The write operation is performed to maintain disk-level fault tolerance. When you perform a write operation, the writes are always replicated to two different disks on the same node (in a redundancy factor 2 configuration). This process ensures that the cluster tolerates any single SSD or HDD failure.

> Note: Redundancy factor of 3 is not supported on the single-node replication target cluster.

Disk Failure

The single-node replication target cluster provides redundancy of data at the disk level. Hence, if a disk fails, you have the data available on another disk. Metadata is replicated across the two SSDs. In an event of an SSD failure, data and metadata is available on the second SSD.

Disk failure can be categorized into two categories:

- SSD failure: In an event of an SSD failure or when an SSD is marked for removal in a single-node replication target (SNRT) cluster with only two SSDs, the SNRT cluster becomes a read-only cluster and stops accepting any incoming replication traffic. An alert is displayed in both the source cluster and single-node replication target cluster. Also, the resiliency status of the cluster is changed to the critical state. The recovery process also slows down. Hence, Nutanix recommends that you 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).

  > Note: If the SNRT cluster has more than two SSDs, the failure of a single SSD does not cause the cluster to become a read-only cluster.

- 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 if there is HDD failure.

  > Note: For more information about recovering from disk failure see the boot/metadata drive documentation for *NX-1155-G5*.

Data Resiliency Status

The Data Resiliency Status window displays more detailed cluster resiliency status information. This window provides information about the number and type of failures the cluster can withstand safely. For single-node replication target clusters, the **Failures Tolerable** column indicates the number of simultaneous failures of that component type (disk and host instead of host and block) that is tolerated (0, 1, or 2). When no failure is tolerated, a message is displayed to highlight the limitation.

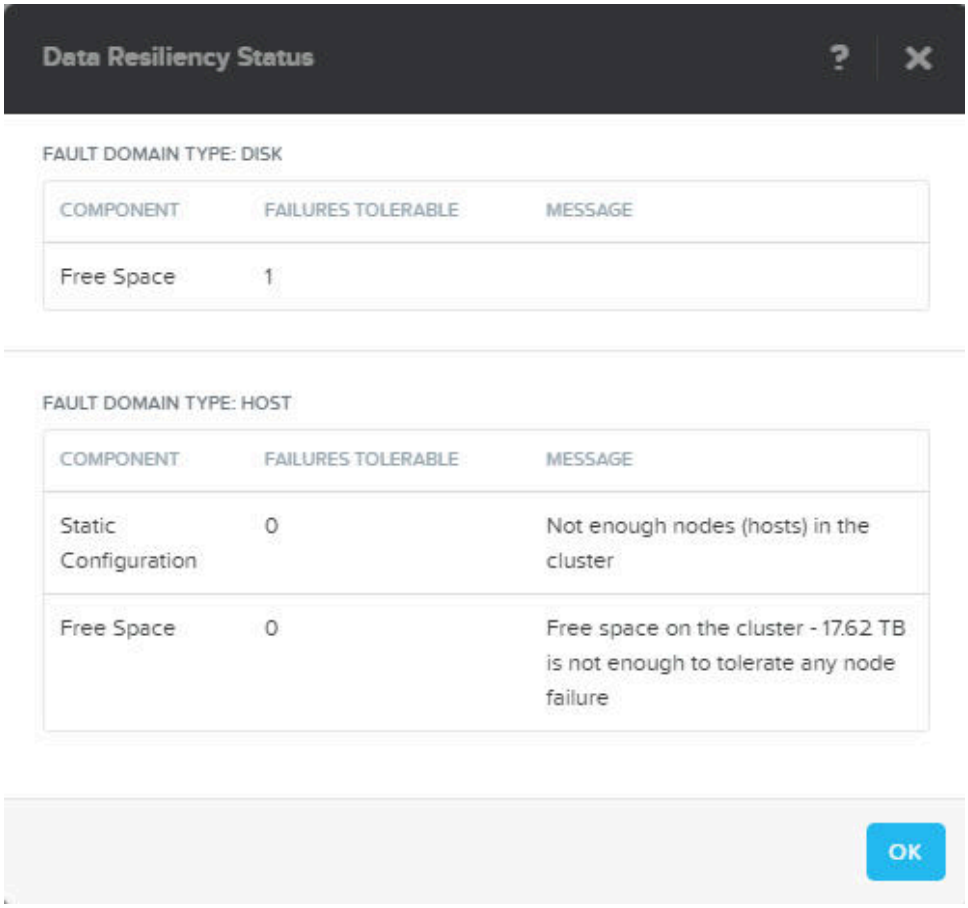

Figure 56: Data Resiliency Status

Single-Node Replication Target Requirements and Limitations

Following are requirements and limitations of the single-node replication target clusters.

Requirements and Limitations

It is recommended that you:

- Apply only 1:1 replication (not n:1 replication).
- Use any of the following platforms:

Table 13: Platform Support

Platform	Single-Node Replication Target (SNRT, 1-node)
NX-1175S-G8	Yes
NX-1175S-G7	Yes
NX-8155-G8	Yes

Platform	Single-Node Replication Target (SNRT, 1-node)
NX-8155-G7	Yes
NX-1175S-G6	Yes
NX-8155-G6	Yes
NX-5155-G6	Yes
NX-1175S-G5	No
NX-1155-G5	Yes

Ensure that:

- Each platform have two SSDs (400 GB each minimum), two HDDs, and enough memory to run VMs on a single node cluster.
- Nutanix recommends that you run only a maximum of five guest VMs on single node cluster like SNRT.

  Nutanix recommends configuring backup for all the five guest VMs running on a single-node cluster to protect the guest VMs in a node failure scenario.
- That IOPS on single node cluster does not to exceed a maximum of 1000 IOPS.
- There is only one VM or VG for each protection domain.
- You use only AHV and ESXi that are supported on the single-node replication target cluster. However, the primary cluster can run any hypervisor (ESXi, Hyper-V, or AHV).
- You do not enable de-duplication on storage containers having protected VMs since it lowers the replication speed.

  Nutanix does not recommend replication of protection domains on such containers to single-node replication target cluster.
- You configure only Nutanix native snapshots. Other types of snapshots are not supported.

Nutanix supports network segmentation on single node clusters for disaster recovery with a minimum AOS version. See the *What's New* section of the *AOS Release Notes*.

The following are not supported:

- Erasure coding (EC-X) on this cluster type.
- LCM on single-node clusters, 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).
- Nutanix Files on this cluster type.
- The firmware updates (BMC or BIOS) on the single-node replication target clusters.
- Metro availability on this cluster type.
- Expanding this cluster to add more nodes.
- Using the **Repair Host Boot Disk** from Prism web console to repair SATA DOM.

- Setting this cluster as a target for third party provides like Veeam or Commvault.

  Using this cluster as a block based storage target (using Volumes) for third party providers like Veeam or Commvault. However, you can use this cluster as an S3 target (using Objects) for such applications.

Conditions for Single-Node Replication Target Clusters

Some conditions for single-node replication target clusters are as follows:

- RPO for single-node replication target cluster is 6 hours with daily change rate of approximately around or less than 1% with a 5 TB of data set.
- You can configure schedules of seven daily, four weekly, and three monthly snapshots .
- RTO for single-node replication target cluster is 4–5 TB within 24 hours.

Licensing Requirements

Table 14: Licensing Requirements

Cluster	Description
Primary cluster	Primary cluster can be at the Starter license while replicating to the single-node replication target cluster. Note: Even if the primary cluster is replicating simultaneously to the single-node replication target cluster and a regular Nutanix cluster in the main datacenter, which is set up for either backup or disaster recovery purpose, you can still have Starter license on the primary cluster.
Single-node replication target cluster	Regular Nutanix licensing model is followed on the single-node replication target clusters. For example, Ultimate license will be required if SED Encryption is a requirement on the single-node replication target clusters.

Imaging Single-Node Replication Target Nodes

You can image single-node replication target nodes by using both Controller VM Foundation and Standalone Foundation.

> Note: Imaging by using Standalone Foundation is restricted to Nutanix sales engineers, support engineers, and partners. Contact Nutanix customer support or your partner for help with this procedure.

Controller VM Foundation

- With Controller VM foundation, you can map the primary cluster nodes (three nodes or more) and the single-node replication target cluster node (one node) during the foundation process itself.

  > Note: Ensure that primary and single-node replication target cluster is in the same broadcast domain or VLAN.

- You can also create a standalone single-node replication target cluster. If you decide not to map the primary cluster and replication target cluster, you need to add the replication target cluster as a remote site to the primary cluster. For information on the adding the single-node replication target cluster as a remote site, see Configuring a Remote Site (Physical Cluster) on page 152.

Standalone Foundation

With standalone foundation, you cannot map primary cluster and replication target cluster. You need to add the replication target cluster as a remote site to the primary cluster. For more information, see Configuring a Remote Site (Physical Cluster) on page 152.

> Note: See the *Field Installation Guide* for more information about imaging the nodes.

Setting Up Single-Node Replication Target Clusters

If you have not mapped the primary cluster and replication target clusters during the Foundation process, you can manually set this cluster as a remote site target to store data replications for the protected domains.

About this task

The process of setting up the remote site in single-node replication target cluster is same as the way to set up a remote site for the normal cluster except that you have to set up the remote site with backup only capability.

Procedure

Login to the Prism web console and follow the steps described in Configuring a Remote Site (Physical Cluster) on page 152.

Recovery Procedures by Using Single-Node Replication Target Clusters

For some reason, if the entities on the primary cluster fail, you can recover the entities by using the snapshots that are replicated to the single-node replication target cluster.

About this task

The process of restoring protected entities from the single-node replication target clusters is same as restoring entities from a normal cluster.

> Note: You can restore data stored on the single-node replication target cluster to the primary site.

Procedure

1. To recover the entities, log on to the Prism Web Console and follow the steps described in Restoration of Protected Entities on page 178.
2. To recover the cluster, follow the steps described in Recovering from the Remote Snapshots on a Backup or DR Site on page 182 for more information.

8

SELF-SERVICE RESTORE

The self-service restore (also known as file-level restore) feature allows virtual machine administrators to do a self-service recovery from the Nutanix data protection snapshots with minimal administrator intervention.

The Nutanix administrator must deploy NGT on the VM before enabling this feature. For more information about enabling and mounting NGT, see *Enabling and Mounting Nutanix Guest Tools* in the *Prism Web Console Guide*. After the feature is enabled and a disk is attached by the guest VM administrators by logging into the VM, the guest VM administrator can recover files within the guest operating system. If the guest VM administrator fails to detach the disk, it gets automatically detached from the VM after 24 hours.

Note:

- The Nutanix administrator can enable this feature for a VM through web interface or nCLI and in-guest actions can be performed by using NGT.
- Self-service supports hourly snapshots (recovery points), irrespective of whether the disaster recovery workflow is Async DR or NearSync.

  Async-DR workflows (one hour or greater) support the self-service restore feature with its hourly snapshots.

  For the NearSync workflow (upto 15 minutes), self-service restore uses the last hourly full snapshot taken. NearSync replication generates hourly recovery points to avoid too many delta recovery points.

Self-Service Enabled Disks Impact on Disaster Recovery

For disks with this feature enabled running on ESXi, the disaster recovery workflows back up the disks. However, the back-up has only the original disks of the VMs. If you replicate the VMs with self-service enabled disks, the following scenarios occur.

- If you attempt to attach a self-service disk for a VM that is part of a protection domain that the workflow replicated to a remote site running AOS version earlier than AOS 4.5, an error message is displayed during the disk attach operation.
- If a snapshot with the attached self-service disk is replicated to a remote site running AOS version earlier than AOS 4.5, an alert message is displayed during the replication process.

Requirements and Limitations of Self-Service Restore

Requirements specific for self-service restore are as follows.

General Requirements and Limitations of Self-Service Restore

To use self-service restore you need the following:

- This feature requires a Pro (or Ultimate) license.

- The Nutanix administrator must deploy NGT on the VM before enabling this feature. The Nutanix administrator can enable this feature for a VM through web interface or nCLI and in-guest actions can be performed by using NGT.

  For more information about enabling and mounting NGT, see *Enabling and Mounting Nutanix Guest Tools* in the *Prism Web Console Guide*.

- Set disk.enableUUID=true in the .vmx file for the VMs on ESXi.
- Configure a protection domain for the Guest VM to generate Nutanix snapshots configured by adding the VM to a protection domain.

  Self-service restore is not supported for the snapshots that you take from the **VM table** view.

- After the feature is enabled and a disk is attached by the guest VM administrators by logging into the VM, the guest VM administrator can recover files within the guest operating system. If the guest VM administrator fails to detach the disk, it gets automatically detached from the VM after 24 hours.

Self-service restore has the following limitations:

- Self-service supports hourly snapshots (recovery points), irrespective of whether the replication method used is Async DR or NearSync.

  Async DR replication schedules (snapshot frequency of one hour or greater) support the self-service restore feature with its hourly snapshots.

  For the NearSync replication schedules (snapshot frequency of up to 15 minutes), self-service restore uses the last hourly full snapshot taken. NearSync replication generates hourly recovery points to avoid too many delta recovery points.

- Self-service restore is not supported for vStore protection domains.
- Volume groups are not supported.
- Snapshots that are created in AOS 4.5 or later releases are only supported.
- IDE/SCSI and SATA disks are only supported. PCI and delta disks are not supported.

Requirements and Limitations for VMs Running Windows Operating System

> Note: For information about the operating systems that support the AOS version you have deployed, see the Compatibility Matrix.

- Ensure that sufficient logical drive letters are available to bring the disk online.
- File Systems. Dynamic disks comprising of NTFS on simple volumes, spanned volumes, striped volumes, mirrored volumes, and RAID-5 volumes are not supported.
- Only 64-bit operating system is supported.
- Following operating systems are supported.
  - Windows Server 2008 R2 or later versions
  - Windows 7 through Windows 10
- Disks created as Microsoft Storage Space devices by using Microsoft Windows Server 2016 or later are not supported.

Requirements and Limitations for VMs Running Linux Operating System

> Note: For information about the operating systems that support the AOS version you have deployed, see the Compatibility Matrix.

- File Systems. Only extended file system (ext2, ext3, and ext4) and XFS file systems are supported.
- Logical Volume Manager (LVM) disks for which the volume group corresponds to only a single physical disk are mounted.
  - Whenever the snapshot disk has an inconsistent file system (as indicated by the fsck check), disk is only attached and not mounted.
- Following Operating Systems are supported.
  - CentOS 6.5 through 6.9 and 7.0 through 7.3
  - Red Hat Enterprise Linux (RHEL) 6.5 through 6.9 and 7.0 through 7.3
  - Oracle Linux 6.5 and 7.0
  - SUSE Linux Enterprise Server (SLES) 11 SP1 through 11 SP4 and 12 SP1 through 12 SP3
  - Ubuntu 14.04 for both AHV and ESXi and 16.10 (AHV only)

Enabling Self-service Restore

Before you begin

Verify that NGT is installed and enabled on the VM. For more information, see *Enabling and Mounting Nutanix Guest Tools* in the *Prism Web Console Guide*.

Procedure

1. To enable self-service restore, click **Manage Guest Tools**.
2. To enable self-service restore feature, click **Self Service Restore (SSR)** check box.
   The Self-Service Restore feature is enabled on the VM. The guest VM administrator can restore the desired file or files from the VM.
3. Click **Submit**.

Self-Service Restore for Windows VMs

The Windows guest VM administrator can restore the desired file or files from the VM through the web interface or by using the ngtcli.

Restoring a File by Using Web Interface for Windows VMs

After you install the NGT software inside the Windows guest VM, the Windows guest VM administrator can restore the desired file or files from the VM through the web interface.

Procedure

1. Log into the guest VM by using administrator credentials.
2. Click the **Nutanix SSR** icon on the desktop.

3. Type the administrator credentials of the VM.

   - If you use NETBIOS domain name in username field (for example, domain\username), then you will be able to log on to SSR only if your account is explicitly added to Administrators group on the server. If username is added to any domain group, which is then added to Administrators group, then logon will fail. Also, you need to type NETBIOS domain name in capital letters (domain name has to be written in the same way as you see in the output of command net localgroup administrators).
   - If you use FQDN in username (for example domain.com\username), then you will only be able to logon if username user is a member of the domain admins group.
   - The snapshots that are taken for that day is displayed. You also have an option to select the snapshots for the week, month, and the year. In addition, you can also define a custom range of dates and select the snapshot.

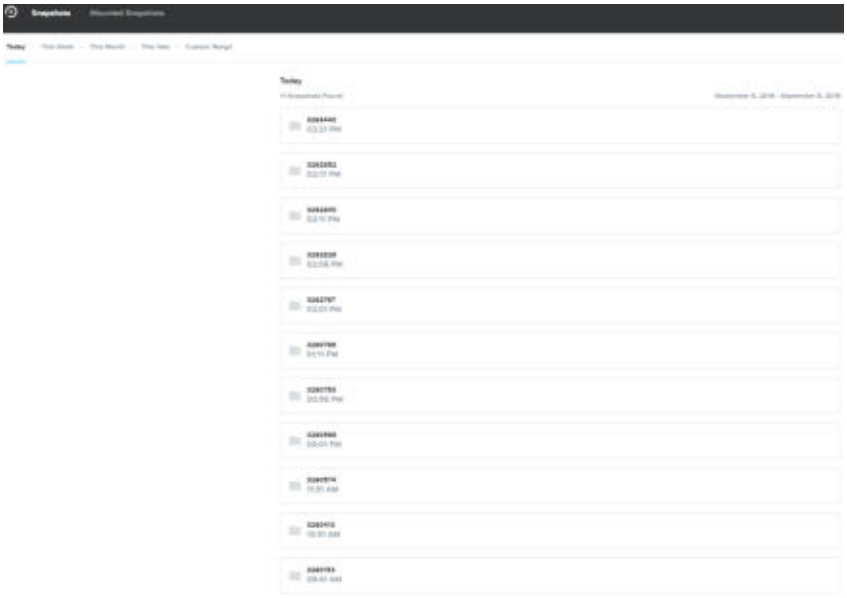

Figure 57: Snapshot Selection

4. Select the appropriate tab, **This Week**, **This Month**, **This Year**.
   You can also customize the selection by clicking **Custom Range** tab and selecting the date range in the **From** and **To** fields.
5. Select the check box of the disks that you want to attach from the snapshot.
6. From the **Disk Action** drop-down menu, select **Mount**.

Figure 58: Mounting of Disks

The selected disk or disks are mounted and the relevant disk label is displayed.

7. Go to the attached disk label drive in the VM and restore the desired files.
8. To view the list of all the mounted snapshots, select **Mounted Snapshots**.
   This page displays the original snapshot drive letters and its corresponding current drive letters. The original drive letters gets assigned to the disk at the time of the snapshot. Mounted drive letters are on which the snapshotted disk is mounted right now.

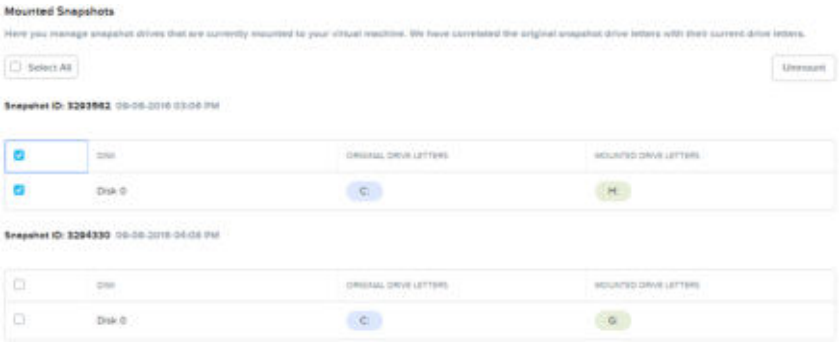

Figure 59: List of Mounted Snapshots

   a. To detach a disk, click the disk label and click **Unmount**.
      You can unmount all the disks at once by clicking **Select All** and then clicking **Unmount**.
9. To detach a disk, select the check box of the disk that you want to unmount and then from the **Disk Action** drop-down menu, select **Unmount**.

Restoring a File by using ngtcli for Windows VMs

After you install the NGT software inside the guest VM, the guest VM administrator can restore the desired file or files from the VM.

Before you begin

Verify that you have configured your Windows VM to use NGT. For more information, see Installing NGT on Windows Machines in the *Prism Web Console Guide*.

Procedure

1. Login to the guest VM.
2. Open the command prompt as an administrator.
3. Go to the ngtcli directory in **Program Files** > **Nutanix**.

```
cd c:\Program Files\Nutanix\ngtcli
```

4. Run the `ngtcli.cmd` command.

5. List the snapshots and virtual disks that are present for the guest VM by using the following command.

   ```
   ngtcli> ssr ls-snaps
   ```

   The snapshot ID, disk labels, logical drives, and create time of the snapshot is displayed. The guest VM administrator can use this information and take a decision to restore the files from the relevant snapshot that has the data.

   To list the snapshots with a specific number, use the following command.

   ```
   ngtcli> ssr ls-snaps snapshot-count=count_value
   ```

   Replace *count_value* with the number that you want to list.

6. Attach the disk from the snapshots.

   ```
   ngtcli> ssr attach-disk disk-label=disk_label snapshot-id=snap_id
   ```

   For example, to attach a disk with snapshot ID 16353 and disk label scsi0:1, type the command.

   ```
   ngtcli> ssr attach-disk snapshot-id=16353 disk-label=scsi0:1
   ```

   After successfully running the command, a new disk with label gets attached to the guest VM.

   If sufficient logical drive letters are not present, bringing disks online action fails. In this case, you should detach the current disk, create enough free slots by detaching other self-service disks and re-attach the disk again.

7. Go to the attached disk label drive and restore the desired files.

8. To detach a disk, use the following command.

   ```
   ngtcli> ssr detach-disk attached-disk-label=attached_disk_label
   ```

   For example, to remove the disk with disk label scsi0:3, type the command.

   ```
   ngtcli> ssr detach-disk attached-disk-label=scsi0:3
   ```

   If the disk is not removed by the guest VM administrator, the disk is automatically removed after 24 hours.

9. To view all the attached disks to the VM, use the following command.

   ```
   ngtcli> ssr list-attached-disks
   ```

Self-Service Restore for Linux VMs

The Linux guest VM user with sudo privileges can restore the desired file or files from the VM through the web interface or by using the ngtcli.

Restoring a File by Using Web Interface for Linux VMs

As Linux guest VM user, after you use the sudo privileges and install the NGT software inside the Linux guest VM, you can restore the desired file or files from the VM through the web interface.

Procedure

1. Log into the guest VM as a user with sudo privileges.

2. Click the **Nutanix SSR** icon on the desktop.

3. Type the root or sudo user credentials of the VM.

   The snapshots that are taken for that day is displayed. You also have an option to select the snapshots for the week, month, and the year. In addition, you can also define a custom range of dates and select the snapshot. For example, in the following figure snapshot taken on this month is displayed.

Figure 60: Snapshot Selection

4. Select the appropriate tab, **This Week**, **This Month**, **This Year**.

   You can also customize the selection by clicking **Custom Range** tab and selecting the date range in the **From** and **To** fields.

5. Select the check box of the disks that you want to attach from the snapshot.

6. From the **Disk Action** drop-down menu, select **Mount**.

   The selected disk or disks are mounted and the relevant disk label is displayed.

Figure 61: Mounting of Disks

7. Go to the attached disk label partitions in the VM and restore the desired files.

> Note: If the disk gets updated between the snapshots, the restore process may not work as expected. If this scenario occurs, you need to contact support to help with the restore process.

8. To view the list of all the mounted snapshots, select **Mounted Snapshots**.
   This page displays the original snapshot drive letters and its corresponding current drive letters. The original drive letters gets assigned to the disk at the time of the snapshot. Mounted drive letters are on which the snapshotted disk is mounted right now.

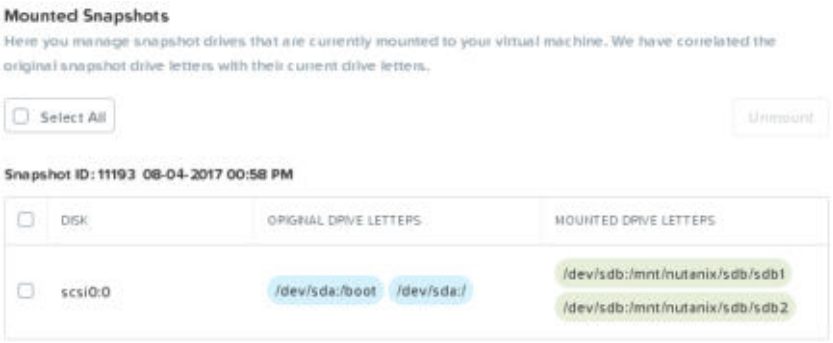

Figure 62: List of Mounted Snapshots

   a. To detach a disk, click the disk label and click **Unmount**.
      You can unmount all the disks at once by clicking **Select All** and then clicking **Unmount**.

9. To detach a disk, select the check box of the disk that you want to unmount and then from the **Disk Action** drop-down menu, select **Unmount**.

Restoring a File by using ngtcli for Linux VMs

As Linux guest VM user, after you use the sudo privileges and install the NGT software inside the guest VM, the user can restore the desired file or files from the VM.

Before you begin

- Mount NGT for a VM. For more information about enabling NGT, see Enabling and Mounting Nutanix Guest Tools in the *Prism Web Console Guide*.
- Verify that you have configured your Linux VM to use NGT. For more information, see Installing NGT on Windows Machines in the *Prism Web Console Guide*.

Procedure

1. Login to the guest VM with sudo or root user credentials.
2. Go to the ngtcli directory in `/usr/local/nutanix/ngt/ngtcli`.
3. Execute the `ngtcli.py` file as follows:
   » For NGT versions prior to 2.0: `/usr/local/nutanix/ngt/python27/python ngtcli.py`
   » For NGT 2.0 or later versions: `/usr/local/nutanix/ngt/python36/python3 ngtcli.py`

4. List the snapshots and virtual disks that are present for the guest VM by using the following command.

   ```
   ngtcli> ssr ls-snaps
   ```

   The snapshot ID, disk labels, logical drives, and create time of the snapshot is displayed. The guest VM administrator can use this information and take a decision to restore the files from the relevant snapshot that has the data.

   To list the snapshots with a specific number, use the following command.

   ```
   ngtcli> ssr ls-snaps snapshot-count=count_value
   ```

   Replace *count_value* with the number that you want to list.

5. Attach the disk from the snapshots.

   ```
   ngtcli> ssr attach-disk disk-label=disk_label snapshot-id=snap_id
   ```

   For example, to attach a disk with snapshot ID 1343 and disk label scsi0:2, type the command.

   ```
   ngtcli> ssr attach-disk snapshot-id=1343 disk-label=scsi0:2
   ```

   After successfully running the command, a new disk with new label is attached to the guest VM.

6. Go to the attached disk label partition and restore the desired files.

   > Note: If the disk gets updated between the snapshots, the restore process may not work as expected. If this scenario occurs, you need to contact support to help with the restore process.

7. To detach a disk, use the following command.

   ```
   ngtcli> ssr detach-disk attached-disk-label=attached_disk_label
   ```

   For example, to remove the disk with disk label scsi0:3, type the command.

   ```
   ngtcli> ssr detach-disk attached-disk-label=scsi0:3
   ```

   If the disk is not removed by the guest VM administrator, the disk is automatically removed after 24 hours.

8. To view all the attached disks to the VM, use the following command.

   ```
   ngtcli> ssr list-attached-disks
   ```

9

RESTORATION OF PROTECTED ENTITIES

In-Place Restore

In-place restoration creates entities—VMs and their attached VGs—on the recovery site, and overwrites the existing ones. The newly created entities replace the older entities and their attachments. For example, if you select the in-place restoration of a VM and its attached VG in a protection domain, a new VM attached with a VG is created on the recovery site. The attachment of the VM and the VG is retained on the recovery site only if you have installed NGT on the protected VM.

Out-of-Place Restore

Out-of-place restoration creates entities—VMs and their attached VGs—on the recovery site, but does not overwrite the existing ones. The newly created entities remain separate from the older entities and their attachments. For example, if you select the out-of-place restoration of a VM and its attached VG in a protection domain, a new VM attached with a VG is created on the recovery site. The attachment of the VM and the VG is retained on the recovery site only if you have installed NGT on the protected VM.

> Note: If you have not installed NGT or if you are using Hyper-V, the entities—VMs and their attached VGs—restore but are detached from each other. Manually reconfigure the attachments after the restoration is complete. For information about attaching volume groups to VMs, see *Volume Group Configuration* in the *Prism Web Console Guide*
>
> .

You can perform either an in-place restore or an out-of-place restore of a protected entity. The steps that you need to perform after restoration depend on which hypervisor is installed on the host. Consider that you rewrite an entity in a protection domain that contains VMs and volume groups, and the volume groups are attached to the VMs. The following table describes the possible scenarios, results, and steps to perform after restoration:

Table 15: Steps to Perform After Restoring Protected Entities

Scenario	Result	Follow-Up Steps
VMware ESXi and Microsoft Hyper-V		
You overwrite one or more of the VMs in the protection domain.	Volume groups are not detached from the restored VMs, but follow-up steps are required. Other VMs attached to the volume groups are not affected.	Sign in to the restored VMs and configure the in-guest iSCSI attachments.

Scenario	Result	Follow-Up Steps
You overwrite a volume group.	The volume group is detached from all VMs.	Log on to all the VMs to which the volume group was attached and configure in-guest iSCSI attachments.
AHV		
You overwrite one or more of the VMs in the protection domain.	Volume groups are detached from the restored VMs. Other VMs attached to the volume groups are not affected.	Sign in to the web console and reattach the volume groups to the VMs. Alternatively, log on to the VMs and configure in-guest iSCSI attachments.
You overwrite a volume group.	The volume group is detached from all VMs.	Sign in to the web console and reattach the volume groups to the VMs. Alternatively, log on to the VMs and configure in-guest iSCSI attachments.

Out-of-place restoration processes do not detach the existing entities, but the new entities are created without attachments, and any desired attachments must be manually reconfigured. For example, if a VM and a volume group in a protection domain are restored to a different location, the existing volume group remains attached to the VM, but the new VM and volume group are not attached to each other. Manually attach the volume group to the VM.

For information about attaching volume groups to VMs, see *Volume Group Configuration* in the *Prism Web Console Guide*.

> Note: The entities of AHV/ESXi/HyperV clusters restore to a different path on ESXi clusters if their files conflict with the existing files on the recovery cluster. For example, there is a file name conflict if a VM (VM1) migrates to a recovery site that already has a VM (VM1) in the same container.
>
> However, the entities restore at a different path with `VmRecoveredAtAlternatePath` alert if the following conditions are met.
>
> - Both the primary and the recovery clusters (Prism Element) are AOS 5.17 or later versions.
> - A path for the entity recovery is not defined while initiating the restore operation.
> - The entities do not have shared disk/s.
> - The entities are not protected in a vStore protection domain.
>
> If these conditions are not satisfied, the migrate/restore operation fails.

Restoring an Entity from a Protection Domain

Before you begin

If you are restoring only volume groups, log on to the attached VMs, stop the applications that are using the volume groups, and take the volume group disks offline.

About this task

In case of single or multiple VM failures occurring due to any adverse event, you can restore one or more VMs to an earlier version if those VMs are protected. You can also restore the protected volume groups.

To restore a protected entity, do the following:

Procedure

1. In the Data Protection dashboard (see Data Protection Dashboard on page 27), click the **Table** tab and then the **Async DR** tab.

   See Data Protection Table View on page 29 for screen shots and additional information about the following three steps.

2. Select the target protection domain in the table (upper screen).

3. Do one of the following.

   » If the desired snapshot is stored locally, click the **Local Snapshots** tab, select the target snapshot from the displayed list, and then click the **Restore** link.

   » If the desired snapshot is stored on a remote site, click the **Remote Snapshots** tab, select the target snapshot from the displayed list, and then click the **Retrieve** link (and verify the request in the dialog box) to copy the snapshot from the remote site to the local cluster.

   > Note:
   >
   > Check the **Include Expired** check box in the **Local Snapshots** or the **Remote Snapshots** tab to see the expired snapshots. Select the expired snapshot and click the **Restore** link.
   >
   > By design, at least one snapshot is retained as a reference snapshot for the next replication. To age out and delete an old snapshot, create a new replication to the same remote site or create a new snapshot.

   The Restore Snapshot window appears.

4. Do the following in the indicated fields:

   a. **What to Restore**: Check the boxes for the entities to restore.

      Click the check box next to **Entity Name** to select all the entities in the list. When you select a protected entity (VM or volume group) for restoration from a snapshot, the user interface selects all related entities so that they can be restored along with the selected entity. If you want their attachment configurations to be restored automatically after recovery, you must restore all related entities. If you choose to restore only a subset of the selected entities, only the selected entities are restored. Their attachment configurations are not restored, and you must manually reattach the restored entities after recovery.

      The following changes occur after related entities are restored.

      - The IQN of a VM changes if the snapshot includes in-guest iSCSI attachments. A string of the form `ntnx-timestamp`, where `timestamp` is the recovery time stamp, is appended to the previous IQN. For example, if the previous IQN was

`iqn.1991-04.com.an-authority:1234`, the IQN after recovery is `iqn.1991-04.com.an-authority:1234:ntnx-timestamp`.

- A VM connects to volume groups by using the data services IP address as the iSCSI target IP address. All previous iSCSI target IP addresses are cleared.
- The volume groups lose the configured static IP address on restore.
- Volume groups are attached to the new IQNs assigned to the VMs.
- After a VM starts, the Nutanix Guest Agent (NGA) on the guest VM reads updated iSCSI configuration from the Nutanix Guest Tools (NGT) CD, and it updates the configuration on the guest VM so that no manual configuration is required on the VM.

b. **How to Restore**: Do one of the following.

- Click the **Overwrite existing entities** option to overwrite the selected entities (in-place restore). Clicking this powers off each selected VM, deletes the entities from the inventory, copies the entity files from the chosen snapshot to the current location, and registers the entities. The VMs do not power on automatically; you must start them manually. You must also manually reattach any volume groups that were attached to the VMs (see Volume Group Configuration in the *Prism Web Console Guide*). This operation does not affect the UUIDs of restored entities.

  CAUTION: If you restore an entity on the source cluster (failback) using third-party backup solutions (for example Veeam), two copies of the same entity conflict in the same Prism Central. The backup solution creates an entity with the same UUID as the UUID of the actual entity that failed over to a different cluster registered with the same Prism Central. You must delete one of the two conflicting entities.

  Note:

  - If you have configured the NearSync schedule, you cannot use this option to restore the entities. If you want to restore an entity by using this method, delete the schedule and then perform the operation.
  - If a VM was part of a protection domain, but got subsequently removed from that protection domain will be added back to the protection domain if an in-place restore is performed on the snapshot that contains that VM.

- Click the **Create new entities** button (Out-of-place restore). Enter the following:

  - (ESXi only) Desired path in the **Path** field
  - Optional name prefix in the **VM Name Prefix** field.
  - Optional name prefix in the **Volume Group Name Prefix** field.

  Click **OK** to create a clone of each selected entity in the path location (path location is only for ESXi) with the same name and the specified prefix. If you do not specify a prefix, the prefix `Nutanix-Clone-` is used. For example, if the original VM name is `myvm` and a prefix is not specified, the cloned VM name is `Nutanix-Clone-myvm` (or `Nutanix-Clone-myvm-yymmdd-hhmmss` if `Nutanix-Clone-myvm` exists). The VM clone is not powered on automatically; you must start it manually. Manually attach volume groups to the new

VM (see Volume Group Configuration in the *Prism Web Console Guide*). This operation created entities with new UUIDs.

Note:

- Entered paths are relative to the current entity location. For example, to clone the VM in `NewDirectoryX`, provide the path as `//NewDirectoryX`. Entities cannot be cloned to a different datastore.
- (ESXi only) Before restoring an entity in ESXi, ensure that you have all the storage containers mounted in the ESXi hosts. This helps in registering the newly created entity.
- (ESXi only) Snapshot recovery fails on VMs that contain split VMDK files as these files are not supported by VMware.

c. When the fields are correct, click the **Restore** button.

What to do next

If you restored only volume groups, log on to the attached VMs, bring the volume group disks back online, and start the application.

Important:

- If a VM does not connect automatically to a volume group after a clone, migrate, or activate operation, reboot the VM, discover the newly restored iSCSI target, and log in manually.
- The VM fails to power on after the restore operation if the source VM is configured with a pass-through device.

Recovering from the Remote Snapshots on a Backup or DR Site

When you have an existing backup or DR setup in your environment and the primary cluster is unavailable, then you can recover the cluster from the snapshots that are residing on a remote backup or DR site.

About this task

- If you have configured a backup site, for example Cloud Connect or single-node replication Target, follow the procedure in this section.
- If you have configured a DR site, see Performing Failover on page 77 for recovery procedures.

Note: On AHV, if you are recovering from a snapshot that is retrieved from a cloud remote site, the VMs are recovered without any network attached.

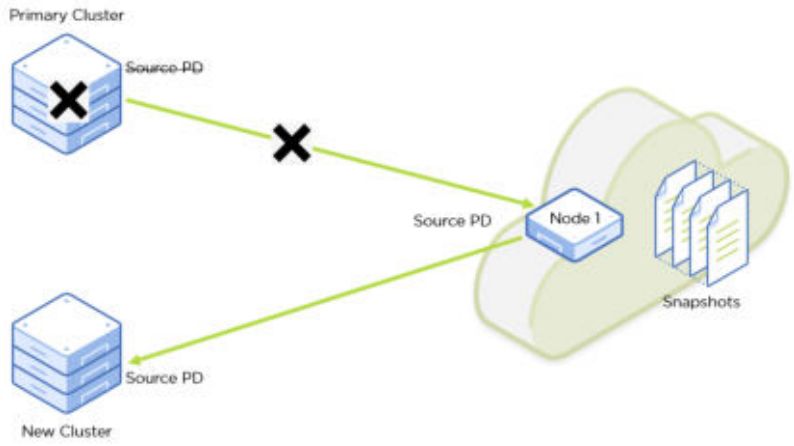

Figure 63: Cluster Recovery

Procedure

1. Create a new cluster.
   For more information about creating a new cluster, see the *Field Installation Guide*.

2. To configure the existing backup site (Cloud Connect or single-node replication Target) on the new cluster, go to **Data Protection** > **Remote Site** and select **Physical Cluster**. For more information about adding a remote site, see Remote Site Configuration on page 152.

   » For single-node replication target cluster as backup, you must perform the same procedure. Log on to single-node replication target clusters and configure the new cluster as a remote site for that cluster. For more information about adding a remote site, see Remote Site Configuration on page 152.

   » For Cloud Connect clusters (AWS) as backup, existing Controller VM located on the AWS cloud must be configured as remote site as follows.

   a. Sign in to the Controller VM of the new cluster and run the following command.

   ```
   nutanix@cvm$ ncli rs create name=remote_site_name_AWS \
   address-list=cvm_ip_AWS_instance capabilities=backup \
   vstore-map=vstore_mapping enable-compression=1 \
   cloud-type=AWS region=region_code_AWS
   ```

   For example, for AWS:

   ```
   nutanix@cvm$ ncli rs create name=amazon \
   address-list=cvm_ip (private IP) of AWS instance capabilities=backup \
   vstore-map=Container-1:backup_ctr enable-compression=1 cloud-type=AWS region=us-east-1
   ```

   b. Sign in to the Controller VM of the AWS instance and run the following command.

   ```
   nutanix@cvm$ ncli rs create name=new_cluster_name \
   address-list=virtual_ip_new_cluster capabilities=backup \
   vstore-map=vstore_mapping enable-compression=1
   ```

   For example,

   ```
   nutanix@cvm$ ncli rs create name=test address-list=virtual IP address of new cluster \
   capabilities=backup vstore-map=backup_ctr:Container-1 enable-compression=1
   ```

3. Create a protection domain on the new cluster with the same name as in the primary cluster.
   For more information about creating a protection domain, see Configuring Data Protection with Asynchronous Replication on page 49.
4. Select the protection domain, and click the **Remote Snapshots** tab.
   The snapshots that you have replicated to the backup site are visible. You can then retrieve and restore from the snapshots on the new cluster.

   Note: You can also create a protection domain on the backup single-node replication target cluster and then send snapshots to the new cluster that you created in step 1.

   - 1. Configure the new cluster as the remote site for that single-node replication target cluster. For more information about adding a remote site, see Remote Site Configuration on page 152.
     2. Create a protection domain on the single-node replication target cluster with the same name as in the primary cluster. For more information about creating a protection domain, see Configuring Data Protection with Asynchronous Replication on page 49.
     3. Select the protection domain, and click the **Remote Snapshots** tab. All the replicated snapshots are listed and you can then retrieve and restore from the snapshots on the new cluster.

Recovery of vGPU-enabled VMs

When you perform recovery of vGPU console-enabled guest VMs, the VMs recovers with the vGPU console.

With AHV, you can recover vGPU console-enabled guest VMs efficiently. When you perform DR of vGPU console-enabled guest VMs, the VMs recovers with the vGPU console. The guest VMs fail to recover when you perform cross-hypervisor disaster recovery (CHDR).

For AHV with minimum AOS versions 6.1, 6.0.2.4 and 5.20:

- vGPU-enabled VMs can be recovered when protected by protection domains in PD-based DR or protection policies in Leap based solutions using asynchronous, NearSync, or Synchronous (Leap only) replications.

  Note: GPU Passthrough is not supported.

- If both site A and site B have the same GPU boards (and the same assignable vGPU profiles), failovers work seamlessly. However, with protection domains, no additional steps are required. GPU profiles are restored correctly and vGPU console settings persist after recovery. With Leap DR, vGPU console settings do not persist after recovery.
- If site A and site B have different GPU boards and vGPU profiles, you must manually remove the vGPU profile before you power on the VM in site B.

For ESXi with minimum AOS versions 6.1, 5.20.4 and 5.15.5 and minimum SRA version 3.0:

- Nutanix DR solutions with protection domains and Leap are not supported.
- You need licenses for VMware SRM 8.3 or later versions.
- You can use vSphere based replication to recover vGPU enabled ESXi VMs.

The vGPU console settings are persistent after recovery and failover is supported for the following:

Table 16: Persistent vGPU Console Settings with Failover Support

Recovery using	For vGPU enabled AHV VMs	For vGPU enabled ESXi VMs
Protection domain based DR	Yes	No
VMware SRM with Nutanix SRA	Not applicable	Yes

For SRA and SRM support, see the Nutanix SRA documentation.

10

SHUTTING DOWN AND STARTING CLUSTERS IN BOTH PRIMARY AND REMOTE SITES

In a metro availability or any other replication configuration for disaster recovery with a cluster each at the primary site and the remote site, there can be circumstances where both sites need to be stopped for simultaneous maintenance.

> Note:
>
> Data replication for metro availability is not supported for AHV clusters. Therefore, Ignore any steps related to AHV clusters and hosts if you are running shutdown and restart operations on clusters with metro availability configurations.

Shutting Down the Clusters

About this task

Data replication for metro availability is not supported for AHV clusters. Therefore, ignore any steps related to AHV clusters and hosts if you are running shutdown and restart operations on clusters with metro availability configurations.

Before you begin

Before you start the shut down, sign in to the Prism Web Console for the clusters at both sites and ensure the following.

- Check the Data Resiliency status. If it is `OK`, proceed to the shut down procedure. If it is not `OK`, wait until it changes to `OK`.

  If it does not change to `OK`, contact Nutanix Support.

- Check Acropolis cluster health by running the NCC checks.

  Select **Health** > **Actions** > **Run NCC Checks** > **All Checks** and click **Run**.

  If there is a critical alarm or alert, fix it or contact Nutanix Support.

- Check and note the IP address of the IPMI console of CVM.

  SSH to the CVM console and use the `ncc hardware_info show_hardware_info` command to display the hardware information including the IPMI IP address.

Procedure

1. Primary site: Shutdown all guest VMs at the primary site.

2. Primary site: Do the following.
   a. (For AHV only) SSH into one CVM at the primary site.
      - Run the `cluster stop` command to shutdown the cluster. Type **Yes** to confirm the shutdown.
      - Run the `sudo cvm_shutdown -P now` command to shutdown the CVM.

        Note:

        Data replication for metro availability is not supported for AHV clusters. Ignore this step if you are running shutdown and restart operations on clusters with metro availability configurations.

   b. (For AHV only) SSH to each AHV host one by one of Nutanix AHV cluster and run the `shutdown -P now` command.
   c. (For ESXi only) Do the following.
      - Use the vSphere client and login to each ESXi host and shutdown the CVM.
      - Use the vSphere client to shutdown the ESXi hosts.
3. Remote site: Shutdown all guest VMs at the remote site.
4. Remote site: Do the following.
   a. (For AHV only) SSH into one CVM at the remote site.
      - Run the `cluster stop` command to shutdown the cluster. Type **Yes** to confirm the shutdown.
      - Run the `sudo cvm_shutdown -P now` command to shutdown the CVM.

        Note:

        Data replication for metro availability is not supported for AHV clusters. Ignore this step if you are running shutdown and restart operations on clusters with metro availability configurations.

   b. (For AHV only) SSH to each AHV host one by one of Nutanix AHV cluster and run the `shutdown -P now` command.
   c. (For ESXi only) Do the following.
      - Use the vSphere client and login to each ESXi host and shutdown the CVM.
      - Use the vSphere client to shutdown the ESXi hosts.

Starting the Clusters

About this task

This section provides the procedure to start the clusters at both sites - primary and remote - after they were shut down.

Before you begin

Ensure that all the hosts are powered off at both the sites.

Procedure

1. Remote site: Do the following.
   a. Manually power on Nutanix hosts at the remote site.

      You can power on the Nutanix AHV nodes using the IPMI Console by opening the IPMI IP address on your browser.

   b. Wait for 10-15 minutes for the nodes (including host and CVM) to boot up then verify that you can ping all management and CVM IPs.

   c. SSH into one CVM at the remote site and start the CVM.

      Check the status of the services. Wait for all services to show as UP.

2. Primary site: Do the following.
   a. Manually power on Nutanix hosts at the primary site.

   b. Wait for 10-15 minutes for the nodes (including host and CVM) to boot up then verify that you can ping all management and CVM IP addresses.

   c. SSH into one CVM at primary site and verify you can ping all secondary site CVM IP addresses.

   d. Start the CVM.

      Check the status of the services. Wait for all services to show as UP.

   e. Check Acropolis cluster health on Prism Web Console by running the NCC checks after Prism initializes.

      Select **Health** > **Actions** > **Run NCC Checks** > **All Checks** and click **Run**.

   f. Data replication with Metro Availability: Change the Distributed Resource Scheduler (DRS) rules to assign the VM affinities to the hosts of the primary site before powering on the VMs.

      Open the Data Protection view and check the **Metro Availability** configuration. Shutting down the remote site triggers the default automatic replication break timeout of 10 seconds for the protection domains for which the failure handling is set to Automatic Resume or Witness. Re-enable **Metro Availability** for the protection domains that are in disabled state.

   g. Power on the other guest VMs as required.

11

NUTANIX CROSS HYPERVISOR DISASTER RECOVERY

Cross hypervisor disaster recovery (CHDR) provides an ability to migrate the VMs from one hypervisor to another (ESXi to AHV or AHV to ESXi) by using the protection domain semantics of protecting VMs, generating snapshots, replicating the snapshots, and then recovering the VMs from the snapshots.

To perform these operations, you must install and configure NGT on all the VMs. For more information, see *Enabling and Mounting Nutanix Guest Tools* in the *Prism Web Console Guide*.

NGT configures the VM with all the required drivers for VM portability. After you prepare the source VMs, You can recover them in a different hypervisor type.

Note:

- When you recover VMs from AHV hypervisor on ESXi hypervisor, the VMs come up with E1000 NIC attached by default. However, if you want to bring the VMs up with VMXNET NIC attached, you must install VMware Tools.
- Static IP address preservation is not supported for CHDR workflows. If a VM has a static IP address configured, you must manually reconfigure the VM IP address after its recovery on the target site.
- Hypervisor-specific properties like multi-writer flags, independent persistent and independent non-persistent disks, or changed block tracking (CBT) are not preserved in the CHDR operation.

For more information about general NGT requirements, see *Nutanix Guest Tools Requirements and Limitations* in the *Prism Web Console Guide*.

Requirements of CHDR

- Nutanix supports only VMs with flat files. Nutanix does not support vSphere snapshots or delta disk files. If you have delta disks attached to the VM and you proceed with the migration, the VMs are lost.
- Nutanix does not support VMs with attached volume groups or shared virtual disks.
- Nutanix supports IDE/SCSI and SATA disks only. PCI/NVMe disks are not supported.
- Nutanix supports CHDR on both the sites with all the AOS versions that are not EOL.
- Set the SAN policy to **OnlineAll** for all the Windows VMs for all the non-boot SCSI disks so that they can automatically be brought online. For more information about setting SAN policy, see Bringing Multiple SCSI Disks Online on page 193.
- (For migration from AHV to ESXi) Automatic reattachment of iSCSI-based volume groups (VGs) fails after VMs are migrated to ESXi because vCenter and, in turn, the Nutanix cluster

do not obtain the IP addresses of the VMs. Therefore, after migration, manually reattach the VGs to the VMs.

- Do not enable VMware vSphere Fault Tolerance (FT) on VMs that are protected with CHDR. If already enabled, disable VMware FT. VMware FT does not allow registration of VMs enabled with FT on any ESXi node after migration. This also results in a failure or crash of Uhura service.
- See UEFI and Secure Boot Support for CHDR on page 190 for OS supporting UEFI and Secure Boot.

Operating Systems Supported

> Note: For information about the operating systems that support the AOS version you have deployed, see the Compatibility Matrix.

Table 17: Operating System Requirements for Legacy BIOS VM CHDR

Operating System	Version	Requirements and Limitations
Windows	• Windows 2008 R2 or later versions • Windows 7 or later versions	• Only 64-bit operating system is supported.
Linux	• CentOS 6.5 and 7.0 • Red Hat Enterprise Linux (RHEL) 6.5 or later and RHEL 7.0 or later. • Oracle Linux 6.5 and 7.0 • Ubuntu 14.04	• SLES operating system is not supported.

UEFI and Secure Boot Support for CHDR

Nutanix supports CHDR migrations of guest VMs having UEFI and Secure Boot.

Table 18: Nutanix Software - Minimum Requirements

Nutanix Software	Minimum Supported Version
Minimum AOS	5.19.1
Minimum PC	pc.2021.1
Minimum NGT	2.1.1

Table 19: Applications and Operating Systems Requirements - UEFI

Operating Systems	Versions
Microsoft Windows	• Microsoft Windows 10 • Microsoft Windows Server 2016 • Microsoft Windows Server 2019
Linux	• CentOS Linux 7.3 • Ubuntu 18.04 • Red Hat Enterprise Linux Server versions 7.1 and 7.7

Table 20: Applications and Operating Systems Requirements - Secure Boot

Operating Systems	Versions
Microsoft Windows	• Microsoft Windows Server 2016 • Microsoft Windows Server 2019
Linux	• CentOS Linux 7.3 • Red Hat Enterprise Linux Server versions 7.7

Table 21: Recovery Limitations

System Configuration	Limitation
Microsoft Windows Defender Credential Guard	VMs which have Credential Guard enabled cannot be recovered with CHDR recovery solution.
IDE + Secure Boot	VMs on ESXi which have IDE Disks or CD-ROM and Secure Boot enabled cannot be recovered on AHV.
UEFI VMs on CentOS 7.4, CentOS 7.5 and Ubuntu 16.04 may fail to boot after CHDR migration.	CentOS 7.4, CentOS 7.5 and Ubuntu 16.04 UEFI VMs do not boot after cross-hypervisor disaster recovery migrations. See KB-10633 for more information about this limitation. Contact Nutanix Support for assistance with this limitation.

System Configuration	Limitation
UEFI VM may fail to boot after failback.	When a UEFI VM is booted on AHV for the first time, UEFI firmware settings of the VM are initialized. The next step is to perform a guest reboot or guest shutdown to fully flush the settings into persistent storage in the NVRAM. If this UEFI VM is failed over to an ESXi host without performing the guest reboot/ shutdown, the UEFI settings of the VM remain partial. Although the VM boots on ESXi, it fails to boot on AHV when a failback is performed. See KB-10631 for more information about this limitation. Contact Nutanix Support for assistance with this limitation.

Recovering Virtual Machines in the Remote Cluster

You can recover the VMs by performing the following procedure.

Procedure

1. Install and configure NGT on the VMs. For installing and configuring NGT and its requirements and limitations, see Nutanix Guest Tools in the Prism Web Console Guide.

   .

   - (Optional) Run `ncli ngt list` to verify communication is active or not.
   - If your VM has SCSI disks, set the SAN policy to *online all*. Otherwise, only the boot-SCSI disk will be automatically brought online. If the policy is not configured properly and if the VM has an IDE boot disk, all SCSI disks will be offline if you try to power on the VM. To bring multiple disks online after recovery, see Bringing Multiple SCSI Disks Online on page 193.
   - Snapshots taken after NGT is installed are only supported for recovery operations. You can validate all the snapshots for VM recovery that are replicated to the remote site. Check the **Remote Snapshots** tab on the local cluster and the **Local Snapshots** tab on the remote cluster.
   - If the SCSI devices are enabled for multipathing and fstab contains the device name `/dev/sd*`, then while migrating a VM from ESXi to AHV, VM fails to boot and a message similar to the following `Device /dev/sd* in use e2fsck can't continue` is displayed. See Blacklisting Devices for Multipathing on page 194 for more information.

2. Create a protection domain and protect the VMs that you want to recover. For more information about creating a protection domain, see Configuring Data Protection with Asynchronous Replication on page 49.

3. Configure a remote site. During the addition of a remote site, you need to add the L2 network mapping. For more information about configuring a remote site and adding network mapping on a remote site, see Configuring a Remote Site (Physical Cluster) on page 152.

   Note: If you do not define any network mapping then the VM gets registered without any network configuration.

4. Migrate or take the snapshot of the protection domain and replicate it to the remote site. For more information about migrating, taking and replicating a snapshot, see Modifying Asynchronous or NearSync Replication Configurations on page 73.

5. Recover the VMs on the remote site by using one of the following methods.

   » By using *Planned Failover* procedure (protection domain migration operation) as described in Performing Failover on page 77.

   Appropriate warnings related to the VMs that cannot be recovered on the remote site are displayed when you try to migrate the protection domain. Click the **Learn more** link for additional information.

   » By using *Unplanned Failover* procedure (protection domain activate operation) as described in Performing Failover on page 77.

   > Note: Before activating the protection domain, verify which VMs are recoverable for the latest snapshot in the **Local Snapshots** tab.

   The **Recovery Details** link of the **Local Snapshots** tab provides information whether the VM is recoverable or not on the local cluster. The **Recovery Details** link of the **Remote Snapshots** tab provides information whether the snapshot that is present on the remote cluster is recoverable on the local cluster or not.

   » By restoring an individual VM from a snapshot that you have replicated. For more information, see Restoring an Entity from a Protection Domain on page 179.

   The **Recovery Details** tab provides the information whether the VM is recoverable or not.

   After performing CHDR, take a snapshot of the protection domain on the target site before you perform in-place restore operation. Otherwise, the in-place restore fails.

   > Note:
   >
   > - If you start the Windows 7 or Windows 2008 R2 VMs after migrating the VMs from ESXi to AHV, a prompt that asks for restart is displayed. It recommended to restart the VMs to install all the required drivers.
   > - Do not uninstall Nutanix VM mobility drivers on the VMs post migration as the VMs become unusable after uninstalling mobility drivers.
   > - (For migration from AHV to ESXi) If you migrate VMs to which (volume groups) VGs are attached, you must perform the following tasks after the migration is complete:
   >
   > 1. Install the following tools on each VM to enable the Nutanix cluster to learn the IP addresses obtained by the VMs:
   >
   >   1. VMware Tools.
   >   2. Nutanix Guest Tools (NGT). If NGT is already installed, reinstall NGT.
   >
   > 2. Manually reattach the VGs to the VMs.

   On successfully recovering the VMs on the remote site, they are ready for use.

   In case of a failure, see the **Alerts** tab for more information regarding the cause of the failure.

Bringing Multiple SCSI Disks Online

This tasks describes how to bring multiple SCSI disks online after recovery.

About this task

In Windows, the SAN policy determines whether a newly discovered disk is brought online or remains offline, and whether it becomes read/write, or remains read-only. If the SAN policy is not configured properly, only the boot disks are brought online.

When the disks are offline, the partitions or volumes will not be available and drive letters are not assigned. Hence you have to make the disks online.

When you recover a VM with multiple SCSI disks, the non-boot disks are not attached by Windows, but are visible. You need to bring these disks online after recovery.

Procedure

Perform one of the following steps to bring the disks online.

- Use the Windows Disk Management tool to bring a disk online.
  a. Go to **Start** > **Control Panel** > **System and Security** > **Administrative Tools** > **Computer Management** > **Storage** > **Disk Management**.
  b. Right-click the offline disk and select **Online**.
- Use the `diskpart` command-line utility to bring a disk online.
  a. Open the Windows command prompt and run `diskpart.exe` command.
  b. List the disks to confirm their status by running the following command.

```
DISKPART> list disk
```

  Output might look similar to the following.

```
Disk ###  Status      Size     Free     Dyn  Gpt
          --------  ----------  -------  -------  ---  ---
          Disk 0    Online        34 GB    12 MB
          Disk 1    Offline     1024 MB      0 B
```

  c. Bring the disk online by entering the following commands.

```
DISKPART> select disk disk number
DISKPART> ATTRIBUTES DISK CLEAR READONLY
DISKPART> Online
```

  Output might look similar to the following.

```
DiskPart successfully onlined the selected disk
```

- Use PowerShell to bring a disk online and make the read-only disk writable.

```
> Get-Disk | Where-Object IsOffline –Eq $True | Set-Disk –IsOffline $False
> Get-Disk | Where-Object IsOffline -Eq $True |Set-Disk -IsReadOnly $False
```

Blacklisting Devices for Multipathing

If the SCSI devices are enabled for multipathing and fstab contains the device name `/dev/sd*`, then while migrating a VM from ESXi to AHV, VM fails to boot and a message similar to the following `Device /dev/sd* in use e2fsck can't continue` is displayed.

About this task

To start the VMs, the SCSI devices with name in `fstab` must be blacklisted from multipathing. You can accomplish this by going into the maintenance mode on the console and performing the following procedure.

Procedure

1. Mount root file system as a root (rw).

   ```
   $ mount / -o rw,remount
   ```

2. Blacklist the devices that are causing the issues by adding the following lines to `multipath.conf` file.

   ```
   blacklist{
   devnode "^sd[a-z]"
   }
   ```

3. Restart the VMs.

12

ASYNCHRONOUS REPLICATION USING CLOUD CONNECT (ON-PREM TO CLOUD)

You can configure AWS as a remote site for backup and restore operations.

The Cloud Connect feature helps you back up and restore copies of virtual machines and files to and from an on-premise cluster and a Nutanix Controller VM located on an Amazon Web Service (AWS) cloud. The Nutanix Controller VM is created on an AWS cloud in a geographical region of your choice. It is a single-node cluster with a 30 terabyte (TB) disk attached to the node, with a usable disk capacity of 20 TB.

Amazon customers are charged only for capacity that is used (not charged for the full capacity). Once configured through the web console, the remote site cluster is managed and monitored through the Data Protection dashboard like any other remote site you have created and configured. To log on to the Nutanix Controller VM on the cloud, get the IP address from the Data Protection dashboard of the remote site cluster and SSH to the IP address from the source cluster Controller VM.

```
nutanix@cvm$ ssh nutanix@cloud_controller_vm
```

Amazon S3 is used to store data (extents) and Amazon Elastic Block Store (EBS) is used to store metadata. You can then use the Amazon management tools to manage and monitor billing and related usage. When the AWS Remote feature replicates a snapshot data to AWS, the Nutanix Controller VM on AWS creates a bucket on S3 storage. The bucket name is ntnx-`cluster_id-cluster_incarnation_id-disk_id`.

Implementation Considerations for Cloud Connect

The Cloud Connect feature enables you to back up and restore copies of virtual machines and files to and from an on-prem cluster and a Nutanix Controller VM on the Amazon Web Service (AWS) cloud. The Nutanix Controller VM is created on an AWS cloud in a geographical region of your choice. It is a single-node cluster with a 30 terabyte (TB) disk attached to the node, with a usable disk capacity of 20 TB.

Amazon customers are charged only for capacity that is used (not charged for the full capacity). Once configured through the web console, the remote site cluster is managed and monitored through the Data Protection dashboard like any other remote site you have created and configured. To log in to the Nutanix Controller VM on the cloud, get the IP address from the Data Protection dashboard of the remote site cluster and SSH to the IP address from the source cluster Controller VM.

```
nutanix@cvm$ ssh nutanix@cloud_controller_vm
```

Amazon S3 is used to store data (extents) and Amazon Elastic Block Store (EBS) is used to store metadata. Users then use the Amazon management tools to manage and monitor billing and related usage. When the AWS Remote feature replicates a snapshot data to AWS, the

Nutanix Controller VM on AWS creates a Nutanix Object on S3 storage. The Nutanix Object name is ntnx-`cluster_id-cluster_incarnation_id-disk_id`.

> Note: To avoid data corruption, do not modify the AWS CVM deployment configuration after the replication starts.

Table 22: Network and Port Requirements and Supported Features

Requirement/Feature	Description/Recommendation
Local cluster	The local cluster requires Internet connectivity and must be able to communicate with the public cloud.
Firewall port access	See Firewall Requirements in the *Security Guide*.
Name server	This feature requires that you configure a name server in the Prism web console.
VPN	Ensure that you have configured a VPN connection between AWS subnet and Nutanix local cluster before configuring the cloud remote site through Prism.
AWS Direct Connect	Supported

Table 23: Cloud Connect Features

Feature	Description/Recommendation
Compression	You can choose to enable compression on the local storage container. Compression on wire is enabled by default and is also enabled by default when the remote AWS storage container is created.
Disaster recovery	Not currently supported.
Hypervisors	AWS is supported on all the three hypervisors (ESXi, Hyper-V, and AHV).
Replication type	You can configure multiple AWS Nutanix Controller VMs, where each Controller VM acts as a separate remote site. You can create instances in multiple AWS regions.

General Limitations and Recommendations

- Nutanix recommends to have only one VM for each protection domain that is configured to replicate to a cloud remote site.
- Do not configure vStore protection domains and replication for backup with Cloud Connect feature.
- Do not use the Nutanix Web console of the Controller VM that is running on the cloud as it is not supported.

- Do not enable deduplication on the cloud instance.
- It is not recommended to enable deduplication on the source cluster.
- Do not use the *DoNotUse-sp* storage pool on the cloud instance to create containers.
- The credentials that are used for creating the cloud remote site cannot be deleted without deleting the cloud remote site.
- If your networks are behind the proxy servers, you cannot use the Cloud Connect feature.
- Changing a vStore mapping causes associated protection domains to initiate full replication of protected entities to the newly specified target container as if it were initial replication. A mapping change therefore results in the overconsumption of storage resources at the remote site. Contact Nutanix Support for help with cleaning up snapshots in the previously specified container.

AWS Account Requirements

Following are the account requirements for the AWS cloud. Create an AWS policy and grant the following permissions in the IAM. For more information on the policies, see *AWS Identity and Access Management* in the *AWS documentation*.

> Note: Nutanix does not support AWS Government cloud.

Table 24: AWS Account Requirements

Requirement	Description/Recommendation

Credentials and permissions	Nutanix recommends that you create an AWS account for the remote site.
	AWS remote site/cloud connect feature requires AWS credentials (access key and secret access key) for the corresponding AWS user to be configured. You can configure one of the following user types.
	• AWS root user having all the permissions.
	• Create an AWS user with the following permissions.
	  • List regions
	  • List subnets
	  • List AWS AMIs
	  • Access the VPC or subnet information
	  • Create instance, query instance, and list instances
	  • Create buckets
	  • Create EBS volume snapshots
	  • Read/Write/Delete data in S3 bucket
	  • Manage security groups
	  After the instance is deployed, you can modify the permissions of the AWS user from the AWS console to give access to perform only the following tasks.
	  • Read, write, and delete data to S3 buckets created by the deploy operation.
	  • Create EBS volume snapshots.
	• Before you upgrading the cloud Nutanix Controller VM in the Cloud for AWS, ensure that the AWS account credentials is set to AmazonEC2FullAccess.

Requirement	Description/Recommendation

Recommended access type	Prism web console **Use VPN Connection**
	• Nutanix recommends Virtual Private Cloud (VPC) access through an Amazon VPN Connection to the Nutanix Controller VM.
	• Nutanix recommends Internet access for VPC during initial configuration described in Configuring a Remote Site in AWS on page 201.
	• The Amazon VPC/VPN connect assigns a private static IP address to the Nutanix Controller VM, which is used with Cloud Connect and persists if the remote Nutanix cluster restarts. The public IP address might change, but it does not affect any behavior because this address is only used for outbound S3 access.
	• If you select the **Use VPN Connection** option, ensure that you have configured:
	  • A VPN connection as part of your Amazon VPC service
	  • A VPN application, device, or server (known as a customer gateway) for your on-prem cluster
	You can also create and use VPC endpoints to simplify access to the Amazon S3 from within a VPC. These endpoints are easy to configure, highly reliable, and provide a secure connection to Amazon S3 that does not require the traffic to go through a public IP address. You do not need to make any change on the cloud Controller VM or local Nutanix cluster.
	If you have disabled the option to assign a public IP address on the subnet you are using, the Nutanix cluster cloud instance is created without a public IP address. S3 and EC2 requests are routed through the VPC end point.
	For more information about configuring VPC endpoint for Amazon S3, see *New – VPC Endpoint for Amazon S3* topic.

Alternate access type	Prism web console **Create SSH Tunnel** • This scenario is not supported for production use. • You can use a secure session shell (SSH) to access the AWS instance. • If you select SSH connectivity, the Nutanix Controller VM is reachable through its public IP address by using SSH.

Requirement	Description/Recommendation
NTP and timezone	Time on your source cluster and AWS instance should be correct. If NTP and timezone are not properly configured, connection to AWS might fail. For more information on configuring timezone and NTP, see *AWS documentation*.

Instance Types for AWS

Nutanix supports the following instance types. If you are replicating more than 10TB (maximum of 20 TB) of data on AWS, you need to update your instance type.

Table 25: Supported Instance Types

Cloud Type	Default		Limit
AWS	For all the supported regions	m4.2xlarge	Up to 20 TB.

Note: If you have upgraded the cloud instance, it is recommended to upgrade the instance type to the one mentioned in the table. For more information about upgrading instance type, see KB 3861.

Configuring a Remote Site in AWS

Before you begin

- See Implementation Considerations for Cloud Connect on page 196.
- See AWS Account Requirements on page 198.
- Ensure that you have configured a name server on the local Nutanix cluster through the web console. See Configuring Name Servers in the *Prism Web Console Guide*.

Procedure

1. In the Data Protection dashboard, click **+Remote Site** > **Cloud**.
   The Cloud Site dialog box appears.
2. Select the **AWS** option, and click **Next**.
3. Click **Add New Key** to set the Amazon credentials for the remote site.
   a. Type a name for this set of credentials in the Name text box.
   b. Type the **Access Key ID** and **Secret Access Key** in their respective fields.
   c. Click **Save**.
   d. (Optional) Click **Add New Key** to enter another set of Amazon credentials for each additional dedicated Nutanix Controller VM (for example, if you are configuring multiple AWS Nutanix Controller VMs, where each Controller VM is managed individually as a separate remote site).
4. Click **Next** to go to the Remote Site Settings dialog box.

5. If you select **VPN/Direct-connect** to configure a VPC/VPN or AWS direct connect connection to the Amazon cloud as follows.

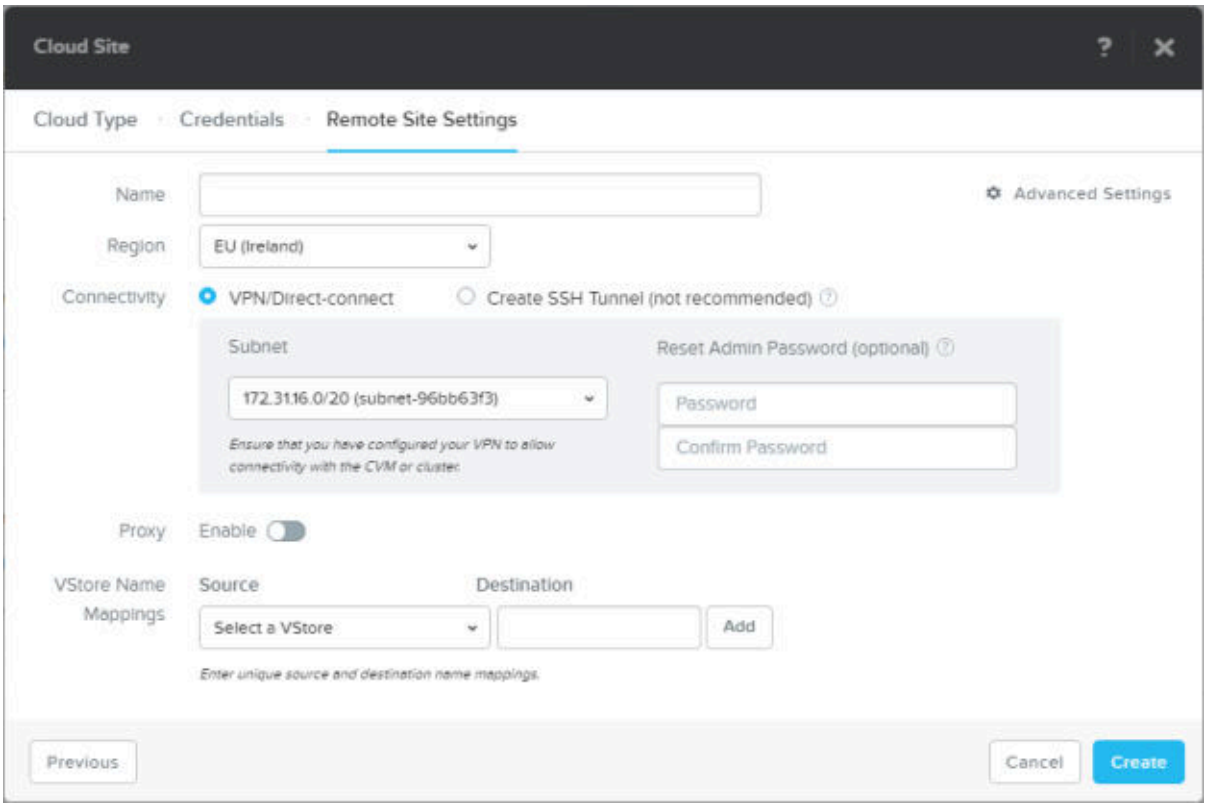

Figure 64: Remote Site Configuration Wizard for VPN to Amazon AWS

a. **Name**: Type a remote site name.

> Note: This entity has the following naming restrictions:
>
> - The maximum length is 75 characters.
> - Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

b. **Region**. From the drop-down menu, select the Amazon region where your Amazon VPC/VPN location is configured.

c. **Subnet**. This field automatically loads all the subnets on your AWS region. Select the subnet in which you want to create the Controller VM.

d. **Reset Admin Password (optional)**. Create a new administrator password for the `nutanix` user account on the cloud-based Nutanix cluster.

e. **Proxy**. Enable to allow addresses in the subnet to be used as a proxy to communicate with other Nutanix components on the remote site. Enabling the proxy is recommended when you have different subnets on primary and secondary sites.

f. **VStore Name Mappings**. Select a local storage container from the **Source** drop-down menu, type the remote storage container name, and then click **Add**. This entry maps

the association between a local source storage container and a remote target storage container at the remote site. Multiple storage containers are supported.

If you perform restore of snapshots on vStore protection domains, you need to modify the VMDK paths in the VM files and point it to the restored path that is specified as part of the restore operation.

g. (Optional) Click **Advanced Settings** to configure compression of data.

   - **Compress on wire**. Enabled by default. Enable this setting if you want to compress data on the cluster before transmitting it. Data is decompressed on the remote cluster.

h. Click **Create**.

6. If you select **Create SSH Tunnel** for configuring a SSH tunnel connection to the Amazon cloud, do the following.

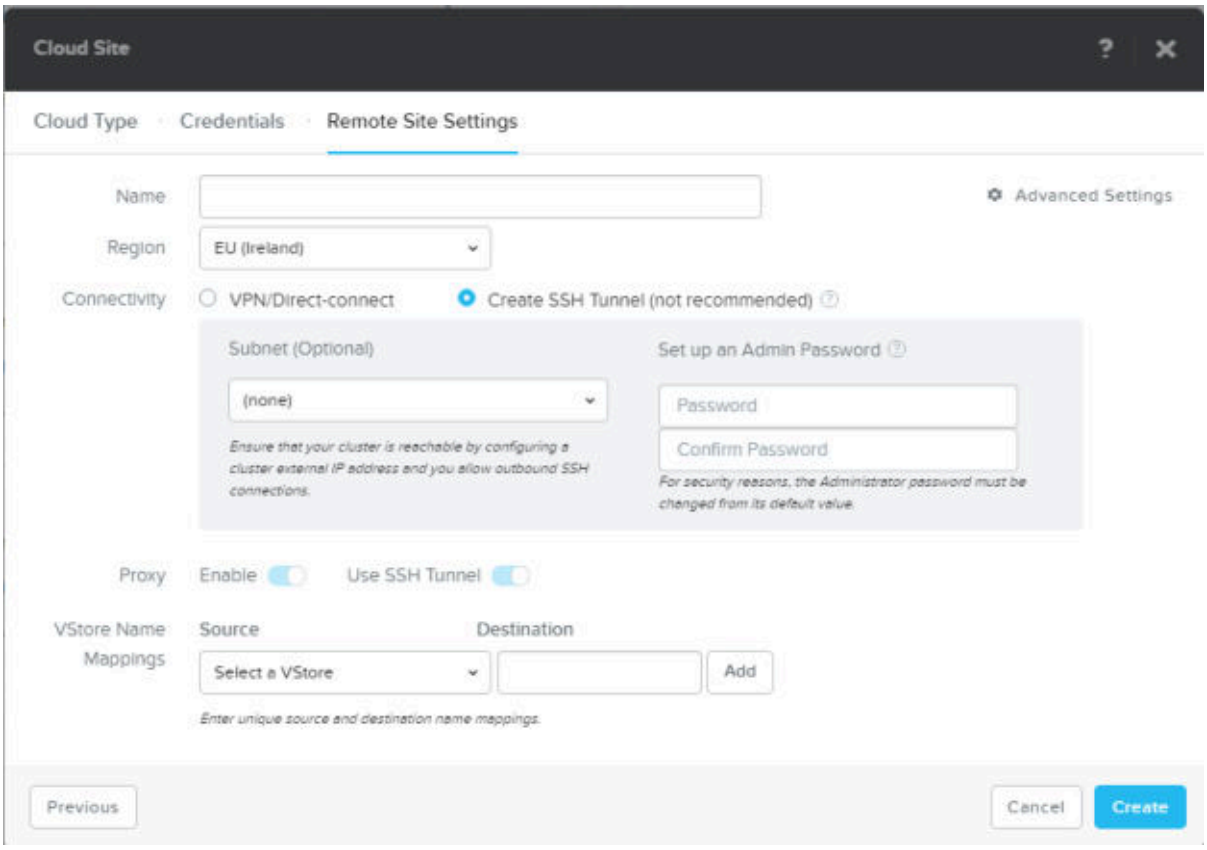

Figure 65: Remote Site Configuration Wizard for SSH to Amazon AWS

a. **Name**: Type a remote site name.

> Note: This entity has the following naming restrictions:
>
> - The maximum length is 75 characters.
> - Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

b. **Region**. From the drop-down, select the Amazon region where your Amazon location is configured.

c. **Subnet**. This field automatically loads all the subnets on your AWS region. Select the subnet in which you want to create the Controller VM.

d. **Set up an Admin Password**. Create a new administrator password for the `nutanix` user account on the cloud-based Nutanix cluster. This field is required.

e. **Proxy**. Enable to allow addresses in the subnet to be used as a proxy to communicate with other Nutanix components on the remote site. Enabling the proxy is recommended when you have different subnets on primary and secondary sites.

f. **Use SSH Tunnel**. Enable an SSH tunnel.

g. **SSH Port.** Type an SSH port number for the SSH tunnel. Valid range is 3000 to 3099.

h. **VStore Name Mappings**. Select a local storage container from the **Source** drop-down, type the remote storage container name, then click **Add**. This entry maps the association between a local source storage container and a remote target storage container at the remote site. Multiple storage containers are supported.

i. (Optional) Click **Advanced Settings** to configure compression of data.

- **Compress on wire**. Enabled by default. Enable this setting if you want to compress data on the cluster before transmitting it. Data is decompressed on the remote cluster.

j. Click **Create**.

The Prism web console displays a message similar to `Remote site remote_site_name has been initiated successfully`.

The remote site initiation and creation takes about 10 minutes. You can see the initiation, creation, and results in the web console Progress Monitor.

What to do next

You have now created a remote site that can be used just like any other remote site for use with data protection.

Recovery Procedures by using Cloud Connect

If for some reason the entities on the primary cluster fails, you can recover the entities by using the snapshots that are being replicated to the AWS. If your primary cluster fails altogether, you can create a new cluster and use the snapshots residing on the AWS to recover the cluster.

About this task

The process of restoring protected entities from the AWS is same as the way you restore entities for the normal cluster.

Procedure

1. To recover the entities, login to the Prism and follow the steps described in Restoration of Protected Entities on page 178.
2. To recover the cluster, see Recovering from the Remote Snapshots on a Backup or DR Site on page 182 for more information.

Upgrading a Nutanix Controller VM in the Cloud

If you are planning to upgrade the cloud Nutanix Controller VM in the Cloud to 5.1.x or later versions, use the following procedure.

About this task

Nutanix supports replications between the all the latest supported LTS and STS released AOS versions. To check the list of the latest supported AOS versions, see KB-5505. Check the EOL document to determine if the AOS versions currently running on the clusters are EOL.

Upgrade the version to the next available supported LTS/STS release. Before you upgrade the AOS, check the Upgrade Paths page to determine if an upgrade path is supported.

> Note: If both clusters have different AOS versions that are EOL, upgrade the cluster with lower AOS version to match the cluster with higher AOS version and then perform the upgrade to the next supported LTS version as advised above.

For example, the clusters in a metro configuration are running AOS versions 5.5.x and 5.10.x respectively. Upgrade the cluster on 5.5.x to 5.10.x. After both the clusters are on 5.10.x, proceed to upgrade each cluster to 5.15.x (supported LTS). Once both clusters are on 5.15.x you can upgrade the clusters to 5.20.x or later.

Nutanix recommends that both the primary and the replication clusters or sites run the same AOS version.

Before you begin

- Before you start the upgrade process, delete any existing install directory in **/home/nutanix** by running the `nutanix@cvm$ rm -rf /home/nutanix/install` command.
- Download the latest AOS build and metadata JSON file from the Nutanix support portal at the **Downloads** link and copy it to cloud Controller VM at the **/home/nutanix** directory.
- Before you upgrading the cloud Nutanix Controller VM in the Cloud for AWS, ensure that the AWS account credentials is set to AmazonEC2FullAccess.

Procedure

1. Run the following command to unzip and extract the upgrade script to the **install/bin** directory.

   ```
   nutanix@cvm$ tar zxf build.tar.gz ./install/bin/upgrade_cloud_cvm
   ```

2. Run the following command to upgrade the cloud Controller VM.

   ```
   nutanix@cvm$ install/bin/upgrade_cloud_cvm --installer_file=build_file.tar.gz \
   --metadata_file=json_file.json
   ```

   The upgrade procedure may take up to one hour to complete. The cloud Controller VM restarts at the end of the upgrade process. Run the **upgrade_status** command to verify that upgrade is successfully completed.

13

SUPPORT SERVICES

Nutanix provides support services in several ways.

- Nutanix technical support can monitor your clusters and provide assistance when problems occur (see Controlling Remote Connections on page 211, Configuring HTTP Proxy on page 212, and Configuring Pulse on page 209).
- Nutanix technical support maintains a portal that you can access to request assistance, download AOS updates, or view documentation (see Accessing the Nutanix Support Portal on page 215).
- Nutanix supports a REST API that allows you to request information or run administration scripts for a Nutanix cluster (see Accessing the REST API Explorer on page 218).

Pulse Health Monitoring

The feature known as Pulse provides diagnostic system data 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 a Nutanix cluster. Information includes:

- system alerts
- current Nutanix software version
- Nutanix processes and Controller VM information
- hypervisor details such as type and version

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 Pulse Access Requirements on page 211). For the complete list of required ports, see Port Reference.

> Note: When logging into Prism the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a screen appears to enable Pulse (see *Logging Into the Web Console* topic in the *Prism Web Console Guide*).

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" in the *Prism Central Guide*), 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 HTTP Proxy on page 212). 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 Pulse Access Requirements on page 211). 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 207 for Pulse configuration recommendations.
- Pulse sends messages through ports 80/8443/443. For the complete list of required ports, see Port Reference.
- When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse. To enable Pulse, click the **Continue** button in the message and follow the prompts; to continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

You can enable (or disable) Pulse at any time. To configure Pulse, do the following:

Procedure

1. Go to the Settings menu (see *Settings Menu* in the *Prism Web Console Guide*) and select **Pulse**.

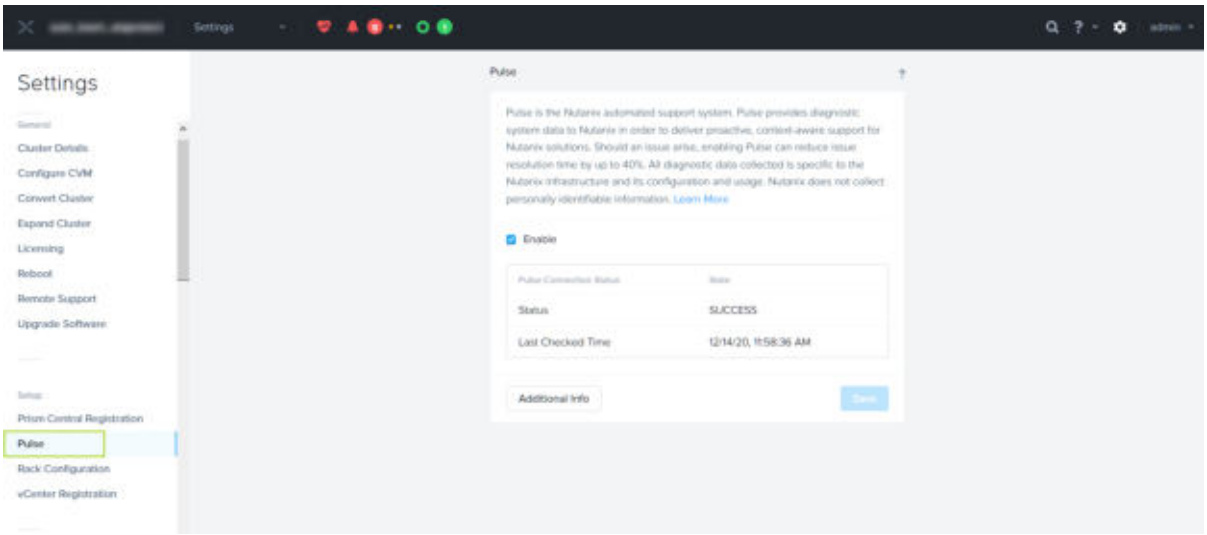

Figure 66: Pulse Window

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 to display additional connection and recipient information.

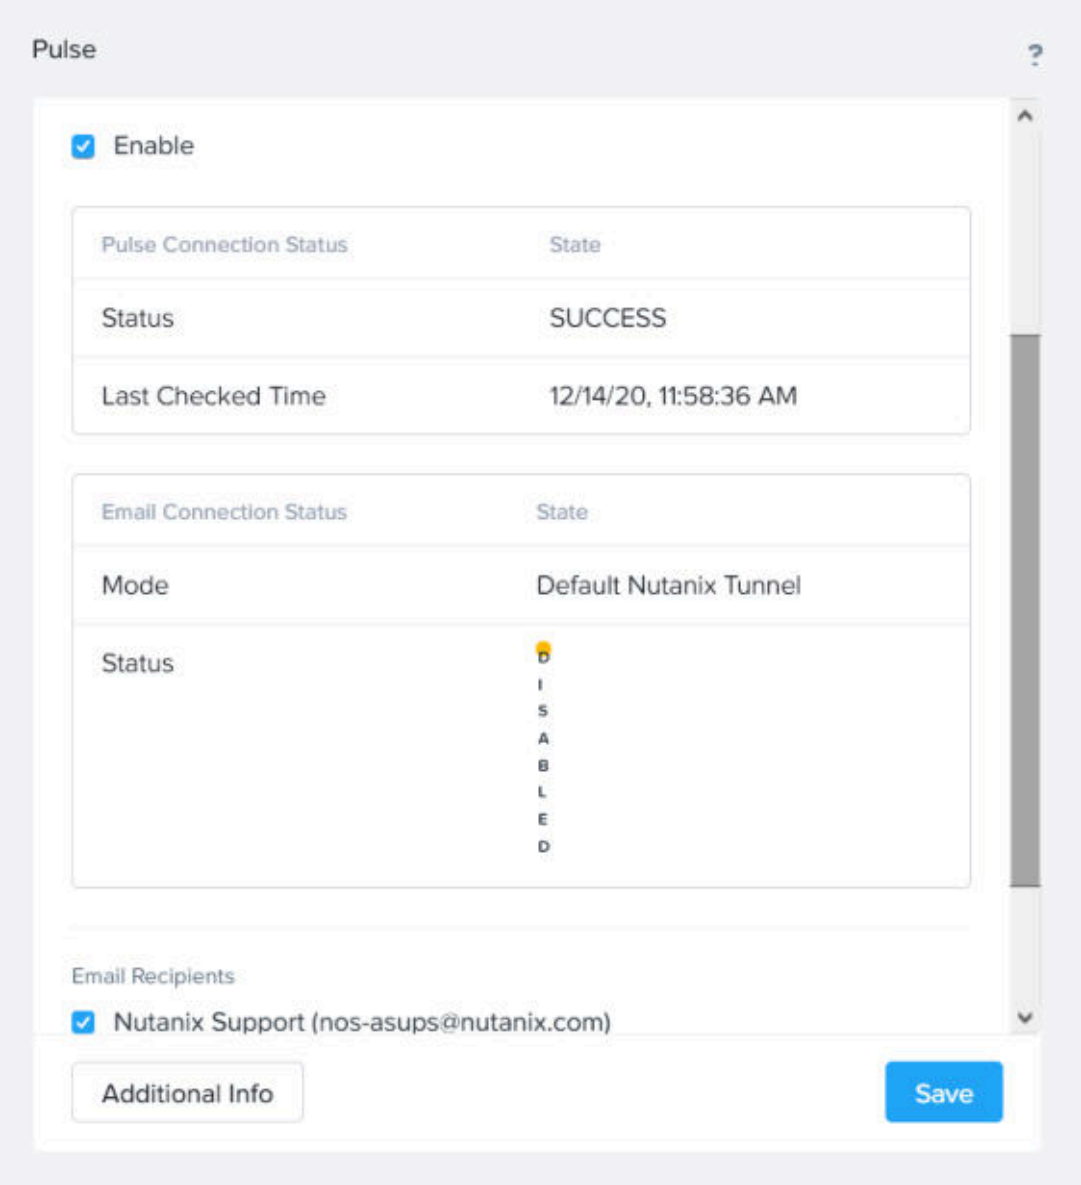

Figure 67: Pulse Additional Info

4. Check the **Pulse Connection Status** field for connection information.

   - **Status**: Displays the transport mechanism status, which is **Success** when the HTTP client can access the Pulse server successfully or **Failure** when it cannot (or **Unknown** when the status is not known).
   - **Last Checked Time**: The last time Pulse sent data and checked the connectivity with the Pulse server.

5. Check the **Email Connection Status** field for connection 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 **Disabled** when neither is available (or **Unknown** when the status is not known).

     If the status is **Success** the following two fields also appear.
   - **Service Center**: Displays the address of the Nutanix support service center to which the Pulse data is sent.
   - **Established Since**: Displays the date when the current Pulse session was established.

6. To send the cluster information to additional recipients, enter one or more e-mail addresses in the **Comma Separated Email List** field.

   You can e-mail the information to yourself or others by entering recipient e-mail addresses in this field. Use a comma to separate multiple addresses.

7. Click the **Save** button to save the new settings.

Pulse Access Requirements

In order to successfully send Pulse and alert messages from a cluster to the Nutanix support servers, Pulse requires the following access:

- Messages are sent from the Zeus (cluster configuration manager) leader, so the firewall must allow the Zeus leader IP address. Because the Zeus leader can change, it is recommended that the IP addresses for all Controller VMs in the cluster be open in the firewall.
- Required firewall rules configured to allow traffic from the cluster to Nutanix support servers. See Firewall Requirements in the *Security Guide*.

Pulse data is collected through the following methods:

- By default, emails are sent daily through port 8443 via nsc01.nutanix.net or nsc01.nutanix.net to nos-asups@nutanix.com.
  - If a SMTP server is configured on a cluster, emails will be sent via the SMTP server to nos-asups@nutanix.com.
- Throughout the day, each Controller VM in the cluster sends Pulse data to insights.nutanix.com over port 443 using the HTTPS REST endpoint.

> Note:
>
> Firewall port requirements are automatically met if you have a Prism Central deployment, see "Prism Central Proxy for Pulse Data" in the *Prism Central Guide*.

Controlling Remote Connections

About this task

Nutanix technical support can remotely assist with problems by logging into the cluster through an SSH connection. This remote connection is disabled by default, but it may be enabled during the site installation procedure or at a later date through the web console.

You must configure rules in your external firewall to allow traffic from the cluster to Nutanix remote support servers. See Firewall Requirements in the *Security Guide*.

To enable (or disable) Nutanix technical support remote access to your cluster through this connection, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Remote Support** in the **Settings** page. The Remote Support dialog box appears.
2. Select (click the radio button for) the desired access state.
   - To allow remote access (temporarily) for Nutanix support, select **Enable** and enter the number and duration (hours or minutes) in the two boxes. When remote support is enabled temporarily, an icon appears in the screen header that includes a dynamic countdown until remote support is disabled.

   - To prevent remote access by Nutanix support, select **Disable**.
3. Click the **Save** button to save the new setting and close the window.

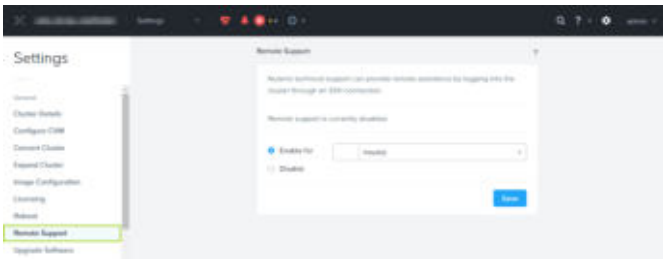

Figure 68: Remote Support Window

Configuring HTTP Proxy

About this task

If the customer site 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 Proxy dialog box appears.

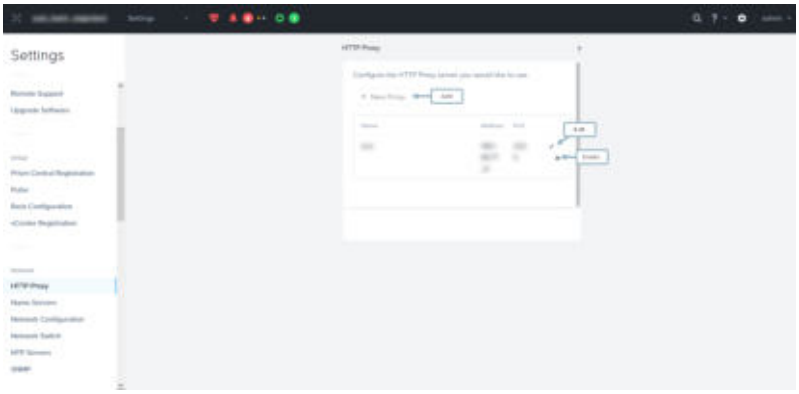

Figure 69: HTTP Proxy Window

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**: Check the box(es) for the protocol to proxy (HTTP, HTTPS, or both).

   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.

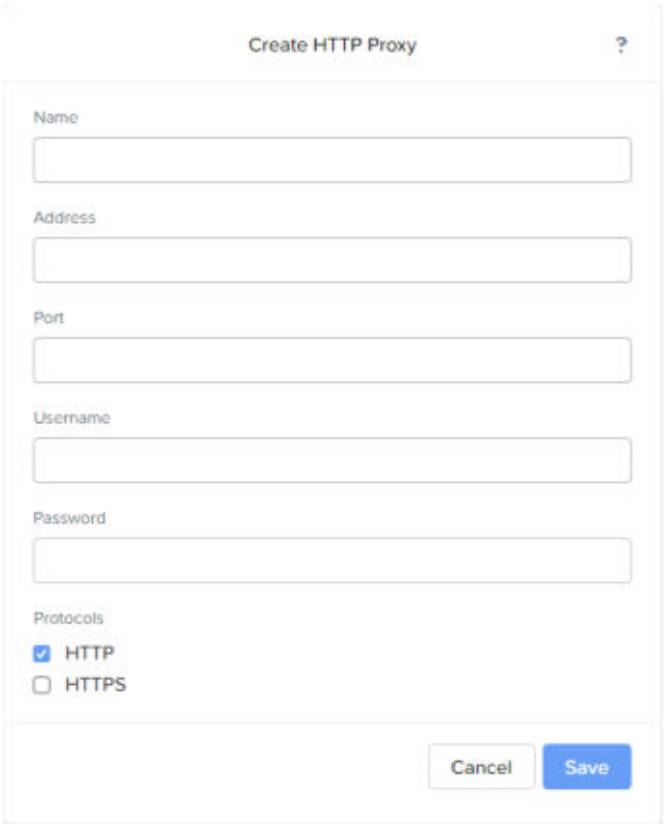

Figure 70: Create HTTP Proxy Window

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.

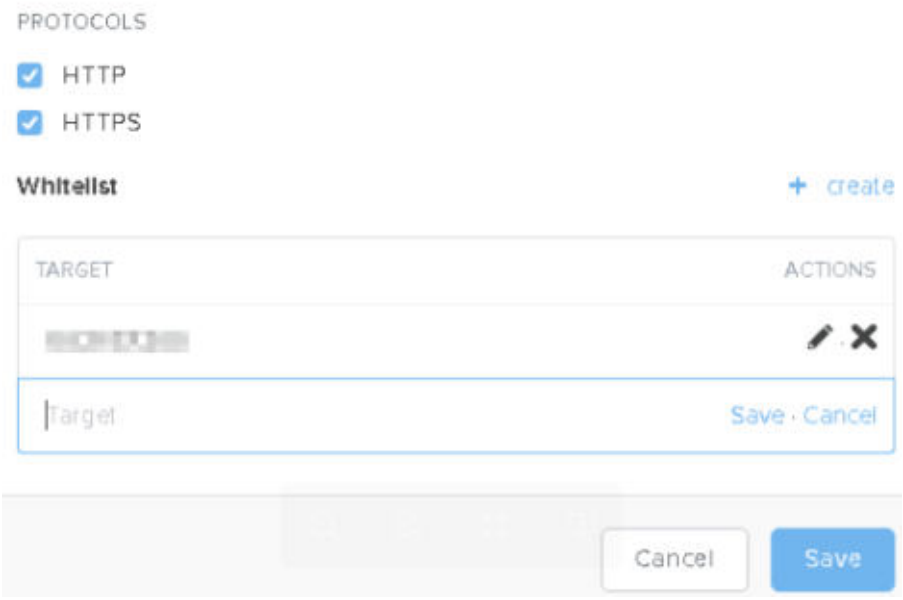

Figure 71: Whitelist Targets

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.

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, select **Support Portal** from the question mark icon

pull-down list.
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 26: Main Menu Options

Category	Option	Description
Documentation	Product Documentation	Displays a page from which you can view the Nutanix product 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.

Category	Option	Description
	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.
	Acropolis Upgrade Paths	Displays a table of 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 & Forums	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.
	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) and Prism Central VM.
	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.

Nutanix REST API

The Nutanix REST APIs allow you to create scripts that run system administration commands against the Nutanix cluster. The API enables the use of HTTP requests to get information about the cluster as well as make changes to the configuration. Output from the commands are returned in JSON format.

There are two versions of the Nutanix REST API.

- v1: The original Nutanix REST API.
- v2: An update of the v1 API. Users of the v1 API are encouraged to migrate to v2.

A complete list of REST API functions and parameters is available in the REST API Explorer section of the Nutanix web console. In addition, the complete reference for the v2 Nutanix API, including code samples in multiple languages, and tutorials are available at http://developer.nutanix.com/.

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 v1 and v2 APIs can both be viewed in the REST API Explorer.

Procedure

1. Open the explorer for the desired version of the API.

   » v1: Connect to the Nutanix web console, click the user icon in the upper-right corner of the web console, and select **REST API Explorer**. In the explorer, select **Version 1** from the menu.

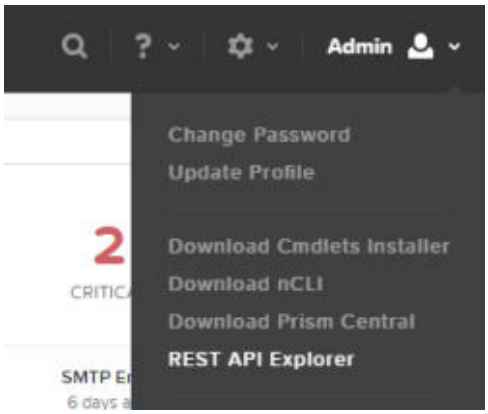

Figure 72: User menu

   » v2: Connect to the Nutanix web console, click the user icon in the upper-right corner of the web console, and select **REST API Explorer**. In the explorer, select **Version 2** from the menu.

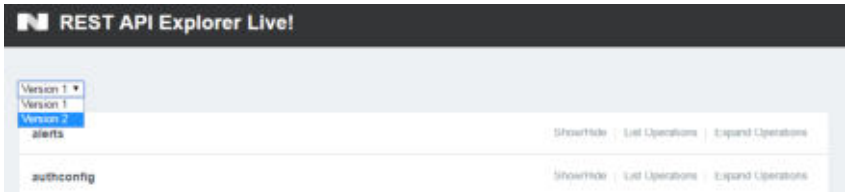

Figure 73: API version selection

The REST API Explorer displays a list of the cluster 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

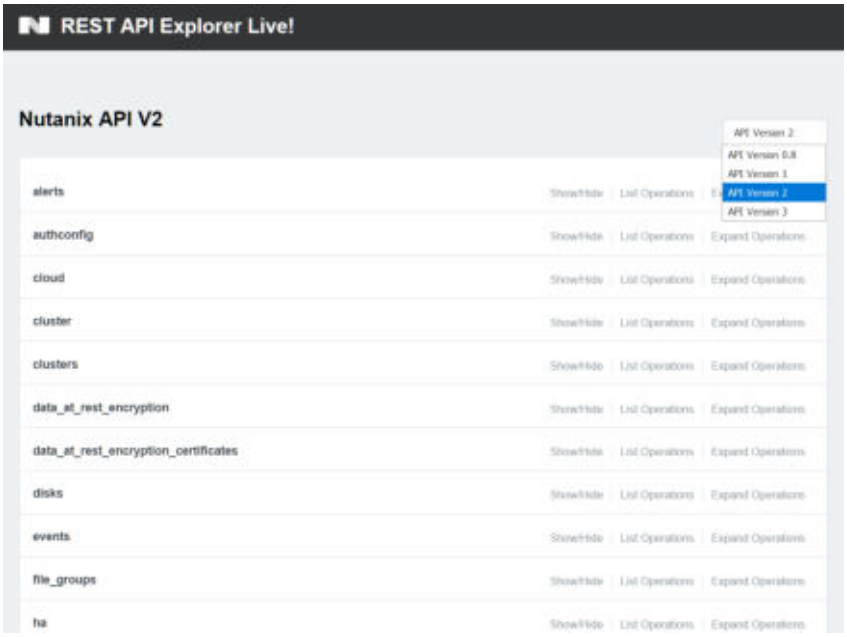

Figure 74: REST API Explorer

Tip: The objects are listed by a relative path that is appended to the base URL `https://management_ip_addr:9440/PrismGateway/services/rest/v[1,2,3]/api`, where *`management_ip_addr`* is the IP address of any Nutanix Controller VM in the cluster.

2. Find the line for the object you want to explore and click **Expand Operations**. For this example, you will operate on a storage pool.

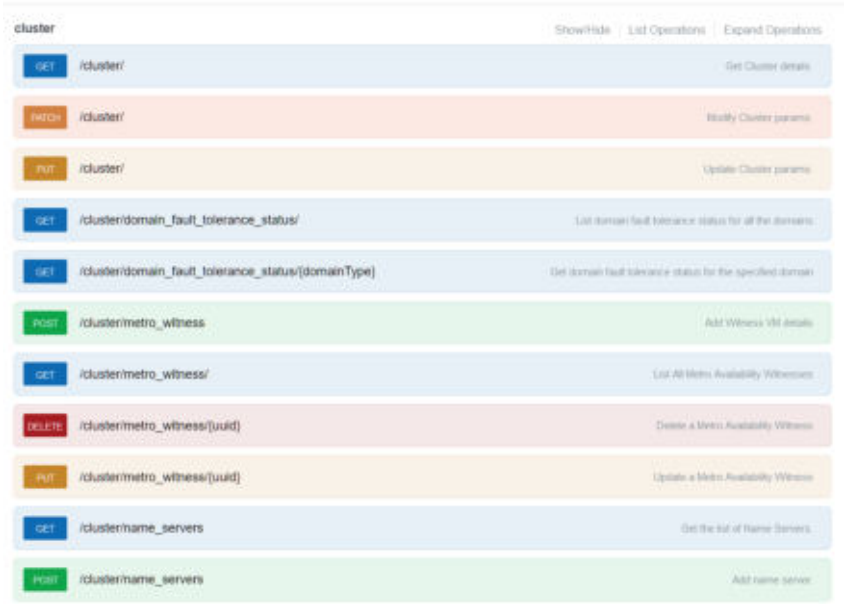

Figure 75: Cluster Operations

3. Click **GET** on the first line to show the details for this API call.

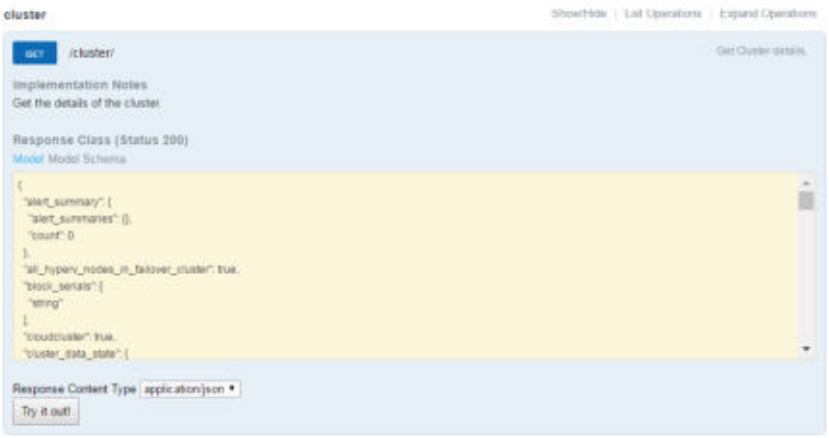

Figure 76: Get Cluster Details

The explorer displays the parameters that can be passed when this action is used.

4. Click **Try it out!** to test the API call when used with your cluster.

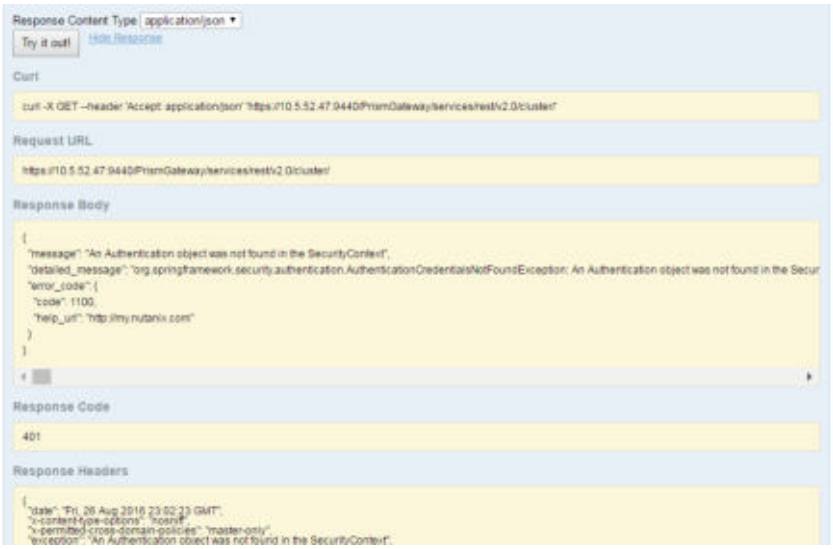

Figure 77: Get Cluster Details Response

The test displays the request URL required for scripting, as well as sample responses expected from the API call.

14

HELP RESOURCES

There are several information sources that you can access at any time when you need help:

- Context-sensitive help documentation (see Accessing Online Help on page 222).
- Health dashboard tutorial (see *Health Dashboard* in the *Prism Web Console Guide*).
- Customer support portal (see Accessing the Nutanix Support Portal on page 215).
- Nutanix community forum (see Accessing the Nutanix Next Community on page 223).
- REST API explorer (see Accessing the REST API Explorer on page 218).
- Glossary of terms (see Glossary on page 223).

Accessing Online Help

About this task

The web console 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 (see Main Menu Options in Cluster Management in the *Prism Web Console Guide*):

   » 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 **General Help** 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.

2. To select a topic from the table of contents, click the 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.

3. To display all the help contents as a single document (*Web Console Guide*), click the **epub** or **pdf** button in the upper right.

   You can view the *Web Console 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 **Other** 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.

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, select **Nutanix Next Community** from the question mark icon pull-down list of the main menu (see Main Menu Options in Cluster Management in the Prism Web Console Guide).
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.

Glossary

aCLI

Acropolis command-line interface.

Acropolis

The Nutanix converged software fabric for virtualization and storage management. It consists of the Acropolis base software, Acropolis Distributed Storage Fabric, AHV, App Mobility Fabric, Prism, and Acropolis APIs.

Acropolis App Mobility Fabric

Provides virtualization management, volume management, and other distributed management functions for Acropolis.

AHV

Nutanix-specific hypervisor host providing core server virtualization and optimized for Nutanix cluster and guest VM use.

block

A set of Nutanix nodes contained in a single enclosure.

block fault tolerance

When certain conditions are met, redundant copies of any data required to serve I/O are placed on nodes that are not in the same block. In the case where multiple nodes in a block fail, the cluster can continue to run because cluster configuration data exists on other blocks. Guest VMs can continue to run because redundant copies of guest VM data and metadata exist on other blocks.

clone

A writeable copy of a vDisk.

Cloud Connect

A feature that enables you to back up and restore copies of virtual machines and files to and from an on-premise cluster and a Nutanix Controller VM located on the Amazon Web Service (AWS) cloud. The cloud-based cluster is managed as a remote site from the Prism Web Console or nCLI.

cluster

A group of nodes contained in one or more Nutanix blocks.

cold data

Data that did not have many rewrites or overwrites for a long time. For example, snapshots, file servers, archives, or backups.

compression

An optional storage container setting that enables one of two types of compression.

compression, inline

Data is compressed as it is written. This type of compression is recommended for workloads that perform batch processing.

compression, post-process

Data is compressed after it is written. The delay time between write and compression is configurable. Nutanix has recommends a delay value of 60. However, every workload has a different I/O profile, so you might require a different delay value otherwise. This type of compression is recommended for most workloads.

Consistency Group

A subset of VMs in a Protection Domain. All VMs within a Consistency Group for that Protection Domain will be snapshotted in a crash-consistent manner. For all VMs in a Consistency Group, a snapshot creates one snapshot for all VMs in the group.

Controller VM

A Nutanix VM that manages storage and other cluster functions on a node.

data resiliency

A measure of the types of failures a cluster can withstand; affected by block awareness and redundancy factor.

datastore

A logical storage container for files necessary for VM operations.

deduplication

The sharing of identical guest VM data on premium tiers (RAM and Flash) for improved performance or on capacity tiers (HDD) for storage space savings. Enabled by properties of a storage container or vDisk.

Distributed Storage Fabric

All storage functionality, including snapshots and clones, data protection, disaster recovery, data path redundancy, replication factors, deduplication, compression, erasure coding, and so on.

Enterprise Cloud Platform

The Nutanix solution that natively converges compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence

erasure coding

Optional algorithm included in the Acropolis base software to help reduce the storage used for fault tolerance. It helps to increase the effective or usable capacity on a cluster, depending on cluster size and data coldness.

failback

Planned failover initiated from recovery site.

failover

Moving VMs from a primary site to a recovery site.

failover, disaster

Failover when the primary site is down.

failover, planned

Failover when both sites are up.

guest VM

A VM running on a Nutanix cluster that executes a workload, such as VDI or Exchange, as opposed to a VM that is involved in cluster operations, such as a Controller VM.

host

An instance of the hypervisor that runs on a Nutanix node.

image service

A workflow in the Prism web console that enables a user to upload ISO or disk images (in ESXi or Hyper-V format) to a Nutanix AHV cluster by specifying a remote repository URL or by uploading a file from a local machine.

local replication

Multiple copies of data within a storage container. These copies exist for fault tolerance: if a physical disk fails, the cluster can recover data from another copy. The cluster manages the replicated data, and the copies are not visible to the user.

local snapshot

Snapshots stored on the same cluster where they originated.

nCLI

Nutanix command-line interface.

node

A physical server contained in a Nutanix block; runs a hypervisor host.

oplog

A write cache on a faster, more expensive storage tier.

Prism

Web-based management interface for managing Nutanix clusters.

Prism Central

Centralized management tool that runs as a separate VM configured as a single-node cluster to monitor and manage multiple clusters through a single web console.

Prism Element

A single cluster being managed by and available through the Prism Central web console.

Protection Domain

A group of VMs to be backed up locally on a cluster or replicated on the same schedule to one or more clusters. Protection Domains may be associated with remote sites. Protection Domain names must be unique across sites. A VM can be in at most one Protection Domain.

Protection Domain, active

A Protection Domain that manages live VMs and makes, replicates, and expires snapshots.

Protection Domain, inactive

A Protection Domain that receives snapshots from a remote cluster.

redundancy factor

The number of nodes plus 1 that the cluster can tolerate being down at one time. By default, Nutanix clusters have a redundancy factor of 2, which means that they can tolerate 1 node being down. They are configurable to redundancy factor 3 to enable tolerating 2 nodes being down.

remote replication, one-to-one

Replicates a production cluster with one idle cluster as a DR target.

remote replication, reciprocal

Cross replication within running (non-idle) production clusters.

remote site

A pair of clusters that can replicate data to each other. A configured remote site can also be located in the cloud (based on Amazon AWS, for example).

remote snapshot

A snapshot copied asynchronously from one cluster to another.

replication factor

The number of copies of data kept by a storage container. The default is 2. Storage Containers on clusters with redundancy factor of 3 can have replication factor of 3.

reserved capacity

A property of a storage container or vDisk that guarantees that a certain amount of storage space is available.

schedule

A property of a Protection Domain that specifies the intervals to take snapshots and how long the snapshots should be retained. A schedule optionally specifies which remote site to replicate to.

Self-Service Restore

Allows a user to restore a file within a virtual machine from the Nutanix protected snapshot with minimal Nutanix administrator intervention.

shadow clone

A cache of a vDisk on all the nodes in the cluster. When a vDisk is read by multiple VMs (such as the base image for a VDI clone pool), the cluster creates shadow clones of the vDisk.

snapshot

A read-only copy of the state and data of a VM at a point in time. Snapshots for a VM are crash consistent, which means that the VMDK on-disk images are consistent with a single point in time. That is, the snapshot represents the on-disk data as if the VM crashed. The snapshots are not, however, application consistent, meaning that application data is not quiesced at the time of snapshot, unless the Protection Domain is configured to use application-consistent snapshots.

storage container

A subset of available storage within a storage pool.

storage pool

A group of physical disks from one or more tiers.

storage replication adapter

A Nutanix-provided module that allows VMware Site Replication Manager (SRM) to use native remote replication.

tier

A type of physical storage in a Nutanix node.

vDisk

Data associated with a VM represented as a set of files on a datastore.

VM high availability

In virtualization management VM high availability, when a node becomes unavailable, VMs that are running on that node are restarted on another node in the same cluster.

VM mobility

The ability to export your existing VMs from one non-AHV cluster to an AHV cluster. This option requires that you install the Nutanix VM Mobility installer on all the VMs. The Nutanix VM Mobility installer deploys the drivers that are required at the destination AHV cluster. After you prepare the source VMs, they can be exported to the AHV cluster.

vStore

A separate mount point within a storage container which has its own NFS namespace. This namespace maps to a Protection Domain. Each vStore is exported as a device through the Nutanix SRA.

COPYRIGHT

Copyright 2022 Nutanix, Inc.

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

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

License

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

Conventions

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

Version

Last modified: March 23, 2022 (2022-03-23T22:05:30+05:30)

SOLUTION BRIEF

Simplifying security for any app at any scale on hybrid cloud and multi-cloud deployments

Securing Virtual Networks with Nutanix Flow and FortiGate-VM

The Challenge

Organizations of various sizes have cloud adoption at the center of their Digital transformation strategy. The movement of applications to hybrid cloud and multi-cloud brings immense benefits but increases the cyberattack surface. They face a double-edged sword with a tightening regulatory environment and rapidly worsening cyber threat landscape. Costs of penalties are high and risk to brand reputation is even higher. There is a critical need to identify, inspect, visualize and secure ingress and egress application traffic flows across hybrid cloud and multi-cloud networks.

The Joint Solution

Fortinet and Nutanix form a software-defined solution with secure, scalable, hyperconverged infrastructure for virtualized applications. The FortiGate-VM running on Nutanix AHV provides deeper visibility into applications and prevents known and Zero-day cyber threats, while Nutanix enables you to meet the growing demands of enterprise applications, hybrid-cloud, and multi-cloud environments.

FortiGate-VM enables hybrid cloud and multi-cloud traffic to be segmented based on the application and then inspected for known and Zero-day threats such as C&C, malware, ransomware etc. Policy based redirection from Flow™ gives NetOps and SecOps teams granular control over traffic inspection and optimizes resource consumption. Automation features and centralized management enable you to ensure that security policies can keep pace with any contextual changes in your enterprise cloud environment.

Nutanix and Fortinet have validated two modes for deployment flexibility. Choose layer 3 routed mode for FortiGate-VM on Nutanix AHV to inspect traffic while acting as a layer 3 gateway. Use layer 2 vWire mode as part of Nutanix Flow traffic redirection to transparently direct VM traffic through a FortiGate-VM firewall running on every AHV host.

- L7/ DPI / Advanced Security with FortiGate-VM
- Cluster wide deployment of service chains with Nutanix Calm
- Multiple Services Per Chain
- Policy Based Selective Traffic Redirection

Joint Solution Components

- Fortinet FortiGate Next-Generation Firewall, FortiManager, FortiAnalyzer, FortiADC, FortiMail, FortiSandbox, FortiSIEM, FortiAuthenticator, FortiNAC, FortiProxy

Joint Solution Benefits

- Segment and control East-West VM traffic
- Advanced L7 security protection and automated remediation
- Add policy based Deep Packet Inspection (DPI) between virtual machines (VMs)
- Identify and control traffic flowing into your secure environment; limit application access based on users and groups; block known and unknown threats
- Securely deploy applications faster and scale at the pace of modern business
- Predictable scale-out architecture to seamlessly expand desktop and server virtualization pilots to full production deployment

FORTINET FABRIC-READY

NUTANIX READY AHV

Diagram of Joint Solution

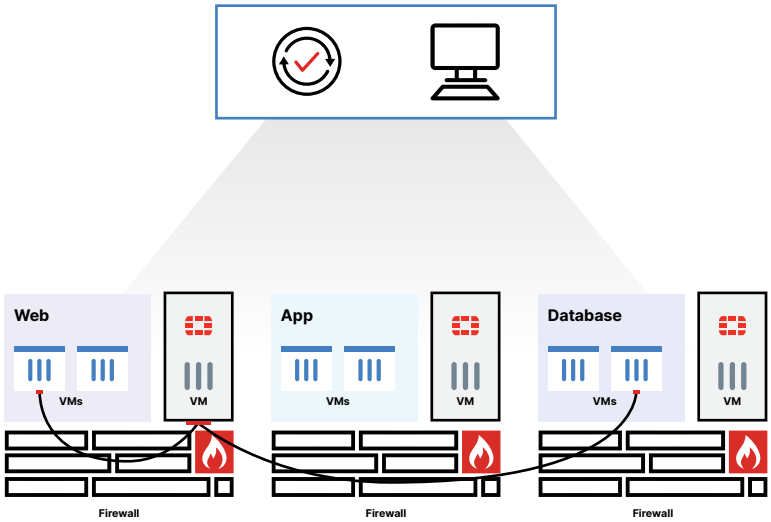

Figure 1: Service chaining Nutanix deployments with Nutanix Flow™ and FortiGate VM.

Solution Components

Fortinet Security Fabric

The Security Fabric allows security to dynamically expand and adapt as more and more workloads and data are added. Security seamlessly follows and protects data, users, and applications as they move between Internet of Things (IoT), devices, and cloud environments throughout the network.

FortiGates are the foundation of Security Fabric, expanding security via visibility and control by tightly integrating with other Fortinet security products and Fabric-Ready Partner solutions.

Fortinet FortiGate VM Next-generation Firewall

Your virtualized data center assets need advanced protection from evolving threats, both known and unknown.

FortiGate VM is a virtualized form factor of our market leading, high performance FortiGate next-generation firewall (NGFW) that delivers advanced security protection for north-south and east-west traffic in software-defined data centers and cloud. Based on the powerful FortiOS operating system and FortiGuard Threat Intelligence services, FortiGate virtual NGFW delivers industry-leading performance and layered threat protection of your software-defined data center traffic, with a single pane-of-glass to manage your physical and virtual network.

Nutanix Flow Networking

Nutanix Flow™ is a software-defined network (SDN) policy engine built into AHV virtualization, with nothing to install and no requirements to alter or reconfigure your physical networking. Flow™ provides VM level application microsegmentation, virtual network segmentation, and policy-based network service insertion.
Flow makes it easy to visualize and discover VM communications and create policy for all virtualized network traffic.

Using service insertion, Flow enables policy-based VM traffic redirection through FortiGate VM firewalls, expanding visibility and control.

About Nutanix

Nutanix, the leader in hypervconverged infrastructure (HCI), makes datacenter infrastructure and clouds invisible, elevating IT to focus on business applications and services. Its Enterprise Cloud OS software converges private, public and distributed clouds, bringing one-click simplicity and agility to infrastructure and application management. This enables IT to rapidly deliver against business needs at a favorable TCO, while retaining hardware and virtualization technology that best suit their skills.

Learn more at www.nutanix.com.

Copyright © 2021 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

July 30, 2021 10:22 AM
401319-B-0-EN

Platform ANY

NX-1065-G8 System Specifications

NX-1065-G8

March 4, 2022

Contents

1. Nutanix Hardware Naming Convention....3

2. System Specifications....7
Node Naming (NX-1065-G8)....7
NX-1065-G8 System Specifications....8

3. Component Specifications....13
Controls and LEDs for Multinode Platforms....13
LED Meanings for Network Cards....15
Power Supply Unit (PSU) Redundancy and Node Configuration....17
Nutanix DMI Information....18

4. Memory Configurations....20
Supported Memory Configurations....20

Copyright....25
License....25
Conventions....25
Default Cluster Credentials....25
Version....26

1

NUTANIX HARDWARE NAMING CONVENTION

Every Nutanix block has a unique name based on the standardized naming convention followed by Nutanix.

The Nutanix hardware model name translates to the following naming convention:

Prefix-Body-Suffix or *NX-WXYZ(G|C|S)-G(4|5|6|7|8)*

Where:

- *NX* is the Nutanix appliance
- *W* is the product series
- *X* is the number of nodes
- *Y* is the chassis form-factor
- *Z* is the disk form-factor
- *G|C|S* indicates that a platform uses GPUs, a platform is an AHV-only node, or a platform supports only one CPU respectively
- *G(4|5|6|7|8)* is the CPU generation

Example: *NX-3460-G8*

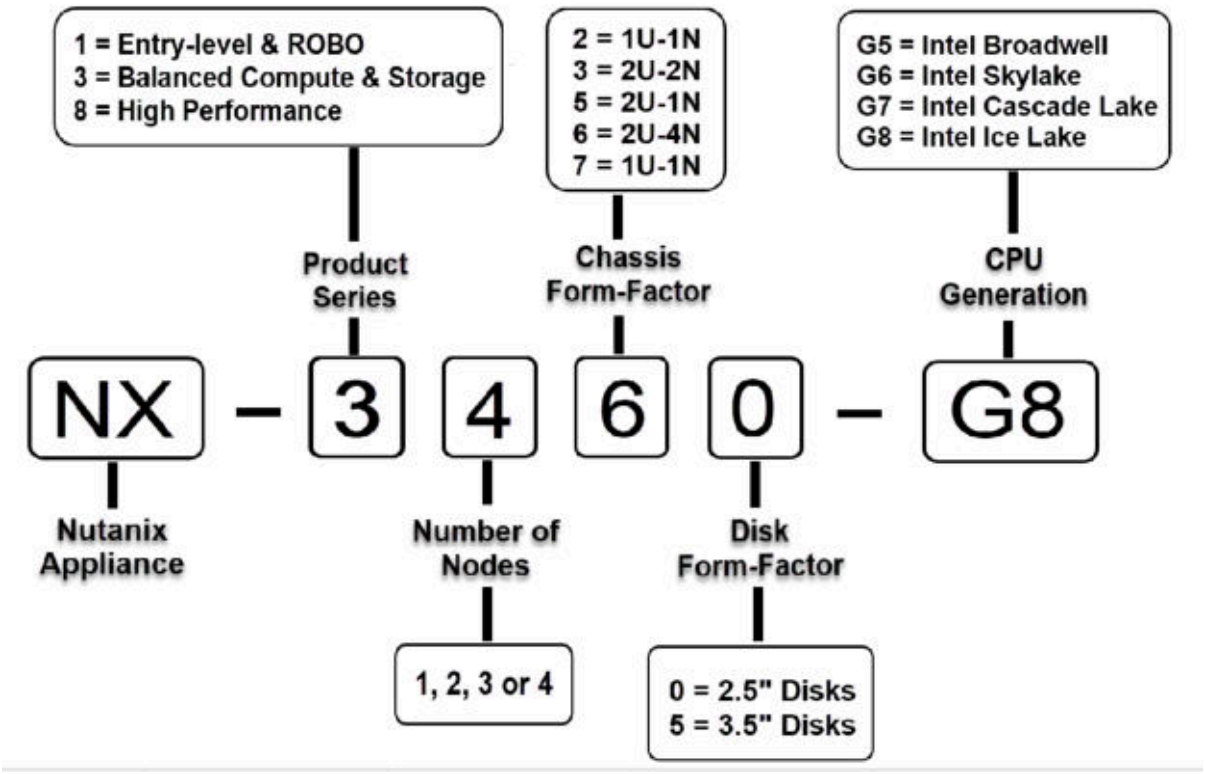

Figure 1: Nutanix Hardware Naming Convention

The following section provides a detailed description of the Nutanix hardware naming convention:

Table 1: Prefix

Prefix	Description
NX	Indicates that the platform is sold directly by Nutanix and support calls are handled by Nutanix. *NX* stands for *Nutanix*.

Table 2: Body

Body	Description
W	Indicates the product series and have one of the following values: • *1* is for small or Remote Office/Branch Office (ROBO) businesses • *3* is for heavy compute • *5* is for dedicated Files and Buckets • *6* is for heavy storage • *8* is for high-performance • *9* is for high performance and experimental businesses
X	Indicates the number of nodes and have one of the following values: • *1* is for single-node platforms • *2* is for multinode platforms • *3* is for multinode platforms • *4* is for multinode platforms Note: Though multinode platforms can have two, three, or four nodes, the number is always mentioned as a generic zero, *0* in the documentation.
Y	Indicates the chassis form-factor and have one of the following values: • *2* for 1U-1N (one rack unit high with one node) • *3* for 2U-2N (two rack units high with two nodes) • *5* for 2U-1N (two rack units high with one node) • *6* for 2U-4N (two rack units high with four nodes) • *7* for 1U-1N (one rack unit high with one node)
Z	Indicates the drive form-factor and have one of the following values: • *0* for 2.5 inch drives • *5* for 3.5 inch drives

Body	Description
• *G* • *C* • *S*	Indicates one of the following: • *G* at the end of the body stands for "graphics" and that the platform is optimized for using the Graphics Processing Unit (GPU) cards. For example, the NX-3155G-G7 platform. • *C* at the end of the body stands for "cold node" and that the platform is AHV-only. Such platforms should only be used for storage. Cold nodes cannot run VMs. Hence, cold nodes are the only exception where we allow mixing of the nodes with different hypervisors in the same cluster. For example, the NX-6035C and NX-6035C-G5 platforms. • *S* at the end of the body stands for "single socket" where the motherboard has only one CPU instead of the usual two that we support. For example, the NX-1175S-G7 platform.

Table 3: Suffix

Suffix	Description
G4	Indicates that the platform uses the Intel Haswell CPU
G5	Indicates that the platform uses the Intel Broadwell CPU
G6	Indicates that the platform uses the Intel Skylake CPU
G7	Indicates that the platform uses the Intel Cascade Lake CPU
G8	Indicates that the platform uses the Intel Ice Lake CPU

2

SYSTEM SPECIFICATIONS

Node Naming (NX-1065-G8)

Nutanix assigns a name to each node in a block, which varies by product type.

NX-1065-G8 platforms have one, two, three, or four nodes per block.

- Node A
- Node B
- Node C
- Node D

The following figure shows the arrangement of drives in the chassis. The first drive in each node contains the Controller VM and metadata:

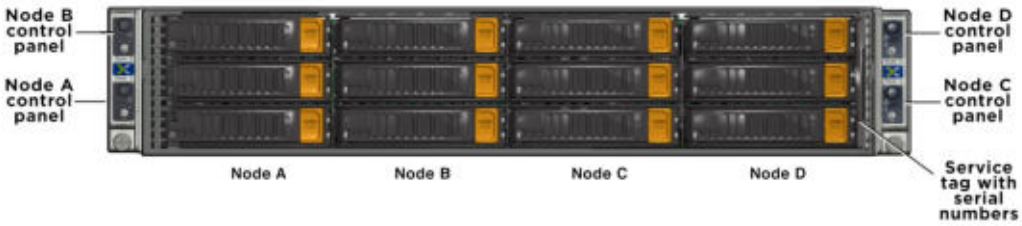

Figure 2: NX-1065-G8 Front Panel

Table 4: Supported Drive Configurations

Hybrid (SSD and HDD)	One SSD and two HDDs per node

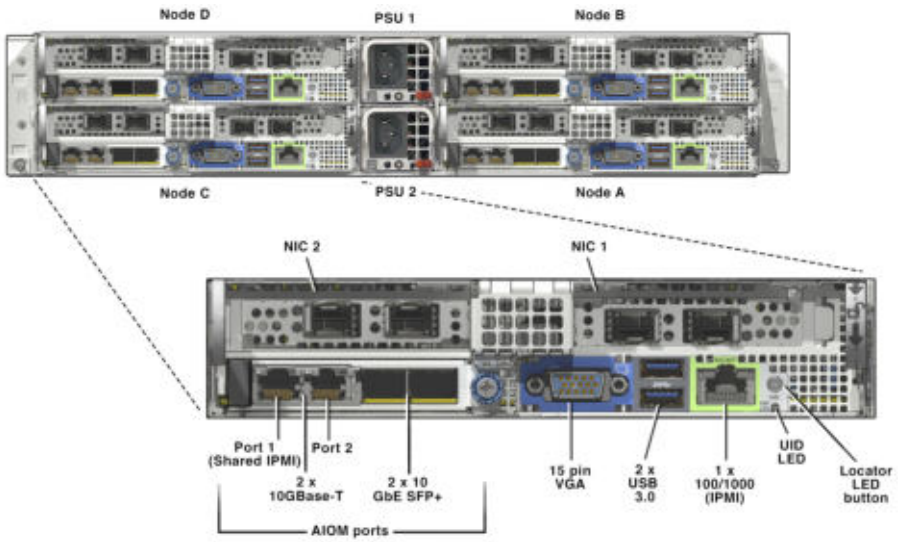

Figure 3: NX-1065-G8 Back Panel

NX-1065-G8 System Specifications

Table 5: System Characteristics

Boot Device	Dual M.2 RAID1 • 2 x 512GB M.2 Boot Device
Chassis	System Chassis • X12 2U4N LFF Chassis
CPU	Note: • 4316, 5315Y, 5317, 5320T supports a maximum operating temperature of 30C. Processor • 2 x Intel Xeon® Silver 4309Y [8 cores / 2.80 GHz] • 2 x Intel Xeon® Silver 4310 [12 cores / 2.10 GHz] • 2 x Intel Xeon® Silver 4310T [10 cores / 2.30 GHz] • 2 x Intel Xeon® Silver 4314 [16 cores / 2.40 GHz] • 2 x Intel Xeon® Silver 4316 [20 cores / 2.30 GHz] • 2 x Intel Xeon® Gold 5315Y [8 cores / 3.20 GHz] • 2 x Intel Xeon® Gold 5317 [12 cores / 3.00 GHz] • 2 x Intel Xeon® Gold 5320T [20 cores / 2.30 GHz]

Memory	32GB RDIMM • 4 x 32 GB = 128 GB • 8 x 32 GB = 256 GB • 12 x 32 GB = 384 GB • 16 x 32 GB = 512 GB 64GB DIMM • 4 x 64 GB = 256 GB • 8 x 64 GB = 512 GB • 12 x 64 GB = 768 GB • 16 x 64 GB = 1.0 TB
Network	Serverboard • 1x 1GbE Dedicated IPMI • 2x 10GbE SFP+ • 2x 10GBaseT (IPMI Failover) NICs in PCIe Slots • 0, 1 or 2 x 10GBaseT 2P NIC • 0, 1 or 2 x 10GBaseT 4P NIC • 0, 1 or 2 x 10GbE 4P NIC • 0, 1 or 2 x 25GbE 2P NIC
Network Cables	Network Cable • OPT,CBL,SFP28,1M,CU • OPT, CBL, 1M, SFP+ TO SFP+ • OPT,CBL,SFP28,3M,CU • OPT, CBL, 3M, SFP+ TO SFP+ • OPT,CBL,SFP28,5M,CU • OPT, CBL, 5M, SFP+ TO SFP+
Power Cable	Power Cable • 2 x C13/14 4ft Power Cable
Power Supply	Power Supply • 2 x 2200W Titanium PSU

Server	Server • 1, 2, 3 or 4 x NX-1065-G8 Server
Storage	
Storage : SSD +HDD	1 x SSD • 1.92TB • 3.84TB • 7.68TB 2 x HDD • 6.0TB • 8.0TB • 12.0TB • 18.0TB
Storage : SSD +HDD SED	1 x SSD • 1.92TB • 3.84TB 2 x HDD • 6.0TB • 8.0TB • 12.0TB
TPM	TPM • 0 or 1 x Unprovisioned Trusted Platform Module
Transceiver	Transceiver • SR SFP+ Transceiver
Chassis fans	4x 80 mm heavy duty fans with PWM fan speed controls

Table 6: Block, power and electrical

Block	Note: • Maximum measurements are shown. Width is from ear to ear; depth is from ears to pull rings. Weight : 44 kg Width : 449 mm Height : 88.1 mm Rack Units : 2 U Depth : 764.75 mm
Package	Weight : 56 kg
Shock	Operating : 2.5 ms Non-Operating : 10 ms
Thermal Dissipation	Typical : 5733 BTU/hr Maximum : 7644 BTU/hr
Vibration (Random)	Operating : 0.4 Grms Non-Operating : 0.98 Grms
Power consumption	3N Max Config • Maximum: 1740 VA • Typical: 1305 VA 1N Max Config • Maximum: 700 VA • Typical: 519 VA 4N Max Config • Maximum: 2250 VA • Typical: 1680 VA 2N Max Config • Maximum: 1211 VA • Typical: 908 VA

Operating environment	Note: • 4316, 5315Y, 5317, 5320T supports a maximum operating temperature of 30C. Operating temperature : 10-35C Non-Operating temperature : -45-70C Operating relative humidity : 20-90% Non-operating relative humidity : 5-95%

Certifications

- Energy Star
- CSAus
- FCC
- CSA
- UL
- cUL
- ICES
- CE
- UKCA
- KCC
- RCM
- VCCI-A
- BSMI
- EAC
- SABS
- S-MARK
- UKRSEPRO
- BIS
- SII
- RoHS
- Reach
- WEEE

3

COMPONENT SPECIFICATIONS

Controls and LEDs for Multinode Platforms

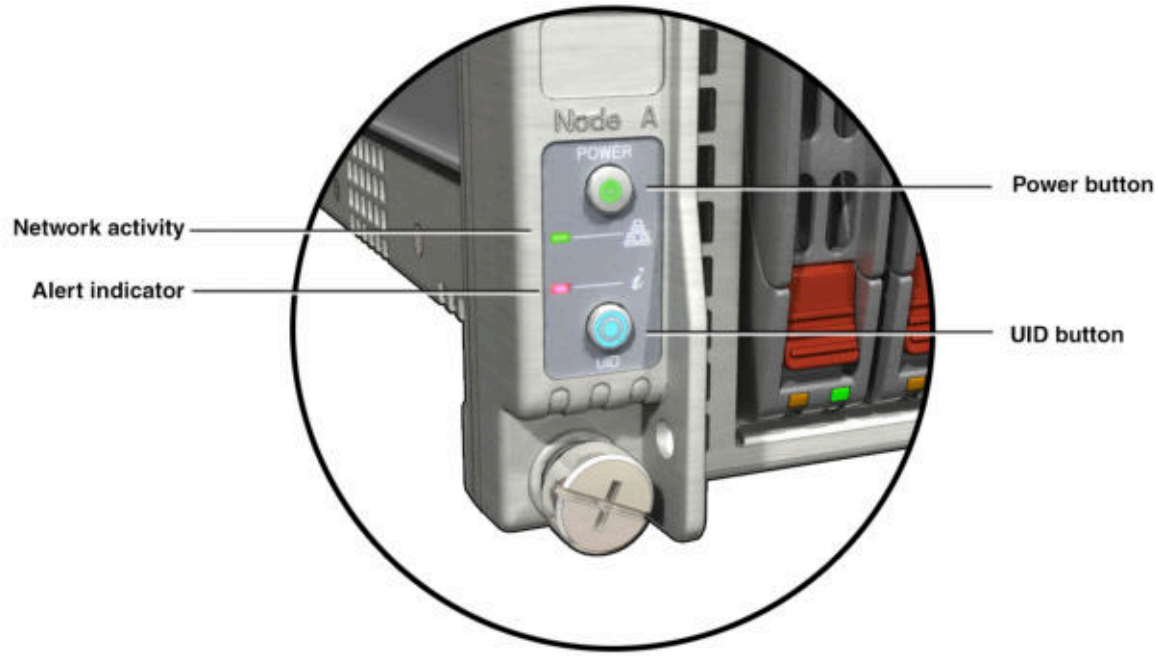

Figure 4: Front of Chassis LEDs

Table 7: LEDs on the Front of the Chassis

Name	Color	Function
Power button	Green	Power On/Off
Network activity	Green	1 GB LAN1, LAN 2 activity
Alert indicator	Red	• Solid - Overheating condition • Flashing every second - Fan failure • Flashing every 4 seconds - Power failure

Name	Color	Function
UID button	Blue	Blinking Blue - Identified (on the UID button and at the rear of the node)
HDD activity (right LED)	Blue or green	HDD/SSD activity
HDD failure (left LED)	Red	HDD/SSD failure by Nutanix software

Table 8: Drive LEDs

Top LED: Activity	Blue or green: Blinking = I/O activity, off = idle
Bottom LED: Status	Solid red = failed drive, On 5 seconds after boot = power on

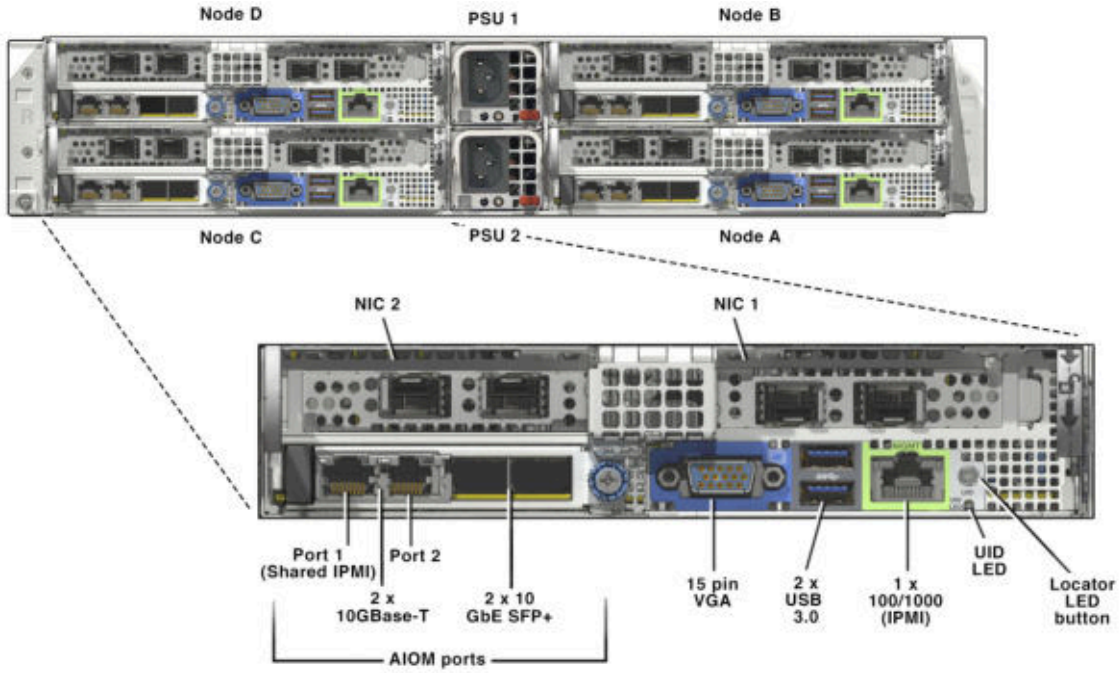

Figure 5: Back Panel LEDs of 2U4N Platforms

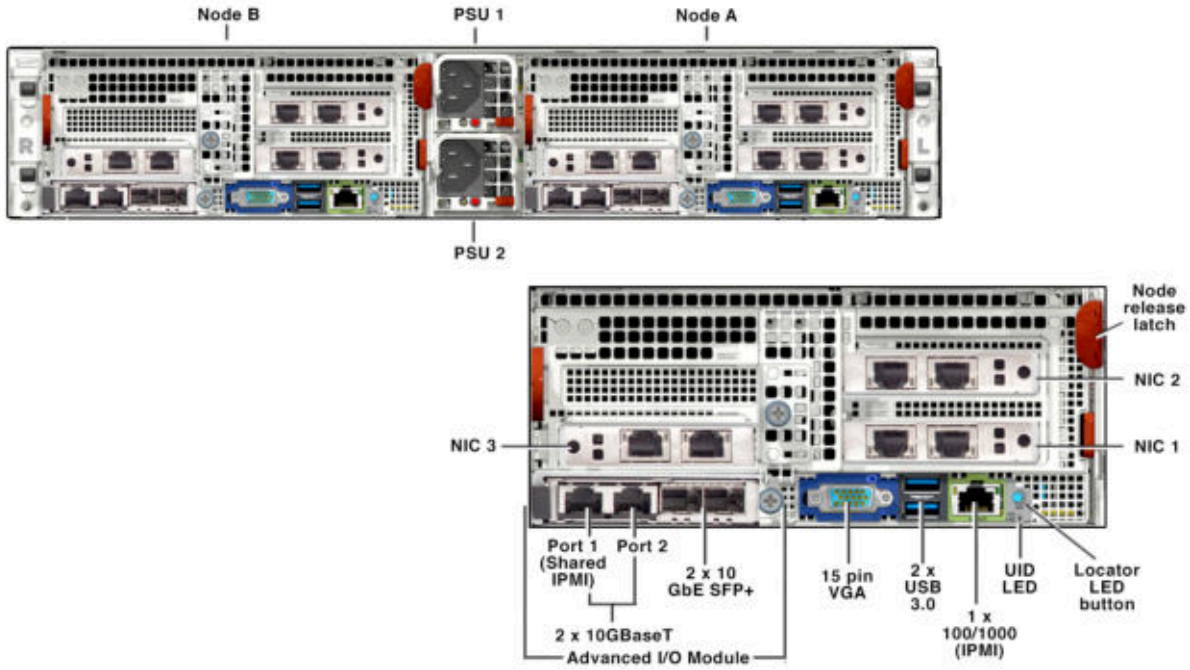

Figure 6: Back Panel LEDs of 2U2N Platform

Name	Color	Function
IPMI, left LED	Green	100 Mbps
	Amber	1 Gbps
IPMI, right LED	Yellow	Flashing - Activity
1 GbE, right LED	Off	No link or 10 Mbps
	Green	100 Mbps
	Amber	1 Gbps
1 GbE, left LED	Yellow	Flashing - Activity
10 GbE, top	On/off	Link
10 GbE, bottom	On/off	Activity
Locator LED (UID)	Blue	Blinking - Node identified

For LED states for add-on NICs, see LED Meanings for Network Cards on page 15.

LED Meanings for Network Cards

Descriptions of LEDs for supported NICs.

Different NIC manufacturers use different LED colors and blink states. Not all NICs are supported for every Nutanix platform. See the system specifications for your platform to verify which NICs are supported.

Table 9: On-Board Ports

NIC	Link (LNK) LED	Activity (ACT) LED
1 GbE dedicated IPMI	Green: 100 Mbps Yellow: 1 Gbps	Blinking yellow: activity
1 GbE shared IPMI	Green: 1 Gbps Yellow: 100 Mbps	Blinking yellow: activity
i Unit identification LED	Blinking blue: UUID has been activated.	

Table 10: SuperMicro NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 1 GbE	Green: 100 Mbps Yellow: 1 Gb/s OFF: 10 Mb/s or No Connection	Blinking yellow: activity
Dual-port 10G SFP+	Green: 10 Gb Yellow: 1 Gb	Blinking green: activity

Table 11: Silicom NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 10G SFP+	Green: all speeds	Solid green: idle Blinking green: activity
Quad-port 10G SFP+	Blue: 10 Gb Yellow: 1 Gb	Solid green: idle Blinking green: activity
Dual-port 10G BaseT	Yellow: 1 Gb/s Green: 10 Gb/s	Blinking green: activity

Table 12: Mellanox NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 10G SFP+ ConnectX-3 Pro	Green: 10 Gb speed with no traffic Blinking yellow: 10 Gb speed with traffic Not illuminated: no connection	Blinking yellow and green: activity
Dual-port 40G SFP+ ConnectX-3 Pro	Solid green: good link	Blinking yellow: activity Not illuminated: no activity
Dual-port 10G SFP28 ConnectX-4 Lx	Solid yellow: good link Blinking yellow: physical problem with link	Solid green: valid link with no traffic Blinking green: valid link with active traffic
Dual-port 25G SFP28 ConnectX-4 Lx	Solid yellow: good link Blinking yellow: physical problem with link	Solid green: valid link with no traffic Blinking green: valid link with active traffic

Power Supply Unit (PSU) Redundancy and Node Configuration

Note: Nutanix recommends that you carefully plan your AC power source needs, especially in cases where the cluster consists of mixed models.

Nutanix recommends that you use 180 V ~ 240 V AC power source to secure PSU redundancy. However, according to the following table, and depending on the number of nodes in the chassis, some NX platforms can work with redundant 100 V ~ 210 V AC power supply units.

Table 13: PSU Redundancy and Node Configuration

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208 V
NX-1065-G8	1 or 2	YES	YES
	3 or 4	NO	YES
NX-1175S-G8	1	YES	YES
NX-3060-G8	1 or 2	YES	YES
	3 or 4	NO	YES
NX-3155G-G8	1, with GPU	NO	YES
	1, without GPU	YES	YES
NX-3170-G8	1	NO	YES

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208 V
NX-8035-G8	1	YES	YES
	2	No	YES
NX-8150-G8	1	NO	YES
NX-8155-G8	1	For CPUs with a thermal design profile equal to or less than 130 W: redundant at 110 V over the entire supported ambient temperature range of 10° C to 35° C. For all other CPUs, use the following rule: • Ambient temperature is 25° C or less: YES • Ambient temperature is greater than 25° C: NO	YES
NX-8170-G8	1	NO	YES

Nutanix DMI Information

Format for Nutanix DMI strings.

VMware reads model information from the direct media interface (DMI) table.

For platforms with Intel Icelake CPUs, Nutanix provides model information to the DMI table in the following format:

`NX-motherboard_idNIC_id-HBA_id-G8`

motherboard-id has the following options:

Argument	Option
T	X12 multi-node motherboard
U	X12 single-node motherboard
W	X12 single-socket single-node motherboard

NIC_id has the following options:

Argument	Option
D1	dual-port 1G NIC
Q1	quad-port 1G NIC
DT	dual-port 10GBaseT NIC

Argument	Option
QT	quad-port 10GBaseT NIC
DS	dual-port SFP+ NIC
QS	quad-port SFP+ NIC

HBA_id specifies the number of nodes and type of HBA controller. For example:

Argument	Option
1NL3	single-node LSI3808
2NL3	2-node LSI3808
4NL3	4-node LSI3808

Table 14: Examples

DMI string	Explanation	Nutanix model
NX-TDT-4NL3-G8	X12 motherboard with dual-port 10GBase-T NIC, 4 nodes with LSI3808 HBA controllers	NX-1065-G8, NX-3060-G8
NX-UDT-1NL3-G8	X12 motherboard with dual-port 10GBase-T NIC, 1 node with LSI3808 HBA controller	NX-3155G-G8, NX-3170-G8, NX-8150-G8, NX-8155-G8, NX-8170-G8

4

MEMORY CONFIGURATIONS

Supported Memory Configurations

DIMM installation information for all Nutanix G8 platforms.

DIMM Restrictions

DIMM type

Each G8 node must contain only DIMMs of the same type. So, for example, you cannot mix RDIMM and LRDIMM in the same node.

DIMM capacity

Each G8 node must contain only DIMMs of the same capacity. So, for example, you cannot mix 32 GB DIMMs and 64 GB DIMMs in the same node.

DIMM speed

G8 nodes ship with 3200 MHz DIMMs. 3200 MHz is the highest speed Nutanix currently supports, so you cannot currently mix DIMM speeds in any G8 node.

DIMM manufacturer

You can mix DIMMs from different manufacturers in the same *node*, but not in the same *channel*.

Multi-node platforms

Multi-node G8 platforms contain only one active DIMM slot per channel, so mixing DIMMs in the same channel is not possible.

Single-socket platforms

The single-socket NX-1175S-G8 platform contains only one DIMM slot per channel, so mixing DIMMs in the same channel is not possible.

Single-node platforms

- Single-node G8 platforms contain two DIMM slots per channel. Within a channel, all DIMMs must be from the same manufacturer.
- When replacing a failed DIMM, if there are two DIMMs in the channel, either replace the failed DIMM with a new DIMM form the same manufacturer, or else replace both DIMMs in the channel and make sure that both new DIMMs are from the same manufacturer.
- When adding new DIMMs to a node, if the new DIMMs and the original DIMMs have different manufacturer part numbers, arrange the DIMMs so that the original DIMMs and the new DIMMs are not mixed in the same channel.
    - EXAMPLE: You have an NX-8155-G8 node that has sixteen 32GB DIMMs for a total of 512 GB. You decide to upgrade to thirty-two 32GB DIMMs for a total of 1024 GB. When you remove the node from the chassis and look at the server board, you see that each CPU has eight DIMMs. Remove all DIMMs from one CPU and place them in the empty DIMM slots for the other CPU.

Then place all the new DIMMs in the DIMM slots for the first CPU, filling all slots. This way you can ensure that the original DIMMs and the new DIMMs do not share channels.

Note: You do not need to balance numbers of DIMMs from different manufacturers within a node, so long as you never mix them in the same channel.

Memory Installation Order for Multi-Node G8 Platforms

A memory channel is a group of DIMM slots.

For G8 multi-node platforms, each CPU is associated with eight active memory channels that contain one blue slot each, plus two inactive black slots.

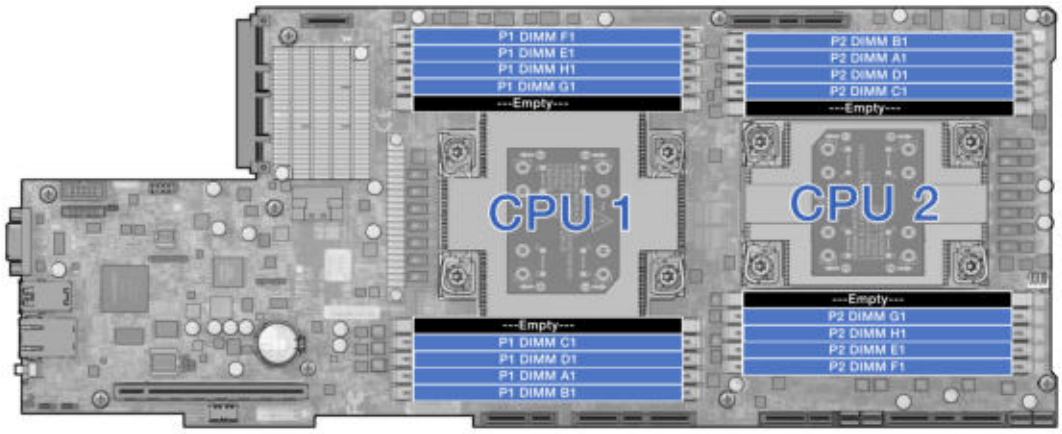

Note: The black slots (C2 and G2 on each CPU) are inactive.

Figure 7: DIMM Slots for a G8 Multinode Serverboard

Table 15: DIMM Installation Order for Multi-node G8 Platforms

Number of DIMMs	Slots to use	Supported capacities (NX-1065-G8)	Supported capacities (NX-3060-G8 and NX-8035-G8)
4	CPU1: A1, E1 CPU2: A1, E1	32 GB, 64 GB	32 GB, 64 GB, 128 GB
8	CPU1: A1, C1, E1, G1 CPU2: A1, C1, E1, G1	32 GB, 64 GB	32 GB, 64 GB, 128 GB
12	CPU1: A1, B1, C1, E1, F1, G1 CPU2: A1, B1, C1, E1, F1, G1	32 GB, 64 GB	32 GB, 64 GB, 128 GB

Number of DIMMs	Slots to use	Supported capacities (NX-1065-G8)	Supported capacities (NX-3060-G8 and NX-8035-G8)
16	Fill all blue slots.	32 GB, 64 GB	32 GB, 64 GB

Memory Installation Order for Single-node G8 Platforms

A memory channel is a group of DIMM slots.

For G8 single-node platforms, each CPU is associated with eight memory channels that contain one blue slot and one black slot each, for a total of 32 DIMM slots.

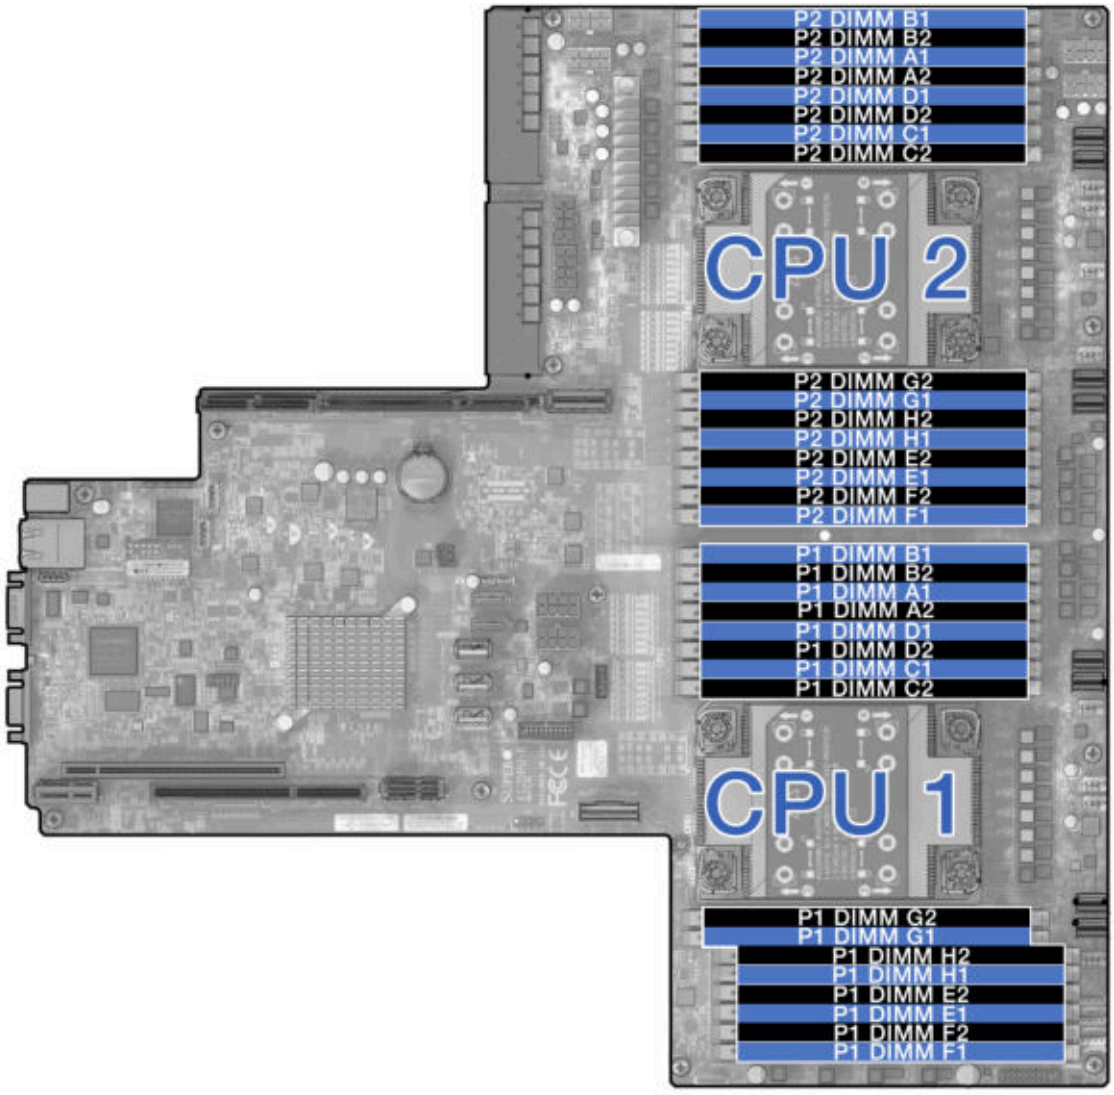

Figure 8: DIMM Slots for a G8 Single-Node Serverboard

Table 16: DIMM Installation Order for Single-Node G8 Platforms

Number of DIMMs	Slots to use	Supported capacities
4	CPU1: A1, E1 (blue slots) CPU2: A1, E1 (blue slots)	32 GB, 64 GB, 128 GB
8	CPU1: A1, C1, E1, G1 (blue slots) CPU2: A1, C1, E1, G1 (blue slots)	32 GB, 64 GB, 128 GB
12	CPU1: A1, B1, C1, E1, F1, G1 (blue slots) CPU2: A1, B1, C1, E1, F1, G1 (blue slots)	32 GB, 64 GB, 128 GB
16	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots)	32 GB, 64 GB, 128 GB
24	CPU1: A1, B1, C1, E1, F1, G1 (blue slots) CPU1: A2, B2, C2, E2, F2, G2 (black slots) CPU2: A1, B1, C1, E1, F1, G1 (blue slots) CPU2: A2, B2, C2, E2, F2, G2 (black slots)	32 GB, 64 GB, 128 GB
32	Fill all slots.	32 GB, 64 GB, 128 GB

Memory Installation Order for the Single-socket NX-1175S-G8 Platform

A memory channel is a group of DIMM slots.

On the NX-1175S-G8, the single CPU is associated with eight memory channels that contain one slot each.

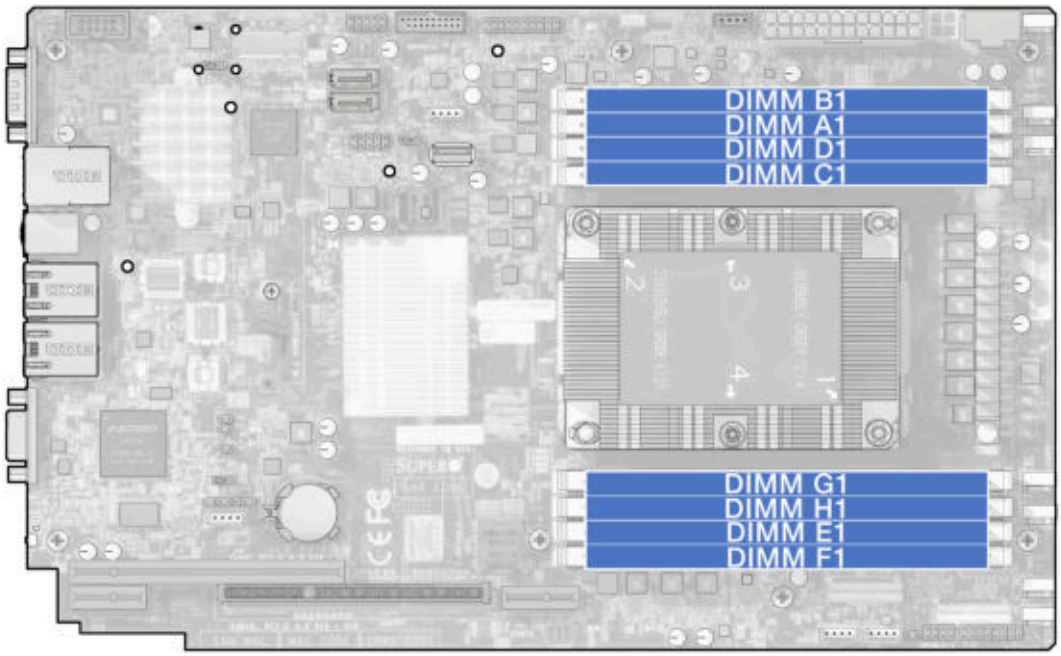

Figure 9: DIMM slots for an NX-1175S-G8 server board

Table 17: DIMM Installation Order for NX-1175S-G8

Number of DIMMs	Slots to use	Supported DIMM capacities
2	A1, E1	32 GB, 64 GB, 128 GB
4	A1, C1, E1, G1	32 GB, 64 GB, 128 GB
8	Fill all slots.	32 GB, 64 GB, 128 GB

COPYRIGHT

Copyright 2022 Nutanix, Inc.

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

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

License

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

Conventions

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

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
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

Interface	Target	Username	Password
SSH client or console	Acropolis OpenStack Services VM (Nutanix OVM)	root	admin

Version

Last modified: March 4, 2022 (2022-03-04T10:22:59+05:30)

Platform ANY

NX-8035-G8 System Specifications

NX-8035-G8

March 4, 2022

Contents

1. Nutanix Hardware Naming Convention....3

2. System Specifications....7
Node Naming (NX-8035-G8)....7
NX-8035-G8 System Specifications....8

3. Component Specifications....15
Controls and LEDs for Multinode Platforms....15
LED Meanings for Network Cards....17
Power Supply Unit (PSU) Redundancy and Node Configuration....19
Nutanix DMI Information....20

4. Memory Configurations....22
Supported Memory Configurations....22

Copyright....27
License....27
Conventions....27
Default Cluster Credentials....27
Version....28

1

NUTANIX HARDWARE NAMING CONVENTION

Every Nutanix block has a unique name based on the standardized naming convention followed by Nutanix.

The Nutanix hardware model name translates to the following naming convention:

Prefix-Body-Suffix or *NX-WXYZ(G|C|S)-G(4|5|6|7|8)*

Where:

- *NX* is the Nutanix appliance
- *W* is the product series
- *X* is the number of nodes
- *Y* is the chassis form-factor
- *Z* is the disk form-factor
- *G|C|S* indicates that a platform uses GPUs, a platform is an AHV-only node, or a platform supports only one CPU respectively
- *G(4|5|6|7|8)* is the CPU generation

Example: *NX-3460-G8*

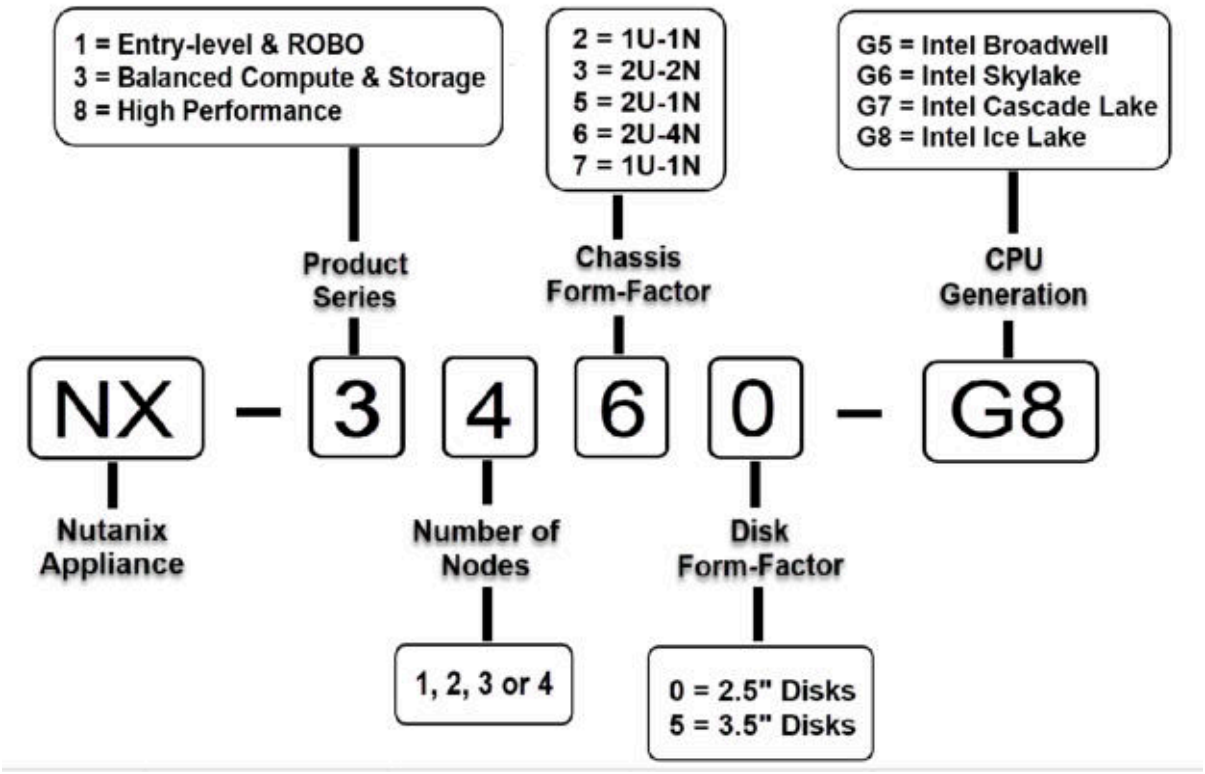

Figure 1: Nutanix Hardware Naming Convention

The following section provides a detailed description of the Nutanix hardware naming convention:

Table 1: Prefix

Prefix	Description
NX	Indicates that the platform is sold directly by Nutanix and support calls are handled by Nutanix. *NX* stands for *Nutanix*.

Table 2: Body

Body	Description
W	Indicates the product series and have one of the following values: • *1* is for small or Remote Office/Branch Office (ROBO) businesses • *3* is for heavy compute • *5* is for dedicated Files and Buckets • *6* is for heavy storage • *8* is for high-performance • *9* is for high performance and experimental businesses
X	Indicates the number of nodes and have one of the following values: • *1* is for single-node platforms • *2* is for multinode platforms • *3* is for multinode platforms • *4* is for multinode platforms Note: Though multinode platforms can have two, three, or four nodes, the number is always mentioned as a generic zero, *0* in the documentation.
Y	Indicates the chassis form-factor and have one of the following values: • *2* for 1U-1N (one rack unit high with one node) • *3* for 2U-2N (two rack units high with two nodes) • *5* for 2U-1N (two rack units high with one node) • *6* for 2U-4N (two rack units high with four nodes) • *7* for 1U-1N (one rack unit high with one node)
Z	Indicates the drive form-factor and have one of the following values: • *0* for 2.5 inch drives • *5* for 3.5 inch drives

Body	Description
• *G* • *C* • *S*	Indicates one of the following: • *G* at the end of the body stands for "graphics" and that the platform is optimized for using the Graphics Processing Unit (GPU) cards. For example, the NX-3155G-G7 platform. • *C* at the end of the body stands for "cold node" and that the platform is AHV-only. Such platforms should only be used for storage. Cold nodes cannot run VMs. Hence, cold nodes are the only exception where we allow mixing of the nodes with different hypervisors in the same cluster. For example, the NX-6035C and NX-6035C-G5 platforms. • *S* at the end of the body stands for "single socket" where the motherboard has only one CPU instead of the usual two that we support. For example, the NX-1175S-G7 platform.

Table 3: Suffix

Suffix	Description
G4	Indicates that the platform uses the Intel Haswell CPU
G5	Indicates that the platform uses the Intel Broadwell CPU
G6	Indicates that the platform uses the Intel Skylake CPU
G7	Indicates that the platform uses the Intel Cascade Lake CPU
G8	Indicates that the platform uses the Intel Ice Lake CPU

2

SYSTEM SPECIFICATIONS

Node Naming (NX-8035-G8)

Nutanix assigns a name to each node in a block, which varies by product type.

For the NX-8035-G8, the node names are:

- Node A
- Node B

Node A physical drives are located in the left side of the chassis while Node B drives are located in the right side. The SSDs contain the Controller VM and metadata.

The NX-8035-G8 supports hard disk drives (HDD), solid-state drives (SSD), and non-volatile memory express drives (NVMe).

The NX-8035-G8 supports partial population of drive slots with SSDs (not HDDs.) Any configuration that includes NVMe drives must also include four SSDs, so partial population is not supported for configurations including NVMe drives.

> Note: If you order a hybrid SSD and HDD configuration, you can later convert the platform to an all-SSD configuration. However, if you begin with either a hybrid SSD and HDD or an all-SSD configuration, you cannot later convert to an SSD with NVMe configuration.

Figure 2: Front Panel of NX-8035-G8 Block

> Note: NVMe drives, if used, must be placed in the two drive slots on the lower left of each node.

Table 4: Supported drive configurations

Hybrid SSD and HDD	Two SSDs and four HDDs per node
All-SSD Note: When partially populating a node with SSDs, you must load the drive slots in order from left to right.	Two SSDs and four empty slots per node
	Four SSDs and two empty slots per node
	Six SSDs per node
SSD with NVMe	Two NVMe drives and four SSDs per node Note: The NVMe drives must go in the two bottom drive slots in the left column in each node.

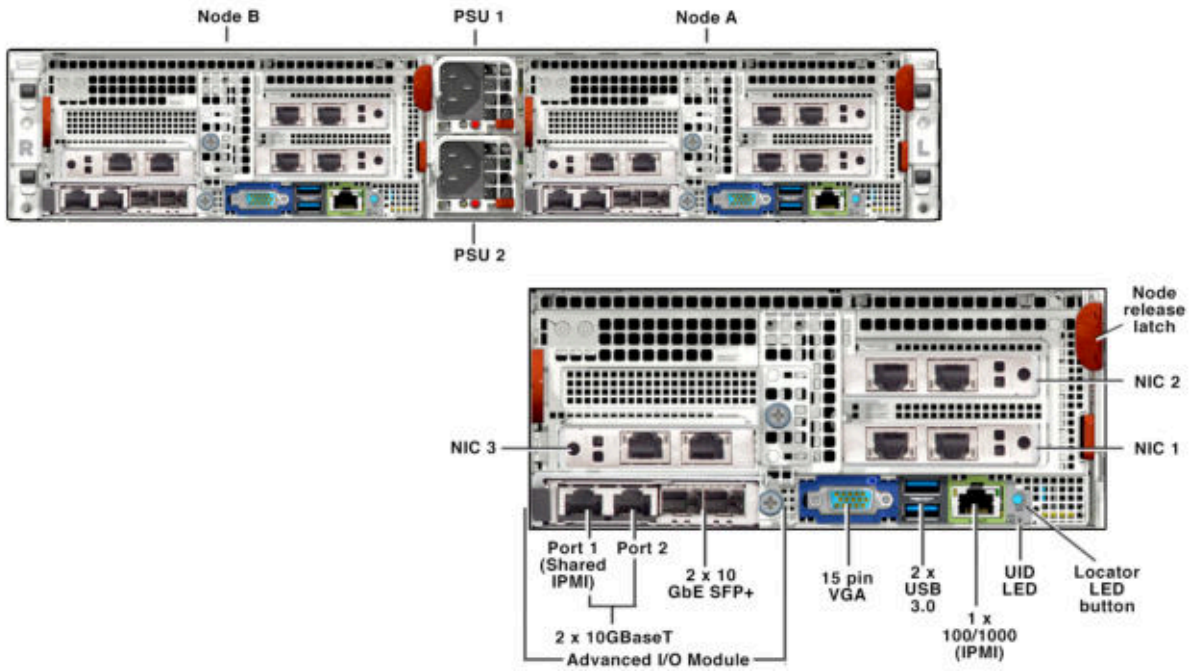

Figure 3: NX-8035-G8 Back Panel

NX-8035-G8 System Specifications

Table 5: System Characteristics

Boot Device	Dual M.2 RAID1 • 2 x 512GB M.2 Boot Device
Chassis	System Chassis • X12 2U2N LFF Chassis

CPU	Processor • 2 x Intel Xeon® Silver 4309Y [8 cores / 2.80 GHz] • 2 x Intel Xeon® Silver 4310 [12 cores / 2.10 GHz] • 2 x Intel Xeon® Silver 4310T [10 cores / 2.30 GHz] • 2 x Intel Xeon® Silver 4314 [16 cores / 2.40 GHz] • 2 x Intel Xeon® Silver 4316 [20 cores / 2.30 GHz] • 2 x Intel Xeon® Gold 5315Y [8 cores / 3.20 GHz] • 2 x Intel Xeon® Gold 5317 [12 cores / 3.00 GHz] • 2 x Intel Xeon® Gold 5318Y [24 cores / 2.10 GHz] • 2 x Intel Xeon® Gold 5320 [26 cores / 2.20 GHz] • 2 x Intel Xeon® Gold 5320T [20 cores / 2.30 GHz] • 2 x Intel Xeon® Gold 6326 [16 cores / 2.90 GHz] • 2 x Intel Xeon® Gold 6330 [28 cores / 2.00 GHz] • 2 x Intel Xeon® Gold 6334 [8 cores / 3.60 GHz] • 2 x Intel Xeon® Gold 6336Y [24 cores / 2.40 GHz] • 2 x Intel Xeon® Gold 6338 [32 cores / 2.00 GHz] • 2 x Intel Xeon® Gold 6342 [24 cores / 2.80 GHz] • 2 x Intel Xeon® Gold 6346 [16 cores / 3.10 GHz] • 2 x Intel Xeon® Gold 6348 [28 cores / 2.60 GHz] • 2 x Intel Xeon® Gold 6354 [18 cores / 3.00 GHz]

Memory	32GB RDIMM • 4 x 32 GB = 128 GB • 8 x 32 GB = 256 GB • 12 x 32 GB = 384 GB • 16 x 32 GB = 512 GB 64GB DIMM • 4 x 64 GB = 256 GB • 8 x 64 GB = 512 GB • 12 x 64 GB = 768 GB • 16 x 64 GB = 1.0 TB 128GB RDIMM • 4 x 128 GB = 512 GB • 8 x 128 GB = 1.0 TB • 12 x 128 GB = 1.5 TB • 16 x 128 GB = 2.0 TB
Network	Serverboard • 1x 1GbE Dedicated IPMI • 2x 10GbE SFP+ • 2x 10GBaseT (IPMI Failover) NICs in PCIe Slots • 0, 1, 2 or 3 x 10GBaseT 2P NIC • 0, 1, 2 or 3 x 10GBaseT 4P NIC • 0, 1, 2 or 3 x 10GbE 4P NIC • 0, 1, 2 or 3 x 25GbE 2P NIC
Network Cables	Network Cable • OPT,CBL,SFP28,1M,CU • OPT, CBL, 1M, SFP+ TO SFP+ • OPT,CBL,SFP28,3M,CU • OPT, CBL, 3M, SFP+ TO SFP+ • OPT,CBL,SFP28,5M,CU • OPT, CBL, 5M, SFP+ TO SFP+

Power Cable	Power Cable • 2 x C13/14 4ft Power Cable
Power Supply	Power Supply • 2 x 2200W Titanium PSU
Server	Server • 1 or 2 x NX-8035-G8 Server
Storage	Note: • When using all-SSD storage you must use only two, four, or six drives. Configurations of three or five drives are not supported.
Storage : NVMe +SSD	2 x NVMe • 1.92TB • 3.84TB • 7.68TB 4 x SSD • 1.92TB • 3.84TB • 7.68TB
Storage : All SSD	2, 4 or 6 x SSD • 1.92TB • 3.84TB • 7.68TB
Storage : All SSD SED	2, 4 or 6 x SSD • 1.92TB • 3.84TB

Storage : SSD +HDD	2 x SSD • 1.92TB • 3.84TB • 7.68TB 4 x HDD • 6.0TB • 8.0TB • 12.0TB • 18.0TB
Storage : SSD +HDD SED	2 x SSD • 1.92TB • 3.84TB 4 x HDD • 6.0TB • 8.0TB • 12.0TB
TPM	TPM • 0 or 1 x Unprovisioned Trusted Platform Module
Transceiver	Transceiver • SR SFP+ Transceiver
Chassis fans	4x 80 mm heavy duty fans with PWM fan speed controls

Table 6: Block, power and electrical

Block	Weight : 36.6 kg Height : 88.1 mm Width : 449 mm Rack Units : 2 U Depth : 764.75 mm
Package	Weight : 49 kg
Shock	Non-Operating : 10 ms Operating : 2.5 ms

Thermal Dissipation	Typical : 5118 BTU/hr Maximum : 6824 BTU/hr
Vibration (Random)	Operating : 0.4 Grms Non-Operating : 0.98 Grms
Power consumption	1N Max Config • Maximum: 1000 VA • Typical: 750 VA 2N Max Config • Maximum: 2000 VA • Typical: 1500 VA
Operating environment	Operating temperature : 10-35C Non-Operating temperature : -45-70C Operating relative humidity : 20-90% Non-operating relative humidity : 5-95%

Certifications

- Energy Star
- CSAus
- FCC
- CSA
- UL
- cUL
- ICES
- CE
- UKCA
- KCC
- RCM
- VCCI-A
- BSMI
- EAC
- SABS
- S-MARK
- UKRSEPRO
- BIS

- SII
- RoHS
- Reach
- WEEE

3

COMPONENT SPECIFICATIONS

Controls and LEDs for Multinode Platforms

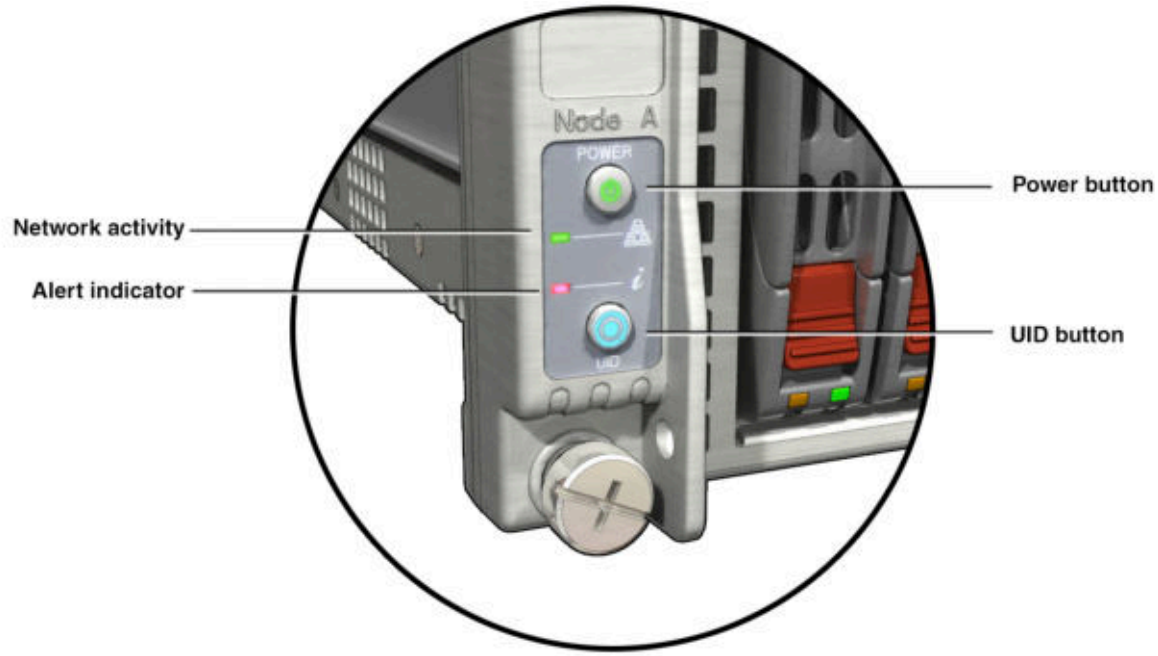

Figure 4: Front of Chassis LEDs

Table 7: LEDs on the Front of the Chassis

Name	Color	Function
Power button	Green	Power On/Off
Network activity	Green	1 GB LAN1, LAN 2 activity
Alert indicator	Red	• Solid - Overheating condition • Flashing every second - Fan failure • Flashing every 4 seconds - Power failure

Name	Color	Function
UID button	Blue	Blinking Blue - Identified (on the UID button and at the rear of the node)
HDD activity (right LED)	Blue or green	HDD/SSD activity
HDD failure (left LED)	Red	HDD/SSD failure by Nutanix software

Table 8: Drive LEDs

Top LED: Activity	Blue or green: Blinking = I/O activity, off = idle
Bottom LED: Status	Solid red = failed drive, On 5 seconds after boot = power on

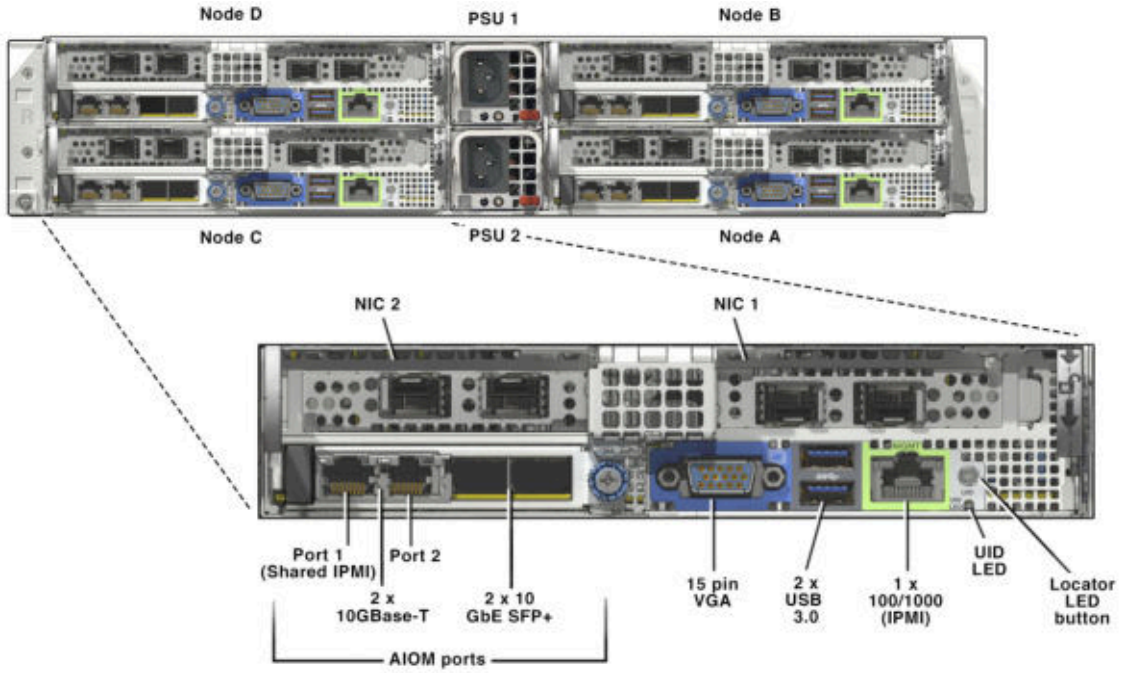

Figure 5: Back Panel LEDs of 2U4N Platforms

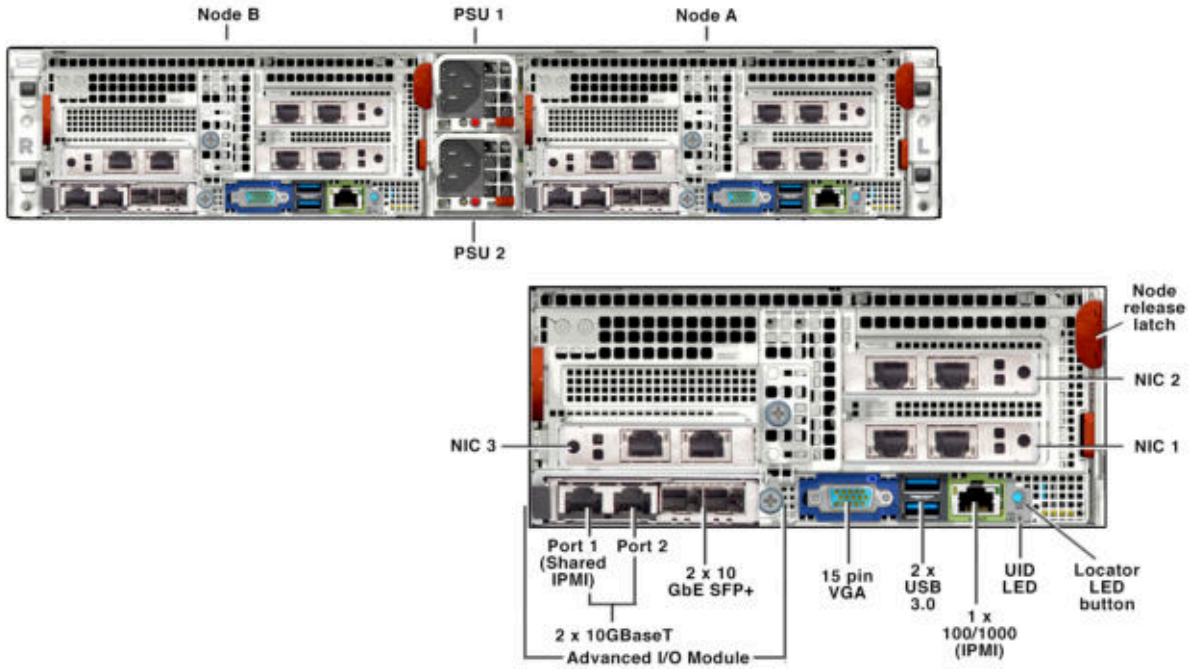

Figure 6: Back Panel LEDs of 2U2N Platform

Name	Color	Function
IPMI, left LED	Green	100 Mbps
	Amber	1 Gbps
IPMI, right LED	Yellow	Flashing - Activity
1 GbE, right LED	Off	No link or 10 Mbps
	Green	100 Mbps
	Amber	1 Gbps
1 GbE, left LED	Yellow	Flashing - Activity
10 GbE, top	On/off	Link
10 GbE, bottom	On/off	Activity
Locator LED (UID)	Blue	Blinking - Node identified

For LED states for add-on NICs, see LED Meanings for Network Cards on page 17.

LED Meanings for Network Cards

Descriptions of LEDs for supported NICs.

Different NIC manufacturers use different LED colors and blink states. Not all NICs are supported for every Nutanix platform. See the system specifications for your platform to verify which NICs are supported.

Table 9: On-Board Ports

NIC	Link (LNK) LED	Activity (ACT) LED
1 GbE dedicated IPMI	Green: 100 Mbps Yellow: 1 Gbps	Blinking yellow: activity
1 GbE shared IPMI	Green: 1 Gbps Yellow: 100 Mbps	Blinking yellow: activity
i Unit identification LED	Blinking blue: UUID has been activated.	

Table 10: SuperMicro NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 1 GbE	Green: 100 Mbps Yellow: 1 Gb/s OFF: 10 Mb/s or No Connection	Blinking yellow: activity
Dual-port 10G SFP+	Green: 10 Gb Yellow: 1 Gb	Blinking green: activity

Table 11: Silicom NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 10G SFP+	Green: all speeds	Solid green: idle Blinking green: activity
Quad-port 10G SFP+	Blue: 10 Gb Yellow: 1 Gb	Solid green: idle Blinking green: activity
Dual-port 10G BaseT	Yellow: 1 Gb/s Green: 10 Gb/s	Blinking green: activity

Table 12: Mellanox NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 10G SFP+ ConnectX-3 Pro	Green: 10 Gb speed with no traffic Blinking yellow: 10 Gb speed with traffic Not illuminated: no connection	Blinking yellow and green: activity
Dual-port 40G SFP+ ConnectX-3 Pro	Solid green: good link	Blinking yellow: activity Not illuminated: no activity
Dual-port 10G SFP28 ConnectX-4 Lx	Solid yellow: good link Blinking yellow: physical problem with link	Solid green: valid link with no traffic Blinking green: valid link with active traffic
Dual-port 25G SFP28 ConnectX-4 Lx	Solid yellow: good link Blinking yellow: physical problem with link	Solid green: valid link with no traffic Blinking green: valid link with active traffic

Power Supply Unit (PSU) Redundancy and Node Configuration

Note: Nutanix recommends that you carefully plan your AC power source needs, especially in cases where the cluster consists of mixed models.

Nutanix recommends that you use 180 V ~ 240 V AC power source to secure PSU redundancy. However, according to the following table, and depending on the number of nodes in the chassis, some NX platforms can work with redundant 100 V ~ 210 V AC power supply units.

Table 13: PSU Redundancy and Node Configuration

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208 V
NX-1065-G8	1 or 2	YES	YES
	3 or 4	NO	YES
NX-1175S-G8	1	YES	YES
NX-3060-G8	1 or 2	YES	YES
	3 or 4	NO	YES
NX-3155G-G8	1, with GPU	NO	YES
	1, without GPU	YES	YES
NX-3170-G8	1	NO	YES

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208 V
NX-8035-G8	1	YES	YES
	2	No	YES
NX-8150-G8	1	NO	YES
NX-8155-G8	1	For CPUs with a thermal design profile equal to or less than 130 W: redundant at 110 V over the entire supported ambient temperature range of 10° C to 35° C. For all other CPUs, use the following rule: • Ambient temperature is 25° C or less: YES • Ambient temperature is greater than 25° C: NO	YES
NX-8170-G8	1	NO	YES

Nutanix DMI Information

Format for Nutanix DMI strings.

VMware reads model information from the direct media interface (DMI) table.

For platforms with Intel Icelake CPUs, Nutanix provides model information to the DMI table in the following format:

`NX-motherboard_idNIC_id-HBA_id-G8`

motherboard-id has the following options:

Argument	Option
T	X12 multi-node motherboard
U	X12 single-node motherboard
W	X12 single-socket single-node motherboard

NIC_id has the following options:

Argument	Option
D1	dual-port 1G NIC
Q1	quad-port 1G NIC
DT	dual-port 10GBaseT NIC

Argument	Option
QT	quad-port 10GBaseT NIC
DS	dual-port SFP+ NIC
QS	quad-port SFP+ NIC

HBA_id specifies the number of nodes and type of HBA controller. For example:

Argument	Option
1NL3	single-node LSI3808
2NL3	2-node LSI3808
4NL3	4-node LSI3808

Table 14: Examples

DMI string	Explanation	Nutanix model
NX-TDT-4NL3-G8	X12 motherboard with dual-port 10GBase-T NIC, 4 nodes with LSI3808 HBA controllers	NX-1065-G8, NX-3060-G8
NX-UDT-1NL3-G8	X12 motherboard with dual-port 10GBase-T NIC, 1 node with LSI3808 HBA controller	NX-3155G-G8, NX-3170-G8, NX-8150-G8, NX-8155-G8, NX-8170-G8

4

MEMORY CONFIGURATIONS

Supported Memory Configurations

DIMM installation information for all Nutanix G8 platforms.

DIMM Restrictions

DIMM type

Each G8 node must contain only DIMMs of the same type. So, for example, you cannot mix RDIMM and LRDIMM in the same node.

DIMM capacity

Each G8 node must contain only DIMMs of the same capacity. So, for example, you cannot mix 32 GB DIMMs and 64 GB DIMMs in the same node.

DIMM speed

G8 nodes ship with 3200 MHz DIMMs. 3200 MHz is the highest speed Nutanix currently supports, so you cannot currently mix DIMM speeds in any G8 node.

DIMM manufacturer

You can mix DIMMs from different manufacturers in the same *node*, but not in the same *channel*.

Multi-node platforms

Multi-node G8 platforms contain only one active DIMM slot per channel, so mixing DIMMs in the same channel is not possible.

Single-socket platforms

The single-socket NX-1175S-G8 platform contains only one DIMM slot per channel, so mixing DIMMs in the same channel is not possible.

Single-node platforms

- Single-node G8 platforms contain two DIMM slots per channel. Within a channel, all DIMMs must be from the same manufacturer.
- When replacing a failed DIMM, if there are two DIMMs in the channel, either replace the failed DIMM with a new DIMM form the same manufacturer, or else replace both DIMMs in the channel and make sure that both new DIMMs are from the same manufacturer.
- When adding new DIMMs to a node, if the new DIMMs and the original DIMMs have different manufacturer part numbers, arrange the DIMMs so that the original DIMMs and the new DIMMs are not mixed in the same channel.
    - EXAMPLE: You have an NX-8155-G8 node that has sixteen 32GB DIMMs for a total of 512 GB. You decide to upgrade to thirty-two 32GB DIMMs for a total of 1024 GB. When you remove the node from the chassis and look at the server board, you see that each CPU has eight DIMMs. Remove all DIMMs from one CPU and place them in the empty DIMM slots for the other CPU.

Then place all the new DIMMs in the DIMM slots for the first CPU, filling all slots. This way you can ensure that the original DIMMs and the new DIMMs do not share channels.

Note: You do not need to balance numbers of DIMMs from different manufacturers within a node, so long as you never mix them in the same channel.

Memory Installation Order for Multi-Node G8 Platforms

A memory channel is a group of DIMM slots.

For G8 multi-node platforms, each CPU is associated with eight active memory channels that contain one blue slot each, plus two inactive black slots.

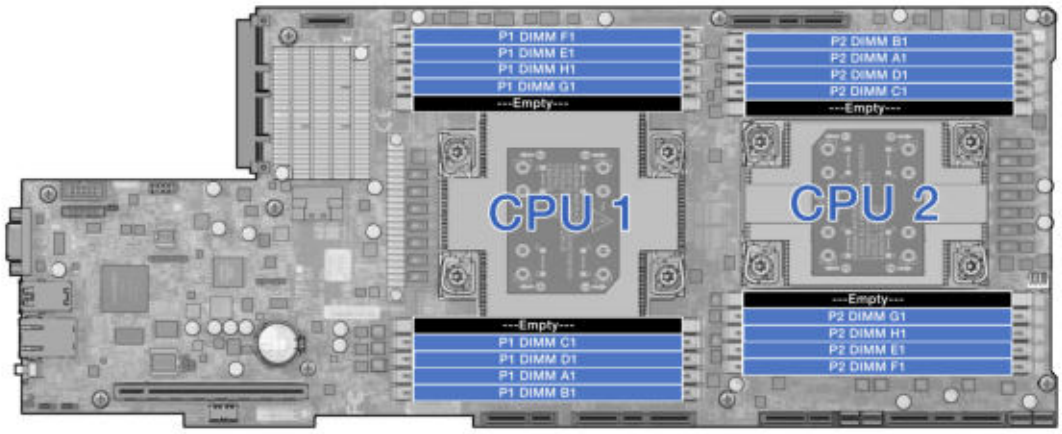

Note: The black slots (C2 and G2 on each CPU) are inactive.

Figure 7: DIMM Slots for a G8 Multinode Serverboard

Table 15: DIMM Installation Order for Multi-node G8 Platforms

Number of DIMMs	Slots to use	Supported capacities (NX-1065-G8)	Supported capacities (NX-3060-G8 and NX-8035-G8)
4	CPU1: A1, E1 CPU2: A1, E1	32 GB, 64 GB	32 GB, 64 GB, 128 GB
8	CPU1: A1, C1, E1, G1 CPU2: A1, C1, E1, G1	32 GB, 64 GB	32 GB, 64 GB, 128 GB
12	CPU1: A1, B1, C1, E1, F1, G1 CPU2: A1, B1, C1, E1, F1, G1	32 GB, 64 GB	32 GB, 64 GB, 128 GB

Number of DIMMs	Slots to use	Supported capacities (NX-1065-G8)	Supported capacities (NX-3060-G8 and NX-8035-G8)
16	Fill all blue slots.	32 GB, 64 GB	32 GB, 64 GB

Memory Installation Order for Single-node G8 Platforms

A memory channel is a group of DIMM slots.

For G8 single-node platforms, each CPU is associated with eight memory channels that contain one blue slot and one black slot each, for a total of 32 DIMM slots.

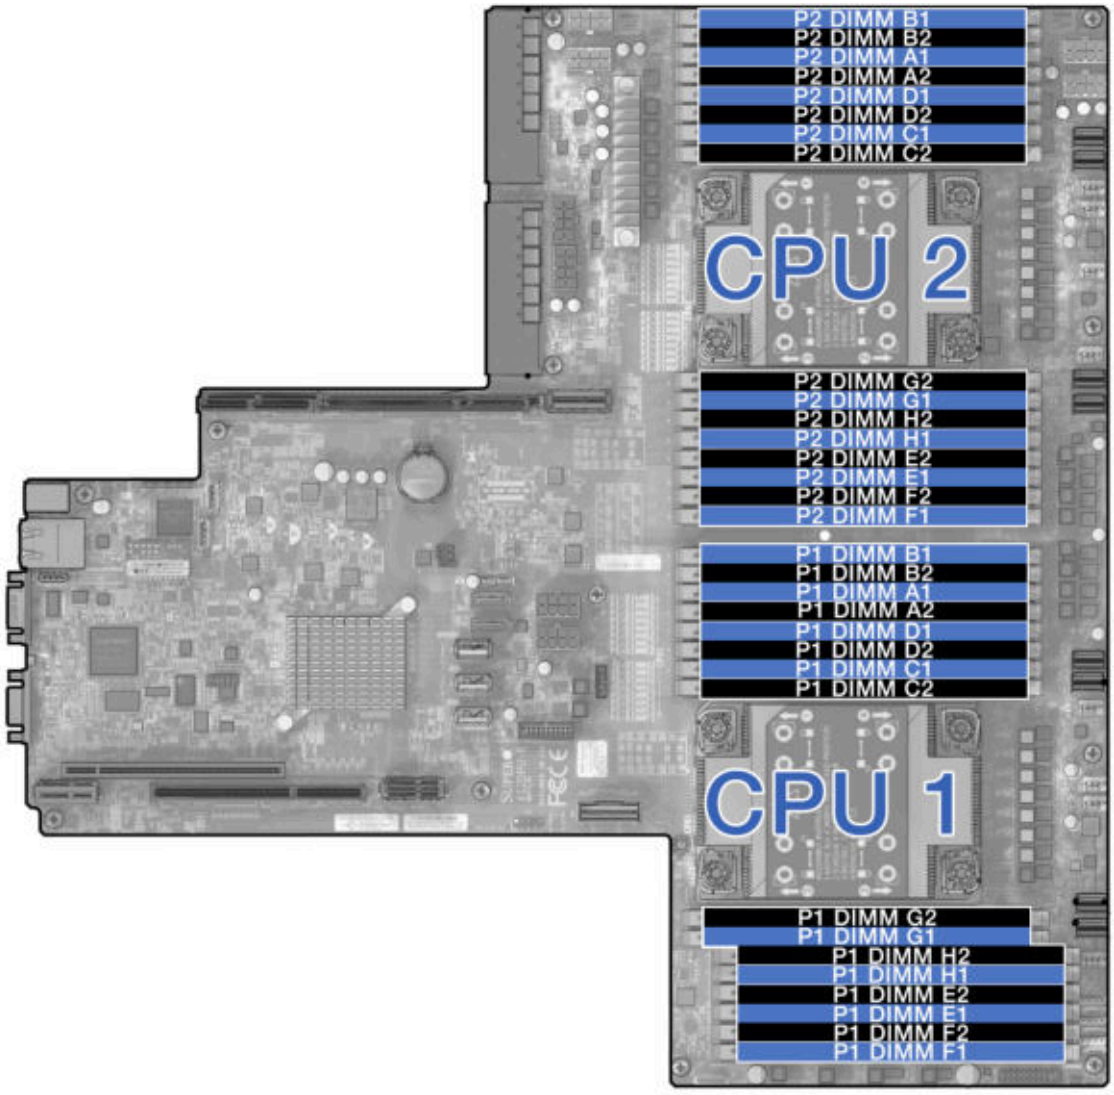

Figure 8: DIMM Slots for a G8 Single-Node Serverboard

Table 16: DIMM Installation Order for Single-Node G8 Platforms

Number of DIMMs	Slots to use	Supported capacities
4	CPU1: A1, E1 (blue slots) CPU2: A1, E1 (blue slots)	32 GB, 64 GB, 128 GB
8	CPU1: A1, C1, E1, G1 (blue slots) CPU2: A1, C1, E1, G1 (blue slots)	32 GB, 64 GB, 128 GB
12	CPU1: A1, B1, C1, E1, F1, G1 (blue slots) CPU2: A1, B1, C1, E1, F1, G1 (blue slots)	32 GB, 64 GB, 128 GB
16	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots)	32 GB, 64 GB, 128 GB
24	CPU1: A1, B1, C1, E1, F1, G1 (blue slots) CPU1: A2, B2, C2, E2, F2, G2 (black slots) CPU2: A1, B1, C1, E1, F1, G1 (blue slots) CPU2: A2, B2, C2, E2, F2, G2 (black slots)	32 GB, 64 GB, 128 GB
32	Fill all slots.	32 GB, 64 GB, 128 GB

Memory Installation Order for the Single-socket NX-1175S-G8 Platform

A memory channel is a group of DIMM slots.

On the NX-1175S-G8, the single CPU is associated with eight memory channels that contain one slot each.

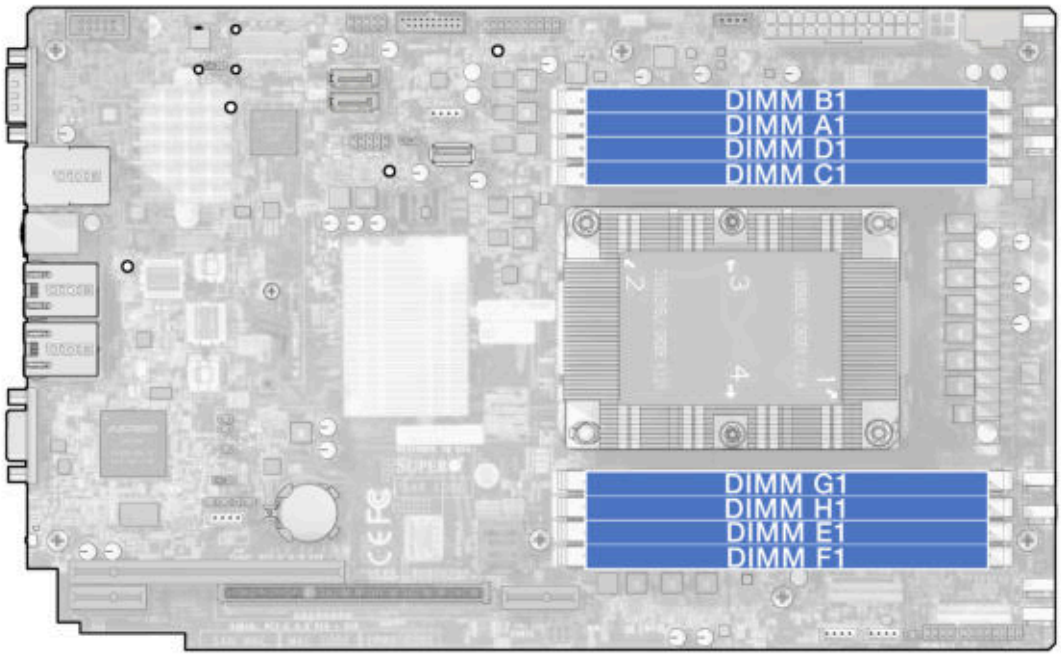

Figure 9: DIMM slots for an NX-1175S-G8 server board

Table 17: DIMM Installation Order for NX-1175S-G8

Number of DIMMs	Slots to use	Supported DIMM capacities
2	A1, E1	32 GB, 64 GB, 128 GB
4	A1, C1, E1, G1	32 GB, 64 GB, 128 GB
8	Fill all slots.	32 GB, 64 GB, 128 GB

COPYRIGHT

Copyright 2022 Nutanix, Inc.

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

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

License

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

Conventions

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

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
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

Interface	Target	Username	Password
SSH client or console	Acropolis OpenStack Services VM (Nutanix OVM)	root	admin

Version

Last modified: March 4, 2022 (2022-03-04T09:34:42+05:30)

Platform ANY

NX-8155-G8 System Specifications

NX-8155-G8

March 4, 2022

Contents

1. Nutanix Hardware Naming Convention....3

2. System Specifications....7
Node Naming (NX-8155-G8)....7
NX-8155-G8 System Specifications....9

3. Component Specifications....17
Controls and LEDs for Single-node Platforms....17
LED Meanings for Network Cards....18
Power Supply Unit (PSU) Redundancy and Node Configuration....19
Nutanix DMI Information....20

4. Memory Configurations....22
Supported Memory Configurations....22

Copyright....27
License....27
Conventions....27
Default Cluster Credentials....27
Version....28

1

NUTANIX HARDWARE NAMING CONVENTION

Every Nutanix block has a unique name based on the standardized naming convention followed by Nutanix.

The Nutanix hardware model name translates to the following naming convention:

Prefix-Body-Suffix or *NX-WXYZ(G|C|S)-G(4|5|6|7|8)*

Where:

- *NX* is the Nutanix appliance
- *W* is the product series
- *X* is the number of nodes
- *Y* is the chassis form-factor
- *Z* is the disk form-factor
- *G|C|S* indicates that a platform uses GPUs, a platform is an AHV-only node, or a platform supports only one CPU respectively
- *G(4|5|6|7|8)* is the CPU generation

Example: *NX-3460-G8*

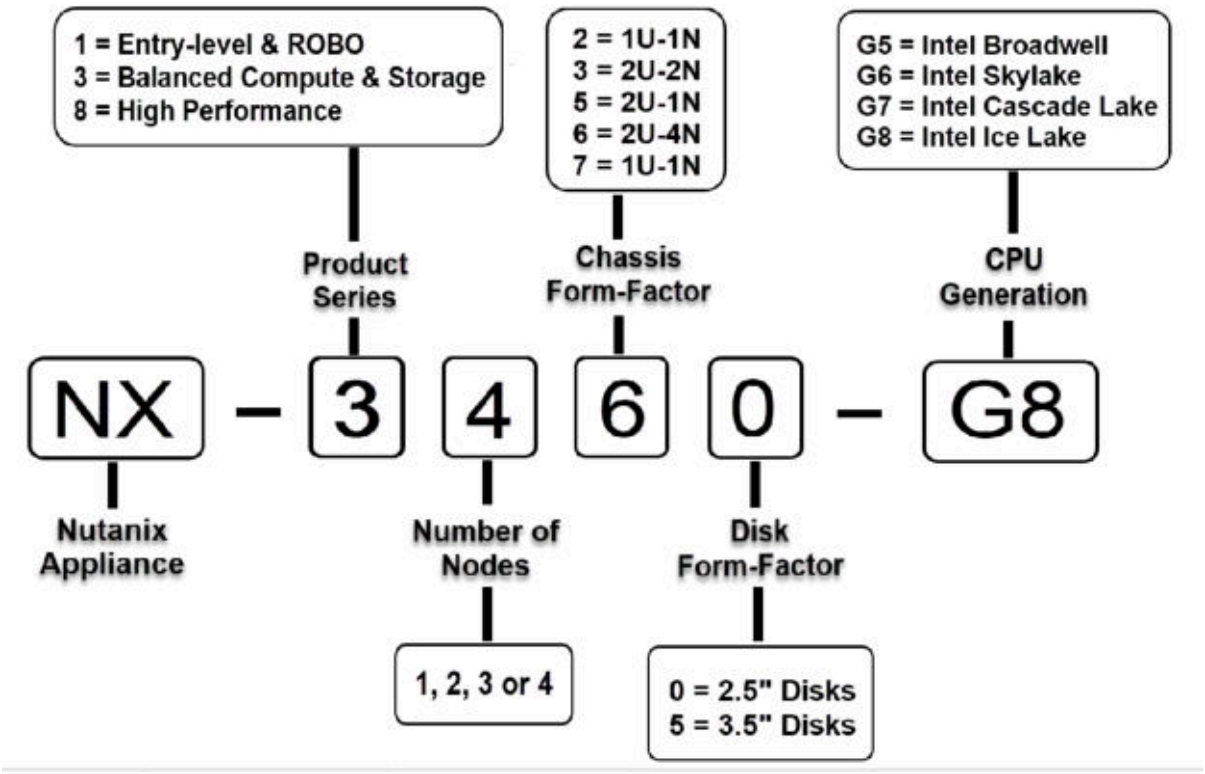

Figure 1: Nutanix Hardware Naming Convention

The following section provides a detailed description of the Nutanix hardware naming convention:

Table 1: Prefix

Prefix	Description
NX	Indicates that the platform is sold directly by Nutanix and support calls are handled by Nutanix. *NX* stands for *Nutanix*.

Table 2: Body

Body	Description
W	Indicates the product series and have one of the following values: • *1* is for small or Remote Office/Branch Office (ROBO) businesses • *3* is for heavy compute • *5* is for dedicated Files and Buckets • *6* is for heavy storage • *8* is for high-performance • *9* is for high performance and experimental businesses
X	Indicates the number of nodes and have one of the following values: • *1* is for single-node platforms • *2* is for multinode platforms • *3* is for multinode platforms • *4* is for multinode platforms Note: Though multinode platforms can have two, three, or four nodes, the number is always mentioned as a generic zero, *0* in the documentation.
Y	Indicates the chassis form-factor and have one of the following values: • *2* for 1U-1N (one rack unit high with one node) • *3* for 2U-2N (two rack units high with two nodes) • *5* for 2U-1N (two rack units high with one node) • *6* for 2U-4N (two rack units high with four nodes) • *7* for 1U-1N (one rack unit high with one node)
Z	Indicates the drive form-factor and have one of the following values: • *0* for 2.5 inch drives • *5* for 3.5 inch drives

Body	Description
• *G* • *C* • *S*	Indicates one of the following: • *G* at the end of the body stands for "graphics" and that the platform is optimized for using the Graphics Processing Unit (GPU) cards. For example, the NX-3155G-G7 platform. • *C* at the end of the body stands for "cold node" and that the platform is AHV-only. Such platforms should only be used for storage. Cold nodes cannot run VMs. Hence, cold nodes are the only exception where we allow mixing of the nodes with different hypervisors in the same cluster. For example, the NX-6035C and NX-6035C-G5 platforms. • *S* at the end of the body stands for "single socket" where the motherboard has only one CPU instead of the usual two that we support. For example, the NX-1175S-G7 platform.

Table 3: Suffix

Suffix	Description
G4	Indicates that the platform uses the Intel Haswell CPU
G5	Indicates that the platform uses the Intel Broadwell CPU
G6	Indicates that the platform uses the Intel Skylake CPU
G7	Indicates that the platform uses the Intel Cascade Lake CPU
G8	Indicates that the platform uses the Intel Ice Lake CPU

2

SYSTEM SPECIFICATIONS

Node Naming (NX-8155-G8)

Nutanix assigns a name to each node in a block, which varies by product type.

The NX-8155-G8 block contains a single node named Node A.

The NX-8155-G8 has twelve drive slots. The first two SSDs contain the Controller VM and metadata. All HDDs (if used) are data-only drives.

> Note: Certain processors (6342, 6348, 8358, 8360Y, and 8380) have a maximum operating temperature of 30 degrees C.

> Note: The 18 TB HDD is available only for Files and Objects use cases.

Figure 2: NX-8155-G8 Front Panel (NVMe Configuration Shown)

Partial Population

The NX-8155-G8 supports partially populating the chassis with fewer than the maximum number of drives, according to the following restrictions:

- Partially populated configurations support only even numbers of drives.
- SSD with NVMe configurations must include a minimum of four SSDs.
- Populate the drive slots in order from bottom to top and left to right.

When partially populating a node, Nutanix recommends that you put an empty drive carrier in each unpopulated drive slot, for air flow and to prevent dust contamination.

Drive Configurations

> Note: If you order a hybrid SSD and HDD configuration, you can later convert the platform to an all-SSD configuration. However, if you begin with either a hybrid SSD and HDD or an all-SSD configuration, you cannot later convert to an SSD with NVMe configuration.

- Two SSDs in slots 1 and 2; four HDDs in slots 3 through 6; all other slots empty
- Two SSDs in slots 1 and 2; six HDDs in slots 3 through 8; all other slots empty
- Two SSDs in slots 1 and 2; eight HDDs in slots 3 through 10; all other slots empty

- Two SSDs in slots 1 and 2; 10 HDDs in slots 3 through 12
- Four SSDs in slots 1 through 4; eight HDDs in slots 5 through 12

2x SSD + 4, 6, 8, or 10 HDD

HDD1	HDD4	HDD7	HDD10
SSD2	HDD3	HDD6	HDD9
SSD1	HDD2	HDD5	HDD8

Figure 3: Hybrid Configuration 1

4x SSD + 8x HDD

SSD3	HDD2	HDD5	HDD8
SSD2	HDD1	HDD4	HDD7
SSD1	SSD4	HDD3	HDD6

Figure 4: Hybrid Configuration 2

- Two SSDs in slots 1 and 2; all other slots empty
- Four SSDs in slots 1 through 4; all other slots empty
- Six SSDs in slots 1 through 6; all other slots empty
- Eight SSDs in slots 1 through 8; all other slots empty
- 10 SSDs in slots 1 through 10; all other slots empty
- 12 SSDs (fill all slots)

SLOT 3: SSD3	SLOT 6: SSD6	SLOT 9: SSD9	SLOT 12: SSD12
SLOT 2: SSD2	SLOT 5: SSD5	SLOT 8: SSD8	SLOT 11: SSD11
SLOT 1: SSD1	SLOT 4: SSD4	SLOT 7: SSD7	SLOT 10: SSD10

Populate the drive slots two by two, in numerical order. For a four-drive partial configuration, put the drives in slots 1 through 4; for a six-drive partial configuration, put the drives in slots 1 through 6; and so on.

Figure 5: All-SSD Configuration

- Four SSDs in slots 1 through 4; four NVMe drives in slots 9 through 12; all other slots empty
- Six SSDs in slots 1 through 6; four NVMe drives in slots 9 through 12; all other slots empty
- Eight SSDs in slots 1 through 8; four NVMe drives in slots 9 through 12

Slot 3: SSD3	Slot 6: SSD6	Slot 9: NVMe1	Slot 12: NVMe4
Slot 2: SSD2	Slot 5: SSD5	Slot 8: SSD8	Slot 11: NVMe3
Slot 1: SSD1	Slot 4: SSD4	Slot 7: SSD7	Slot 10: NVMe2

- Put four NVMe drives in slots 9 through 12. No other NVMe configurations are supported.
- All NVMe capacities can mix with all SSD capacities.
- When using NVMe drives, you must include a minimum of four SSDs. Populate the drive slots two by two, in numerical order.

Figure 6: SSD with NVMe Configuration

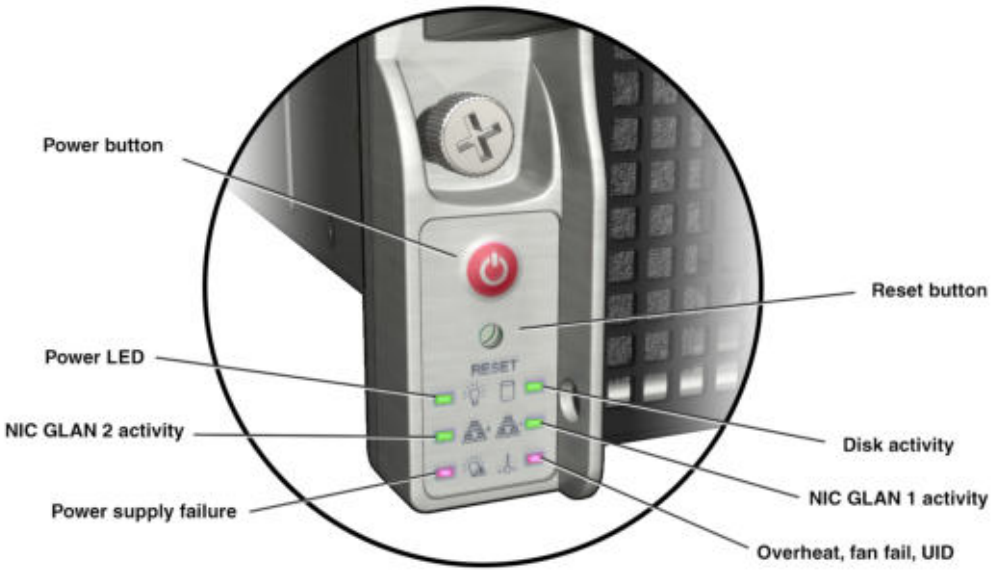

Figure 7: NX-8155-G8 Control Panel

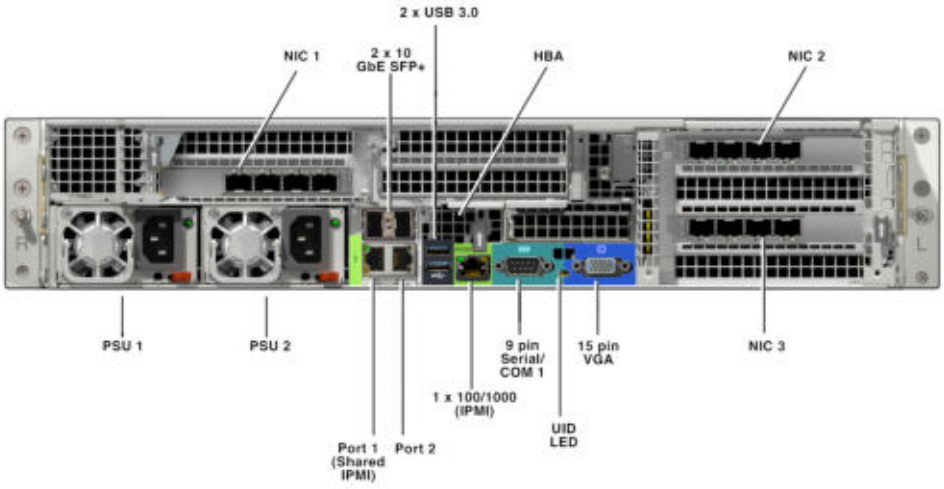

Figure 8: NX-8155-G8 Back Panel

NX-8155-G8 System Specifications

Table 4: System Characteristics

Boot Device	Dual M.2 RAID1 • 2 x 512GB M.2 Boot Device
CPU	Note: • Processor models 6342, 6348, 8358, 8360Y, 8380 have a maximum operating temperature of 30C. Processor • 2 x Intel Xeon® Silver 4310 [12 cores / 2.10 GHz] • 2 x Intel Xeon® Silver 4314 [16 cores / 2.40 GHz] • 2 x Intel Xeon® Silver 4316 [20 cores / 2.30 GHz] • 2 x Intel Xeon® Gold 5315Y [8 cores / 3.20 GHz] • 2 x Intel Xeon® Gold 5317 [12 cores / 3.00 GHz] • 2 x Intel Xeon® Gold 5318Y [24 cores / 2.10 GHz] • 2 x Intel Xeon® Gold 5320 [26 cores / 2.20 GHz] • 2 x Intel Xeon® Gold 6326 [16 cores / 2.90 GHz] • 2 x Intel Xeon® Gold 6334 [8 cores / 3.60 GHz] • 2 x Intel Xeon® Gold 6338 [32 cores / 2.00 GHz] • 2 x Intel Xeon® Gold 6342 [24 cores / 2.80 GHz] • 2 x Intel Xeon® Gold 6346 [16 cores / 3.10 GHz] • 2 x Intel Xeon® Gold 6348 [28 cores / 2.60 GHz] • 2 x Intel Xeon® Gold 6354 [18 cores / 3.00 GHz] • 2 x Intel Xeon® Platinum 8358 [32 cores / 2.60 GHz] • 2 x Intel Xeon® Platinum 8360Y [36 cores / 2.40 GHz] • 2 x Intel Xeon® Platinum 8380 [40 cores / 2.30 GHz]

Memory	32GB RDIMM • 4 x 32 GB = 128 GB • 8 x 32 GB = 256 GB • 12 x 32 GB = 384 GB • 16 x 32 GB = 512 GB • 24 x 32 GB = 768 GB • 32 x 32 GB = 1.0 TB 64GB DIMM • 4 x 64 GB = 256 GB • 8 x 64 GB = 512 GB • 12 x 64 GB = 768 GB • 16 x 64 GB = 1.0 TB • 24 x 64 GB = 1.5 TB • 32 x 64 GB = 2.0 TB
Network	Serverboard • 1x 1GbE Dedicated IPMI • 2x 10GbE SFP+ • 2x 10GBaseT (IPMI Failover) NICs in PCIe Slots • 0, 1, 2 or 3 x 10GBaseT 2P NIC • 0, 1, 2 or 3 x 10GBaseT 4P NIC • 0, 1, 2 or 3 x 10GbE 4P NIC • 0, 1, 2 or 3 x 25GbE 2P NIC • 0, 1, 2 or 3 x 100GbE 2P NIC

Network Cables	Network Cables • OPT,CBL,100GbE,QSFP28,1M • OPT,CBL,QSFP+ to QSFP+,1M,CU • OPT,CBL,SFP28,1M,CU • OPT, CBL, 1M, SFP+ TO SFP+ • OPT,CBL,100GbE,QSFP28,3M • OPT,CBL,QSFP+ to QSFP+,3M,CU • OPT,CBL,SFP28,3M,CU • OPT, CBL, 3M, SFP+ TO SFP+ • OPT,CBL,100GbE,QSFP28,5M • OPT,CBL,QSFP+ to QSFP+,5M,CU • OPT,CBL,SFP28,5M,CU • OPT, CBL, 5M, SFP+ TO SFP+
Power Cable	Power Cable • 2 x C13/14 4ft Power Cable
Power Supply	2000W PSU • 2 x 2000W Titanium PSU
Server	Server • NX-8155-G8 Server
Storage	Note: • The number of HDDs should be twice or more the number of SSDs. • 18TB HDD configurations are for Files and Objects use cases only.
Storage : NVMe +SSD	4 x NVMe • 1.92TB • 3.84TB • 7.68TB 4, 6 or 8 x SSD • 1.92TB • 3.84TB • 7.68TB

Storage : All SSD	4, 6, 8, 10 or 12 x SSD • 1.92TB • 3.84TB • 7.68TB
Storage : All SSD SED	4, 6, 8, 10 or 12 x SSD • 1.92TB • 3.84TB
Storage : SSD +HDD	2 x SSD • 1.92TB • 3.84TB • 7.68TB 4, 6, 8 or 10 x HDD • 6.0TB
Storage : SSD +HDD	2 x SSD • 3.84TB • 7.68TB 4, 6, 8 or 10 x HDD • 8.0TB • 12.0TB
Storage : SSD +HDD	2 x SSD • 3.84TB • 7.68TB 4, 6, 8 or 10 x HDD • 18.0TB

Storage : SSD +HDD	4 x SSD • 1.92TB • 3.84TB • 7.68TB 8 x HDD • 6.0TB • 8.0TB • 12.0TB
Storage : SSD +HDD	2 or 4 x SSD • 3.84TB • 7.68TB 8 x HDD • 18.0TB
Storage : SSD +HDD SED	2 x SSD • 3.84TB 4, 6, 8 or 10 x HDD • 8.0TB • 12.0TB
Storage : SSD +HDD SED	4 x SSD • 1.92TB • 3.84TB 8 x HDD • 8.0TB • 12.0TB
TPM	TPM • 1 x Unprovisioned Trusted Platform Module
Transceiver	Transceiver • SR SFP+ Transceiver
Chassis fans	4x 80 mm heavy duty fans with PWM fan speed controls

Table 5: Block, power and electrical

Block	Depth : 741.1 mm Width : 440.9 mm Rack Units : 2 U Height : 88.1 mm Weight : 27 kg
Package	Weight : 36.29 kg
Shock	Non-Operating : 10 ms Operating : 2.5 ms
Thermal Dissipation	Maximum : 4007 BTU/hr Typical : 2806 BTU/hr
Vibration (Random)	Operating : 0.4 Grms Non-Operating : 0.98 Grms
Power consumption	Max Config • Maximum: 1175 VA • Typical: 822 VA
Operating environment	Note: • Processor models 6342, 6348, 8358, 8360Y, 8380 have a maximum operating temperature of 30C. Operating temperature : 10-35C Non-Operating temperature : -45-70C Operating relative humidity : 20-90% Non-operating relative humidity : 5-95%

Certifications

- Energy Star
- CSAus
- FCC
- CSA
- UL
- cUL
- ICES

- CE
- UKCA
- KCC
- RCM
- VCCI-A
- BSMI
- EAC
- SABS
- S-MARK
- UKRSEPRO
- BIS
- SII
- RoHS
- Reach
- WEEE

3

COMPONENT SPECIFICATIONS

Controls and LEDs for Single-node Platforms

Table 6: Front Panel Controls and Indicators

LED or Button	Function
Power button	System on/off (Press and hold 4 seconds to power off.)
Reset button	System reset
Power LED	Solid green when block is powered on
Disk Activity LED	Flashing orange on activity
10 GbE LAN 1 LED (Port 1)	Activity = green flashing
10 GbE LAN 2 LED (Port 2)	Activity = green flashing
Multiple Function ***i*** LED	
i Unit identification	Solid blue
i Overheat condition	Solid red
i Power supply failure	Flashing red at 0.25Hz
i Fan failure	Flashing red at 1Hz
Unit Identifier button (UID)	Press button to illuminate the ***i*** LED (blue)

Table 7: Drive LEDs

Top LED: Activity	Blue or green, blinking = I/O activity, off = idle
Bottom LED: Status	Red solid = failed drive, On five seconds after boot = power on

Table 8: Back Panel Controls and Indicators

LED	Function
PSU1 and PSU2 LED	Green = OK (PSU normal) Yellow = no node power or PSU is not inserted completely

LED	Function
	Red for Fail
10 GbE LAN 1 LED (Port 1)	Activity = green flashing, Link = solid green
10 GbE LAN 2 LED (Port 2)	Activity = green flashing, Link = solid green
1 GbE dedicated IPMI	Left LED green for 100 Mbps, Left amber for 1 Gbps, Right yellow flashing for activity
i Unit identification LED	Solid blue
2 x 10 GbE and 4 x 10 GbE NIC LEDs	Link and Activity = green flashing 10 Gbps, Amber = 1 Gbps

For LED states for add-on NICs, see LED Meanings for Network Cards on page 18.

LED Meanings for Network Cards

Descriptions of LEDs for supported NICs.

Different NIC manufacturers use different LED colors and blink states. Not all NICs are supported for every Nutanix platform. See the system specifications for your platform to verify which NICs are supported.

Table 9: On-Board Ports

NIC	Link (LNK) LED	Activity (ACT) LED
1 GbE dedicated IPMI	Green: 100 Mbps Yellow: 1 Gbps	Blinking yellow: activity
1 GbE shared IPMI	Green: 1 Gbps Yellow: 100 Mbps	Blinking yellow: activity
i Unit identification LED	Blinking blue: UUID has been activated.	

Table 10: SuperMicro NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 1 GbE	Green: 100 Mbps Yellow: 1 Gb/s OFF: 10 Mb/s or No Connection	Blinking yellow: activity
Dual-port 10G SFP+	Green: 10 Gb Yellow: 1 Gb	Blinking green: activity

Table 11: Silicom NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 10G SFP+	Green: all speeds	Solid green: idle Blinking green: activity
Quad-port 10G SFP+	Blue: 10 Gb Yellow: 1 Gb	Solid green: idle Blinking green: activity
Dual-port 10G BaseT	Yellow: 1 Gb/s Green: 10 Gb/s	Blinking green: activity

Table 12: Mellanox NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual-port 10G SFP+ ConnectX-3 Pro	Green: 10 Gb speed with no traffic Blinking yellow: 10 Gb speed with traffic Not illuminated: no connection	Blinking yellow and green: activity
Dual-port 40G SFP+ ConnectX-3 Pro	Solid green: good link	Blinking yellow: activity Not illuminated: no activity
Dual-port 10G SFP28 ConnectX-4 Lx	Solid yellow: good link Blinking yellow: physical problem with link	Solid green: valid link with no traffic Blinking green: valid link with active traffic
Dual-port 25G SFP28 ConnectX-4 Lx	Solid yellow: good link Blinking yellow: physical problem with link	Solid green: valid link with no traffic Blinking green: valid link with active traffic

Power Supply Unit (PSU) Redundancy and Node Configuration

Note: Nutanix recommends that you carefully plan your AC power source needs, especially in cases where the cluster consists of mixed models.

Nutanix recommends that you use 180 V ~ 240 V AC power source to secure PSU redundancy. However, according to the following table, and depending on the number of nodes in the chassis, some NX platforms can work with redundant 100 V ~ 210 V AC power supply units.

Table 13: PSU Redundancy and Node Configuration

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208 V
NX-1065-G8	1 or 2	YES	YES
	3 or 4	NO	YES
NX-1175S-G8	1	YES	YES
NX-3060-G8	1 or 2	YES	YES
	3 or 4	NO	YES
NX-3155G-G8	1, with GPU	NO	YES
	1, without GPU	YES	YES
NX-3170-G8	1	NO	YES
NX-8035-G8	1	YES	YES
	2	No	YES
NX-8150-G8	1	NO	YES
NX-8155-G8	1	For CPUs with a thermal design profile equal to or less than 130 W: redundant at 110 V over the entire supported ambient temperature range of 10° C to 35° C. For all other CPUs, use the following rule: • Ambient temperature is 25° C or less: YES • Ambient temperature is greater than 25° C: NO	YES
NX-8170-G8	1	NO	YES

Nutanix DMI Information

Format for Nutanix DMI strings.

VMware reads model information from the direct media interface (DMI) table.

For platforms with Intel Icelake CPUs, Nutanix provides model information to the DMI table in the following format:

```
NX-motherboard_idNIC_id-HBA_id-G8
```

motherboard-id has the following options:

Argument	Option
T	X12 multi-node motherboard
U	X12 single-node motherboard
W	X12 single-socket single-node motherboard

NIC_id has the following options:

Argument	Option
D1	dual-port 1G NIC
Q1	quad-port 1G NIC
DT	dual-port 10GBaseT NIC
QT	quad-port 10GBaseT NIC
DS	dual-port SFP+ NIC
QS	quad-port SFP+ NIC

HBA_id specifies the number of nodes and type of HBA controller. For example:

Argument	Option
1NL3	single-node LSI3808
2NL3	2-node LSI3808
4NL3	4-node LSI3808

Table 14: Examples

DMI string	Explanation	Nutanix model
NX-TDT-4NL3-G8	X12 motherboard with dual-port 10GBase-T NIC, 4 nodes with LSI3808 HBA controllers	NX-1065-G8, NX-3060-G8
NX-UDT-1NL3-G8	X12 motherboard with dual-port 10GBase-T NIC, 1 node with LSI3808 HBA controller	NX-3155G-G8, NX-3170-G8, NX-8150-G8, NX-8155-G8, NX-8170-G8

4

MEMORY CONFIGURATIONS

Supported Memory Configurations

DIMM installation information for all Nutanix G8 platforms.

DIMM Restrictions

DIMM type

Each G8 node must contain only DIMMs of the same type. So, for example, you cannot mix RDIMM and LRDIMM in the same node.

DIMM capacity

Each G8 node must contain only DIMMs of the same capacity. So, for example, you cannot mix 32 GB DIMMs and 64 GB DIMMs in the same node.

DIMM speed

G8 nodes ship with 3200 MHz DIMMs. 3200 MHz is the highest speed Nutanix currently supports, so you cannot currently mix DIMM speeds in any G8 node.

DIMM manufacturer

You can mix DIMMs from different manufacturers in the same *node*, but not in the same *channel*.

Multi-node platforms

Multi-node G8 platforms contain only one active DIMM slot per channel, so mixing DIMMs in the same channel is not possible.

Single-socket platforms

The single-socket NX-1175S-G8 platform contains only one DIMM slot per channel, so mixing DIMMs in the same channel is not possible.

Single-node platforms

- Single-node G8 platforms contain two DIMM slots per channel. Within a channel, all DIMMs must be from the same manufacturer.
- When replacing a failed DIMM, if there are two DIMMs in the channel, either replace the failed DIMM with a new DIMM form the same manufacturer, or else replace both DIMMs in the channel and make sure that both new DIMMs are from the same manufacturer.
- When adding new DIMMs to a node, if the new DIMMs and the original DIMMs have different manufacturer part numbers, arrange the DIMMs so that the original DIMMs and the new DIMMs are not mixed in the same channel.
    - EXAMPLE: You have an NX-8155-G8 node that has sixteen 32GB DIMMs for a total of 512 GB. You decide to upgrade to thirty-two 32GB DIMMs for a total of 1024 GB. When you remove the node from the chassis and look at the server board, you see that each CPU has eight DIMMs. Remove all DIMMs from one CPU and place them in the empty DIMM slots for the other CPU.

Then place all the new DIMMs in the DIMM slots for the first CPU, filling all slots. This way you can ensure that the original DIMMs and the new DIMMs do not share channels.

Note: You do not need to balance numbers of DIMMs from different manufacturers within a node, so long as you never mix them in the same channel.

Memory Installation Order for Multi-Node G8 Platforms

A memory channel is a group of DIMM slots.

For G8 multi-node platforms, each CPU is associated with eight active memory channels that contain one blue slot each, plus two inactive black slots.

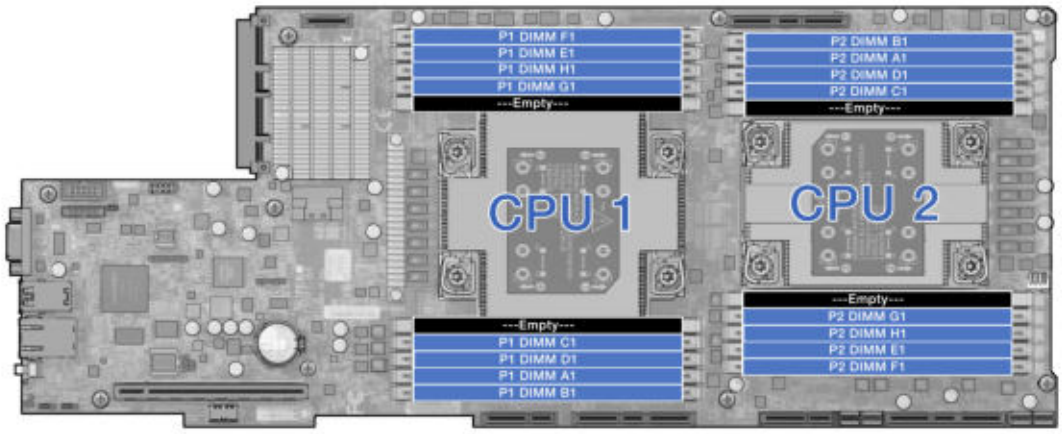

Note: The black slots (C2 and G2 on each CPU) are inactive.

Figure 9: DIMM Slots for a G8 Multinode Serverboard

Table 15: DIMM Installation Order for Multi-node G8 Platforms

Number of DIMMs	Slots to use	Supported capacities (NX-1065-G8)	Supported capacities (NX-3060-G8 and NX-8035-G8)
4	CPU1: A1, E1 CPU2: A1, E1	32 GB, 64 GB	32 GB, 64 GB, 128 GB
8	CPU1: A1, C1, E1, G1 CPU2: A1, C1, E1, G1	32 GB, 64 GB	32 GB, 64 GB, 128 GB
12	CPU1: A1, B1, C1, E1, F1, G1 CPU2: A1, B1, C1, E1, F1, G1	32 GB, 64 GB	32 GB, 64 GB, 128 GB

Number of DIMMs	Slots to use	Supported capacities (NX-1065-G8)	Supported capacities (NX-3060-G8 and NX-8035-G8)
16	Fill all blue slots.	32 GB, 64 GB	32 GB, 64 GB

Memory Installation Order for Single-node G8 Platforms

A memory channel is a group of DIMM slots.

For G8 single-node platforms, each CPU is associated with eight memory channels that contain one blue slot and one black slot each, for a total of 32 DIMM slots.

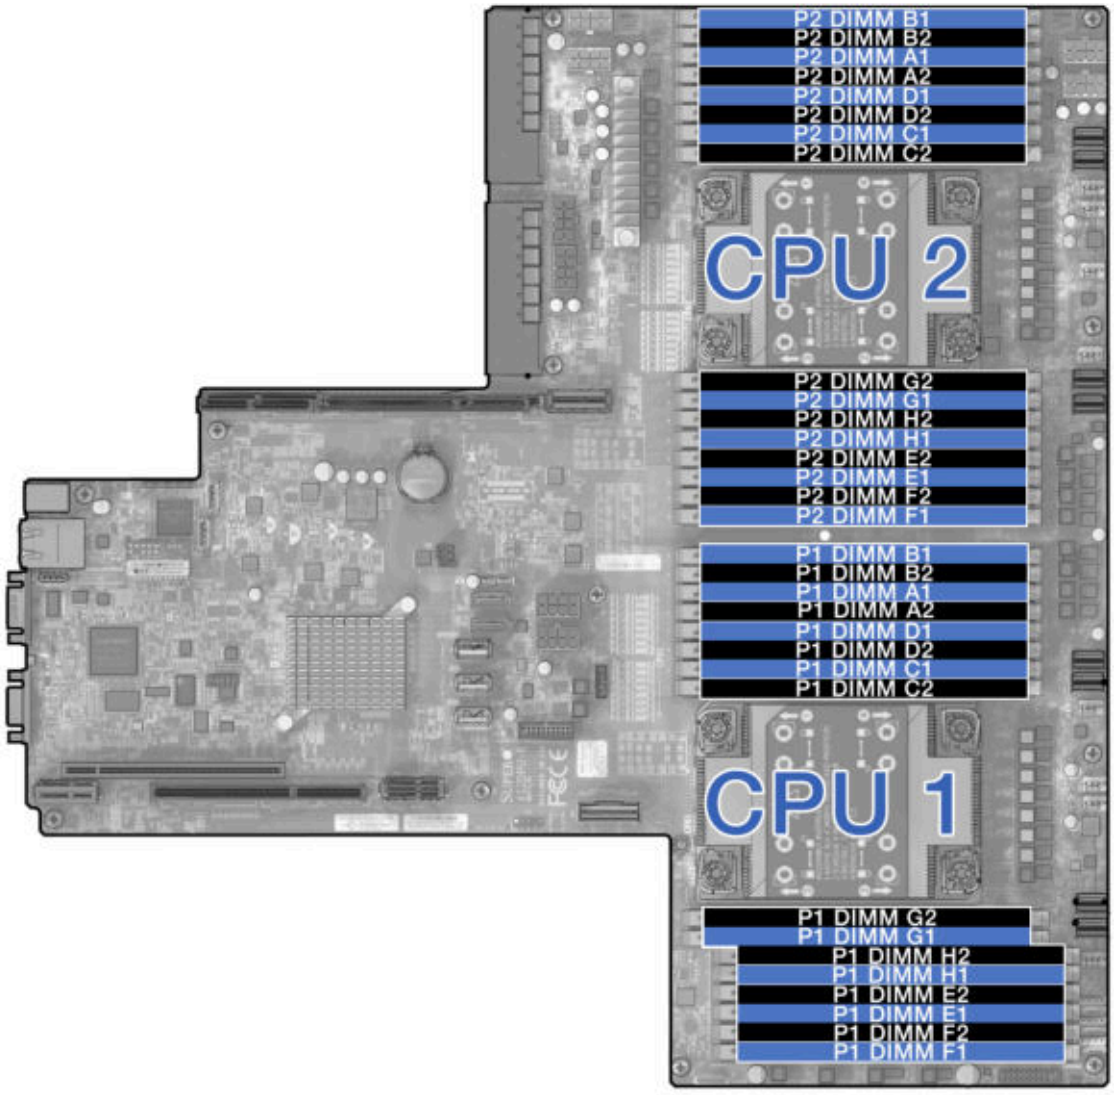

Figure 10: DIMM Slots for a G8 Single-Node Serverboard

Table 16: DIMM Installation Order for Single-Node G8 Platforms

Number of DIMMs	Slots to use	Supported capacities
4	CPU1: A1, E1 (blue slots) CPU2: A1, E1 (blue slots)	32 GB, 64 GB, 128 GB
8	CPU1: A1, C1, E1, G1 (blue slots) CPU2: A1, C1, E1, G1 (blue slots)	32 GB, 64 GB, 128 GB
12	CPU1: A1, B1, C1, E1, F1, G1 (blue slots) CPU2: A1, B1, C1, E1, F1, G1 (blue slots)	32 GB, 64 GB, 128 GB
16	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots)	32 GB, 64 GB, 128 GB
24	CPU1: A1, B1, C1, E1, F1, G1 (blue slots) CPU1: A2, B2, C2, E2, F2, G2 (black slots) CPU2: A1, B1, C1, E1, F1, G1 (blue slots) CPU2: A2, B2, C2, E2, F2, G2 (black slots)	32 GB, 64 GB, 128 GB
32	Fill all slots.	32 GB, 64 GB, 128 GB

Memory Installation Order for the Single-socket NX-1175S-G8 Platform

A memory channel is a group of DIMM slots.

On the NX-1175S-G8, the single CPU is associated with eight memory channels that contain one slot each.

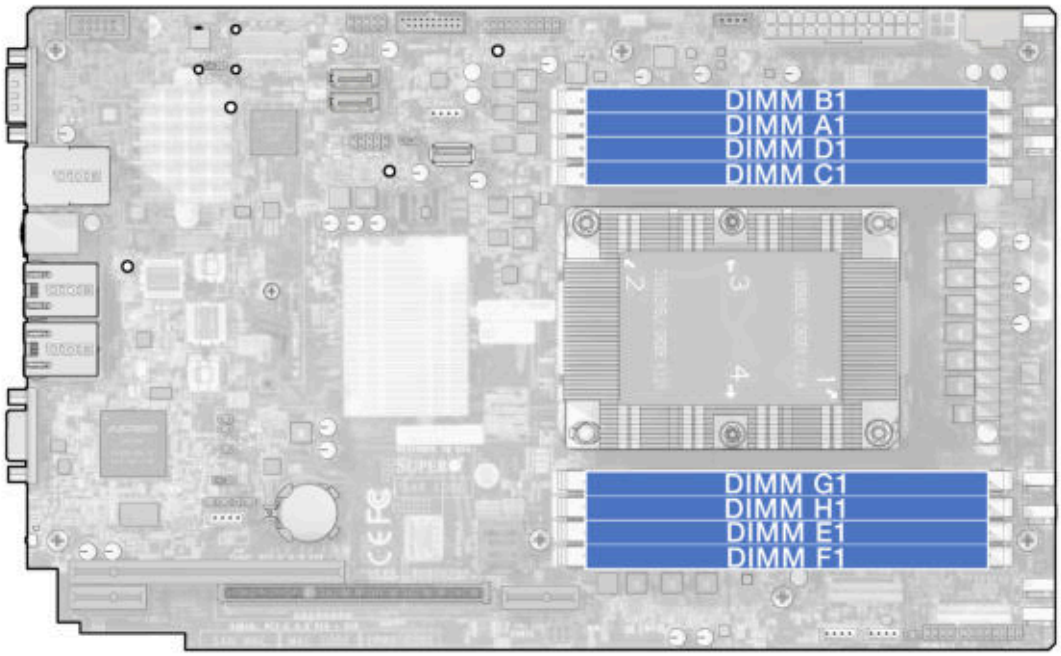

Figure 11: DIMM slots for an NX-1175S-G8 server board

Table 17: DIMM Installation Order for NX-1175S-G8

Number of DIMMs	Slots to use	Supported DIMM capacities
2	A1, E1	32 GB, 64 GB, 128 GB
4	A1, C1, E1, G1	32 GB, 64 GB, 128 GB
8	Fill all slots.	32 GB, 64 GB, 128 GB

COPYRIGHT

Copyright 2022 Nutanix, Inc.

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

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

License

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

Conventions

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

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
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

Interface	Target	Username	Password
SSH client or console	Acropolis OpenStack Services VM (Nutanix OVM)	root	admin

Version

Last modified: March 4, 2022 (2022-03-04T09:37:26+05:30)

v7.0 | October 2021 | TN-2041

TECH NOTE

Nutanix Files

NUTANIX
YOUR ENTERPRISE CLOUD

Copyright

Copyright 2021 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.

Contents

1. Executive Summary 5

2. Audience and Purpose 6

3. Nutanix Enterprise Cloud Overview 8
- Nutanix HCI Architecture 9

4. Nutanix Files Architecture 11
- File Server Virtual Machine 11
- Exports and Shares 16
- Load Balancing and Scaling 25
- High Availability 28
- Active Directory and SMB Operations 31
- Active Directory, LDAP, and NFS Operations 38
- Multiprotocol 42
- Quotas 45
- Selective File Blocking 47
- File Analytics 49
- Smart Tier 57
- Hypervisor-Specific Support 63

5. Backup and Disaster Recovery 65
- Self-Service Restore 65
- Protection Domains and Consistency Groups 66
- Cluster Migration, Failure, and Restoration 67
- Cloning 67
- Files Smart Disaster Recovery 68

6. Third-Party Integration 74
- Antivirus 74
- File Operations Monitoring 77
- Intelligent Backup 79

7. Conclusion....84

Appendix....85
About Nutanix....85

List of Figures....86

List of Tables....89

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 Nutanix.

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, which unifies and simplifies 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 native remote replication and 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 Nutanix distributed storage for intracluster data resiliency. AOS distributed storage also provides data efficiency techniques such as erasure coding (EC-X).

Nutanix Files includes File Analytics, which gives you a variety of useful insights into your data, including full audit trails, anomaly detection, ransomware detection and intelligence, data age analytics, and custom reporting.

© 2021 Nutanix, Inc. All rights reserved

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.

© 2021 Nutanix, Inc. All rights reserved

Version Number	Published	Notes
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.
5.3	October 2020	Updated for version 3.7
6.0	April 2021	Updated for Files 3.8, File Analytics 3.0, and Files Manager 2.0.
7.0	October 2021	Updated for Files 4.0 and Data Lens. Added Smart Tier and SSR Interoperability sections. Updated Shares and Exports, Notifications, Selective File Block, and Files Operations Monitoring sections. Updated Files Smart DR figure.

© 2021 Nutanix, Inc. All rights reserved

3. Nutanix Enterprise Cloud Overview

Nutanix delivers a web-scale, hyperconverged infrastructure solution purpose-built for virtualization and both containerized and private 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 the core HCI product families—Nutanix AOS and Nutanix Prism management—along with other software products that automate, secure, and back up cost-optimized infrastructure.

Available attributes of the Nutanix enterprise cloud OS stack include:

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

Nutanix provides services and can be broken down into three main components: an HCI-based distributed storage fabric, management and operational intelligence from Prism, and AHV virtualization. Nutanix Prism furnishes one-click infrastructure management for virtual environments running on AOS. AOS is hypervisor agnostic, supporting two third-party hypervisors

© 2021 Nutanix, Inc. All rights reserved

—VMware ESXi and Microsoft Hyper-V—in addition to the native Nutanix hypervisor, AHV.

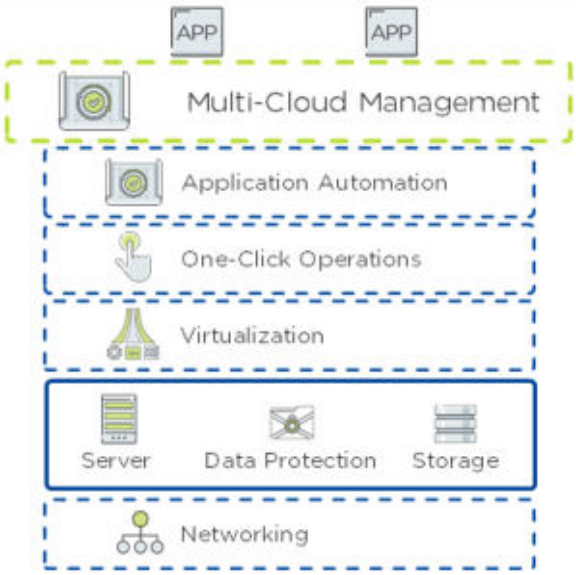

Figure 1: Nutanix Enterprise Cloud OS Stack

Nutanix HCI Architecture

Nutanix does not rely on traditional SAN or network-attached storage (NAS) 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 entirely solid-state storage with NVMe for optimal performance or a hybrid combination of SSD and HDD storage that provides a combination of performance and additional capacity. The storage fabric automatically tiers data across the cluster to different classes of storage devices using intelligent data placement algorithms. For best

© 2021 Nutanix, Inc. All rights reserved

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 Nutanix enterprise cloud software, visit the Nutanix Bible and Nutanix.com.

© 2021 Nutanix, Inc. All rights reserved

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 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.

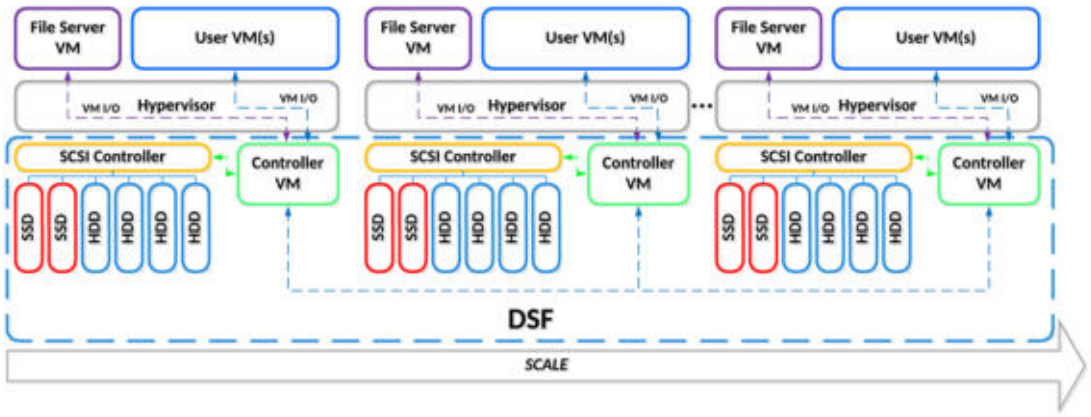

Figure 2: Nutanix Files Instances Run as VMs

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 basic configuration: four vCPU and 12 GiB of RAM. You can add more vCPU, 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 AOS 5.10.1 or later.

© 2021 Nutanix, Inc. All rights reserved

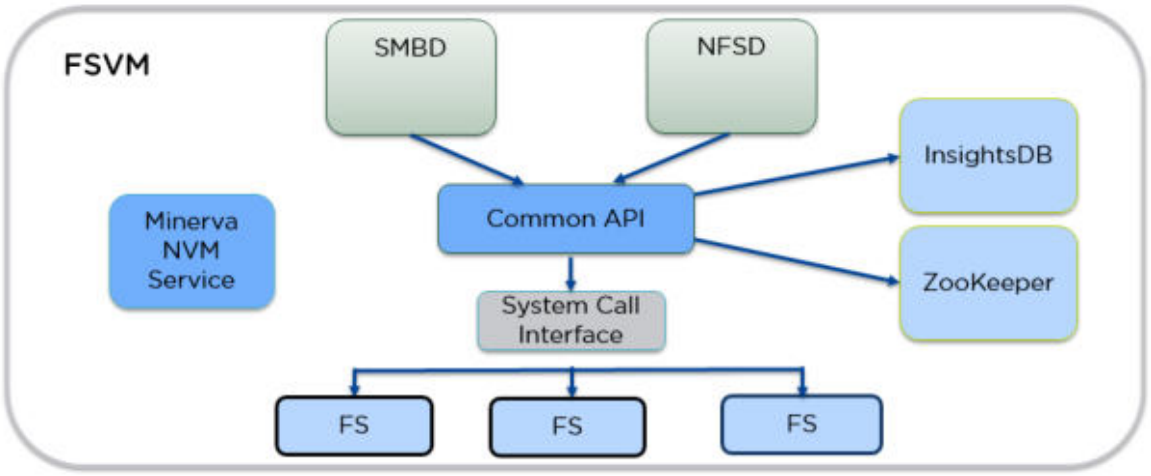

Figure 3: Data Path Architecture of Nutanix Files

Nutanix Files can support SMB and NFS from the same FSVM, and beginning with Files version 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 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.

© 2021 Nutanix, Inc. All rights reserved

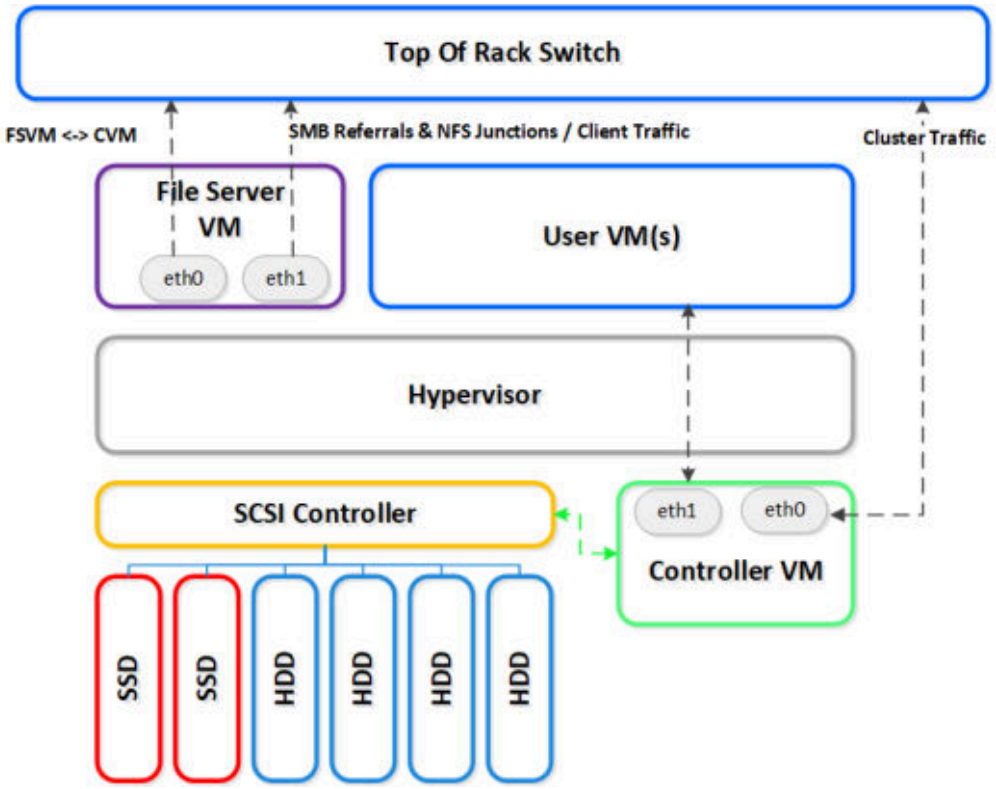

Figure 4: FSVM Internal Communication on One Node

The previous 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 that you place the FSVM storage interface on the same network VLAN as the Nutanix CVM management interface (eth0 of the CVM) to help ensure maximum performance for the iSCSI sessions.

© 2021 Nutanix, Inc. All rights reserved

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

Nutanix File Server VM

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.

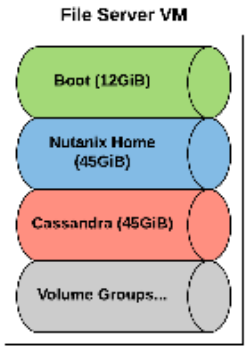

Figure 5: FSVM vDisks and Volume Groups

Starting with the Files 3.7 release, each FSVM also has a volume group that helps maintain auditing events and persistent handles for SMB Transparent Failover.

Shares and Exports

A Nutanix Files cluster is a single namespace that includes a collection of file systems used to support the SMB and NFS shares. The file systems support dynamic metadata, which enables you to store an unlimited number of files. The file system also supports variable block length allocations up to the default size of 64 KB. The variable block length matches the size of the file; for example, a 1 KB file allocates 1 KB on the file system. Any file over 64 KB allocates in 64

© 2021 Nutanix, Inc. All rights reserved

KB increments. Starting with Files 3.7, you can specify the maximum allocation size to improve performance. For environments with more random access patterns, you can choose a maximum allocation size of 16 KB. For environments with sequential patterns, you can choose 1 MB. These file systems use Nutanix volume groups as storage.

> Tip: For guidance on whether you should use the random or sequential file system option, see the Nutanix Files Performance Tech Note (requires Nutanix Portal account).

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 a fault occurs, one of the surviving FSVMs takes over volume group ownership and continues servicing requests.

With the 3.x release of Nutanix Files, volume groups contain 16 vDisks. Files 3.x versions prior to 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 can be used by the assigned shares.

With the 4.x release of Nutanix Files, volume groups contain 17 vDisks. We added an additional vDisk as a separate intent log (SLOG) to accept synchronous writes, as the SLOG device helps with long-duration synchronous random write performance.

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

Figure 6: 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.

© 2021 Nutanix, Inc. All rights reserved

Figure 7: 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 triggers when you've used 80 percent of the space in the current pool.

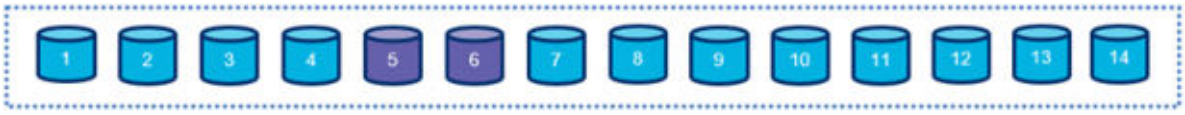

Figure 8: 120 TB File System Pool After Second Expansion

Figure 9: 140 TB File System Pool After Final Expansion

The Nutanix Files 3.7 release introduced a scale-up function to the file system pool. Once you've scaled a file system out to use all 16 disks, if that pool reaches 80 percent capacity, it uses additional space on the existing disks in the volume group. This vertical expansion allows support for a file system pool of up to 280 TB.

In summary, the maximum sizes for individual volume groups based on Nutanix Files version are as follows:

- Nutanix Files 2.x to 3.2: 40 TB
- Nutanix Files 3.2 to 3.6: 140 TB
- Nutanix Files 3.7 forward: 280 TB

> Tip: You must have created the File server with a minimum version of 3.x to take advantage of horizontal scale-out to 140 TB and vertical scale-up to 280 TB.

Exports and Shares

There are two types of SMB shares and NFS exports.

© 2021 Nutanix, Inc. All rights reserved

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.

Starting with Files version 3.7, you can store files in the root of distributed NFS shares. Distributed SMB shares 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 distributed file system (DFS) referrals for SMB and junctions for NFSv4 to make sure that clients can connect to the right top-level directories.

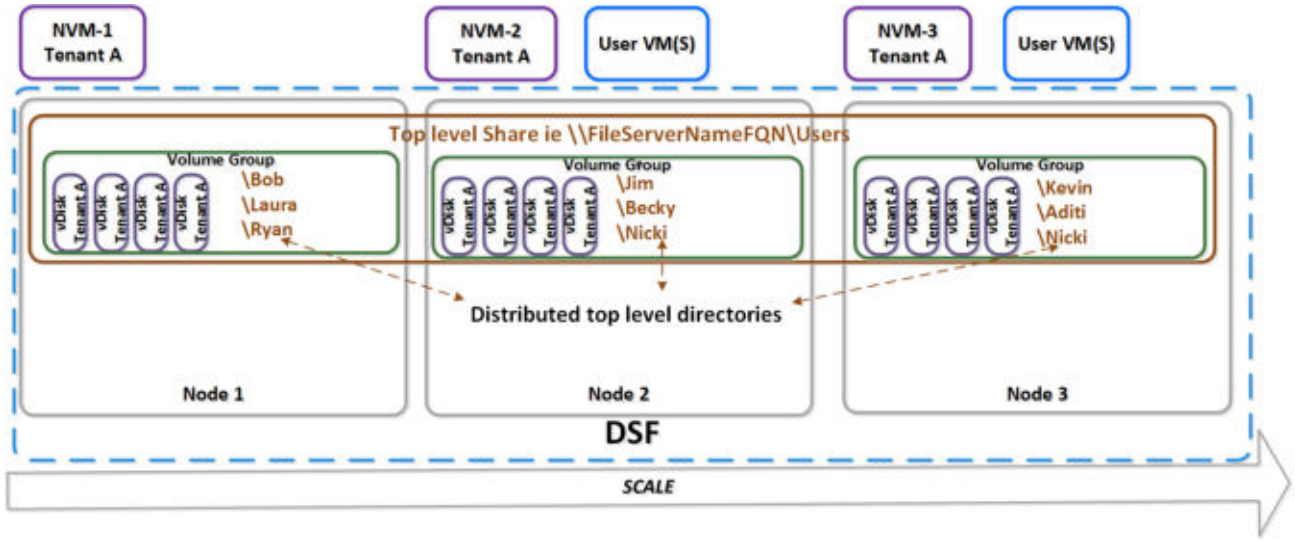

Figure 10: Distributed Directory Shares

© 2021 Nutanix, Inc. All rights reserved

Distributed share directories work well for home shares and exports because Nutanix Files automatically spreads the workload over multiple FSVMs per user (see the previous figure). If a user creates a share called "\\FileServer1\Users" that contains the 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 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. We recommend that you set permissions at the share or export root before you deploy 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 don't 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.

© 2021 Nutanix, Inc. All rights reserved

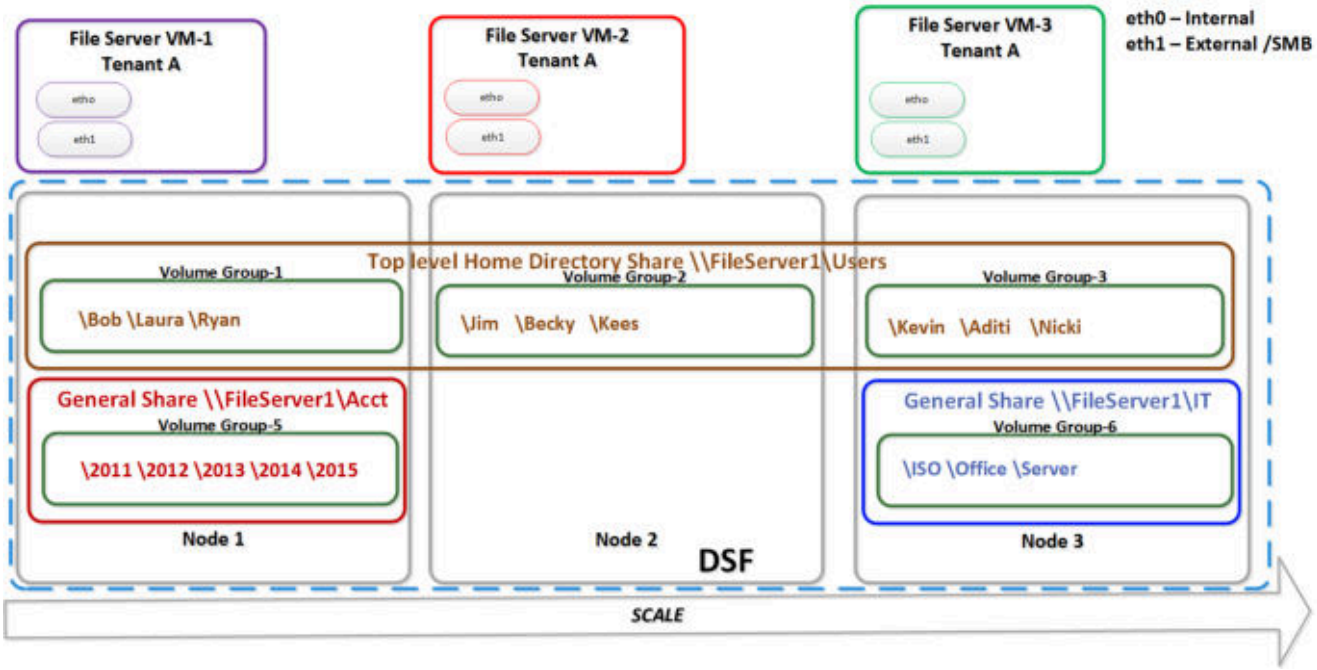

Figure 11: Two Standard Shares on the Same File Server

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 someone accesses the directory. Because these are exports rather than standard directories, it takes a few steps to remove a top-level directory. In this section we demonstrate how distributed top-level directories behave and then walk you 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:

© 2021 Nutanix, Inc. All rights reserved

```
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 you delete a top-level directory. You can delete the project2 directory from any NFS client with the export mounted if you sign 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 don't automatically mount the top-level directories of distributed shares. Because the top-level directories aren't submounted, you can rename and delete them without the additional steps required for NFSv4.

© 2021 Nutanix, Inc. All rights reserved

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

© 2021 Nutanix, Inc. All rights reserved

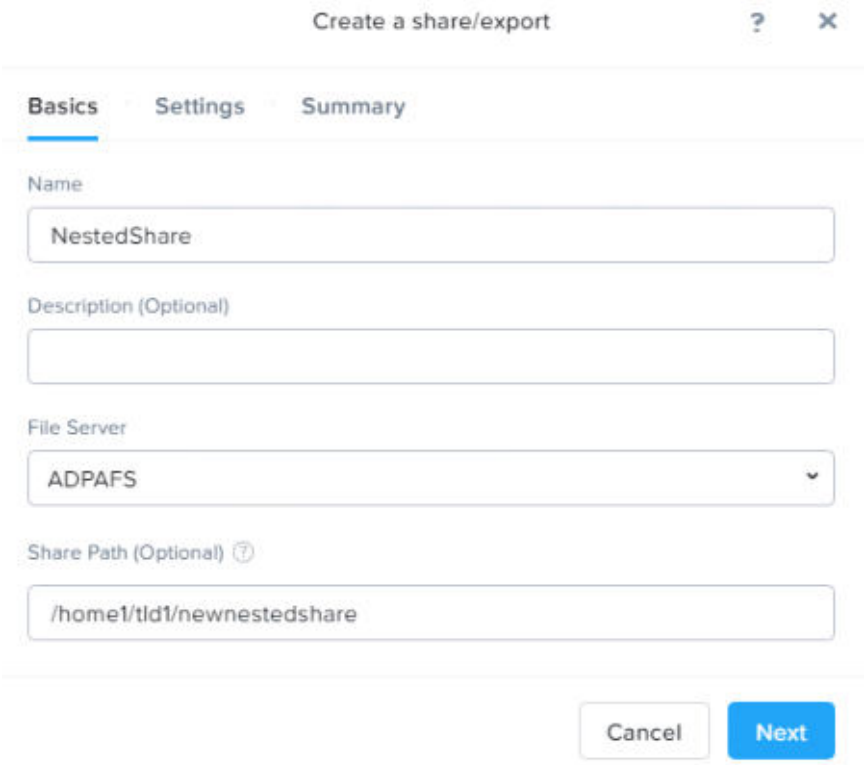

Figure 12: Create Share Menu with Nested Share Path

Connected Shares

The connected shares function allows you to mount shares and exports as subfolders within other shares and exports. You can create a folder at any level of the folder hierarchy and mount either a standard or distributed share into that folder. Starting with Files 3.8, the share used as the parent hosting the folder can be a standard, distributed, or nested share.

© 2021 Nutanix, Inc. All rights reserved

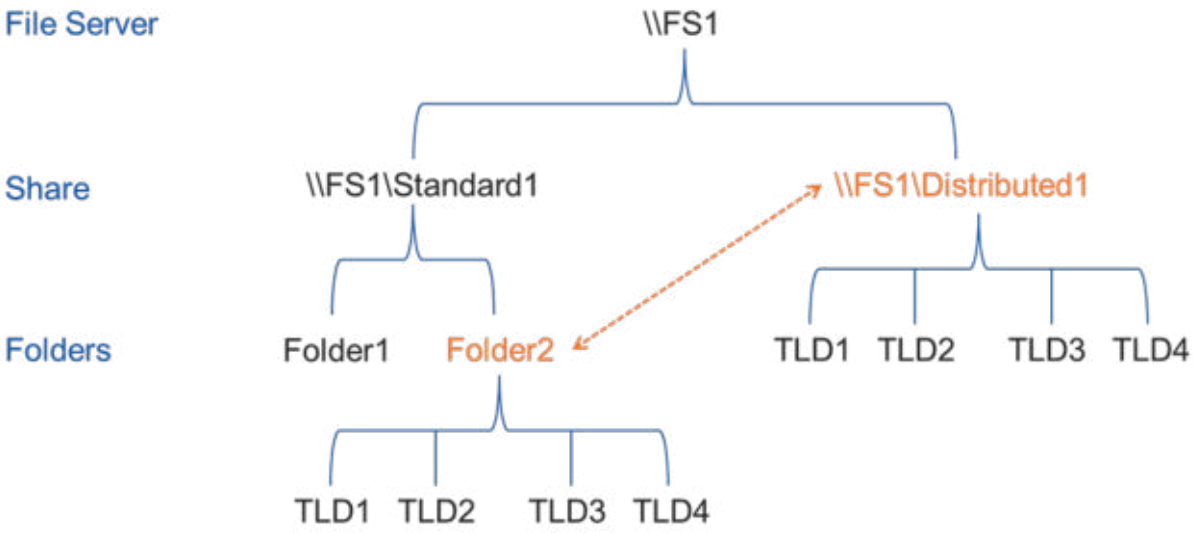

Figure 13: Submounting a Distributed Share into a Standard Share

There are several benefits of submounting shares and exports:

- You can use a distributed share at any level of a directory structure. Some applications have the folder paths where they store data or create directories hardcoded. If you submount, you can establish a distributed share at the level that matches these application requirements.
- A folder adopts the size limit of the submounted shares, which enables folder-level quotas.
- A folder adopts the self-service restore setting for a submounted share, which enables folder-level snapshots.

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.

© 2021 Nutanix, Inc. All rights reserved

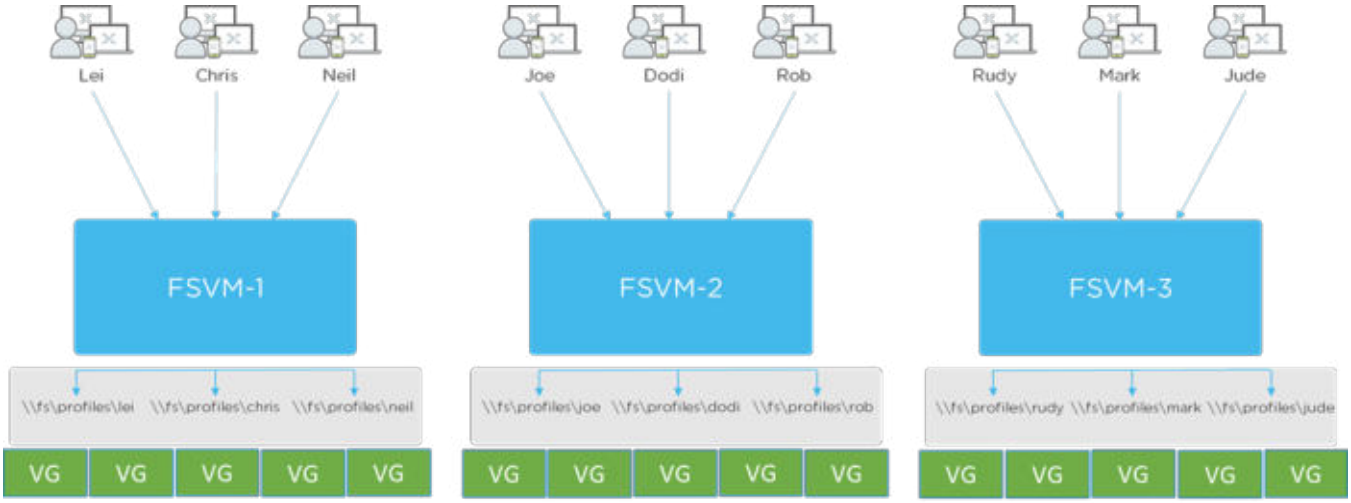

Figure 14: Distributed Share and Top-Level Directories

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

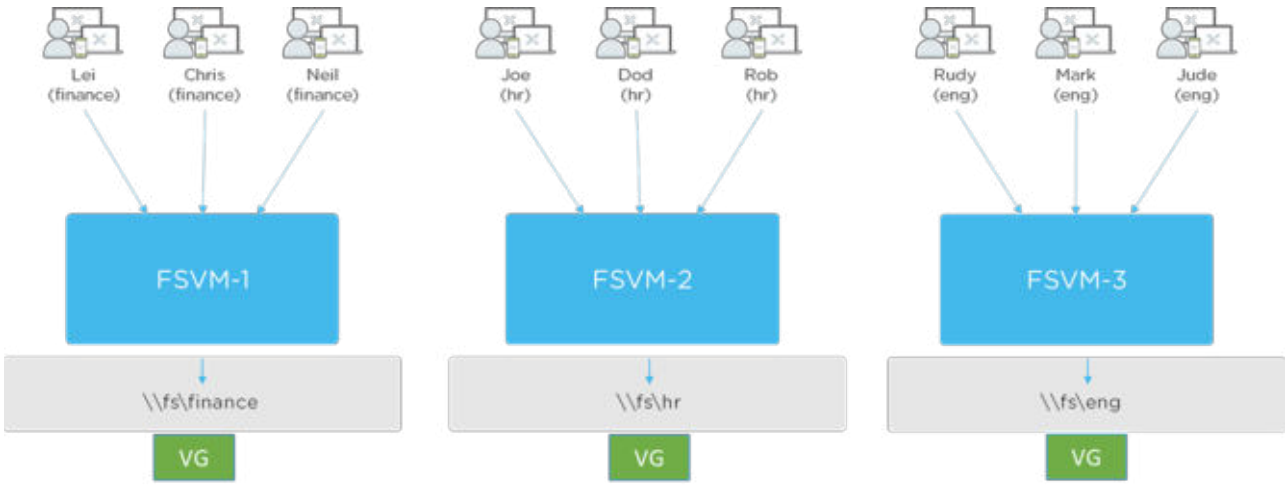

Figure 15: Three Standard Shares

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 deploy 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.

© 2021 Nutanix, Inc. All rights reserved

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. We don't recommend turning on deduplication.

> 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 in a single share and billions of files across a multinode cluster with multiple shares. To achieve good response times for environments with high file and directory counts, give some thought to directory design. Placing millions of files or directories in a single directory slows file enumeration, which may impact some applications.

For performance-sensitive applications, Nutanix recommends that you limit 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 that you limit 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 shouldn't 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.

Load Balancing and Scaling

Initial load balancing for a Files cluster is based on the number of standard shares or the number of top-level directories you have in a distributed share. As explained in the Exports and Shares section, a standard share resides on a single volume group owned by one FSVM at a time. A distributed share is a collection of volume groups containing top-level directories. If some standard

© 2021 Nutanix, Inc. All rights reserved

shares or top-level directories encounter a large workload and the volume groups supporting those workloads share an FSVM, a performance bottleneck can occur. Nutanix Files helps load balance workloads automatically when it detects potential performance problems by redistributing volume groups to different FSVMs for better load balancing across nodes.

Load balancing may occur in the following situations:

1. When an administrator removes an FSVM from the cluster.
2. When the distribution shares or top-level directories become poorly balanced during normal operation because of changing client usage patterns or suboptimal initial placement.
3. When increased user demand necessitates adding a new FSVM and its volume groups are initially empty.

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.

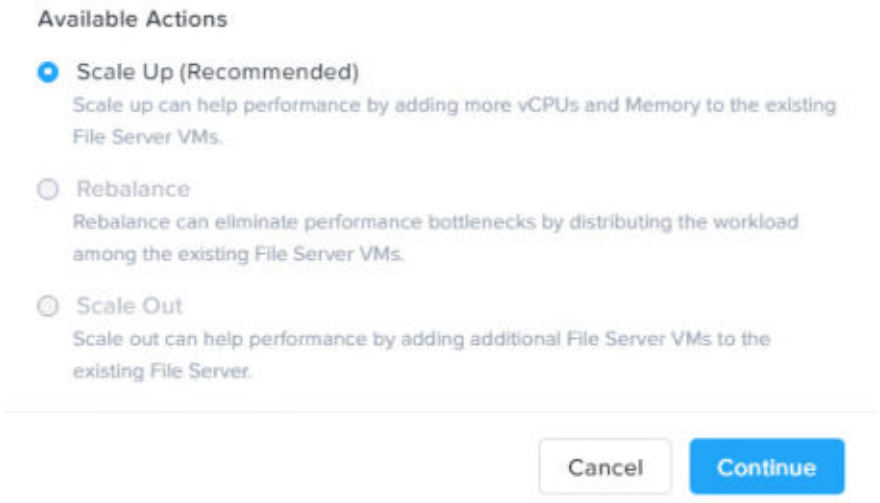

Figure 16: Performance Optimization Recommendation

© 2021 Nutanix, Inc. All rights reserved

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

> Note: We recommend a scale-up operation for performance optimization if SMB connection limits reach 95 percent utilization over a two-hour time window.

There is a brief outage during volume group migration and FSVM scale-out if you aren't using shares with continuous availability enabled. The file share or export requires a client reconnect after migration and scaling out. Today most clients try to reconnect for 50–60 seconds, which limits the overall impact.

Load balancing also occurs for each vDisk in a volume group. Nutanix Volumes requires the administrator to configure an iSCSI data service IP address. The data service IP address is a highly available virtual IP address used as an iSCSI discovery portal. Each vDisk in a volume group represents its own iSCSI target that any 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 the Acropolis Dynamic Scheduler (ADS).

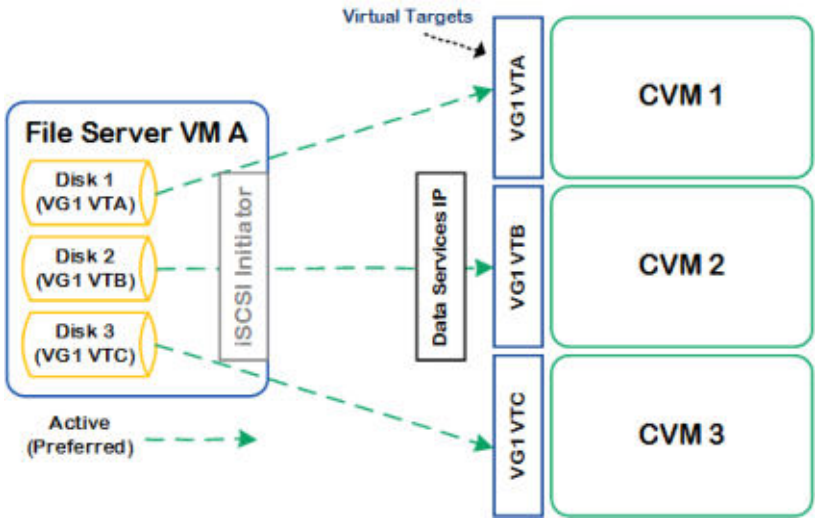

Figure 17: Load Balancing Volume Groups with Nutanix Volumes

© 2021 Nutanix, Inc. All rights reserved

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

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 sign on again. The new authentication occurs through the external data services IP, which redirects the session to a healthy CVM. The following figure shows this general process.

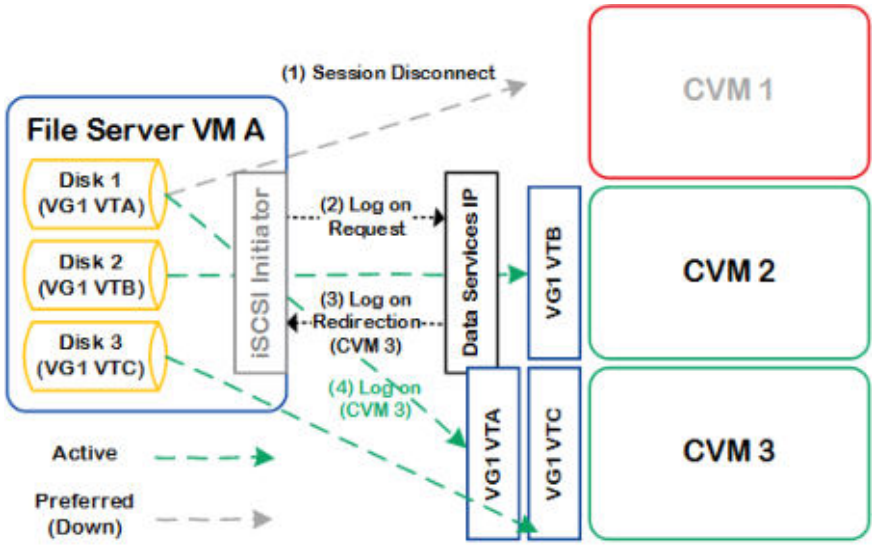

Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups

When the failed CVM returns to operation, the iSCSI session fails back. In the case of a failback, the system signs the FSVM off and redirects it to the appropriate CVM.

© 2021 Nutanix, Inc. All rights reserved

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.

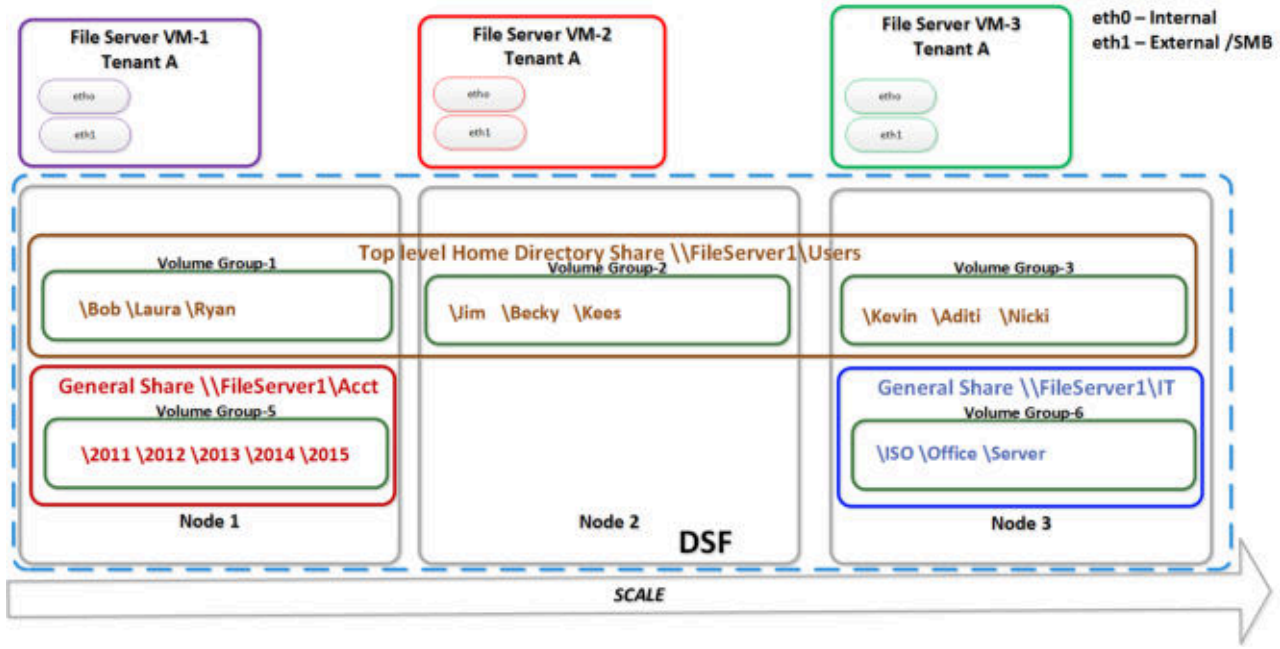

Figure 19: Each FSVM Controls Its Own Volume Groups

When an FSVM is unavailable, the remaining FSVMs volunteer for ownership of the shares and exports 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

© 2021 Nutanix, Inc. All rights reserved

immediately. This arrangement ensures distribution of volume groups even if multiple FSVMs fail.

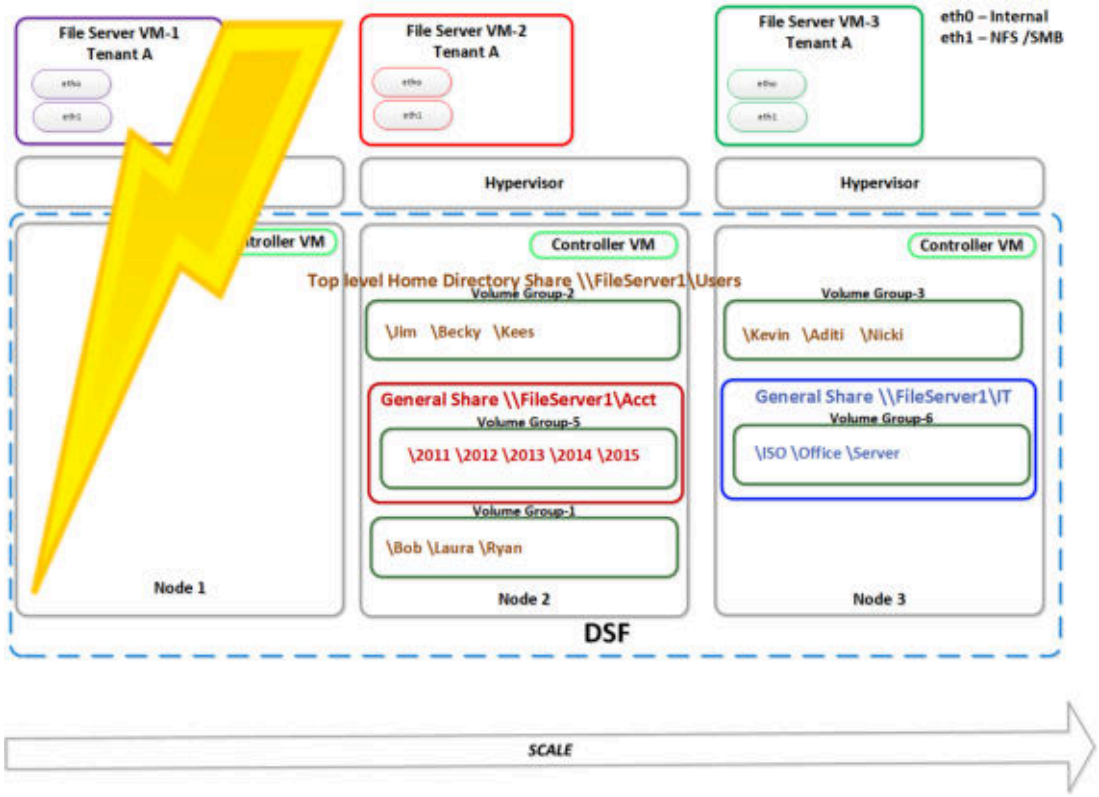

Figure 20: FSVM-1 Failure

The Nutanix Files Zookeeper instance tracks the original FSVM's ownership using the storage IP address (eth0), which doesn't 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.

The failover process typically takes between 30 and 180 seconds, and client operations can be impacted during this time. SMB shares enabled for continuous availability and hard-mounted NFS shares experience I/O delays but maintain connections during high-availability events.

© 2021 Nutanix, Inc. All rights reserved

If an FSVM is offline for 30 minutes or more, the Nutanix Files cluster redistributes core cluster services (such as Zookeeper) owned by the offline node. The cluster needs at least three online FSVMs in order to move services. Moving cluster services allows the Files cluster to rebuild itself smaller and survive subsequent node failures.

When FSVM-1 comes back online and finds its shares and exports locked, it assumes that a high-availability event has occurred. After the grace period expires, FSVM-1 regains control of the volume group through the Minerva CVM service. Once the failed FSVM comes back online the cluster automatically switches back to the larger size.

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

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.

Single-Node Environments

Nutanix Files clusters that only have one FSVM rely on hypervisor-based high availability to maintain services. If the physical node that owns the FSVM fails, the VM restarts on another physical node. Client connections are disrupted until the FSVM restarts.

Active Directory and SMB Operations

Nutanix Files SMB works with Active Directory. To deploy a cluster, you need domain privileges to create the machine account and the DNS entries for DFS referrals. The file server doesn't store these credentials.

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

- Create computer objects.

© 2021 Nutanix, Inc. All rights reserved

- 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 3.7 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.

© 2021 Nutanix, Inc. All rights reserved

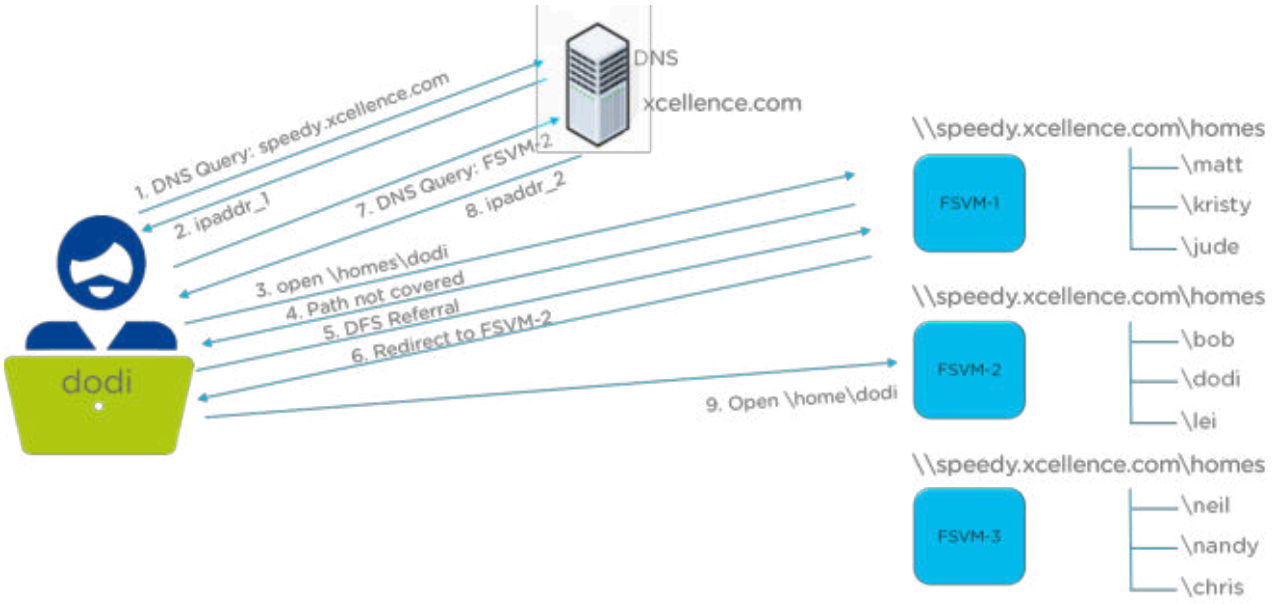

Figure 21: DNS Request for SMB

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 the same way 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. NTFS-

© 2021 Nutanix, Inc. All rights reserved

level permissions, also called Windows access control lists (ACLs) or NTACLs, manage all file and folder access.

The distributed share has some special requirements because of how 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 can't just delete the top-level directory because it's a separate share created internally by Nutanix Files. Because you can't directly delete a top-level directory, removing it requires additional steps.

There are two options for renaming or deleting distributed share folders:

1. Find out which FSVM is hosting the folder and rename or delete the folder directly using the FSVM as the UNC path.
2. Use the Nutanix Files MMC snap-in to manage top-level directories. Any file server administrator can perform the MMC operations; you don't need to assign privileges manually. Establish a connection to the Files namespace and perform the following SMB share management tasks:

   a. Create, delete, rename, and change permissions for top-level directories.
   b. Change NTFS permissions for shares.

> Tip: Use the Nutanix Files MMC when you modify NTFS 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 modify share-level permissions 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. By default, all SMB shares have share-level ACLs set to Everyone with Full Control.

> Note: We recommend leaving share permissions at Full Control and managing access with NTFS permissions.

You can also manage open file locks and user sessions with the Windows Shared Folders MMC.

© 2021 Nutanix, Inc. All rights reserved

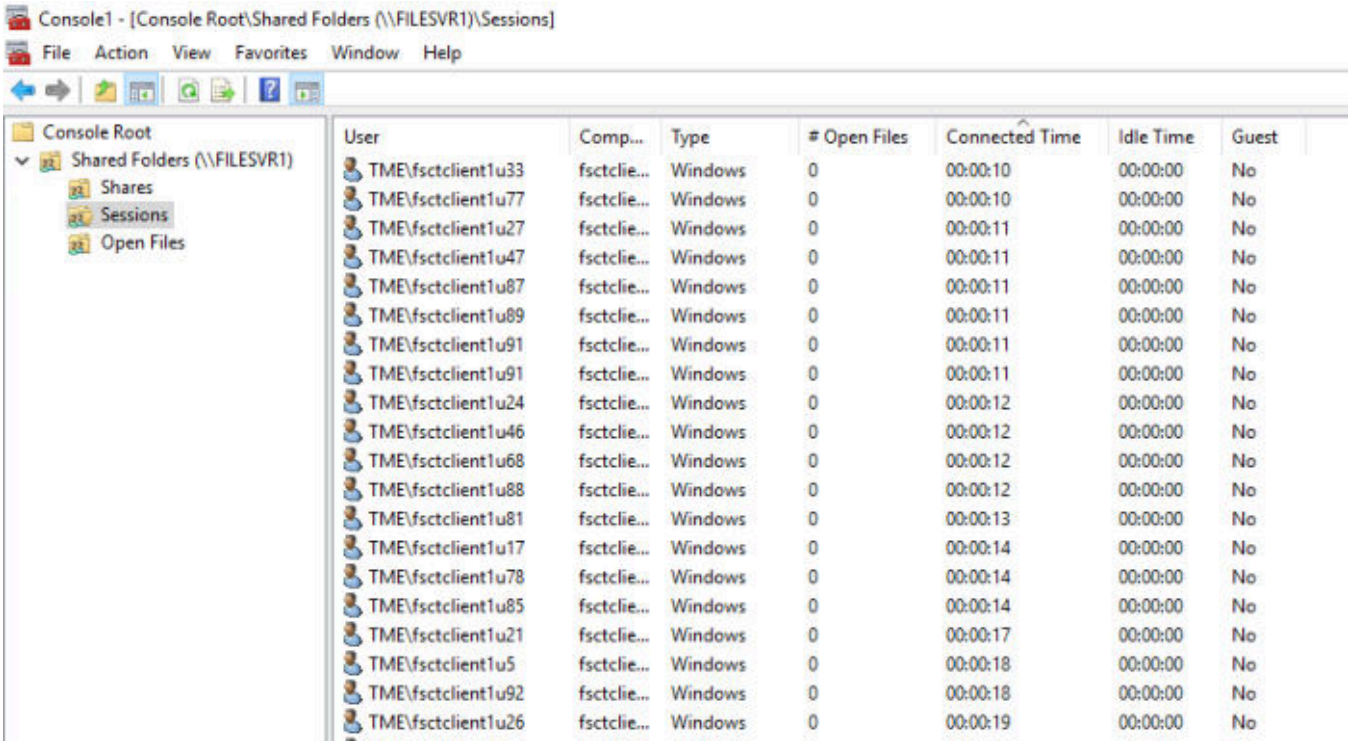

Figure 22: Shared Folders MMC

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 doesn't control security permissions, and running ABE has associated overhead. Every time a user requests a browse operation, ABE must filter out objects the user doesn't have permission for. 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 (using a distributed share) are a great example of where you shouldn't enable ABE. As most users get their home mapping when they sign on and always have access to their own content, we don't recommend enabling ABE for home directory use cases.

© 2021 Nutanix, Inc. All rights reserved

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 eavesdropping 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 introduced 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.

© 2021 Nutanix, Inc. All rights reserved

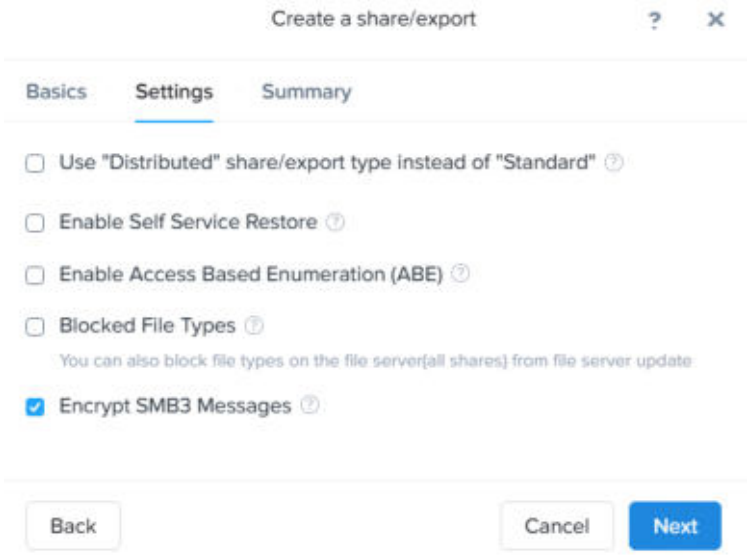

Figure 23: Enable SMB Encryption

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 doesn't maintain durable handles through file server events such as FSVM failures or upgrades.

SMB 3.0 Transparent Failover

SMB 3.0 Transparent Failover is a feature of the SMB 3.0 protocol that enables fully nondisruptive operations against SMB shares. Transparent failover is often called continuously available file shares. File server–side failures or upgrade events persist file handles so that clients transparently reconnect to another

© 2021 Nutanix, Inc. All rights reserved

FSVM without impacting applications. Nutanix Files versions 3.7.1 and later support SMB 3.0 Transparent Failover. Continuously available shares are intended for specific use cases, like Citrix App Layering and FSLogix, that require nondisruptive operations.

> Note: SMB Transparent Failover enforces synchronous write operations, which may decrease the maximum possible performance for SMB shares.

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 when you use Nutanix Files as a member server and Nutanix Files shares as folder targets in the namespace. When you use 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 the PeerGFS and Nutanix Files datasheet.

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.

© 2021 Nutanix, Inc. All rights reserved

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 doesn't 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 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.

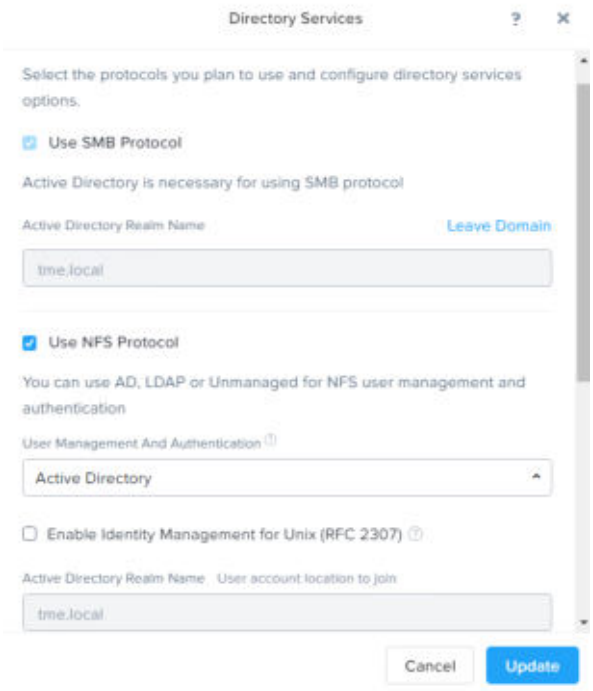

Figure 24: Nutanix Files User Management with NFS

© 2021 Nutanix, Inc. All rights reserved

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

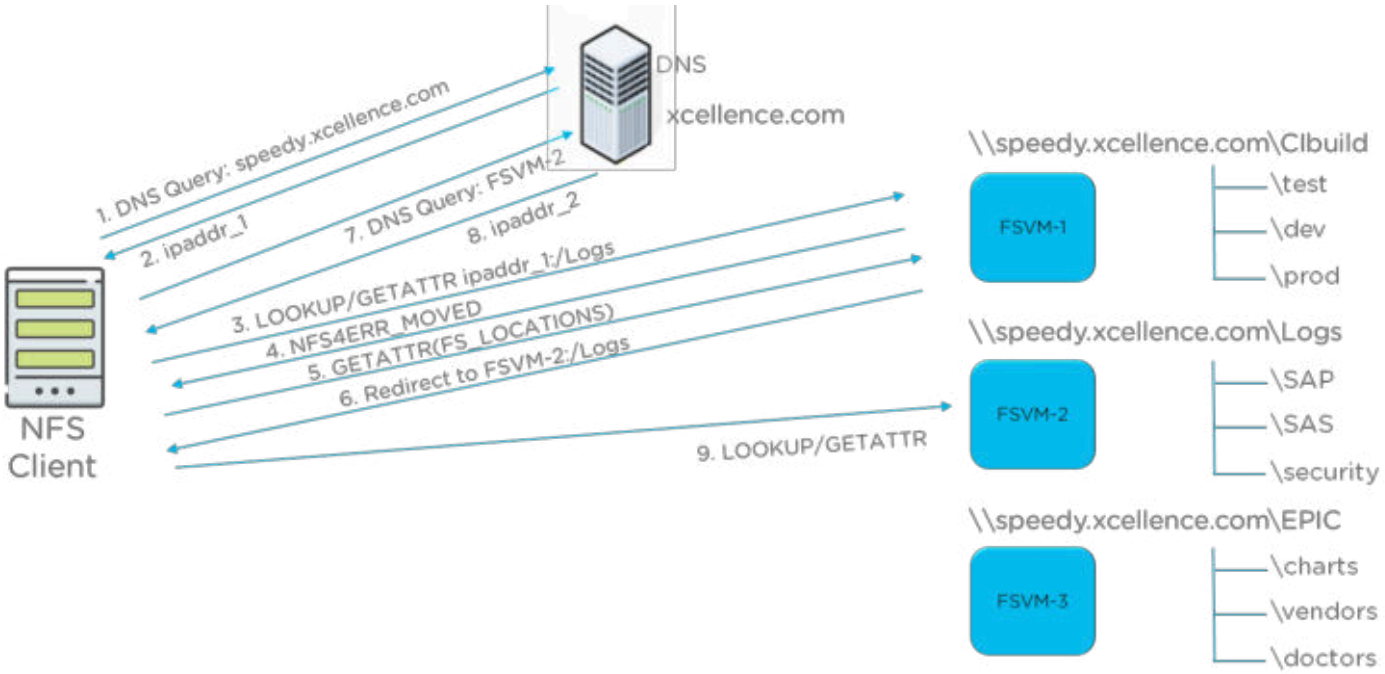

Figure 25: DNS Request for NFSv4

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.

© 2021 Nutanix, Inc. All rights reserved

NFSv3

Nutanix Files supports NFSv3. Each file server has NFSv3 enabled by default, 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 doesn't 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 doesn't support DFS-like referrals. Clients aren't 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.

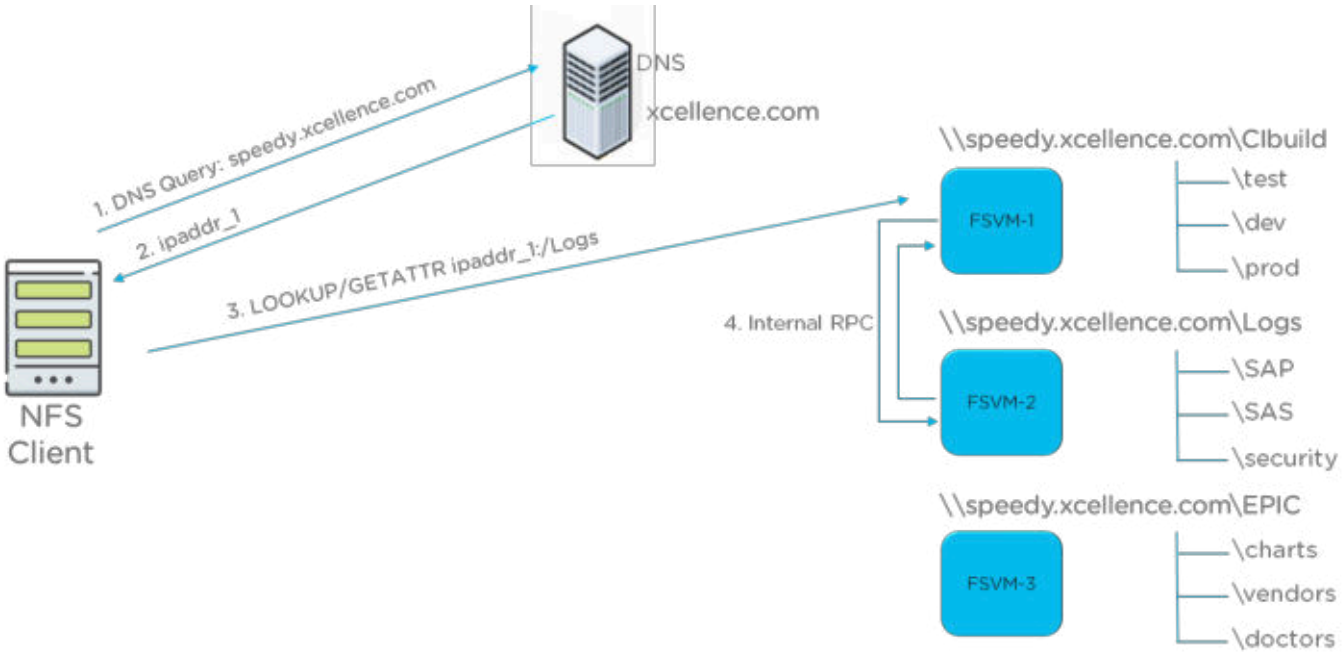

Figure 26: DNS Request for NFSv3

© 2021 Nutanix, Inc. All rights reserved

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 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).

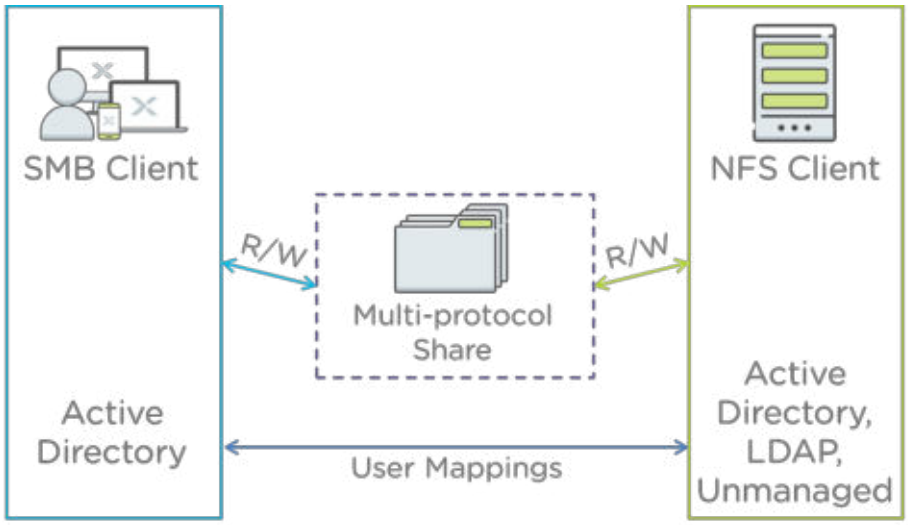

Figure 27: Multiprotocol Shares

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.

© 2021 Nutanix, Inc. All rights reserved

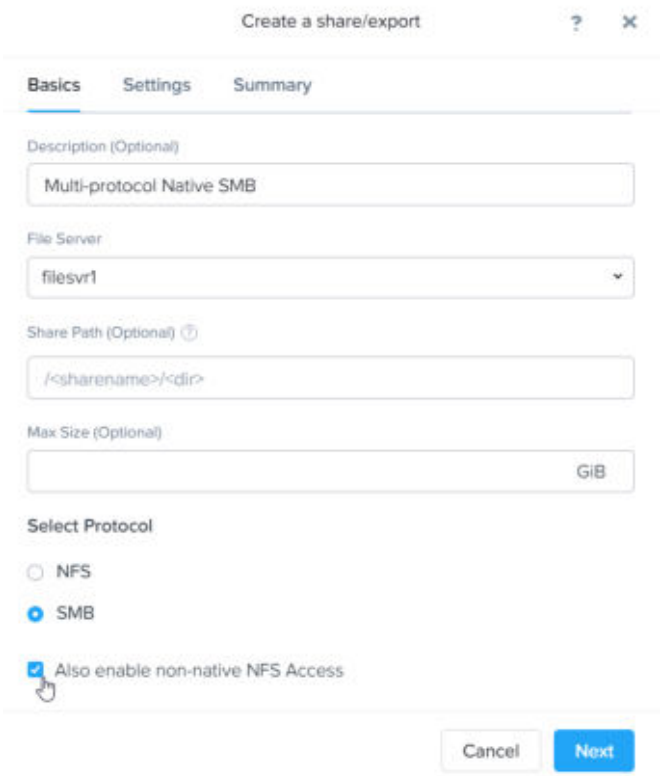

Figure 28: Enable Non-Native Protocol Access

You manage all access control for a multiprotocol share using the native protocol. For SMB, the native protocol is Windows 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.

© 2021 Nutanix, Inc. All rights reserved

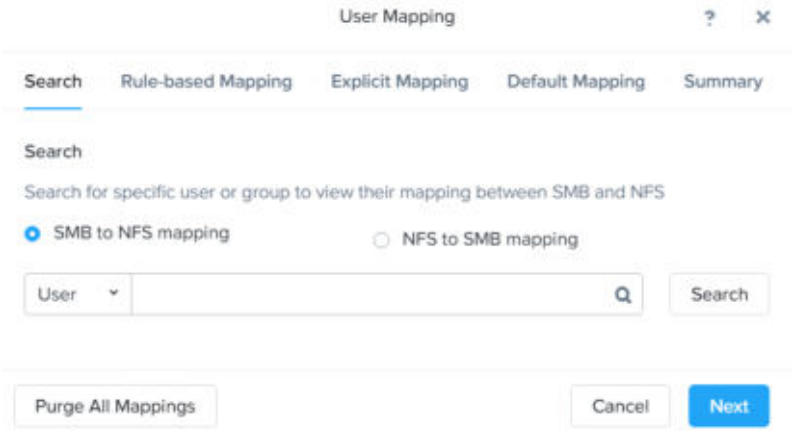

Figure 29: Multiprotocol User Mapping

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 Active Directory 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)

© 2021 Nutanix, Inc. All rights reserved

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

* Kerberos 5, 5i, and 5p

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 can't 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.

© 2021 Nutanix, Inc. All rights reserved

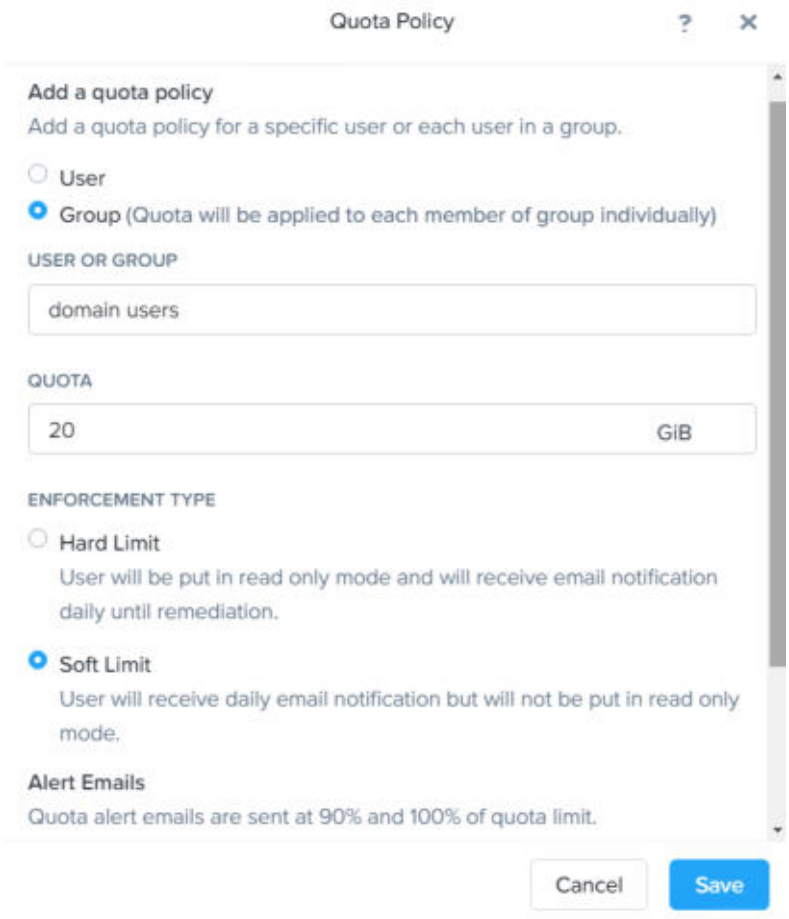

Figure 30: Setting a Share Quota

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 know they may need to take action. With the Files 4.0 release you can customize the email template used for notification.

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

© 2021 Nutanix, Inc. All rights reserved

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. 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 they resolve the issue.

Selective File Blocking

You can define a list of file names and file extensions to block from storage 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. These entries can include wildcards for both the file name and file extension. For example, you can block the file pattern encrypt *.* xt.

© 2021 Nutanix, Inc. All rights reserved

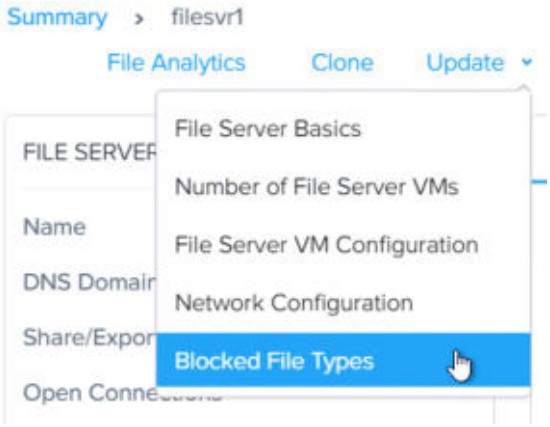

Figure 31: Blocked File Types for the File Server

A comma-separated list of file names and 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.

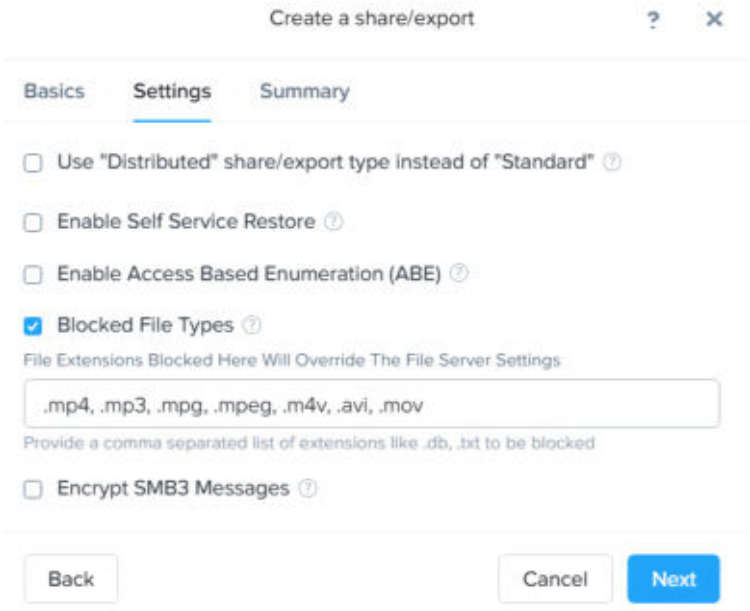

Figure 32: Blocked File Types for a Share

When you attempt to create a file with or rename a file to a blocked pattern, an access denied message appears. You can read, edit, or delete existing files

© 2021 Nutanix, Inc. All rights reserved

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 you apply the policy. With version 3.0 or later, File Analytics can also apply a list of file patterns matching known ransomware variants.

> Note: When you use Data Lens to manage your ransomware file name patterns, you can't see the patterns from within the file server or share specific blocking lists.

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 for these APIs is forwarding such events to a syslog server for retention and audit trails. While logging audit trails is useful, we needed to simplify insights into this data, so 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: 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 form a repository of information so administrators can review what operations have occurred against specific data and by specific users. File Analytics analyzes logged events 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.

© 2021 Nutanix, Inc. All rights reserved

- 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 in 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.

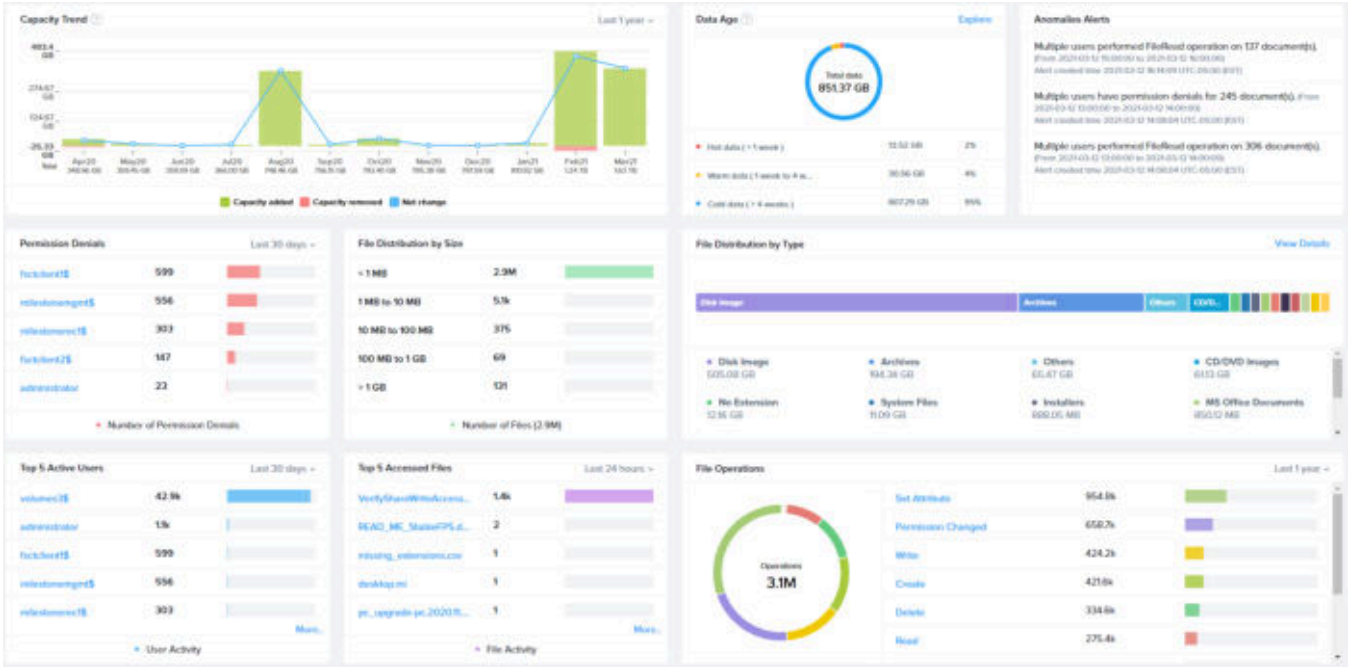

Figure 33: File Analytics Dashboard

© 2021 Nutanix, Inc. All rights reserved

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.

Data Age Analytics

Starting with version 3.0, File Analytics added an extension called data age analytics to the data age widget on the dashboard. Data age analytics shows you how frequently users access your data over time ranges you define. Once you customize the data age ranges to match your requirements, data age analytics shows you the hot, warm, or cold data trends over the specified timeframe, such as the last week, month, or longer interval. You can also see the growth percentage of a given category in your designated range.

Figure 34: Data Age Analytics

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.

© 2021 Nutanix, Inc. All rights reserved

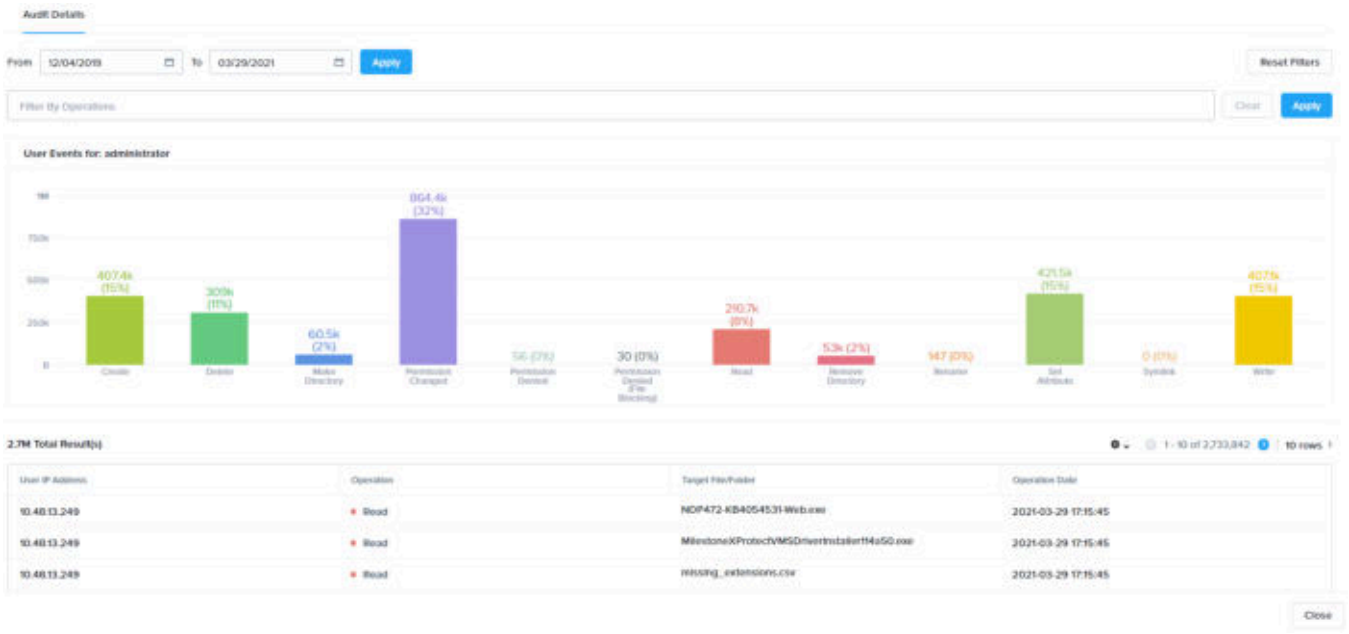

Figure 35: File Analytics Audit Trails

The ability 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 file for further reporting activities.

Anomaly Detection

File Analytics allows you to define anomaly alerts that represent specific operations as run by an individual or against the file server as a whole.

© 2021 Nutanix, Inc. All rights reserved

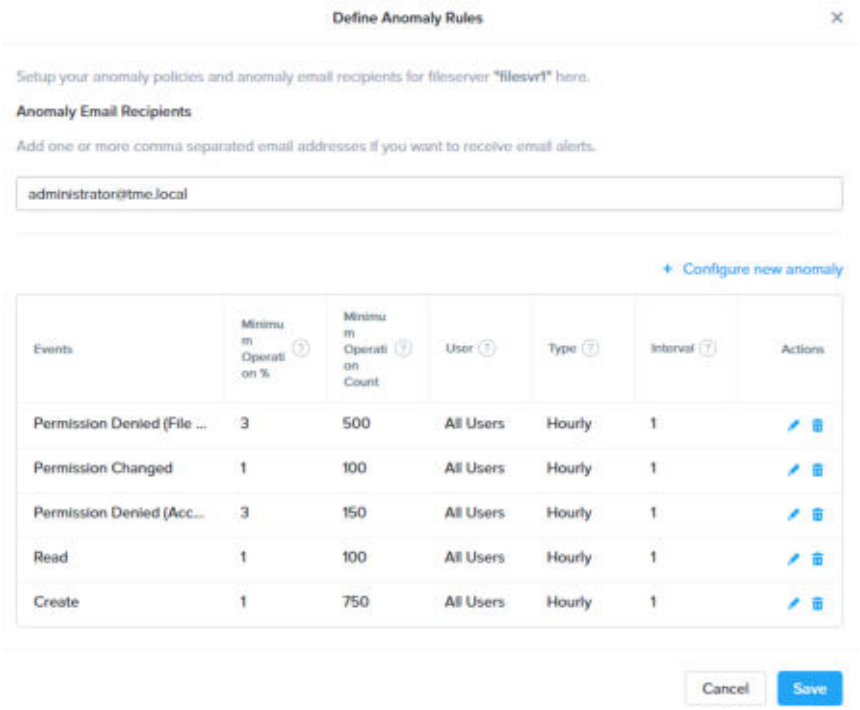

Figure 36: File Analytics Anomaly Detection

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.

© 2021 Nutanix, Inc. All rights reserved

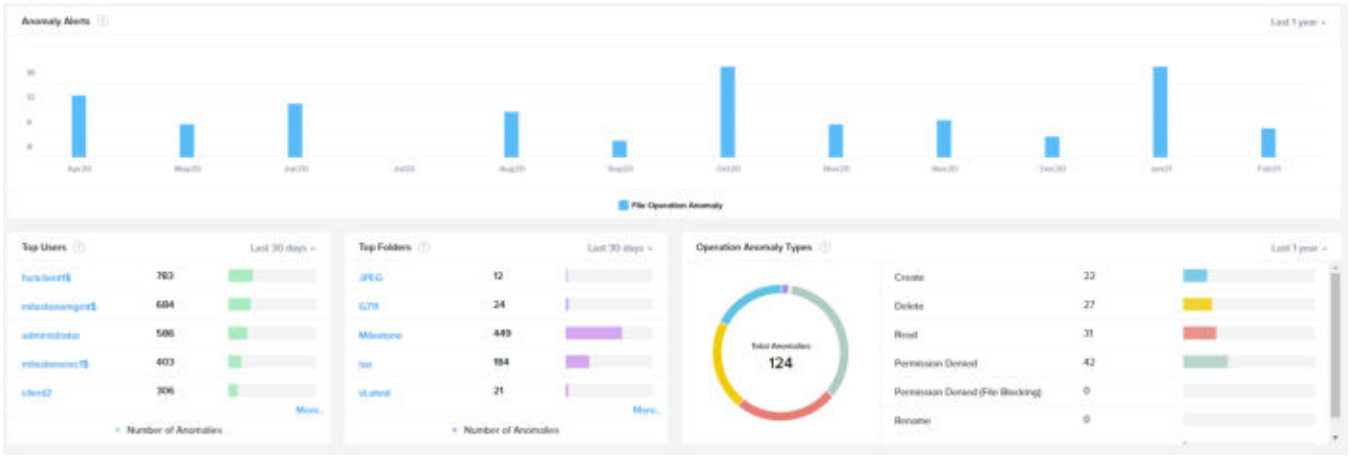

Figure 37: File Analytics Usage Anomalies Dashboard

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.

Ransomware Intelligence

Ransomware is a persistent concern that requires multiple security controls and software layers to mitigate. Nutanix Files has long supported centralized antivirus scanning through ICAP (Internet Content Adaptation Protocol) with security vendors like Trend Micro, McAfee, BitDefender, and Symantec. Nutanix offers a comprehensive approach to ransomware across our portfolio of products that you can read about in more detail in our Ransomware Threat solution brief. To seriously respond to ransomware, a solution should prevent the infiltration of ransomware and malware, detect any infection attempts, alert the organization and initiate defensive measures, and, if the worst happens, provide a comprehensive strategy for recovery.

Nutanix Files and File Analytics have many of the core features required to help detect, protect, analyze, and recover from ransomware. File Analytics 3.0 begins the journey to combine these technologies into a comprehensive interface, built so you can manage your ransomware strategy with Nutanix Files. It starts with a dedicated ransomware dashboard in File Analytics that summarizes any detected vulnerabilities, including the impacted shares and

© 2021 Nutanix, Inc. All rights reserved

clients that may be compromised. The dashboard also shows you whether your shares are protected with SSR snapshots. You can enable SSR against the unprotected shares from this interface.

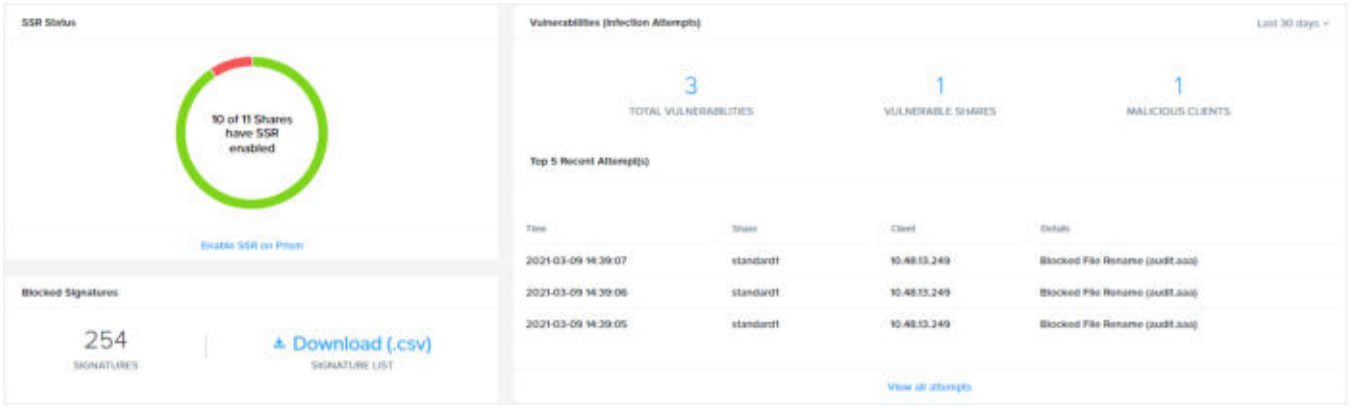

Figure 38: File Analytics Ransomware Page

Nutanix Files can block file creation or file rename operations for specific file extensions. Files 3.8 extended this feature to support wildcards with file names and file extensions. When you enable ransomware protection in File Analytics, Files automatically adds the names and extensions of known ransomware variants to the blocking list. If anyone attempts any file creation or file rename event involving these blocked file types, the event appears as a vulnerability that Files reports on the dashboard and emails to the designated users.

Custom Reporting

File Analytics captures real-time user audit data and file metadata for Nutanix Files environments. The 3.0 release of File Analytics enabled you to mine this data more effectively by introducing a custom reporting page, so you no longer needed to rely on the information available on the dashboard or through audit search.

When creating a new report, you first select an entity, such as files, folders, audit events, or users. You can then choose which attributes to filter with customizable values. For example, you can choose attributes like file size greater than or less than a given value, file age based on access time or creation date, or all audit events within a given time range. Finally, you can

© 2021 Nutanix, Inc. All rights reserved

choose which columns of data to include in the report, like user and file names, paths, clients, operation types, and other details associated with the entity.

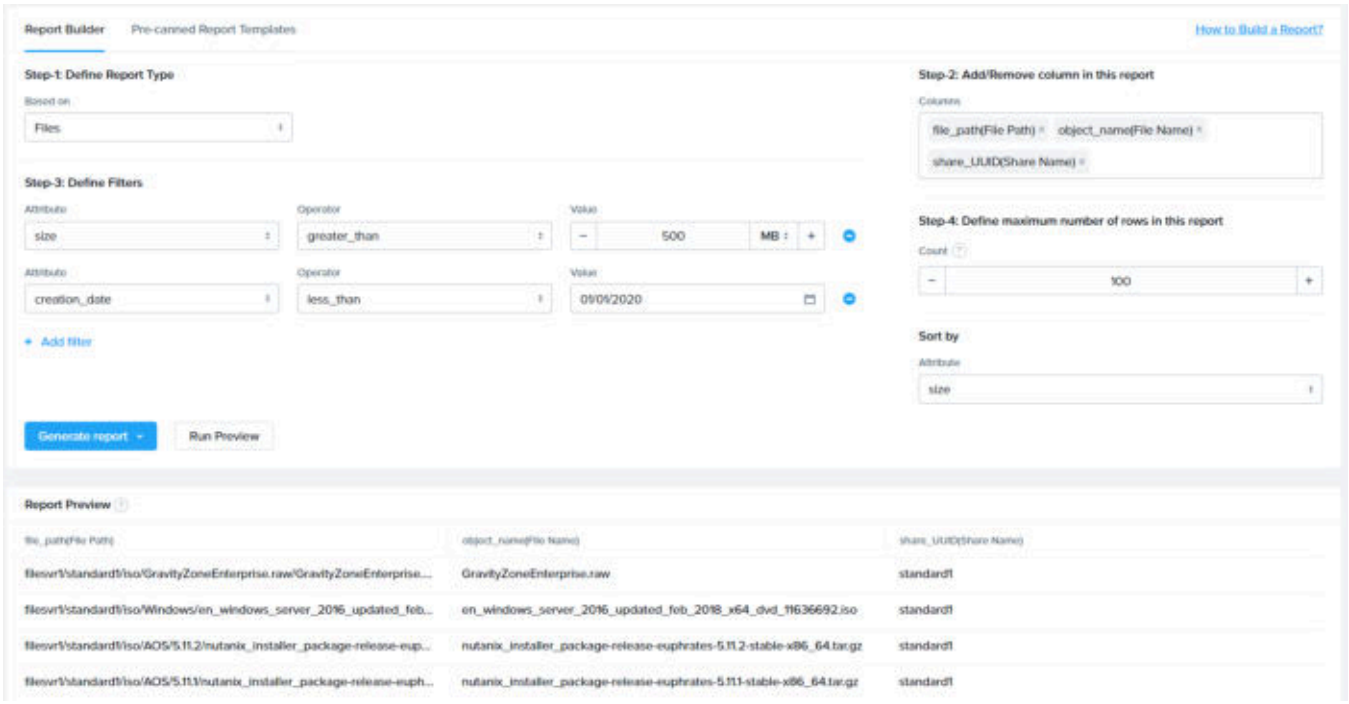

Figure 39: Report Builder

Once you define the report, you can save it and run it again at any time.

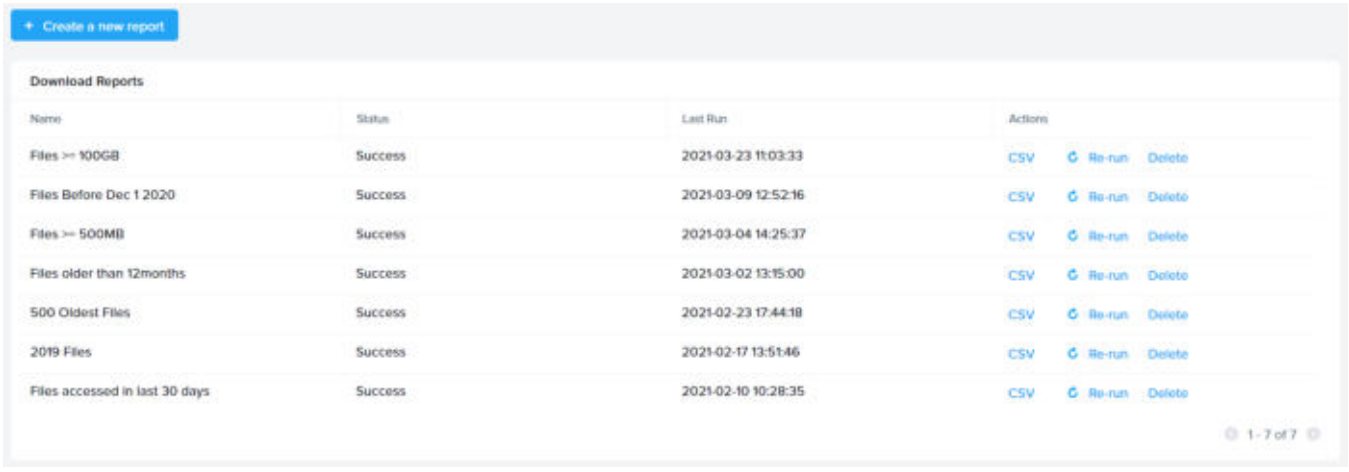

Figure 40: Saved Reports

© 2021 Nutanix, Inc. All rights reserved

Smart Tier

Administrators with large scale Network Attached Storage (NAS) environments frequently look for ways to improve storage efficiency, simplify administration, and decrease cost. With Nutanix Files 4.0, we added a native tiering framework called Smart Tier to address these requirements. Smart Tier enables you to move cold or rarely used data to lower-cost storage while maintaining a single namespace. You can tier data to any qualified S3 API compliant target, including on-premises solutions or the public cloud. Smart Tier enables hybrid multicloud environments to handle unstructured data to retain long term, save money, or provide a virtually limitless pool of local storage for your Nutanix Files workloads.

Smart Tier Architecture

The Nutanix Files 4.0 tiering engine uses APIs over HTTPS to accept tiering and recall requests and to move data to S3-compliant targets. Three targets have been validated for the initial release: Nutanix Objects, Amazon S3 Standard and IA tiers, and Wasabi Cloud Storage. When you tier files, Nutanix Files maintains the metadata for the file while it moves the data to the S3 target.

© 2021 Nutanix, Inc. All rights reserved

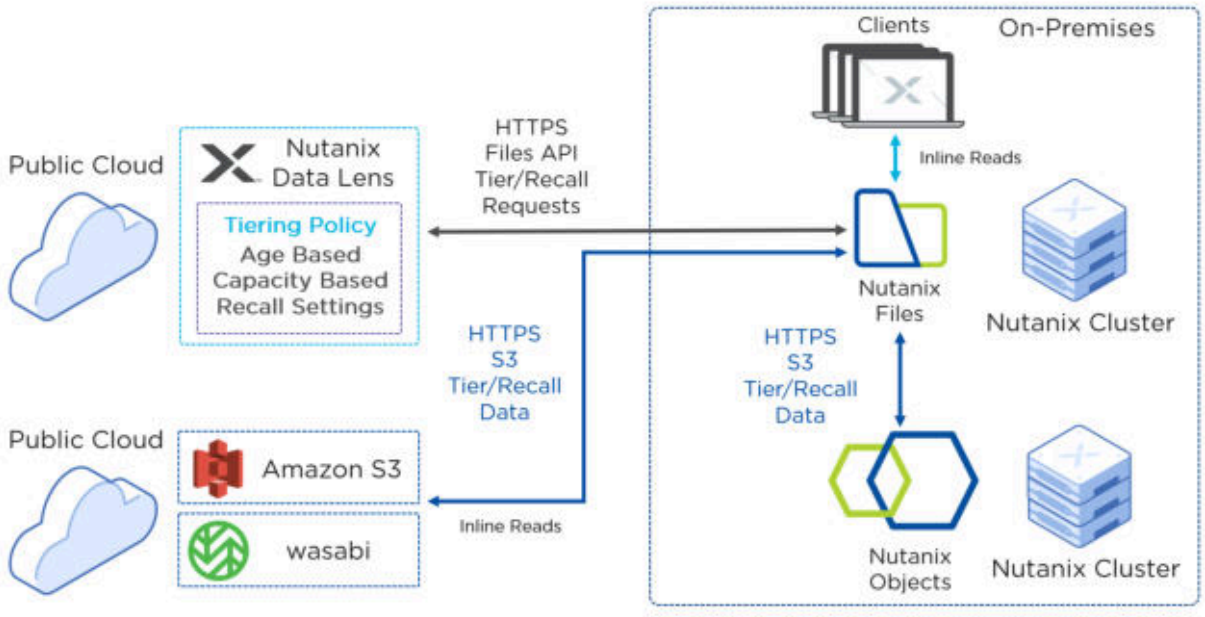

Figure 41: Smart Tier Architecture

Smart Tier maintains a file stub in the SMB or NFS share path where clients can perform inline reads to access the data. You can also recall files automatically, based on access patterns, or with manual recall operations. Configure tiering policies through Nutanix Data Lens, which allows you to set age- and capacity-based thresholds, manual or automatic tiering schedules, and recall settings.

Nutanix Data Lens

Data Lens is a Nutanix software as a service (SaaS) that provides file analytics and reporting, anomaly detection, audit trails, ransomware protection features, and tiering management for your entire Nutanix Files environment.

© 2021 Nutanix, Inc. All rights reserved

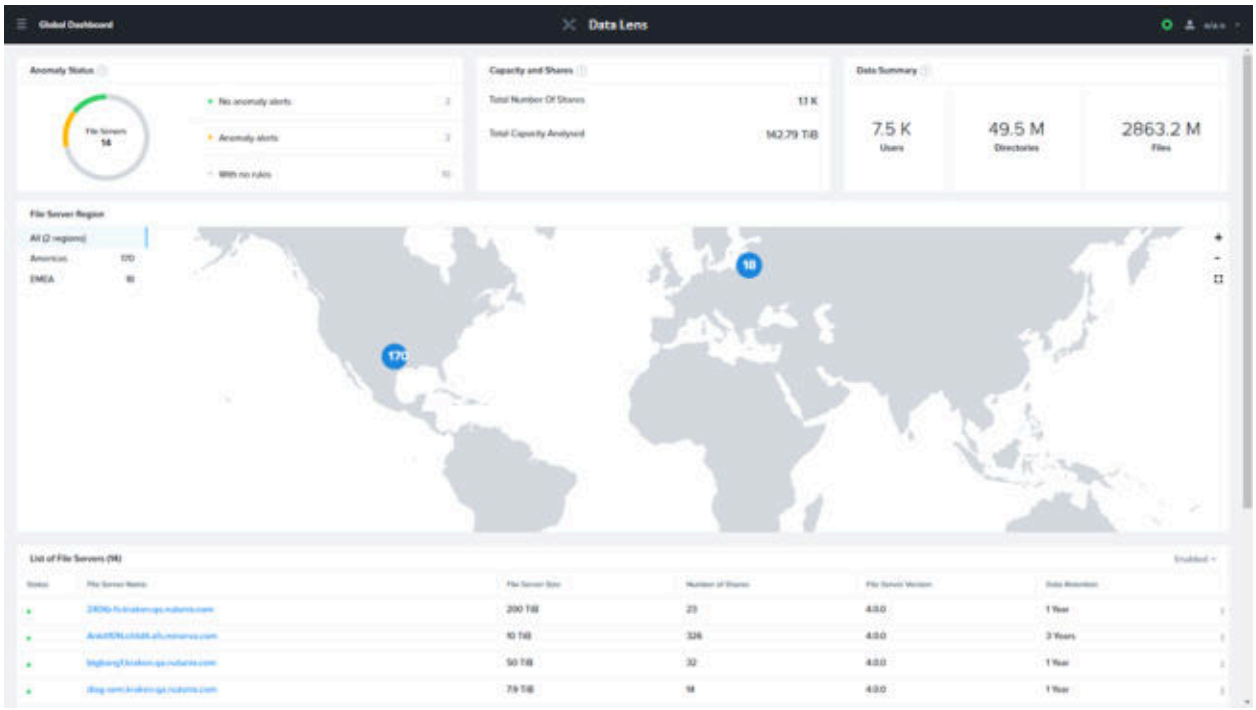

Figure 42: Data Lens Dashboard

Data Lens tracks the access patterns for all shares and exports to determine the age of each file for a given file server instance. You can define hot, warm, and cold data age categories to match your requirements. The data age explorer lets you view the file age on a share-by-share basis to understand what data is included in a given tiering policy.

© 2021 Nutanix, Inc. All rights reserved

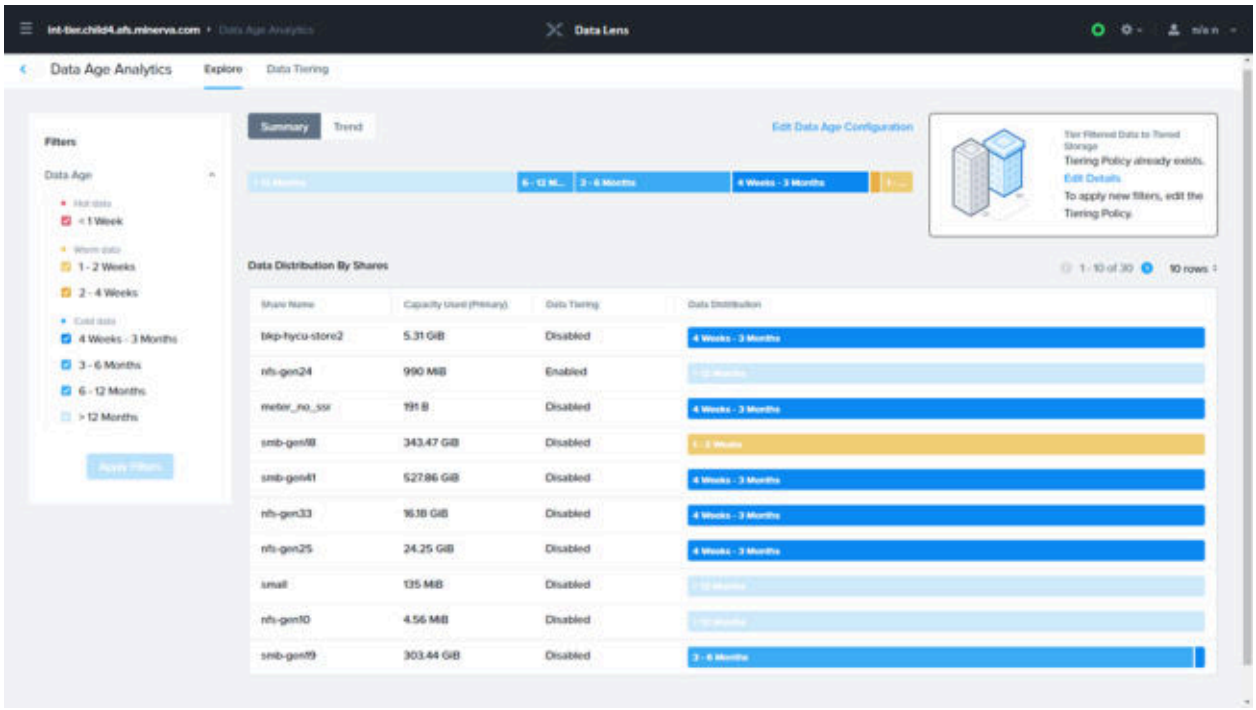

Figure 43: Data Lens Age Explorer

© 2021 Nutanix, Inc. All rights reserved

Smart Tier Configuration

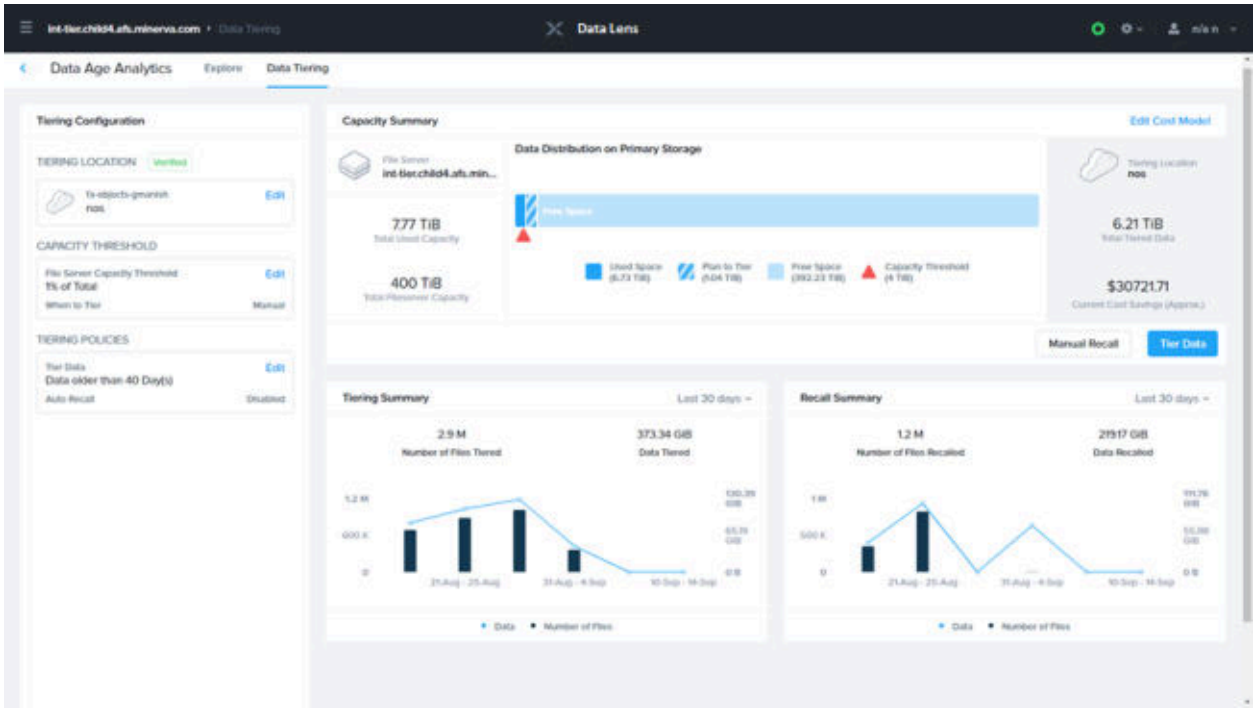

Figure 44: Data Lens Tiering

To configure Smart Tier, define a tiering location (a supported S3-compliant target). Configure the following:

- Target URL over HTTPS
- Bucket name
- Access and secret keys
- Retention period
- Certificate (optional)

The retention period determines when Nutanix Files deletes a tiered object after removing the stub from the file system.

After you configure your location, define a capacity threshold and when-to-tier policy. The capacity threshold represents the percent of allocated space consumed before the system considers tiering. You also define whether the

© 2021 Nutanix, Inc. All rights reserved

system performs the tiering policy manually or on an automated basis, and the time windows in which automated tiering occurs.

The third step is to define the formal tiering policy, including:

- Age of the data (based on the last file read or write operation)
- Minimum file size (64 KB or larger)
- Shares to exclude from the policy
- Automatic recall settings

The automatic recall setting lets you define whether to recall files automatically and under what conditions, specifically the number of times the file is accessed over a given period. If you choose to recall files manually, Data Lens lets you search for and choose individual files, folder paths, or shares to recall.

Tiered Data

When data reaches the configured age and the file server is at its capacity threshold, Smart Tier moves file data to the target. Smart Tier marks the previously consumed space as free, creating additional space in the share for new or recalled data.

Tiered files appear as regular files in their folder paths. SMB shares have an offline attribute set, which in Windows Explorer shows an X beside the file icon.

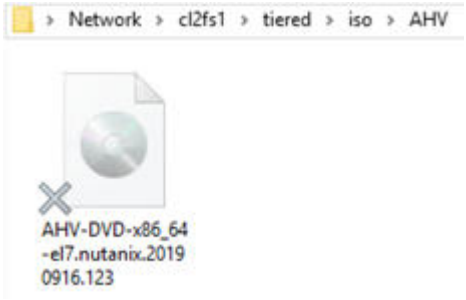

Figure 45: Tiered File

You can see the file size (size on disk), which represents just the metadata, and the actual file size, which represents the tiered data. You can read (which here means retrieve the data inline from) tiered files at any time, but you can't write

© 2021 Nutanix, Inc. All rights reserved

to a tiered file directly, so you need to either retrieve the file or make a separate copy to edit it.

Through Data Lens you can see the hot, warm, cold, and archived data associated with tiering. You can also define cost thresholds that help you see the current and potential cost savings based on your tiering policies.

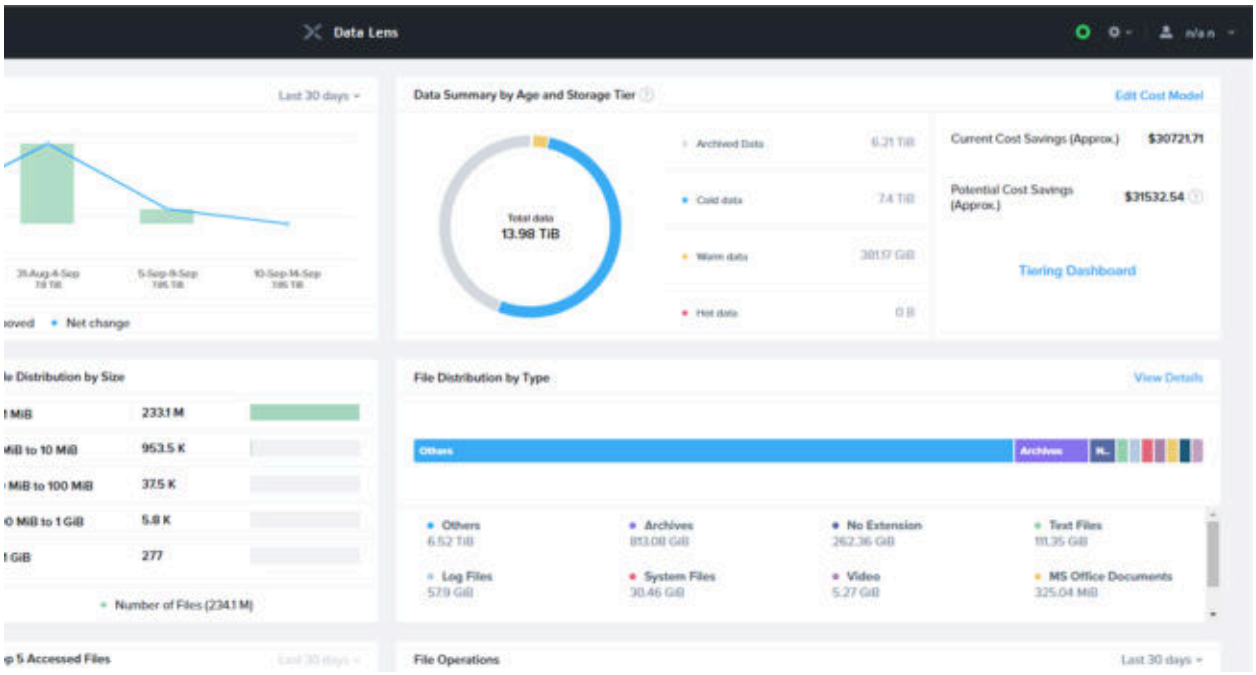

Figure 46: Data Summary by Age and Storage Tier

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

© 2021 Nutanix, Inc. All rights reserved

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
```

© 2021 Nutanix, Inc. All rights reserved

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.

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.

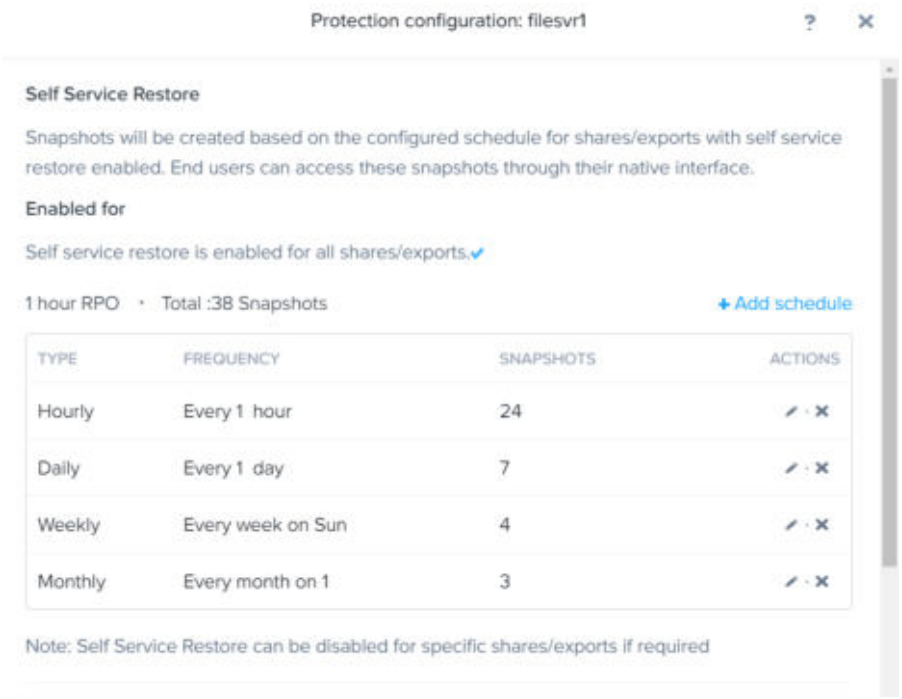

Figure 47: Protection Using SSR

SSR allows you to view snapshots of shares while the share is in use. The share snapshots are read-only, point-in-time copies.

© 2021 Nutanix, Inc. All rights reserved

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.

You can change schedule frequency 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.

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. Use AOS node sizing guidelines when you determine your RPO requirements.

When you create a file server, Prism automatically sets up a corresponding protection domain, which it annotates with the Nutanix Files cluster name.

© 2021 Nutanix, Inc. All rights reserved

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.

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.

Cloning

Because Nutanix Files cloning doesn't affect the original Files cluster, it offers improved support for a variety of use cases:

- Backups at the primary and secondary sites.
- Disaster recovery 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.

© 2021 Nutanix, Inc. All rights reserved

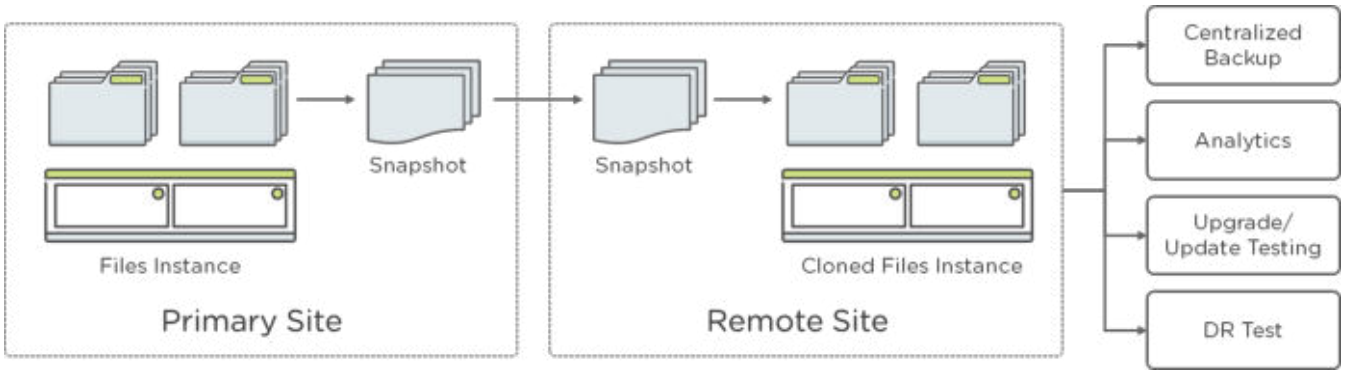

Figure 48: Nutanix Files Cloning Use Cases

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.

Files Smart Disaster Recovery

Nutanix Files long relied on the core snapshot and remote replication capabilities of Nutanix AOS software. Using the core Nutanix software had several benefits, including simplified and consolidated administration that aligned with the core hyperconverged infrastructure (HCI) environment. But there were also some drawbacks around granularity, node density, and the active-passive nature of share access and failover orchestration. Files Smart DR helps address these challenges while maintaining the same simple and consolidated administrative experience.

Prism Central Integration

Files Smart DR moves remote replication management and orchestration into Prism Central. With recent releases of Prism Central, Nutanix Files has supported an integrated service called Files Manager. The first release of Files Manager discovers all Files instances running on clusters registered in Prism Central and provides views on all alerts and events across your file server farm. You can also view the file server configurations and launch Prism Element to manage the file servers.

© 2021 Nutanix, Inc. All rights reserved

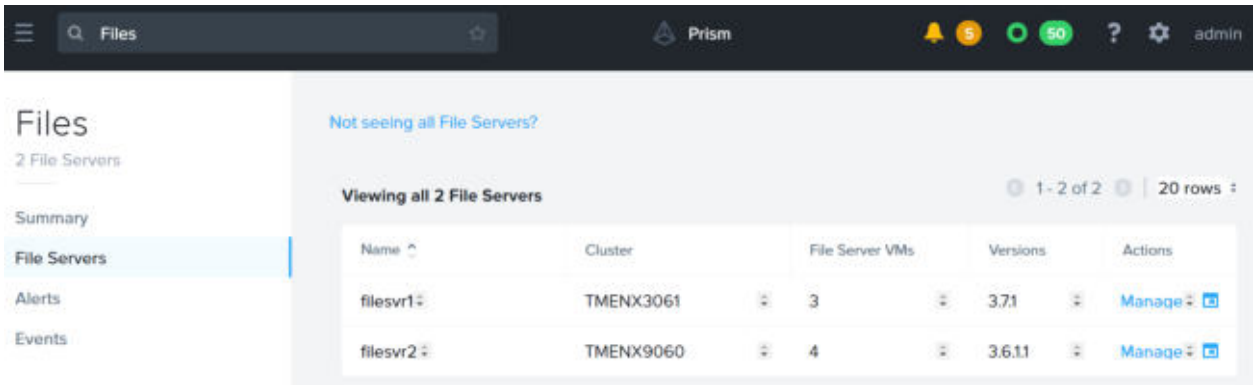

Figure 49: File Manager in Prism Central

Files Manager 2.0 provides a data protection menu where you can configure, monitor, and orchestrate failover and failback operations for Smart DR.

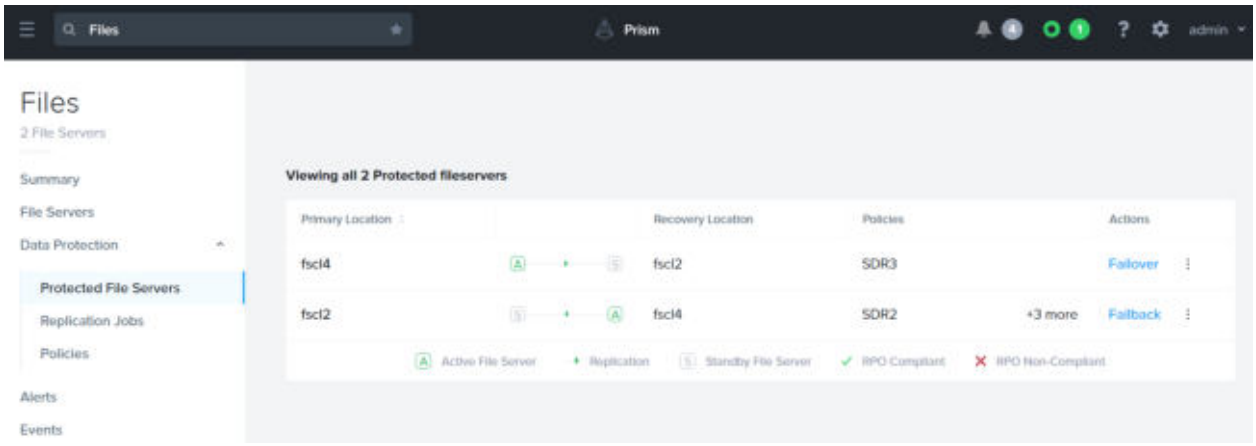

Figure 50: Files Manager Data Protection

Files Smart DR Architecture

In addition to moving management into Prism Central, Smart DR changes several key areas of Nutanix Files remote replication. First, the AOS clusters no longer manage the replication engine itself using native Nutanix protection domains; instead, Nutanix Files manages replication directly. As with Files Self-Service Restore (SSR), Smart DR takes snapshots at the share (file system) level, replicating block-level incremental changes between source and target.

Replication occurs between active file servers running on their respective Nutanix clusters. Shares that function as replication targets are available in a

© 2021 Nutanix, Inc. All rights reserved

read-only state. Replicating between active file servers helps shorten failover times, reduce your recovery time objective (RTO), and simplify use of the replicated data for use cases like backup consolidation or reporting.

Because Files now manages replication, node density limits specific to AOS snapshots and replication no longer apply. You can now use our most storage-dense nodes, which support up to 350 TB of hybrid storage, with the benefits of native remote replication.

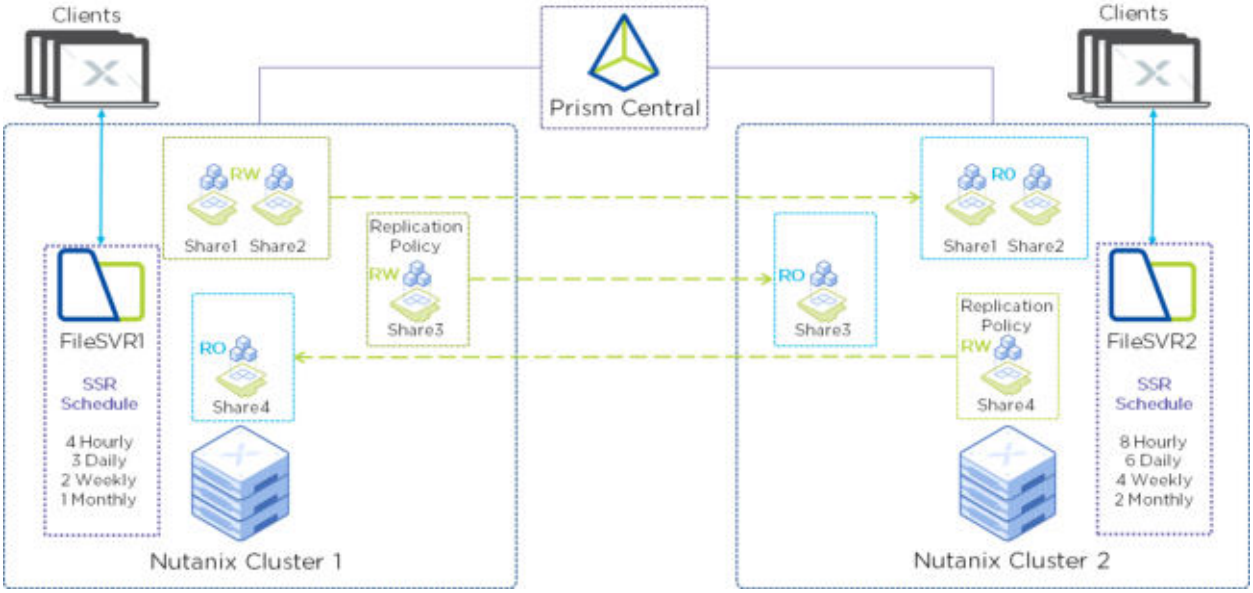

Figure 51: Files Smart DR

Further, for the file systems that support the shares performing replication, you can now set policies on a share-by-share basis to manage your recovery point objective (RPO) at the share level instead of at the file server level.

Replication Policies

A replication policy defines the share or group of shares you want to replicate. The policy also defines the source and target file servers and the replication frequency. You can create multiple policies as needed and specify a default policy for any newly created shares.

© 2021 Nutanix, Inc. All rights reserved

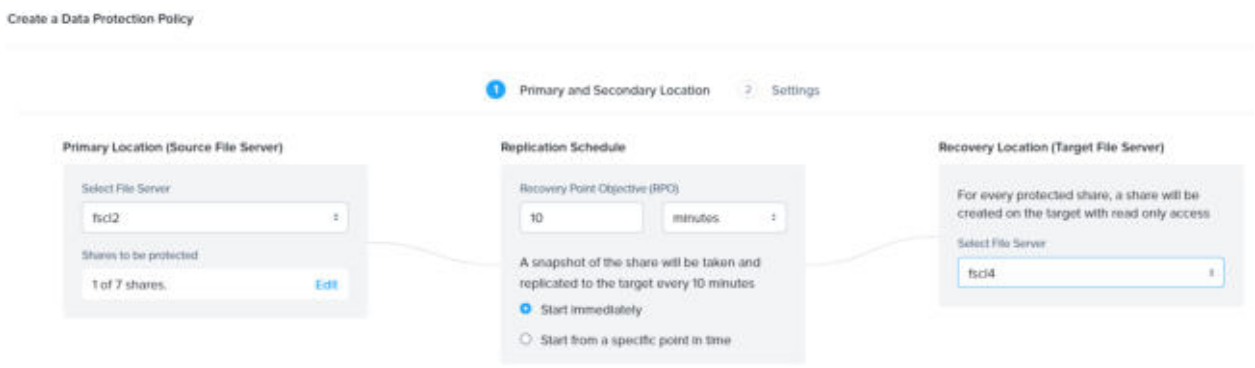

Figure 52: Data Protection Policy

You can configure your replication schedule to be as short as one minute.

Monitoring

Files Manager in Prism Central provides SLA monitoring and job replication status at several levels. You can use the RPO compliance overview provided on the summary page to quickly understand if replication is keeping up with the defined policy.

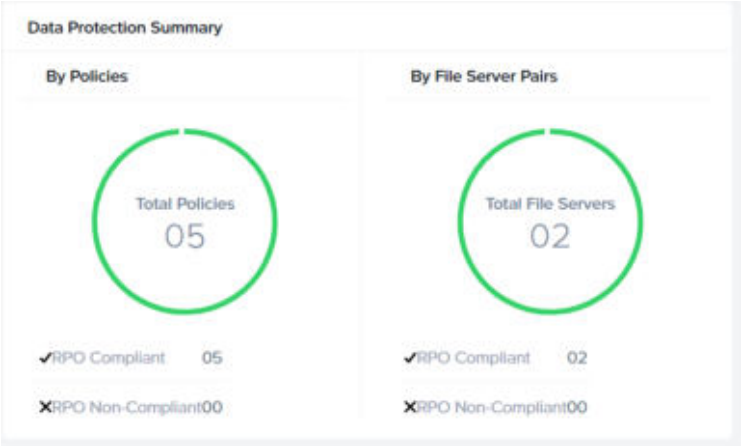

Figure 53: Data Protection Compliance

You can also view each replication job to monitor its completion percentage, start and end times, amount of data synchronized, and average network bandwidth usage.

© 2021 Nutanix, Inc. All rights reserved

Failover and Failback

You can manage planned or unplanned failover and failback from Prism Central. During failover operations, Prism Central orchestrates the required updates to the DNS and Active Directory service principal names (SPNs) to move the file server instance name from the source to the target.

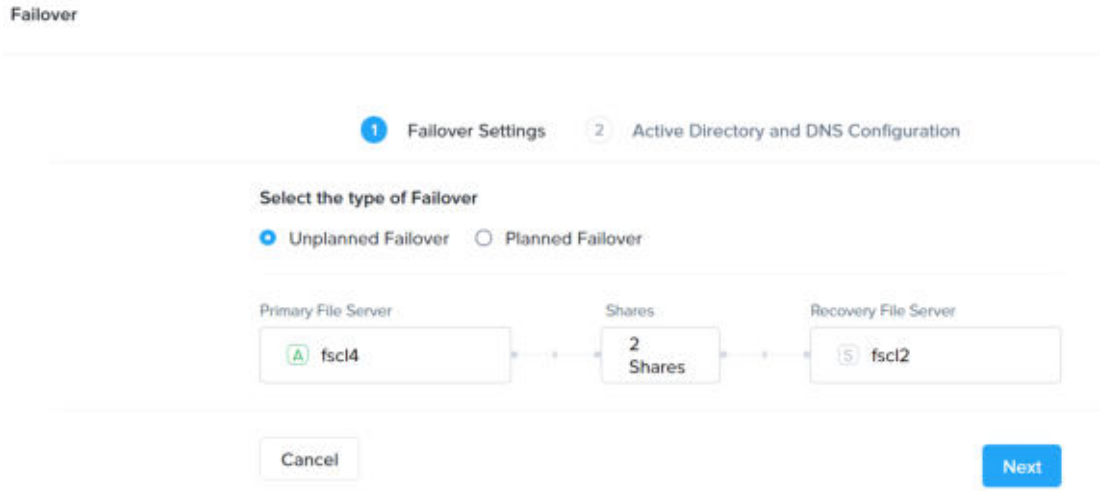

Figure 54: Files Smart DR Failover

With a planned failover, you can choose to begin replicating in the opposite direction automatically. Replicating after failover helps maintain service level agreements (SLAs) and RPOs during your failover testing or disaster avoidance operations.

SSR Interoperability

With Files 4.0, Smart DR supports replicating the snapshots between the source share and its target. The target share retention schedule doesn't have to match the source retention schedule. For example, if the file server is configured on the source to maintain the last two hourly snapshots, you can configure the target file server to maintain the last ten hourly snapshots to provide a longer retention window. Alternatively, you can match retention schedules to ensure that the same SSR copies exist in both sites for any failover event.

© 2021 Nutanix, Inc. All rights reserved

SSR schedules must be aligned by their frequency type between the source and target. For example, if you have daily snapshots scheduled on your source, you need daily snapshots scheduled on your target if you want them to be replicated and retained.

Files Smart DR Summary

Files Smart DR is a smart, simple, and effective way to replicate between Files instances, either on-premises or running on Nutanix Clusters in AWS (NCA).

In summary, Files Smart DR gives you the following advantages—all while maintaining simple and streamlined administration for your Nutanix environment:

- Prism Central integration.
- Share-level replication policies.
- Replication to an active read-only file server target.
- Faster RTOs.
- Support for storage-dense nodes.
- SSR compatibility.

© 2021 Nutanix, Inc. All rights reserved

6. Third-Party Integration

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 you to integrate file and web servers with security products. Nutanix chose this method to give customers wide latitude in selecting the antivirus solution that works best for their specific environments.

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 isn't infected, it returns the file handle to the SMB client.

Figure 55: ICAP Workflow

© 2021 Nutanix, Inc. All rights reserved

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 FSVM performance, one-click optimization recommends increasing the virtual CPU resources or scaling out the FSVMs. 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.

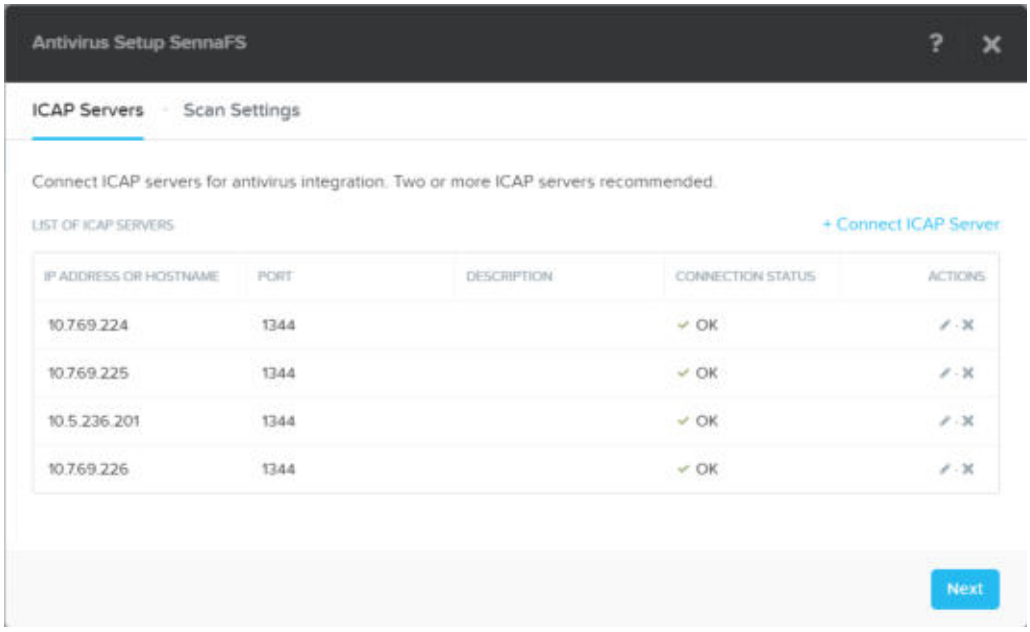

Figure 56: Configure Multiple ICAP Servers

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.

© 2021 Nutanix, Inc. All rights reserved

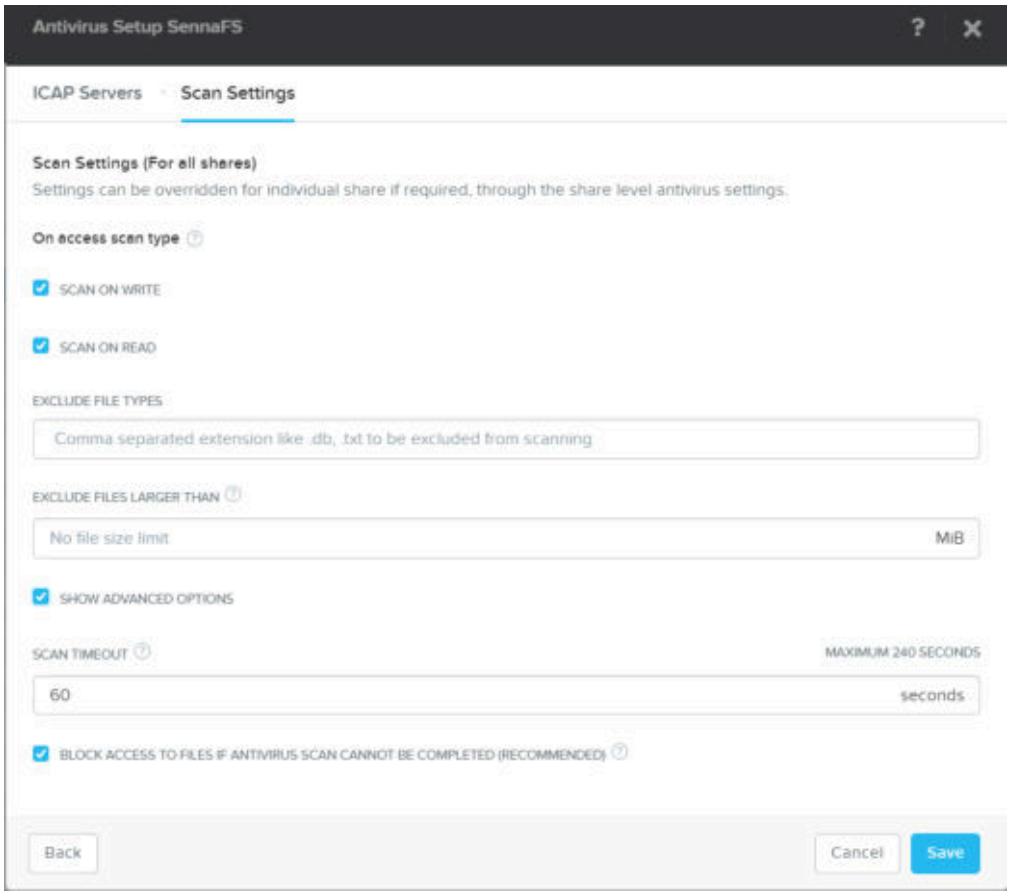

Figure 57: Default Scan Settings

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

© 2021 Nutanix, Inc. All rights reserved

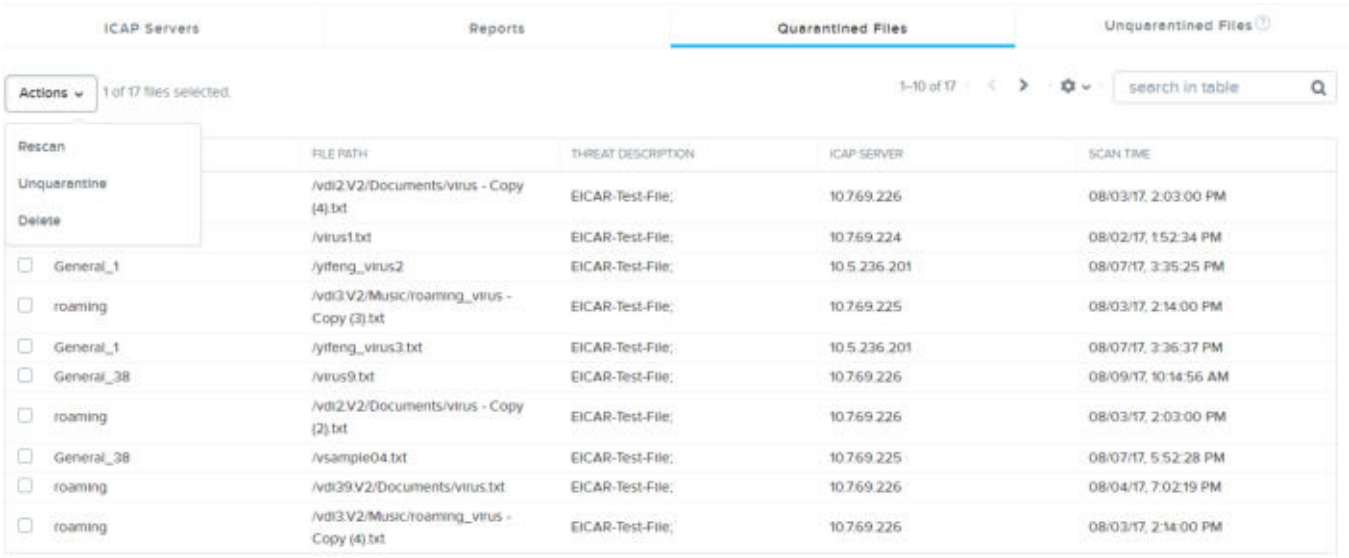

Figure 58: Quarantined Files

Once Nutanix Files quarantines a file, the administrator can rescan the file, remove it from quarantine, or delete it. 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 Compatibility and Interoperability Matrix on the Nutanix portal for the latest list of qualified ICAP server vendors with Nutanix Files.

File Operations Monitoring

File operations monitoring is a native Nutanix Files API used to forward all SMB and NFS operations run 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

© 2021 Nutanix, Inc. All rights reserved

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, see the PeerGFS and Nutanix Files datasheet.

- 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 an add-on for Nutanix Files. See netwrix.com for more details.

- Varonis

  Varonis version 8.6 introduces support for Nutanix Files version 3.7.1 and higher. This support includes Varonis's entire platform, DatAdvantage, DatAlert/DatAlert Analytics, DataPrivilege, Data Classification Engine, DatAnswers, Automation Engine, and Data Transport Engine.

© 2021 Nutanix, Inc. All rights reserved

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 shortens backup times because it doesn't 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 ensures that 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.

© 2021 Nutanix, Inc. All rights reserved

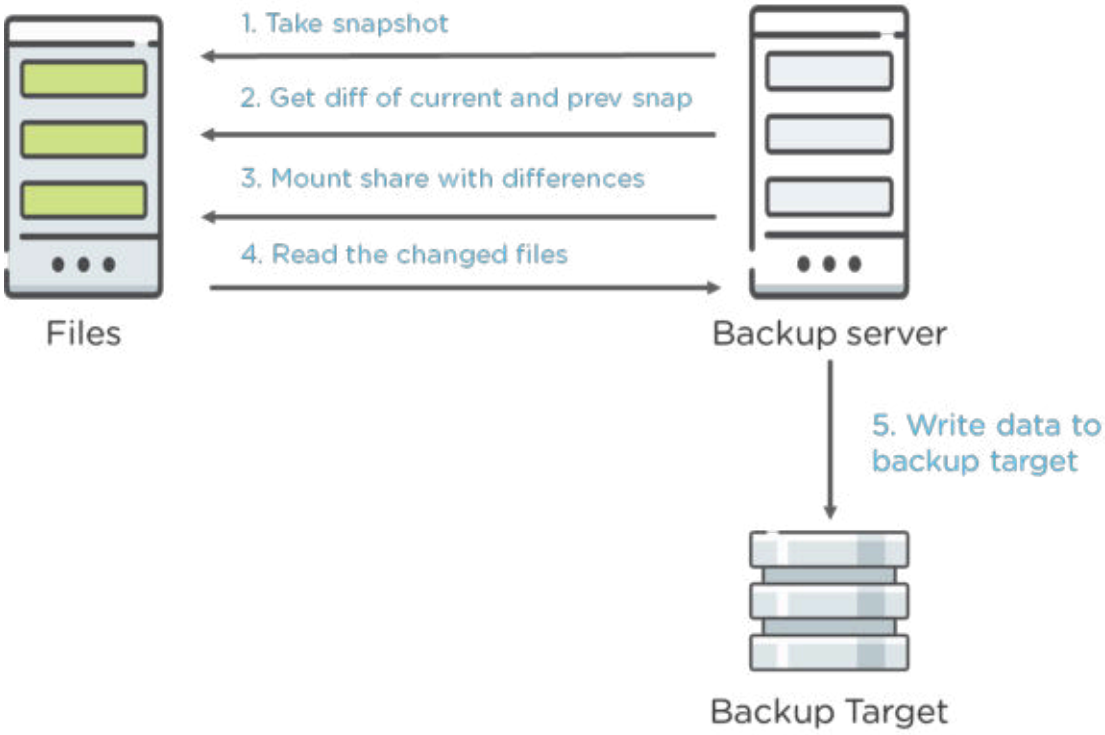

Figure 59: CFT Backup Process

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. Commvault version 11 service pack 15 is integrated with the Nutanix Files CFT API as of the 3.5 release. Veritas Netbackup 9.1 is fully integrated with the Nutanix Files CFT API.

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.

© 2021 Nutanix, Inc. All rights reserved

There are also vendors that provide support for backing up file shares without mounting to a guest VM. These applications can read directly from the UNC path.

Both Veeam and Rubrik are validated as non-CFT backup solutions for Nutanix Files. Refer to the Compatibility and Interoperability Matrix for the full list of validated backup applications.

> Tip: Don't try 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.

© 2021 Nutanix, Inc. All rights reserved

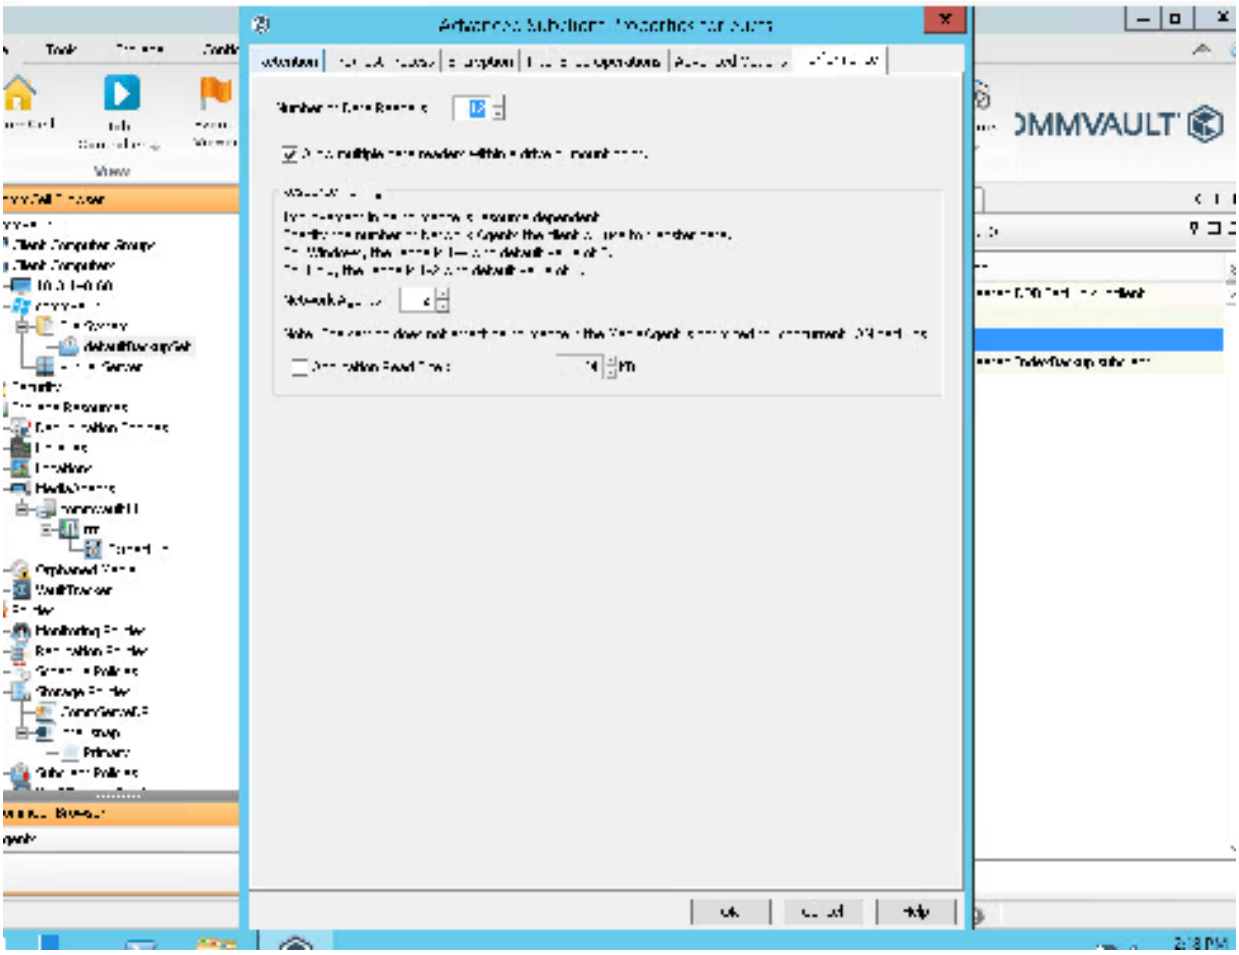

Figure 60: Adding More Readers to Drive Throughput on a Home Share

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 vCPU. Adding more vCPUs to the media agent achieved a shorter backup time.

© 2021 Nutanix, Inc. All rights reserved

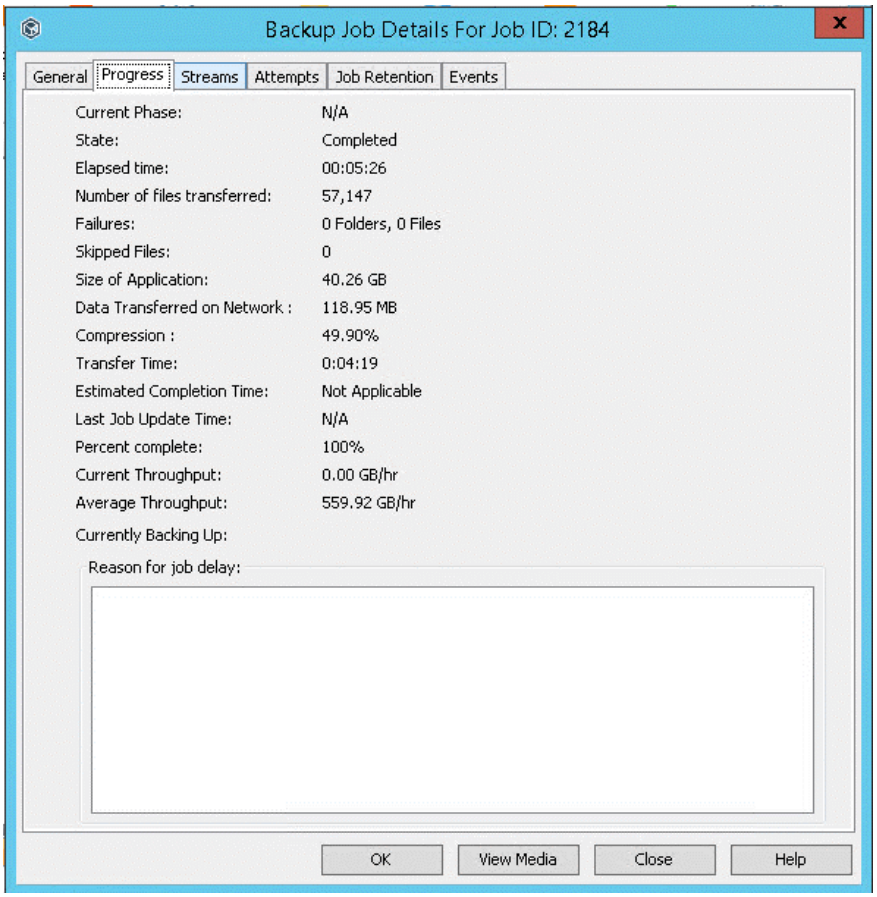

Figure 61: Backup Job Details

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

© 2021 Nutanix, Inc. All rights reserved

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.

AOS distributed storage and Prism make hosted file services highly resilient and easy to use—for example, you can configure the native data efficiency features of AOS, such as 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.

File Analytics gives you deep insight into your data, including storage consumption trends, file sizes, file counts, and data age. Analytics also provides end-to-end audit trails and anomaly detection. Ransomware intelligence helps you detect, protect against, analyze, and recover from ransomware threats.

Because you can deploy Files on a new or existing Nutanix environment, as well as in environments that are dedicated or mixed (with other services), you have the flexibility to choose your architecture to take advantage of Files services.

Nutanix Files streamlines design, implementation, and maintenance, providing an unrivaled experience for key use cases like end-user computing, video capture, medical imaging archives, application data, and the many other projects that store data in shared environments.

© 2021 Nutanix, Inc. All rights reserved

Appendix

About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix enterprise cloud software 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.

© 2021 Nutanix, Inc. All rights reserved

List of Figures

Figure 1: Nutanix Enterprise Cloud OS Stack..... 9
Figure 2: Nutanix Files Instances Run as VMs..... 11
Figure 3: Data Path Architecture of Nutanix Files..... 12
Figure 4: FSVM Internal Communication on One Node..... 13
Figure 5: FSVM vDisks and Volume Groups..... 14
Figure 6: Initial 40 TB File System Pool..... 15
Figure 7: 80 TB File System Pool Following First Expansion..... 16
Figure 8: 120 TB File System Pool After Second Expansion..... 16
Figure 9: 140 TB File System Pool After Final Expansion..... 16
Figure 10: Distributed Directory Shares..... 17
Figure 11: Two Standard Shares on the Same File Server..... 19
Figure 12: Create Share Menu with Nested Share Path..... 22
Figure 13: Submounting a Distributed Share into a Standard Share..... 23
Figure 14: Distributed Share and Top-Level Directories..... 24
Figure 15: Three Standard Shares..... 24
Figure 16: Performance Optimization Recommendation..... 26
Figure 17: Load Balancing Volume Groups with Nutanix Volumes..... 27
Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups..... 28
Figure 19: Each FSVM Controls Its Own Volume Groups..... 29
Figure 20: FSVM-1 Failure..... 30
Figure 21: DNS Request for SMB..... 33
Figure 22: Shared Folders MMC..... 35
Figure 23: Enable SMB Encryption..... 37

Figure 24: Nutanix Files User Management with NFS....39
Figure 25: DNS Request for NFSv4....40
Figure 26: DNS Request for NFSv3....41
Figure 27: Multiprotocol Shares....42
Figure 28: Enable Non-Native Protocol Access....43
Figure 29: Multiprotocol User Mapping....44
Figure 30: Setting a Share Quota....46
Figure 31: Blocked File Types for the File Server....48
Figure 32: Blocked File Types for a Share....48
Figure 33: File Analytics Dashboard....50
Figure 34: Data Age Analytics....51
Figure 35: File Analytics Audit Trails....52
Figure 36: File Analytics Anomaly Detection....53
Figure 37: File Analytics Usage Anomalies Dashboard....54
Figure 38: File Analytics Ransomware Page....55
Figure 39: Report Builder....56
Figure 40: Saved Reports....56
Figure 41: Smart Tier Architecture....58
Figure 42: Data Lens Dashboard....59
Figure 43: Data Lens Age Explorer....60
Figure 44: Data Lens Tiering....61
Figure 45: Tiered File....62
Figure 46: Data Summary by Age and Storage Tier....63
Figure 47: Protection Using SSR....65
Figure 48: Nutanix Files Cloning Use Cases....68
Figure 49: File Manager in Prism Central....69

Figure 50: Files Manager Data Protection..... 69
Figure 51: Files Smart DR..... 70
Figure 52: Data Protection Policy..... 71
Figure 53: Data Protection Compliance..... 71
Figure 54: Files Smart DR Failover..... 72
Figure 55: ICAP Workflow..... 74
Figure 56: Configure Multiple ICAP Servers..... 75
Figure 57: Default Scan Settings..... 76
Figure 58: Quarantined Files..... 77
Figure 59: CFT Backup Process..... 80
Figure 60: Adding More Readers to Drive Throughput on a Home Share..... 82
Figure 61: Backup Job Details..... 83

List of Tables

Table 1: Document Version History....6
Table 2: Supported Active Client Connections....32
Table 3: User Mapping Requirements....44
Table 4: Order of Precedence for Quotas....47

Prism 6.1

Prism Web Console Guide

March 21, 2022

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 16
  - Network Link Failure 18
  - Rack Fault Tolerance 18
  - 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
  - Monitoring Disk Rebuild Progress 61
  - Monitoring Node Rebuild Progress 63
  - Understanding Displayed Statistics 64
  - Modifying Cluster Details 66
  - Modifying UI Settings 68
  - Finding the AHV Version on Prism Element 72
  - Finding the AOS Version Using Prism Element 73
- Licensing 74
- Software and Firmware Upgrades 74
- Life Cycle Management (LCM) 75
- Multi-Cluster Management 75

Register (Unregister) Cluster with Prism Central....75
Installing Prism Central (1-Click Internet)....82
Prism Central Disaster Recovery....91
Single-node Clusters....94
Prerequisites and Requirements....95
Overriding Read-Only Mode....96
Two-Node Clusters....98
Two-Node Cluster Guidelines....99
Witness for Two-node Clusters....100
Failure and Recovery Scenarios....100
Increasing the Cluster Fault Tolerance Level....102
Replication Factor 1 Overview....107
Replication Factor 1 Recommendations and Limitations....107
Enabling Replication Factor 1....111
Creating a Storage Container with Replication Factor 1....111
Disabling Replication Factor 1....116
CVM Memory Configuration....116
Increasing the Controller VM Memory Size....118
Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro)....120
Restarting the Nodes in the Cluster....122

3. Storage Management....124
Storage Components....124
Compression....126
Deduplication....127
Erasure Coding....128
Capacity Reservation Best Practices....133
Storage Dashboard....133
Storage Overview View....134
Storage Diagram View....139
Storage Table View....146
Modifying a Storage Pool....156
Creating a Storage Container....157
Modifying a Storage Container....162
Creating a Storage Container with Replication Factor 1....163
Configuring a Warning Threshold for Resilient Capacity....167
Rebuild Capacity Reservation....167
Reserving Rebuild Capacity....173
Volume Group Configuration....174
Concurrent Access from Multiple Clients....175
Creating a Volume Group....175
Modifying a Volume Group....177
Flash Mode for Virtual Machines, Disks, and Volume Groups....177
Removing Flash Mode for Virtual Disks of a Volume Group....179
Recycle Bin....179
Recycle Bin Limitations and Considerations....180
Enabling and Disabling the Recycle Bin....181
Viewing Space Used by the Recycle Bin....181
Clearing Storage Space Used by the Recycle Bin....182

4. Network Management....183
Network Configuration for Cluster....183
Creating or Updating a Virtual Switch....185

Deleting a Virtual Switch ... 191
Migrating Bridges after Upgrade ... 192
Network Configuration for VM Interfaces ... 194
Configuring a Virtual Network for Guest VM Interfaces ... 194
Network Segmentation ... 202
Configuring Network Switch Information ... 202
Enabling LAG and LACP on the ToR Switch (AHV Only) ... 206
Creating SNMP Profiles ... 206
Modifying Switch Information ... 209
Network Visualization ... 210
Prerequisites ... 211
Network Visualizer ... 211
Viewing the Network Visualizer ... 212
Customizing the Topology View ... 212
Viewing VM Information ... 214
Viewing Host Information ... 215
Viewing Switch Information ... 217

5. Hardware Management ... 219
Hardware Dashboard ... 219
Hardware Overview View ... 221
Hardware Diagram View ... 223
Hardware Table View ... 231
Expanding a Cluster ... 243
Prerequisites and Requirements ... 250
Node Maintenance ... 255
Putting a Node into Maintenance Mode (AHV) ... 257
Viewing a Node under the Maintenance Mode ... 261
Exiting a Node from the Maintenance Mode (AHV) ... 263
UVM Status when Node under Maintenance Mode ... 265
Starting Host Boot Disk Repair ... 268
Modifying a Cluster ... 269
Removing a Disk ... 270
Removing a Node from a Cluster ... 271
Adding a Node ... 276
Compute-Only Node Configuration (AHV Only) ... 277
Adding a Compute-Only Node to an AHV Cluster ... 279
Storage-only Node Configuration ... 281
Imaging and Adding Storage-only Nodes ... 283

6. Nutanix Volumes ... 284

7. File Server Management ... 285

8. Data Protection ... 286

9. Health Monitoring ... 287
Health Dashboard ... 287
Configuring Health Checks ... 292
Configuring NCC Frequency ... 295
Running Checks by Using Web Console ... 295

Collecting Logs by Using Web Console....296

10. Virtual Machine Management....300
VM Dashboard....300
VM Overview View....302
VM Table View....303
VM Management....316
Creating a VM (AHV)....316
Managing a VM (AHV)....326
Virtual Machine Snapshots....338
Adding Multiple vGPUs to the Same VM....342
Migrating Live a vGPU-enabled VM Within the Cluster....344
Creating a VM (ESXi)....347
Managing a VM (ESXi)....350
Configuring Images....358
Virtual Machine Customization....361
Customizing Linux Virtual Machines with Cloud-Init....362
Cloud-Init Limitations and Guidelines....362
Customization of Windows Virtual Machines with System Preparation....363
VM High Availability in Acropolis....364
Enabling High Availability for the Cluster....366
Nutanix Guest Tools....367
Nutanix Guest Tools Requirements and Limitations....368
Enabling NGT and Mounting the NGT Installer in a VM....369
NGT Installation....371
Enabling NGT and Mounting the NGT Installer on Cloned VMs....375
Regenerating NGT Certificates for Guest VMs....376
Nutanix Guest Agent and Controller VM Communication....377
CD-ROM Eject Functionality of NGT....377
Nutanix Guest Tools Usage in Disaster Recovery....378
NGT Metrics Collection for System Performance (Windows Perfmon)....378
Upgrading NGT....383
Reconfiguring NGT....383
Disabling, Uninstalling, and Removing Nutanix Guest Tools....384
Affinity Policies Defined in Prism Element....387
Configuring VM-VM Anti-Affinity Policy....388
Connect to Citrix Cloud....389
Connecting to the Citrix Cloud....389
Guest VM Cluster Configuration (AHV Only)....394
Creating a Guest VM Cluster by Directly Attaching a Volume Group (AHV Only)....394

11. Performance Monitoring....396
Analysis Dashboard....396
Creating an Entity Chart....398
Creating a Metric Chart....399
Chart Metrics....400
Exporting Performance Data....410

12. Alert and Event Monitoring....412
Alerts Dashboard....412
Alerts Summary View....412
Events Summary View....418
Configuring Alert Emails....420

Configuring Alert Policies 424
Alerts/Health checks 426
Availability Zone 426
Cluster 426
Controller VM 540
DR 558
Entity Sync 615
Guest VM 615
Hardware 627
License 645
Node 646
Network 671
Nutanix Files 688
Other 713
Storage 726
System Indicator 745
VM 746

13. View Task Status 747
View Task Status Dashboard 748

14. System Management 750
Configuring a Filesystem Whitelist 750
Configuring Name Servers 751
Cluster Time Synchronization 752
Recommendations for Time Synchronization 752
Configuring NTP Servers 753
Configuring an SMTP Server 754
Configuring SNMP 755
Nutanix MIB 762
Configuring a Banner Page 769
VM Management through Prism Element (ESXi) 769
Registering a Cluster to vCenter Server 771
Managing vCenter Server Registration Changes 773
In-Place Hypervisor Conversion 774
Requirements and Limitations for In-Place Hypervisor Conversion 775
In-Place Hypervisor Conversion Process 778
Converting Cluster (ESXi to AHV) 779
Converting Cluster (AHV to ESXi) 780
Stopping Cluster Conversion 782
Internationalization (i18n) 783
Localization (L10n) 783
Changing the Language Settings 784
Hyper-V Setup 785
Joining the Cluster and Hosts to a Domain 785
Creating a Failover Cluster for Hyper-V 787
Manually Creating a Failover Cluster (SCVMM User Interface) 788
Enabling Kerberos for Hyper-V 790

15. Security and User Management 795

16. Support Services 796

Pulse Health Monitoring....796
Configuring Pulse....798
Pulse Access Requirements....800
Controlling Remote Connections....800
Configuring HTTP Proxy....801
Accessing the Nutanix Support Portal....804
Nutanix REST API....807
Accessing the REST API Explorer....807

17. Help Resources....811
Accessing Online Help....811
Accessing the Nutanix Next Community....812
Glossary....812

Copyright....817
License....817
Conventions....817
Version....817

Booklist....818

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 75 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 124 describes how to monitor storage use in a Nutanix cluster and how to create storage containers.
- Hardware Management on page 219 describes how to monitor hardware configurations in a Nutanix cluster and how to expand the cluster.
- Health Monitoring on page 287 describes how to monitor the health of VMs, hosts, and disks across a Nutanix cluster.
- Virtual Machine Management on page 300 describes how to monitor status of the VMs across a Nutanix cluster.
- Performance Monitoring on page 396 describes how to monitor and analyze performance in a Nutanix cluster.
- Alert and Event Monitoring on page 412 describes how to monitor activity in a Nutanix cluster and how to configure alert policies and notification.
- System Management on page 750 describes how to configure various system settings such as for SNMP, NTP, and SMTP.
- Security and User Management on page 795 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 796 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.

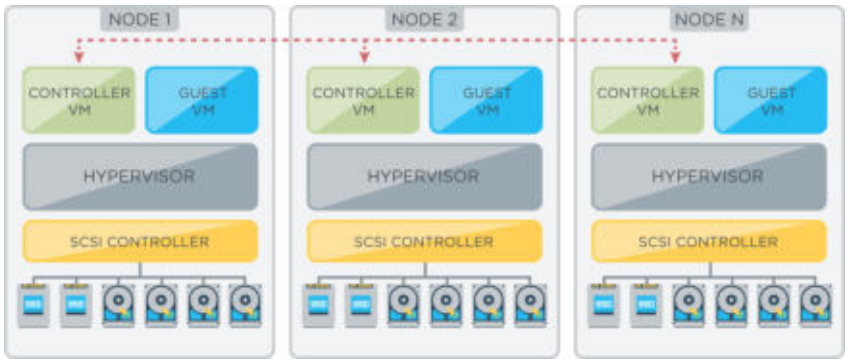

Guest VM Data Management

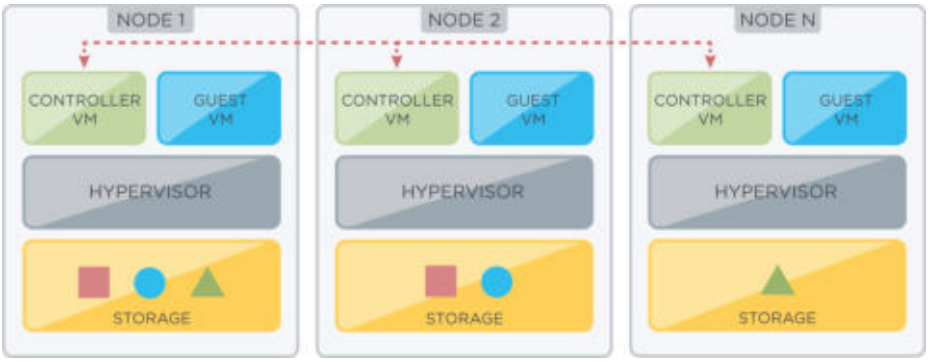

Hosts read and write data in shared Nutanix datastores as if they were connected to a SAN. From the perspective of a hypervisor host, the only difference is the improved performance that results from data not traveling across a network. VM data is stored locally, and replicated on other nodes for protection 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

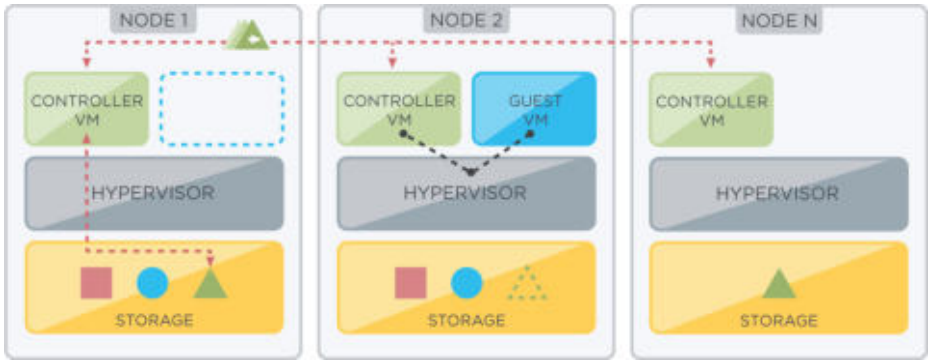

Live migration of VMs, whether it is initiated manually or through an automatic process like vSphere DRS, is fully supported by the Nutanix Enterprise Cloud Computing Platform. All hosts within the cluster have visibility into shared Nutanix datastores through the Controller VMs. Guest VM data is written 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

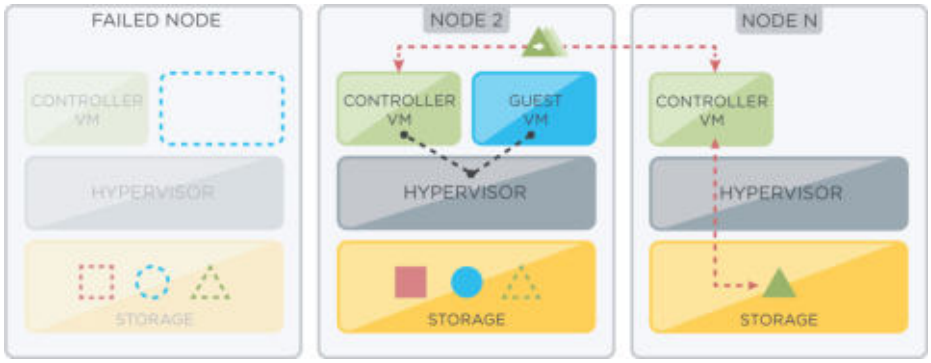

The built-in data redundancy in a Nutanix cluster supports high availability provided by the hypervisor. If a node fails, all HA-protected VMs can be automatically restarted on other nodes in the cluster. The hypervisor management system, such as vCenter, selects a new host for the VMs, which may or 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

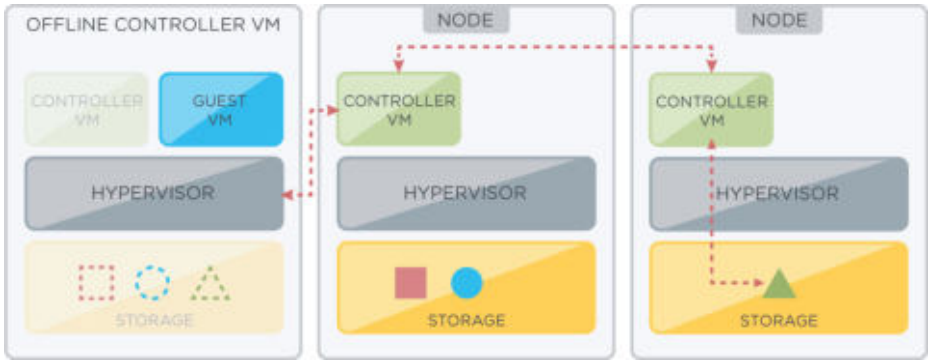

The Nutanix cluster automatically selects the optimal path between a hypervisor host and its guest VM data. The Controller VM has multiple redundant paths available, which makes the cluster more resilient to failures.

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

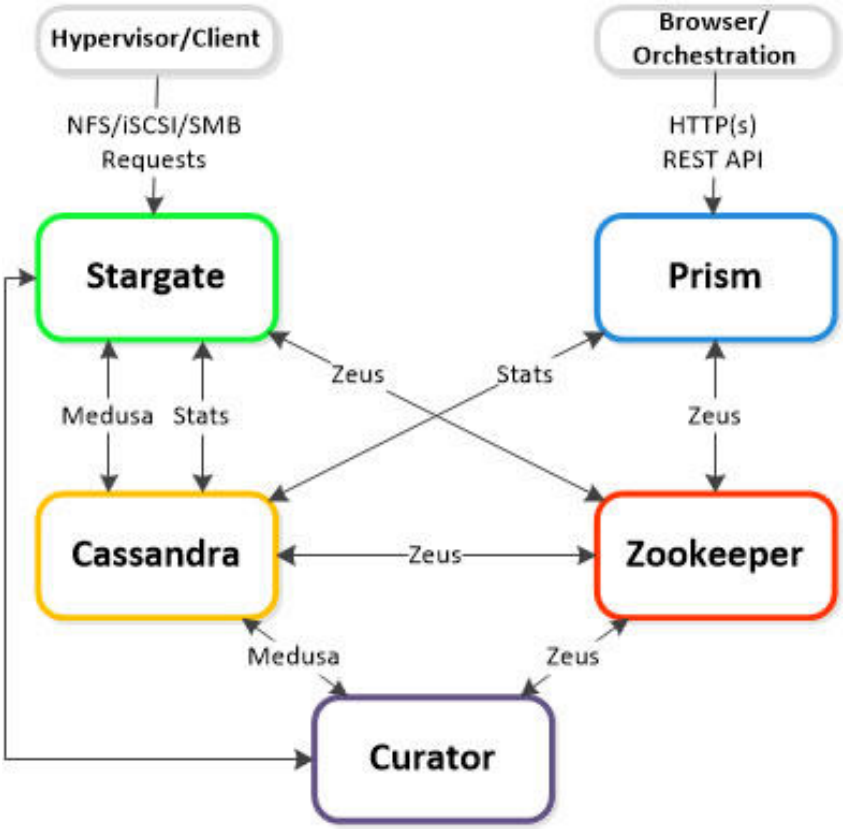

The Nutanix cluster has a distributed architecture, which means that each node in the cluster shares in the management of cluster resources and responsibilities. Within each node, there are software components that perform specific tasks during cluster operation.

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 18
- Block Fault Tolerance on page 22
- Configuring Node Fault Tolerance on page 27

The following terms appear on various screens of Prism web console and Prism Central:

Failure Domain

Also known as Fault Tolerance Level and defined as the entity such as a node, block or rack, the failure of which the cluster can tolerate while still running guest VMs and responding to commands through the management console. You can configure failure domain as a node in the cluster, a block or a rack.

Fault Tolerance

It is defined as the number of failures of the Failure Domain entity that the cluster can tolerate tolerate while still running guest VMs and responding to commands through the management console. For example, if Failure Domain is defined as a node, then fault tolerance must be defined as the number of failed nodes that the cluster can

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

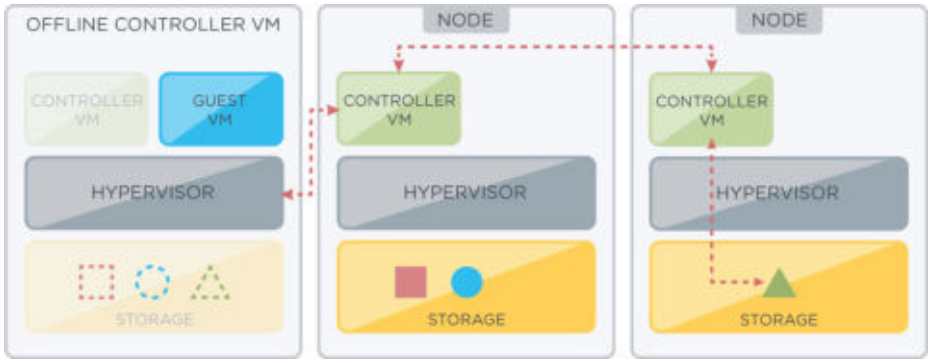

The Nutanix cluster monitors the status of Controller VMs in the cluster. If any Stargate process fails to respond two or more times in a 30-second period, another Controller VM redirects the storage path on the related host to another Stargate. Reads and writes occur over the 10 GbE network until the 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.

A CVM failure may include a user powering down the CVM, a CVM rolling upgrade, or any event, which might bring down the CVM. In any of these cases the storage traffic is served by another CVM in the cluster. The hypervisor and CVM communicate using a private network on a dedicated virtual switch. This means that the entire storage traffic is routed through an internal IP address on the CVM. The external IP address of the CVM is used for remote replication and for CVM to CVM communication.

In the event of a local CVM failure, the local addresses previously used by the local CVM become unavailable. In this case, the NDFS automatically detects the outage and redirects the storage traffic to another CVM in the cluster over the network. The re-routing is done transparently to the hypervisor and to the VMs running on the host. This indicates that even if a CVM is powered-off, the VMs continue to perform the I/O operations. The Nutanix Distributed File System (NDFS) is also self-healing, which means that NDSF detects when a CVM is powered-off and it automatically reboots the local CVM. Once the local CVM is available, the traffic is seamlessly transferred back to be served by the local CVM.

The NDFS uses replication factor (RF) and checksum to ensure data redundancy and availability in the case of a node or disk failure or corruption. In the case of a node or disk failure the data is then re-replicated among all nodes in the cluster to maintain the RF, which is called re-protection. Re-protection may result after a CVM is down.

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

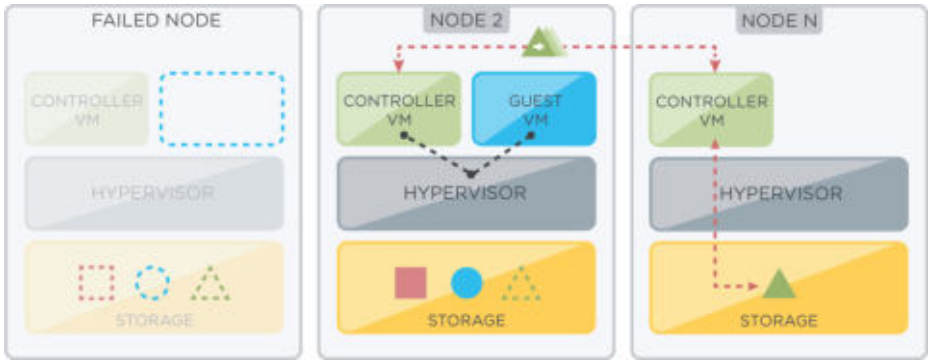

The built-in data redundancy in a Nutanix cluster supports high availability provided by the hypervisor. If a node fails, all HA-protected VMs can be automatically restarted on other nodes in the cluster.

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 might restart under certain rare conditions on dual SSD nodes if a boot drive fails, or if you unmount 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, ESXi, and HyperV.

  Note: – Prism Central workflows are not applicable for rack fault tolerance using HyperV since HyperV support is not available for Prism Central

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
- Network Latency Limit - The round trip latency for communication between Controller VMs in a Nutanix cluster should be less than or equal to 1 ms (millisecond).

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**.

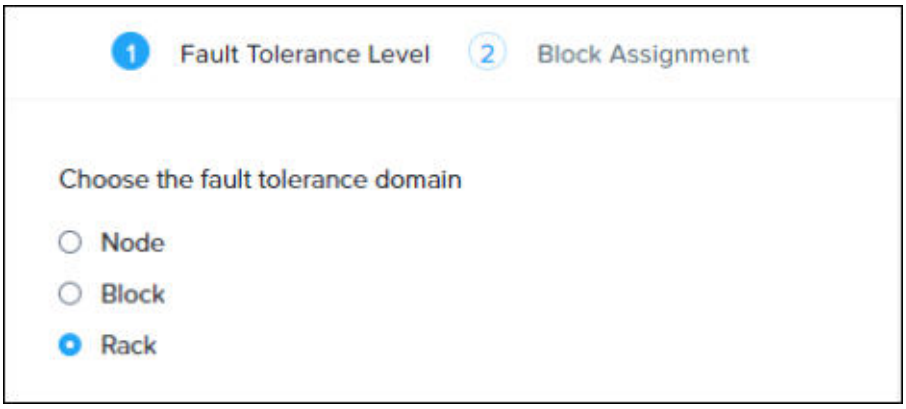

Figure 1: Select Fault Tolerance Level

   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.

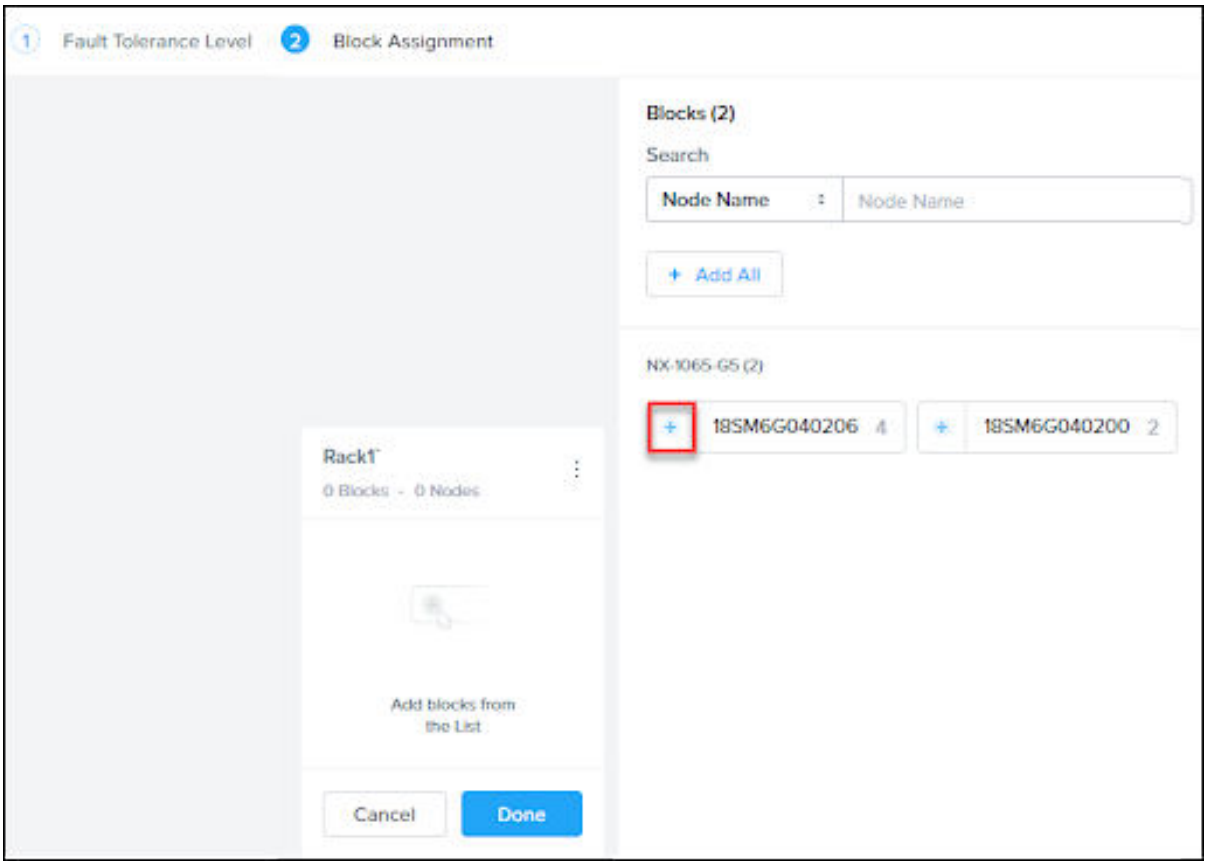

Figure 2: Add blocks to rack

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.

Note: The fault tolerance has three states.

- **OK**: This state indicates that the fault tolerance domain is highly resilient to safely handle a node failure.
- **Warning**: This state indicates that the fault tolerance level is almost reaching to 0. The fault tolerance domain may not handle a node failure gracefully.
- **Critical**: This state indicates that the fault tolerance level is 0, and the fault tolerance domain cannot handle a node failure.

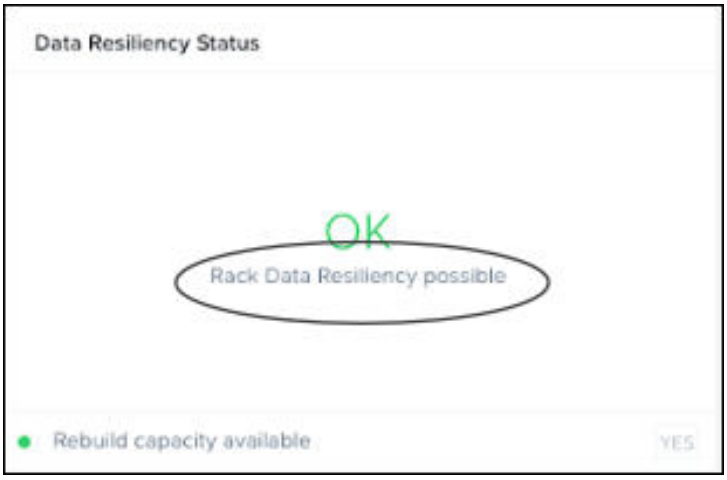

Figure 3: Data Resiliency widget

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

Fault Domain	Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Minimum Number of Racks	Data Resiliency
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
	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.

> CAUTION: If rack awareness is enabled on a cluster with erasure coding enabled containers, the cluster may temporarily lose some space savings due to the EC strip unpacking. Any storage savings lost will be reclaimed after 7 days when the data is write cold.

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.

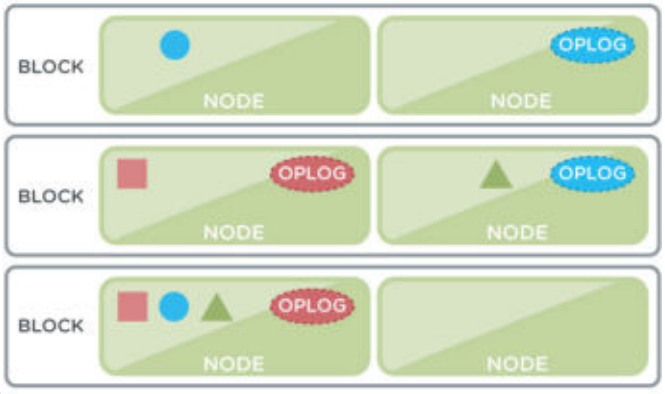

Figure 5: Block-aware placement of guest VM data

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.

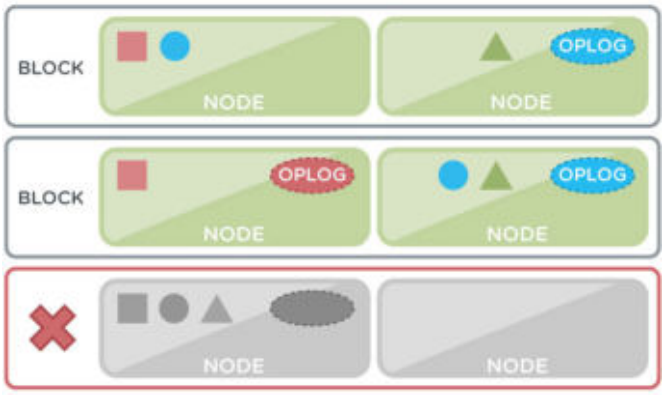

Figure 6: Block-aware placement of guest VM data with block failure

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 .

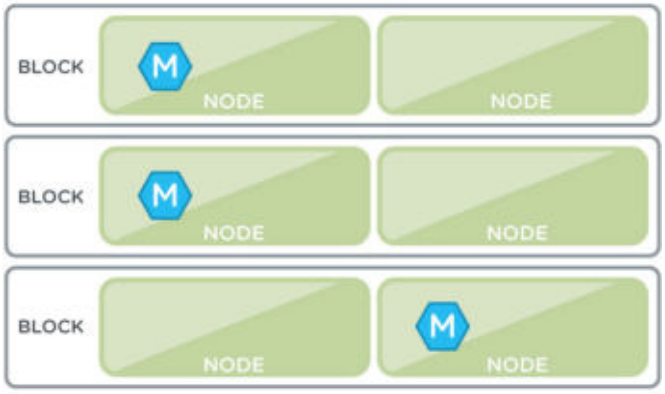

Figure 7: Block-aware placement of metadata

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.

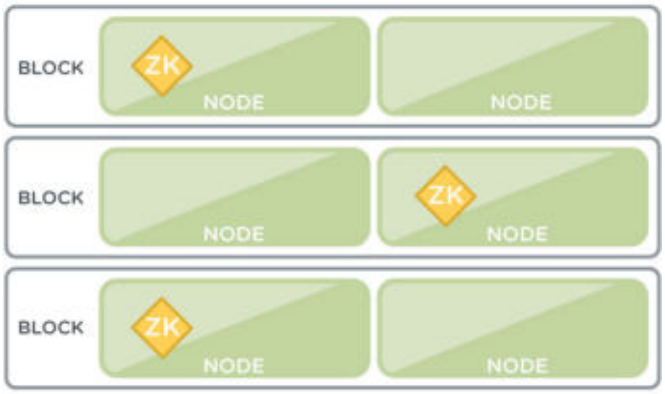

Figure 8: Block-aware placement of configuration data

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
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.

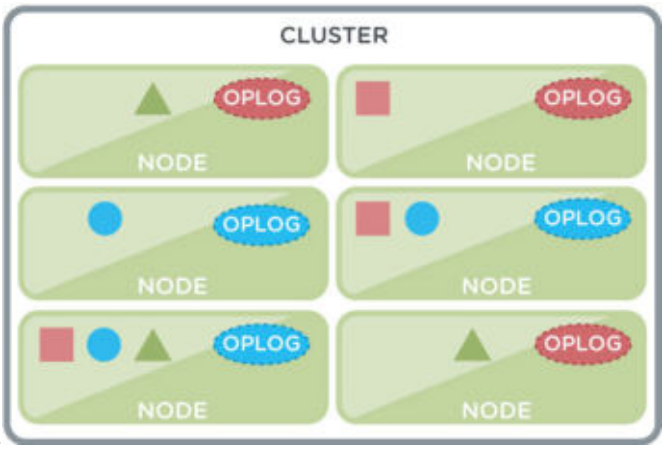

Figure 10: Replication factor 3 placement of guest VM data

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.

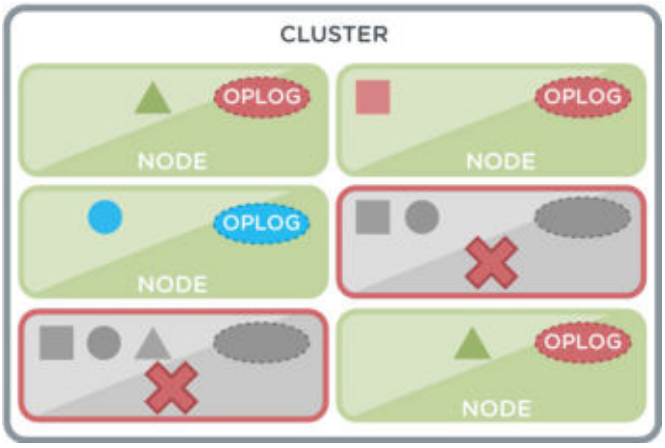

Figure 11: Replication factor 3 placement of guest VM data with failure of 2 nodes

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.

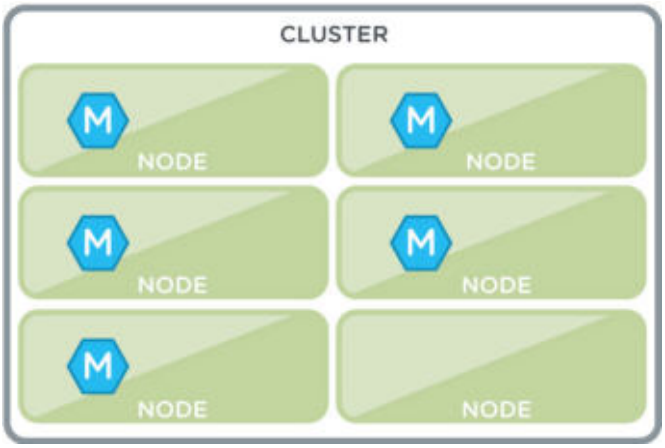

Figure 12: Redundancy factor 3 metadata

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.

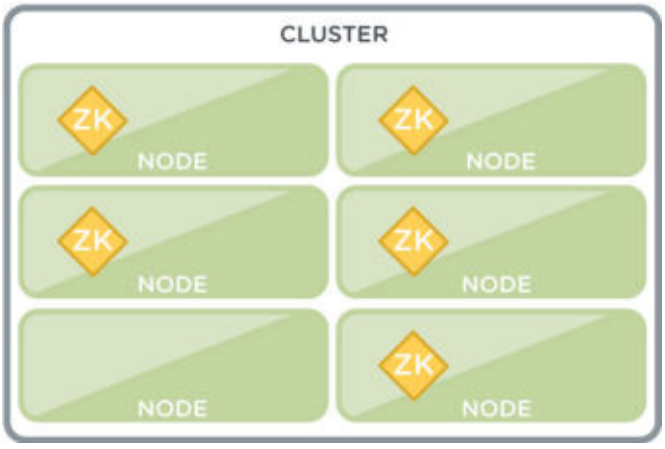

Figure 13: Redundancy factor 3 configuration data

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.

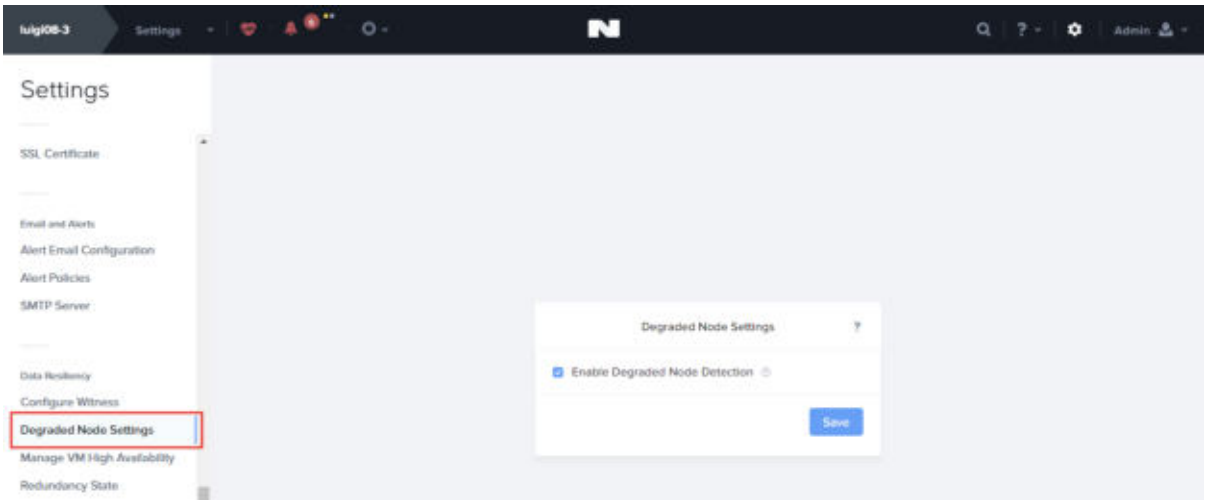

Figure 14: Degraded Node Settings Window

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.

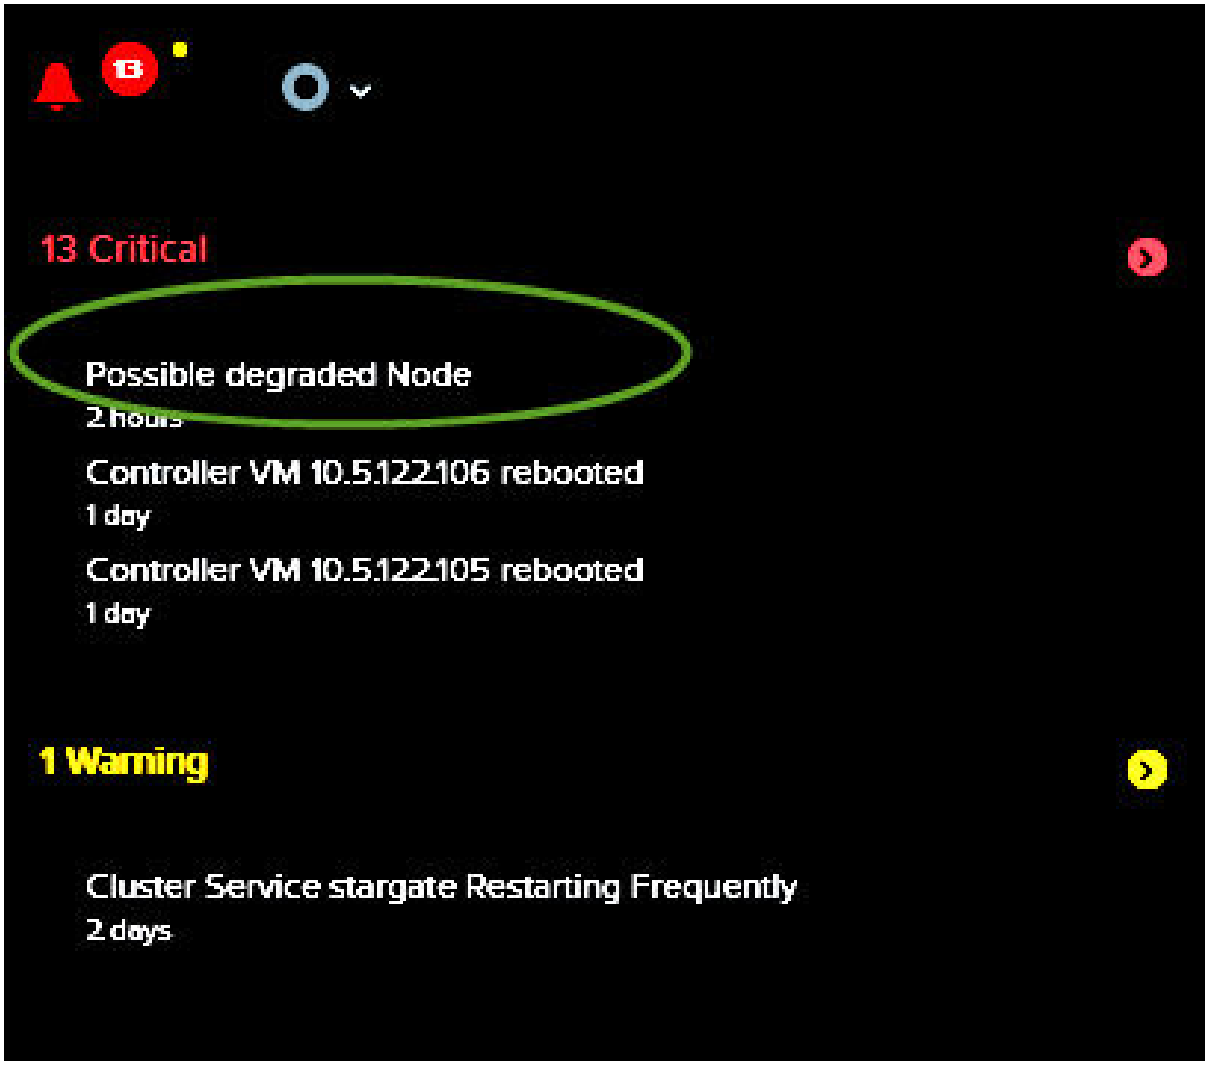

Figure 15: Possible Degraded Node in Alerts

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.

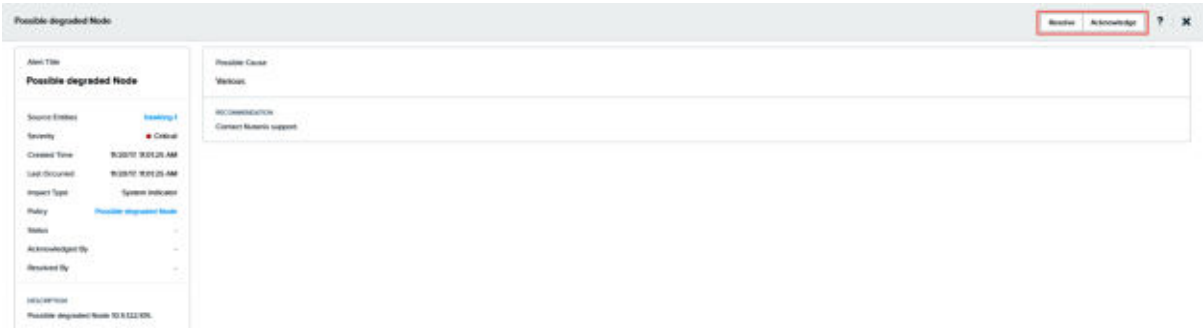

Figure 16: Possible Degraded Node Window

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.

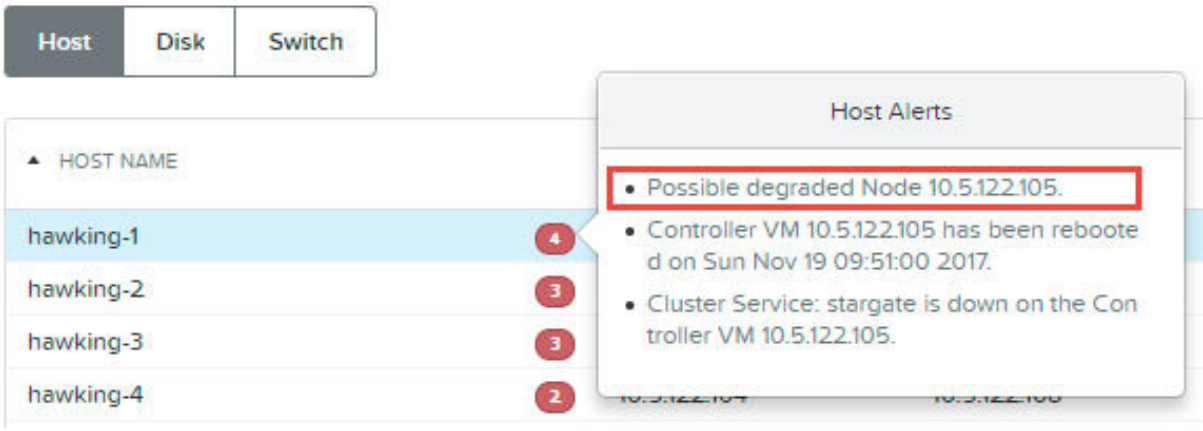

Figure 17: Degraded Node Pop-up Alert

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 “” () . • 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 later = 62 TB
	ESXi 5.1 and previous = 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 = 6 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
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

Entity	Supported Maximum
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
	128 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 74).
- You can upgrade a cluster when a new AOS release is available through the web console (see Software and Firmware Upgrades on page 74.
- If you have multiple clusters, you can manage them all through a single web interface (see Multi-Cluster Management on page 75).

> 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 412).

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 769), 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.

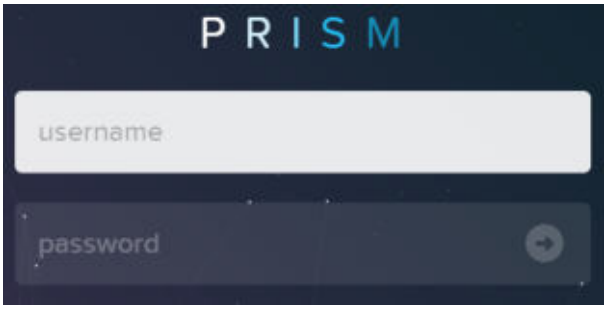

Figure 19: Login Screen

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 68.

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.

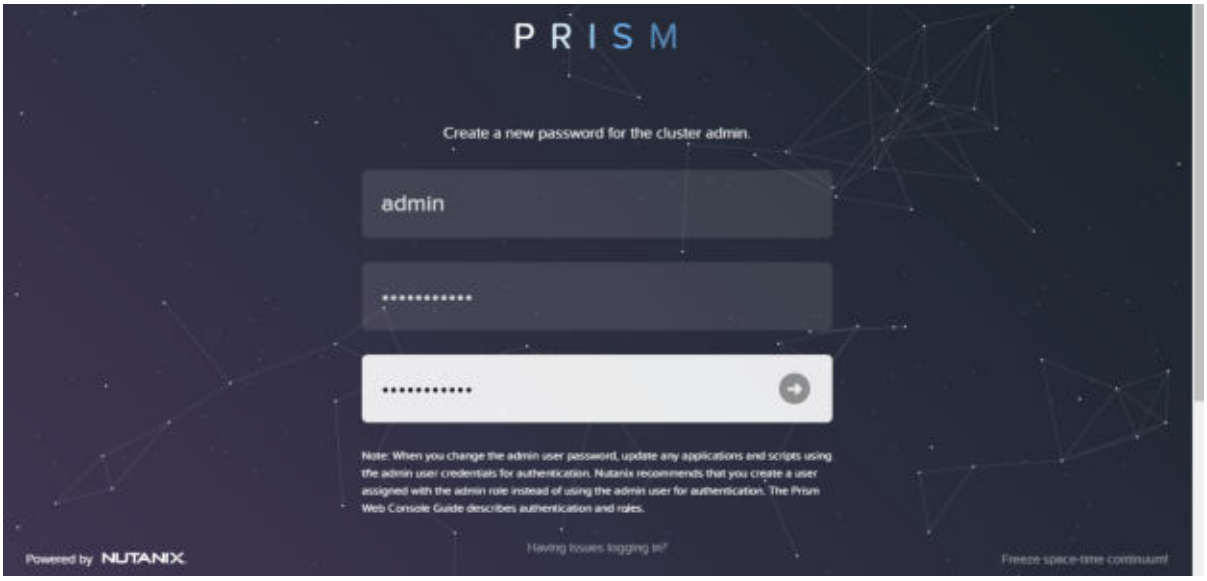

Figure 20: Login Screen (first admin login)

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.

- The default password expiration age for the admin user is 60 days. You can configure the minimum and maximum password expiration days based on your security requirement.
  - nutanix@cvm$ sudo chage -M <MAX-DAYS> admin
  - nutanix@cvm$ sudo chage -m <MIN-DAYS> admin
- 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.

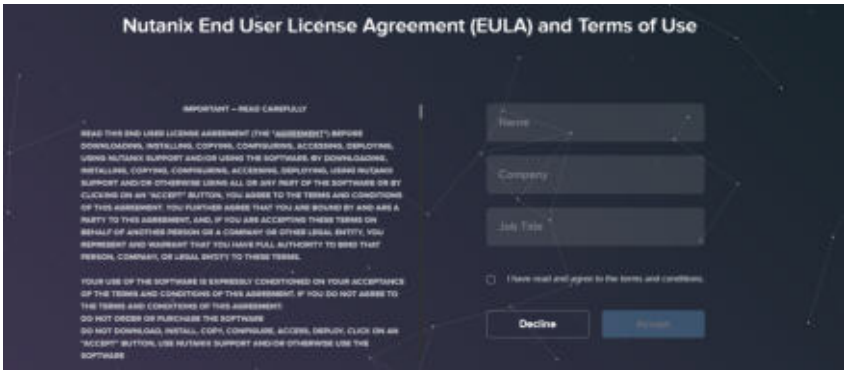

Figure 21: EULA Screen

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 798), 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 798). Enabling Pulse is recommended unless providing cluster information to Nutanix customer support violates your security policy.

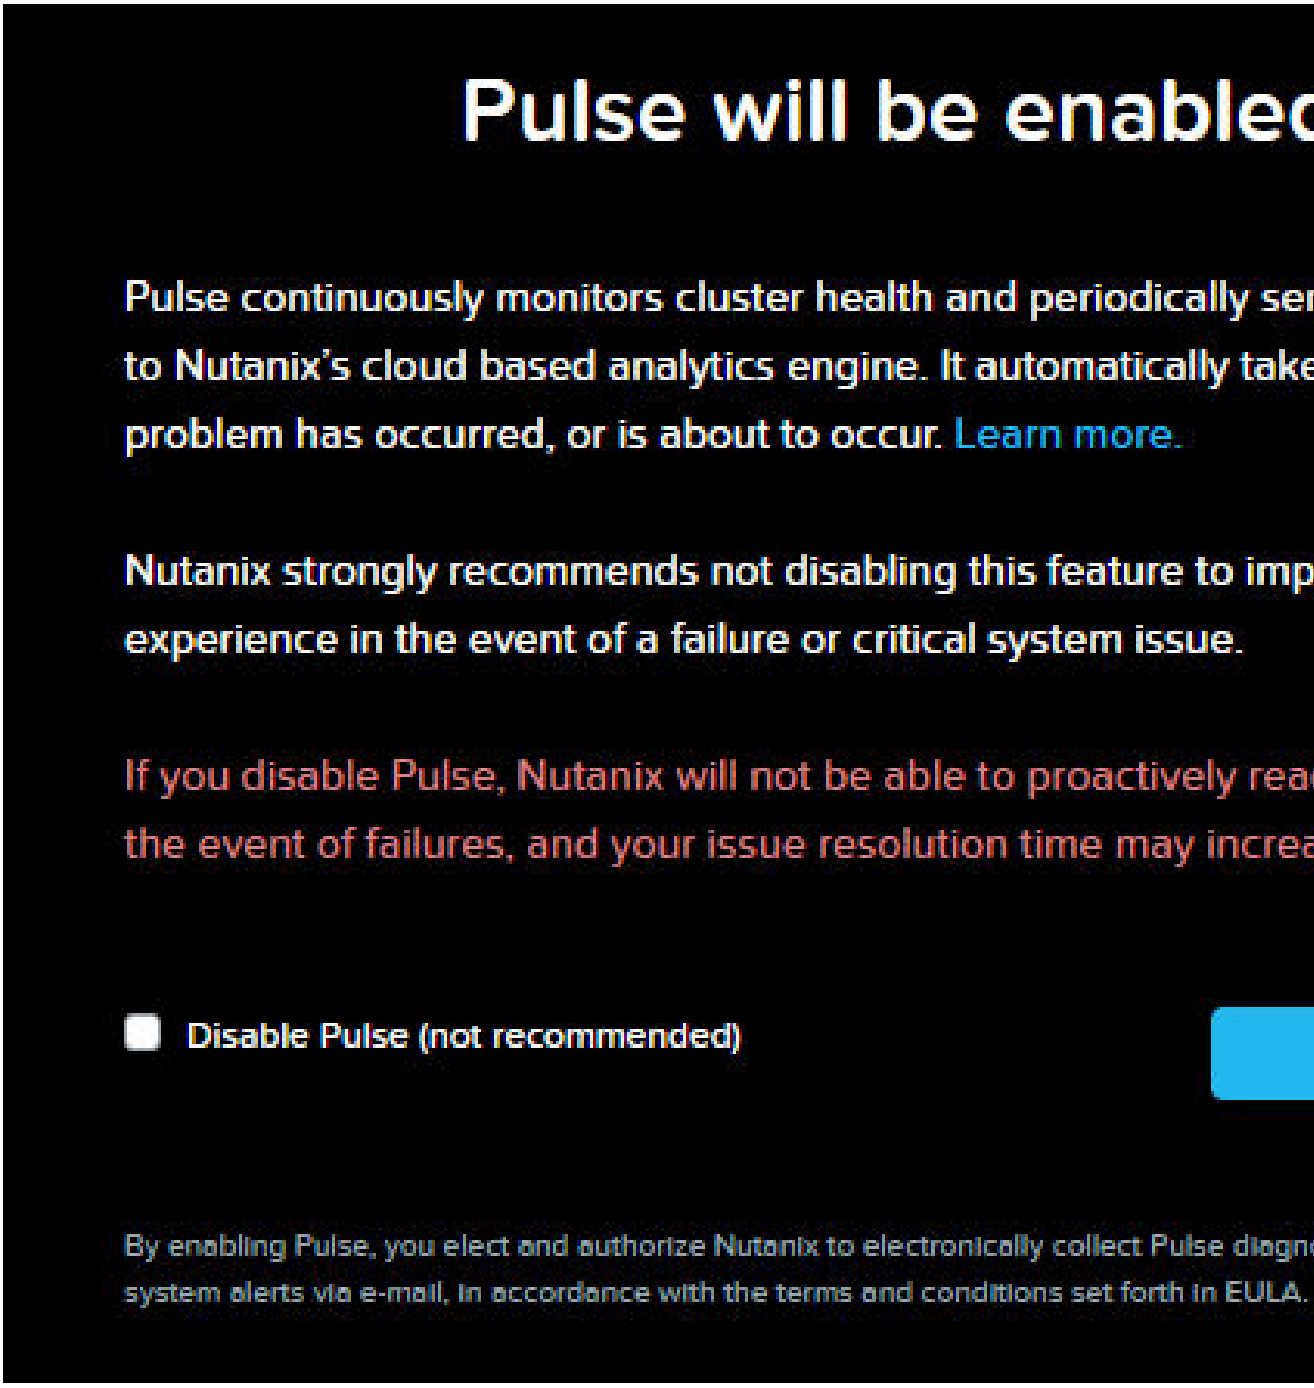

Figure 22: Pulse Screen

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.

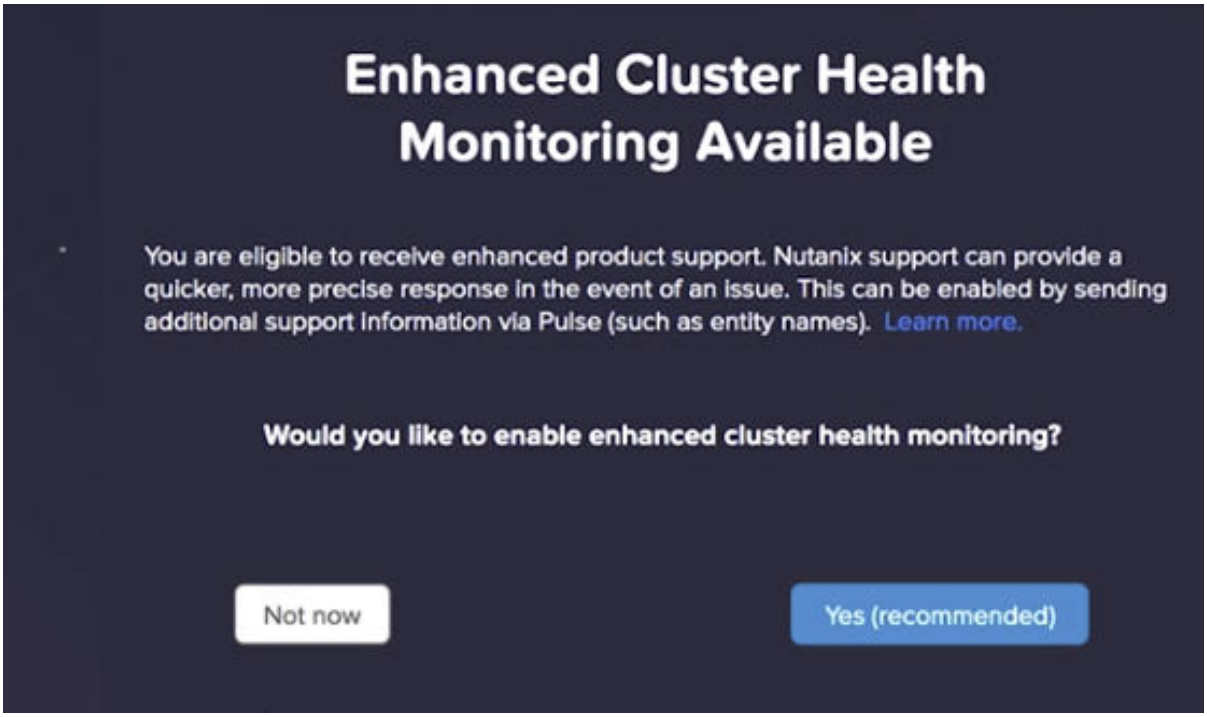

Figure 23: Enhanced Monitoring Screen (upgrade example)

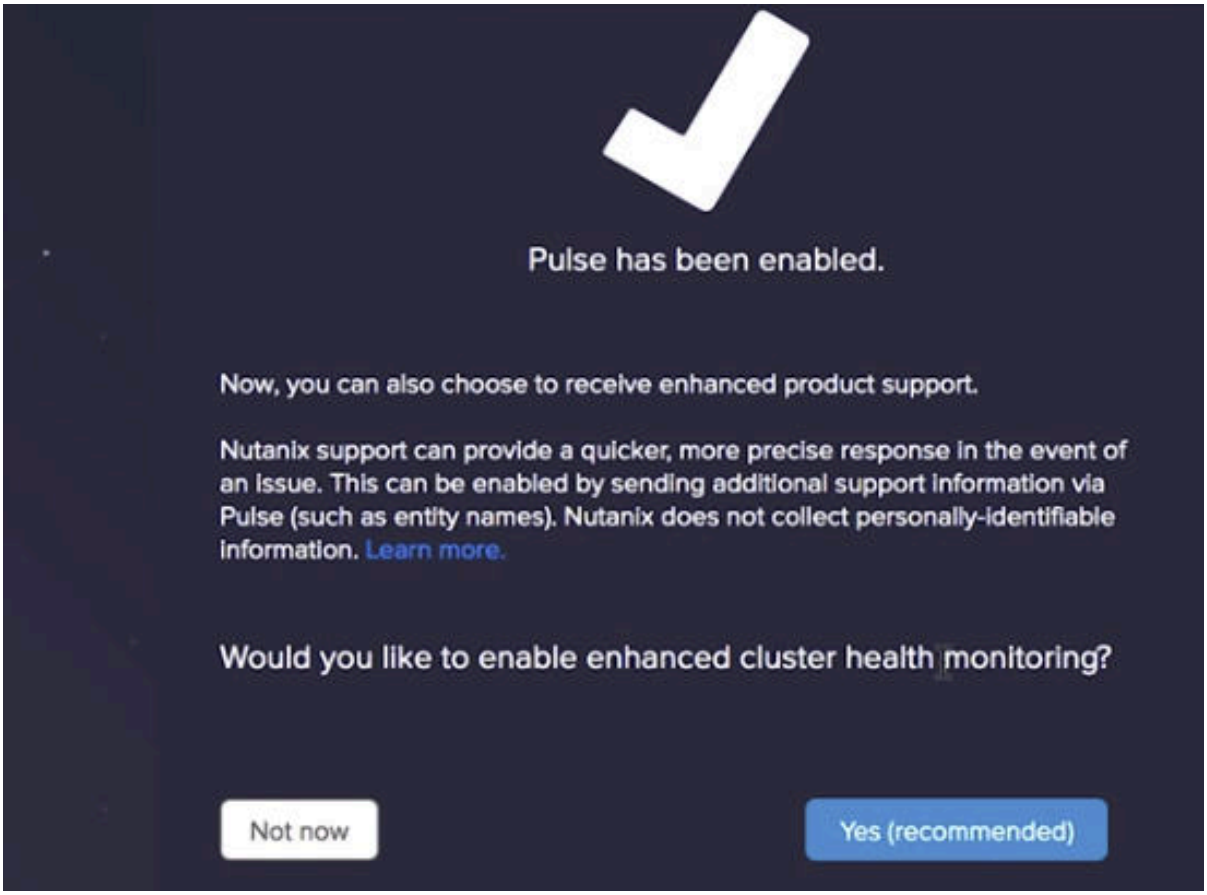

Figure 24: Enhanced Monitoring Screen Example (initial example)

9. [2-node clusters only] If a screen about registering with a Witness appears, see the *Witness Option* topic in the *Data Protection and Recovery with Prism Element* guide.

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.

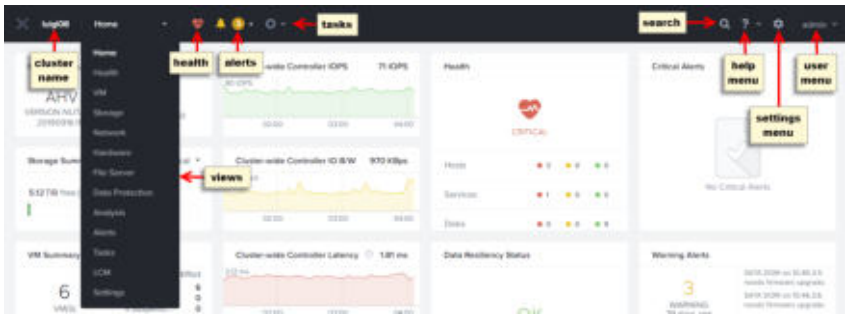

Figure 25: Main Menu with Expanded View List

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 66).

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 287).
- **VM**. Displays a virtual machine information dashboard (see VM Dashboard).
- **Storage**. Displays a storage information dashboard (see Storage Dashboard on page 133).
- **Network** (AHV only). Displays the network visualiser.
- **Hardware**. Displays a hardware information dashboard (see Hardware Dashboard on page 219).
- **File Server**. Displays a file server dashboard (see Nutanix Files Guide).
- **Data Protection**. Displays a data protection information dashboard (see “” ()).
- **Analysis**. Displays a screen to create and run performance monitors (see Analysis Dashboard on page 396).
- **Alerts**. Displays a screen of alert and events messages (see Alerts Dashboard).
- **Tasks**. Displays a screen of task messages (see View Task Status on page 747).
- **LCM**. Displays the life cycle manager (LCM) dashboard (see Life Cycle Management (LCM) on page 75).
- **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 287).
- 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 an alert 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 View Task Status on page 747).
- 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.

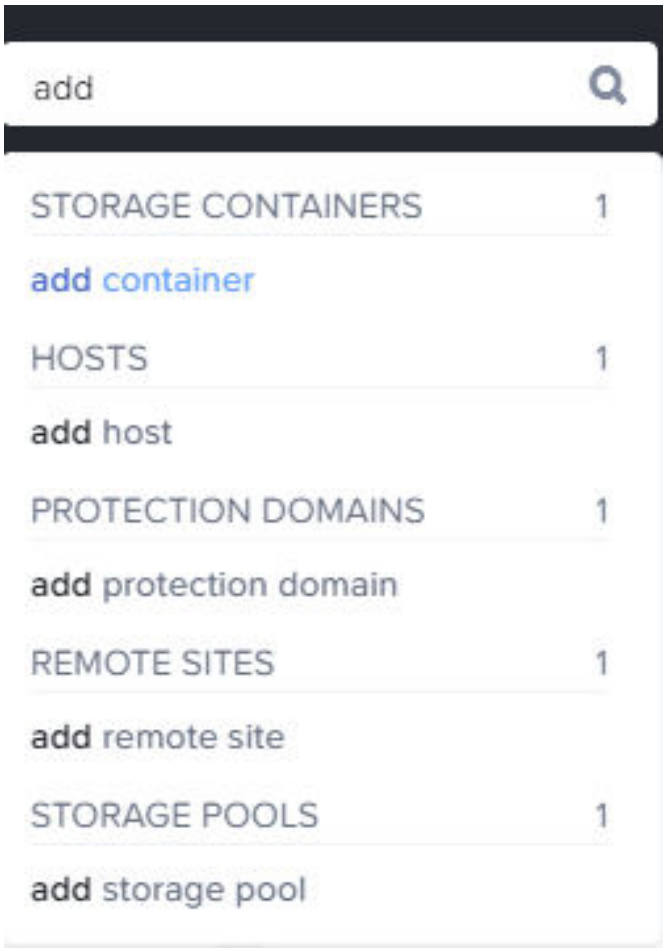

Figure 26: Search Box Example

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: Help Menu Options

Name	Description
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help on page 811).
Health Tutorial	Opens the Health dashboard tutorial that takes you through a guided tour of the health analysis features (see Health Dashboard on page 287).
General Help	Opens the online help at the introduction page.

Name	Description
Support Portal	Opens a new browser tab (or window) at the Nutanix support portal login page (see Accessing the Nutanix Support Portal on page 804).
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 812. 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 Options

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 807).
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 *PowerShell Cmdlets Reference* in the *Acropolis V1 API Reference* for information about the cmdlets.
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 Finding the AOS Version Using Prism Element on page 73).
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.

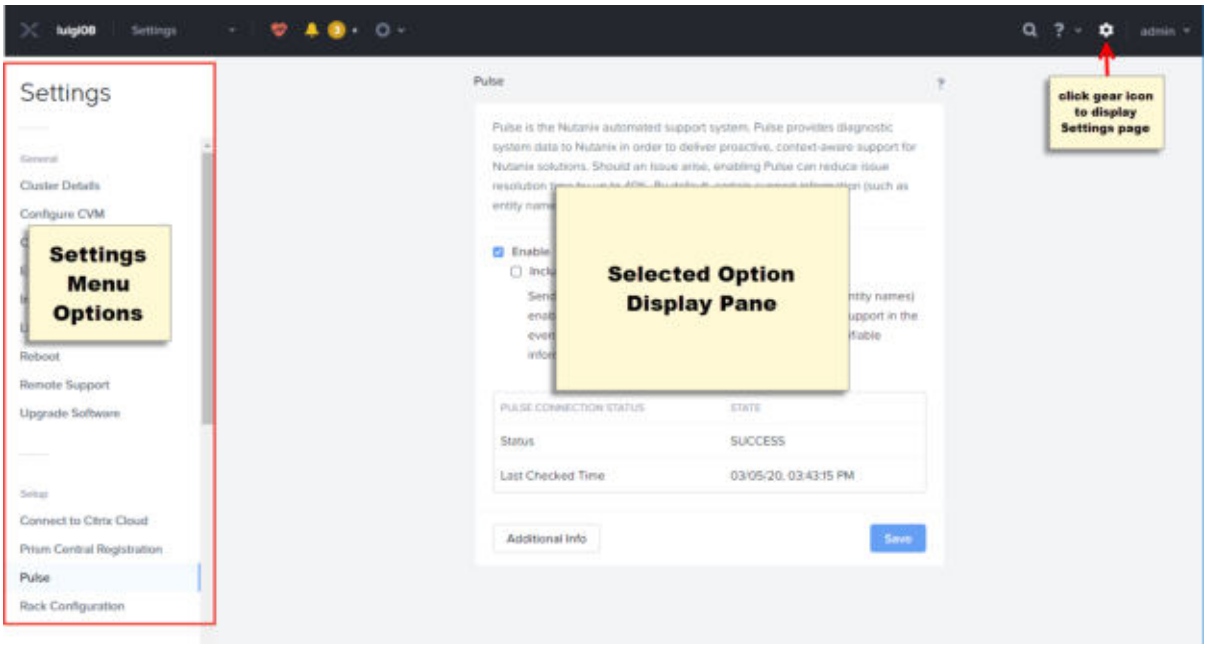

Figure 27: Settings Page

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 66).
Configure CVM	Opens the **Configure CVM** window to increase the Controller VM memory size (see Increasing the Controller VM Memory Size on page 118).
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 774).
Expand Cluster	Opens the **Expand Cluster** window to add new nodes to the cluster (see Expanding a Cluster on page 243).

Name	Description
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 358).
Licensing	Opens the **Licensing** window to install or update the cluster license that enables entitled Nutanix features (see Licensing on page 74).
Reboot [AHV]	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 on page 122).
Remote Support	Opens the **Remote Support Services** window, which enables (or disables) Nutanix remote support access (see Controlling Remote Connections on page 800).
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 74).
vCenter Registration [ESXi only]	Opens the **vCenter Registration** window to register (or unregister) the cluster with the vCenter instance (see Registering a Cluster to vCenter Server on page 771).
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 389).
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 75).
Pulse	Opens the **Pulse** window to enable the sending of cluster information to Nutanix customer support for analysis (see Configuring Pulse on page 798).
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 18).
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 801).
Name Servers	Opens the **Name Servers** window to configure name servers for the cluster (see Configuring Name Servers on page 751).
Network Configuration [AHV only]	Opens the **Network Configuration** window to configure network connections for the cluster (see Network Configuration for VM Interfaces on page 194).

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 202). 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 753).
SNMP	Opens the **SNMP Configuration** window to enable and configure SNMP for the cluster (see Configuring SNMP on page 755).
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 750).
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 420).
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 424).
SMTP Server	Opens the **SMTP Server Settings** window to configure an SMTP server (see Configuring an SMTP Server on page 754).
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 the *Witness Option* topic in the *Data Protection and Recovery with Prism Element* guide.
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 366).
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 102).
Appearance	
Language Settings	Opens the **Languages** window, which allows you to select the language of the web console (see Localization (L10n) on page 783).
UI Settings	Opens the **UI Settings** window to to configure Prism UI background themes, disable (or re-enable) the login screen (see Modifying UI Settings on page 68).
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 769).

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 the options 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, health 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.

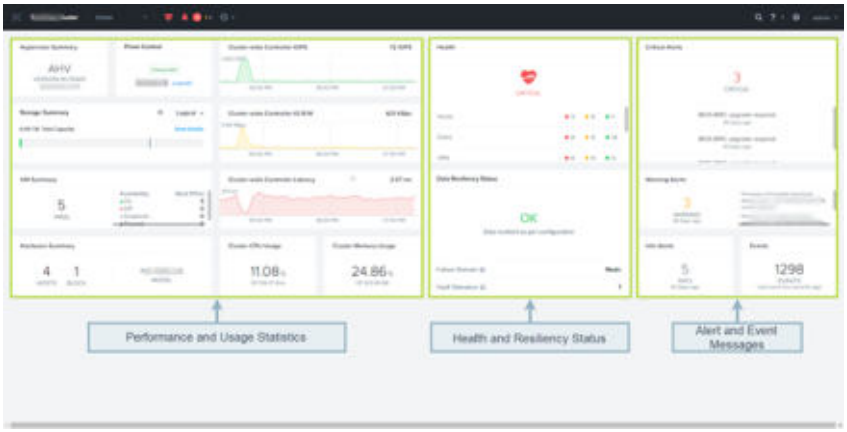

Figure 28: Home Dashboard

Note: See Understanding Displayed Statistics on page 64 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 **Launch** to launch the Prism Central instance in a new tab of your browser.

Name	Description
Storage Summary	Displays information about the storage space utilization (in GiB or TiB) and resilient capacity of the cluster. Placing the cursor anywhere on the horizontal axis displays a breakdown view of the storage capacity usage. You can select **Physical** or **Logical** from the drop-down list on the right. **Physical**: Displays information about the physical capacity, which is the total addressable space. **Logical**: Displays information about the logical capacity, which is the usable space after accounting for the data replication factor. The **View Details** link displays the resiliency status and storage information of all the individual nodes in the cluster (see Storage Details Page on page 137). You can also configure a threshold warning for the resilient capacity utilization in the cluster by clicking the gear icon (see Configuring a Warning Threshold for Resilient Capacity on page 167). 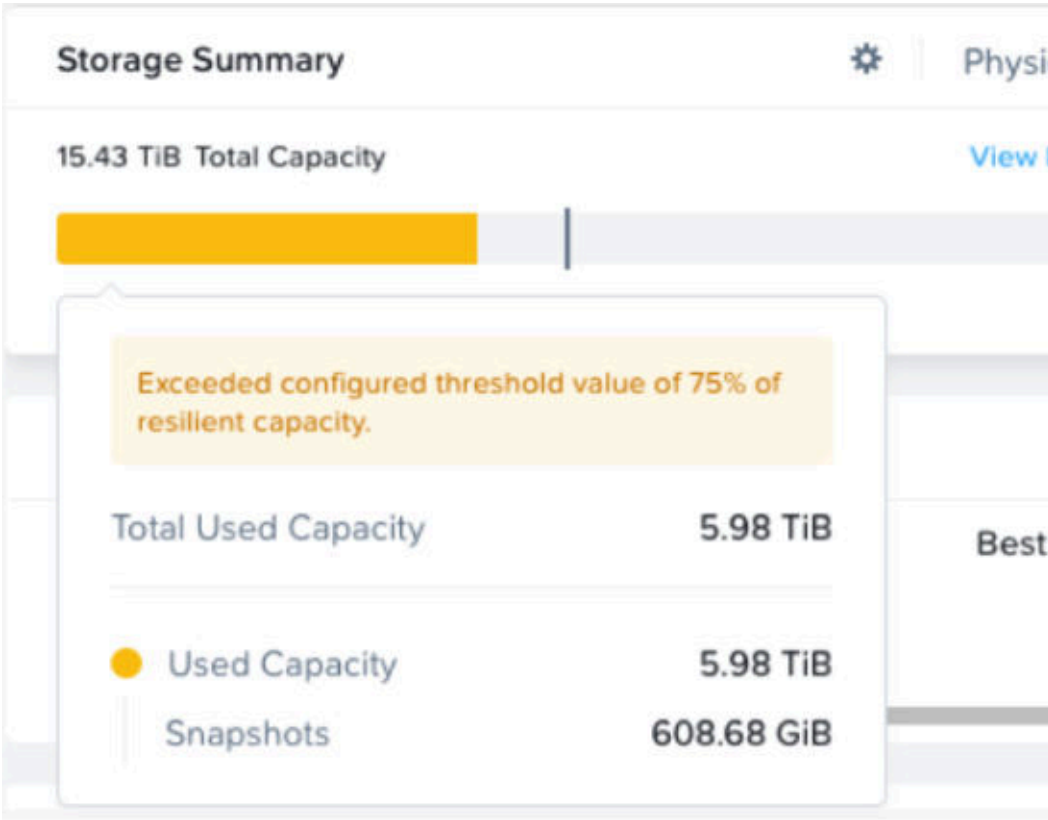
VM Summary	Displays the total number of VMs in the cluster broken down by on, off, and suspended states.
Hardware Summary	Displays the number of hosts and blocks in the cluster, plus one or more Nutanix block model numbers.

Name	Description
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 396).
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).
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 the VMs, hosts, or disks line displays detailed information about that object in the Health page (see Health Dashboard on page 287).

Name	Description
Data Resiliency Status	Displays information indicating whether the cluster is protected currently from potential data loss due to a component failure. Clicking anywhere in this field displays the Data Resiliency Status dialog box. For more information about the Data Resiliency Status, 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 Progress**. Indicates data rebuild status for offline nodes and disks marked for removal. For more information about the Rebuild Progress indicator, see the following "Rebuild Progress" section. • **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.

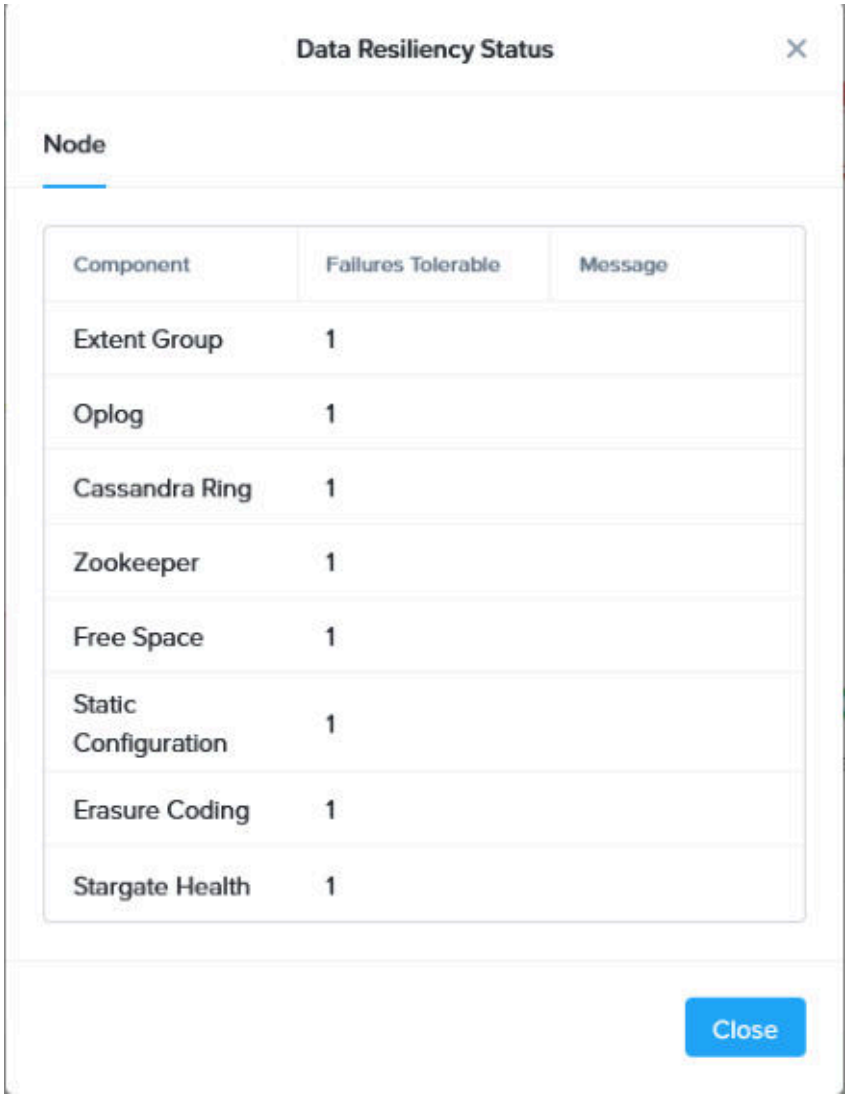

Figure 29: Data Resiliency Status Window

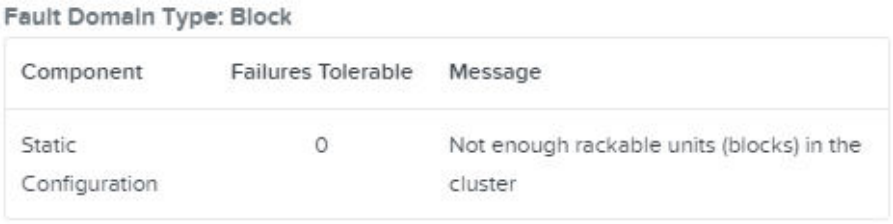

Rebuild Progress (Data Resiliency Widget)

You can monitor the data rebuild status of cluster entities like disks and nodes through the **Rebuild Progress** indicator. This rebuild progress bar reflects in the **Data Resiliency Status** widget on the Prism Element **Home** page within a maximum duration of 90 seconds when a data rebuild event occurs due to a disk removal (or a node removal which triggers disk removals of the individual disks).

For example, if you mark a disk for removal, it takes a maximum of 90 seconds for its rebuild progress to appear in the **Data Resiliency Status** widget.

The main objective of this feature is to provide the admin with information about the ongoing data rebuilds for disks and nodes. Along with the overall ETA and rebuild progress percentage for participating entities, you can also monitor the individual progress of the entities.

Data Rebuild Scenarios

You can monitor the data rebuild for the following scenarios.

- You mark a disk for removal.

  A disk is removed in any of the following cases.

  - You manually perform the **Remove Disk** operation.
  - A Stargate process internally encounters that a disk has gone down and marks it for removal.
- You mark a node offline.

The **Data Resiliency Status** widget displays the data rebuild progress and ETA as an aggregate of all overlapping rebuilds. If a data rebuild operation intersects other rebuild operations, they are considered to belong to a single rebuild generation. In such a scenario, the Prism Element web console displays the weighted average of the progress percentage of all the entities and maximum time of the ETA for all rebuilds in that generation.

See the Monitoring Node Rebuild Progress on page 63 and Monitoring Disk Rebuild Progress on page 61 sections for information on how to monitor the overall data rebuild ETA, overall progress percentage, and detailed progress information for nodes and disks respectively.

> Note: **Data Resiliency Status** displays only a single generation data rebuild at any point in time.

The following examples show the overlapping data rebuilds for entities.

Example 1: Node and Disk

The following screenshot displays the overall rebuild progress percentage and ETA for a node and a disk.

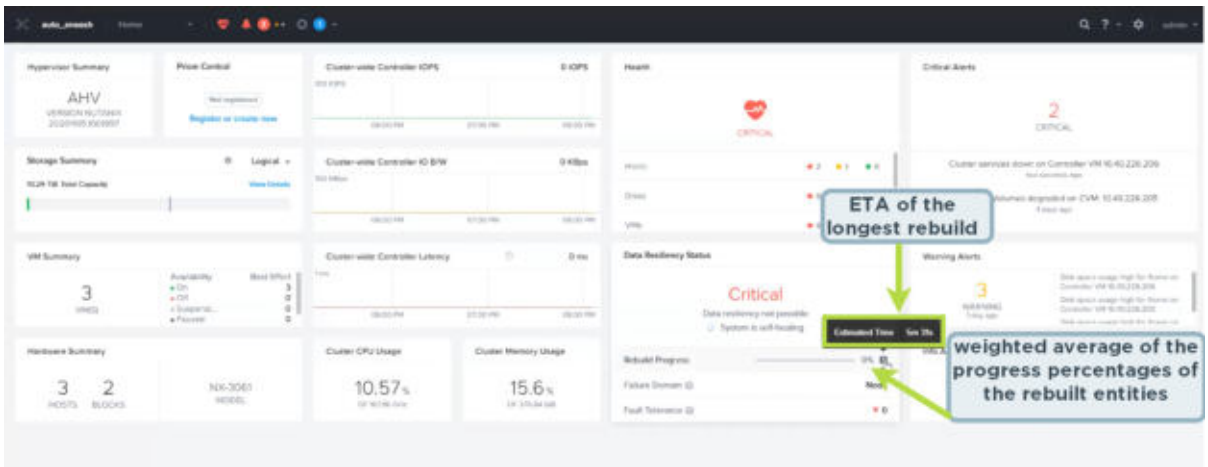

The following screenshot displays the rebuild information for a node.

Figure 31: Rebuild Progress Details (Node)

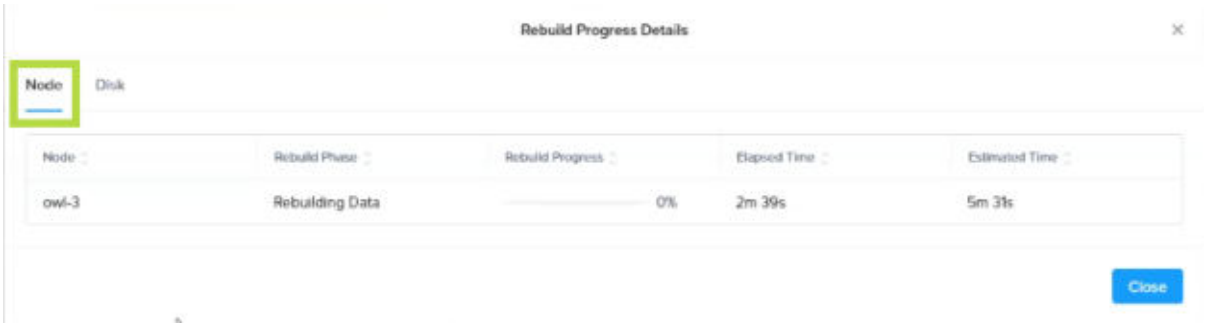

Figure 30: Overall Estimated Time

The following screenshot displays the rebuild information for a disk.

Figure 32: Rebuild Progress Details (Disk)

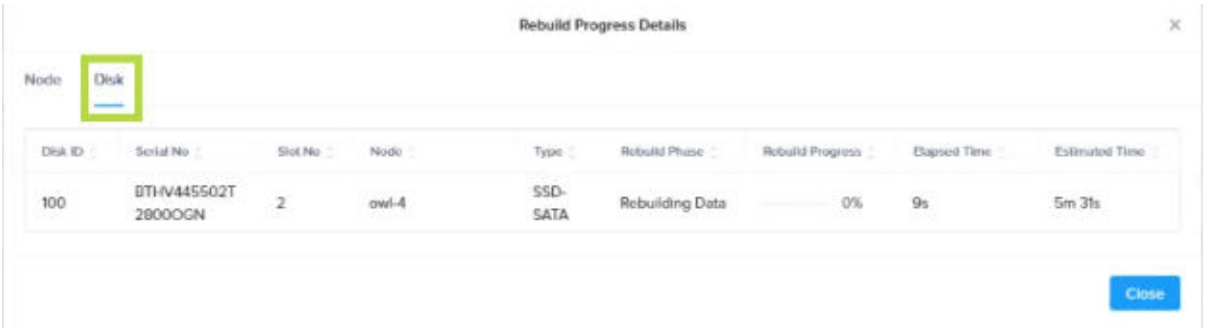

Example 2: Two Disks

The following screenshot displays the overall rebuild progress percentage and ETA for two disks.

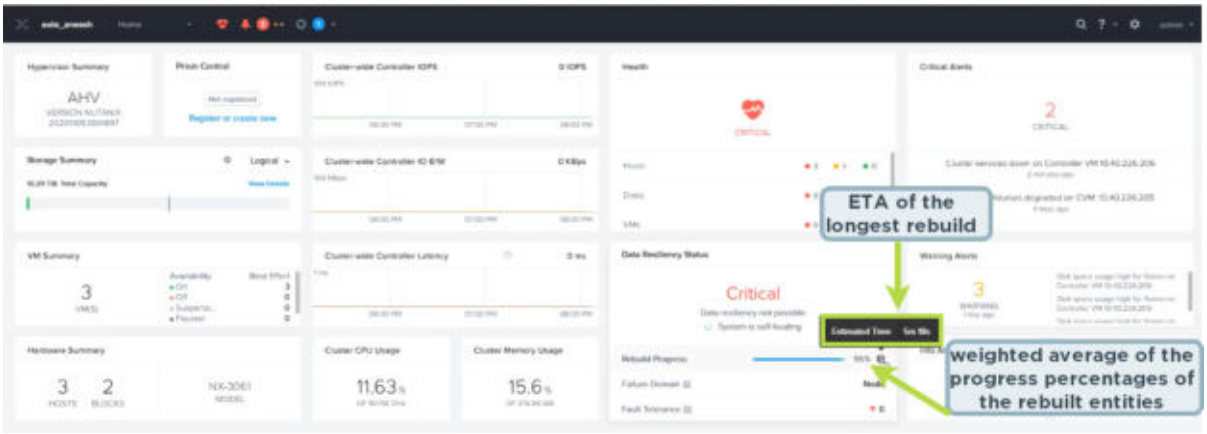

Figure 33: Overall Estimated Time (Data Rebuild)

The following screenshot displays the rebuild information for both disks.

Figure 34: Rebuild Progress Details (Disks)

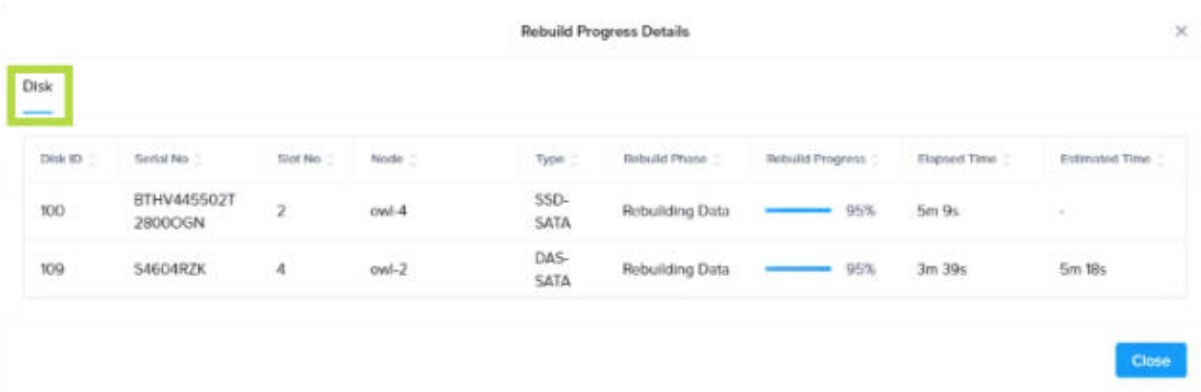

Rebuild Phases

A data rebuild has the following phases.

- **Rebuilding Data**. Indicates that the data rebuilt progress up to 100% . This includes time to rebuild data and also validate that all data has actually been rebuilt.
- **Data Rebuild Complete**. Indicates the success of the data rebuild operation.
- **Aborted**. Indicates an aborted rebuild operation. The rebuild operation aborts when a node that went offline comes back online.

Data Rebuilt

The following data rebuilds when a rebuild scenario triggers.

> Note: The data rebuild percentage does not account for metadata rebuild. For example, if the metadata rebuild is at 10% and the data rebuild is 100%, the **Data Resiliency Status** widget displays the rebuild percentage as 100%. The metadata rebuild progress currently appears on the **Tasks** page of the Prism Element web console.

- Extent store data
  - Extent groups on victim disk.
- Oplog episodes
  - Unflushed oplog episodes on the victim disk.
- Near Sync LWS episodes.

> Note: The progress percentage of the extent store data rebuild restricts to 95% until the rebuild completes for oplog and NearSync LWS episodes

Monitoring Disk Rebuild Progress

About this task

Perform the following procedure to monitor the data rebuild progress and overall rebuild completion ETA for a disk. Curator periodically sends the data rebuild information to Insights. Prism Element queries this data rebuild information from the Insights DB. The rebuild progress bar for the disk reflects in the **Data Resiliency** widget within a maximum duration of 90 seconds.

Procedure

1. Go to the Prism Element **Home** > **Data Resiliency Status** widget.
2. Hover on the (?) symbol to view the rebuild completion ETA for the current generation.

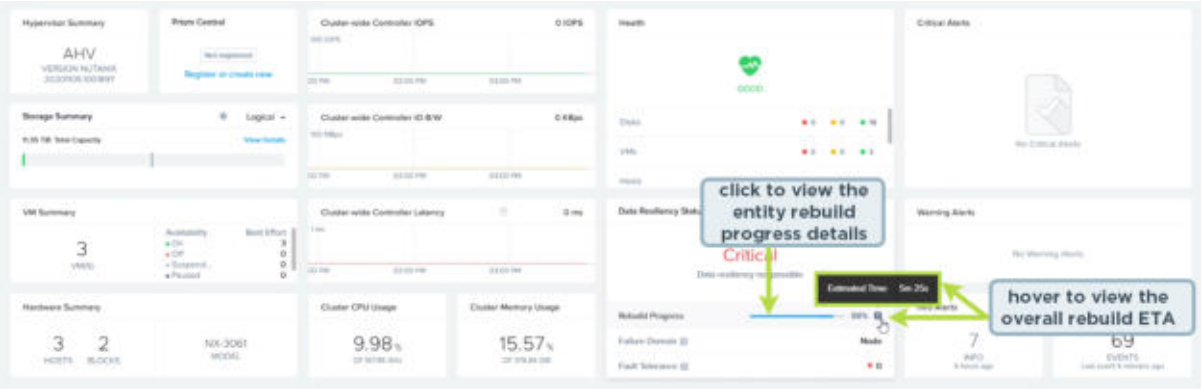

Figure 35: Overall Estimated Time and Progress Percentage (Data Rebuild)

3. Monitor the **Rebuild Progress** progress bar to view the rebuild completion percentage.
4. Click anywhere on the **Rebuild Progress** field to view the following rebuild metrics for the individual entities.

   - **Disk ID**. Displays the disk identification number (ID number).
   - **Serial No**. Displays the disk serial number (serial number).
   - **Slot No**. Displays the slot number associated with the disk (slot number).
   - **Node**. Displays the node name.
   - **Type**. Displays the disk type (tier name). The Nutanix platform 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 platform type [SSD-PCIe | SSDSATA | DAS-SATA].

- **Rebuild Phase**. Displays the current data rebuild status. See Home Dashboard on page 53 for information about the data rebuild phases.
- **Rebuild Progress**. Displays the rebuild progress percentage.
- **Elapsed Time**. Displays the time elapsed since the data rebuild operation started.
- **Estimated Time**. Displays the estimated time for the data rebuild operation to complete.

The **Rebuild Progress Details** dialog box appears.

The following screenshot displays the data rebuild progress details after you remove the disk.

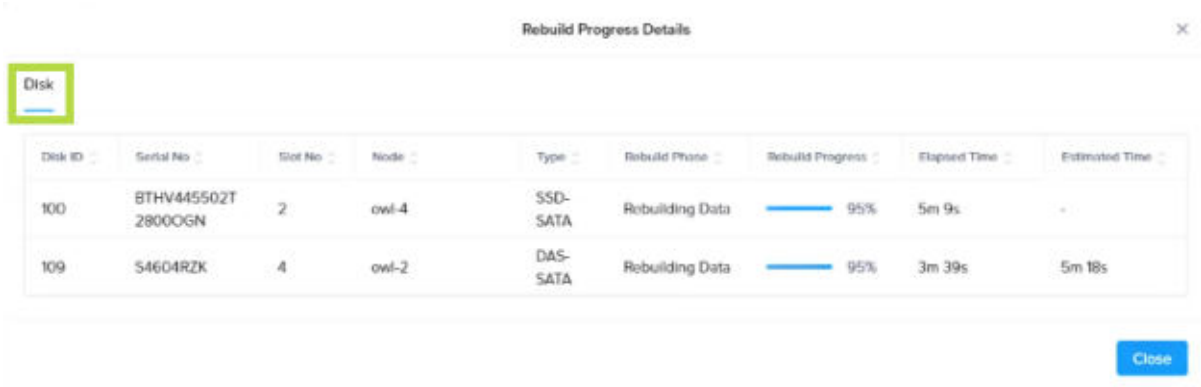

Figure 36: Rebuild Progress Details (Disk)

5. You can verify the rebuild event status through the **Tasks** option.

Monitoring Node Rebuild Progress

About this task

Perform the following procedure to monitor the data rebuild progress and overall rebuild completion ETA for a node. Curator periodically sends the data rebuild information to Insights. Prism Element queries this data rebuild trigger from the Insights DB. The rebuild progress bar for the node reflects in the **Data Resiliency** widget within a maximum duration of 90 seconds.

> Note: The offline node progress appears only if the Stargate health is down.

Procedure

1. Go to the Prism Element **Home** > **Data Resiliency Status** widget.

2. Hover on the (?) symbol to view the rebuild completion ETA for the current generation.

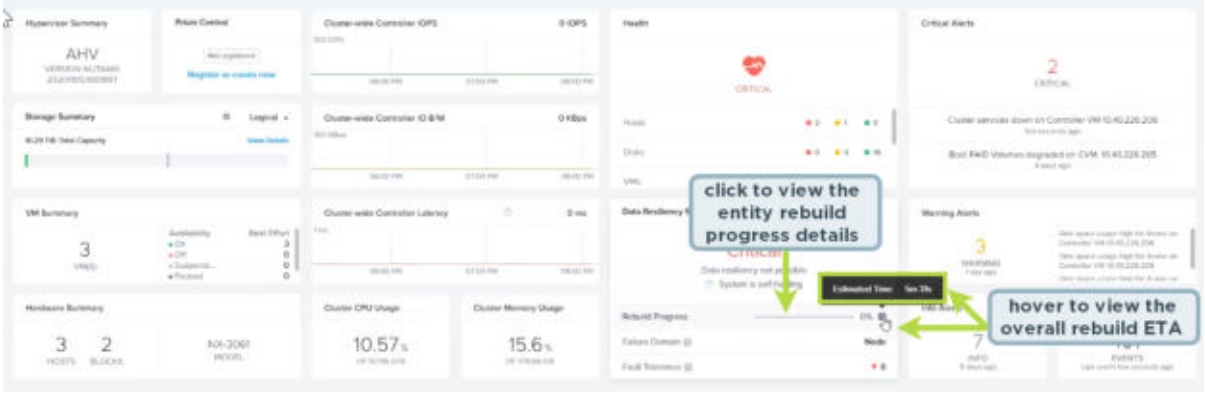

Figure 37: Overall Estimated Time and Progress Percentage (Data Rebuild)

3. Monitor the **Rebuild Progress** progress bar to view the rebuild completion percentage.
4. Click anywhere on the **Rebuild Progress** field to view the following rebuild metrics for the individual entities.

   - **Node**. Displays the node name.
   - **Rebuild Phase**. Displays the current data rebuild status. See Home Dashboard on page 53 for information about the data rebuild phases.
   - **Rebuild Progress**. Displays the rebuild progress percentage.
   - **Elapsed Time**. Displays the time elapsed since the data rebuild operation started.
   - **Estimated Time**. Displays the estimated time for the data rebuild operation to complete.

   The **Rebuild Progress Details** dialog box appears.

   The following screenshot displays the data rebuild progress details after you remove the disk.

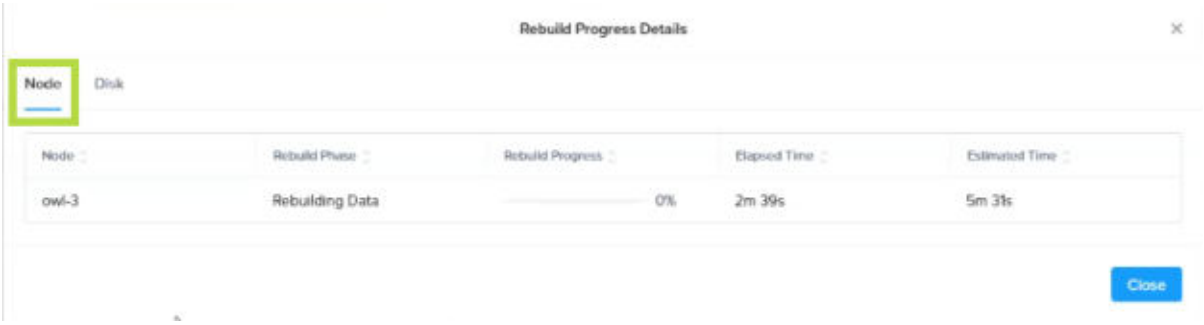

Figure 38: Rebuild Progress Details (Node)

5. You can verify the rebuild event status through the **Tasks** option.

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

Hypervisor Type	Overview, VM, and Storage	Hardware	Analysis
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 or modify cluster parameters (cluster name, virtual IP address, and data services IP address) through the **Settings** menu.

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 read-only fields for the cluster UUID (universally unique identifier), ID, incarnation ID, subnet, and encryption status values. It also contains editable fields for cluster name, FQDN, virtual IP configuration, and iSCSI data services IP address. 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.

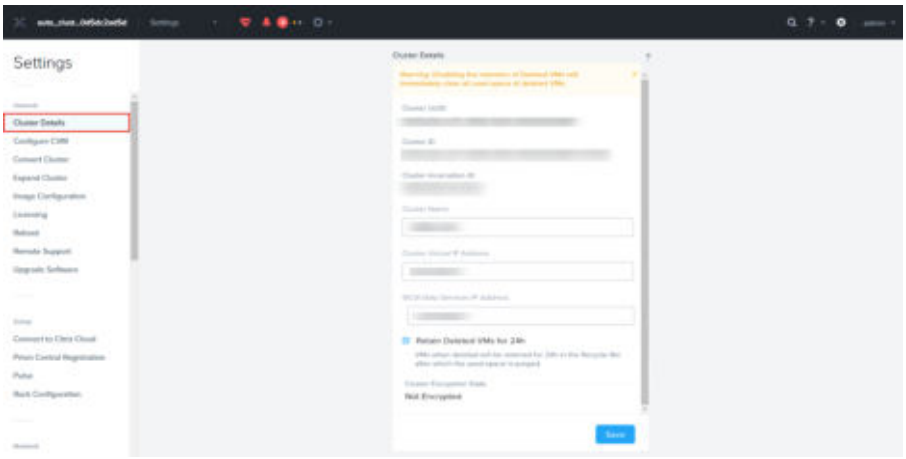

Figure 39: Cluster Details Window

2. In the **Cluster Name** field, enter (or update) a name for the cluster.
   The default name is simply Unnamed. Providing a custom name is optional but recommended.

3. In the **FQDN** field, enter the fully qualified domain name (FQDN) for the cluster.

   The virtual IP (next field) always resolves to the current Controller VM leader in the cluster. As a result, there is no load balancing of requests. In contrast, specifying a domain name allows requests to be balanced across the Controller VMs to improve performance and resiliency. However, this requires an administrator configure the domain name in the DNS server as the set of IP addresses associated with the Controller VMs in the cluster. For larger clusters, the domain name should resolve to at least three Controller VMs, but it is

recommended that all the Controller VMs be specified to maximize the performance and resiliency benefits.

When a FQDN is specified, it is the first choice for IP address resolution. Make sure the virtual IP (next field) is not mapped to multiple DNS entries as this can lead to resolution problems. If the FQDN fails to resolve in DNS, the existing Controller VM IP addresses are used.

4. In the **Virtual IP** field, enter an IP address that will be used as a virtual IP for the cluster.

   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.

   > 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 67 for more information.

   > Note: You are not required to configure either a virtual IP or a domain name, but configuring both provides better addressing flexibility and cluster performance.

5. In the **Virtual IPv6** field, enter an IPv6 address that will be used as a virtual IPv6 address for the cluster.

   A Controller VM runs on each node and has its own IPv6 address, but this field sets a logical IPv6 address that always connects to one of the active Controller VM in the cluster (assuming at least one is act

   ive), which removes the need to enter the address of a specific Controller VM. The virtual IPv6 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.

6. In the **ISCSI Data Services IP Address** field, enter (or update) an IP address for use with Nutanix Volumes and other data services applications.

   > 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 68 for more information.

7. To disable the Recycle Bin, clear (unselect) **Retain Deleted VMs for 24h**.

   When you disable the Recycle Bin, AOS automatically deletes any entities in the Recycle Bin after 24 hours. Any entities you delete after disabling the Recycle Bin are marked for deletion as soon as possible and are not stored in the storage container Recycle Bin folder.See also Recycle Bin on page 179.

8. To enable the Recycle Bin, select **Retain Deleted VMs for 24h**.

   If you enable the Recycle Bin and then delete a guest VM or volume group vDisk, it retains its contents (deleted vDisks and configuration information) for up to 24 hours, unless the cluster free storage space reaches critical thresholds.

9. 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 383.
- 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 66.

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 logon page includes background animation, and Prism logs out the users automatically after the users have been 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.

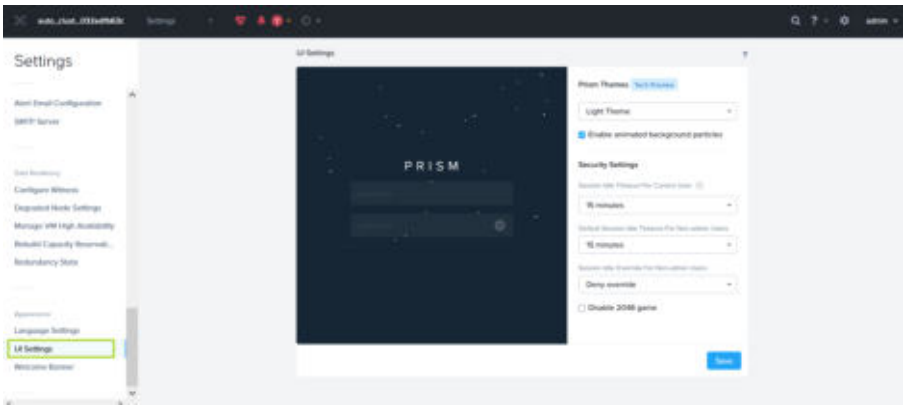

Figure 40: UI Settings Window

2. To set different variations of the Prism UI background themes, select any of the following options from the **Prism Themes** drop-down menu. The UI changes cosmetically to reflect the selection.

Important:

- The Prism themes feature is currently in technical preview. You may encounter visual anomalies in few settings/views where the prism themes are not applied.

For example, Licensing Settings. Nutanix recommends that you register a case on the support portal to report any anomalies.

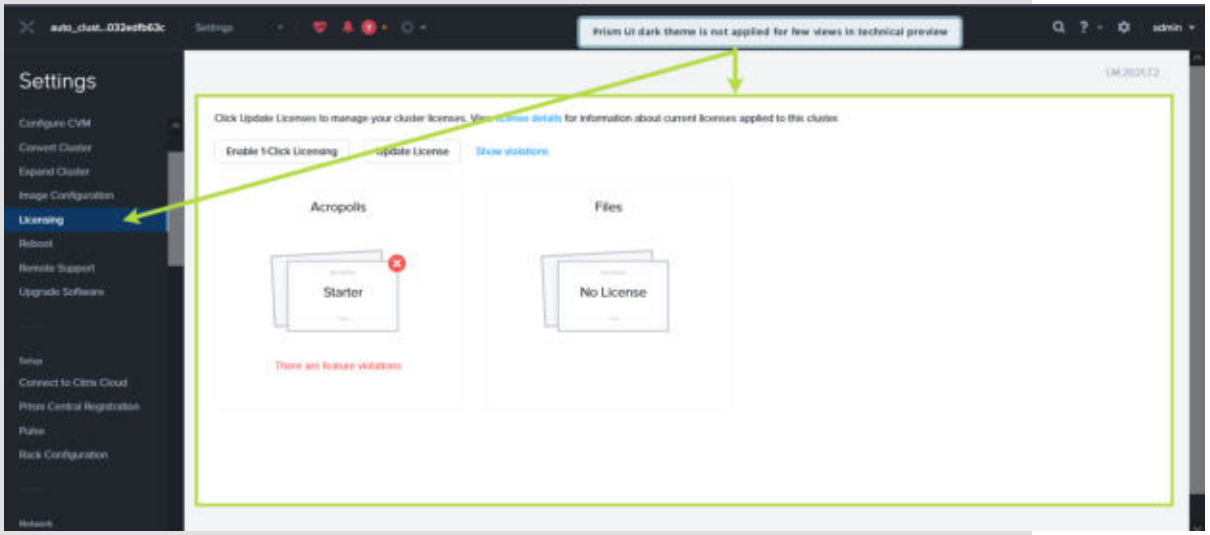

Figure 41: Example: Licensing View (Dark Theme Exception)

- You are required to Save your selection and refresh any open Prism tabs after changing a background theme.
- The Prism themes setting is not a universal feature between Prism Central and Prism Web Console. The background themes need to be configured in Prism Central and Prism Web Console respectively.

- Select **Light Theme** for light background with high contrast view. This is the default Prism background theme.
- Select **Dark Theme** for dark background with high contrast view. A pop-up appears prompting you to click **Continue** to proceed.

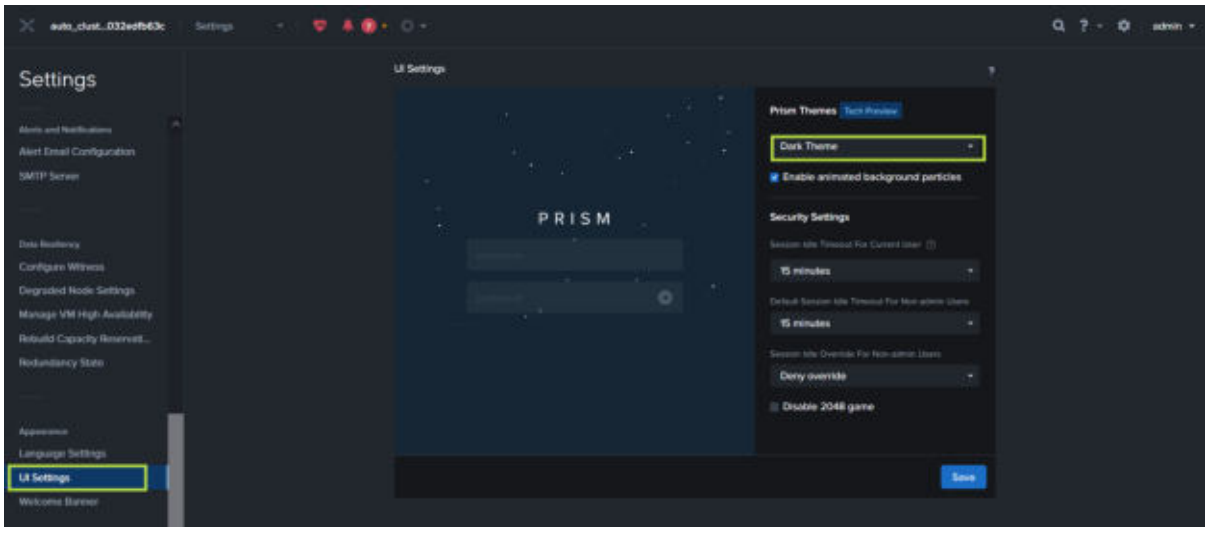

Figure 42: Prism Theme (Dark Mode)

- Select **Auto (OS defined)** to apply background themes defined in the OS. A pop-up appears prompting you to click **Continue** to proceed. For example, if you have defined **Dark** mode in the OS setting, the Prism UI honors the setting and sets a dark background theme.

3. To disable the logon page background animation, uncheck the **Enable animated background particles** box (or check it to enable).

   Unchecking the **Enable animated background particles** box in the Prism UI Settings dialog box disables the creation or drawing of particles entirely. This action stops the drawing of the particles on the Prism Element logon page.

   > Note: This setting is not persistent. In other words, if the Prism service restarts, this setting is lost and must be disabled again.

   Disabling the particles allows you to conserve critical CPU resources that are used in creating and maintaining the particles.

   > Note: Disabling or enabling this setting in Prism Web Console does not propagate to Prism Central or vice versa. The setting must be disabled in Prism Web Console and Prism Central UI separately.

   You can disable the particle animation from the logon page by clicking **Freeze space-time continuum!** at the right bottom of the logon page. This action stores a setting in the local browser to stop the animation. However, this action does not stop creation or drawing of the particles itself.

   > Note: You can enable the particle animation by clicking **Engage the warp drive!** .

4. To customize the theme, background color, title text, or blurb text on the logon page, do the following:

   - Click on the top bar (displayed in the following image) in the **UI Settings** dialog box and simultaneously press the option key on the MAC system or Alt key on the Windows system. Options for customizing the theme, title text, and blurb text are displayed.
   - Select the theme from the options displayed for **Theme**. You can change the HEX codes to create your own custom gradient background color for the logon page.
   - In the **Title Text** field, enter the text to create your custom title.
   - In the **Blurb Text** field, enter the text to create your custom blurb text. This text is displayed below the password field.

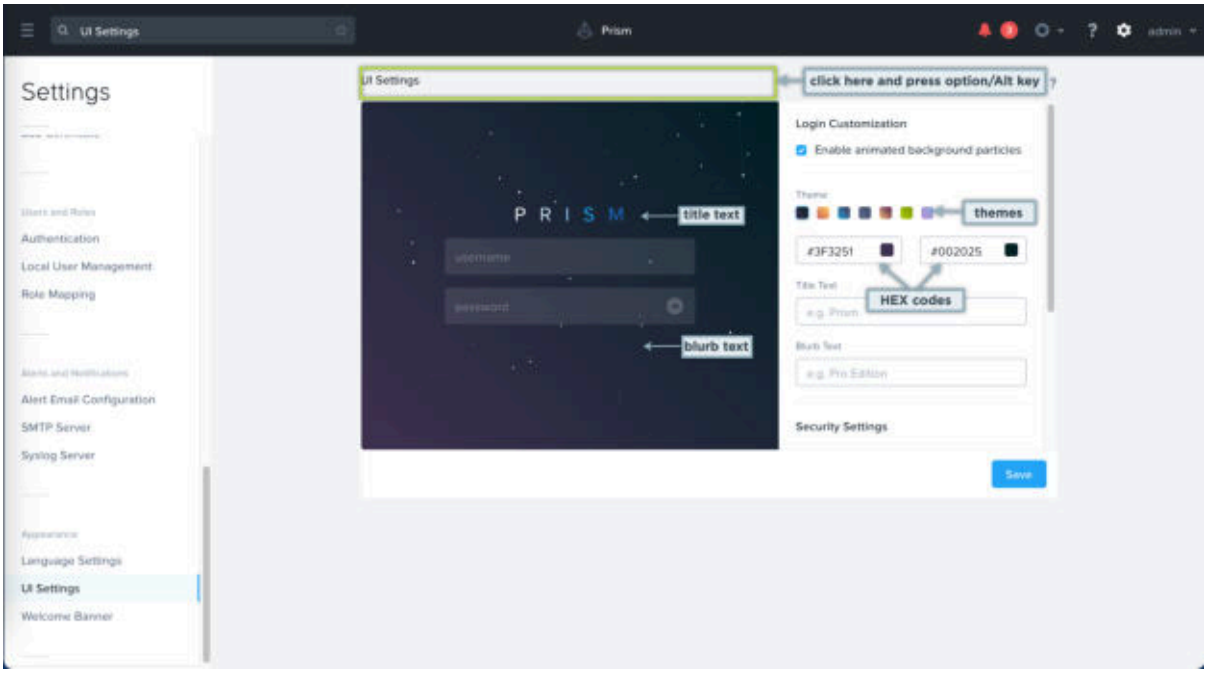

Figure 43: UI Settings Window for customizing the theme, title text, and blurb text

5. To configure session timeout, do the following under **Security Settings**:

   - Select the session timeout for the current user from the **Session Timeout For Current User** drop-down list.
   - Select the default session timeout for all 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.

6. Uncheck the **Disable 2048 game** box to disable the 2048 game.
7. Click **Save** to save your changes.

Finding the AHV Version on Prism Element

You can see the installed AHV version in the Prism Element web console.

About this task

To view the AHV version installed on the host, do the following.

Procedure

1. Log on to Prism Web Console
2. The **Hypervisor Summary widget** widget on the top left side of the Home page displays the AHV version.

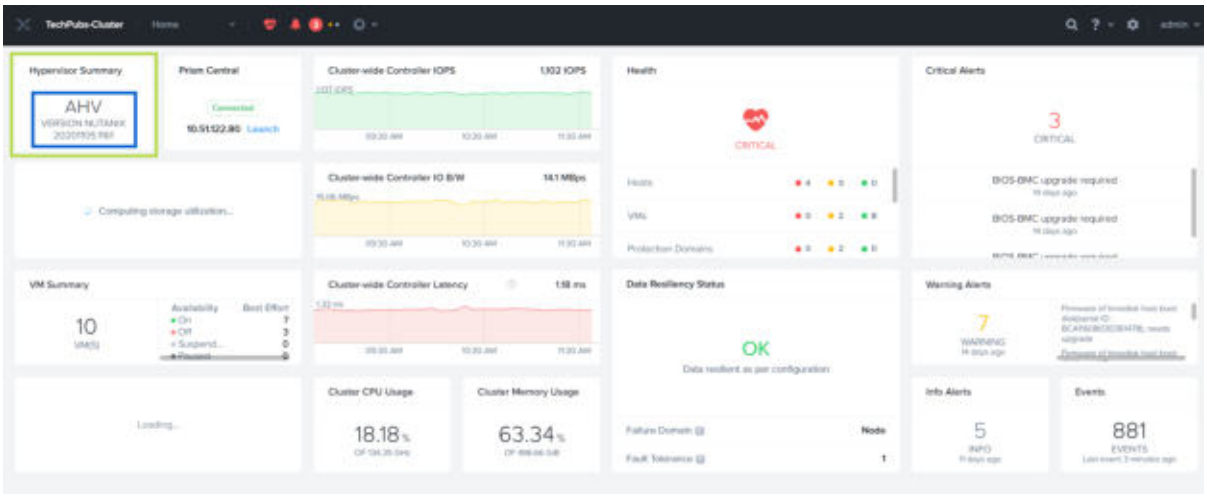

Figure 44: LCM Page Displays AHV Version

Finding the AOS Version Using Prism Element

To view the Nutanix AOS version running in the cluster, do the following:

Procedure

1. Log on to the Prism Web Console.
2. Click the drop-down next to the username in the top right corner.

3. Select the **About Nutanix** option.
   The About Nutanix window displays the AOS version along with the Nutanix cluster check (NCC) and Life Cycle Manager (LCM) version numbers.

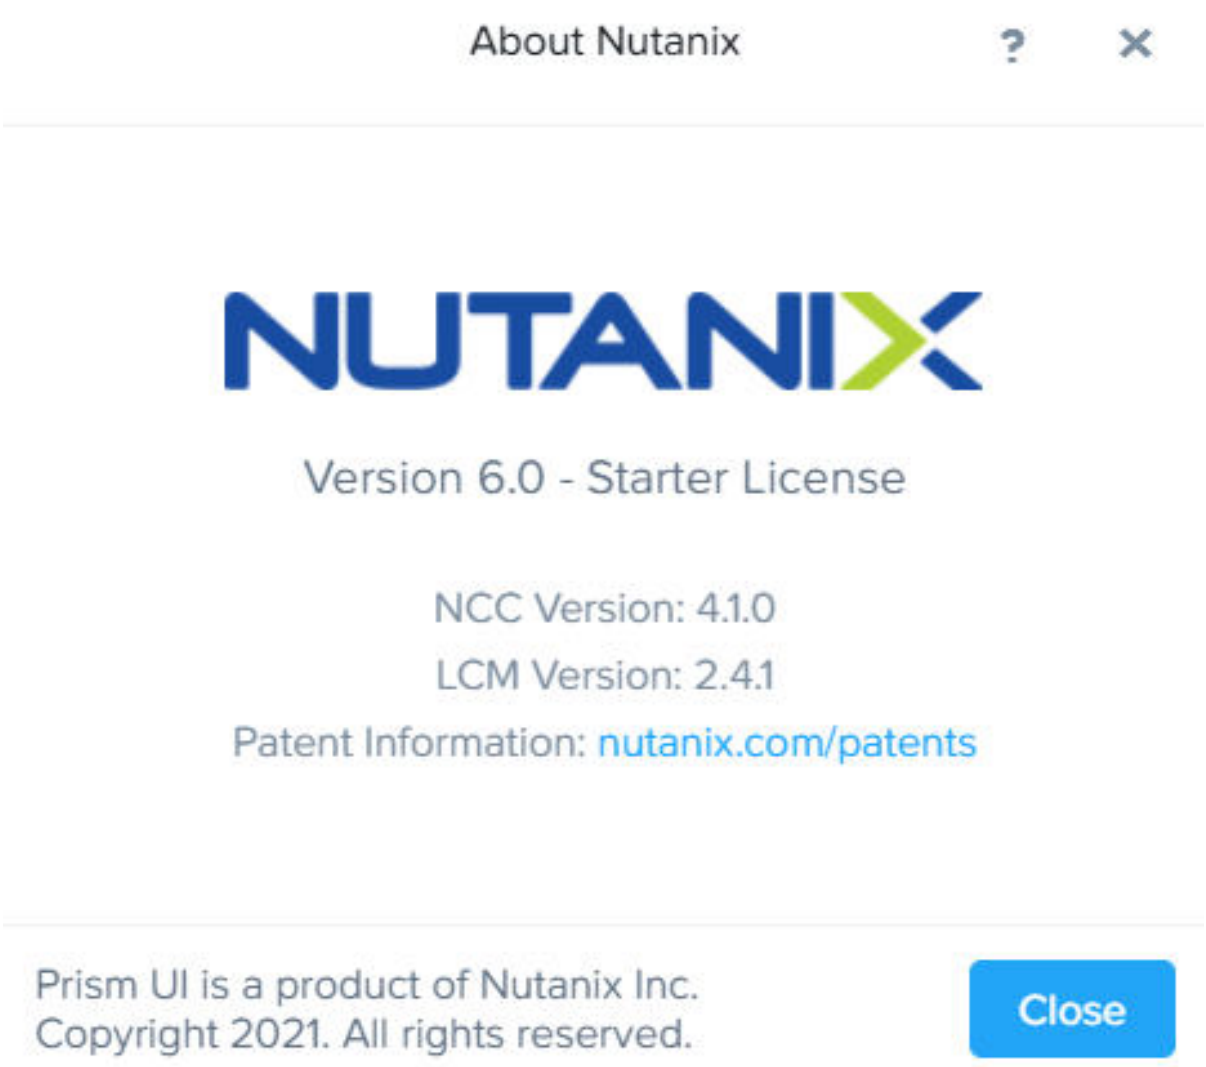

Figure 45: About Nutanix Window

4. Click the **Close** button to close the window.

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the Nutanix License Manager Guide for information about how to license your cluster, apply add-on licenses, and manage your licenses in a variety of situations.

Software and Firmware Upgrades

You can upgrade the AOS, non-AHV 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.
- Do not enable client authentication in combination with ECDSA certificates on a registered cluster since it causes interference when communicating with Prism Central.
- Ports 9440 and 80 need to be open in both directions between the Prism Central VM and all the Controller VMs (and the cluster virtual IP address if configured) in each registered cluster. For the complete list of required ports, see Port Reference.
- A cluster can register with just one Prism Central instance at a time. To register with a different Prism Central instance, first unregister the cluster.

About this task

To register a cluster with Prism Central, do the following:

Procedure

1. Log in to the Prism web console on the target cluster as the user admin.
2. In the Health dashboard, from the **Actions** drop-down menu, select **Run Checks**. to run Nutanix Cluster Checks.
3. Do one of the following:

   » On the **Home** dashboard, click **Register or create new** from the **Prism Central** widget.

   » Click the gear icon and select **Prism Central Registration** from the **Settings** menu.

Figure 46: Home Dashboard

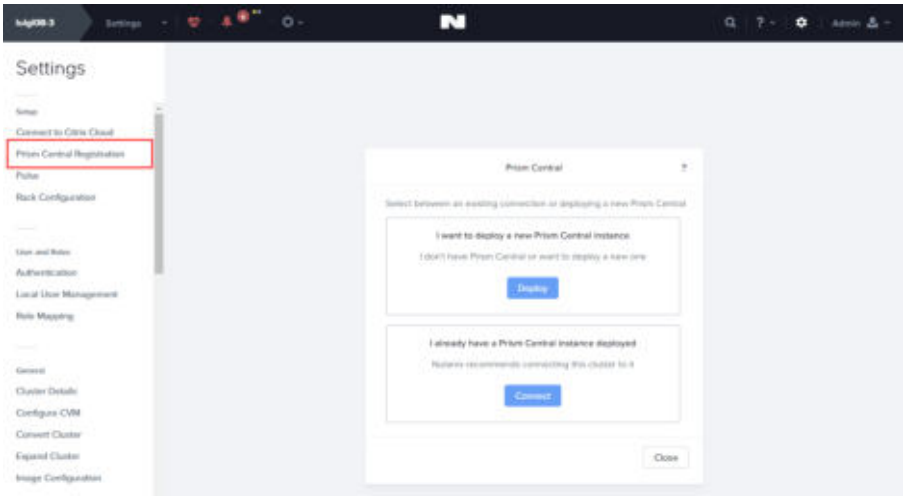

Figure 47: Settings Menu

4. In the first screen of the Prism Central dialog box, click the **Connect** button.

   This screen includes two options, **Deploy** and **Connect**. This procedure describes how to connect to an existing Prism Central instance. See *Installing Prism Central (1-Click Method)*

section in Acropolis Upgrade Guide for instructions on how to deploy a new Prism Central instance.

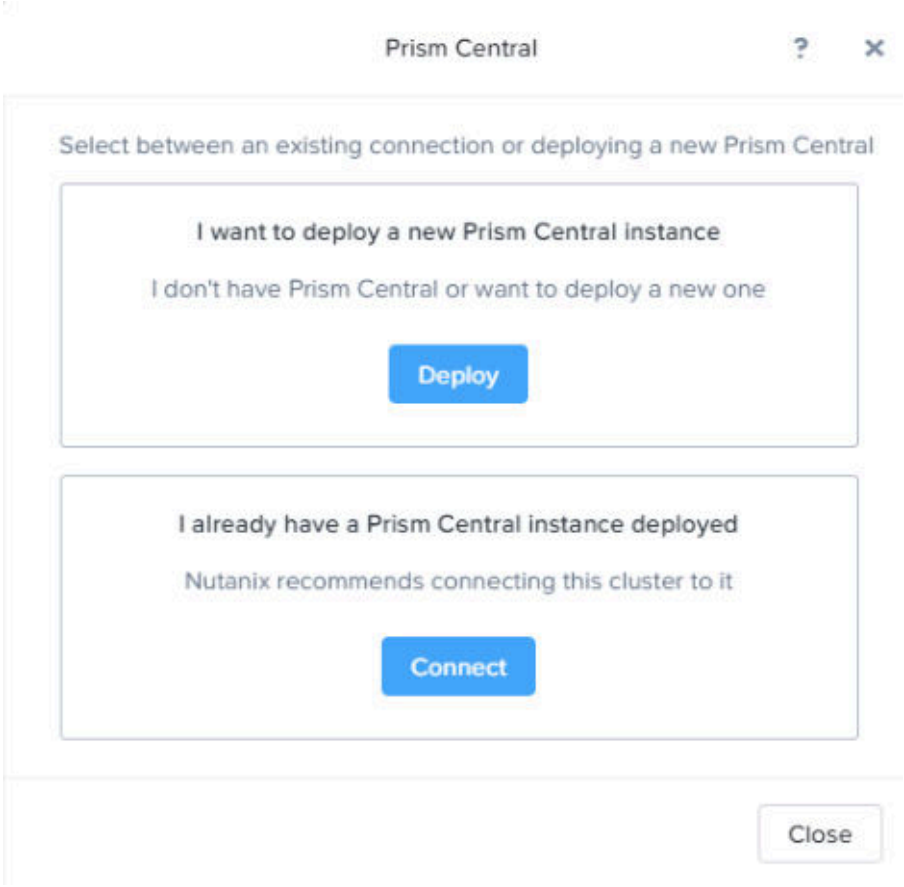

Figure 48: Prism Central Dialog Box (deploy/connect screen)

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

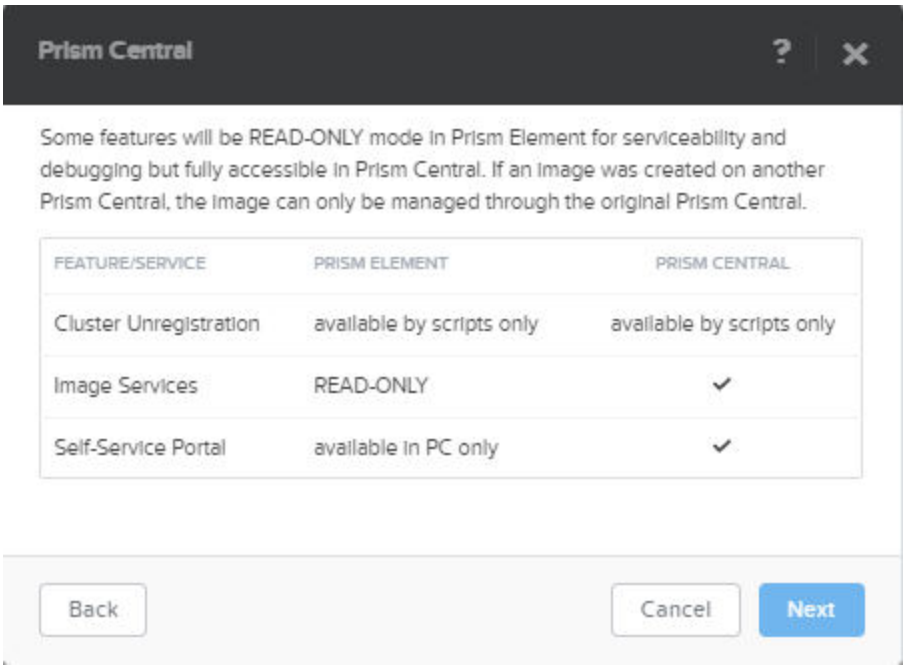

Figure 49: Prism Central Dialog Box (services screen)

6. In the registration screen, do the following in the indicated fields:
   a. **Prism Central IP**: Enter the IP address of the Prism Central VM.
   b. **Port**: The default port number is 80. This is an optional field. For the complete list of required ports, see Port Reference.
   c. **Username**: Enter admin as the Prism Central user name.
   d. **Password**: Enter the Prism Central admin user password.

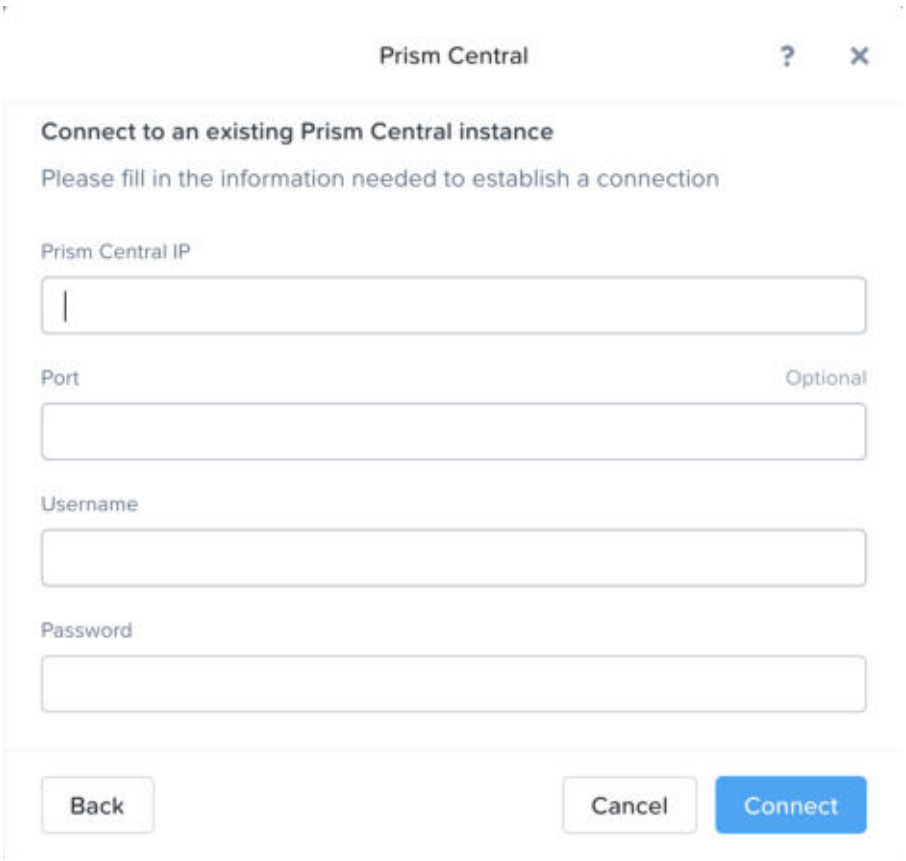

Figure 50: Prism Central Dialog Box (registration screen)

7. When all the fields are correct, click the **Connect** button to save the values and close the window.
   This registers the cluster on the specified Prism Central VM and allows the passing of information between the cluster and Prism Central.

   Note: The user credentials provided when registering a cluster (Prism Element) with Prism Central are only used once. After registration, modifying the admin password would not impact any communication between Prism Element and Prism Central.

Unregistering from Prism Central

About this task

To unregister a cluster from an existing Prism Central instance, do the following:

Note: Unregistering a cluster through the Prism GUI is no longer available. This option was removed to reduce the risk of accidentally unregistering a cluster because several features

(including role-based access control, application management, micro-segmentation policies, and self-service capability) require Prism Central to run your clusters. If a cluster is unregistered from Prism Central, not only will these features not be available but the configuration for them may also be erased. Therefore, only the following procedure is available to unregister a cluster. See KB 4944 for additional details if you have enabled Self Service, Calm, or other special features in Prism Central.

Procedure

1. Log on to any Controller VM of the registered cluster through an SSH session.
2. Run the **cluster status** command and verify that all services are in a healthy state.

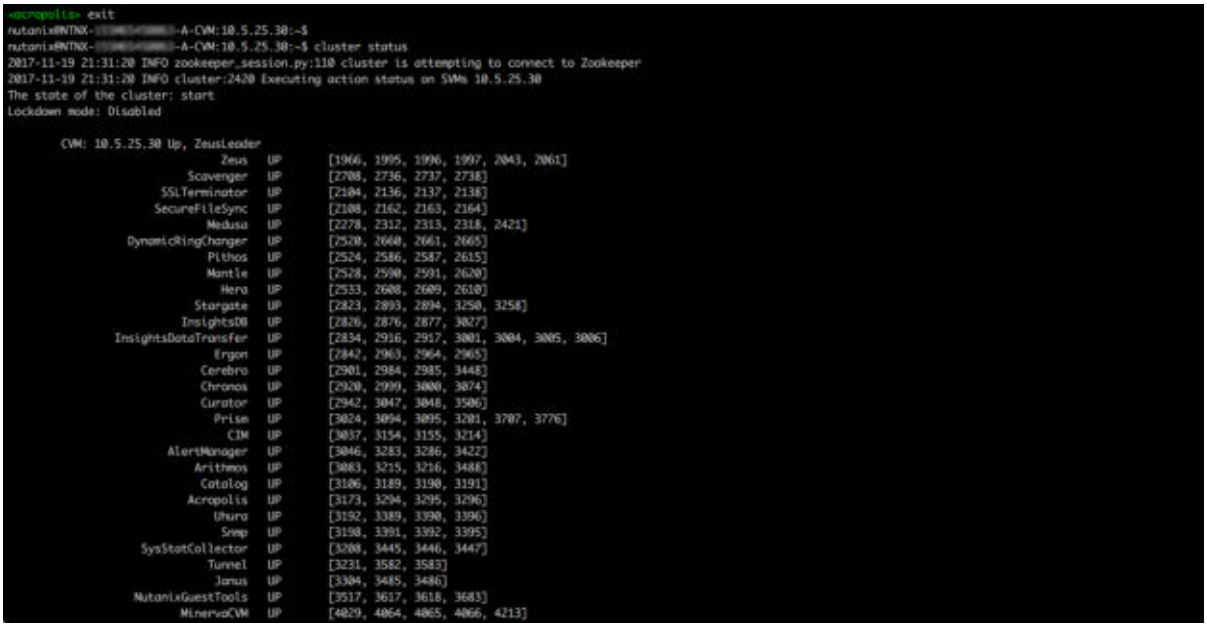

Figure 51: Cluster Status Command

3. Unregister the cluster from Prism Element.

   nutanix@cvm$ ncli multicluster remove-from-multicluster external-ip-address-or-svm-ips=*pc-name-or-ip* username=*pc-username* password=*pc-password* force=true

   Replace *pc-name-or-ip* with the Prism Central name or IP address and *pc-username* and *pc-password* with the login credentials for your Prism Central administrator account. This step can take some time (though typically just a few seconds). To check if the unregistration has completed, run the following command.

   nutanix@cvm$ ncli multicluster get-cluster-state

4. Get the UUID for the cluster.

   ```
   nutanix@cvm$ ncli cluster info
   ```

   Find the "Cluster UUID" value in the displayed information.

   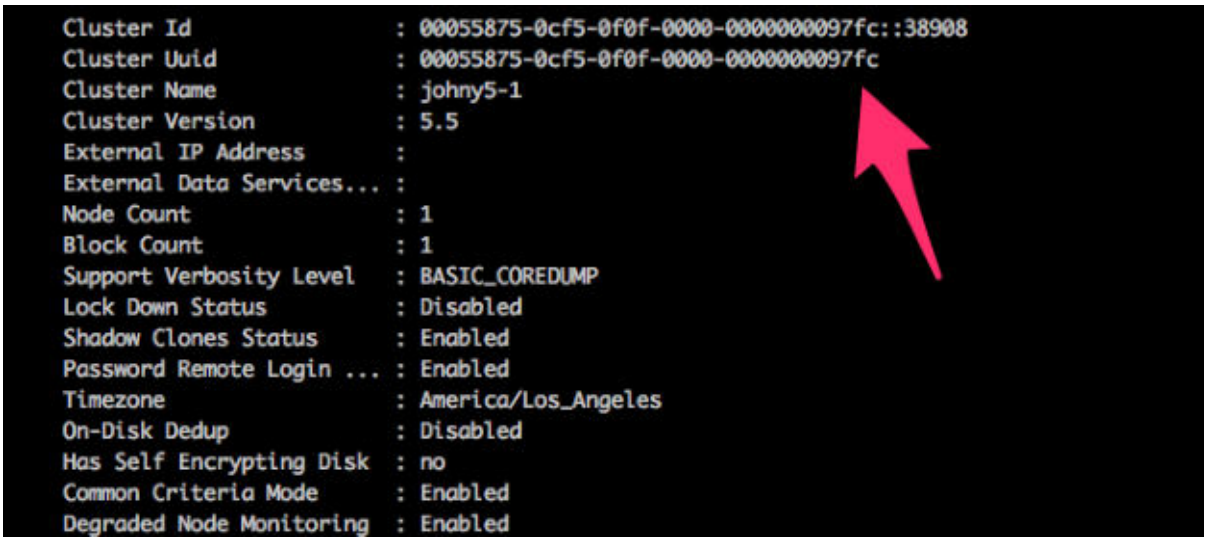

   Figure 52: Cluster Info Output

5. Log on to the Prism Central VM through an SSH session (as the "nutanix" user) and do the following:

   a. Run the unregistration clean-up script.

      ```
      [pcvm]$ python /home/nutanix/bin/unregistration_cleanup.py uuid
      ```

      Replace *uuid* with the value you obtained in step 4. This script removes all remaining registration information about that cluster and completes the unregistration process for the Prism Central VM.

   b. Get the UUID for Prism Central.

      ```
      [pcvm]$ ncli cluster info
      ```

      Find the "Cluster UUID" value in the displayed information (see step 4), which in this case is the UUID for Prism Central.

6. Go back to the Controller VM and run the **unregistration_cleanup.py** script to complete the unregistration process on the cluster.

   ```
   nutanix@cvm$ python /home/nutanix/bin/unregistration_cleanup.py uuid
   ```

   In this case the *uuid* is the Prism Central UUID obtained in step 5b.

What to do next

After performing these steps you can re-register the cluster with a new or re-created Prism Central instance.

If the clean up process does not complete successfully, try the following:

- Check the logs to indicate if there are any input errors when calling the script. The logs for the unregistration cleanup script can be found under **~/data/logs/unregistration_cleanup.log**.
- If errors occur during script execution, run the **cluster status** command and check that the cluster services are up and running. Rerun the script and check if it succeeds.

Note: If Prism Self Service is enabled and the unregistered cluster is set as the default cluster, you must update the default project to select a new (registered) cluster. See the topic *Modifying A Project* in the Prism Central Guide.

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.
- Ensure the port TCP port 2100 is open from the Prism Element cluster to the Prism Central VM IP address. For the complete list of required ports, see Port Reference.
- Ensure network connectivity between the VM VLAN and portgroup of the Prism Element cluster Controller VM and the Prism Central VM VLAN and portgroup.
- 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 Element web console for a cluster as the user admin.
2. Run NCC.
   a. In the Health dashboard, from the **Actions** drop-down menu, select **Run Checks**.
   b. Select **All checks**. Select this option to run all the checks at once.

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 53: Home Dashboard

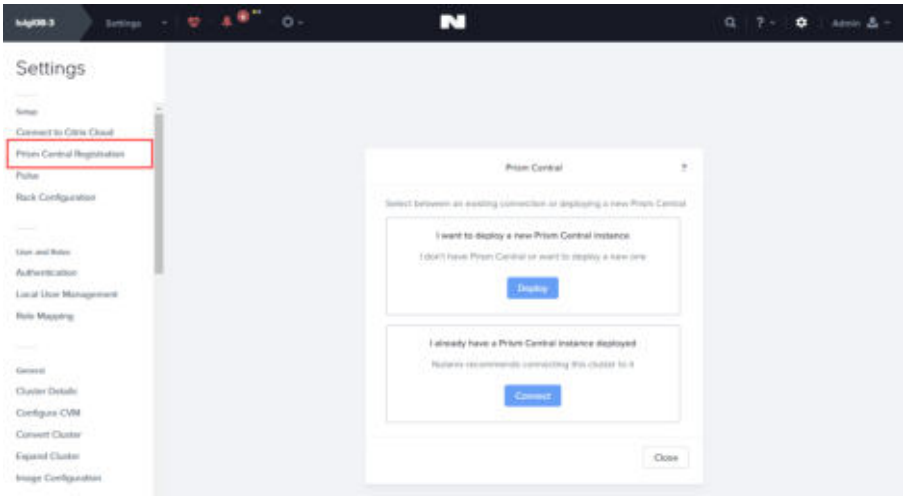

Figure 54: Settings Menu Launched from Gear Icon

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 "Register (Unregister) Cluster with Prism Central"

topic 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.

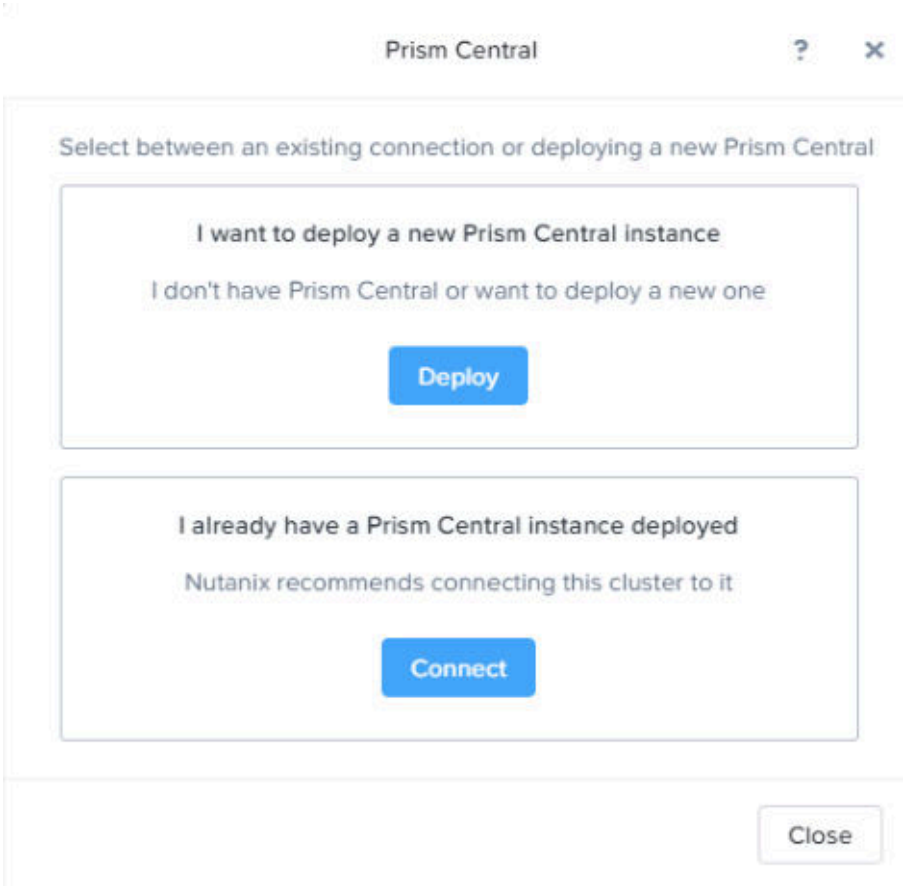

Figure 55: Prism Central Dialog Box (deploy/connect screen)

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.

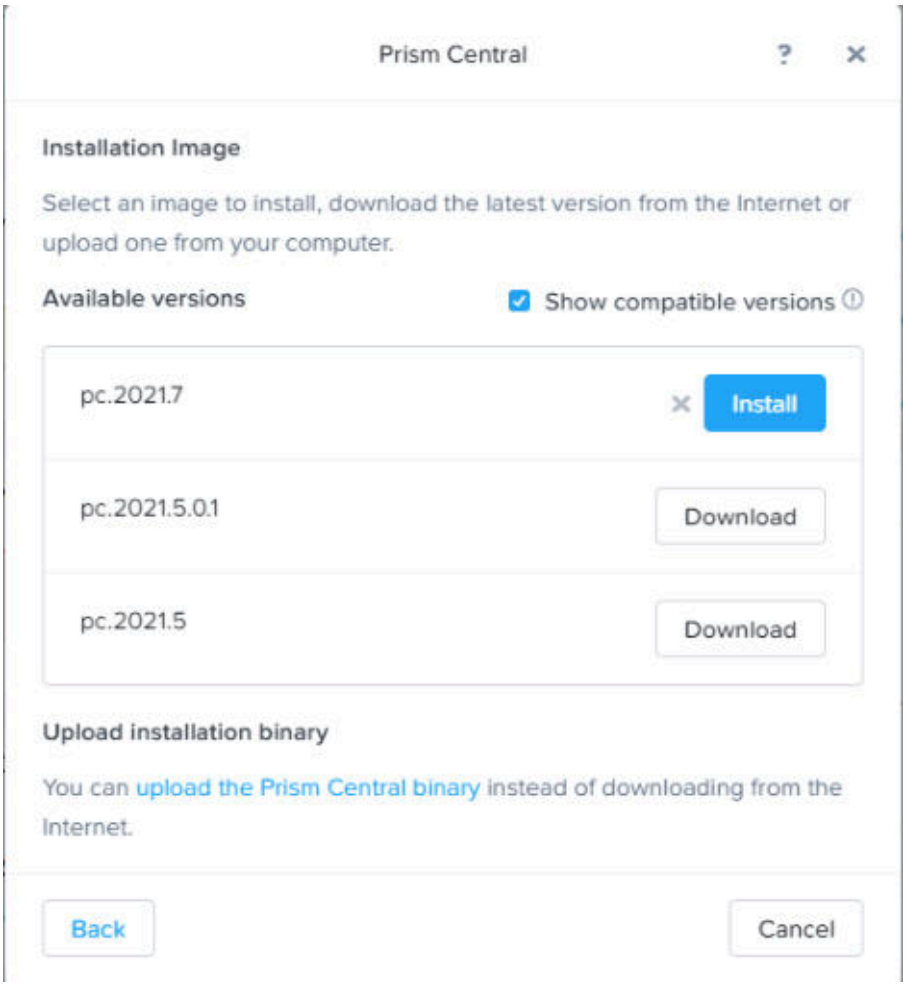

Figure 56: Prism Central Dialog Box (installation image screen)

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 "small", "large", or "x-large" in size. Thus, you may choose from six configurations. (The x-large size is available only on

AOS 6.0.1 or later.) For information about supported capacities (guest VMs, nodes, and clusters) for each of the six configurations, see the "Prism Central Scalability" section in the release notes for the Prism Central version to be installed.

A small size VM is created with a single vDisk, but large and x-large VMs are created with three additional vDisks in order to shard Cassandra metadata to improve performance. While the listed disk capacity of a Prism Central VM includes all the vDisks, the additional vDisks do not impact overall capacity. During the installation you can monitor progress ("VM disk attach" tasks) from the Tasks dashboard.

Table 10: Prism Central Instance Configurations

VM Size	1-VM Instance	3-VM Instance
small	6 vCPUs, 26 GB of memory, and 500 GiB of storage.	18 vCPUs, 78 GB of memory, and 1500 GiB of storage (combined)
large	10 vCPUs, 44 GB of memory, and 2500 GiB of storage	30 vCPUs, 132 GB of memory, 7500 GiB of storage (combined)
x-large	14 vCPUs, 60 GB of memory, and 2500 GiB of storage.	42 vCPUs, 180 GB of memory, and 7500 GiB of storage (combined)

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
      - **X-LARGE** - manage up to 12500 VMs (same as large) with more vCPUs and memory to account for the ANC (Atlas Network Controller) service

   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.

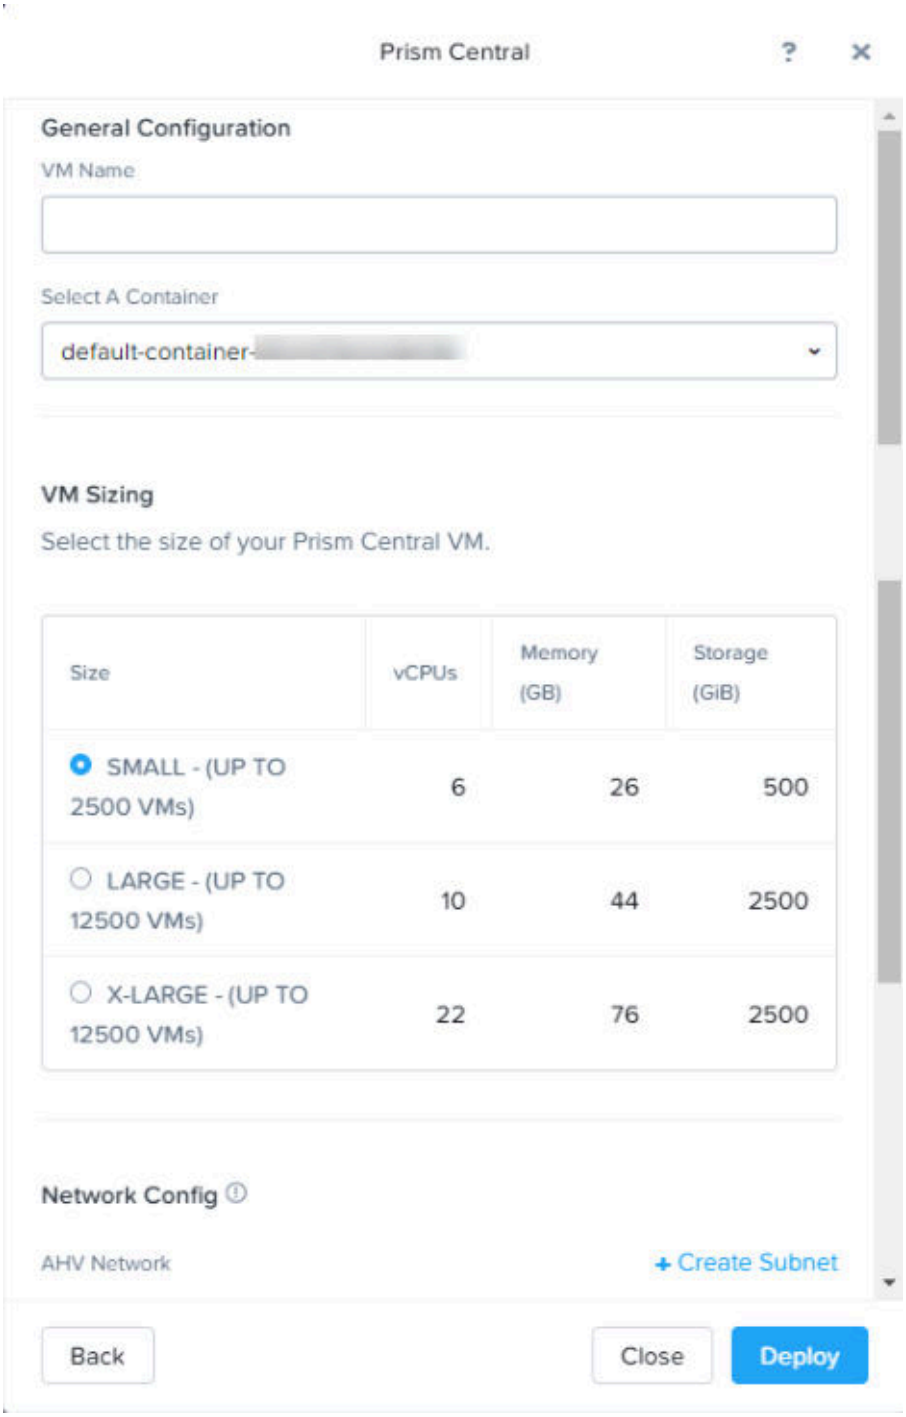

Figure 57: Prism Central Dialog Box (1-VM general configuration screen)

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** (small) - manage up to 5000 VMs
      - **Up to 25K VMs** (large) - manage up to 25000 VMs
      - **Up to 25K VMs** (x-large) - manage up to 12500 VMs (same as large) with more vCPUs and memory to account for the ANC (Atlas Network Controller) service

   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 58: Edit IP Field

      2. Repeat this step for the other two VMs.

   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.

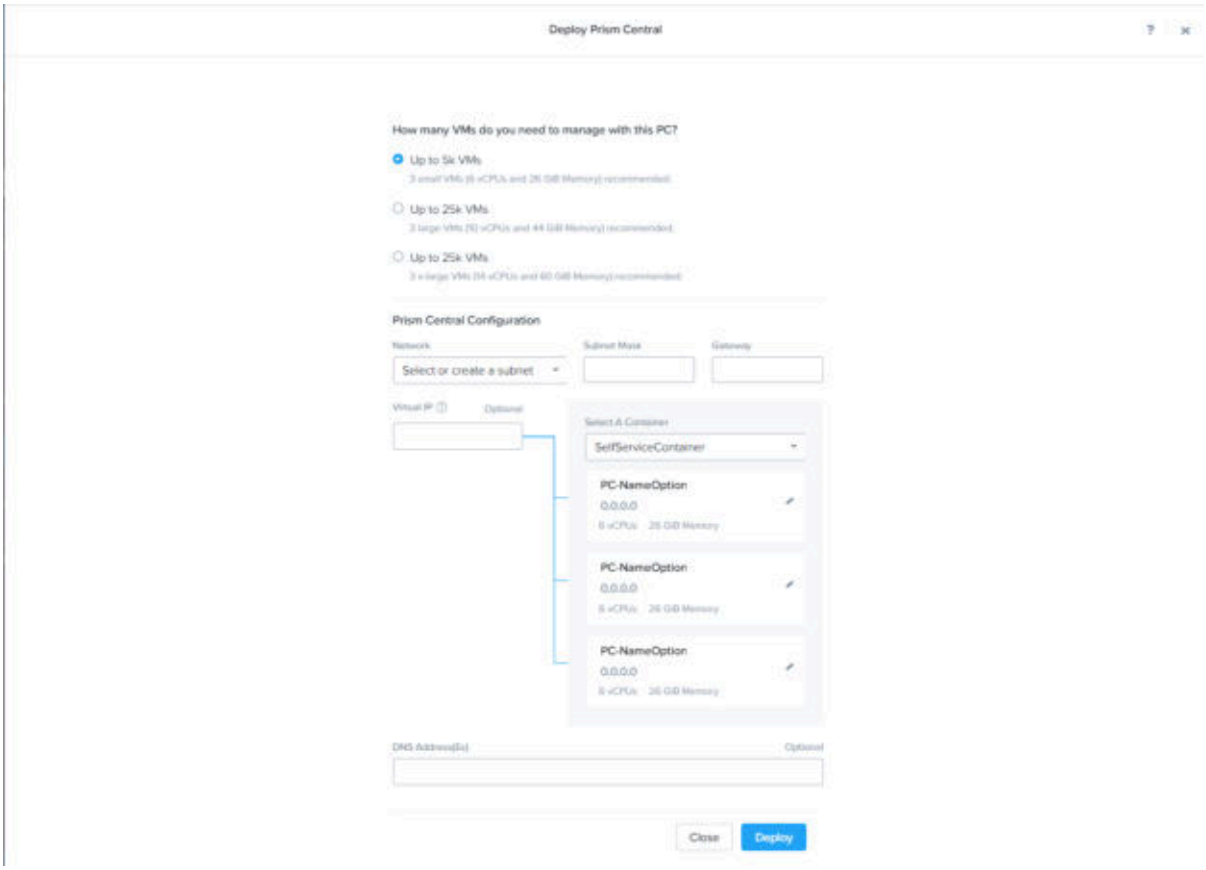

Figure 59: Deploy Prism Central Page (3 VM)

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 60: Tasks Page (Prism Central deployment)

Figure 61: 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. See "Register (Unregister) Cluster with Prism Central" in the *Prism Central Guide* for instructions on how to connect to an existing Prism Central instance.

Prism Central Disaster Recovery

You can protect your Prism Central deployment (including scaleout deployment) and recover it when an event like a natural calamity, network outage, or power failure occurs and makes Prism Central unavailable. This feature enables you to automate continuous backup of your Prism Central VM (and various service configurations within it) to up to three AHV or ESXi clusters registered to Prism Central. The seamless 1-Click recovery restores important service configurations along with your Prism Central deployment.

The following services can be protected and data associated with them recovered.

- Leap
- Flow microsegmentation
- Prism Pro (AIOps)
- VM management
- Cluster management
- Identity and Access Management (IAMv1)
- Categories
- Networking
- Licensing

Although the following services continue to run, they are not protected and therefore data associated with them not recovered.

- Flow networking
- Calm
- LCM
- Karbon
- Objects
- Files
- Catalog
- Images
- VM templates
- Reporting templates

For Leap Recovery Jobs in progress, see this Note.

> Note: Identity and Access Management (IAMv2): Prism Central disaster recovery does not support automatic recovery of data for Identity and Access Management (IAMv2). Therefore, the data associated with IAMv2 must be manually recovered using the **restore_iamv2.sh** script. This script uses the commands used for microservices infrastructure.

Therefore, ensure that you enable microservices infrastructure before running this script on the recovered Prism Central VM.

```
nutanix@pcvm$ cd /home/nutanix/cluster/bin
nutanix@pcvm:~/cluster/bin$ sh restore_iamv2.sh
```

Run this script within one hour from enabling the microservices infrastructure.

The script output allows you to select a backup copy from a list of backups available displayed in the output. Enter the number of the latest backup (as indicated by time stamp of the backup). The data associated with IAM is restored from the backup you select.

You can protect a Prism Central deployment either with the CLI-based Prism Central disaster recovery solution or the GUI-based Prism Central disaster recovery solution described in this document. For information about CLI-based Prism Central disaster recovery solution, contact Nutanix Support.

Note: Nutanix does not recommend testing 1-Click Prism Central Disaster Recovery. Prism Central Disaster Recovery must be manually triggered only when the Prism Central VM becomes unavailable due to an unplanned event like a natural calamity, network outage, or power failure.

Protecting Prism Central

You can protect a Prism Central deployment to up to three clusters running AHV or ESXi. Clusters running Hyper-V are not supported.

Before you begin

Ensure that you meet the following requirements before you can protect your Prism Central.

- Prism Central must be running on version pc.2021.7 or newer.
- Prism Central must have NTP configuration to synchronize time between Prism Central and the registered clusters.
- The clusters used to protect (back up) Prism Central must be registered to Prism Central you want to protect.
- The clusters used to protect (back up) Prism Central must be running AOS 6.0 or newer.

About this task

Procedure

1. Log on to Prism Central.
2. Click the settings button (gear icon) at the top-right corner of the window.
3. Click **Prism Central Management** in the **General** section on the left pane.
4. Click **Protect Now** in the **Disaster Recovery** widget.
   The **Protect Now** option is available only when one or more AHV or ESXi clusters are registered to Prism Central.
   The Protect Prism Central window displays the services that will be protected and the services that will not be protected.
5. Click **Continue**, select one or more clusters to back up Prism Central, and then click **Proceed**. Clicking **Proceed** starts synchronizing Prism Central configuration data to the selected clusters. Creating the first backup on the selected clusters takes at least 30 minutes. After

the first backup, the system keeps synchronizing Prism Central configuration data with the selected clusters every 30 minutes.

6. (optional) After the configuration, if you want to back up Prism Central to more clusters, click **+ Add Backup**. However, the number of clusters backing up Prism Central must not exceed three.
7. (optional) After the configuration, if you want to remove a cluster from backing up Prism Central, click **Remove**.

Recovering Prism Central (1-Click Disaster Recovery)

If you have protected your Prism Central deployment, you can use any registered AHV or ESXi cluster to recover Prism Central.

Before you begin

Ensure that you meet the following requirements before you can recover your Prism Central.

- The cluster used to recover Prism Central must be registered to the protected Prism Central.
- The ESXi cluster used to recover Prism Central must be registered to the vCenter Server.
- The cluster used to recover Prism Central must be running AOS 6.0 or newer.
- The cluster used to recover Prism Central must have ISCSI data service IP address for efficient recovery of Leap or Calm service configurations.

About this task

Procedure

1. Log on to any Prism Element registered to Prism Central you want to recover.
   The Prism Element dashboard shows Prism Central widget. Prism Central information (IP address and connection status) is available in the widget.
2. Click **Recover Now** in the Prism Central widget.
   The **Recover Now** option is available only when Prism Central is **Disconnected**.
   The Recover Prism Central window displays the services that will be recovered and the services that will not be recovered.
3. Click **Continue** and specify the following information.
   a. Select the cluster where you want to recover the Prism Central deployment.
   b. Verify the version of Prism Central that would recover on the selected cluster.
   c. Select the network where you want to recover and install Prism Central.
      The **Subnet Mask**, **Gateway**, and **DNS Address(s)** fields show the relevant information associated with the selected network.
   d. Enter details (name, IP address) for the Prism Central deployment you want to recover and click **Save**.
   e. Review the summary and click **Recover**.

   Prism Central recovers in at least 1 hour depending upon the configuration data it hosts. The recovered Prism Central shows all the guest VMs, disks, and metrics in at least 30–40 minutes more after the recovery. Do not perform any actions on the recovered Prism Central

until all the recovery tasks are completed on the cluster. You can see the status of recovery and the related processes in the **Tasks** window.

What to do next

Consider the following after Prism Central recovery.

- Use the newly recovered Prism Central only.

  If the old Prism Central becomes available, shut down or delete the old Prism Central deployment because running the old Prism Central can cause data corruption.

- Reset the credentials.

  Prism Central recovers with the default credentials. Nutanix recommends changing the default credentials.

- Reconfigure the proxy server.

  If the old Prism Central had a proxy server, reconfigure the proxy server so that the recovered Prism Central maps to the correct IP address.

- Reconfigure the fully qualified domain name (FQDN).

  If the old Prism Central had an FQDN, reconfigure the FQDN so that the recovered Prism Central maps to the correct IP address.

- Leap Recovery Plan Jobs (RPJ) in progress: Perform the steps mentioned in KB-10962.

  If the old Prism Central had a failover task running (Leap) or protection policy with guest VMs protected with Synchronous replication schedule, perform the steps mentioned in KB-10962 to ensure that all the failover tasks stuck in the running state are terminated and a script is executed for efficient recovery of Prism Central.

- Re-enable Microservices Infrastructure.

  If the old Prism Central had Microservices infrastructure enabled, see *Enabling Microservices Infrastructure* in the *Prism Central Guide* to enable Microservices infrastructure on the recovered Prism Central.

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 (see Prerequisites and Requirements on page 95).

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

Feature	Description
Hypervisors supported	AHV and ESXi
Network segmentation	You can segment the network on a single-node cluster by using Prism Element. For minimum AOS version that supports network segmentation on single-node clusters, see the *What's New* section of the *AOS 5.19.2 Release Notes*. For more information about network segmentation, see Securing Traffic Through Network Segmentation section in the Security Guide.
Unsupported features	• Cluster expansion • Deduplication • Erasure coding • Volumes • NearSync DR and Metro Availability (asynchronous DR is supported) • Capacity analysis (in Prism)

Prerequisites and Requirements

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:

- Single-node clusters are supported only on a select set of hardware models. See KB 5943 for information about supported models.
- As a best practice, Nutanix recommends configuring a maximum of five guest VMs. Ensure that the cluster platform has the minimal physical resources to cater to the compute and disk requirements. Also, ensure that the CVM resources are optimally consumed.

  Nutanix also recommends configuring backup for all the five guest VMs running on a single-node cluster to protect the guest VMs in a node failure scenario.

  > Important: Failing to configure backup for guest VMs may result in data loss as data cannot be recovered from a single-node cluster. The data loss can be observed when there is any metadata inconsistency or file system corruption.

- 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 supported for software updates, but not for firmware updates.
- Async DR is supported for 6 hour RPO only.

- Use external DNS and NTP servers. Nutanix recommends that you host the DNS and NTP servers on a different cluster.
- 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 96). 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.

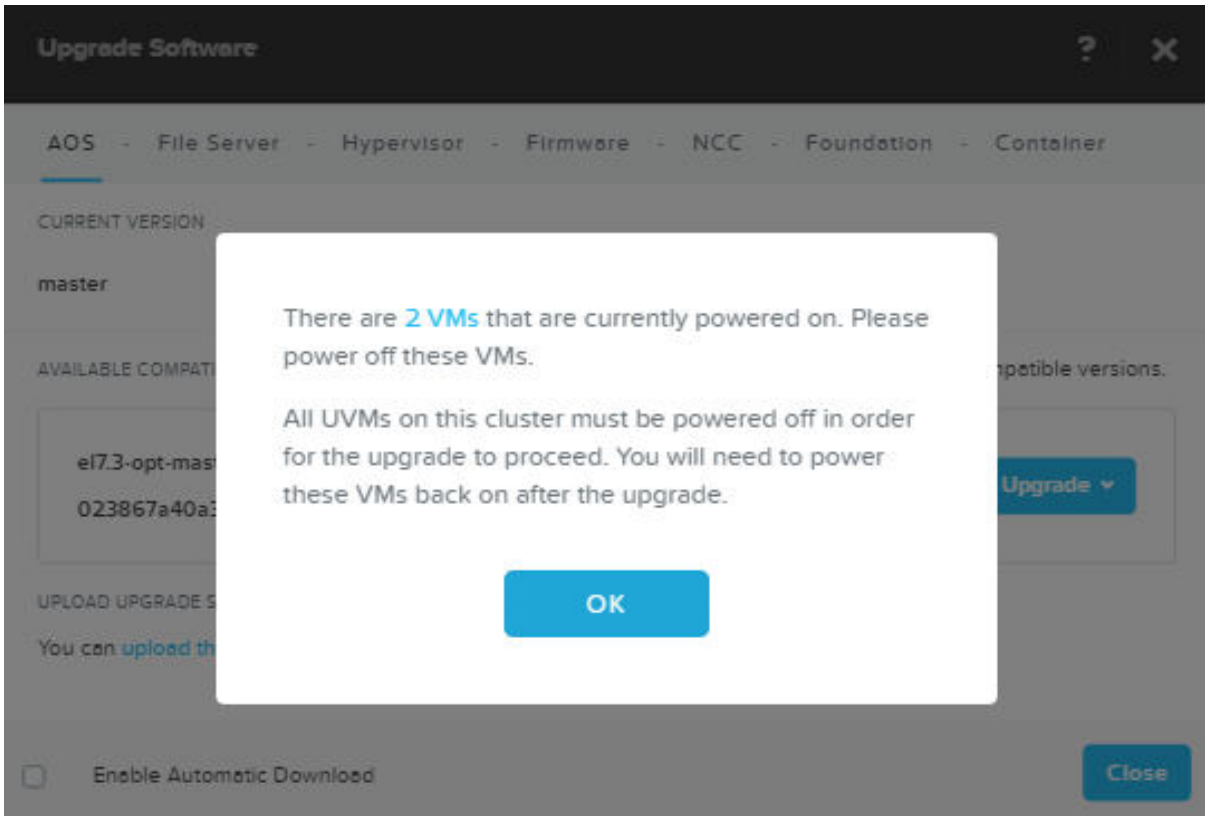

Figure 62: Upgrading Single-Node Cluster Warning

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.

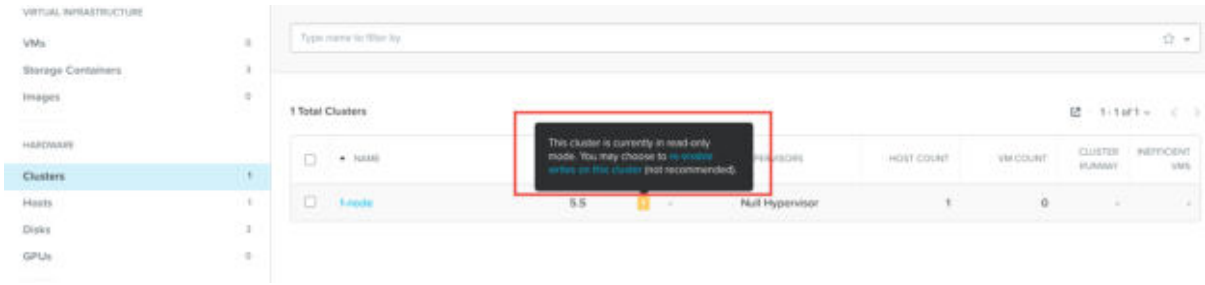

Figure 63: Re-enable Write Mode Message

   A dialog box appears to confirm that you want to override read-only mode.

   b. Click **Override** to confirm.

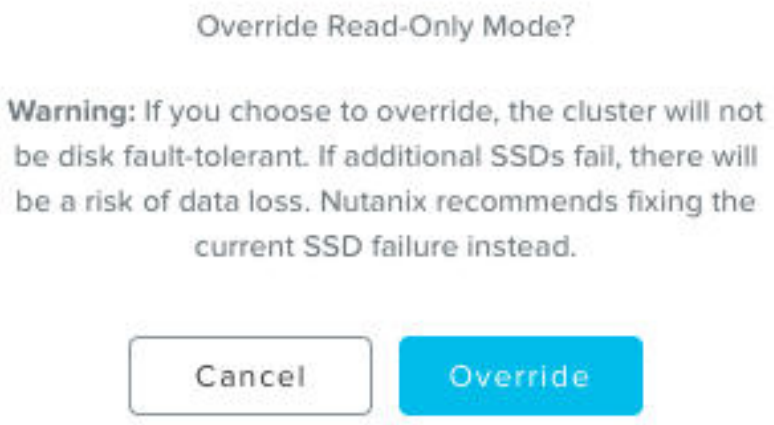

Figure 64: Confirm Override Mode

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 94), 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)* in the *Data Protection and Recovery with Prism Element* Guide). 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.

Figure 65: 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

Feature	Description
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:

- Two-node clusters are supported only on a select set of hardware models. See KB 5943 for information about supported models.
- 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 100). 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.
- Use external DNS and NTP servers. Nutanix recommends that you host the DNS and NTP servers on a different cluster.
- 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 100 instances (any combination of two-node clusters and Metro Availability protection domains).
  - 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.

- Node or disk failures are handled as follows:
  - Node failure: When a two-node cluster is operating normally, data is replicated across the nodes. If a node fails, data is replicated across disks on the healthy node to maintain resiliency.
  - HDD failure: An HDD failure is handled in the same way as in a three-node cluster, that is data is rebuilt in the background. If there is insufficient rebuild capacity when an HDD fails, the cluster can fill up and I/O can fail.
  - SSD failure: If a node loses an SSD, the remedial action is the same as for an HDD failure, that is the cluster remains normal and data is rebuilt in the background for the disk. Occasionally, an SSD failure on one node can result in the other node in the cluster transitioning to a standalone state because the SSD failure caused a critical storage service to restart, which caused user I/O to stall. If the nodes have only a single SSD left and if either node fails, the healthy node goes into read-only state.

Witness for Two-node Clusters

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. For information about configuring a witness for a two-node cluster, see *Witness Option* topic in the *Data Protection and Recovery with Prism Element* guide.

The Witness option provides a dashboard of status information about all the registered two-node clusters. (The two-node view appears only in a 5.6 or later version of the Witness.) For information about Witness dashboard for a two-node cluster, see *Witness VM Dashboard (Two-node Cluster)* topic in the *Data Protection and Recovery with Prism Element* guide.

Failure and Recovery Scenarios

There are several potential failure scenarios between the nodes and the Witness VM. The steps for recovering from a failure depend on the nature of the failure. This section describes the steps needed (or not needed) when a failure occurs. Each failure scenario generates one or

more alerts that you can review. See Alerts Dashboard on page 412 for information about viewing alert messages.

Node Failure

When a node goes down, the live node sends a leadership request to the Witness VM and goes into single-node mode. In this mode RF2 is still retained at the disk level, meaning data is copied to two disks. (Normally, RF2 is maintained at the node level normally meaning data is copied to each node.) If one of the two metadata SSDs fails while in single-node mode, the cluster (node) goes into read-only mode until a new SSD is picked for metadata service. When the node that was down is back up and stable again, the system automatically returns to the previous state (RF2 at the node level). No user intervention is necessary during this transition.

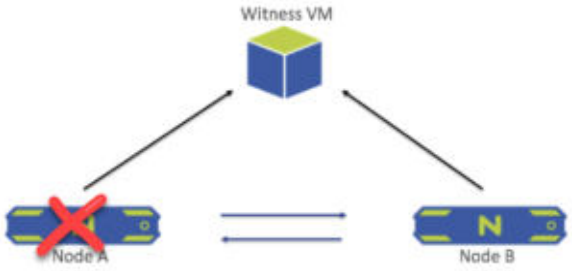

Figure 66: Failure Diagram: Node Down

Network Failure Between The Nodes

When the network connection between the nodes fails, both nodes send a leadership request to the Witness VM. Whichever node gets the leadership lock stays active and goes into single-node mode. All operations and services on the other node are shut down, and the node goes into a waiting state. When the connection is re-established, the same recovery process as in the node failure scenario begins.

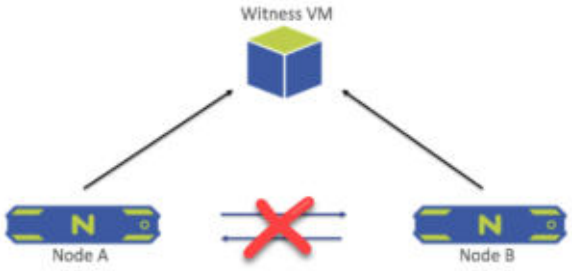

Figure 67: Failure Diagram: Node-to-Node Communication Down

Network Failure Between Node and Witness VM

When the network connection between a single node (Node A in this example) and the Witness fails, an alert is generated that Node A is not able to reach the Witness. The cluster is otherwise unaffected, and no administrator intervention is required.

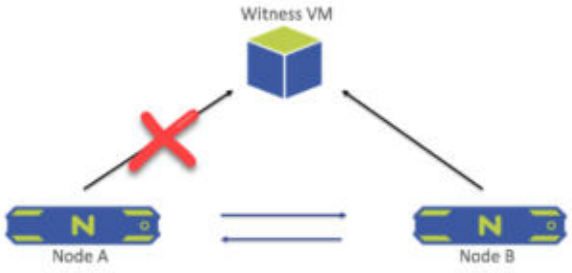

Figure 68: Failure Diagram: Node-to-Witness Communication Down

Witness VM Failure

When the Witness goes down (or the network connections to both nodes and the Witness fail), an alert is generated but the cluster is otherwise unaffected. When connection to the Witness is re-established, the Witness process resumes automatically. No administrator intervention is required.

If the Witness VM goes down permanently (unrecoverable), follow the steps for configuring a new Witness through the **Configure Witness** option of the Prism web console (see *Configuring a Witness (Two-node Cluster)* in the *Data Protection and Recovery with Prism Element* Guide).

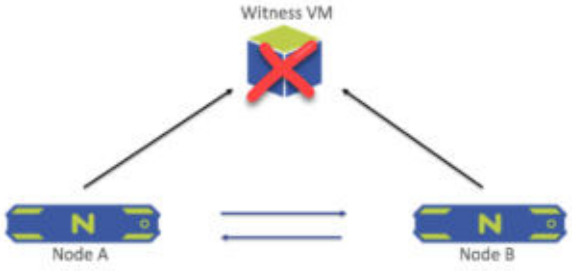

Figure 69: Failure Diagram: Witness Down

Increasing the Cluster Fault Tolerance Level

If you have a cluster set to Fault Tolerance one (FT1), you have the option to increase the cluster to FT2. Redundancy factor 3 is only supported in clusters that have FT2. This topic describes how to increase the cluster FT level.

About this task

- If you enable block awareness in addition to increasing the fault tolerance (FT) level, you need a minimum of 5 blocks.
- Increasing the cluster FT level might require at least 30 percent of your disk space.
- Each cluster must have a minimum of five nodes.
- Changes to Fault Tolerance cannot be reverted. You also cannot reduce the cluster fault tolerance level. For example, you cannot change the redundancy factor from 3 to 2 or 2 to 1. If you attempt to reduce this setting, the web console displays an error.

You can also enable Replication Factor 1 at the **Settings** > **Redundancy State** page. For more information about this setting, see Replication Factor 1 Overview on page 107, Replication

Factor 1 Recommendations and Limitations on page 107, and Enabling Replication Factor 1 on page 111.

Procedure

1. To view the number of host failures that Zookeeper can tolerate, open a web browser and log on to the Prism web console. Click **Data Resiliency Status** on the main dashboard to display the Data Resiliency Status window.

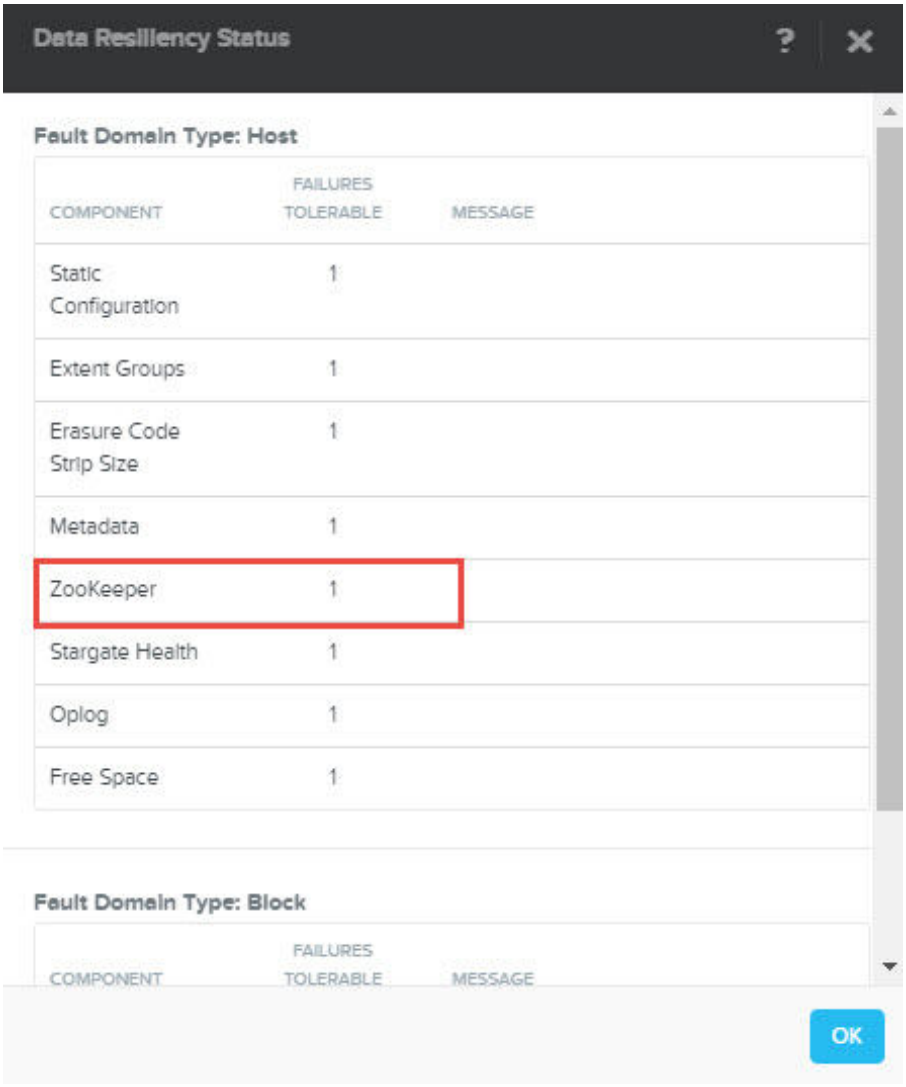

Figure 70: Data Resiliency Status window (RF 2)

2. Click the gear icon in the main menu and then select **Redundancy State** in the **Settings** page. The Redundancy Factor Readiness page appears, displaying the existing Redundancy Factor of the cluster.

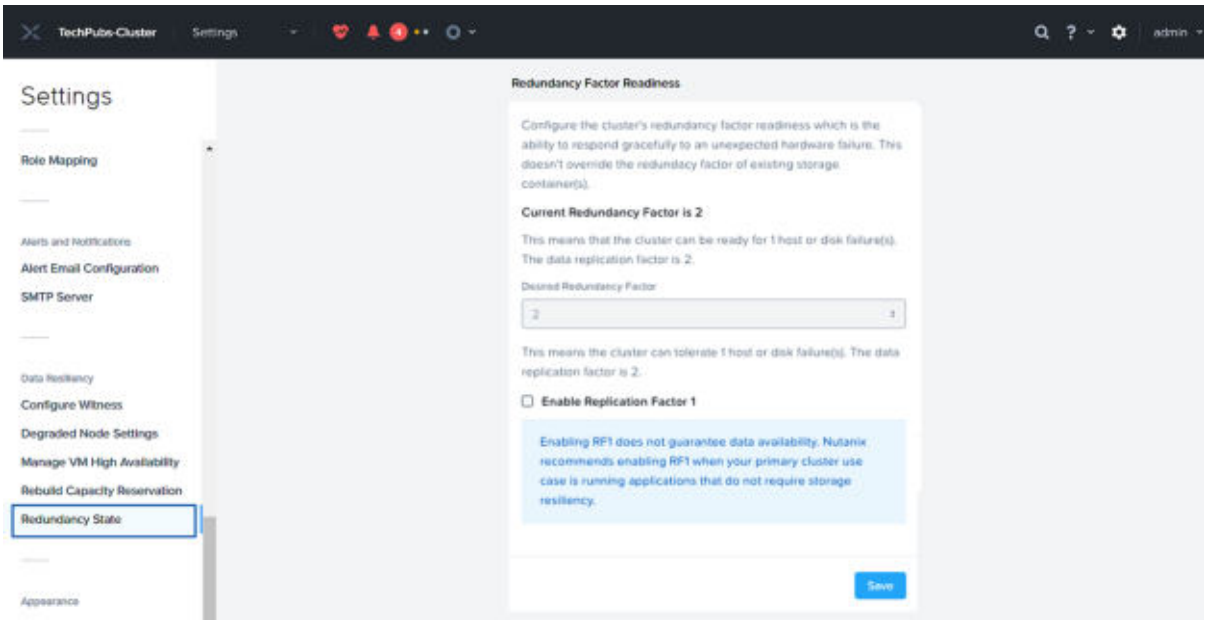

Figure 71: Redundancy Factor Readiness and Replication Factor 1

3. Select **3** in the **Desired Redundancy Factor** drop-down list to set the cluster to redundancy factor 3, then click **Save**.

   You can also enable Replication Factor 1 at the **Settings** > **Redundancy State** page. For more information about this setting, see Replication Factor 1 Overview on page 107, Replication Factor 1 Recommendations and Limitations on page 107, and Enabling Replication Factor 1 on page 111.

4. Click **Data Resiliency Status** on the main dashboard to verify that the redundancy factor is now 3.

   The process of increasing the cluster redundancy factor can take hours to complete because it must replicate metadata across the ring. When the value of **Failures Tolerable**

for Zookeeper increases to 2 as displayed in the following example, it confirms that the redundancy factor is now 3.

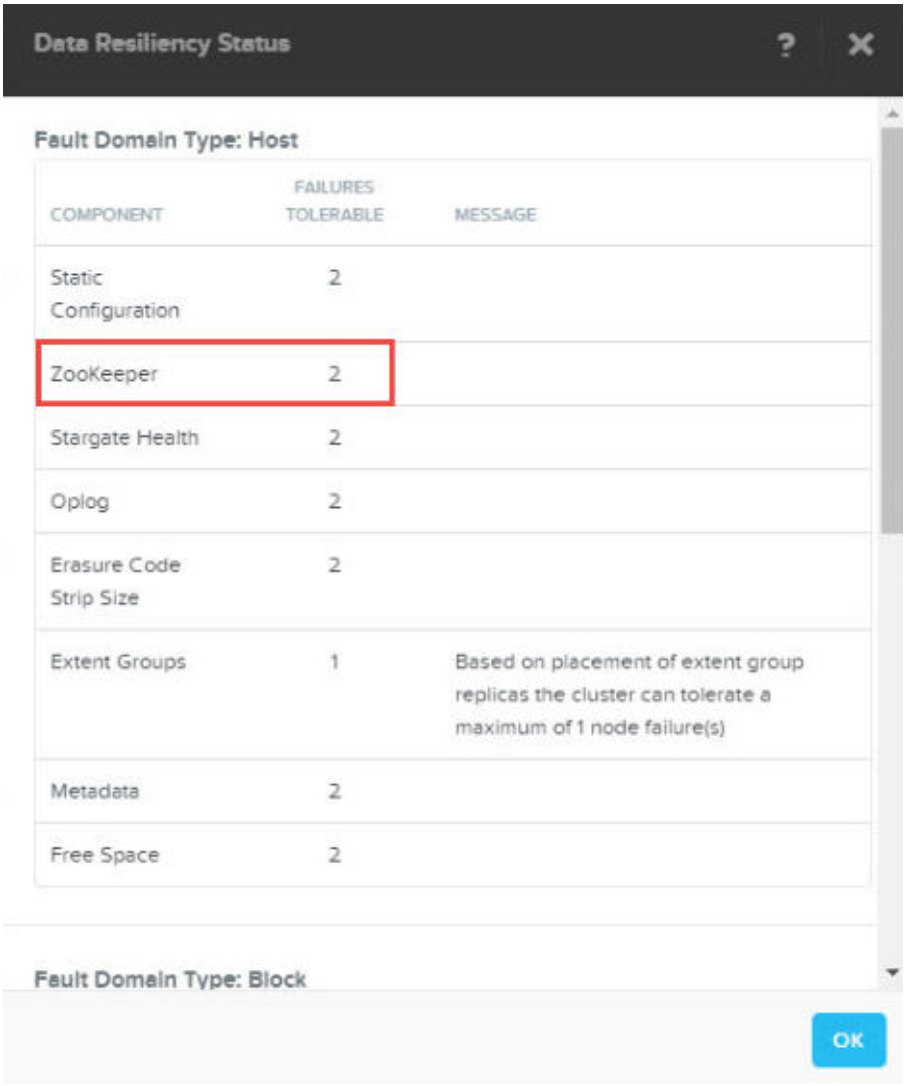

Figure 72: Data Resiliency Status window (RF 3)

5. Set the replication factor to 3 for every storage container you want to have three copies of the data.

   Increasing the redundancy factor for the cluster does not automatically increase the replication factor for any container. This gives you granular control in case you have a container where you do not want to incur the overhead of a third copy. However, this also

means no container has a third copy of the data until you explicitly increase the replication factor for that container to 3.

> Note: Be sure to change the replication factor to 3 for all desired containers, including any created by the system such as the SelfServiceContainer or NutanixManagementShare.
>
> To set the replication factor to 3 for the system storage container *NutanixManagementShare*, run
>
> ncli> ctr edit name=*NutanixManagementShare* rf=3 id=storage_container_ID force=true

a. Log in to the Controller VM and type the following command.

```
nutanix@cvm$ ncli
```

b. Locate the storage container(s) to update.

```
ncli> ctr list
```

Output similar to the following is displayed.

```
Id                    : 00052e5a-bc71-2112-0000-00000000261a::11
    Uuid              : c49eb9af-2eba-41b1-bae5-08227f7cff13
    Name              : storage_container_name
    Storage Pool Id   : 00052e5a-bc71-2112-0000-00000000261a::10
    Storage Pool Uuid : b785ff57-9d53-4f05-bad7-3ac3587f2960
    Free Space (Logical)      : 10.29 TiB (11,316,325,361,581 bytes)
    Used Space (Logical)      : 0 bytes
    Allowed Max Capacity      : 10.29 TiB (11,316,325,361,581 bytes)
    Used by other Containers  : 0 bytes
    Explicit Reservation      : 0 bytes
    Thick Provisioned         : 0 bytes
    Replication Factor        : 2
    Oplog Replication Factor  : 2
    NFS Whitelist Inherited   : true
    Container NFS Whitelist   :
    VStore Name(s)            : default-container-9754
    Random I/O Pri Order      : SSD-PCIe, SSD-SATA, DAS-SATA
    Sequential I/O Pri Order  : SSD-PCIe, SSD-SATA, DAS-SATA
    Compression               : off
    Fingerprint On Write      : off
    On-Disk Dedup             : none

    Erasure Code              : off
```

c. For each target storage container, set the replication factor to 3.

```
ncli> ctr edit name=storage_container_name rf=3
```

Replace *storage_container_name* with the name of the storage container.

Output similar to the following is displayed. This shows that the replication factor is now at the correct state (3) for the storage container.

```
Id                    : 00052e5a-bc71-2112-0000-00000000261a::1381
    Uuid              : a567b66d-b356-4883-90c5-c37b2f3e0fad
    Name              : storage container_name
    Storage Pool Id   : 00052e5a-bc71-2112-0000-00000000261a::10
    Storage Pool Uuid : b785ff57-9d53-4f05-bad7-3ac3587f2960
    Free Space (Logical)      : 6.86 TiB (7,544,216,907,720 bytes)
    Used Space (Logical)      : 0 bytes
    Allowed Max Capacity      : 6.86 TiB (7,544,216,907,720 bytes)
```

```
Used by other Containers  : 0 bytes
Explicit Reservation      : 0 bytes
Thick Provisioned         : 0 bytes
Replication Factor        : 3
Oplog Replication Factor  : 3
NFS Whitelist Inherited   : true
Container NFS Whitelist   :
VStore Name(s)            : aaa
Random I/O Pri Order      : SSD-PCIe, SSD-SATA, DAS-SATA
Sequential I/O Pri Order  : SSD-PCIe, SSD-SATA, DAS-SATA
Compression               : off
Fingerprint On Write      : off
On-Disk Dedup             : none

Erasure Code              : off
```

As with increasing the redundancy factor, increasing the replication factor can take some time to complete. You can verify the status by again going to the Data Resiliency Status window in Prism. When the value of **Failures Tolerable** for Extent Groups (which reflects the container replication level) increases to 2, it confirms that the replication factor is now 3.

Replication Factor 1 Overview

Nutanix recommends that you only enable RF1 when your cluster is running applications that provide their own data protection or high availability.

When you enable replication factor 1 (RF1), the cluster maintains a single copy of data for each RF1-enabled storage container. The storage container can contain multiple vDisks you can directly attach to a VM or implement as part of a volume group. Each RF1-enabled storage container is associated with a specific node that you choose. All the data in the RF1 storage container will be associated with that node only.

For clusters running the ESXi hypervisor, you must associated the storage container with a single host only. Mounting the RF1 storage container on all or multiple hosts is not supported.

As only one copy of data exists in an RF1-enabled storage container, RF1 does not guarantee data availability if a node or disk failure occurs.

Clusters with RF1 storage containers can improve I/O performance for write-intensive applications that do not need storage level data protection. Applications that provide their own data protection and applications that perform at high I/O rates on ephemeral data can benefit from this type of storage. Workloads that can benefit from storing data in an RF1 storage container include artificial intelligence (AI), machine learning (ML), and big data.

To get started, see Replication Factor 1 Recommendations and Limitations on page 107 and Enabling Replication Factor 1 on page 111.

Replication Factor 1 Recommendations and Limitations

Enabling replication factor 1 (RF1) does not guarantee data availability. Nutanix recommends that you only enable RF1 when your cluster is running applications that provide their own data protection or high availability.

Recommendations

Nutanix recommends the following.

- For each cluster where you want to use RF1, create one RF1 storage container per node. The RF1-enabled storage container can contain multiple RF1-enabled vDisks.

- For clusters running the ESXi hypervisor, make sure you mount the RF1 storage container on a single host only. Mounting the RF1 storage container on all or multiple hosts is not supported.
- Do not enable RF1 storage containers on clusters where SSD tier capacity is less than 6 percent of total cluster capacity.

Limitations

Table 13: Unsupported Software, Features, and Operations

Nutanix Software	Description
Nutanix Era	Using an RF1 storage container as part of configuration and deployment is not supported.
Nutanix Files	Using an RF1 storage container as part of configuration and deployment is not supported.
Nutanix Karbon	Using an RF1 storage container as part of configuration and deployment is not supported.
Nutanix Objects	Using an RF1 storage container as part of configuration and deployment is not supported.
Prism Central	Do not deploy Prism Central on RF1-enabled storage containers. This configuration is not supported.
Cluster Hypervisor	**Support or Limitation**
Microsoft Hyper-V	Enabling RF1 on clusters running Microsoft Hyper-V is not supported.
Mixed hypervisor cluster	For VMware ESXi and AHV mixed hypervisor clusters, RF1-enabled storage containers and VMs with attached RF1 vDisks are not supported on the AHV storage-only nodes. You can enable RF1 on the nodes running the VMware ESXi hypervisor.
VMware ESXi	Using vSphere HA with Distributed Resource Scheduler (DRS) is not supported for any VM with at least one RF1-enabled vDisk.
Feature or Operation	**Support or Limitation**
Snapshots and data protection/recovery	For vDisks in RF1-enabled storage containers, disaster recovery operations like snapshots and replication are not supported.

Feature or Operation	Support or Limitation
In-place hypervisor conversion	Not supported when converting a cluster from AHV to ESXi.
Deduplication	Not supported for RF1-enabled storage containers.
Erasure coding	Not supported for RF1-enabled storage containers.
Metro availability	Not supported for RF1-enabled storage containers and VMs with attached RF1 vDisks.
Recycle bin	Not supported. When you delete an RF1 vDisk, the vDisk is marked for deletion as soon as possible and bypasses the recycle bin.
Storage container settings or modification with RF1 enabled	These settings or modifications are not supported. • Increasing the replication factor • Enabling capacity reservation (Reserve Capacity setting) or Reserve Rebuild Capacity • Storage containers with Reserve Capacity or Reserve Rebuild Capacity already enabled
Maintenance mode operations (planned)	The following operations automatically shut down and restart RF1 VMs. • Controller VM shutdown • AOS 1-click upgrade • Memory update • Host boot disk replacement • Nutanix Flow microsegmentation work flows • Request Reboot (rolling reboot operation)
ESXi hypervisor 1-click upgrade	Automatic VM shutdown is not supported. You must manually shut down RF1 VMs before upgrading and restart them after the upgrade is completed.
LCM upgrade operations	The following LCM upgrade operations automatically shut down and restart RF1 VMs. • AOS upgrade • Firmware upgrades for ESXi and AHV clusters • AHV hypervisor upgrade

Feature or Operation	Support or Limitation
Upgrade operations, VM with volume group access	• For any VM with a volume group directly attached to the VM (AHV clusters only), the cluster will automatically shut down and power on the VM as part of the upgrade operation. • For any VM with volume group access provided by Nutanix Volumes (that is, connection to the volume group from the VM guest OS through an iSCSI target IP address), you must manually shut down the VM. You can power on the VM when the upgrade operation is completed.
Unplanned event	These unplanned events might affect I/O operations for VMs with vDisks in RF1-enabled storage containers, as you cannot shut down these VMs in advance of these events. • Disk or node failure • Controller VM Stargate service failure • Controller VM failure or restart • Degraded node
Node removal	When you remove a node that contains an RF1-enabled storage container, the storage container and its vDisks are marked for removal. The storage container is not automatically marked for deletion. As part of node removal (which first fails in this case), you are prompted to delete the storage container, including all data on the storage container. This is effectively the same as deleting the container.
Disk removal or replacement	When you remove or replace a node's disk that is part of an RF1-enabled storage container, the storage container and RF1 vDisks are marked for deletion. As part of this procedure, data in an RF1-enabled storage container is automatically moved to other disks associated with the RF1 node. If there is insufficient disk space on the node, disk removal fails with a message similar to Warning: no disk space for RF1 data.

Enabling Replication Factor 1

To enable RF1, enable the Replication Factor 1 setting as a first step. After you enable this setting, create a dedicated RF1-enabled storage container, then associate that single RF-1 enabled storage container with one node.

About this task

- See Replication Factor 1 Recommendations and Limitations on page 107.
- Once the replication factor for a container is set to 1, you cannot increase the replication factor on that container.

Procedure

1. Open a web browser and log on to the Prism Element web console.
2. Click the gear icon in the main menu and then select **Redundancy State** in the **Settings** page. The Redundancy Factor Readiness page appears, displaying the existing Redundancy Factor of the cluster and the **Enable Replication Factor 1** check box.

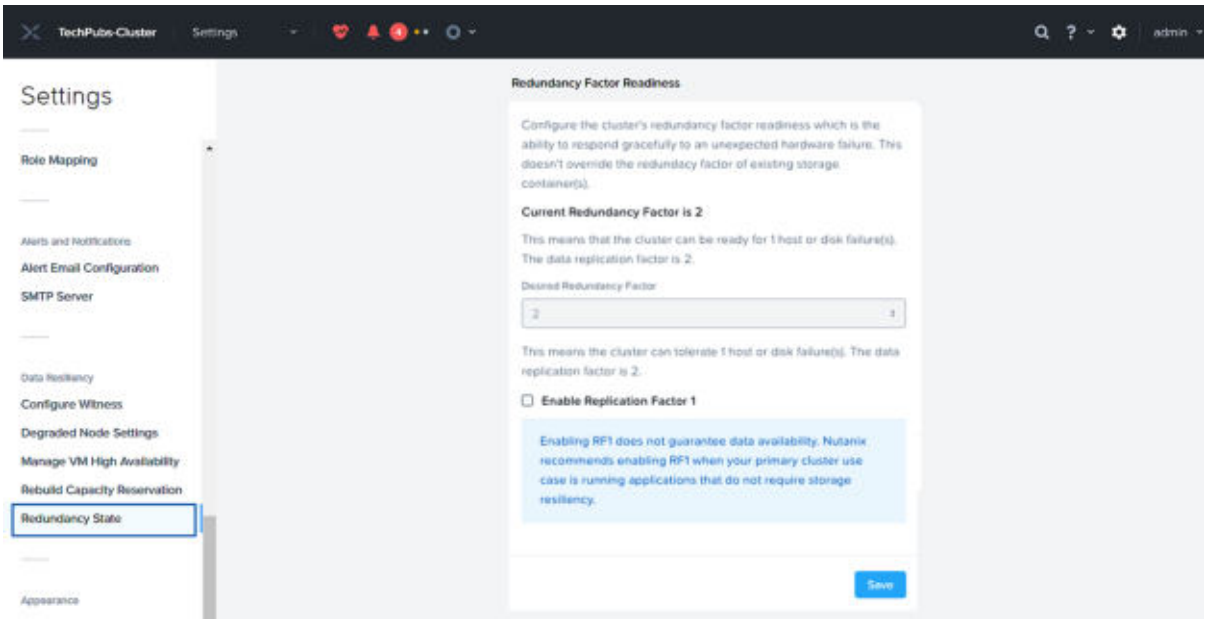

Figure 73: Redundancy Factor Readiness and Replication Factor 1

3. Select **Replication Factor 1**, click **Yes** to confirm the Are you sure you want to enable RF1? message, and then click **Save**.

What to do next

Create an RF1 storage container and assign it to a node. See Creating a Storage Container with Replication Factor 1 on page 111.

Creating a Storage Container with Replication Factor 1

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers on page 753) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail

to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

For replication factor 1 (RF1) storage container recommendations and limitations, see Replication Factor 1 Recommendations and Limitations on page 107. To create a non-RF1 storage container, see Creating a Storage Container on page 157.

About this task

AOS automatically creates the correct type of access to the RF1-enabled storage container for each hypervisor.

- vSphere: The storage container is accessible as an NFS datastore. This requires access to the vSphere APIs. Ensure that you have appropriate vSphere license to access the APIs.
- AHV: The storage container is accessible transparently.

> Note: The NutanixManagementShare storage container is a built-in storage container for Nutanix clusters for use with the Nutanix Files and Self-Service Portal (SSP) features. This storage container is used by Nutanix Files and SSP for file storage, feature upgrades, and other feature operations. To ensure proper operation of these features, do not delete this storage container. Nutanix also recommends that you do not delete this storage container even if you are not using these features. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.

Procedure

1. In the Storage dashboard (see Storage Dashboard on page 133), click the **Storage Container** button.
   The Create Storage Container dialog box appears.
2. **Name**: Enter a name for the storage container.
   The storage container has naming restrictions depending on the cluster hypervisor.

   AHV:

   - Container Name Length: Maximum length is 75 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   - Case Sensitivity: Container names are case sensitive.

   ESXi:

   - Container Name Length: Maximum length is 42 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   - Case Sensitivity: Container names are case sensitive.

3. **Storage Pool**: Select a storage pool from the drop-down list. **Max Capacity** displays the amount of free space available in the selected storage pool.

   a. (vSphere only) **NFS Datastore**: Select the **Mount on the following ESXi hosts** button to mount the storage container on a single host. From the list of host IP addresses, check the box for one host only.

      For clusters running the ESXi hypervisor, make sure you mount the RF1 storage container on a single host only. Mounting the RF1 storage container on all or multiple hosts is not supported.

4. Click the **Advanced Settings** button and configure the storage container.

   a. **Replication Factor**: Select **1**

      Nutanix supports a replication factor (RF) of 1 if you first select the **Enable Replication Factor 1** checkbox at the **Settings** > **Redundancy State** page. If you do not select **Enable Replication Factor 1**, Nutanix supports a replication factor of 2 or 3.

   b. From the **Node** drop-down menu, select a node to associated with the storage container.

      You can create one RF1 storage container per node, so you can only select one node here.

   c. **Reserved Capacity**: Leave this setting at 0.

   d. **Advertised Capacity**: Sets a maximum storage space for this storage container. Enter the amount (in GiB) to reserve in this field.

      This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

   e. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is configurable. For post-process compression,

where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

All data in the storage container is compressed when you select **Compression**. See Compression on page 126 for information about using compression.

f. **Deduplication** and **Erasure Coding** are greyed out and unavailable when you select replication factor 1.

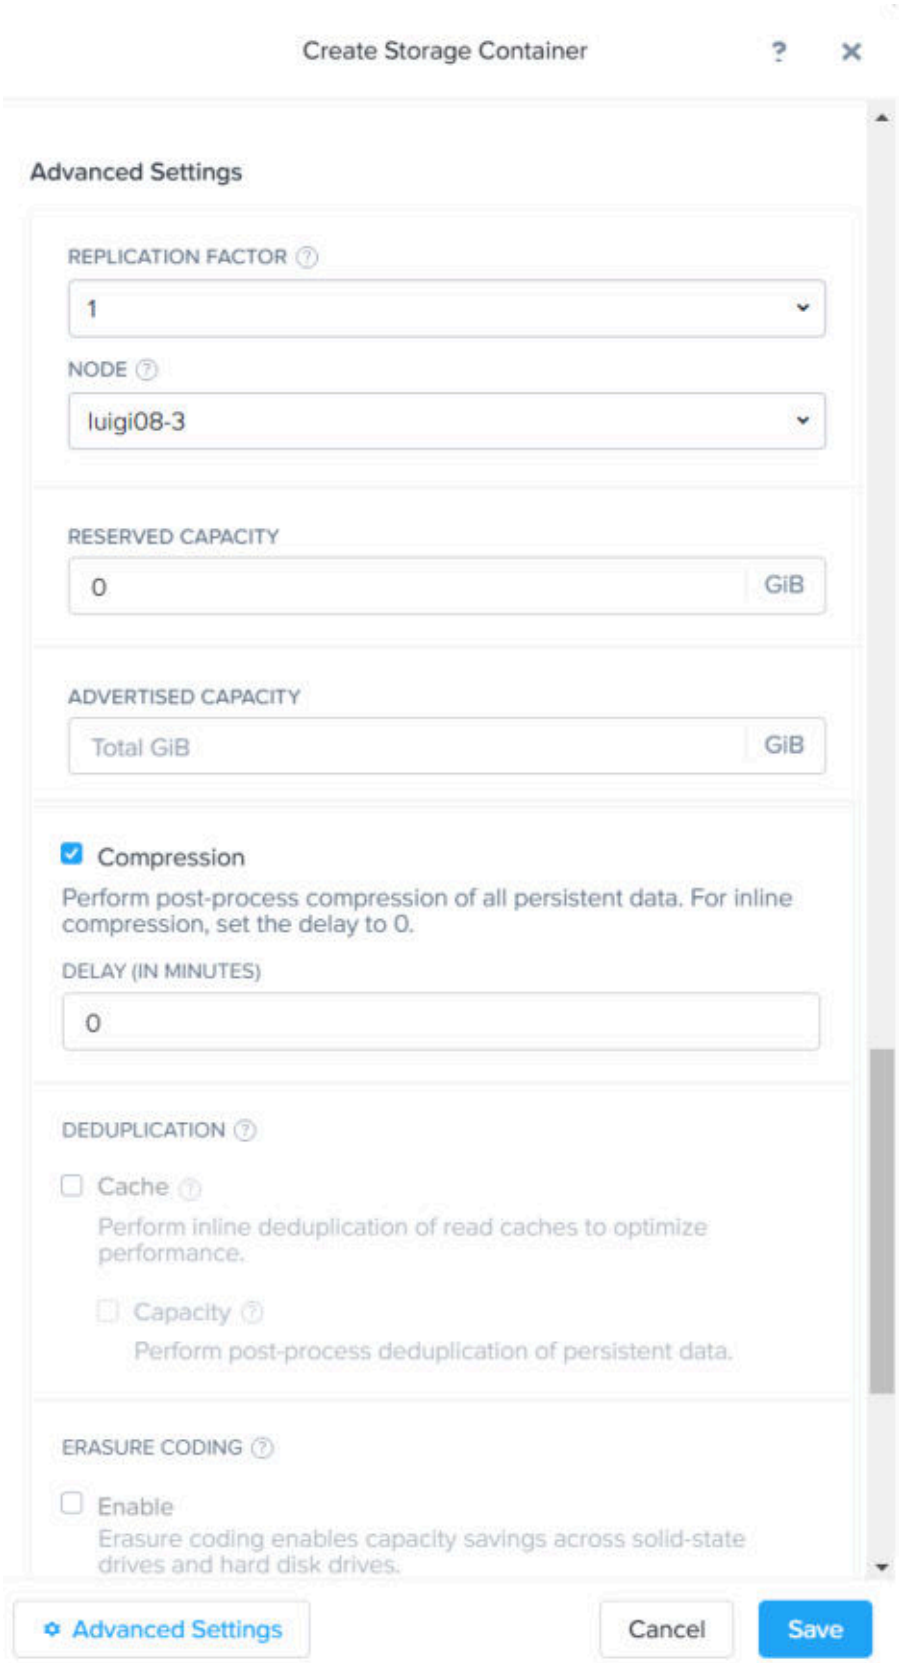

Figure 74: Create Storage Container Advanced Settings

5. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form *ip_address/ netmask*).

   A whitelist, also known as an allowlist, is a set of addresses that are allowed access to this storage container. Allowlists are used to allow appropriate traffic when unauthorized access from other sources is denied.

   > Note: Setting a storage container level allowlist overrides any global allowlist for this storage container.

6. When all the field entries are correct, click **Save**.

Disabling Replication Factor 1

About this task

Disabling Replication Factor 1 prevents you from creating a new RF1-enabled storage container. Any existing RF1-enabled containers with associated vDisks remain in place.

Procedure

1. Open a web browser and log on to the Prism Element web console.
2. Click the gear icon in the main menu and then select **Redundancy State** in the **Settings** page. The Redundancy Factor Readiness dialog box appears. This displays the existing Redundancy Factor of the cluster and the **Enable Replication Factor 1** checkbox.
3. Clear (unselect) **Replication Factor 1** and then click **Save**.
   After disabling Replication Factor 1, if you create or modify a storage container, your replication factor choices are 2 or 3.

CVM Memory Configuration

Foundation provisions the default memory and vCPUs to each Controller VM (CVM) according to your platform model. See the CVM vCPU and vRAM Allocation topic in the *Field Installation Guide* for information about how Foundation allocates memory and vCPUs to your platform model. After Foundation, you can increase the memory reserved for each Controller VM in your cluster by using the 1-click Controller VM Memory Upgrade feature available in the Prism Element web console. You might need to increase the Controller VM memory if your cluster configuration is as one or more of the configurations described in this document.

Default CVM Memory Requirements

The following table provides information on default CVM memory requirements for various configurations.

Table 14: Default CVM Memory Requirements for specific platform types

Configuration	Memory Requirement
vdi: general / VDI / server virtualization	20 GB
storage_heavy or minimal_compute_node	28 GB
Unless the platform has hardware attribute all_flash_low_perf (rare), the above 2 rules are ignored if the platform is All Flash.	32 GB

Configuration	Memory Requirement
high_perf	32 GB
dense, or if the node has 2+ NUMA nodes and one of the following: 1. 32+ TB of HDD 2. 48+ TB of SSD and NVMe combined	40 GB
If there is a hardware attribute called default_cvm_memory_in_gb (rare), all of the above are ignored.	The value defined by default_cvm_memory_in_gb.
If AOS is older than 5.1 or total memory is <62GB.	Subtract 4 GB from the previous value, only if the value is < 32 GB.
If the determined value leaves < 6 GB of remaining memory, all of the above are ignored.	As much memory that leaves 6 GB of remaining memory.

Minimum CVM Memory Requirements

The following tables provide recommendations on optimal memory sizes for various cluster configurations.

> Note: The drive capacity mentioned in the following tables is for each node in the cluster and not the aggregated drive capacity of the cluster.

Table 15: Minimum CVM Memory Requirements for Each Node in Remote Direct Memory Access (RDMA) Clusters

Configuration	Memory Requirement
Two-socket nodes	Each CVM requires 32 GB of minimum memory for each node.
Four-socket nodes	Each CVM requires 48 GB of minimum memory for each node.

Table 16: Minimum CVM Memory Required for Nodes With All Flash

Configuration	Memory Requirement
Nodes With All Flash	Foundation allocates 32 GB for nodes with All Flash disks.

Table 17: Minimum CVM Memory Required for Nodes With Both HDDs and SSDs

Configuration	Memory Requirement
Nodes With All Flash	See Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro) on page 120 for information about CVM memory requirements on nodes with both HDDs and SSDs as well as All-Flash nodes to support various snapshot frequencies.

Table 18: Additional CVM Memory Required for Cluster Features

Configuration	Memory Requirement
Capacity tier deduplication (includes performance tier deduplication)	Add 12 GB
Performance tier deduplication	Add 8 GB
Redundancy factor 3	Add 8 GB

> Note:
>
> Deduplication is supported without additional memory if CVM is already configured with 32GB memory. However, RF3 would require additional 8GB memory.

Table 19: Minimum CVM Memory Required for SPDK on Clusters

Configuration	Memory Requirement
SPDK for nodes with NVMe devices	To ensure that SPDK is enabled, set a minimum memory of 48 GB on each Controller VM of a cluster with Blockstore. For more information about Blockstore and SPDK, see the *Blockstore Support with SPDK* section in the *Acropolis Advanced Administration Guide*.

Increasing the Controller VM Memory Size

About this task

> Note: Nutanix does not support decreasing the Controller VM memory below the recommended minimum requirements.

Perform the following procedure to increase the Controller VM memory in the Prism Element web console.

Procedure

1. Log on to the Prism Element web console.

2. Run NCC as described in Running Checks by Using Web Console on page 295.

3. Click the gear icon in the main menu and then select **Configure CVM** in the **Settings** page. The Configure CVM dialog box is displayed.

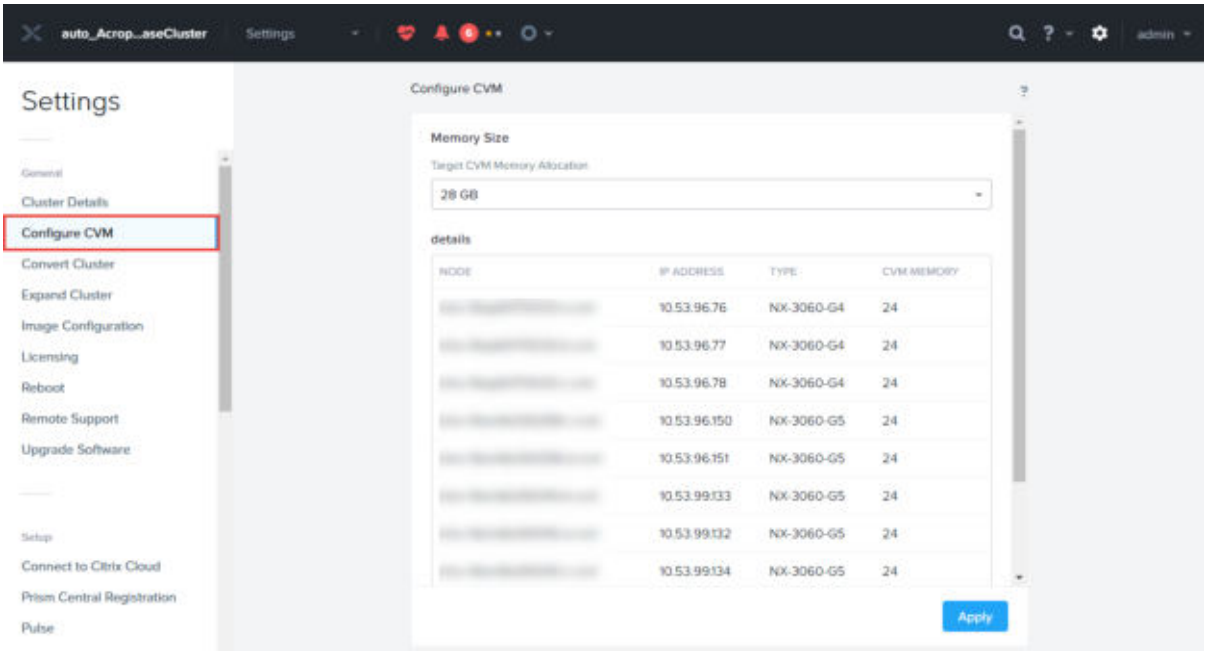

Figure 75: Configure CVM Window

4. If your cluster is running the ESXi hypervisor, you are managing your cluster by using VMware vCenter Server, and you have not authenticated the vCenter Server instance, perform the following steps:

   a. In the **Configure CVM** dialog box, click **Add vCenter**.

   b. Enter the vCenter Server IP address and administrator credentials.

   c. Click **Add**.

5. Select the **Target CVM Memory Allocation** memory size and click **Apply**.

   The values available from the drop-down menu can range from 16 GB to the maximum available memory in GB.

   AOS applies memory to each Controller VM that is below the amount you choose. Resizing memory for a Controller VM requires a reboot as part of the process.

   > Note: The maximum CVM memory that can be allocated through the web console is 64 GB. To upgrade the CVM memory beyond 64 GB, contact Nutanix support.

   The Controller VM can only be allocated a maximum memory of the size available in the drop-down list. For example, if a Controller VM is allocated 64 GB and you choose to upgrade this memory, the value remains at the maximum permissible value of 64 GB. If a Controller VM was already allocated more memory than your choice, it remains at the allocated amount. For example, if a Controller VM is at 20 GB memory and you select 28 GB, the Controller VM memory is upgraded to 28 GB. However, if a Controller VM is at 48 GB memory and you select 28 GB, the Controller VM memory remains unchanged at 48 GB.

Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync and Metro)

For DR solutions with Asynchronous, NearSync, and Synchronous (using metro availability) replication schedules to succeed, the nodes must have certain resources. This section provides information about the node configurations necessary to support the asynchronous, NearSync, and synchronous replication schedules.

- The conditions and configurations provided in this section apply to Local and Remote snapshots.
- Any node configuration with two or more SSDs, each of 1.2 TB or greater capacity, supports NearSync.
- Any node configuration that supports snapshots at a frequency of 6 hours also supports metro availability.
- Both the primary cluster and the replication target cluster must fulfill the same minimum resource requirements for node configurations for the required RPO.
- Ensure that any new node or disk addition meets the minimum requirements provided in this section.
- Features such as Deduplication and RF3 may require additional memory depending on the DR schedules and other workload run on the cluster.

> Note:
>
> Deduplication is supported without additional memory if CVM is already configured with 32GB memory. However, RF3 would require additional 8GB memory

> Note: The default minimum snapshot or recovery point frequency with the default Foundation configuration is 6 hours. To increase the snapshot frequency, you must also modify the Foundation configuration (SSD and CVM) accordingly. For example, an all-flash setup with a capacity between 48 TB to 92 TB has the default snapshot interval is 6 hours. If you want to increase the snapshot interval to 1 hour, you must modify the default Foundation configuration to:
>
> - 14 vCPUs for CVM
> - 40 GB for CVM

The table lists the supported frequency for the snapshots across various hardware configurations.

Table 20: Snapshot Frequency

Type of disk	Capacity per node	Minimum snapshot frequency	Foundation Configuration - SSD and CVM requirements
Hybrid	Total HDD capacity per node: 32 TB or lower. And Total Capacity Per Node: (HDD + SSD) of 40 TB or lower.	• NearSync • Async (Hourly)	No change required —existing Foundation defaults. • 2 x SSDs • Each SSD must be minimum 1.2 TB or more for NearSync.
	Total HDD capacity per node: Between 32-64 TB. And Total Capacity Per Node: (HDD + SSD) Between 40-92 TB. Up to 32 TB SSD (4 x 7.68 TB SSDs)	• NearSync • Async (Hourly)	Modify Foundation configurations to minimum: • 4 x SSDs • Each SSD must be minimum 1.2 TB or more for NearSync. • 14 vCPU for CVM • 40 GB for CVM
	Total HDD capacity per node: Between 32-64 TB. And Total Capacity Per Node: (HDD + SSD) Between 40-92 TB Up to 32 TB SSD (4 x 7.68 TB SSDs)	Async (Every 6 hours)	No change required —Default Foundation configuration.
	Total HDD capacity per node: Between 64-80 TB. And Total Capacity Per Node: (HDD + SSD) 96 TB or lower.	Async (Every 6 hours)	No change required —Default Foundation configuration.

Type of disk	Capacity per node	Minimum snapshot frequency	Foundation Configuration - SSD and CVM requirements
	Total HDD capacity per node: Greater than 80 TB. And Total Capacity Per Node: (HDD + SSD) Between 96-136 TB. Maximum 136 TB.	Async (Every 6 hours)	Modify Foundation configurations to minimum: • 12 vCPU for CVM • 36 GB for CVM
All Flash	Total Capacity Per Node: 48 TB or lower	• NearSync • Async (Hourly)	No change required —Default Foundation configuration. Each SSD must be minimum 1.2 TB or more for NearSync.
	Total Capacity Per Node: Between 48-92 TB	• NearSync • Async (Hourly)	Modify Foundation configurations to minimum: • 14 vCPU for CVM • Each SSD must be minimum 1.2 TB or more for NearSync. • 40 GB for CVM
	Total Capacity Per Node: Between 48-92 TB	Async (Every 6 hours)	No change required —Default Foundation configuration.
	Total Capacity Per Node: Greater than 92 TB	Async (Every 6 hours)	Modify Foundation configurations to minimum: • 12 vCPU for CVM • 36 GB for CVM

Restarting the Nodes in the Cluster

About this task

The Request Reboot operation in the Prism web console gracefully restarts the selected nodes one after the other.

Perform the following procedure to restart the nodes in the cluster.

Procedure

1. Click the gear icon in the main menu and then select **Reboot** in the **Settings** page.

2. In the Request Reboot window, select the nodes you want to restart, and click **Reboot**.

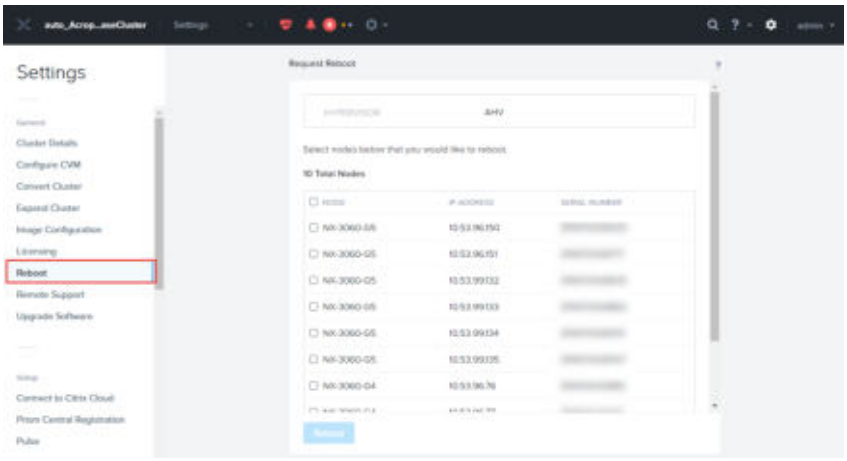

Figure 76: Request Reboot

A progress bar is displayed that indicates the progress of the restart of each node.

3

STORAGE MANAGEMENT

Storage in a Nutanix cluster is organized hierarchically into several components that allow you to manage capacity and performance. Following is a brief outline of the storage features of Nutanix and how to navigate these storage features.

- Nutanix clusters provide storage pool, storage container, volume group, and virtual disk components to organize storage. (see Storage Components on page 124).
- Prism web console helps you monitor storage usage across the cluster (see Storage Dashboard on page 133).
- Prism web console helps you create storage containers (see Creating a Storage Container on page 157) and volume groups (see Creating a Volume Group on page 175).
- Prism web console helps you configure a threshold warning for storage capacity available in the cluster after accounting for the storage space needed to rebuild and restore in case of any component failures (see Configuring a Warning Threshold for Resilient Capacity on page 167).
- Prism web console helps you reserve storage capacity for rebuilding failed nodes, blocks or racks (see Rebuild Capacity Reservation on page 167).

Storage Components

Storage in a Nutanix cluster is organized into the following components.

Storage Tiers

Each type of storage hardware (SSD-PCIe (NVMe), SSD (SATA SSD), and HDD) is placed in a storage tier. You can determine the tier breakdown for disks in a storage pool through the web console (see "Storage Table View" section in the Prism Web Console guide.

Storage Pools

Storage pools are groups of physical disks from one or more tiers. Storage pools provide physical separation because a storage device can only be assigned to a single storage pool at a time. Nutanix recommends creating a single storage pool for each cluster. This configuration allows the cluster to dynamically optimize capacity and performance. Isolating disks into separate storage pools provides physical separation, but can create an imbalance of these resources if the disks are not actively used. When you expand your cluster by adding new nodes, the new disks can also be added to the existing storage pool. This scale-out architecture allows you to build a cluster that grows with your needs.

When you create a cluster, a default predefined storage pool is available. This pool includes the total capacity of all the disks on all the hosts in the cluster.

Storage Containers

A storage container is a subset of available storage within a storage pool (see "Creating a Storage Container" section). Storage containers are created within a storage pool to hold

virtual disks (vDisks) used by virtual machines. By default, storage is thinly provisioned, which means that the physical storage is allocated to the storage container as needed when data is written, rather than allocating the predefined capacity when the storage container is created. Storage efficiency features such as compression, deduplication, and erasure coding are enabled at the container level.

Volume Groups

A volume group is a collection of logically related virtual disks (or volumes). A volume group is attached to VM either directly or using iSCSI. You can add vDisks to a volume group, attach them to one or more consumers, include them in disaster recovery policies, and perform other management tasks. You can also detach a volume group from one VM and attach it to another, possibly at a remote location to which the volume group is replicated.

You manage a volume group as a single unit. When a volume group is attached to a VM, the VM can access all of the vDisks in the volume group. You can add, remove, and resize the vDisks in a volume group at any time.

Each volume group is identified by a UUID, a name, and an iSCSI target name. Each disk in the volume group also has a UUID and a SCSI index that specifies ordering within the volume group. A volume group can be configured for either exclusive or shared access.

You can backup, protect, restore, and migrate volume groups. You can include volume groups in protection domains configured for asynchronous data replication (Async DR), either exclusively or with VMs. However, volume groups cannot be included in a protection domain configured for metro availability, in a protected vStore, or in a consistency group for which application consistent snapshots are enabled.

vDisks

A vDisk is created within a storage container or volume group to provide storage to the virtual machines. A vDisk shows up as a SCSI device when it is mapped to a VM.

Containers for VMware and Hyper-V (Datastores/SMB Shares)

In vSphere, a datastore is a logical container for files necessary for VM operations. Nutanix provides the choice by supporting both iSCSI and NFS protocols when mounting a storage volume as a datastore within vSphere. NFS has many performance and scalability advantages over iSCSI, and it is the recommended datastore type.

In Hyper-V environments, storage containers are mounted as an SMB share.

> Note: Using a Nutanix storage container as a general-purpose NFS or SMB share is not recommended. For NFS and SMB file service, use Nutanix Files.

NFS Datastores. The Distributed Storage Fabric (DSF) reduces unnecessary network chatter by localizing the data path of guest VM traffic to its host. This boosts performance by eliminating unnecessary hops between remote storage devices that is common with the pairing of iSCSI and VMFS. To enable vMotion and related vSphere features (when using ESX as the hypervisor), each host in the cluster must mount an NFS volume using the same datastore name. The Nutanix web console and nCLI both have a function to create an NFS datastore on multiple hosts in a Nutanix cluster.

To correctly map the local ESX datastore to the Nutanix container:

- Map the NFS share with 192.168.5.2 (internal IP address) and not the Controller VM IP address or cluster virtual IP address.
- The name of the datastore should be same as the name of the container.

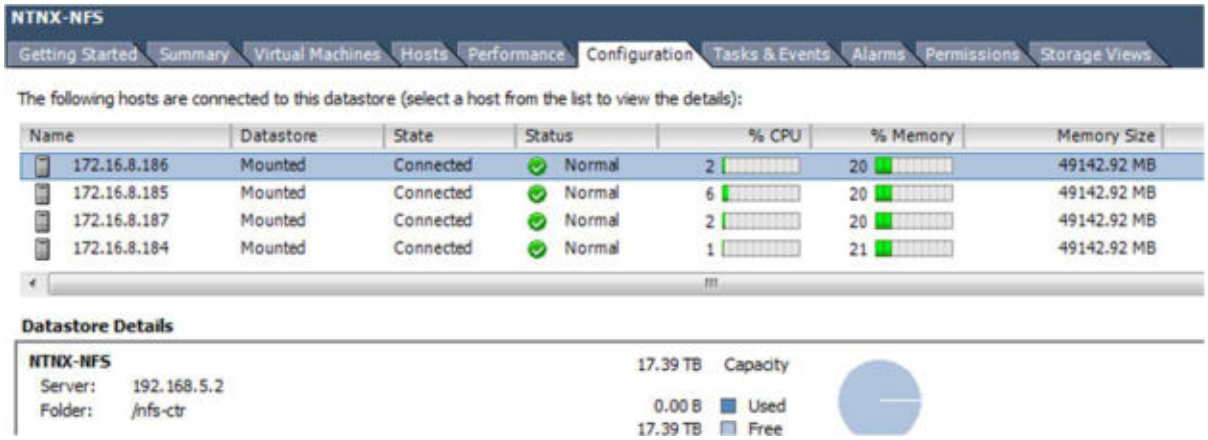

Figure 77: vSphere Configuration of NFS Datastore

SMB Library Share. The Nutanix SMB share implementation is the Hyper-V equivalent of an NFS Datastore with feature and performance parity with a vSphere configuration. The registration of a Nutanix storage container as an SMB Library share can be accomplished through a single powershell script, or through the Virtual Machine Manager GUI.

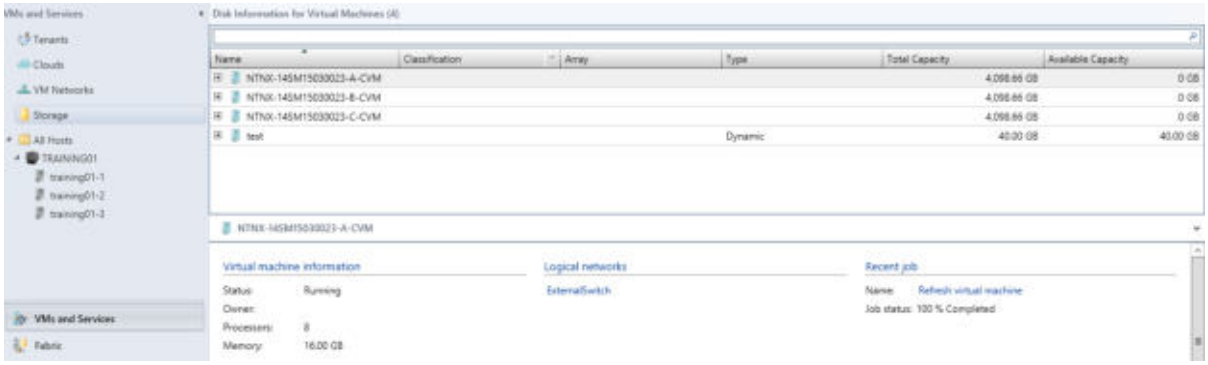

Figure 78: Hyper-V Configuration of an SMB Share

Compression

You can enable compression on a storage container. Compression can save physical storage space and improve I/O bandwidth and memory usage—which may have a positive impact on overall system performance.

> Note: If the metadata usage is high, compression is automatically disabled. If compression is automatically disabled, an alert is generated.

The following types of compression are available.

Post-process compression

Data is compressed after it is written. The delay time between write and compression is configurable, and Nutanix recommends a delay of 60 minutes.

Inline compression

Data is compressed as it is written. When you create a new storage container, inline compression is enabled by default for all license tiers. It is set to a delay of 0, compressing data immediately as it is written.

Compression Ratios

You can view compression ratios and usage savings in the Prism Web Console.

- Cluster

  In the Storage dashboard, under **Capacity Optimization**, click the **After** bar, and hover your mouse over **Compression**.

- Storage container

  In the Storage dashboard **Table** view, on the **Storage Container** tab, click the storage container for which you want to view the compression ratio. You can see the compression ratio for the selected storage container under **Storage Container Details**.

Deduplication

Deduplication reduces space usage by consolidating duplicate data blocks on Nutanix storage. You can enable either cache deduplication only or both cache and capacity deduplication on a storage container.

Cache Deduplication

You can enable cache deduplication of read cache to optimize performance. Cache deduplication means deduplication performed on the data in memory.

> Note: Cache deduplication is not enabled by default.

Capacity Deduplication

You can enable capacity deduplication of persistent data to reduce storage usage. Capcity deduplication means deduplication performed on the data in hard disk storage (HDD). Capacity deduplication is not enabled by default and is available if you have purchased a Pro or higher license. Note that you can enable capacity deduplication only if cache deduplication is enabled.

Enabling Deduplication

Deduplication is enabled at the storage container level (the **Cache** property for cache deduplication and the **Capacity** property for capacity deduplication). These storage container properties can be set in the web console or nCLI.

> Note: Deduplication enabled on storage containers having protected VMs lowers the replication speed.

In addition, Controller VMs in clusters with deduplication enabled need to be configured with more RAM:

- Cache deduplication: 24 GB RAM
- Capacity deduplication: 32 GB RAM

For more information about enabling deduplication, see Creating a Storage Container on page 157.

Deduplication Best Practices

This table shows when deduplication is recommended and when it is not.

Enable deduplication	Do not enable deduplication
• Full clones • Physical-to-virtual (P2V) migration • Persistent desktops	• Linked clones or Nutanix VAAI clones: Duplicate data is managed efficiently by DSF so deduplication has no additional benefit • Server workloads: Redundant data is minimal so may not see significant benefit from deduplication

Erasure Coding

Erasure coding increases the usable capacity on a cluster. Instead of replicating data, erasure coding uses a parity information to rebuild data in the event of a disk failure. The capacity savings of erasure coding is in addition to deduplication and compression savings.

If you have configured redundancy factor 2, two data copies are maintained. For example, consider a 6-node cluster with 4 data blocks (a b c d). In this example, we start with 4 data blocks (a b c d) configured with redundancy factor 2. In the following image, the white text represents the data blocks and the green text represents the copies.

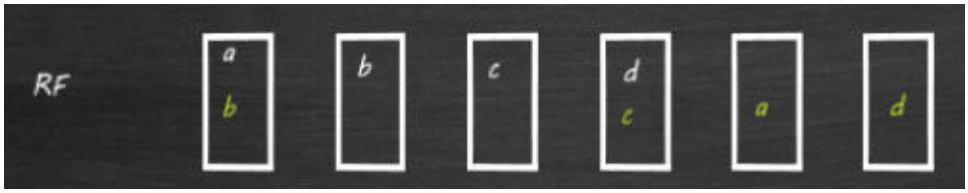

Figure 79: Data copies before Erasure Coding

When the data becomes cold, the erasure code engine computes parity “P” for the data by performing an exclusive OR operation.

Figure 80: Computing Parity

Once parity is computed, the data block copies are removed and replaced with the parity information. Redundancy through parity results in data reduction because the total data on the system is now a+b+c+d+P instead of 2 × (a+b+c+d).

Note: Each block in the stripe is placed on a separate node to protect from a single node failure.

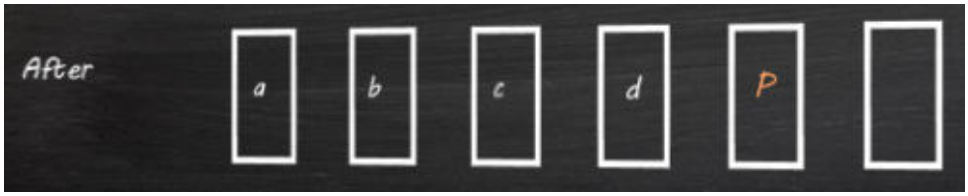

Figure 81: After Computation of Parity

If the node containing data block c fails, block c is rebuilt using the rest of the erasure coded stripe (a b d and P) as displayed in the following image.

Block c is then placed on a node that does not have any other members of this erasure coded stripe.

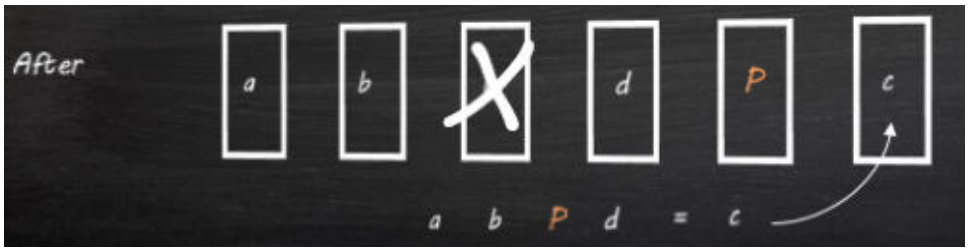

Figure 82: Post Node Failure

Note: When the cluster is configured for the redundancy factor 3, two parity blocks are maintained so that the erasure coded data has the same resiliency as the replicated data. An erasure coded stripe with two parity blocks can survive the failure of two nodes.

Example of Savings from Erasure Coding

The space savings from the erasure coding depends on the cluster size, redundancy setting, and percentage of cold data. You can view the erasure coding usage savings from the storage container summary.

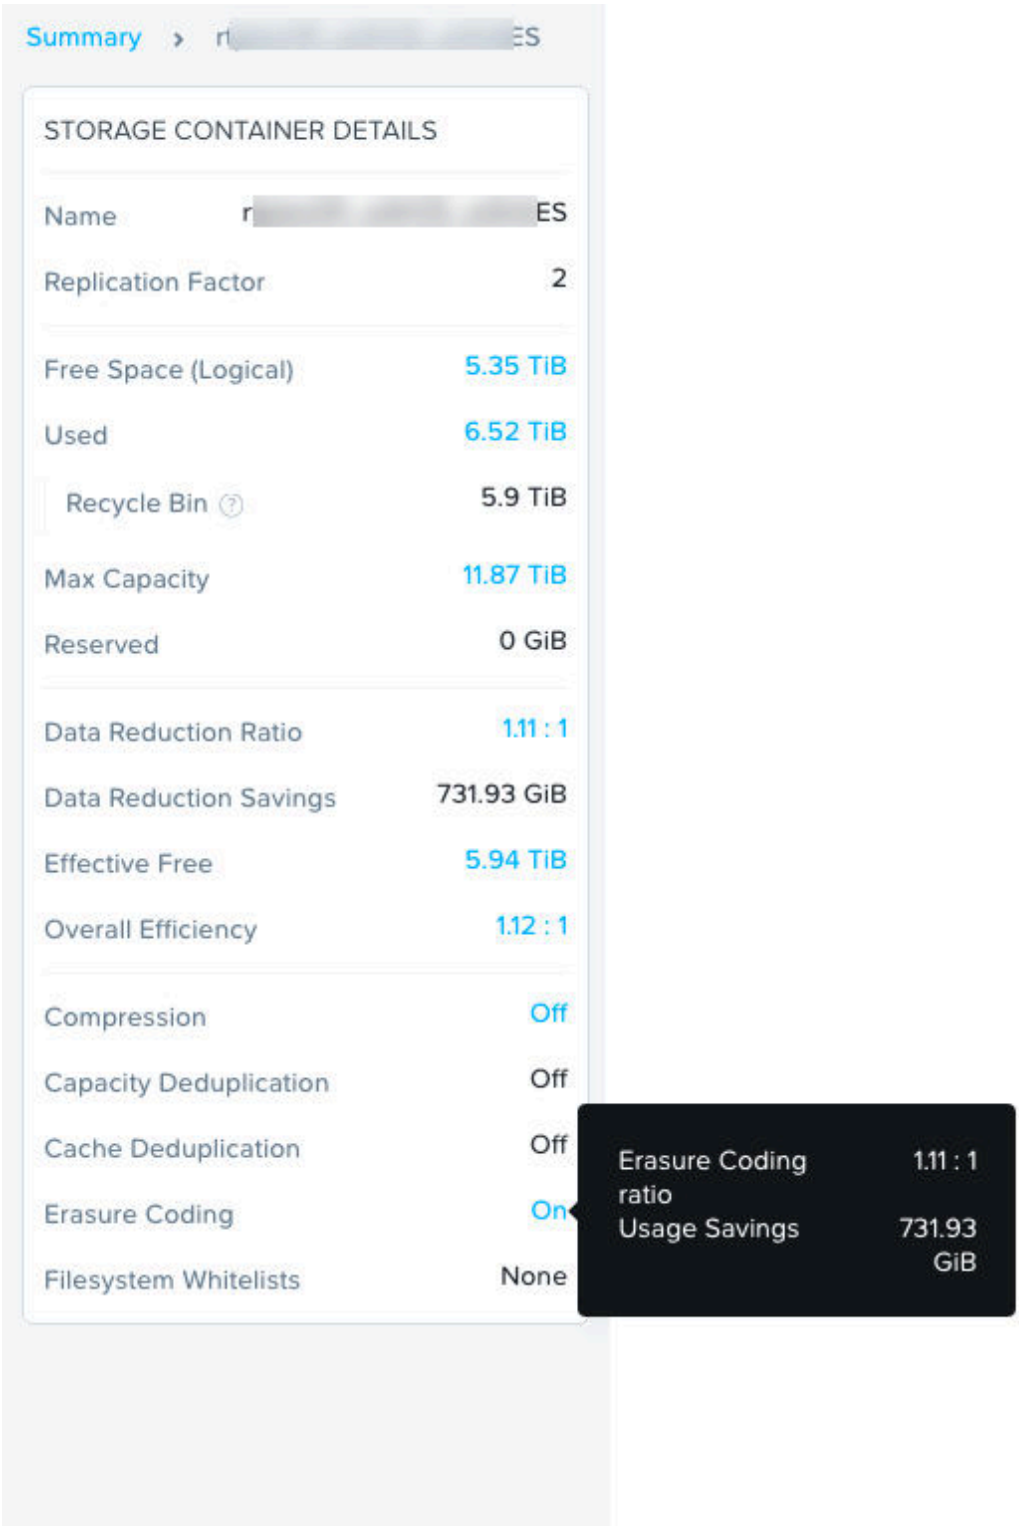

Figure 83: Storage Container Summary: Usage Savings Screen

In a 6-node cluster configured with redundancy factor 2, erasure coding uses a stripe size of 5. 4 nodes for data and 1 node for parity. The sixth node in the cluster ensures that if a node fails, a node is available for rebuild. With a stripe of 4 data to 1 parity, the overhead is 25%. Without erasure coding, the overhead is 100%.

Erasure coding stripe size adapts to the size of the cluster starting with the minimum 4 nodes with a maximum of 5 node stripe width. The following image displays the various configurations of cluster size, possible stripe widths, and approximate savings that might occur when erasure coding is enabled.

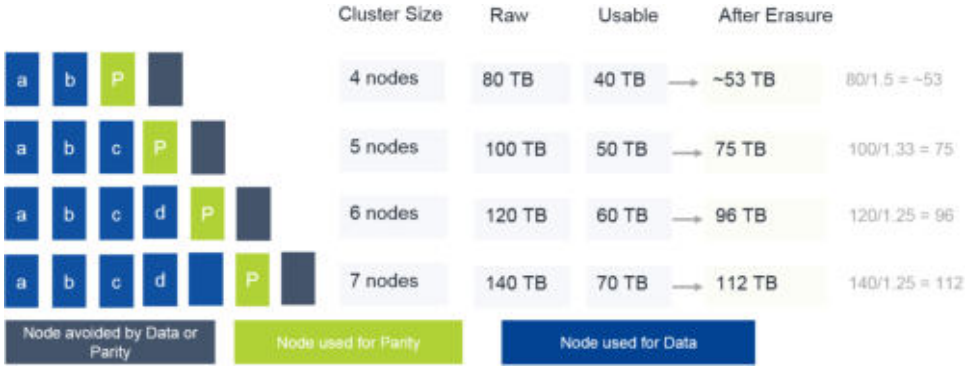

Figure 84: Example of Space Saving from Erasure Coding on 20 TiB Nodes

Erasure Coding Best Practices and Requirements

The following are the recommended best practices and requirements for implementing Erasure Coding:

- A cluster must have at least four nodes/blocks/racks to enable erasure coding. The cluster can have all four flash nodes or a combination of flash and hybrid nodes, or all hybrid nodes. If erasure coding is enabled on a storage container, a minimum of four blocks for RF2 or six blocks for RF3 is required to maintain block awareness.
- The recommended minimum configuration for multiple node removal operations is:

Table 21: Minimum Recommended Configuration for Erasure Coding

Desired Awareness Type	FT Level	Min. Units	Simultaneous Failure Tolerance
Node	1	4 nodes	1 node
Node	2	6 nodes	2 nodes
Block	1	4 blocks	1 block
Block	2	6 blocks	2 blocks
Rack	1	4 racks	1 rack
Rack	2	6 racks	2 racks

Note: We recommend that you always maintain a cluster size that is at least one node greater than the combined strip size (data + parity) to allow space to rebuild the strips if a node fails.

- AOS dynamically calculates the erasure coding strip sizes depending on the number of nodes, blocks, and racks. The maximum supported and recommended strip sizes are (4,1) or (4,2) depending on the nodes, blocks, and racks. We recommend that you should not

change the strip size. Greater strip sizes increases the space savings, however, it increases the cost of rebuild.

- Erasure coding effectiveness (data reduction savings) might reduce on workloads that have many overwrites outside of the erasure coding window. The default value for erasure coding window is seven days for write cold.
- Read performance is affected during rebuild and the amount depends on cluster strip size and read load on the system.
- Erasure coding is an asynchronous process, and hence the time taken to calculate and display space savings depends on the type and coldness of data. A minimum of two full curator scans are required to calculate the data savings.
- We recommend that you have RF+1 storage heavy or storage only nodes for all-flash clusters with storage heavy nodes. For example, if you have a four-node RF2 cluster, then you must add a minimum of three storage heavy nodes for optimum performance.

Inline Erasure Coding

Inline erasure coding encodes and creates erasure coding strips inline by erasure coding data without waiting for the data to become write cold as compared to the regular erasure coding functionality. In inline erasure coding the erasure coded strips are created on the same vDisk data, hence it is recommended to enable inline erasure coding for workloads that do not require data locality.

> Note:
>
> - Inline erasure coding can be enabled only for clusters running AOS version 5.18 or higher.
> - Erasure coding must be enabled on the container before you enable inline erasure coding.
> - If you want to use Inline erasure coding, contact Nutanix support.

Enabling Inline Erasure Coding

Inline erasure coding can be enabled only using nCLI. Inline erasure coding is added as a storage container parameter in Zeus. Note that erasure coding must be enabled on the container to enable inline erasure coding.

About this task

To manually enable erasure coding, run the following nCLI commands.

Procedure

- Enabling erasure coding on the storage container:

  ```
  ncli> container edit name=container_name erasure-code=on
  ```

  Replace *container_name* with the name of the storage container on which you want to enable erasure coding.

- Enabling inline erasure coding on the storage container.

  Run the following nCLI commands:

  ```
  ncli> container create name=container_name sp-id=storage_pool_id erasure-code=on inline-ec-enabled=true
  ```

  Replace *container_name* and *storage_pool_id* with the storage container name and storage pool ID on which you want to enable erasure coding.

  ```
  ncli> container edit inline-ec-enabled=true id=container_id
  ```

  Replace *container_id* with the ID of the storage container.

Capacity Reservation Best Practices

By default, each storage container has access to all of the unused storage in the storage pool. If a storage pool has multiple storage containers, one storage container may take all the remaining storage space and leave others with no available space. To make sure that there is space for a storage container, you can enable capacity reservation.

Capacity reservation allows you to guarantee that a storage container has a minimum amount reserved that is unavailable to other storage containers.

Following are best practices for capacity reservation.

- Reserve capacity for a storage container only if the storage pool has multiple storage containers. Unless there is a specific reason to have multiple storage containers, Nutanix recommends having a single storage pool with a single storage container.
- In total, reserve no more than 90% of the space in the storage pool.
- When setting an advertised capacity for a storage container, be aware that some extra space should be allocated beyond the projected size of any VMs placed in the container to allow room for data that has not yet been garbage collected, which can be substantial depending on the workload (10% or more of the storage capacity in some cases).

Storage Dashboard

The Storage dashboard displays dynamically updated information about the storage configuration in a cluster. To view the Storage dashboard, select **Storage** from the pull-down list on the far left of the main menu.

Menu Options

In addition to the main menu (see Main Menu on page 46), the Storage screen includes a menu bar with the following options:

- View selector. The Storage dashboard provides three viewing modes.
  - Click the **Overview** button on the left to display storage information in a summary view (see Storage Overview View on page 134).
  - Click the **Diagram** button to display a diagram of the storage pools and storage containers in the cluster nodes from which you get detailed information by clicking on a storage pool or storage container of interest (see Storage Diagram View on page 139).
  - Click the **Table** button to display hardware information in a tabular form. The Table screen is further divided into Volume Group, Storage Pool, and Storage Container views; click the **Volume Group** tab to view volume group information, the **Storage Pool** tab to view storage pool information, and the **Storage Container** tab to view storage container information (see Storage Table View on page 146).

- Action buttons. Click the **Volume Group** button on the right to add a volume group to the cluster (see Creating a Volume Group on page 175) in a storage container. Click the **Storage Container** button to add a storage container to a storage pool (see Creating a Storage Container on page 157).
- Page selector. In the Table view, hosts and disks are listed 10 per page. When there are more than 10 items in the list, left and right paging arrows appear on the right, along with the total count and the count for the current page.
- Export table content. In the Table view, you can export the table information to a file in either CSV or JSON format by clicking the gear icon on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.

Storage Overview View

The Storage Overview view displays storage-specific performance and usage statistics on the left plus the most recent storage-specific alert and event messages on the right. 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. The following figure is a sample view, and the table describes each field in this view.

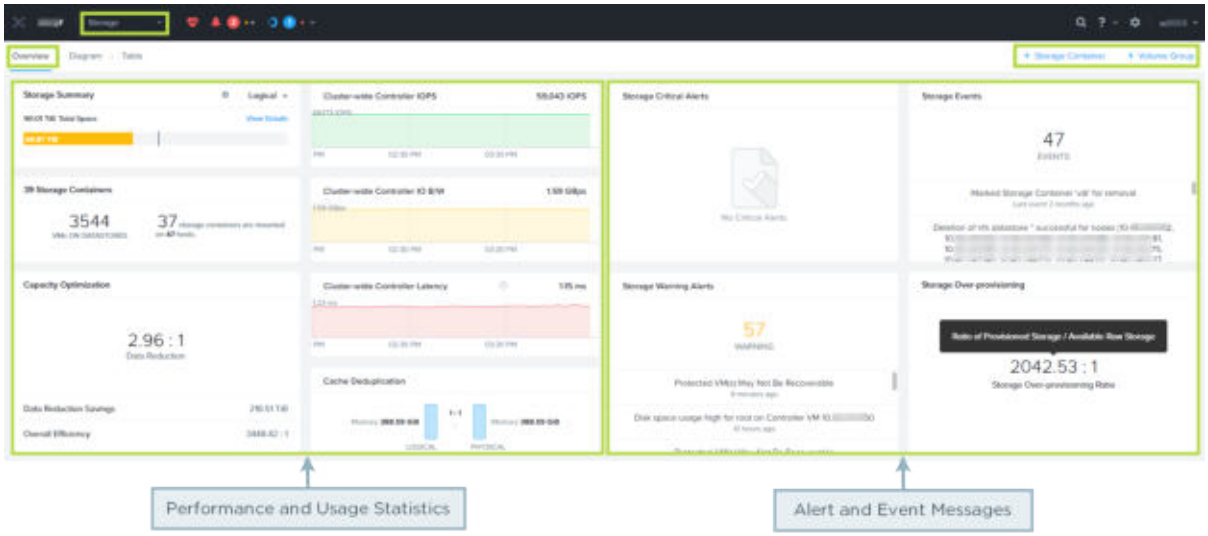

Figure 85: Storage Overview View

Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

Table 22: Storage Overview Widgets

Name	Description
Storage Summary	Displays information about the storage space utilization (in GiB or TiB) and resilient capacity of the cluster. Placing the cursor anywhere on the horizontal axis displays a breakdown view of the storage capacity usage. You can select **Physical** or **Logical** from the drop-down list on the right. **Physical**: Displays information about the physical capacity, which is the total addressable space. **Logical**: Displays information about the logical capacity, which is the usable space after accounting for the data replication factor. The **View Details** link displays the resiliency status and storage information of all the individual nodes in the cluster (see Storage Details Page on page 137). You can also configure a threshold warning for the resilient capacity utilization in the cluster by clicking the gear icon (see Configuring a Warning Threshold for Resilient Capacity on page 167). 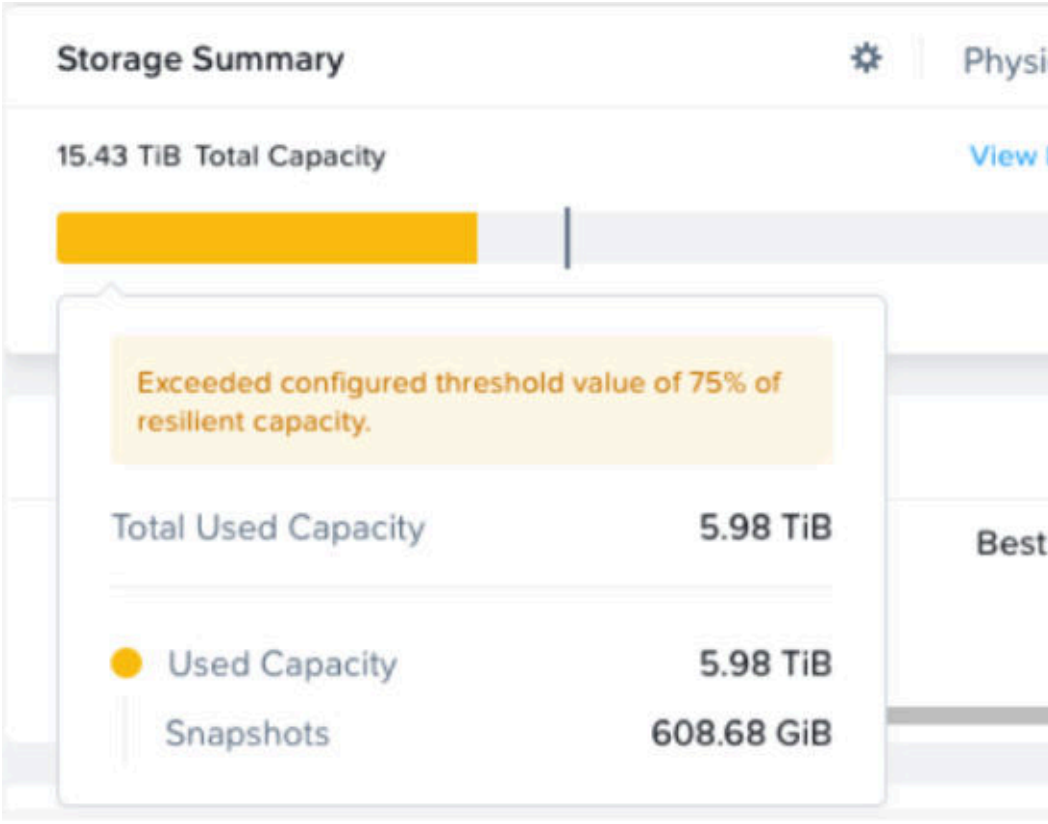
Storage Containers	Displays the number of storage containers, number of VMs, and number of hosts on which the storage containers are mounted in the cluster.

Name	Description
Capacity Optimization	Displays the data reduction ratio (compression, deduplication, and erasure coding), data reduction savings (compression, deduplication, and erasure coding), and the current gained by enabling compression, deduplication, and erasure coding features.
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.)
Cluster-wide Controller IO B/W	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. 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 396).
Cluster-wide Controller Latency	Displays the average I/O latency (in milliseconds) in the cluster.
Cache Deduplication	Displays summary statistics (example below) for memory and flash (SSD) storage used and saved through deduplication (fingerprint-on-write enabled).
Storage Critical Alerts	Displays the five most recent unresolved storage-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the **view all alerts** button at the bottom of the list (see Alerts Dashboard).
Storage Warning Alerts	Displays the five most recent unresolved storage-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the **view all alerts** button at the bottom of the list.
Storage Events	Displays the ten most recent storage-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the **view all events** button at the bottom of the list.

Name	Description
Storage Over-provisioning	Displays the storage over-provisioning ratio (calculated based on the provisioned storage and the available raw storage) in the cluster. Note that the time taken for the **Storage Over-provisioning Ratio** widget to reflect the changes made in the cluster varies according to the recent storage operations/activities performed.

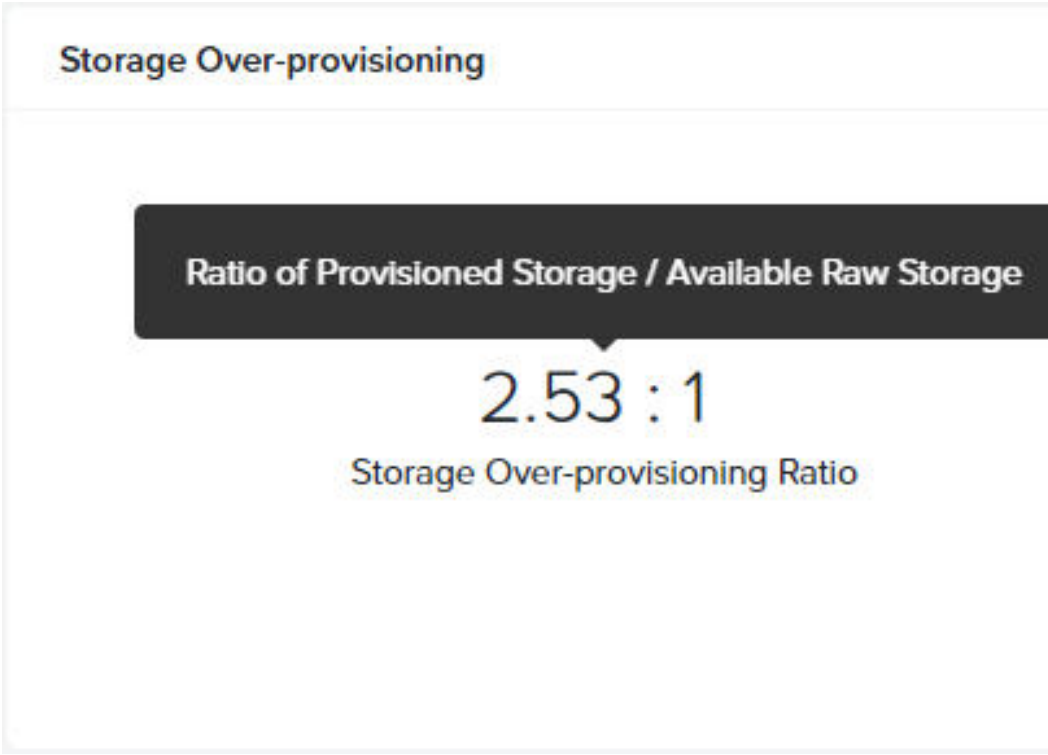

Storage Details Page

The Storage Details page displays the resiliency status and storage information (both logical and physical) of all the individual nodes in the cluster.

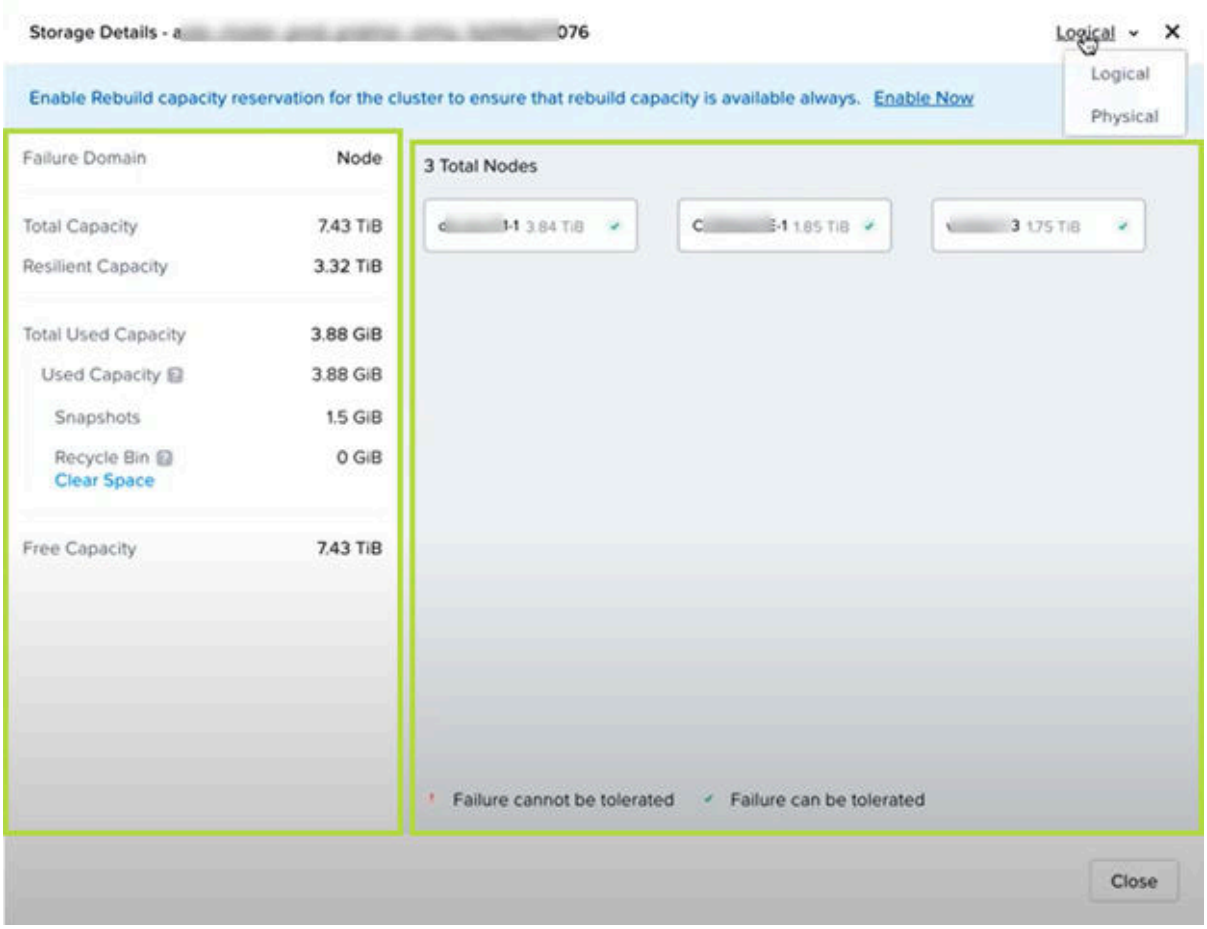

The Storage Details page is divided into two sections:

- The right section displays a diagrammatic representation of the number of nodes present in the cluster along with the respective storage capacity used. Select **Physical** or **Logical** from the drop-down list on the right to switch between the physical storage capacity view and the logical capacity view of the cluster.
- The left section provides detailed storage information of the cluster as follows.

Parameter	Description	Values
Failure Domain	Displays the entity (node, block, or rack) which if failed, the cluster can tolerate while still running guest VMs and responding to commands through the management console.	node, block, or rack
Total Capacity	Displays the total capacity of all the disks on all the hosts in the cluster.	xxx [GB|TB]
Resilient Capacity	Displays the total resilient capacity of the cluster.	xxx [GB|TB]

Parameter	Description	Values
Total Used Capacity	Displays the sum of all the storage space used by the cluster. The total used capacity is calculated based on the following: • **Used Capacity**. The amount of used storage space in the cluster (by user data). • **Snapshots**. The total storage capacity in the cluster consumed by snapshots (sum of both local and remote). • **Recycle Bin**. The total storage capacity in the cluster consumed by VMs in recycle bin.	xxx [GB|TB]
Free Capacity	Displays the available storage capacity on all the disks on all the hosts in the cluster.	xxx [GB|TB]

Storage Diagram View

The Storage Diagram view displays information about storage pools and storage containers. The displayed information is dynamically updated to remain current.

The Storage Diagram view screen is divided into two sections:

- The top section is a cascading diagram of the storage units. Initially, a cluster bar appears with storage information about the cluster (used, provisioned, and available storage) and colored blocks in the bar for each storage pool. Clicking on a storage pool block displays storage information about the storage pool and a bar with blocks for each storage container in that storage pool. Clicking on a storage container block displays storage information about the storage container and a bar for that storage container. The storage container bar includes **Physical** and **Logical** buttons on the right that you can click to display information about the physical capacity, which is the total addressable space, or logical capacity, which is the usable space after accounting for the data replication factor. Because the default replication factor is set to two, the logical space is typically half the physical space. You can edit a storage pool or storage container by clicking the pencil (edit) or X (delete) icon to the right of the name. Clicking the **close** link at the far right hides that storage pool or storage container bar from the display.
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

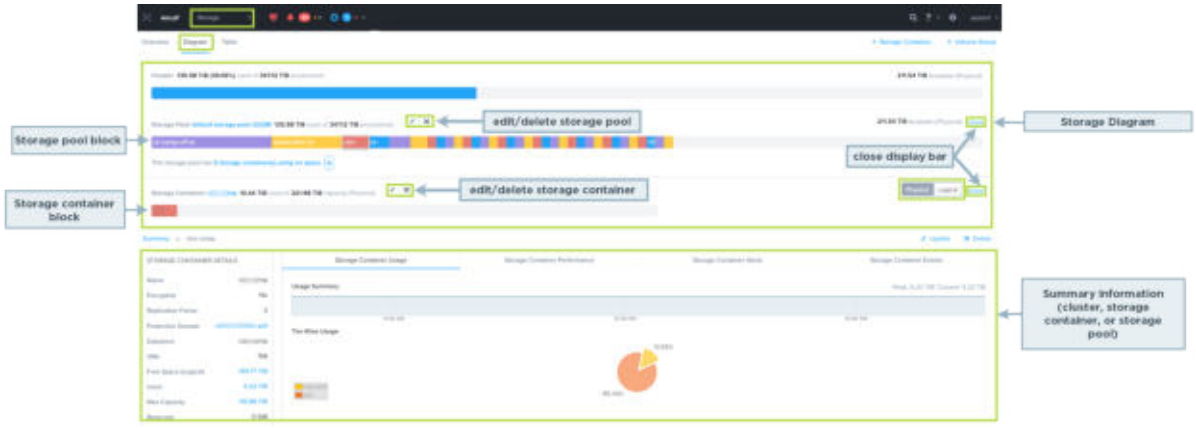

Figure 86: Storage Diagram View

Storage Container Details

Selecting a storage container in the diagram displays information about that storage container in the lower section of the screen (see Creating a Storage Container on page 157).

- When a storage container is selected, **Summary:** ***storage_container_name*** appears below the diagram, and action links appear on the right of this line (see Modifying a Storage Container on page 162 for more information about these actions):
  - Click the **Update** link to update the settings for this storage container.
  - Click the **Delete** link to delete this storage container configuration.
- In the **Summary:** ***storage_container_name*** table, hover your mouse over the value to see additional details of that parameter.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Storage Container Usage**, **Storage Container Performance**, **Storage Container Alerts**, **Storage Container Events**.

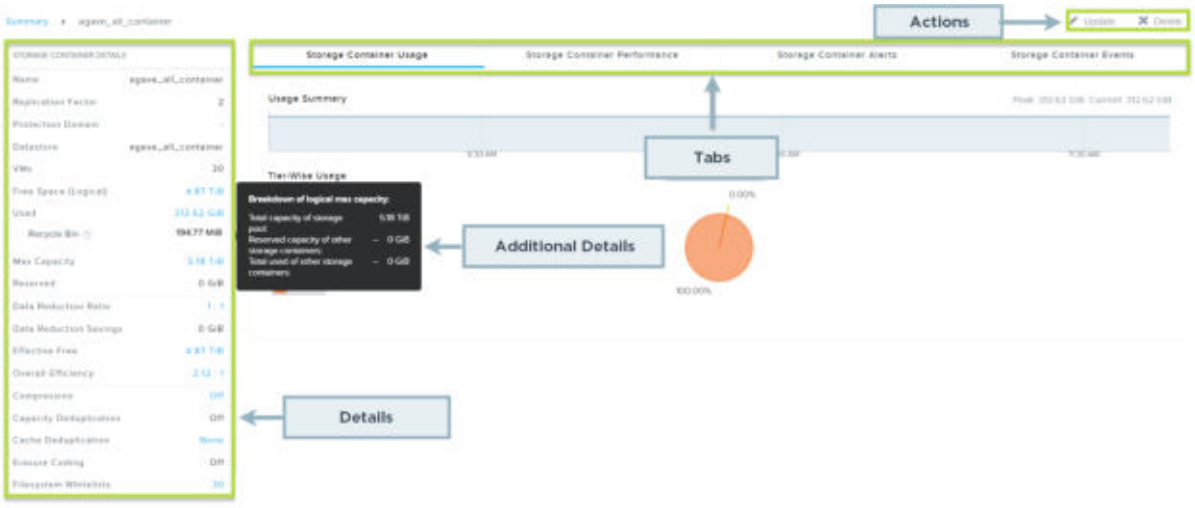

Figure 87: Storage Diagram View: Storage Container

Table 23: Storage Container Details Fields

Parameter	Description	Values
Name	Displays the name of the storage container.	(name)
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Protection Domain	Displays the name of the protection domain if the you have a configured a protection domain for that storage container.	(name)
Datastore	Displays the name of the datastore.	(name)
VMs	Displays the number of VMs	(number)
Free Space (Logical)	Displays the amount of free storage space available to the storage container that is unreserved.	xxx [GB|TB]
Used	Displays the amount of used storage space in the storage container, including space used by Snapshots and Recycle Bin.	xxx [GB|TB]
Max Capacity	Displays the total amount of storage capacity available to 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 the total available storage regardless of how many storage containers are defined. Therefore, when you have two storage containers, it can appear that you have twice as much capacity because maximum capacity for both storage containers show the full amount. Maximum capacity is calculated as the total logical capacity in the storage pool, minus any reserved capacity, minus space used by other storage containers.	xxx [TB]
Reserved	Displays the total reserved storage capacity in the storage container.	xxx [GB|TB]
Data Reduction Ratio	Displays the ratio of how much the data size is reduced by enabling compression, deduplication, and erasure coding.	
Data Reduction Savings	Displays the data reduction savings by enabling compression, deduplication, and erasure coding.	
Effective Free	Displays the amount of usable free space after data reduction.	

Parameter	Description	Values
Overall Efficiency	Displays the overall efficiency that you can achieve by enabling compression, deduplication, and erasure coding features.	
Compression	Displays whether compression is enabled.	[Off|On]
Capacity Deduplication	Displays whether on disk deduplication is enabled on hard disks (HDD). Cache deduplication is a prerequisite for Capacity deduplication.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication when data is read. Data duplication (commonly referred to as dedup) is a specialized data reduction technique for eliminating duplicate copies of data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, Off, On]
Erasure Coding	Displays whether erasure coding is enabled.	[On, Off]
Filesystem Whitelists	Displays whether you have configured filesystem whitelist for this storage container.	[None, On, Off]

Storage Pool Details

Selecting a storage pool in the diagram displays information about that storage pool in the lower section of the screen.

- When a storage pool is selected, **Summary:** ***storage_pool_name*** appears below the diagram, and action links appear on the right of this line (see Modifying a Storage Pool on page 156 for more information about this action):
  - Click the **Update** link to update the settings for this storage pool.
- Four tabs appear that display information about the selected storage pool (see following sections for details about each tab): **Storage Pool Usage**, **Storage Pool Performance**, **Storage Pool Alerts**, **Storage Pool Events**.

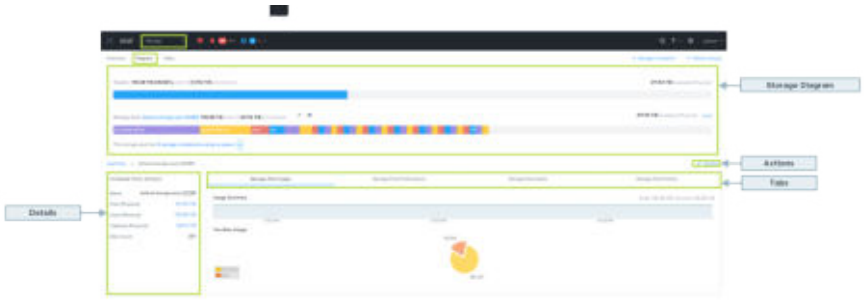

Figure 88: Storage Diagram View: Storage Pool

Table 24: Storage Pool Details Fields

Parameter	Description	Values
Name	Displays the name of the storage pool.	(name)
Free (Physical)	Displays the total amount of physical storage space that is available.	xxx [GB|TB]
Used (Physical)	Displays the total amount of physical storage space used in the storage pool.	xxx [GB|TB]
Capacity (Physical)	Displays the total physical storage space capacity in the storage pool.	xxx [TB]
Disk Count	Displays the number of disks in the storage pool.	(number)

Cluster Summary Information

When a storage container or storage pool is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Storage Summary** column (on the left) includes five fields:
  - **Free (Physical)**. Displays the amount of physical storage space still available in the cluster.
  - **Used (Physical)**. Displays the amount of physical storage space used currently in the cluster, including the Recycle Bin.
  - **Capacity (Physical)**. Displays the total physical storage capacity in the cluster.
  - **Storage Pool(s)**. Displays the names of the storage pool. Clicking on a name displays the detail information for that storage pool in this section.
  - **Storage Container(s)**. Displays the names of the storage containers. Clicking on a name displays the detail information for that storage container in this section.
- Four tabs appear that display cluster-wide information (see following sections for details about each tab): **Usage Summary**, **Performance Summary**, **Storage Alerts**, **Storage Events**.

Usage Tab

The Usage tab displays graphs of storage usage. The tab label varies depending on what is selected in the table:

- **Usage Summary** (no storage pool or storage container selected). Displays usage statistics across the cluster.
- **Storage Container Usage** (storage container selected). Displays usage statistics for the selected storage container.
- **Storage Pool Usage** (storage pool selected). Displays usage statistics for the selected storage pool.

The Usage tab displays the following two graphs:

- **Cluster-wide Usage Summary**: Displays a rolling time interval usage monitor 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 (see Analysis Dashboard on page 396).

- **Tier-wise Usage**: Displays a pie chart divided into the percentage of storage space used by each disk tier in the cluster, storage pool, or storage container. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.

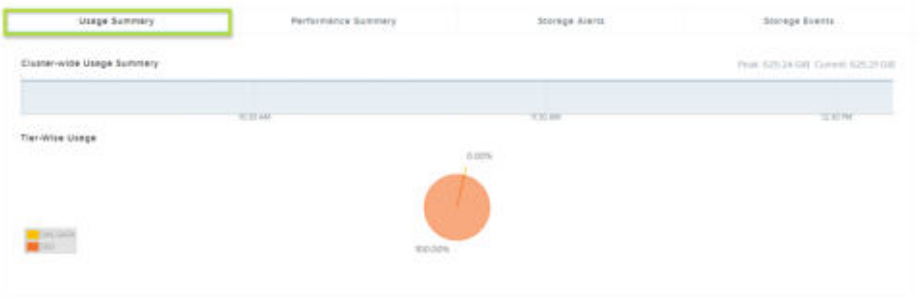

Figure 89: Storage Table View: Usage Tab

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label varies depending on what is selected in the table:

- **Performance Summary** (no storage pool or storage container selected). Displays storage performance statistics across the cluster.
- **Storage Container Performance** (storage container selected). Displays storage performance statistics for the selected storage container.
- **Storage Pool Performance** (storage pool selected). Displays storage performance statistics for the selected storage pool.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see Analysis Dashboard on page 396). The Performance tab includes the following three graphs:

- **[Cluster-wide Hypervisor|Controller|Disk] IOPS**: Displays I/O operations per second (IOPS) for the cluster, selected storage container, or selected storage pool.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected storage container, or selected storage pool.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Latency**: Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected storage container, or selected storage pool.

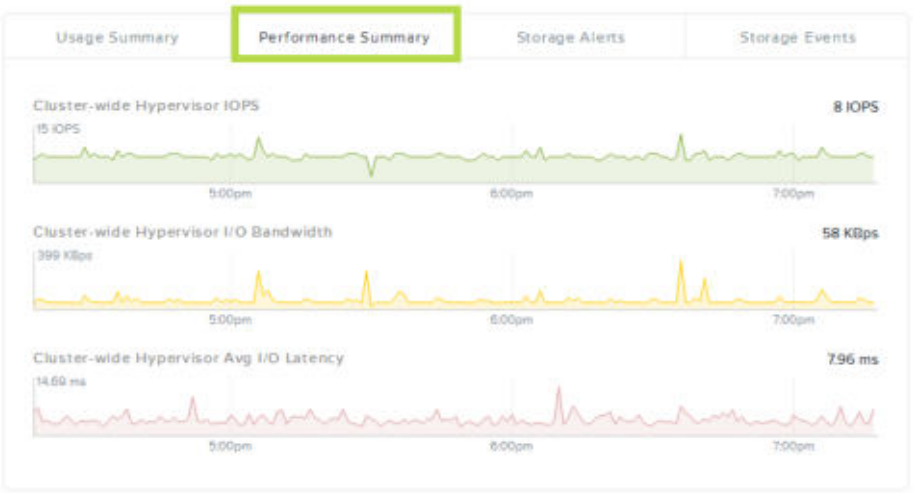

Figure 90: Storage Table View: Performance Tab

Storage Alerts Tab

The Alerts tab displays the unresolved alert messages about storage pools or storage containers in the same form as the Alerts page (see Alerts Summary View on page 412). Click the **Unresolved X** button in the filter field to also display resolved alerts.

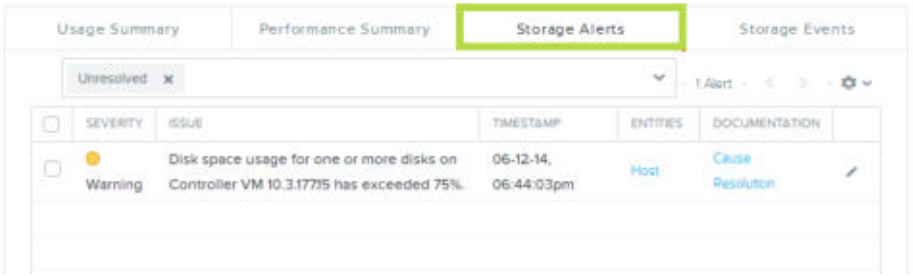

Figure 91: Storage Table View: Alerts Tab

Storage Events Tab

The Events tab displays the unacknowledged event messages about storage pools or storage containers in the same form as the Events page (see Events Summary View on page 418). Click the **Include Acknowledged** button to also display acknowledged events.

Figure 92: Storage Table View: Events Tab

Storage Table View

The Storage Table view displays information about storage pools and storage containers in a tabular form. Click the **Volume Group** tab to display volume group information; click the **Storage Pool** tab in the screen menu bar to display storage pool information; click the **Storage Container** tab to display storage container information. The displayed information is dynamically updated to remain current.

The Storage Table view is divided into two sections:

- The top section is a table. Each row represents a single volume group, storage pool, or storage container and includes basic information about that entity. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

> Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

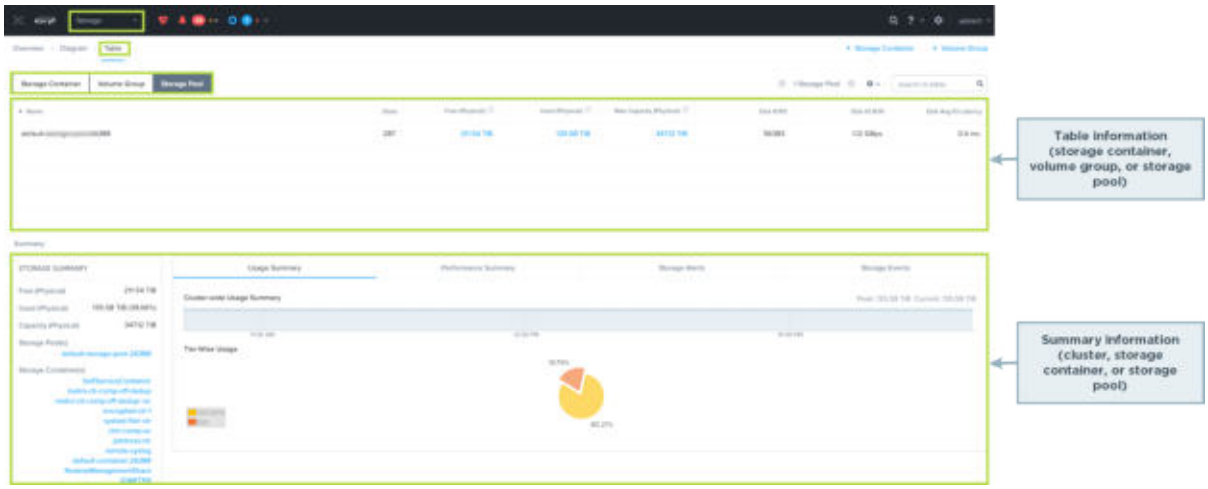

Figure 93: Storage Table View

Volume Group tab

Clicking the Volume Group tab displays information about volume groups in the cluster (see Creating a Volume Group on page 175).

- The table at the top of the screen displays information about all the configured volume groups, and the details column (lower left) displays additional information when a volume group is selected in the table. The following table describes the fields in the volume group table and detail column.

- When a volume group is selected, Summary: *volume_group_name* appears below the table, and action links appear on the right of this line (see Modifying a Volume Group on page 177 for more information about these actions):

  - Click the **Update** link to update the settings for this volume group.
  - Click the **Delete** link to delete this volume group.

  Five tabs appear that display information about the selected volume group (see following sections for details about each tab): Performance Metrics, Virtual Disks, Volume Group Tasks, Volume Group Alerts, and Volume Group Events.

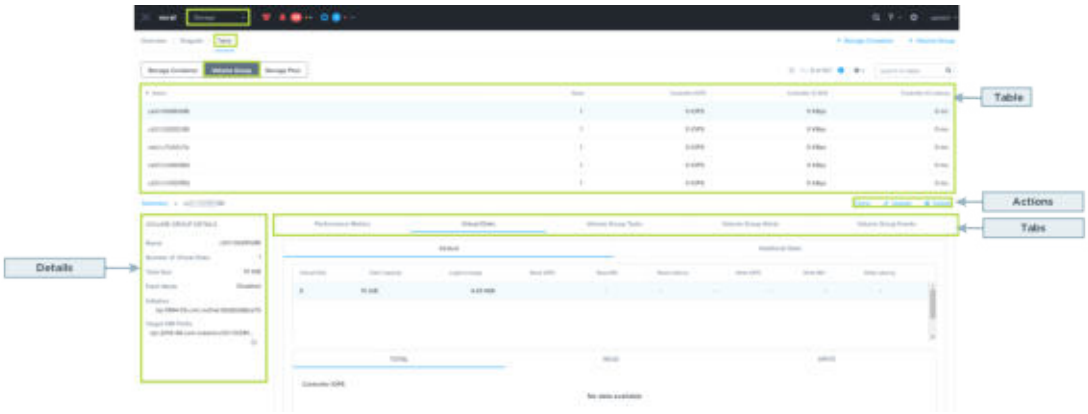

Figure 94: Storage Table View: Volume Group

Table 25: Volume Group Table and Detail Fields

Parameter	Description	Values
Volume Group Table Fields (upper screen)		
Name	Displays the name of the volume group.	(name)
Disks	Displays the number of disks in the volume group.	[0–256]
Controller IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
Controller IO B/W	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]
Controller IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Volume Group Details Fields (lower screen)		
Name	Displays the name of the volume group.	(name)

Parameter	Description	Values
Number of Virtual Disks	Displays the number of virtual disks in the volume group.	[0–256]
Total Size	Displays the total size of the volume group.	xxx [GB|TB]
Flash Mode	Displays whether the flash mode is enabled.	
Shared	Indicates whether the volume group is shared across iSCSI initiators.	[Yes|No]
Initiators	Displays the iSCSI initiators to which the volume group is attached.	(None|List of names)
Target IQN Prefix	Displays the IQN prefix of the target iSCSI.	

Storage Container Tab

Clicking the **Storage Container** tab displays information about storage containers in the cluster (see Creating a Storage Container on page 157).

- The table at the top of the screen displays information about all the configured storage containers, and the details column (lower left) displays additional information when a storage container is selected in the table. The following table describes the fields in the storage container table and detail column.
- When a storage container is selected, **Summary:** ***storage_container_name*** appears below the table, and action links appear on the right of this line (see Modifying a Storage Container on page 162 for more information about these actions):
  - Click the **Update** link to update the settings for this storage container.
  - Click the **Delete** link to delete this storage container configuration.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Storage Container Breakdown**, **Storage Container Usage**, **Storage Container Performance**, **Storage Container Alerts**, **Storage Container Events**.

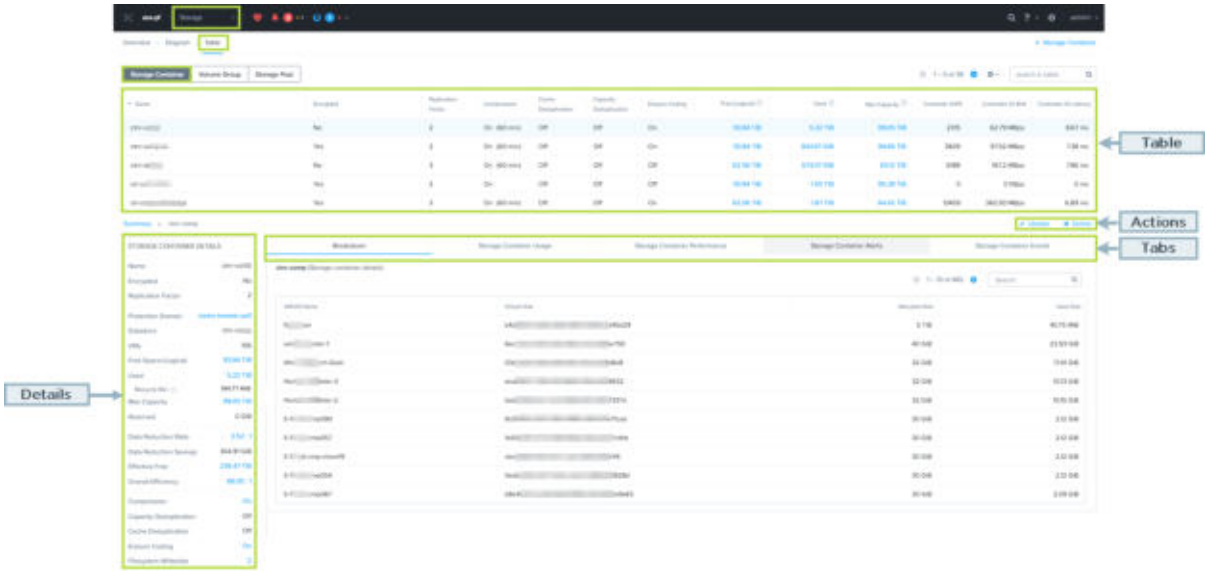

Figure 95: Storage Table View: Storage Container

Table 26: Storage Container Table and Detail Fields

Parameter	Description	Values
Storage Container Table Fields (upper screen)		
Name	Displays the name of the storage container.	(name)
Encrypted	Displays the encryption status of the storage container.	[Yes| No]
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a data reduction technique for eliminating duplicate copies of data. Setting this parameter to **On** causes dedup to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether on disk deduplication is enabled on hard disks (HDD). Cache deduplication is a prerequisite for Capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled for the storage container or not	[On, Off]

Parameter	Description	Values
Free (logical)	Displays the amount of free storage space in the storage container.	xxx [GB|TB]
Used	Displays the amount of used storage space in the storage container, including space used by the Recycle Bin.	xxx [GB|TB]
Max Capacity	Displays the total amount of storage capacity available to 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 maximum capacity for both storage containers show the full amount. Maximum capacity is calculated as the total logical capacity in the storage pool, minus any reserved capacity, minus space used by other storage containers.	xxx [TB]
Controller 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]
Controller IO B/W	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
Controller IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
Storage Container Details Fields (lower screen)		
Name	Displays the name of the storage container.	(name)
Encrypted	Displays the encryption status of the storage container.	
Protection Domain	Displays the data protection domain used for the storage container.	(DR name)
VMs	Displays the number of VMs associated with the storage container.	xxx
Free Space (Logical)	Displays the amount of free storage space available to the storage container that is unreserved.	xxx [GB|TB]
Used	Displays the amount of used storage space for the storage container.	xxx [GB|TB]

Parameter	Description	Values
Max Capacity	Displays the total amount of storage capacity available to the storage container (see the Reserved Capacity description).	xxx [TB]
Reserved	Displays the total reserved storage capacity in the storage container.	xxx [GB|TB]
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]
Data Reduction Ratio	Displays the ratio of how much the data size is reduced by enabling compression, deduplication, and erasure coding.	x.xx : 1
Data Reduction Savings	Displays the data reduction savings by enabling compression, deduplication, and erasure coding.	xxx [GB|TB]
Effective Free	Displays the amount of usable free space after data reduction.	xxx [GB|TB]
Overall Efficiency	Displays the overall efficiency that you can achieve by enabling compression, deduplication, and erasure coding features.	xxx [GB|TB]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether on disk deduplication is enabled, that is dedup compression applied to data on hard disks (HDD). Cache deduplication is a prerequisite for Capacity deduplication.	[On, Off]
Filesystem Whitelists	Displays whether you have configured filesystem whitelist for this storage container.	[None, On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]

Storage Pool Tab

Clicking the **Storage Pool** tab displays information about storage pools in the cluster.

- The table at the top of the screen displays information about the storage pool, and the details column (lower left) displays additional information when a storage pool is selected in the table. The following table describes the fields in the storage pool table and detail column.

- When a storage pool is selected, **Summary:** ***storage_pool_name*** appears below the table, and action links appear on the right of this line (see Modifying a Storage Pool on page 156 for more information about this action):
  - Click the **Update** link to update the settings for this storage pool.
- Four tabs appear that display information about the selected storage pool (see following sections for details about each tab): **Storage Pool Usage**, **Storage Pool Performance**, **Storage Pool Alerts**, **Storage Pool Events**.

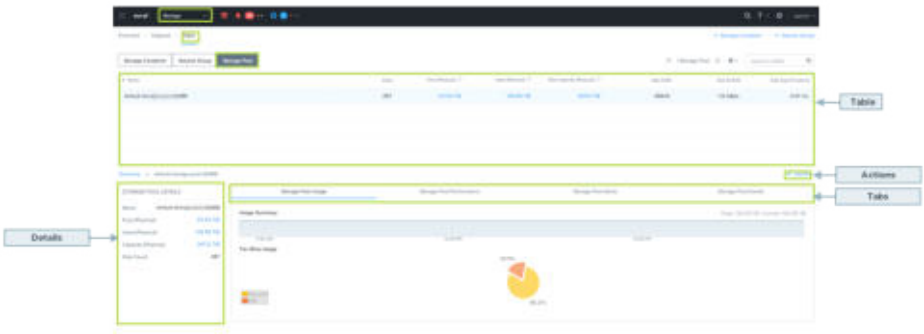

Figure 96: Storage Table View: Storage Pool

Table 27: Storage Pool Table and Detail Fields

Parameter	Description	Values
Storage Pool Table Fields (upper screen)		
Name	Displays the name of the storage pool.	(name)
Disks	Displays the number of disks in the storage pool.	(number)
Free (Physical)	Displays the total amount of physical storage space that is available.	xxx [GB|TB]
Used (Physical)	Displays the total amount of physical storage space used in the storage pool.	xxx [GB|TB]
Max Capacity (Physical)	Displays the total physical storage space capacity in the storage pool.	xxx [TB]
Disk IOPS	Displays the current I/O operations per second (IOPS) for the storage pool. The IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by physical disks across the storage pool.	[0 - unlimited]
Disk IO B/W	Displays the I/O bandwidth used per second for physical disk requests in this storage pool.	xxx [MBps|KBps]
Disk Avg IO Latency	Displays the average I/O latency for physical disk requests in this storage pool.	xxx [ms]
Storage Pool Details Fields (lower screen)		

Parameter	Description	Values
Name	Displays the name of the storage pool.	(name)
Free (Physical)	Displays the total amount of physical storage space that is available.	xxx [GB|TB]
Used (Physical)	Displays the total amount of physical storage space used in the storage pool.	xxx [GB|TB]
Capacity (Physical)	Displays the total physical storage space capacity in the storage pool.	xxx [TB]
Disk Count	Displays the number of disks in the storage pool.	(number)

Cluster Summary Information

When a storage pool, storage container, or volume group is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Storage Summary** column (on the left) includes five fields:
  - **Available (Physical)**. Displays the amount of physical storage space still available in the cluster.
  - **Used (Physical)**. Displays the amount of physical storage space used currently in the cluster.
  - **Capacity (Physical)**. Displays the total physical storage capacity in the cluster.
  - **Storage Pool**. Displays the name of the storage pool in the cluster. Clicking the name displays detailed information about the storage pool in this section.
  - **Storage Container(s)**. Displays the names of the storage containers. Clicking a name displays detailed information about that storage container in this section.
- Four tabs appear that display cluster-wide information (see following sections for details about each tab): **Usage Summary**, **Performance Summary**, **Storage Alerts**, **Storage Events**.

Breakdown Tab

The **Breakdown** tab is displayed in the **Summary** section only when a storage container is selected from the storage container table.

The **Breakdown** tab displays the list of virtual disks, associated VMs or volume groups, and the amount of storage space utilized by each in the selected storage container.

Figure 97: Storage Table View: Breakdown Tab

Usage Tab

The Usage tab displays graphs of storage usage. The tab label varies depending on what is selected in the table:

- **Usage Summary** (no storage pool, storage container, or volume group selected). Displays usage statistics across the cluster.
- **Storage Container Usage** (storage container selected). Displays usage statistics for the selected storage container.
- **Storage Pool Usage** (storage pool selected). Displays usage statistics for the selected storage pool.
- **Volume Group Usage** (volume group selected). Displays usage statistics for the selected volume group.

The Usage tab displays the following two graphs:

- **Usage Summary**: Displays a rolling time interval usage monitor 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 (see Analysis Dashboard on page 396).
- **Tier-wise Usage**: Displays a pie chart divided into the percentage of storage space used by each disk tier in the cluster, storage pool, or storage container. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.

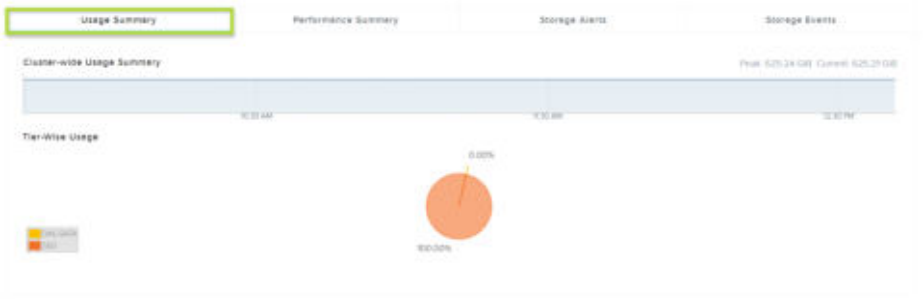

Figure 98: Storage Table View: Usage Tab

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label varies depending on what is selected in the table:

- **Performance Summary** (no storage pool, storage container, or volume group selected). Displays storage performance statistics across the cluster.
- **Storage Container Performance** (storage container selected). Displays storage performance statistics for the selected storage container.
- **Storage Pool Performance** (storage pool selected). Displays storage performance statistics for the selected storage pool.
- **Volume Group Performance** (volume group selected). Displays storage performance statistics for the selected volume group.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see Analysis Dashboard on page 396). The Performance tab includes the following three graphs:

- **[Cluster-wide Hypervisor|Controller|Disk] IOPS**: Displays I/O operations per second (IOPS) for the cluster, selected storage container, selected storage pool, or selected volume group.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected storage container, selected storage pool, or selected volume group.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Latency**: Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected storage container, selected storage pool, or selected volume group.

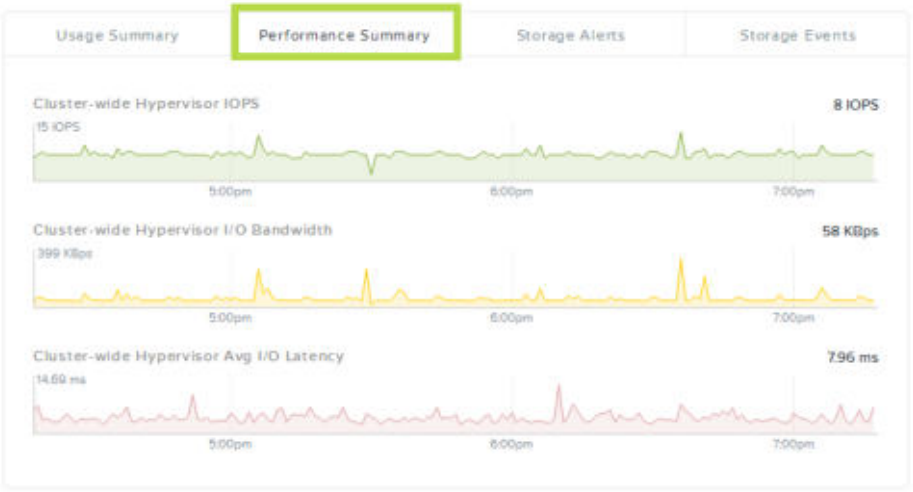

Figure 99: Storage Table View: Performance Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about storage pools, storage containers, or volume groups in the same form as the Alerts page (see Alerts Summary View on page 412). Click the **Unresolved X** button in the filter field to also display resolved alerts.

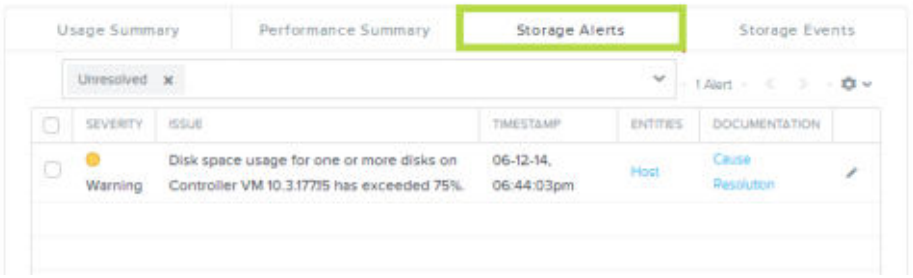

Figure 100: Storage Table View: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about storage pools, storage containers, or volume groups in the same form as the Events page (see Events Summary View on page 418). Click the **Include Acknowledged** button to also display acknowledged events.

Figure 101: Storage Table View: Events Tab

Modifying a Storage Pool

A storage pool is a defined group of physical disks that can be modified as the cluster changes. A storage pool is created automatically when the cluster is created.

About this task

You can only modify the name of the storage pool in the cluster. You cannot delete a storage pool.

To modify the name of the storage pool, do the following.

Procedure

1. In the web console select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Pool** tabs.
2. To update the storage pool, select the target storage pool and then click the **Update** link. The Update Storage Pool window appears displaying the current name and capacity of the storage pool.
3. Enter the new name of the storage pool in the **Name** field.
4. Click **Save**.

Creating a Storage Container

A storage container is a subset of available storage within a storage pool. Storage efficiency features including compression, deduplication and erasure coding are enabled at the container level.

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers on page 753) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor.

- Hyper-V: The storage container is accessible as an SMB share.
- vSphere: The storage container is accessible as an NFS datastore. This requires access to the vSphere APIs. Ensure that you have appropriate license of vSphere to access the APIs.
- AHV: The storage container is accessible transparently.

Note:

- A storage pool and one storage container are created automatically when the cluster is created.
- A storage container is not created if you have not configured the Controller VMs with enough memory. Controller VM memory allocation requirements differ depending on the models and features that are being used. For more information, see *Controller VM Memory Configurations* in *Acropolis Advanced Administration Guide*.
- The NutanixManagementShare storage container is a built-in storage container for Nutanix clusters for use with the Nutanix Files and Self-Service Portal (SSP) features. This storage container is used by Nutanix Files and SSP for file storage, feature upgrades, and other feature operations. To ensure proper operation of these features, do not delete this storage container. Nutanix also recommends that you do not delete this storage container even if you are not using these features. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.

To create a storage container, do the following:

Procedure

1. In the Storage dashboard (see Storage Dashboard on page 133), click the **Storage Container** button.
   The Create Storage Container dialog box appears.

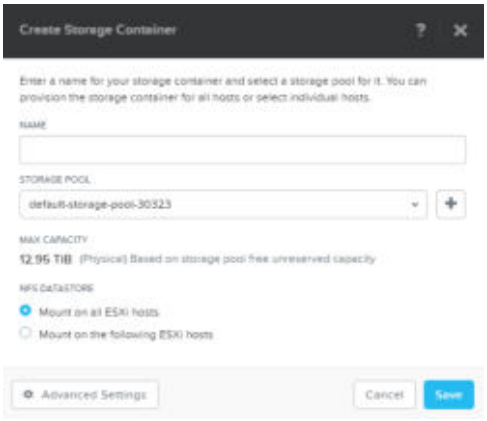

Figure 102: Create Storage Container (vSphere)

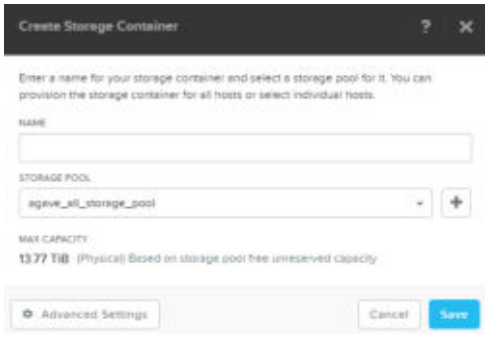

Figure 103: Create Storage Container (AHV)

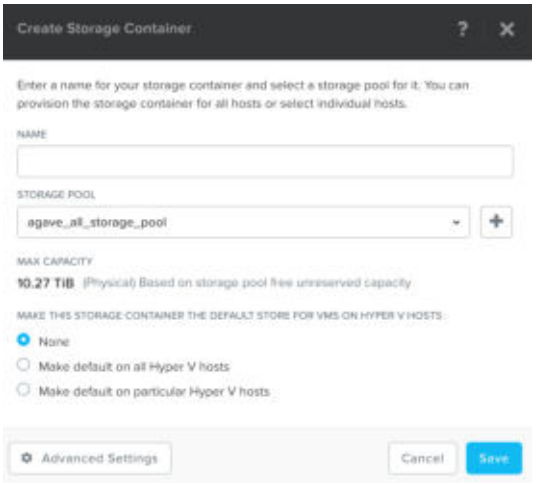

Figure 104: Create Storage Container (Hyper-V)

2. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

   Note: This entity has the following naming restrictions across various hypervisors.

   AHV:

   - Container Name Length: Maximum length is 75 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   - Case Sensitivity: Container names are case sensitive.

   ESXi:

   - Container Name Length: Maximum length is 42 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   - Case Sensitivity: Container names are case sensitive.

   Hyper-V:

   - Container Name Length: Maximum length is 32 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), decimal digits (0-9), dots (.), hyphens (-), and underscores(_).
   - Case Sensitivity: Container names are case insensitive. For example, if there is a container with name ABCD, then creating another container with name abcd is restricted.

   b. **Storage Pool**: Select a storage pool from the drop-down list.

   The following field, **Max Capacity**, displays the amount of free space available in the selected storage pool.

   c. (vSphere only) **NFS Datastore**: Select the **Mount on all ESXi hosts** button to mount the storage container on all hosts. Select the **Mount on the following ESXi hosts** button to mount the storage container on a subset of hosts, which displays a list of host IP addresses below this field. Check the boxes of the hosts to be included.

   d. (Hyper-V only) Set this storage container as default store for VMs on Hyper-V hosts.

   Depending on your selection the Create Virtual Machine Wizard of Hyper-V automatically populates the storage location with the relevant storage container.

Option	Description
Make default on all Hyper-V hosts	Makes this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.
Make default on particular Hyper-V hosts	Provides you with an option to select the hosts on which you want to make this storage container a default location for storing virtual machine configuration and

Option	Description
	virtual hard disk files on all the Hyper-V hosts.

3. To configure additional parameters, click the **Advanced Settings** button.

Figure 105: Create Storage Container Advanced Settings (1)

Do the following in the indicated fields:

a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

   Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > Note: To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See Increasing the Cluster Fault Tolerance Level on page 102.

   Nutanix supports a replication factor (RF) of 1 if you first select the **Enable Replication Factor 1** checkbox at the **Settings** > **Redundancy State** page. For more information about this setting, see Replication Factor 1 Overview on page 107 and Replication Factor 1 Recommendations and Limitations on page 107. If you do not select **Enable Replication Factor 1**, Nutanix supports an RF of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > Note: To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See Increasing the Cluster Fault Tolerance Level on page 102.

b. **Reserved Capacity**: To reserve storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 133 for more information.

c. **Advertised Capacity**: Sets a maximum storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised

capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

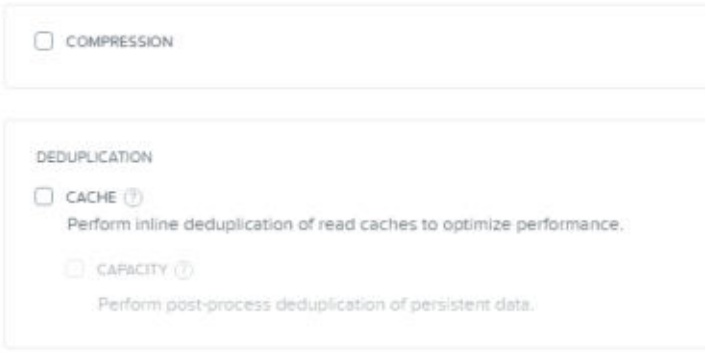

Figure 106: Create Storage Container Advanced Settings (2)

d. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is configurable. For post-process compression, where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

   All data in the storage container is compressed when you select **Compression**. See Compression on page 126 for information about using compression.

e. **Deduplication**: Select the **CACHE** check box to perform inline deduplication of read caches to optimize performance. If you enable this option, the Controller VMs must be configured to have at least 24 GB of RAM. This feature is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. Turning deduplication on for VAAI clone or linked clone environments is not recommended.

   Select the **CAPACITY** check box to perform post-process deduplication of persistent data. This option is recommended primarily for full clone, persistent desktops, and physical to virtual migration use cases that need storage capacity savings (not just performance

savings from deduplication). It is further recommended that the Controller VMs have at least 32GB of RAM and 300GB SSDs for the metadata disk to use this option.

Figure 107: Create Storage Container Advanced Settings (3)

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 128.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form *ip_address/ netmask*).

   A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

   > Note: Setting a storage container level whitelist overrides any global whitelist for this storage container.

4. When all the field entries are correct, click the **Save** button.

Modifying a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be modified as conditions change.

About this task

Storage Containers can be modified to change how the data in that storage container is handled, for example to apply compression.

> Note: The web console does not allow you to rename a storage container in an AHV cluster when modifying container details through the Update Storage Container dialog box. You cannot rename a storage container if it contains vdisks.

To modify (or delete) a storage container, do the following:

Procedure

1. In the web console select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Container** tabs.

2. To update the storage container, select the target storage container and then click the **Update** link.
   The Update Storage Container window appears, which includes the same fields as the Create Storage Container window (see Creating a Storage Container on page 157). Through this

window you can specify NFS datastore mounts, reserve storage capacity, enable (or disable) compression, enable or disable erasure coding, select or deselect fingerprinting on writes which enables deduplication on reads, and configure filesystem whitelist.

Note:

- If the compression policy is changed from compressed to uncompressed (or vice versa), the existing compressed (uncompressed) data in the storage container will be uncompressed (compressed) as a background process when the next data scan detects the data that needs this change.
- The Prism web console does not provide an option to change the container replication factor. That can be done only through the nCLI (see Increasing the Cluster Fault Tolerance Level on page 102).

3. To delete the storage container, select the target storage container and then click the **Delete** link.

Creating a Storage Container with Replication Factor 1

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers on page 753) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

For replication factor 1 (RF1) storage container recommendations and limitations, see Replication Factor 1 Recommendations and Limitations on page 107. To create a non-RF1 storage container, see Creating a Storage Container on page 157.

About this task

AOS automatically creates the correct type of access to the RF1-enabled storage container for each hypervisor.

- vSphere: The storage container is accessible as an NFS datastore. This requires access to the vSphere APIs. Ensure that you have appropriate vSphere license to access the APIs.
- AHV: The storage container is accessible transparently.

Note: The NutanixManagementShare storage container is a built-in storage container for Nutanix clusters for use with the Nutanix Files and Self-Service Portal (SSP) features. This storage container is used by Nutanix Files and SSP for file storage, feature upgrades, and other feature operations. To ensure proper operation of these features, do not delete this storage container. Nutanix also recommends that you do not delete this storage container even if you are not using these features. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.

Procedure

1. In the Storage dashboard (see Storage Dashboard on page 133), click the **Storage Container** button.
   The Create Storage Container dialog box appears.

2. **Name**: Enter a name for the storage container.
   The storage container has naming restrictions depending on the cluster hypervisor.

   AHV:

   - Container Name Length: Maximum length is 75 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   - Case Sensitivity: Container names are case sensitive.

   ESXi:

   - Container Name Length: Maximum length is 42 characters.
   - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   - Case Sensitivity: Container names are case sensitive.

3. **Storage Pool**: Select a storage pool from the drop-down list. **Max Capacity** displays the amount of free space available in the selected storage pool.

   a. (vSphere only) **NFS Datastore**: Select the **Mount on the following ESXi hosts** button to mount the storage container on a single host. From the list of host IP addresses, check the box for one host only.
      For clusters running the ESXi hypervisor, make sure you mount the RF1 storage container on a single host only. Mounting the RF1 storage container on all or multiple hosts is not supported.

4. Click the **Advanced Settings** button and configure the storage container.

   a. **Replication Factor**: Select **1**
      Nutanix supports a replication factor (RF) of 1 if you first select the **Enable Replication Factor 1** checkbox at the **Settings** > **Redundancy State** page. If you do not select **Enable Replication Factor 1**, Nutanix supports a replication factor of 2 or 3.

   b. From the **Node** drop-down menu, select a node to associated with the storage container.
      You can create one RF1 storage container per node, so you can only select one node here.

   c. **Reserved Capacity**: Leave this setting at 0.

   d. **Advertised Capacity**: Sets a maximum storage space for this storage container. Enter the amount (in GiB) to reserve in this field.
      This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

   e. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is configurable. For post-process compression,

where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

All data in the storage container is compressed when you select **Compression**. See Compression on page 126 for information about using compression.

f. **Deduplication** and **Erasure Coding** are greyed out and unavailable when you select replication factor 1.

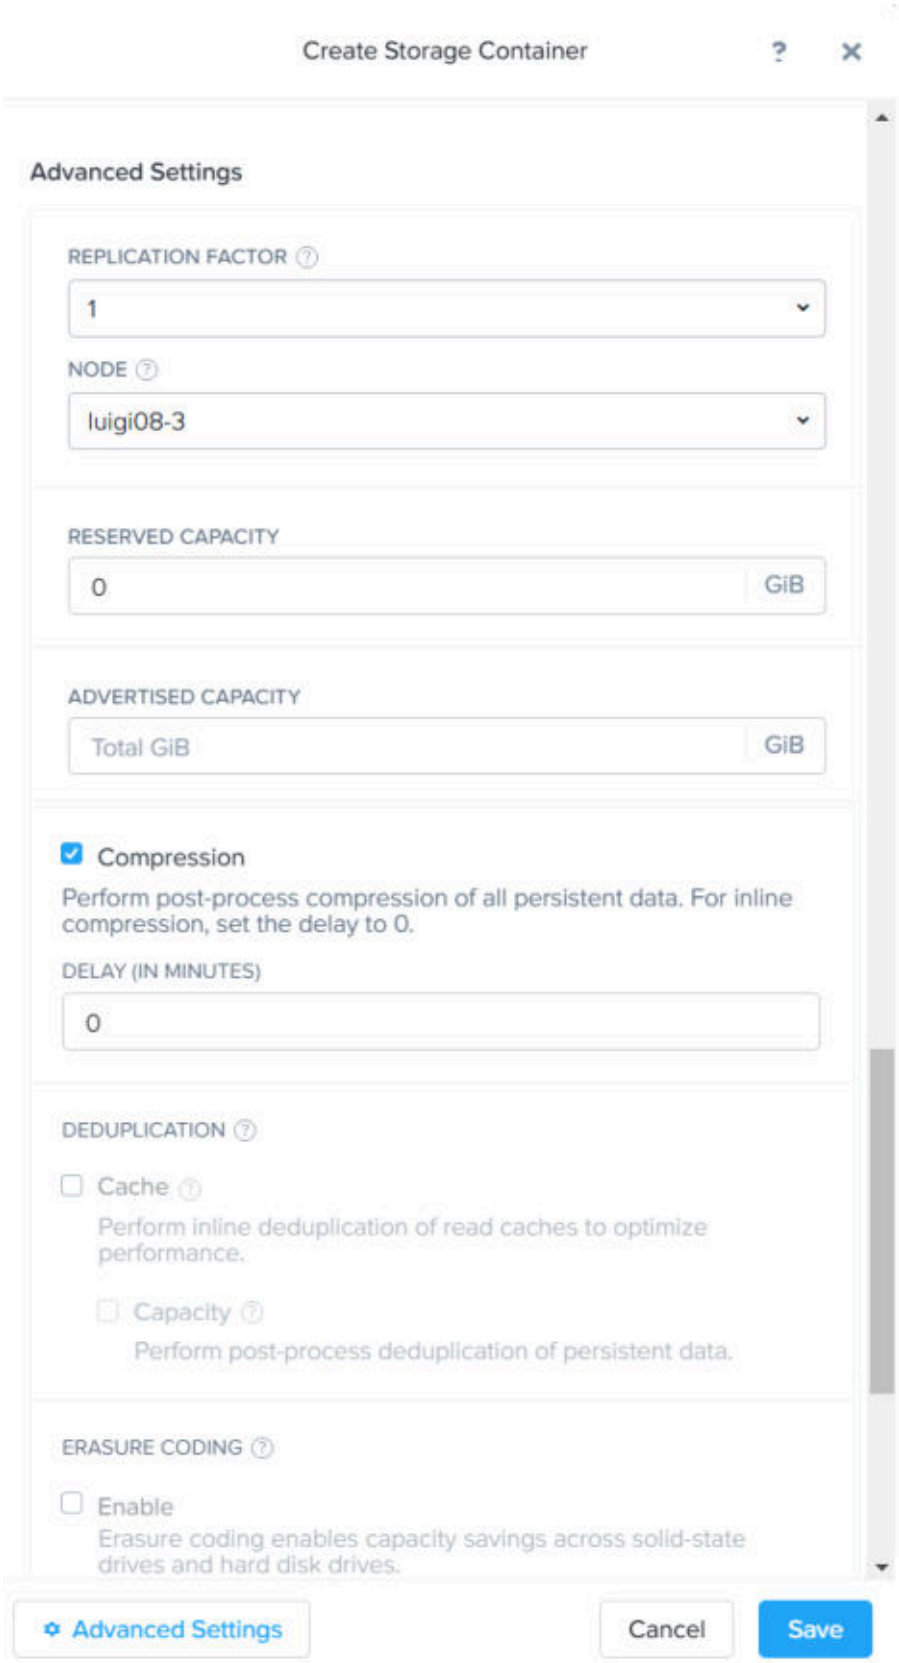

Figure 108: Create Storage Container Advanced Settings

5. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form *ip_address/ netmask*).

   A whitelist, also known as an allowlist, is a set of addresses that are allowed access to this storage container. Allowlists are used to allow appropriate traffic when unauthorized access from other sources is denied.

   > Note: Setting a storage container level allowlist overrides any global allowlist for this storage container.

6. When all the field entries are correct, click **Save**.

Configuring a Warning Threshold for Resilient Capacity

Resilient capacity is the storage capacity available in the cluster after accounting for the storage space needed to rebuild and restore in case of any component failures. The resilient capacity in a cluster depends on the cluster fault tolerance level and the configured cluster failure domain (node, block, or rack).

About this task

You can configure a warning threshold to track the resilient capacity in the cluster. When the used capacity crosses the set threshold, the storage summary widget changes color as a warning.

To configure a warning threshold for resilient capacity, do the following.

Procedure

1. In the web console, select **Storage** from the pull-down main menu and then select **Overview** to display storage information in a summary view.
2. Click the gear icon in the **Storage Summary** widget.
   The Configure Warning Threshold window appears.
3. Select one of the following:
   - **Use default**. Select this option to use the default 75% warning threshold limit for resilient capacity.
   - **Set manually**. Select this option and enter your custom warning threshold limit for resilient capacity in the cluster.
4. Click **Save**.
   A resilient capacity marker is set according to the warning threshold limit on the storage capacity chart, hover over the chart to see more details.

Rebuild Capacity Reservation

Nutanix recommends that you reserve the Rebuild Capacity in the clusters. Reserving rebuild capacity in a cluster enhances the cluster's self-healing mechanism.

Prism web console helps you reserve storage capacity within clusters for use in rebuilding failed nodes, blocks or racks (depending on what is configured as cluster failure domain). The AOS calculates and reserves some capacity to enable self-healing and rebuilding of failed cluster failure domain. This reserved capacity is called Rebuild Capacity.

When you reserve rebuild capacity, the system reserves the necessary capacity within the total cluster storage capacity based on Failure Domain and Fault Tolerance. When there is a failure

of the Failure Domain entity such as failure of a node, then the cluster rebuilds the node's data by replicating the data on the remaining nodes.

For more information about Failure Domain and Fault Tolerance, see Availability Domains on page 14.

Resilient Capacity

Storage capacity available in the cluster after accounting for the storage space needed to rebuild and restore in case of any component failures.

Rebuild Capacity

Storage capacity that can be reserved in the cluster to self-heal from failures. AOS calculates this capacity and reserves it when you enable the **Reserve Rebuild Capacity** option in Prism web console Settings. The cluster will stop accepting Write requests when the Resilient Capacity limit is reached.

Requirements and Limitations

You can reserve rebuild capacity subject to the following limitations:

- Nutanix supports rebuild capacity reservation on the cluster with three or more node. Nutanix does not support rebuild capacity reservation on single-node and two-node clusters.
- It is advisable for the Prism administrator to confirm if a manual container reservation was previously configured on the cluster before configuring a rebuild capacity reservation.
- The cluster must have a single storage pool that is the default storage pool. Do not create storage pools in the cluster when you have reserved rebuild capacity.
- For more information about storage pools, see Storage Components on page 124.
- Total capacity used or consumed must be less than the resilient capacity.

  If the used capacity is close to the resilient capacity and it increases due to large write operations and/or internal background jobs and migration tasks, then the used capacity can overshoot the resilient capacity. If rebuild capacity is already reserved on the cluster, the cluster stops accepting write requests.

  Total Capacity in the cluster is the sum of resilient capacity and reserved rebuild capacity. Capacity usage as the percentage of resilient capacity (Used Capacity) changes the color of the bar displaying the used capacity. Capacity usage also generates alerts. For example, if usage is 95 percent or more, then the cluster generates a critical alert after a specified number of NCC check iterations. When you have reserved rebuild capacity and the usage is 95 percent or more (of the resilient capacity), the cluster stops accepting Write requests.

  > CAUTION: Enabling Rebuild Capacity Reservation on a cluster with current usage (Total Used Space) close to the Resilient Capacity threshold might result in a VM outage. To avoid the issue, before enabling reservation ensure that current usage (Total Used Space) is not more than 90% of the calculated Resilient Capacity threshold.

- When you enable **Reserve Rebuild Capacity**, the usable capacity reduces from the Total Capacity to resilient capacity. Therefore, Recycle Bin usage as a percentage of usable capacity also changes to being a percentage of the resilient capacity after reserving Rebuild Capacity. See Recycle Bin on page 179 for information.

  For example, if Recycle Bin usage as a percentage of usable capacity is four percent (4%) before enabling Reserve Rebuild Capacity, it could change to, for example, 10% of resilient capacity after enabling **Reserve Rebuild Capacity**. When Reserve Rebuild Capacity is enabled, Recycle Bin usage as a percentage of resilient capacity is used as the threshold to auto-disable Recycle Bin. This threshold percentage is five percent of usable (whether

Total or resilient) capacity. In the above example, when Curator identifies that the threshold percentage is exceeded during a full scan, the Recycle Bin is auto-disabled.

> Note: Auto-disabling of Recycle Bin depends on full scans by Curator. Therefore, the excess usage condition of the Recycle Bin could continue for sometime between two full scans before Recycle Bin is disabled.

- Changes in the cluster storage capacity impacts Resilient, Rebuild, Used and Free capacities in the cluster.
  - When rebuild capacity is reserved and you try to remove a host or a disk, the cluster calculates the possible used capacity and resilient capacity after the removal. A confirmation dialog box appears displaying a message that the Resilient capacity will be reduced to the possible reduced capacity after the removal. If the used capacity is greater than the possible resilient capacity, then the host or disk removal fails.

    Thus when rebuild capacity is reserved, Prism allows a node or disk removal only if the cluster's data can rebuild after such removal and at the same time preserve the configured domain's fault tolerance.
  - When rebuild capacity reservation is enabled, the data consumption of failed nodes is not accounted in the total usage. When rebuild capacity reservation is not enabled, the data consumption of failed nodes is added to the total usage thereby inflating it.
  - When rebuild capacity reservation is enabled, Oplog usage is accounted in the total usage. The oplog is a fast write back cache to absorb random writes. Its consumption will also be accounted towards total usage when rebuild capacity reservation is enabled.
- You cannot reserve rebuild capacity if you have enabled redundancy factor (RF) 1. You cannot enable RF 1 if you have reserved rebuild capacity.

  When you reserve rebuild capacity in a cluster, do not enable RF1 using CLI or API.
- If you want to change the failure domain, for example, from **Node** to **Block** you must disable Rebuild Capacity Reservation. You can enable Rebuild Capacity Reservation after changing the failure domain. You cannot change the failure domain when Rebuild Capacity Reservation is enabled.

  Do not modify the failure domain using APIs when the Rebuild Capacity Reservation is enabled.
- Ensure that features such as Erasure Coding, De-duplication and compression are not disabled after you enable rebuild capacity reservation. Disabling such features can lead to drastic increase in used capacity beyond 95 percent and put the cluster in read-only mode when the usage before turning off these features is close to the threshold.
- The rebuild process is completed only when fault tolerance is not exceeded. For example, if failure domain is **Node** and fault tolerance is **2**, the rebuild process that starts after one node failure completes successfully. The rebuild process that starts after a second node failure (concurrent) also completes successfully. However, at this stage, if a third node fails, the rebuild process starts but does not complete.
- The rebuild process does not complete if data replicas are unavailable due to reasons such as link failures or disk failures.

> Note:
>
> Usage due to internal background jobs and migration tasks can increase the used capacity even when there are no write operations running. Heavy write operations with small containers also drastically increases the used capacity. When the used capacity

is close to resilient capacity, increase the capacity of the hosts in the cluster to ensure that the usage does not meet or exceed 95 percent of the total capacity.

Rebuild Capacity Display

When rebuild capacity is not reserved, the Overview tab of **Storage** page displays resilient capacity Warning Threshold in Storage Summary.

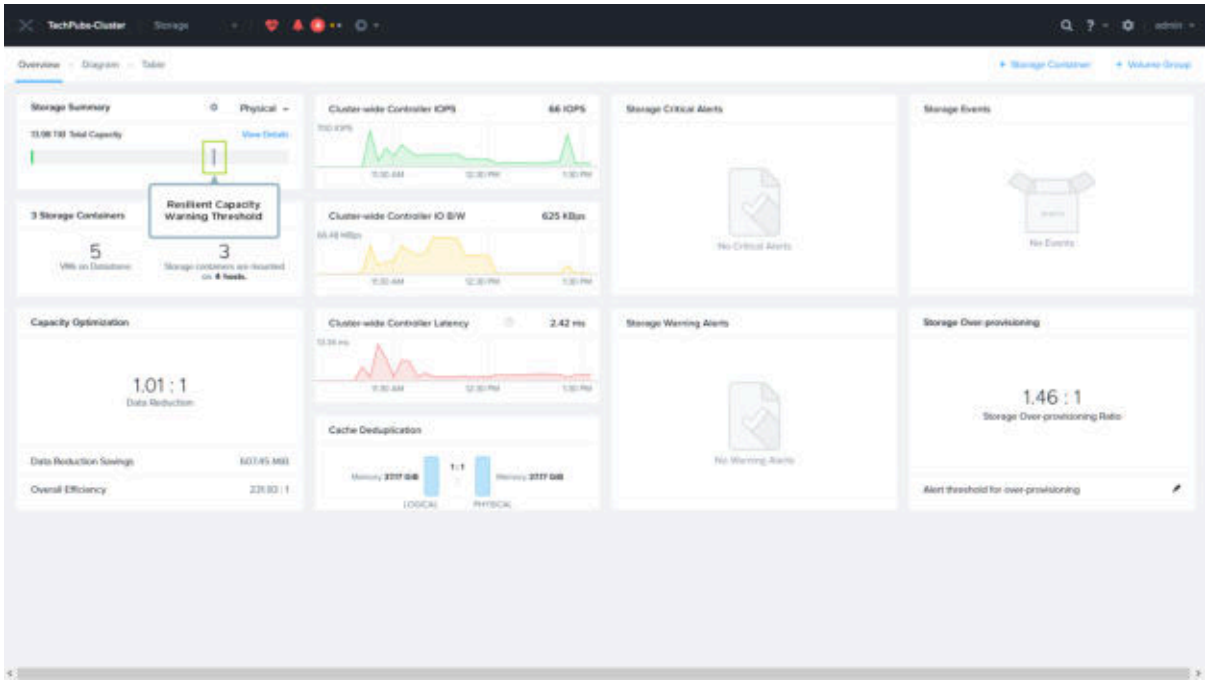

Figure 109: Storage Summary - Rebuild Capacity not reserved

See Configuring a Warning Threshold for Resilient Capacity on page 167 for more information about Resilient Capacity and Warning Threshold configuration.

Click the **View Details** link to open the Storage Details page, after selecting Physical in the drop-down on the right. The Storage Details page displays a banner informing you about rebuild capacity reservation. It also provides the **Enable Now** link that opens the **Rebuild Capacity Reservation** where you can reserve rebuild capacity.

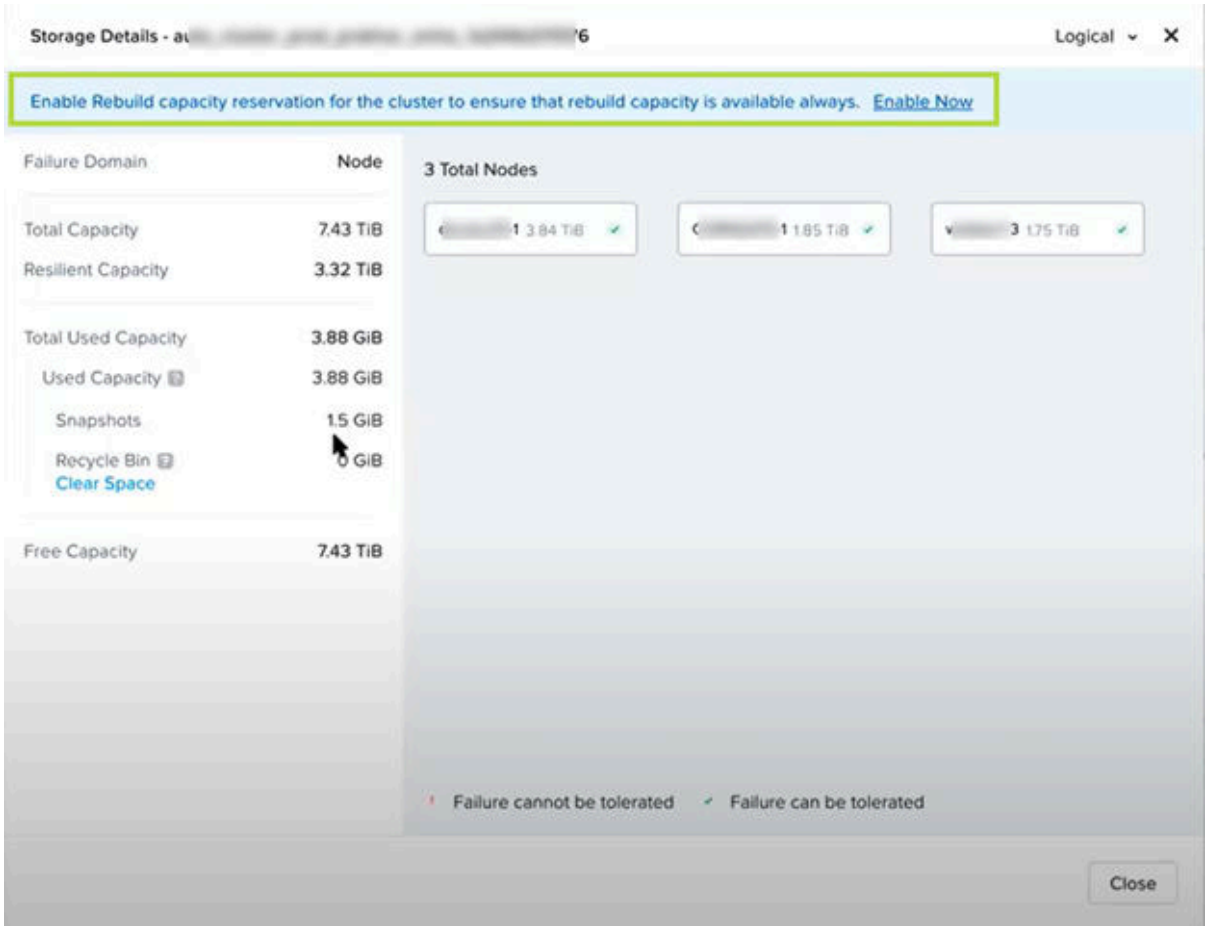

Figure 110: Storage Details - before Reserving Rebuild Capacity

To enable rebuild capacity reservation, see Reserving Rebuild Capacity on page 173.

When you enable rebuild capacity reservation, the cluster calculates the required rebuild capacity based on parameters such as fault tolerance, failure domain, and total storage capacity in the cluster. The reserved rebuild capacity is displayed on the **Rebuild Capacity Reservation** page.

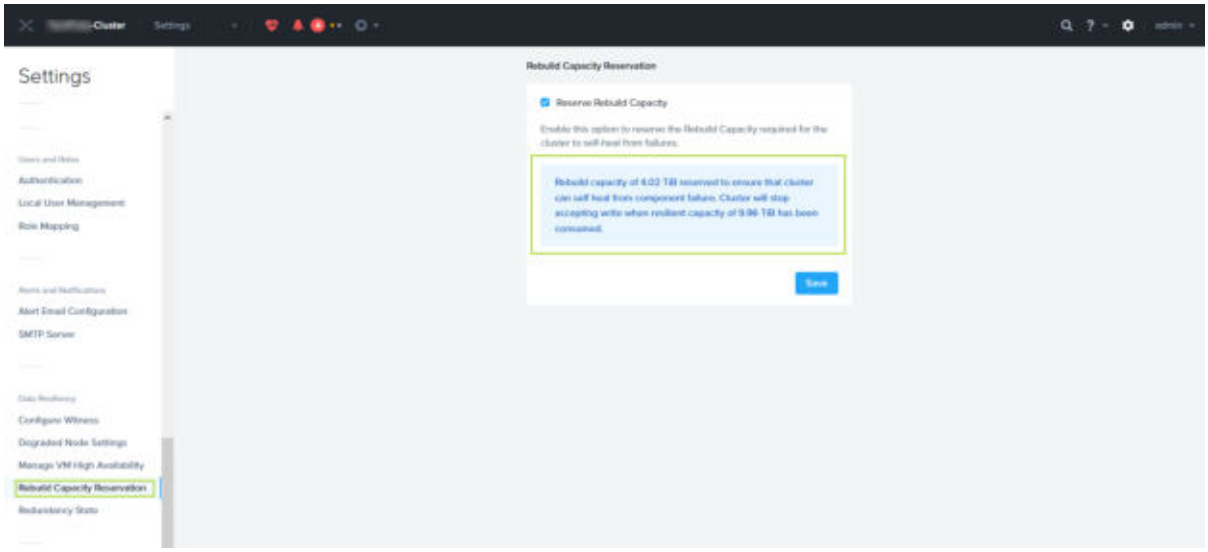

Figure 111: Rebuild Capacity Reserved

After you reserve rebuild capacity, the Overview tab of **Storage** page does not display resilient capacity Warning Threshold in Storage Summary.

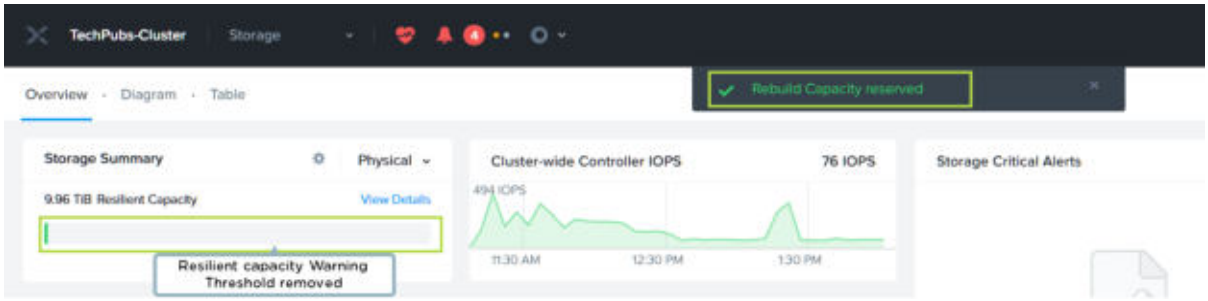

Figure 112: Storage Summary - Rebuild Capacity Reserved

After you reserve rebuild capacity, the Storage Details page displays a banner informing you that the cluster has reserved rebuild capacity. The details show an additional item - Rebuild Capacity with the capacity reserved in TiB. The capacity numbers change in the Storage Details page after the reservation.

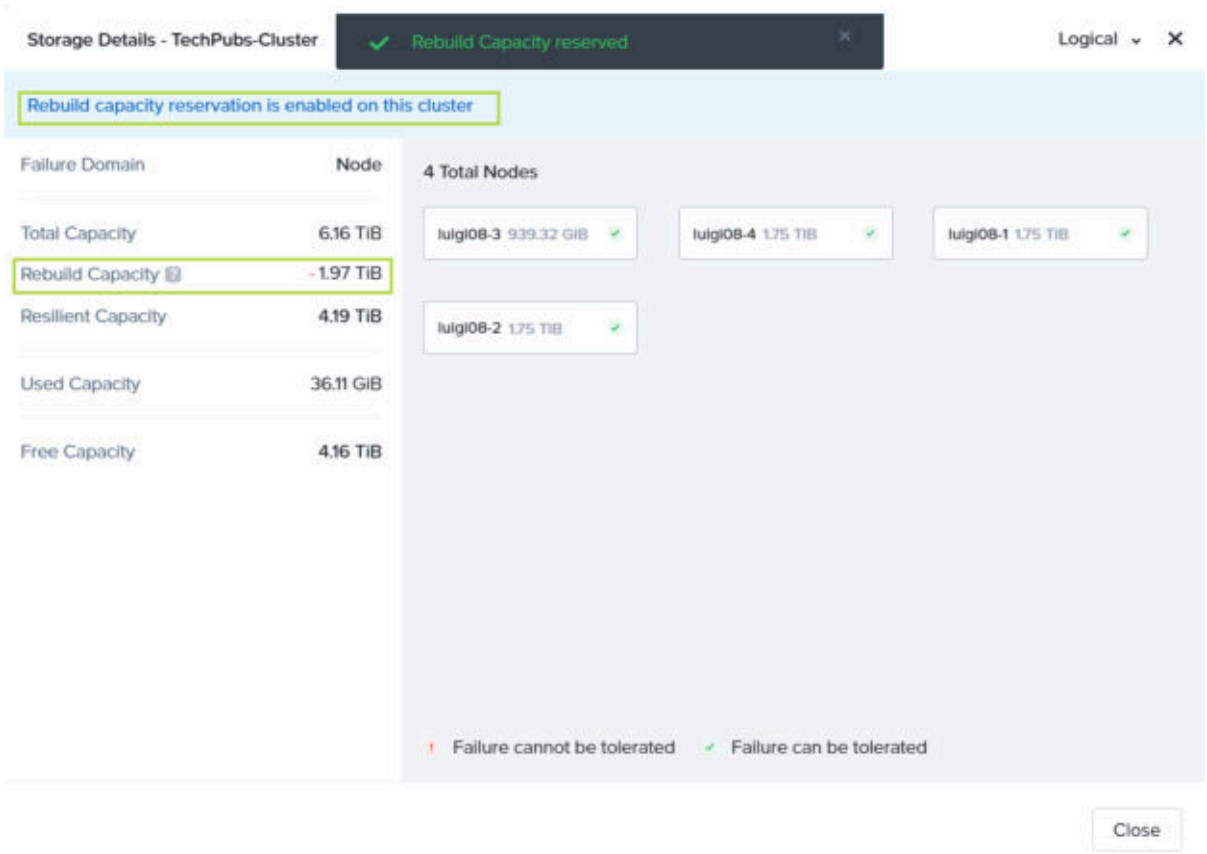

Figure 113: Storage Details - before Reserving Rebuild Capacity

Reserving Rebuild Capacity

This procedure helps you reserve rebuild capacity in the cluster.

About this task

To reserve rebuild capacity, open the **Rebuild Capacity Reservation** page. You can open the page by:

- Clicking the **Enable Now** link in the banner on theStorage Details view.

  Or

- Navigating to it in **Settings** > **Data Resiliency**.

Do the following in the **Rebuild Capacity Reservation** page.

Procedure

- Check the **Reserve Rebuild Capacity** check box.

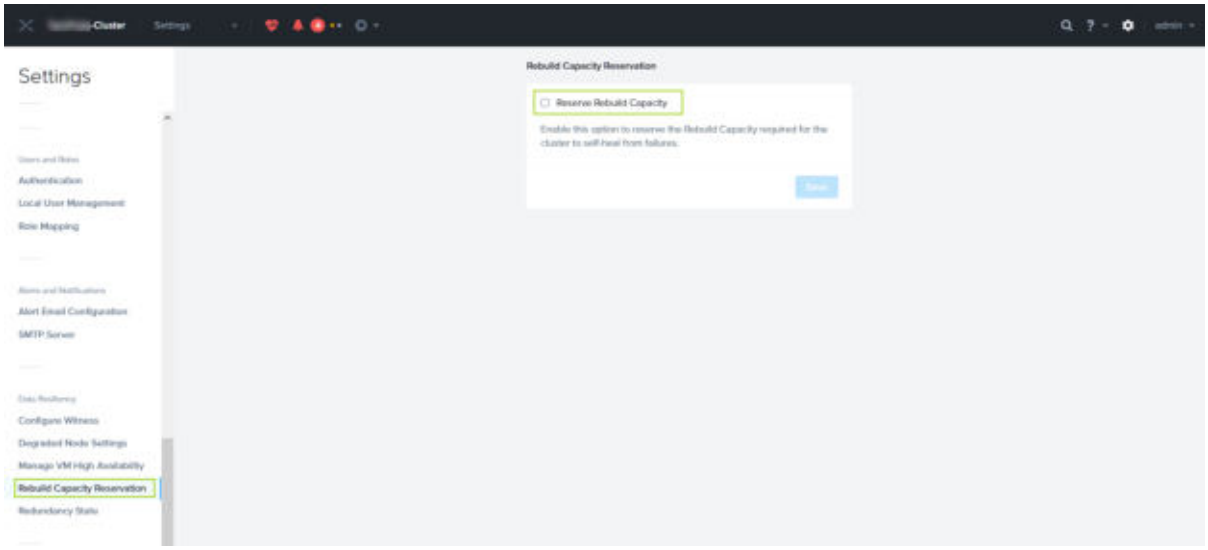

Figure 114: Enable Rebuild Capacity Reservation

Volume Group Configuration

A volume group is a collection of virtual disks (vDisks). Each volume group is identified by a name and UUID. Each vDisk has a UUID and a SCSI index to specify the ordering of disks within the volume group.

You create a volume group on a storage container and enable access to the volume group in one of the following ways:

iSCSI

iSCSI access to a volume group is provided through Nutanix Volumes. See the Nutanix Volumes Guide for more information.

Attach Directly to a VM

If the hypervisor is AHV, you can create a volume group and attach it to one or more VMs as a SCSI disk by using the Prism Element web console, the Nutanix REST API, or the nCLI or aCLI commands.

Once the volume group is attached to the VM by using the Prism Element web console, vDisks appear as SCSI devices to the guest operating system. AOS transparently handles load balancing and path resiliency on the guest VM. Directly attached volume groups support Controller VM failover and do not impact VM migration.

You can attach a volume group to a VM at any time. This option is available only to VMs deployed on AHV.

This method of attaching a volume group to a VM has the following limitations:

- The VM must run on the same cluster as the volume group.

Volume groups are managed independently of the VMs to which they are attached. You can configure a directly attached volume group for access by a single VM (exclusive access) or by multiple VMs (shared access).

CAUTION: Multiple VMs writing to a vDisk that belongs to a shared volume group without additional software to manage the access can lead to data corruption. Use shared access when you are configuring VMs for use with cluster aware software.

Concurrent Access from Multiple Clients

Products, features, or solutions might require concurrent access to volume groups, either by multiple iSCSI initiators or multiple VM attachments.

Multiple iSCSI Initiators

The following products, features, or solutions are supported for concurrent access to volume groups:

- Oracle RAC (bare-metal and virtualized environments)
- Linux VMs
- Windows Failover Clustering
- IBM Spectrum Scale (GPFS)
- Veritas InfoScale Storage

Multiple Directly Attached VMs

The following products, features, or solutions are supported for concurrent access to volume groups:

- Oracle RAC
- Linux VMs
- IBM Spectrum Scale (GPFS)
- Veritas InfoScale Storage
- Windows Failover Clustering
- Linux guest VM clustering (also known as RHEL HA)

Creating a Volume Group

About this task

To create a volume group, do the following:

Procedure

1. In the **Storage** dashboard, click the **Volume Group** button.
   The Create Volume Group dialog box is displayed.
2. In **Name**, enter a name for the volume group.

   Note: This entity has the following naming restrictions:

   - The maximum length is 75 characters.
   - Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), decimal digits (0-9), dots (.), and hyphens (-).

3. The **iSCSI Target Name Prefix** is auto-filled with the volume group **Name**. You can accept this prefix or enter your own target name prefix for the volume group. This entity has the same naming restrictions as **Name**.
4. In **Description**, enter a description for the volume group.
5. To add a disk to the volume group, do the following:
   a. In the **Storage** section, click **Add New Disk**.
   b. In the **Add Disk** dialog box, in **Storage Container**, select the storage container to use from the pull-down list.
   c. In **Size**, enter the disk size in GiBs.
   d. Click **Add**.
   e. Repeat these steps to add another disk for this volume group, if desired.
6. Do the following in the **Access Settings** section:
   a. Select **Enable external client access** if you are whitelisting clients that are external to or not residing in this cluster. Otherwise, leave this clear. If you select this check box, it remains selected the next time you create a volume group.
   b. If you are using one-way CHAP security, select **CHAP Authentication** and type a 12-character to 16-character password (also known as a CHAP secret) in the **Target Password** field.
      Initiators must use the same password to authenticate to the AOS cluster.
   c. Click **Add New Client** to configure the iSCSI initiators, and then enter the client Initiator iSCSI Qualified Name (IQN) in the **Client IQN/IP Address** field to create the whitelist. If you have configured Mutual CHAP authentication on the client, select **CHAP Authentication** and enter the iSCSI client password (secret). Click **Add**.

      > Note: Ensure that you enter the client IQN in the **Client IQN/IP Address** field and not the IP address. AOS does not support a whitelist containing IP addresses in a volume group.

      Access Control displays any configured clients. This list includes any clients attached to volume groups in the cluster. Repeat this step to add more initiators allowed to access this storage. For information about which products, features, or solutions are supported for concurrent access to a volume group, see Concurrent Access from Multiple Clients on page 175.
7. To enable the Flash Mode feature, select **Enable Flash Mode**.

   > Note: Individual virtual disks of a volume group cannot be excluded from Flash Mode by using the Prism Element web console. However, you can exclude individual virtual disks from flash mode by using aCLI. For more information, see Removing Flash Mode for Virtual Disks of a Volume Group on page 179.
8. Click **Save**.

What to do next

If the hypervisor is AHV, you can now attach the volume group to the VM and start using the vDisks. If you want to use iSCSI and you already whitelisted the host IP addresses, log on to the VMs, and configure iSCSI.

Modifying a Volume Group

About this task

To modify a volume group, do the following:

Procedure

1. In the web console, select **Storage** from the pull-down main menu (upper left of screen), and then select the **Table** and **Volume Group** tabs.
2. To update a volume group, select the volume group, and then click the **Update** link.
   The Update Volume Group window appears, which includes the same fields as the Create Volume Group window (see Creating a Volume Group on page 175). In this window you can change the volume group name, add and remove disks, configure the volume group for sharing, and add or remove (by unchecking) entries from the initiator whitelist. On AHV clusters, you can attach the volume group to a VM as a SCSI disk (described later in this procedure). If you attach a volume group to a VM that is part of a protection domain, the volume group is not protected automatically. Add the volume group to the same Consistency Group manually.

   You can increase the size of the volume group. Reducing the size of a volume group is not supported.

   Any unchecked initiator that remains unused by this or any other volume group ultimately disappears from this list.
3. To attach a volume group to a VM, click **Attach to a VM**, and then select the VM from the **Available VMs** list.
4. To delete a volume group, select the volume group, and then click the **Delete** link.
5. To manage the iSCSI client list, click the **Summary** link to deselect any volume groups, and then click the **Manage Initiators** link. (This link appears only when no volume groups are selected.)

   The Manage iSCSI Clients window appears, which includes a list of the clients and available actions. To modify a client (enable/disable CHAP authentication), click the pencil icon for that client, which displays the Edit iSCSI client window. See the Nutanix Volumes Guide for more information.

Flash Mode for Virtual Machines, Disks, and Volume Groups

Flash mode for VM allows you to set the storage tier preference to SSD for a virtual machine or volume group. Flash mode is not enabled by default for all database deployments. But, flash mode allows latency sensitive workloads to remain on the SSD tier and avoid any potential performance impact of down migration of data.

By default, you can use up to 25% of the cluster-wide SSD tier as flash mode space for VMs or VGs. If the data size for flash mode enabled VMs or VGs exceeds 25% of the SSD capacity, the system may down migrate the data. Before performing down-migration, the flash mode feature tries to preserve the excess data on the SSD tier for some reasonable amount of time so that you can take corrective actions on the cluster and bring back to stable state. To reduce flash mode usage, you can disable flash mode on some VMs or VGs or add SSDs.

If you enable this feature on a VM, all the virtual disks that are attached to the VM are automatically placed on the SSD tier. Also, virtual disks added to this VM are automatically placed on SSD. However, you can update the VM configuration to remove the flash mode from any virtual disks.

CAUTION: While enabling the flash mode feature for a VM may increase the performance of that VM, it may also lower the performance of VMs that do not have flash mode feature enabled. Nutanix recommends you to consider the performance impact on other VMs and VGs. To mitigate any impact on the performance, you can update the VM configuration and remove the flash mode on individual virtual disks. For example, you can enable flash mode on the applications data disks and disable it on the log disks.

A node failure may cause a reduction in SSD tier capacity that can drive the flash mode usage above the 25% of the tier capacity. In this event, alerts will report the flash usage exceeding 25%.

Note:

- This feature is supported on ESXi and AHV for VMs and on all the hypervisors for VGs.
- For the cluster created using ESXi hosts, you must register your cluster with the vCenter Server. For more information, see Registering a Cluster to vCenter Server on page 771.

You can enable the feature on the VM, and VM disks only during the VM update workflow from the Prism Element. However, for VGs, you can enable the feature during the creation of VGs. For more information on enabling the feature on VMs, see VM Management on page 316 and for enabling the feature on VGs, see Creating a Volume Group on page 175.

Alerts are raised in the following scenarios.

- When the flash usage exceeds 25% of the SSD tier.
- When the VM has the flash mode feature enabled but is in the powered off state.

Guidelines for the Flash Mode feature

- If a VM or a VG with flash mode feature enabled is cloned, then the flash mode policies are not automatically applied to the cloned VM or VG.
- The flash mode policies are not automatically applied if a VM or VG is restored from the DR snapshot on the local or the remote site.
- The flash mode policies can be used for controlling down migration from the hot tier to the cold tier. The policies do not restrict the usage for the VM or VG on the hot tier. If space is available on the hot tier, more data from the VM or VG can reside on the hot tier.
- If you enable erasure coding on the cluster, the parity resides on the HDD tier.
- If you replicate a VM or VG to a remote site (by using Metro Availability or Async DR), the remote site does not have flash mode feature enabled automatically. You need to manually enable the flash mode feature again.
- Do not use flash mode feature with capacity tier deduplication. Enabling flash mode feature does not affect performance tier deduplication.
- If you perform a storage migration or storage vMotion, the new virtual disk created on the target datastore will not have the flash mode feature enabled. You need to manually enable the flash mode feature again.
- If you try to upgrade from an AOS version lower than 5.0 to version 5.0 or higher with SSD pinning feature enabled, the upgrade process is blocked. Hence, ensure that you have disabled the SSD pinning feature before performing an upgrade. Also, do not enable the SSD pinning feature during the upgrade process.

- Flash mode is not supported with more than one storage pool in the cluster. If you delete and then create a storage pool in your cluster, you cannot configure the flash mode feature for a small duration of time. In this duration there are two storage pools in the cluster temporarily and the flash mode feature cannot be enabled.
- If you enable flash mode on all-NVMe clusters, it may result in the following error.

  ```
  Update failed: NfsConfigVDisk RPC returned kInvalidValue: kInvalidValue: 'SSD-SATA' supplied for tier params has no associated disk.
  ```

  This error can be safely ignored because the flash mode feature only pins data to the SSD tier and prevents down-migration to the HDD tier. Since all-NVMe clusters have neither SSD nor HDD tier, the feature is a no-op on such clusters.

Removing Flash Mode for Virtual Disks of a Volume Group

Individual virtual disks of a volume group cannot be excluded from flash mode by using Prism. However, you can exclude individual virtual disks from flash mode by using aCLI.

Procedure

1. Log in to the Controller VM in your cluster through an SSH session and access the Acropolis command line.
2. (Optional) If you have not enabled the flash mode feature for the VG by using Prism, you can enable it by running the following command.

   ```
   acli> vg.update vg_name flash_mode=true
   ```

   Replace *vg_name* with the name of the VG on which you want to enable the flash mode.
3. Remove the flash mode feature on a particular disk.

   ```
   acli> vg.disk_update vg_name index_value flash_mode=false
   ```

   Replace *vg_name* with the name of the VG and *index_value* with the index value of the VG. For example to disable flash mode for disk at index 0 with name example_vg, use the following command.

   ```
   acli> vg.disk_update example_vg 0 flash_mode=false
   ```

Recycle Bin

Enabled by default and with help from Nutanix Support, the Recycle Bin helps you manage and restore deleted storage entities like guest VMs and volume group vDisks.

> Note: To restore a storage entity, contact Nutanix Support.

> Note: When Replication Factor 1 is enabled for a storage container, the recycle bin is disabled for that storage container. When you delete a storage entity like a guest VM or volume group vDisk, it is marked for deletion as soon as possible and bypasses the recycle bin. See Replication Factor 1 Overview on page 107.

When the Recycle Bin is enabled, AOS creates a Recycle Bin associated with the storage container configured in your cluster. If you then delete a storage entity (like a guest VM or volume group vDisk), the Recycle Bin retains vDisk and configuration data files for up to 24 hours. See the *AOS Family Release Notes* for release information about the Recycle Bin.

If you enable the Recycle Bin and then delete a guest VM or volume group vDisk, it retains its contents (deleted vDisks and configuration information) for up to 24 hours, unless the cluster free storage space reaches critical thresholds. After 24 hours, AOS automatically deletes

these files. AOS deletes the files in less than 24 hours if your cluster is unable to maintain sufficient free disk space. AOS triggers free disk space alerts when your cluster reaches critical thresholds.

After you disable the Recycle Bin, AOS automatically deletes any entities in the Recycle Bin after 24 hours. Then, when you subsequently delete any storage entities, AOS marks them for deletion as soon as possible. They are not stored in the storage container Recycle Bin folder.

After deleting one or more storage entities, the **Free** storage space shown in the Prism Element web console might not update for 24 hours. To see the space used by the Recycle Bin in the Storage Summary, go to the web console **Storage** dashboard and select the **Diagram** or **Table** view.

For more information about Recycle Bin behavior, see Recycle Bin Limitations and Considerations on page 180.

Recycle Bin Limitations and Considerations

Limitations and guidelines to consider when using the Recycle Bin.

- The Recycle Bin stores vDisk and configuration data for up to 24 hours. After 24 hours, these files are deleted. The files are deleted in less than 24 hours if your cluster is unable to maintain sufficient free disk space.

  The Recycle Bin is temporarily disabled when the Recycle Bin is using more than five percent of cluster storage capacity. AOS triggers free disk space alerts when your cluster reaches critical thresholds. In this case, newly-deleted entities are marked them for deletion as soon as possible. They are not stored in the storage container Recycle Bin folder.

  When the Recycle Bin is cleared manually or automatically and the Recycle Bin is using two percent or less of cluster storage capacity, Recycle Bin is automatically re-enabled and available for use after all Curator service scans are completed. (Among other cluster tasks, the Curator service controls Recycle Bin and storage cleanup.)

- If the Recycle Bin contains more than 2000 files, deleted storage entities guest VMs and volume group vDisks bypass the Recycle Bin. AOS marks them for deletion as soon as possible. They are not stored in the storage container Recycle Bin folder.

- The web console displays the **Clear Recycle Bin** (available on the **Storage** dashboard in Summary view) or **Clear Space** options (available in **Storage Details**) only if the Recycle Bin contains deleted items. If these options are available but the space used by the Recycle Bin is 0, AOS has detected the items but not yet calculated the storage space used.

- After you empty the Recycle Bin with **Clear Recycle Bin** or **Clear Space** or AOS empties the Recycle Bin, these options are not displayed again until you delete storage entities like guest VMs and volume group vDisks.

- When Replication Factor 1 is enabled for a storage container, the recycle bin is disabled for that storage container. When you delete a storage entity like a guest VM or volume group vDisk, it is marked for deletion as soon as possible and bypasses the recycle bin. See Replication Factor 1 Overview on page 107

- Recycle Bin is not supported on storage containers where metro availability is enabled.

- Recycle Bin is not available for recovering protection domain snapshots.

- When you delete a VM with a thick provisioned vDisk reservation setting, the setting is removed. After recovering this VM, the reservation is not restored and defaults to thin provisioned.

- If you are using the X-Ray application, an automated testing framework and benchmarking application for enterprise-grade datacenters, you must disable the Recycle Bin on your AOS cluster, which is enabled by default.

Enabling and Disabling the Recycle Bin

The Recycle Bin is enabled by default. You can turn the Recycle Bin on or off through the Cluster Details window in the web console.

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.

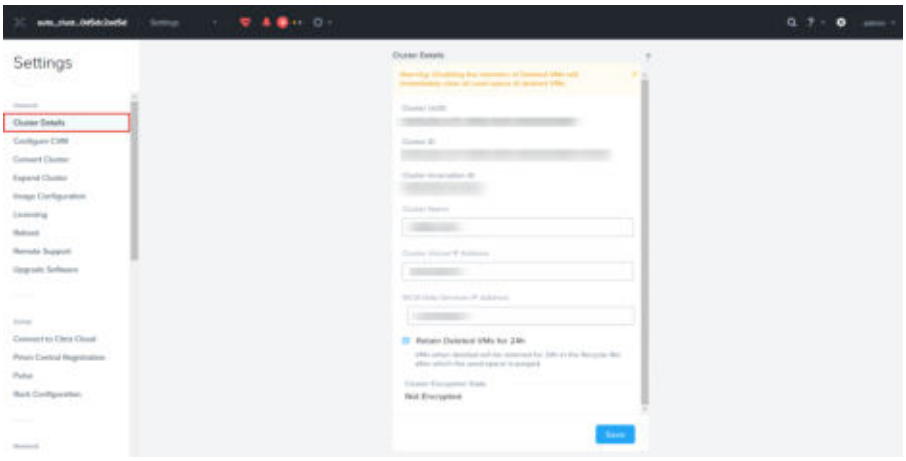

Figure 115: Cluster Details Window

2. To disable the Recycle Bin, clear (unselect) **Retain Deleted VMs for 24h**.

   After you disable the Recycle Bin, AOS automatically deletes any entities in the Recycle Bin after 24 hours. Then, when you subsequently delete any storage entities, AOS marks them for deletion as soon as possible. They are not stored in the storage container Recycle Bin folder. See also Recycle Bin on page 179.

3. To enable the Recycle Bin, select **Retain Deleted VMs for 24h**.

   If you enable the Recycle Bin and then delete a guest VM or volume group vDisk, it retains its contents (deleted vDisks and configuration information) for up to 24 hours, unless the cluster free storage space reaches critical thresholds.

4. Click **Save** to save the values and close the window.

Viewing Space Used by the Recycle Bin

The Prism Element web console shows the space used by the Recycle Bin in the Storage dashboard.

Before you begin

See also Storage Overview View on page 134, Storage Diagram View on page 139, and Storage Table View on page 146.

Procedure

1. In the web console, go to the **Storage** dashboard by clicking **Home** > **Storage**.

2. To view the space used by the Recycle Bin, click the **Diagram** or **Table** tab.
   The Storage Summary panel shows the space used by the Recycle Bin. If the **Clear Recycle Bin** option is available but the space used by the Recycle Bin is 0, AOS has detected the deleted items but not yet calculated the storage space used.

3. Also, from the **Home** dashboard, at the **Storage Summary** widget, click **View Details**.
   The Recycle Bin shows the space used. If the **Clear Space** option is available but the space used by the Recycle Bin is 0, AOS has detected deleted items but not yet calculated the storage space used.

Clearing Storage Space Used by the Recycle Bin

You can clear the storage space used by the Recycle Bin from the Home or Storage dashboards.

Before you begin

See also Storage Table View on page 146. Do these steps from the Prism Element web console.

- The web console displays the **Clear Recycle Bin** (available on the **Storage** dashboard in Summary view) or **Clear Space** options (available in **Storage Details** from the **Storage Summary** widget) only if the Recycle Bin contains deleted items. If these options are available but the space used by the Recycle Bin is 0, AOS has detected the items but not yet calculated the storage space used.
- After you empty the Recycle Bin with **Clear Recycle Bin** or **Clear Space** or AOS empties the Recycle Bin, these options are not displayed again until you delete storage entities like guest VMs and volume group vDisks.

Procedure

1. From the **Home** dashboard: at the **Storage Summary** widget, click **View Details**.
   a. Under Recycle Bin, click **Clear Space**.
   b. In the confirmation pop-up, click **Delete**.
      You cannot undo this action. Any entities in the Recycle Bin are marked for deletion immediately.
   c. Click **Close**.
2. From the **Storage** dashboard: Click the **Table** tab.
   a. To clear the space used by the Recycle Bin, Click **Clear Recycle Bin**.
   b. In the confirmation pop-up, click **Delete**.
      You cannot undo this action. Any entities in the Recycle Bin are marked for deletion immediately.

4

NETWORK MANAGEMENT

Nutanix provides several features to manage and monitor network settings for the cluster.

- To configure network connections in clusters through the web console with Nutanix virtualization management (such as those running AHV as the hypervisor), see Network Configuration for Cluster on page 183.
- To enable LAG and LACP on the T0R switch, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 206.
- To configure the network interfaces for a VM, see Network Configuration for VM Interfaces on page 194.
- To track and record networking statistics for a cluster, the cluster requires information about the first-hop network switches and the switch ports being used. You can configure one or more network switches for statistics collection (see Configuring Network Switch Information on page 202).
- A network visualizer is provided that presents a consolidated graphical representation of the network formed by the VMs and hosts in a Nutanix cluster and first-hop switches. You can use the visualizer to monitor the network and to obtain information that helps you troubleshoot network issues (see Network Visualization on page 210).

Network Configuration for Cluster

You can view and configure network connections for the cluster in the **Network Configuration**. To display the **Network Configuration** page, click the gear icon in the main menu and then select **Network Configuration** in the **Settings** page.

Network Connections

Each VM network interface is bound to a virtual network, and each virtual network is bound to a single VLAN. Information about the virtual networks configured currently appears in the Network Configuration page.

The Network Configuration page includes three tabs.

- The **Subnets** tab displays a list of the configured networks.

Figure 116: Networks Tab

- The **Internal Interfaces** tab displays a list of LAN interfaces.

Figure 117: Internal Interfaces Tab

- The **Virtual Switch** tab displays a list of virtual switches configured, including the default system-generated virtual switch vs0.

Figure 118: Virtual Switch Tab

The following table describes the fields in each tab.

Table 28: Network Configuration Fields

Parameter	Description	Values
Networks Tab		
Name	Displays the name of the virtual network.	(name)
Virtual Switch	Displays the name of the virtual switch in the form **vs#**, for example **vs0** for virtual switch 0 which is the default virtual switch.	(vs<number>)
VLAN ID	Displays the VLAN identification number for the network in the form **vlan.#**, for example **vlan.27** for virtual LAN number 27.	(ID number)
Internal Interfaces Tab		

Parameter	Description	Values
Descriptive Name	Displays a name for the LAN.	(LAN name)
Subnet (Gateway IP / Prefix Length)	Displays the subnet that the internal interface belongs to in the form *<IP Address>*/<number (prefix)>	(IP Address/prefix number)
Features	Displays the features available on the internal interfaces.	
Interface	Displays the interface designation such as eth0 or eth1.	(interface name)
Virtual Switch		
Name	Displays the name of the switch in the form vs#	(vs<number>)
Bridge	Displays the name of the bridge associated with the virtual switch in the form br#, for example br0 for the default bridge.	(br<number>)
MTU (bytes)	Displays the MTU set for the virtual switch in bytes. The default MTU is 1500.	(number)
Bond Type	Displays the uplink bond type associated with the virtual switch. See the *Bond Type* table in . For example, Active-Backup	(<bond_type>)

LAG and LACP on the ToR Switch

For more information, see Enabling LAG and LACP on the ToR Switch (AHV Only) on page 206.

Creating or Updating a Virtual Switch

This section provides the procedure to create or update a virtual switch.

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can create or update a VS using the Create Virtual Switch or Update Virtual Switch dialog box. The Create Virtual Switch and Update Virtual Switch dialog boxes have identical parameters and settings.

To open the Create Virtual Switch dialog box in Prism Web Console, do the following.

Procedure

1. Click **Settings** > **Network Configuration**.
2. On the Network Configuration page, click the **Virtual Switch** > **Create VS** link.

After you click the link, complete the following procedure.

About this task

To create a virtual switch using the Create Virtual Switch dialog box, do the following.

Procedure

1. In the Create Virtual Switch dialog box, provide the necessary information on the **General** tab.

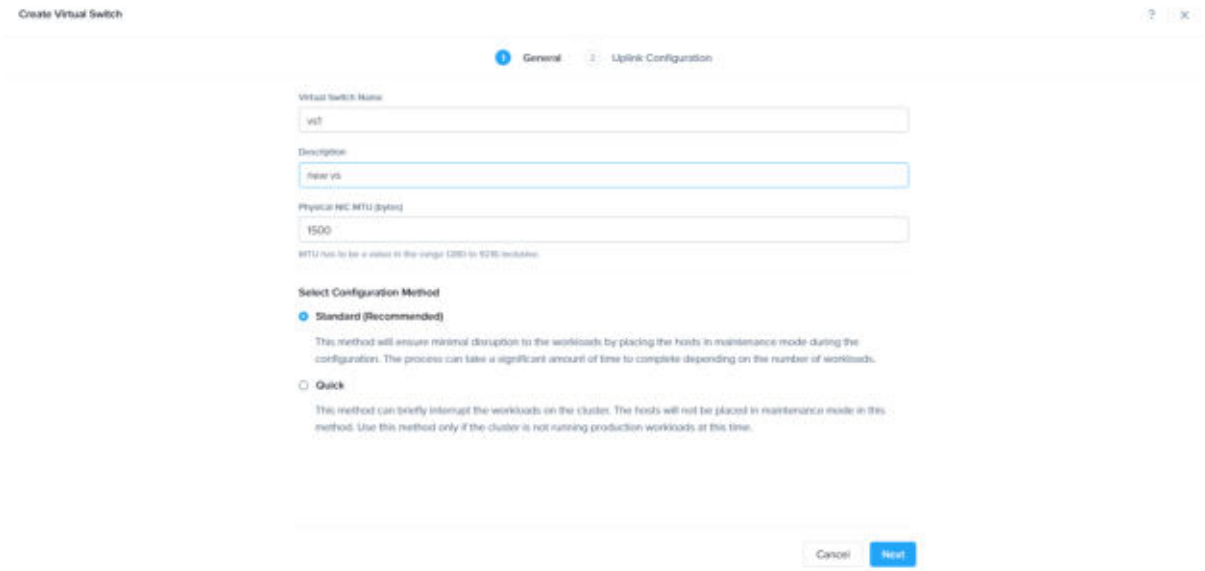

Figure 119: Create Virtual Switch - General tab

Field	Description
Virtual Switch Name	Enter a name for the virtual switch.
Description	Provide a description for the virtual switch that helps identify the virtual switch.
Physical NIC MTU (bytes)	MTU must be a value in the range 1280 to 9216 inclusive.

Field	Description
Select Configuration Method	Select one of the two methods that you can use to implement the VS configuration: • **Standard (Recommended)**: This method ensures no disruptions occur to the workloads by putting the hosts in maintenance mode and migrating the VMs out of the host before applying the configuration. This process requires a longer duration of time to complete. The time required depends on the number and configuration of VMs. In this method, the VS configuration is deployed in the rolling update process. • **Quick**: This method interrupts the workloads running on the hosts. Use this method only if you are not running production workloads because it may result in network interruptions. In this method, the VS configuration is deployed in a rolling update process but the nodes are not put in maintenance mode before modifying the VS configuration on the node. Note: After selecting the **Quick** configuration method, although a task named **Hypervisor rolling restart** appears in the task tray, the hosts are not actually rebooted.

2. Click **Next** to go to the Uplink Configuration tab.
   On the Uplink Configuration tab, provide the following details:

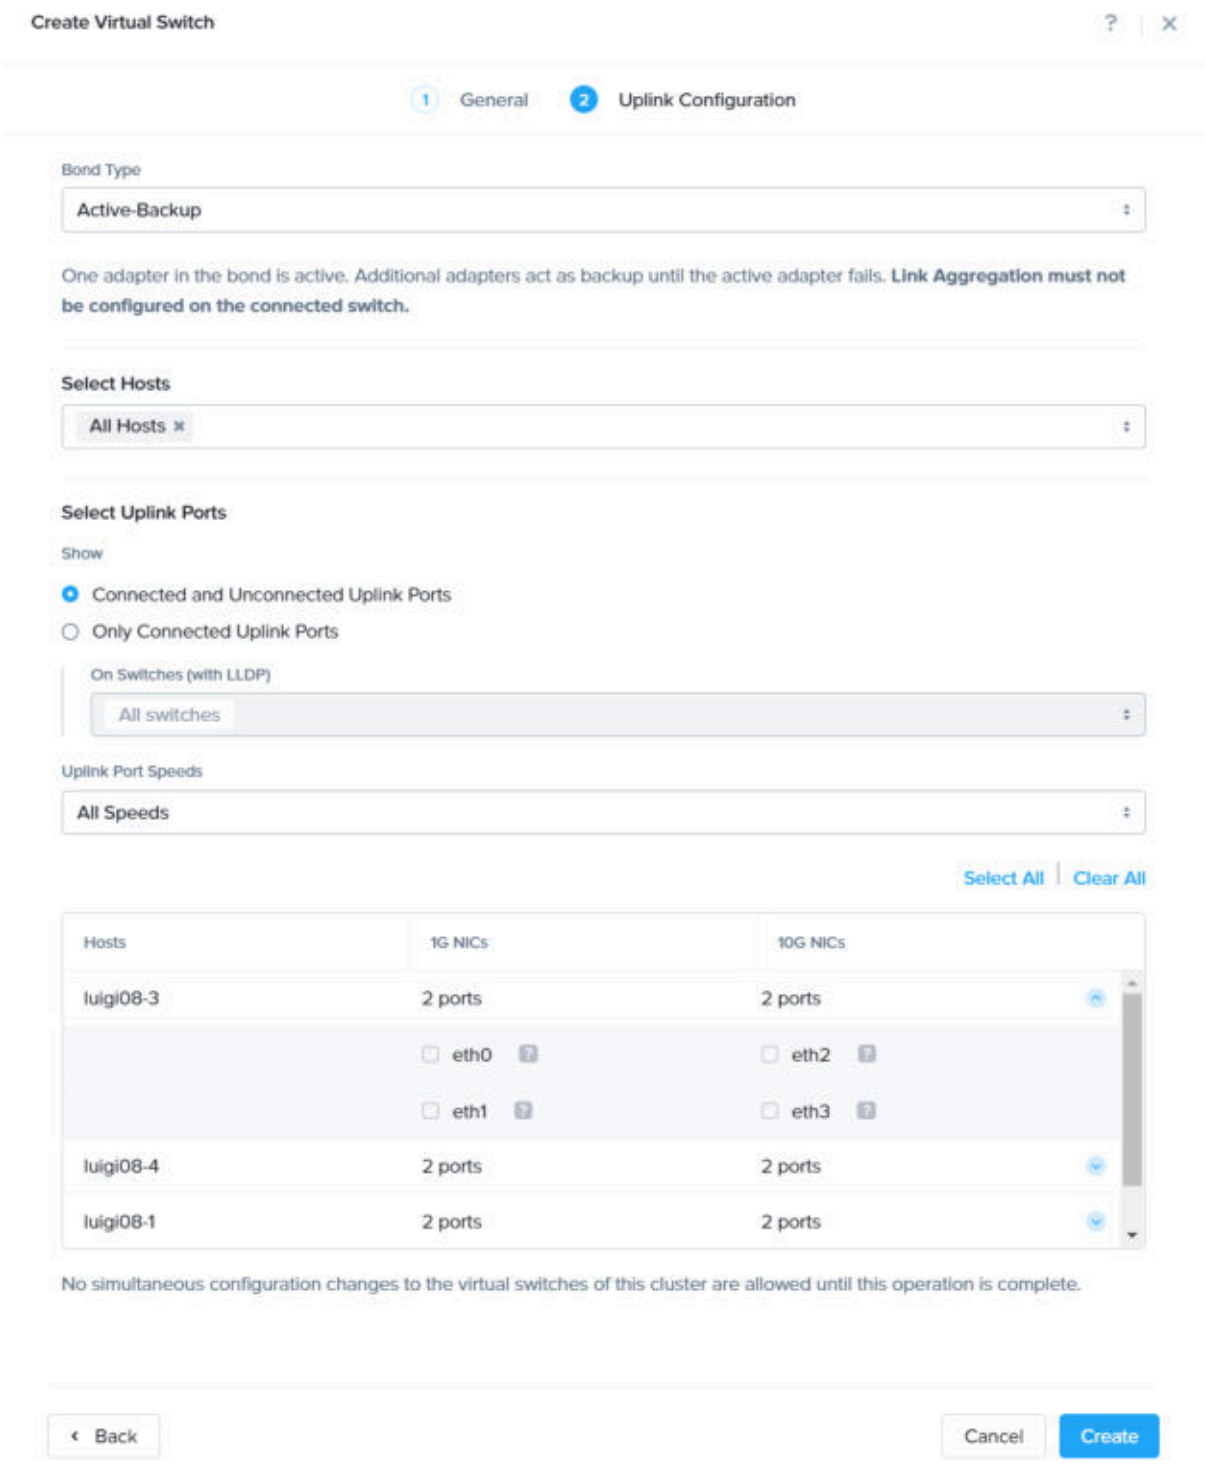

Figure 120: Create Virtual Switch - Uplink Configuration tab

Field	Description and Value
Bond Type	Select an appropriate bond type. See the *Bond Types* table for details about the bond types.
Select Hosts	Select the hosts that can host VMs.
Select Uplink Ports	Select the criteria that need to be satisfied for the uplink ports. The available uplink ports that satisfy the criteria are displayed in the (Host port) table at the bottom of this tab.
Show	(Port Type) Select one of the following: **Connected and Unconnected Uplink Ports**: Select this if you want to use ports that are not currently connected but may be connected later. **Only Connected Uplink Ports**: Select this option if you want to use only the connected ports. You must also select in the **On Switches (with LLDP)** drop down box the switches with LLDP.
Uplink Port Speeds	Select a speed to display the ports that have the selected speed. You can select speeds such as **1G**, **10G** or both (**All Speeds**). The speeds displayed depend on the NIC type that is installed on the host. Based on your selection the columns in the (Host Port) table change dynamically to display the ports with the speeds you selected.

Field	Description and Value
(Host Port) table	Based on the Selections you made in this Select Uplink Ports section, a table displays the hosts that have the uplink ports that satisfy the selected criteria. Select the ports you need for this configuration from the list. Click the down arrow on the right side of the table to display the ports listed for each host. Click the check box of a port to select the port. When the check box is already checked, click check box to clear the selection. Clearing the selection removes the uplink port (NIC) from the virtual switch. Note: A port listing is greyed out if it is unavailable because it is already associated with another virtual switch. Click **Select All** to select all the ports available and listed. Click **Clear All** to unselect all the ports available and listed.

Table 29: Bond Types

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Backup	Recommended. Default configuration, which transmits all traffic over a single active adapter.	10 Gb	10 Gb
Active-Active with MAC pinning Also known as balance-slb	Works with caveats for multicast traffic. Increases host bandwidth utilization beyond a single 10 Gb adapter. Places each VM NIC on a single adapter at a time. Do not use this bond type with link aggregation protocols such as LACP.	10 Gb	20 Gb

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Active Also known as LACP with balance-tcp	LACP and link aggregation required. Increases host and VM bandwidth utilization beyond a single 10 Gb adapter by balancing VM NIC TCP and UDP sessions among adapters. Also used when network switches require LACP negotiation. The default LACP settings are: • Speed—Fast (1s) • Mode—Active fallback-active-backup • Priority—Default. This is not configurable.	20 Gb	20 Gb
No Uplink Bond	No uplink or a single uplink on each host. Virtual switch configured with the No uplink bond uplink bond type has 0 or 1 uplinks. When you configure a virtual switch with any other bond type, you must select at least two uplink ports on every node.	-	-

Note: The Maximum VM NIC Throughput and Maximum Host Throughput values are not restricted to the value provided in this table. The values in the table are indicated for an assumption of 2 x 10 Gb adapters with simplex speed.

For more information about uplink configuration, see the *Uplink Configuration (AHV Only)* topic in the *AHV Administration Guide*.

3. Click **Create** to create the virtual switch.

   Click **Cancel** to exit without creating the virtual switch.

   Click **Back** to go back to the General tab.

Deleting a Virtual Switch

About this task

For more information about virtual switch, see the *Layer 2 Network Management* section in the *AHV Administration Guide*.

You can delete a VS on the Virtual Switch tab in the Network Configuration dialog box. When you delete a virtual switch, the bridge created by the virtual switch is also removed completely.

You can also disable all the virtual switches. However, you can do so only using CLI. Prism does not support disabling virtual switches. Use the net.disable_virtual_switch to disable the virtual switch functionality and remove all the virtual switches including the default virtual switch vs0. However, the bridge and uplink bond or interface configurations are not affected by this command. This action cannot be undone either.

To delete a virtual switch, do the following:

Procedure

1. Click the **Delete** icon (trash icon) next to the virtual switch listed in the list on the Virtual Switch tab.

   The Delete Virtual Switch dialog box opens. Deletion cannot be undone or reversed. Therefore, you must confirm the deletion in this dialog box.

2. Click **Delete** to confirm deletion.

   Click **Cancel** to close the Delete Virtual Switch dialog box without deleting the virtual switch.

What to do next

Check if the virtual switch is deleted from the list on the Virtual Switch tab.

Migrating Bridges after Upgrade

Before you begin

Ensure the following prerequisites are satisfied before converting a bridge to a VS.

- The bridge intended to be converted to a Virtual Switch must have consistent configurations across all the nodes in the cluster in terms of bond-type, MTU and LACP parameters.
- Within a virtual switch, the VLAN IDs must be exclusive across networks. An exception to this configuration is when a network is IPAM enabled and the other network is not IPAM enabled. In a scenario where an existing bridge has duplicate VLAN IDs, only one of these networks (per IPAM state) will get migrated under the virtual switch. The additional networks remains as is with no visible impact to the functionality.

  Note:

  Networks with same the VLAN IDs can exist across different virtual switches.

About this task

You can migrate or convert the bridges other than br0 in the cluster to virtual switches after you upgraded the minimum or compatible version of AOS and AHV. You can convert the other bridges only on Prism Web Console or using aCLI. You can convert only one bridge at a time. You need to repeat the workflow for every bridge that you want to convert to a virtual switch.

Note:

You cannot migrate any bridges using Prism Central.

To migrate the bridges after upgrade, perform the following.

Procedure

1. Click the **Convert Bridges to VS** link on the Network page in Prism Web Console.

   If there are any bridges that you can migrate, then the **Convert Bridges to VS** dialog box is displayed as follows.

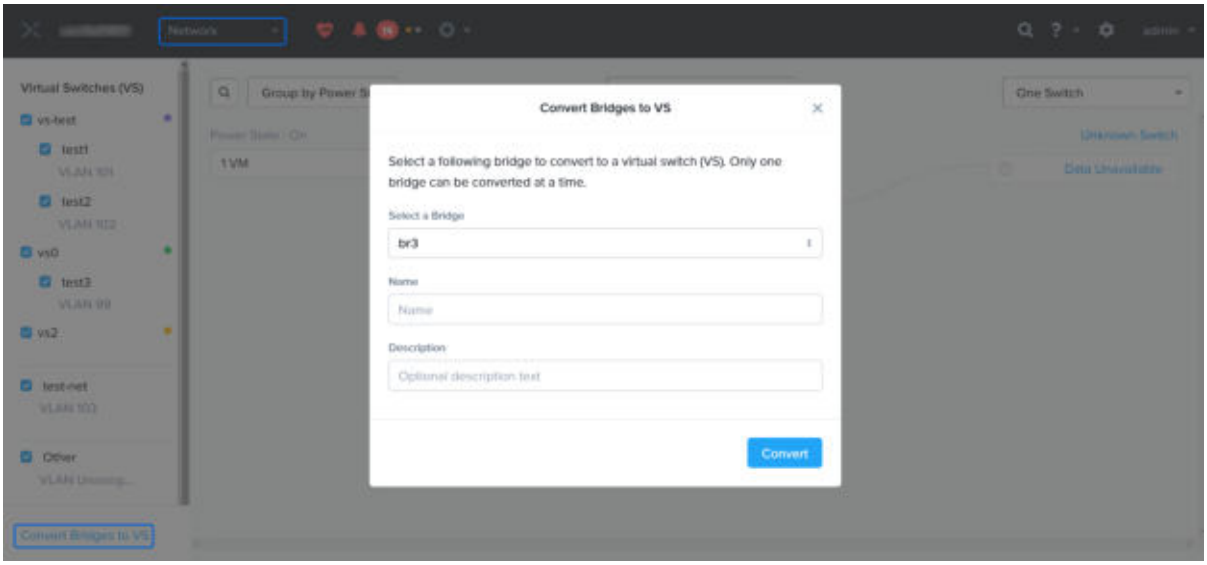

Figure 121: Convert Bridges to VS

If there are no bridges that you can migrate, then the following is displayed.

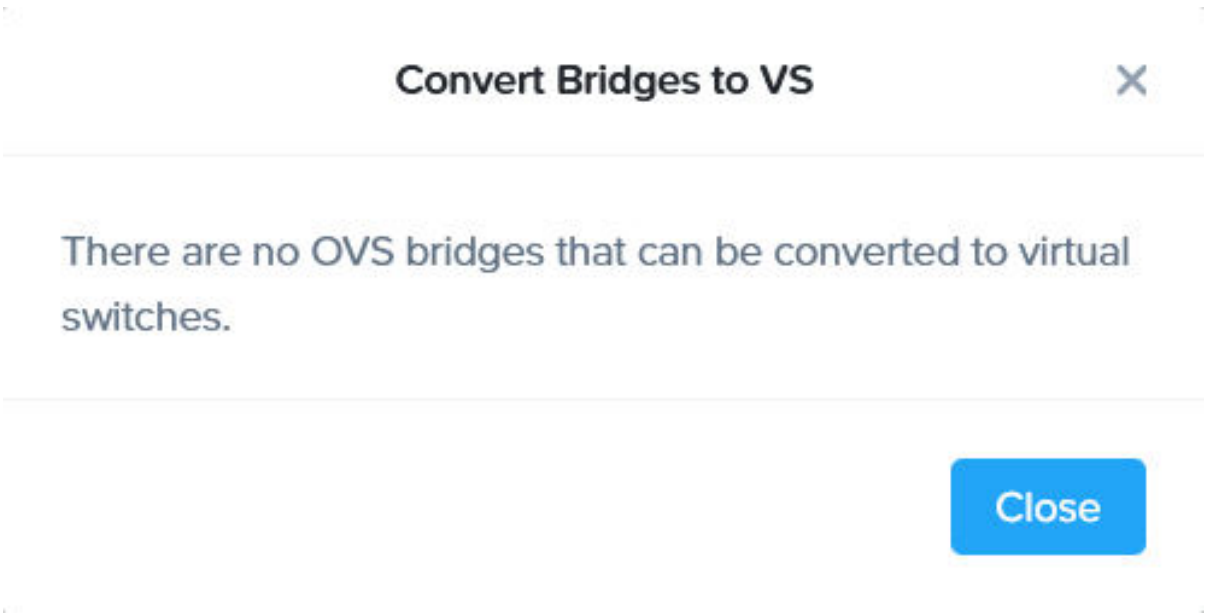

Figure 122: No OVS Bridges to Convert

2. On the **Convert Bridges to VS** dialog box, select the bridge you want to migrate in the **Select a Bridge** drop down field.

   The **Select a Bridge** drop down field provides a list of bridges you can migrate. Select the bridge you want to migrate and complete the following steps.

   a. Provide a **Name** and a **Description** for the virtual switch that you want to create by migrating the bridge.

      For example, if you select *br1* in **Select a Bridge**, you can provide *vs1* as the name for the virtual switch.

3. Click **Convert**.

   The converted virtual switch is available in the Network page in Prism Web Console and the Virtual Switch tab on the Network Configuration page in Prism Central.

Network Configuration for VM Interfaces

In clusters with Nutanix virtualization management (such as those running AHV), you can configure network connections through the web console. Each VM network interface is bound to a virtual network, and each virtual network is bound to a single VLAN. Additionally, on AHV, Hyper-V, and ESXi, you can secure intra-cluster communications by segmenting the virtual network on the cluster.

About this task

To configure a virtual network for VM interfaces, do one of the following:

Procedure

- To configure virtual networks for user VM interfaces, see Configuring a Virtual Network for Guest VM Interfaces on page 194.
- To secure intra-cluster communications by configuring network segmentation, see “” () .

Configuring a Virtual Network for Guest VM Interfaces

Follow this procedure to configure a virtual network for guest VMs. You can associate the network with a VLAN and configure IP address management for interfaces on the network.

About this task

> Note: Do not add any other device (including guest VMs) to the VLAN to which the CVM and hypervisor host are assigned. Isolate guest VMs on one or more separate VLANs.

To configure a virtual network for guest VM interfaces, do the following:

Procedure

1. Click the gear icon in the main menu and then select Network Configuration in the **Settings** page.

   > Note: This option only appears in clusters that support this feature.

   The Network Configuration window appears.

Figure 123: Network Configuration Window

2. Click the **Virtual Networks** tab and then click the **Create Network** button.
   The **Create Network** dialog box appears. Do the following in the indicated fields:

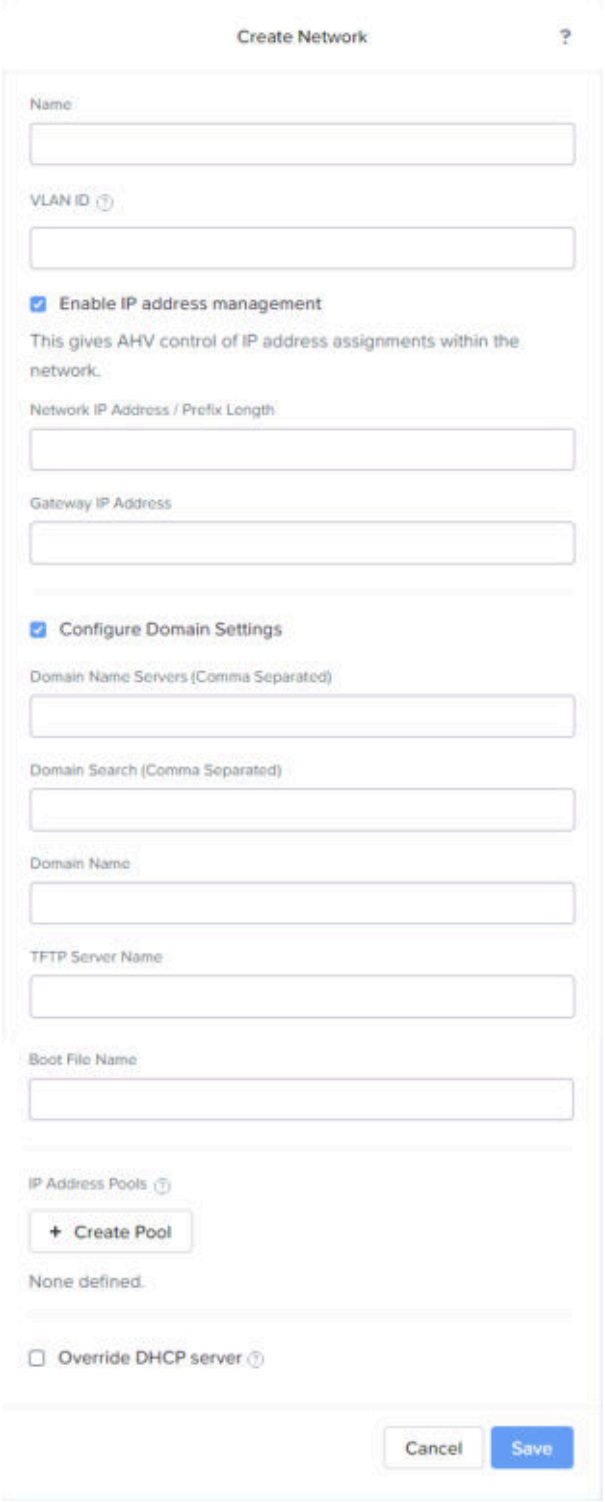

Figure 124: Create Network Dialog Box

   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 displays more fields. If you do not check the box, Files does not attempt network management. In this case, it is assumed management for this virtual LAN is handled outside the cluster.

   > Note: If you do not enable this option while creating a network, you cannot enable or disable IP address management (IPAM).
   >
   > IP Address Management (IPAM) is a feature of AHV that allows it to assign IP addresses automatically to VMs by using DHCP. See *IP Address Management* section in the *AHV Administration Guide*.

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.

3. 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 dialog box:
   If you do not assign a pool, assign IP addresses to VMs manually.

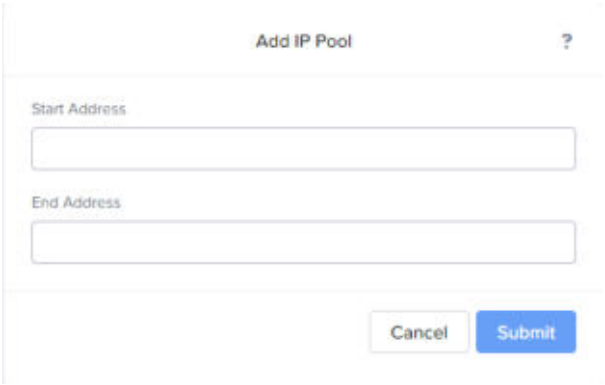

Figure 125: Add IP Pool Dialog Box

   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 window and return to the Create Network dialog box.

4. To configure a DHCP server, click 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 you do not check the box, the **DHCP Server IP Address** field does not display, and the DHCP server IP address generates 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.

Figure 126: DHCP Server IP Field

5. When all the information is correct, click the **Save** button to configure the network connection, close the Create Network dialog box, and return to the Network Configuration window.

   Note:

   - You can also specify network mapping to control network configuration for the VMs when they start on the remote site. For more information about configuring

networking mapping on remote site, see "Configuring a Remote Site (Physical Cluster)" in the *Prism Web Console Guide*.

- Verify (or update as needed) that the physical switch configuration allows traffic for the same VLAN's that are configured for the virtual networks.

Verifying IPAM Address Pool for Sufficient IP Addresses

This procedure verifies the IPAM managed IP address pool for sufficient IP addresses and checks when the IP pool is close to running out of IP addresses.

About this task

To verify the IPAM managed IP address pool for sufficient IP addresses, do the following:

Procedure

1. Click the gear icon in the main menu and then select Network Configuration in the **Setting** page.

2. Select the **Networks** tab (default) to view networks for the VMs in the Network Configuration dialog box. It consists of the following:

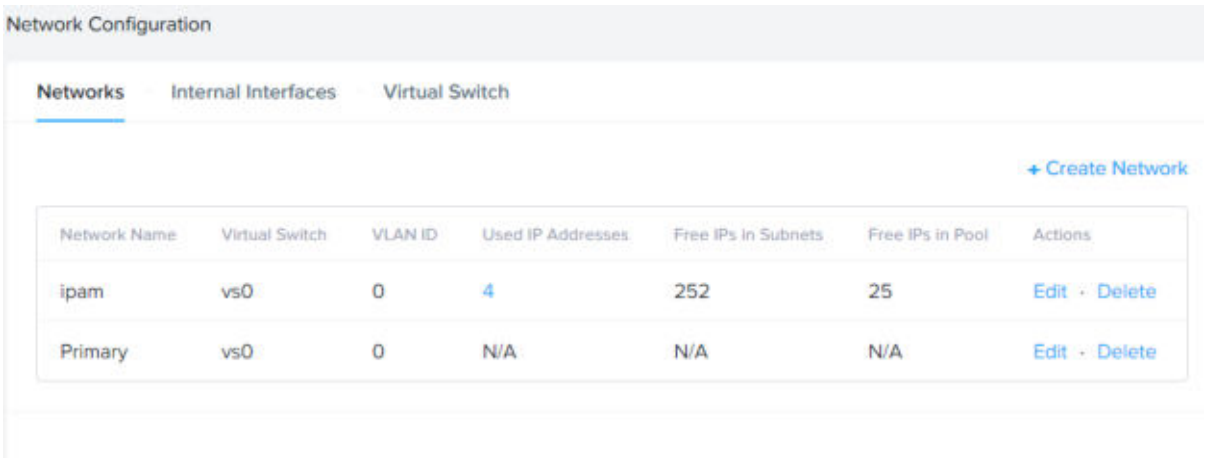

Figure 127: Verify the IPAM Managed IPAddress Pool

a. **Network Name**: Displays the networks configured on the VM.

b. **Virtaul Switch**: Displays the virtual Switch configured on the VM.

c. **VLAN ID**: Displays the VLAN ID configured on the VM.

d. **Used IP Addresses**: Displays the number of IP addresses in use by the IPAM enabled network. Click on the value displayed to view which VM is using which IP addresses from the IPAM enabled network.

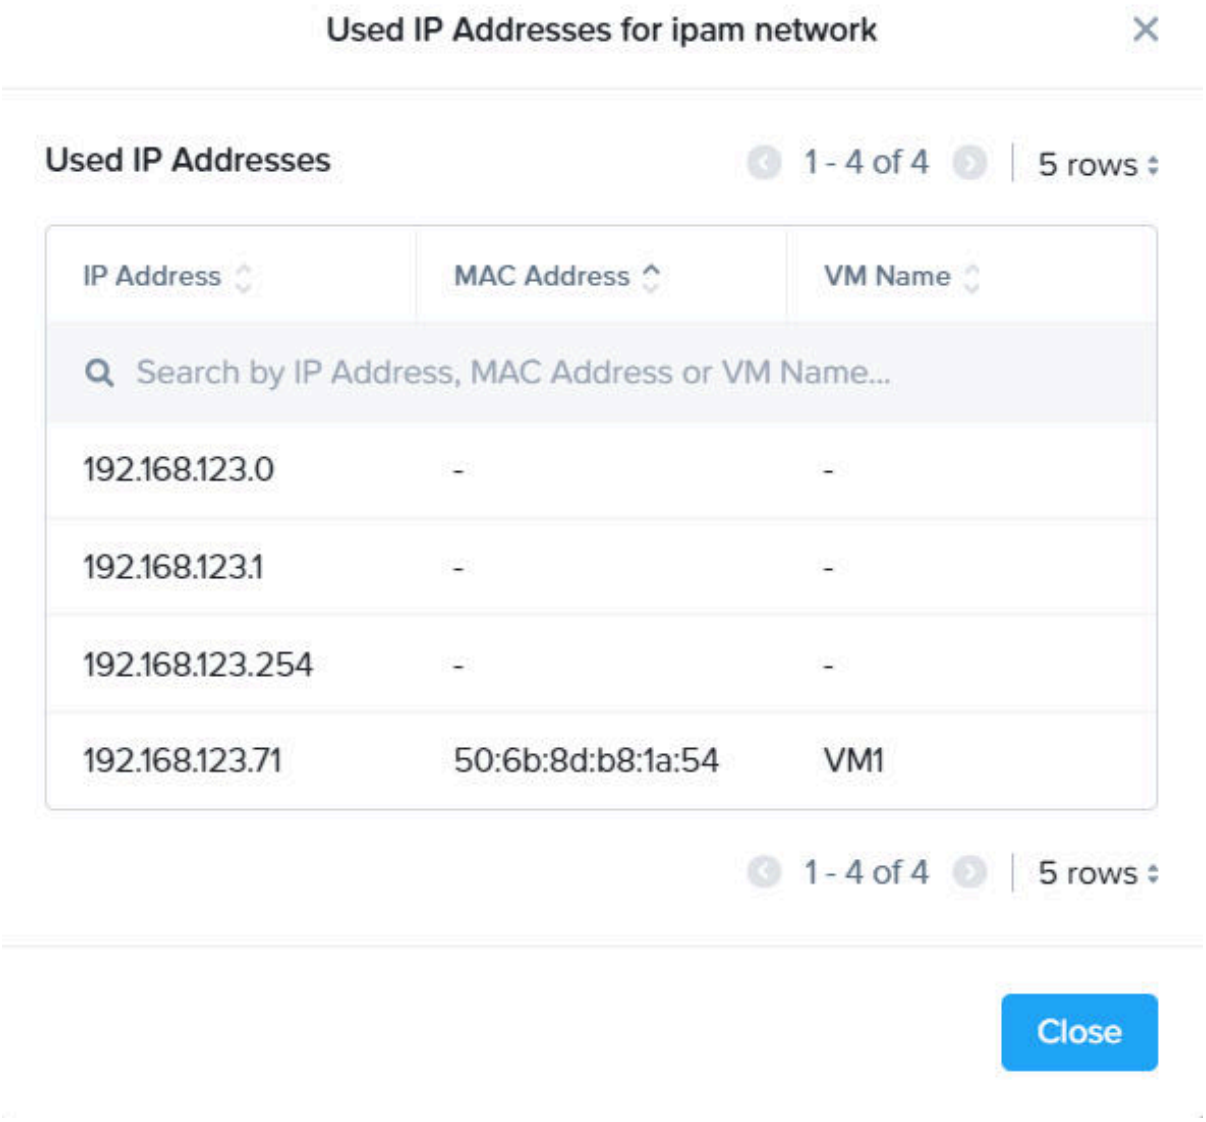

Figure 128: Used IP Address for IPAM Network

e. **Free IPs in Subnets**: Displays the number of IP addresses configured in the subnet.

f. **Free IPs in Pool**: Displays the number of free IPs available in the IP Pool configured for the subnet. The value in this field dynamically changes every time the network dynamically assigns a new IP to a VM.

g. **Actions**: You can select either the ui**Edit** or **Delete** actions to update the Network configuration.

Modifying Network Connections

Follow this procedure to modify network connections defined through the Nutanix virtualization management.

About this task

To modify or delete a network configuration, do the following:

Procedure

1. Do one of the following:
   - » Go to the List tab of the VMs dashboard and click the **Network Config** button.
   - » Go to the Networks dashboard and click the **Network Config** button.
2. To modify a network configuration, in the Network tab, select the target network, and click the pencil icon.
   The Update Network Configuration dialog box appears, which contains the same fields as the Create Network Configuration dialog box (see Configuring a Virtual Network for Guest VM Interfaces on page 194). Do the following:

   > Note: Changing the VLAN ID of a network in use is not allowed; only the name can be changed.

   a. Update the field values as desired.
   b. When the configuration is correct, click the **Save** button to close the page, update the network configuration, and return to the Network Configuration page.
3. To delete a network configuration, 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.

Network Segmentation

You can segment the network on a Nutanix cluster by using Prism Element.

For more information, see *Securing Traffic Through Network Segmentation* section in the *Security Guide*.

Configuring Network Switch Information

About this task

To track and record networking statistics for a cluster, the cluster requires information about the first-hop network switches and the switch ports being used. Switch port discovery is supported with switches that are RFC 2674 compliant. Switch port discovery involves obtaining statistics from the Q-BRIDGE-MIB on the switch and then identifying the MAC address that corresponds to the host. Such discovery is currently best-effort, so it is possible that, at times, an interface might not be discovered.

Before you begin

Before you configure network switch information in the Prism web console, configure the corresponding SNMP settings on the first-hop network switch.

About this task

To configure one or more network switches for statistics collection, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Network Switch** in the **Settings** page.

   Note: Network switch configuration is supported only for AHV clusters.

   The Network Switch Configuration dialog box appears.

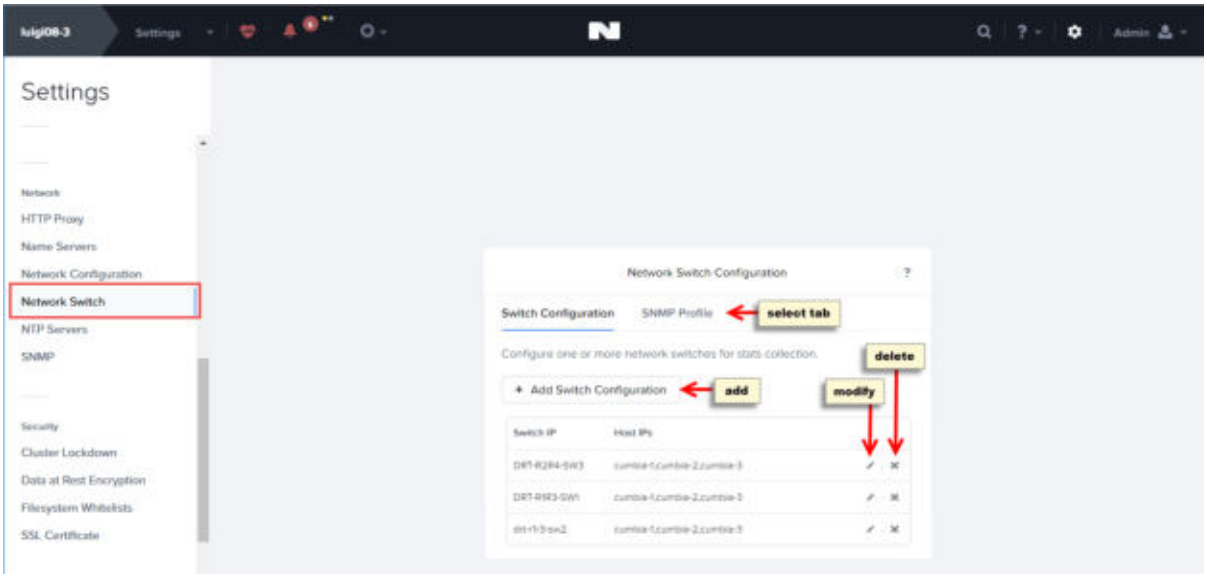

Figure 129: Network Switch Configuration Window

2. Click the **Switch Configuration** tab and then click the **Add Switch Configuration** button.

3. Do the following in the indicated fields:

   a. **Switch Management IP Address**: Enter the Management IP address of the switch or the fully qualified switch name.

   b. **Host IP Addresses or Host Name (separated by commas)**: Enter the IP address or fully qualified host name of each host in the cluster that uses this switch to route traffic.

      When there are multiple hosts, enter the addresses in a comma separated list. Failing to add the host list might result in issues with switch port discovery.

   c. **SNMP Profile**: Select the SNMP profile to apply (or **None**) from the pull-down list.

      Any SNMP profiles you have created appear in the list (see Creating SNMP Profiles on page 206).

      > Note: Selecting a profile populates the remaining fields automatically with the profile values. If you have not created a profile (or select **None** from the list), you must enter the values in the remaining fields manually.

   d. **SNMP Version**: Select the SNMP version to use from the pull-down list.

      The options are **SNMPv2c** and **SNMPv3**.

   e. **SNMP Security Level** (SNMPv3 only): Select the security level to enforce from the pull-down list.

      The options are **No Authorization No Privacy**, **Authorization But No Privacy**, and **Authorization and Privacy**. This field appears only when SNMPv3 is selected as the version.

   f. **SNMP Community Name**: Enter the SNMP community name to use.

   g. **SNMP Username**: Enter the SNMP user name.

   h. **SNMP Authentication Type**: Select the authentication protocol to use from the pull-down list.

      The options are **None** and **SHA**.

      > Note: This field and the following three fields are set to **None** or left blank (and read only) when the version is SNMPv2c or the security level is set to no authorization.

   i. **SNMP Authentication Pass Phrase**: Enter the appropriate authentication pass phrase.

   j. **SNMP Privacy Type**: Select the privacy protocol to use from the pull-down list.

      The options are **None**, **AES**, and **DES**.

   k. **SNMP Privacy Pass Phrase**: Enter the appropriate privacy pass phrase.

   l. When all the fields are correct, click the **Save** button.

      This saves the profile and displays the new entry in the Switch Configuration tab.

      > Note: As a security protection, the **SNMP Authentication Pass Phrase** and **SNMP Privacy Pass Phrase** fields appear blank after saving (but the entered phrases are saved).

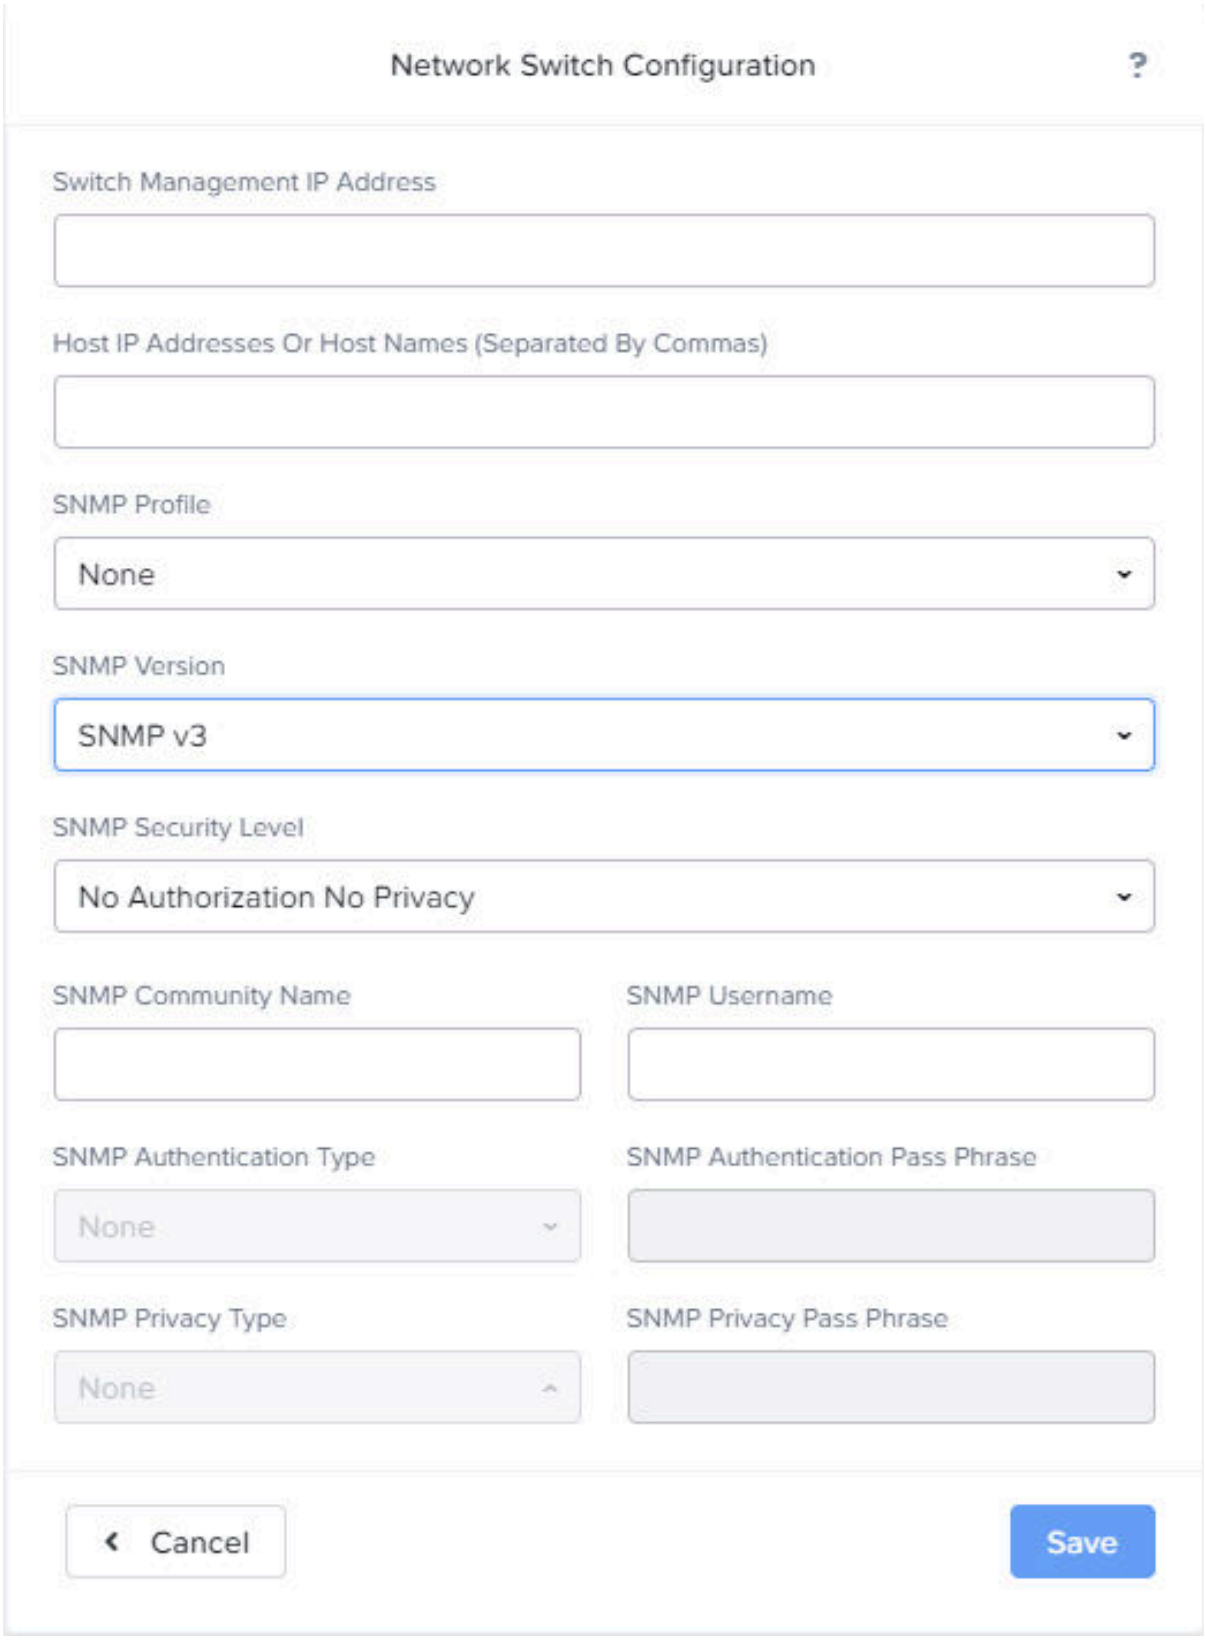

Figure 130: Add Network Switch Configuration

Enabling LAG and LACP on the ToR Switch (AHV Only)

Active-Active with MAC pinning does not require link aggregation since each source MAC address is balanced to only one adapter at a time. If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

About this task

In an **Active-Active** NIC-teaming configuration, network traffic is balanced among the members of the team based on source and destination IP addresses and TCP and UDP ports. With link aggregation negotiated by LACP, the uplinks appear as a single layer-2 link so a VM MAC address can appear on multiple links and use the bandwidth of all uplinks. If you do not enable LAG and LACP in this **Active-Active** configuration, network traffic might be dropped by the switch.

Perform the following procedure to enable LAG and LACP on a TOR switch after enabling the Active-Active policy by using the Prism Element web console.

Procedure

1. Put the node in maintenance mode. This is in addition to the previous maintenance mode that enabled Active-Active on the node.
2. Enable LAG and LACP on the ToR switch connected to that node.
3. Exit maintenance mode after LAG and LACP is successfully enabled.
4. Repeat steps 1 to 3 for every node in the cluster.

Creating SNMP Profiles

About this task

If you need to configure multiple network switches, it may be useful to create one or more SNMP profiles that can be applied when configuring the switches (see Configuring Network Switch Information on page 202). To create an SNMP profile, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Network Switch** in the **Settings** page.

   Note: Network switch configuration is supported only for AHV clusters.

   The Network Switch Configuration dialog box appears.

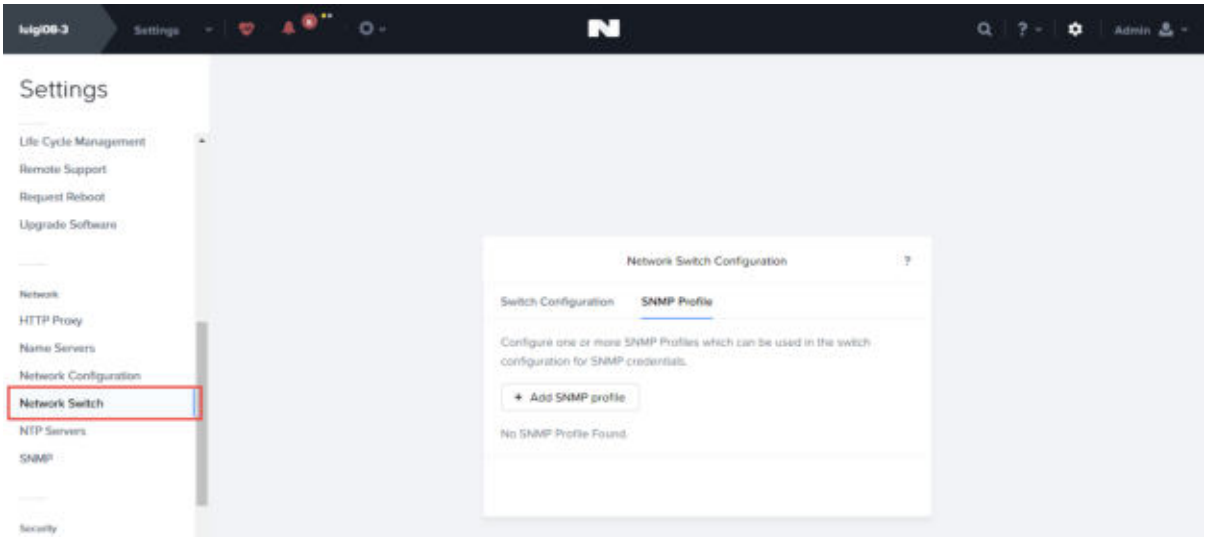

Figure 131: Network Switch Configuration Window

2. Click the **SNMP Profile** tab and then click the **Add SNMP Profile** button.

3. Do the following in the indicated fields:

   a. **Profile Name**: Enter a name for the profile.
   The name can be anything as it is for display purposes only.

   b. **SNMP Version**: Select the SNMP version to use from the pull-down list.
   The options are **SNMPv2c** and **SNMPv3**.

   c. **SNMP Security Level** (SNMPv3 only): Select the security level to enforce from the pull-down list.
   The options are **No Authorization No Privacy**, **Authorization But No Privacy**, and **Authorization and Privacy**. This field appears only when SNMPv3 is selected as the version.

   d. **SNMP Community Name**: Enter the SNMP community name to use.

   e. **SNMP Username**: Enter the SNMP user name.

   f. **SNMP Authentication Type**: Select the authentication protocol to use from the pull-down list.
   The options are **None** and **SHA**.

   > Note: This field and the following three fields are set to **None** or left blank (and read only) when the version is SNMPv2c or the security level is set to no authorization.

   g. **SNMP Authentication Pass Phrase**: Enter the appropriate authentication pass phrase.

   h. **SNMP Privacy Type**: Select the privacy protocol to use from the pull-down list.
   The options are **None**, **AES**, and **DES**.

   i. **SNMP Privacy Pass Phrase**: Enter the appropriate privacy pass phrase.

   j. When all the fields are correct, click the **Save** button.
   This saves the profile and displays the new entry in the SNMP Profile tab.

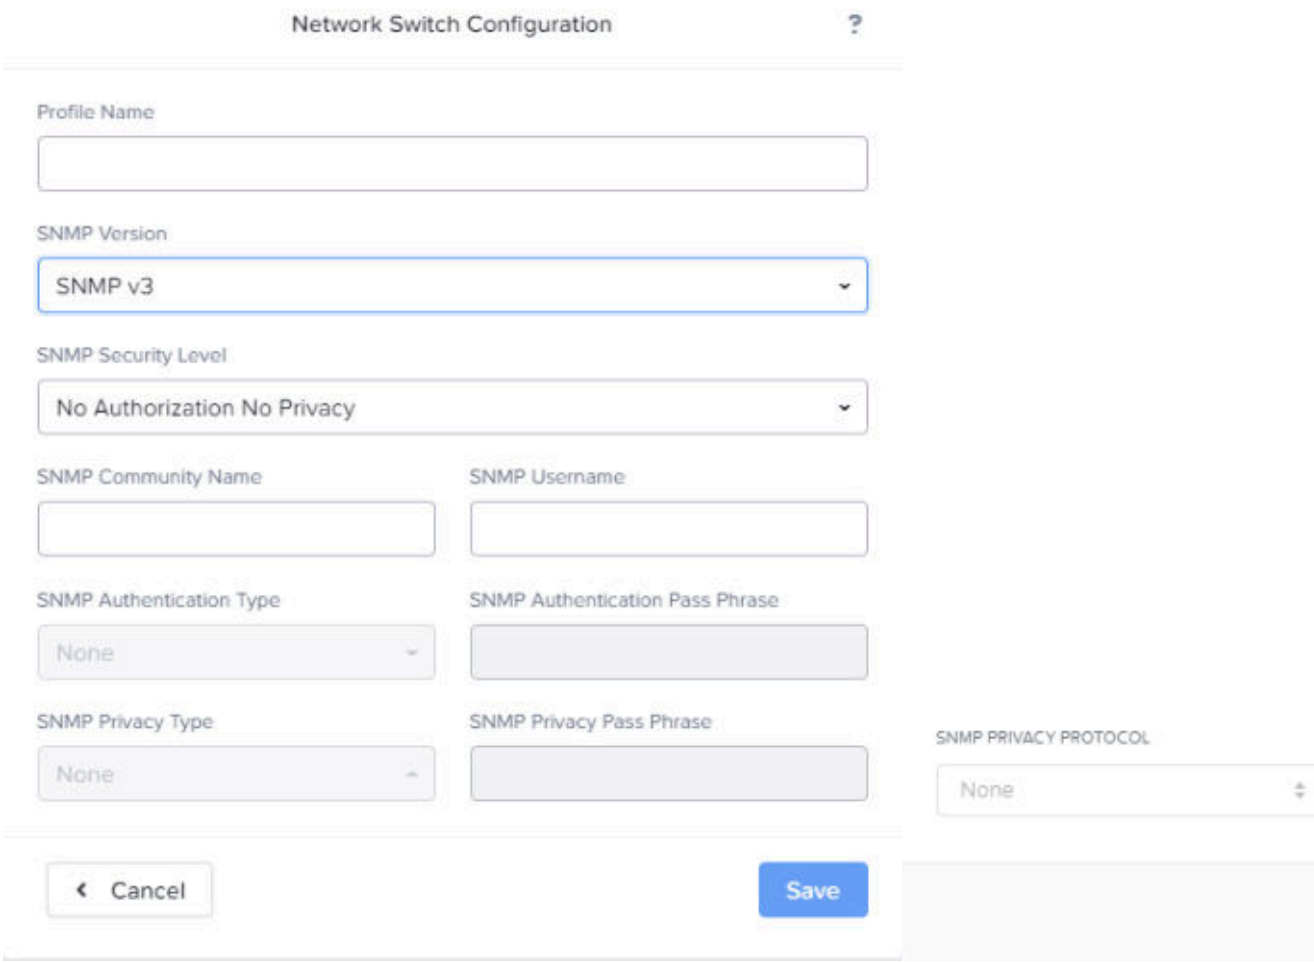

Figure 132: Network Switch Configuration Window (add SNMP profile)

Modifying Switch Information

About this task

To modify the information about a network switch (or the SNMP profile used in the switch information), do the following:

Procedure

1. Click the gear icon in the main menu and then select **Network Switch** in the **Settings** page.

   Note: Network switch configuration is supported only for AHV clusters.

   The Network Switch Configuration dialog box appears.

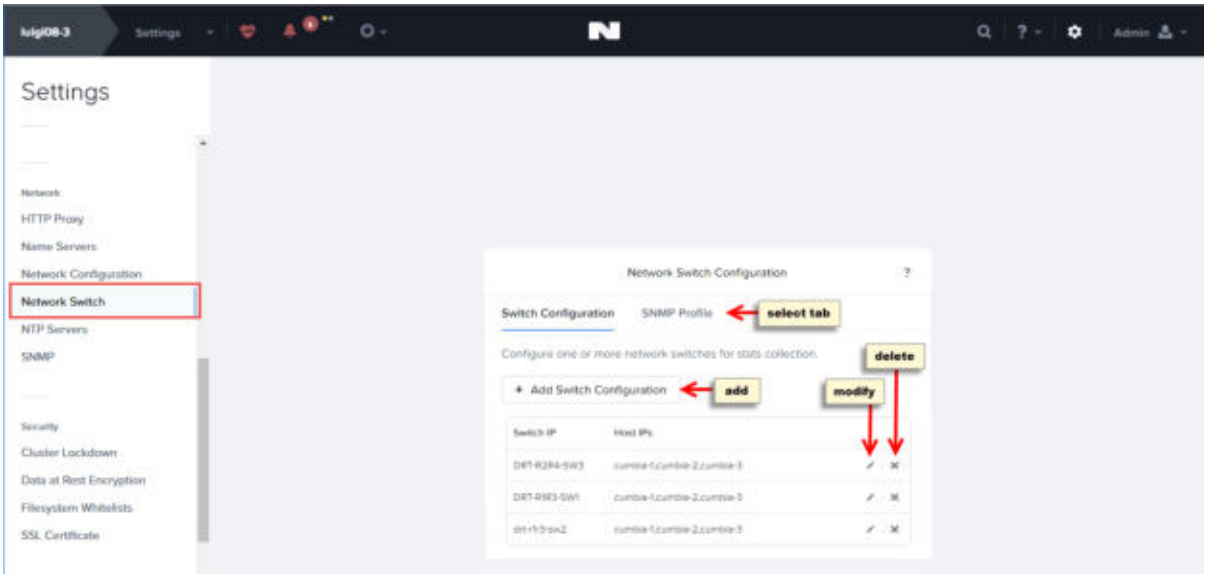

Figure 133: Network Switch Configuration Window

2. To modify a switch configuration (or SNMP profile), select the **Switch Configuration** (or **SNMP Profile**) tab, go to the target switch (or profile) entry in the table, and click the pencil icon.

   This displays the configuration fields (see Configuring Network Switch Information on page 202 or Creating SNMP Profiles on page 206). Edit the entries as desired and then click the **Save** button.

3. To delete a switch configuration (or SNMP profile), select the **Switch Configuration** (or **SNMP Profile**) tab, go to the target switch (or profile) entry in the table, and click the X icon.

   This deletes the entry and removes it from the table.

Network Visualization

The network visualizer is a consolidated graphical representation of the network formed by the VMs and hosts in a Nutanix cluster and first-hop switches. You can use the visualizer to monitor the network and to obtain information that helps you troubleshoot network issues.

You can use the visualizer to view the following:

- Physical and logical network topology.
- Summary of the number and types of devices in the network.
- Network configuration of the devices in the topology and of components such as physical and virtual NICs.
- Real-time usage graphs of physical and virtual interfaces.

CAUTION: You cannot use the visualizer to configure the network. The network visualizer is available only on AHV clusters.

Prerequisites

Before you use the network visualizer, do the following:

- Configure network switch information on the Nutanix cluster. See Configuring Network Switch Information on page 202.
- Enable LLDP or CDP on the first-hop switches. The visualizer uses LLDP or CDP to determine which switch port is connected to a given host interface. If LLDP or CDP is unavailable, SNMP data is used on a best-effort basis. For information about configuring LLDP or CDP, see the switch manufacturer's documentation.
- (Optional) Configure SNMP v3 or SNMP v2c on the first-hop switches if you want the network visualizer to display the switch and switch interface statistics. The visualizer uses SNMP for discovery and to obtain real-time usage statistics from the switches. For information about configuring SNMP, see the switch manufacturer's documentation.

Note: This is not a mandatory requirement to use the network visualizer. This is a prerequisite only if you want the network visualizer to display the switch and switch interface statistics.

Note: Network visualization depends on a functional internal DNS system to map switch hostnames and IP addresses based on the LLDP responses. Incorrect or partial DNS configuration displays inaccurate network details in the UI. To troubleshoot the network visualization issues, see *KB-4085* in the Nutanix Support portal.

Network Visualizer

The network visualizer displays interactive visual elements for the networked devices and for network components such as physical and logical interfaces. It also provides filtering and grouping capabilities that you can use to limit the display to a specific set of networked devices and connections.

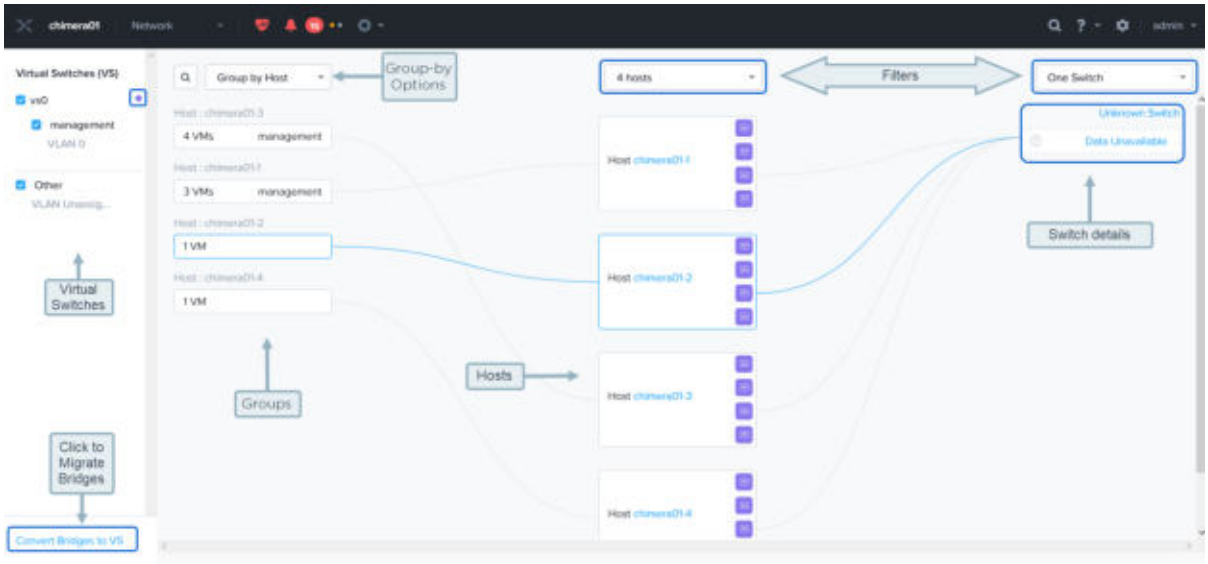

Figure 134: Network Visualizer

The network visualizer includes the virtual networks pane and the topology view.

Virtual Networks Pane

Lists the virtual switches (VSs) configured on the Nutanix cluster. Selecting a VS includes the VMs on that VS in the topology view. Conversely, clearing a check box excludes the VMs on that VS from the topology view. You can show up to five VSs at a time.

The check box titled Other corresponds to VMs that are not on any VS. At a minimum, this check box is associated with the Controller VMs in the cluster.

Topology View

Displays the network topology. The topology view shows the following entities:

Virtual Switch (VS)

VSs are configured on the cluster. The visualizer displays a different color for each VS. It shows the VSs to which a VM or the VMs in a VM group belong. It also shows which VSs are configured on a first-hop switch.

VMs

VMs on the VSs that are selected in the virtual networks pane. Filter and group-by options enable you to customize the topology view.

Hosts

Hosts in the Nutanix cluster. The filter above the hosts enables you to specify which hosts you want to show in the topology view.

Switches

First-hop switches and the VSs configured on each of them. The filter above the switches enables you to specify which switches you want to show in the topology view.

Viewing the Network Visualizer

About this task

To view the network visualizer, do the following:

Procedure

1. Log on to the web console.
2. Click **Network** in the main menu.

Customizing the Topology View

You can group VMs by a VM property or use a search filter to specify which network devices you want to show or exclude.

About this task

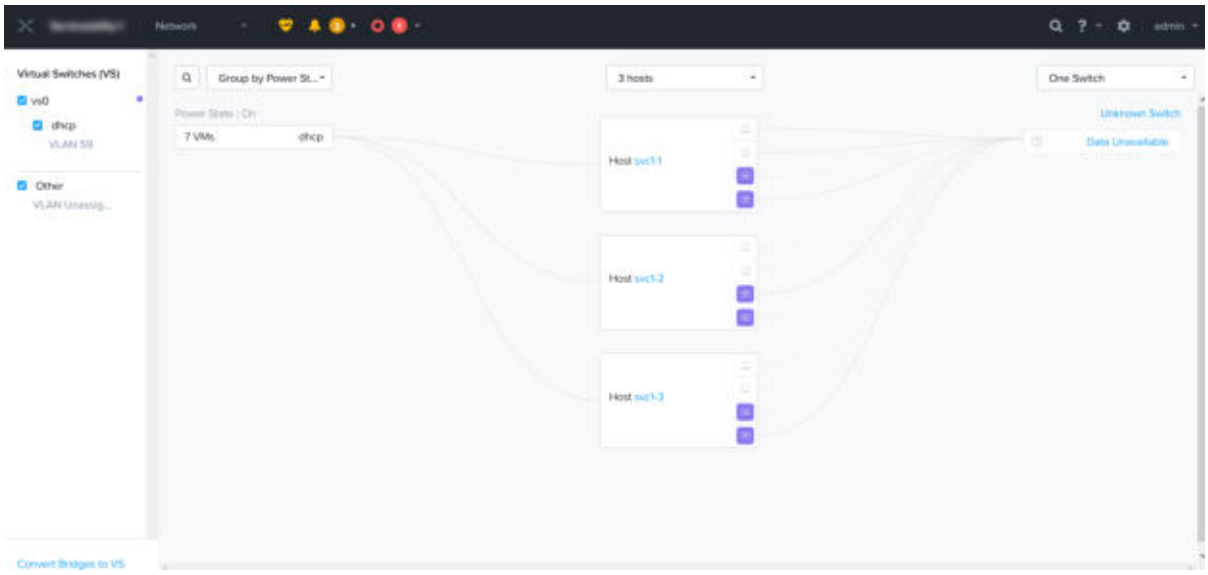

Figure 135: Network page

To customize the topology view, do the following:

Procedure

1. Specify which virtual switches you want to show in the topology view:

   - In the **Virtual Switch (VS)** pane, select the virtual switches and the networks in each virtual switch that you want to display in the topology view and clear those that you do not want to display.
   - Select or clear **Other** if you want to include or exclude, respectively, the VMs that are not on any VLANs.

2. Specify which VMs you want to display in the topology view:

   - Select a group-by option from the menu at the top of the VMs. The following group-by options are available:

     - **Power State**. Groups VMs by states such as On and Off. By default, the VMs are grouped by power state.
     - **Host**. Groups VMs by the host on which they are running.
     - **VM Type**. Groups VMs into guest VMs and Controller VMs.

   - Enter a string in the search filter field to filter VMs that match the search string.
   - If the group-by and filter operations result in a VM group, click the VM group to show the VMs in the group.

     When you click a VM group, the visualizer displays ten VMs at first. To load ten more VMs, click **Load More VMs**.

     To group the VMs again or to clear the filter, click **Back** beside the group-by menu.

3. Specify which Nutanix hosts you want to display in the topology view:
   a. Click the menu above the Nutanix hosts.
   b. Select the hosts that you want to show and clear those that you want to exclude.
4. Specify which switches you want to display in the topology view:
   a. Click the menu above the switches.
   b. Select the switches that you want to show and clear those that you want to exclude.

Viewing VM Information

In the visualizer, you can view the settings and real-time statistics of a virtual NIC.

About this task

To view VM information, do the following:

Procedure

1. Use the group-by and filtering capabilities of the visualizer (see Customizing the Topology View on page 212) to display the VM you want, and then click the name of the VM.

   The VM network information window appears.

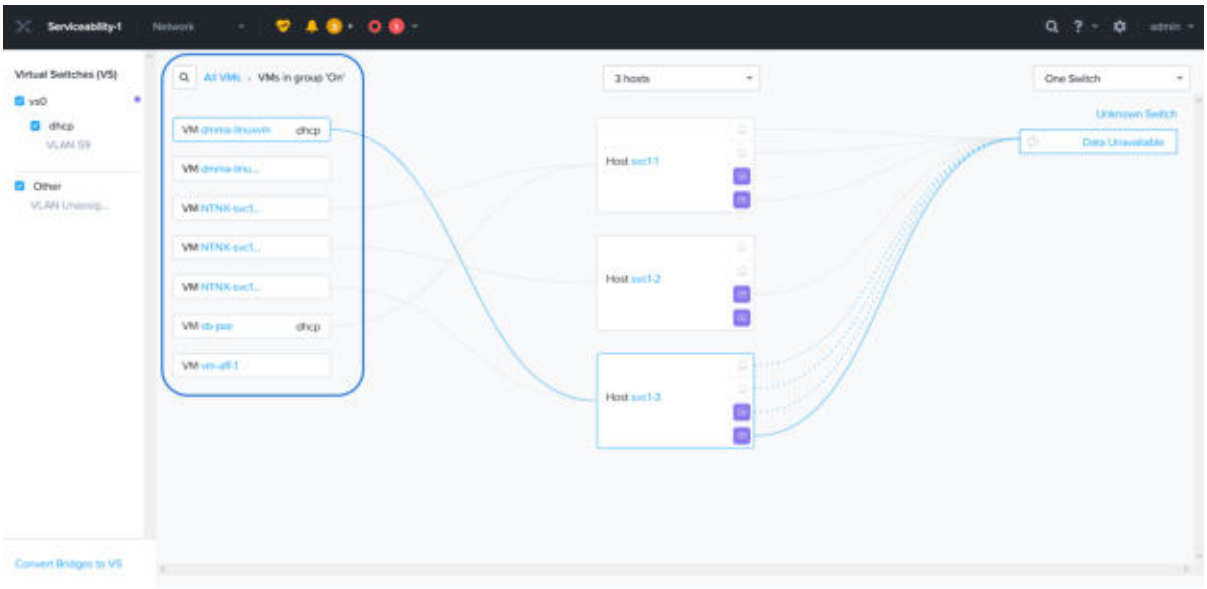

Figure 136: VM Network

2. From the virtual NIC list at the top of the dialog box, select the NIC for which you want show settings and statistics.

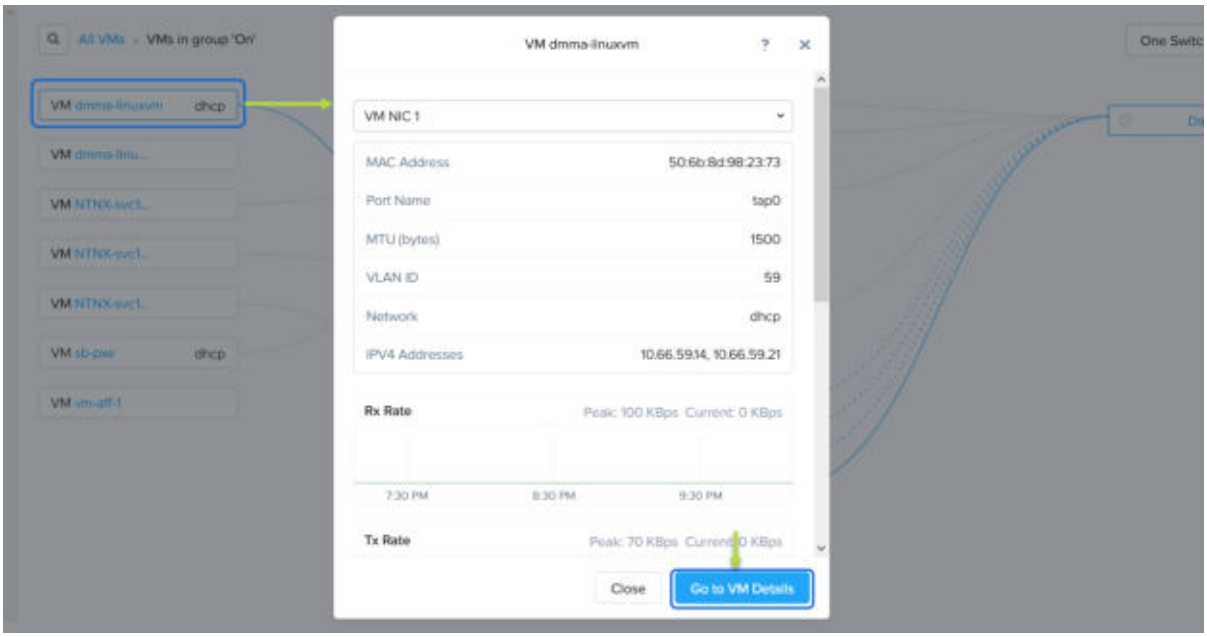

Figure 137: VM Details

For information about the statistics that are displayed for the virtual NICs, see VM NICs Tab on page 311.

3. Optionally, point to a location on a graph to view the value at that point in time.

4. If you want additional information about the VM, click **Go to VM Details**.
The **VM** table view on the **VM** page is displayed.

> Tip: You can return to the visualizer by pressing the back button in your browser.

Viewing Host Information

In the visualizer, you can view the internal network diagram of a Nutanix host. You can click a network component, such as a bridge or bond, to view the settings and statistics of that component.

About this task

To view host information, do the following:

Procedure

1. In the topology view, hover on or click a host interface to view the settings of the interface.

   For a host interface, statistics are shown in addition to interface settings. Point to a location on a graph to view the value at that point in time.

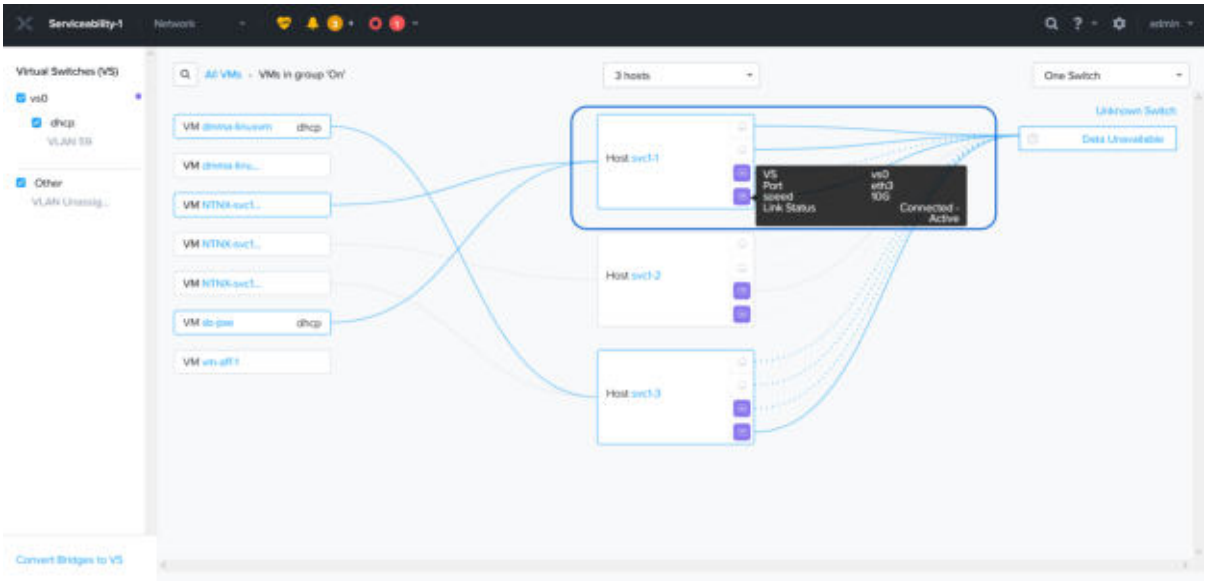

Figure 138: Host Interface Information

For information about the statistics that are displayed for the host NICs, see Host NICs Tab on page 242.

These lines are highlighted because they are related to the host of selected VM. A solid line leading from a bond to a host interface indicates the uplink is connected and is an active connection. A dotted line indicates a connected but unconfigured passive bond interface. A dotted line also indicates an uplink on a virtual switch with the bond type configured as No Uplink Bond.

2. Click the name of a host.
   The host network information window appears. The window shows a diagram of the host's internal network. The right pane displays host network information.

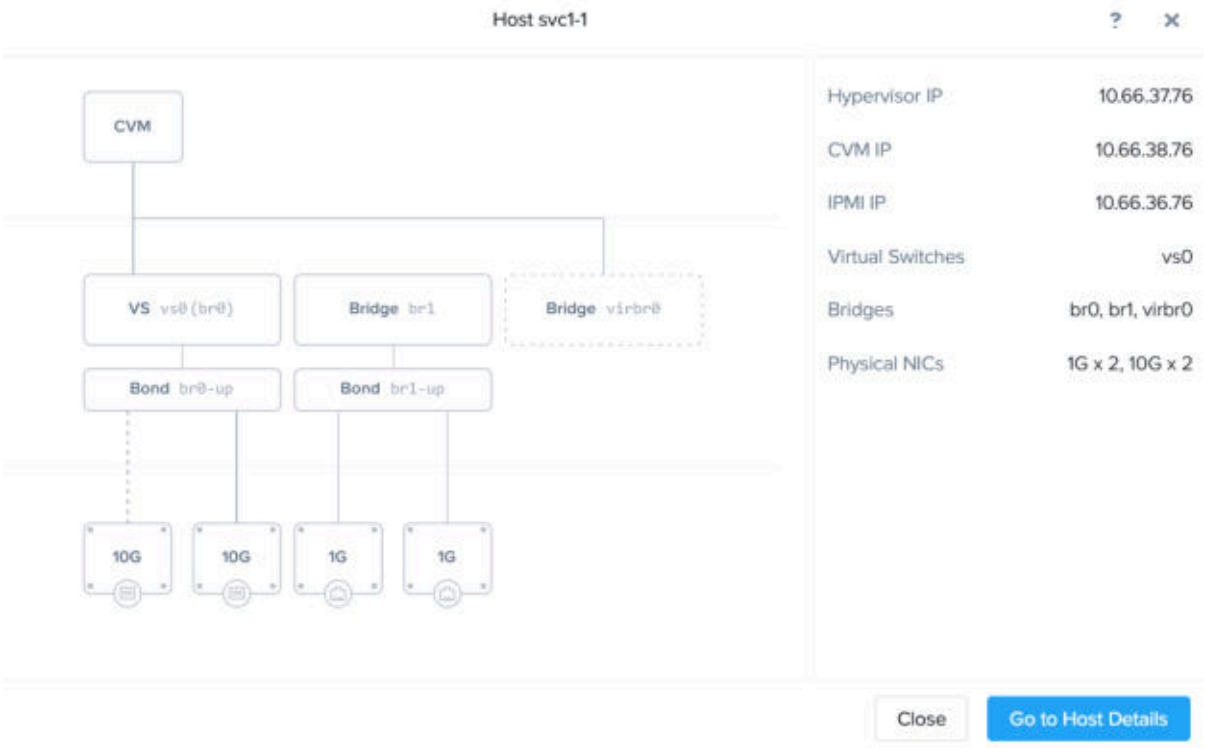

Figure 139: Host Network

3. Click any network component in the diagram and view its network settings and statistics in the right pane.

   » Click the Controller VM, and then, in the right pane, select a virtual NIC from the list to view settings and statistics of that virtual interface. Optionally, point to a location on a graph to view the value at that point in time.

   For information about the statistics that are displayed for the virtual NICs, see VM NICs Tab on page 311.

   » Click a virtual switch to view the configuration details of the virtual switch.

   » Click a bridge to view the configuration details of the bridge.

   A solid rectangle indicates an external bridge. A dotted rectangle indicates an internal bridge.

   » Click a bond to view the settings of the bond.

4. If you want additional information about the host, click **Go to VM Details**.
   The **Host** table view on the **Hardware** page is displayed.

   Tip: You can return to the visualizer by pressing the back button in your browser.

Viewing Switch Information

In the visualizer, you can view both switch-level information and interface-level information.

Viewing Switch-Level Information

Switch-level information includes details such as switch name, vendor, and management IP address.

About this task

To view switch-level information, do the following:

Procedure

1. In the topology view, click the name of a switch.
   The switch information window is displayed.
2. To view additional switch information, click **Go to Switch Details**.
   The **Switch** view on the **Hardware** page is displayed.

   Tip: You can return to the visualizer by pressing the back button in your browser.

Viewing Interface-Level Information

Interface-level information includes details such as interface name, physical address, and interface type. The visualizer also displays statistics for each interface.

About this task

To view interface-level information, do the following:

Procedure

1. Click a switch interface.
   The switch port information window is displayed. The window shows both network settings and interface statistics.

   For information about the statistics that are displayed for switch interfaces, see Switch Tab on page 236.
2. Optionally, to show the value at any given point in time in a graph, point to the location on the graph.

5

HARDWARE MANAGEMENT

A Nutanix block is a 1U or 2U chassis that contains one to four independent nodes, each optimized for high-performance compute, memory, and storage. Each node includes CPUs, DIMMs, SSD and HDD storage, fans, network connectors (multiple 10 and 1GbE ports), and other typical hardware components. Each block also includes dual power supplies. A Nutanix cluster can contain an unlimited number of nodes (and blocks), and you can expand the cluster size at any time.

- You can monitor hardware configurations and status across the cluster through the web console (see Hardware Dashboard on page 219).
- You can expand the cluster through the web console (see Expanding a Cluster on page 243).

Figure 140: Nutanix Block

Hardware Dashboard

The Hardware dashboard displays dynamically updated information about the hardware configuration in a cluster. To view the Hardware dashboard, select **Hardware** from the pull-down list on the far left of the main menu.

Menu Options

In addition to the main menu (see Main Menu on page 46), the Hardware screen includes a menu bar with the following options:

- View selector. The Hardware dashboard provides three viewing modes.
  - Click the **Overview** button on the left to display hardware information in a summary view (see Hardware Overview View on page 221).
  - Click the **Diagram** button to display a diagram of the cluster nodes from which you get detailed hardware information by clicking on a component of interest (see Hardware Diagram View on page 223).
  - Click the **Table** button to display hardware information in a tabular form. The Table screen is further divided into host, disk, and switch views; click the **Host** tab to view host information, the **Disk** tab to view disk information, or the **Switch** tab to view switch information (see Hardware Table View on page 231).
- Expand cluster button. Click the **Expand cluster** button on the right to add nodes to the cluster (see Expanding a Cluster on page 243).
- Repair Host Boot Disk. Click the Repair Host Boot Disk button on the right to repair the boot drive of your hosts.
- Page selector. In the Table view, hosts and disks are listed 10 per page. When there are more than 10 items in the list, left and right paging arrows appear on the right, along with the total count and the count for the current page.
- Export table content. In the Table view, you can export the table information to a file in either CSV or JSON format by clicking the gear icon

on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.

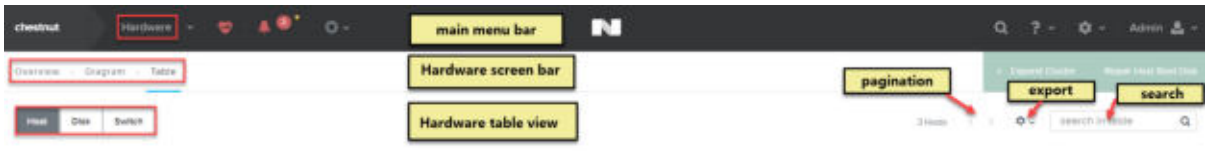

Figure 141: Hardware Dashboard Menu

Hardware Overview View

The Hardware Overview view displays hardware-specific performance and usage statistics on the left plus the most recent hardware-specific alert and event messages on the right. 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.

The following figure is a sample view, and the table describes each field in this view.

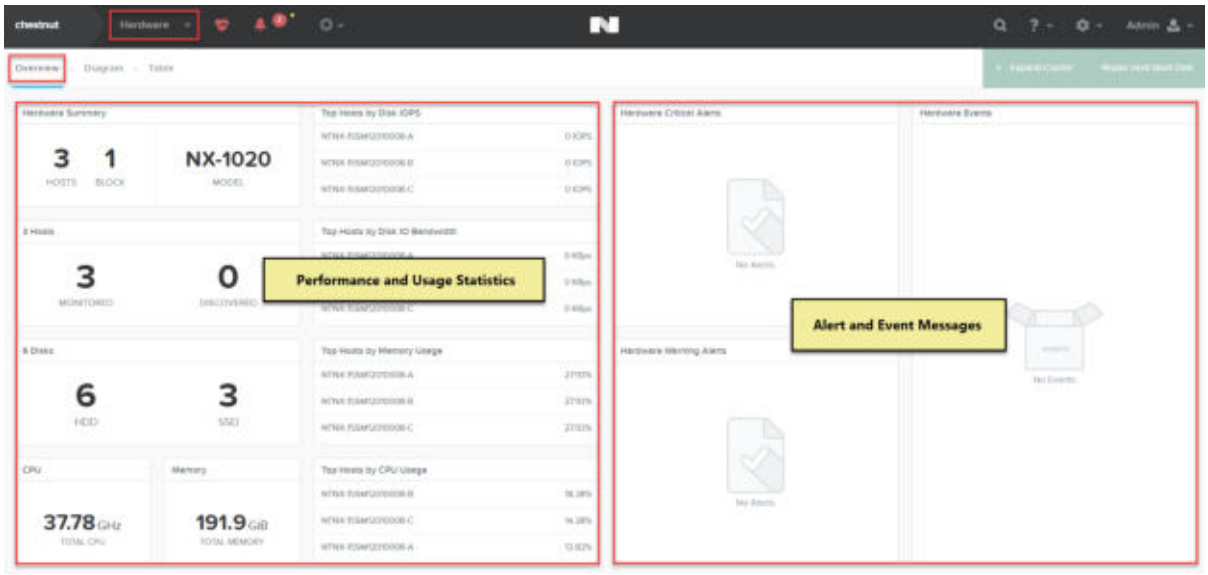

Figure 142: Hardware Overview View

Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

Table 30: Hardware Overview View Fields

Name	Description
Hardware Summary	Displays the number of hosts and blocks in the cluster. It also dispalys the Nutanix model number.
Hosts	Displays the number of hosts in the cluster broken down by on, off, and suspended states. It also displays the number of discovered nodes that have not yet been added to the cluster.
Disks	Displays the total number of disks in the cluster broken down by tier type (SSD-PCIe, SSD-SATA, DAS-SATA).
CPU	Displays the total amount of CPU capacity (in GHz) in the cluster.
Memory	Displays the total amount of memory (in GBs) in the cluster.
Top Hosts by Disk IOPS	Displays I/O operations per host for the 10 most active hosts.
Top Hosts by Disk IO Bandwidth	Displays I/O bandwidth used per host for the 10 most active hosts. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume.
Top Hosts by Memory Usage	Displays the percentage of memory capacity used per host for the 10 most active hosts.
Top Hosts by CPU Usage	Displays the percentage of CPU capacity used per host for the 10 most active hosts.

Name	Description
Hardware Critical Alerts	Displays the five most recent unresolved hardware-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the **view all alerts** button at the bottom of the list (see Alerts Dashboard).
Hardware Warning Alerts	Displays the five most recent unresolved hardware-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the **view all alerts** button at the bottom of the list.
Hardware Events	Displays the ten most recent hardware-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the **view all events** button at the bottom of the list.

Hardware Diagram View

The Hardware Diagram view displays information about hosts and disks. The displayed information is dynamically updated to remain current.

The Hardware Diagram view screen is divided into two sections:

- The top section is an interactive diagram of the cluster blocks. Clicking on a disk or host (node) in the cluster diagram displays information about that disk or host in the summary section.
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

> Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

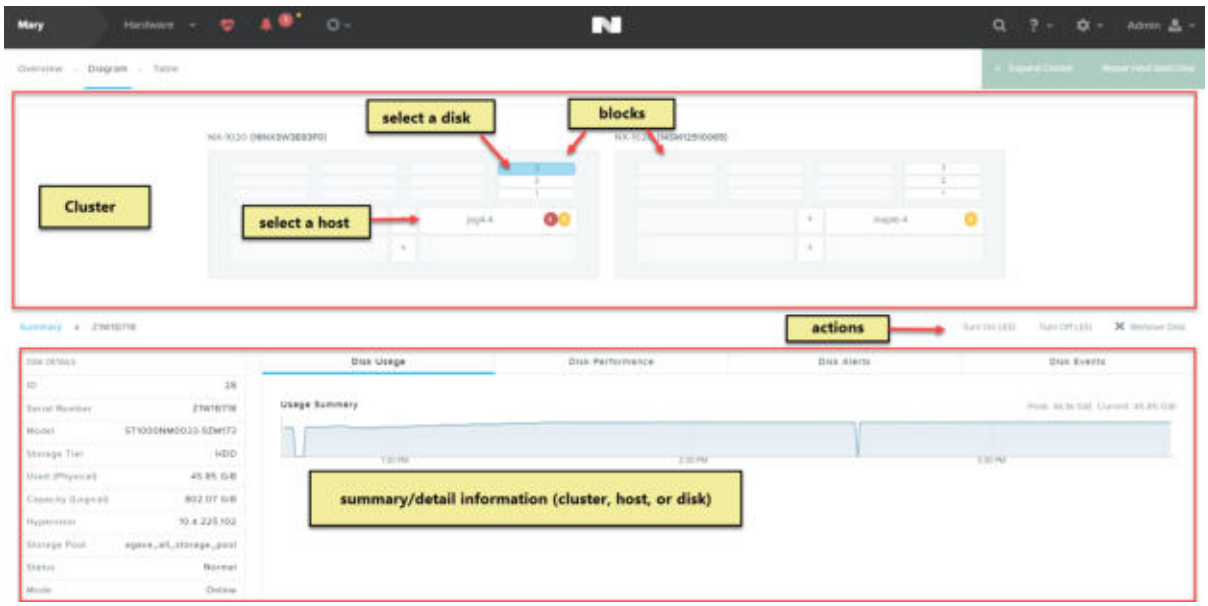

Figure 143: Hardware Diagram View

Host Details

Selecting a host in the diagram displays information about that host in the lower section of the screen.

- When a host is selected, **Summary:** ***host_name*** appears below the diagram, and action links appear on the right of this line:
  - Click the **Turn On LED** link to light up the host LED light on the chassis.
  - Click the **Turn Off LED** link to turn off the host LED light on the chassis.
  - Click the **Remove Host** link to remove this host from the cluster.
- Five tabs appear that display information about the selected host (see following sections for details about each tab): **Host Performance**, **Host Usage**, **Host Alerts**, **Host Events**, **Host NICs**.

Figure 144: Hardware Diagram View: Hosts

Table 31: Host Detail Fields

Parameter	Description	Values
Host Name	Displays the name of the host.	(host name)
Hypervisor IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
Controller VM IP	Displays the IP address assigned to the Controller VM.	(IP address)
IPMI IP	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console. This field does not appear in Prism Central.	(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)
Block Serial	Displays the block serial number.	(block serial number)

Parameter	Description	Values
Block Model	Displays the block model number.	(model series number)
Storage Capacity	Displays the total amount of storage capacity on this host.	xxx [GB|TB]
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)
Memory	Displays the total memory capacity for this host.	xxx [MB|GB]
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CPU Model	Displays the CPU model name	(CPU model name)
No. of CPU Cores	Displays the number of CPU cores on this host.	(number of CPU cores)
No. of Sockets	Displays the number of sockets.	(number of sockets)
No. of VMs	Displays the number of VMs running on this host.	(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 oglog, and you can change the size through the nCLI.	xxx [GB]
Monitored	Displays whether the host is high availability (HA) protected. A **true** value means HA is active for this host. A **false** 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 **true**. A **false** value is likely a sign of a problem situation that should be investigated.	[true|false]
Hypervisor	Displays the name and version number of the hypervisor running on this host.	(name and version #)
Datastores	Displays the names of any datastores. This field does not appear in Prism Central.	(name)

Disk Details

Selecting a disk in the diagram displays information about that disk in the lower section of the screen.

- When a disk is selected, **Summary:** ***disk_name*** appears below the diagram, and action links appear on the right of this line:
  - Click the **Turn On LED** link to light up the LED light on the disk.
  - Click the **Turn Off LED** link to turn off the LED light on the disk.
  - Click the **Remove Disk** link to remove this disk from the cluster.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Disk Usage**, **Disk Performance**, **Disk Alerts**, **Disk Events**.

Figure 145: Storage Diagram View: Disks

Table 32: Disk Detail Fields

Parameter	Description	Values
ID	Displays the disk identification number.	(ID number)
Serial Number	Displays the disk serial number.	(serial number)
Model	Displays the disk model number.	(model number)
Storage 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]
Used (Physical)	Displays the amount of used space on the drive.	xxx [GB|TB]
Capacity (Physical)	Displays the total physical space on the drive.	xxx [GB|TB]
Hypervisor	Displays the IP address of the hypervisor controlling the disk.	(IP address)
Storage Pool	Displays the name of the storage pool in which the disk resides.	(name)

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
Mode	Displays whether the disk is currently online or offline.	[online|offline]
Self Encryption Drive	Displays whether this is a self encrypting drive (SED).	[present|not present]
Password Protection Mode [SED only]	Displays whether data-at-rest encryption is enabled for the cluster. When it is enabled, a key is required to access (read or write) data on the drive. This field appears only when the drive is a SED.	[protected|not protected]

Cluster Summary Information

When a host or disk is not selected in the diagram (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Hardware Summary** column (on the left) includes the following fields:
  - **Blocks**. Displays the number of blocks in the cluster.
  - **Hosts**. Displays the number of hosts in the cluster.
  - **Total Memory**. Displays the total memory capacity (GBs) in the cluster.
  - **Total CPU Capacity**. Displays the total CPU capacity (GHz) in the cluster.
  - **Disks**. Displays the number of disks in each storage tier (DAS-SATA, SSD-SATA, and SSD-PCIe) in the cluster. Tier types vary depending on the Nutanix model type.
  - **Network Switches**. Displays the number of network switches being used in the cluster.
  - **GPUs**. (AHV only) Comma-separated list of GPUs installed on the host. GPU information includes the model name and a count in parentheses if multiple GPUs of the same type are installed on the host. If the firmware on the GPU is in compute mode, the string "compute" is appended to the model name. No string is appended if the GPU is in graphics mode.

    The field is hidden if no GPUs are configured or if the hypervisor is not AHV.
- Three tabs appear that display cluster-wide information (see following sections for details about each tab): **Performance Summary**, **Hardware Alerts**, **Hardware Events**.

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label and number of graphs varies depending on what is selected in the diagram:

- **Performance Summary** (no host or disk selected). Displays resource performance statistics (CPU, memory, and disk) across the cluster.
- **Host Performance** (host selected). Displays resource performance statistics (CPU, memory, and disk) for the selected host.
- **Disk Performance** (disk selected). Displays disk performance statistics for the selected disk.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see Analysis Dashboard on page 396). The Performance tab includes the following graphs:

- **[Cluster-wide] CPU Usage**: Displays the percentage of CPU capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] Memory Usage**: Displays the percentage of memory capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] IOPS**: Displays I/O operations per second (IOPS) for the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Latency**: Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected host, or selected disk.

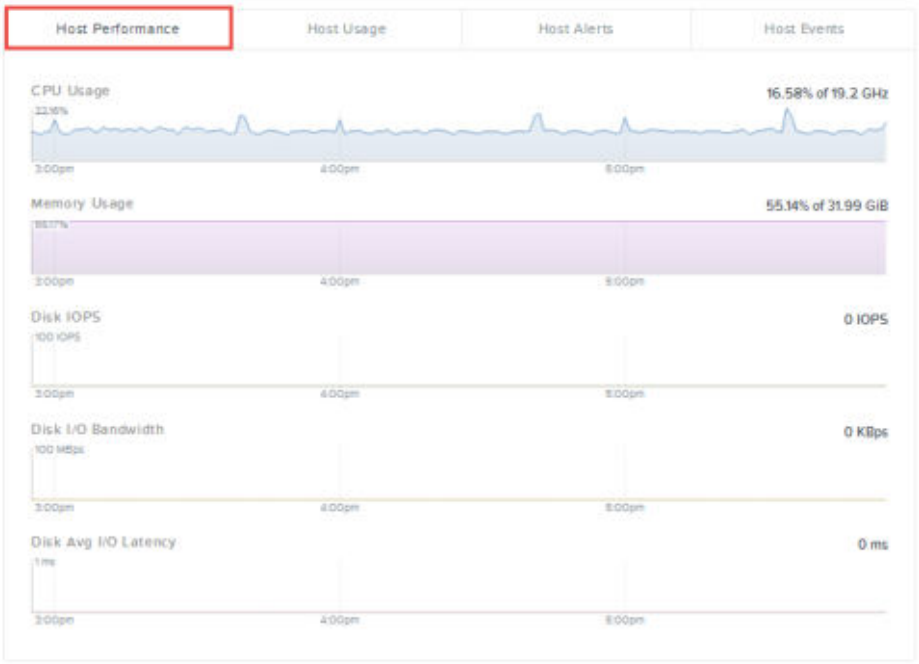

Figure 146: Performance Tab

Usage Tab

The Usage tab displays graphs of storage usage. This tab appears only when a host or disk is selected. The tab label varies depending on what is selected in the diagram:

- **Host Usage** (host selected). Displays usage statistics for the selected host.
- **Disk Usage** (disk selected). Displays usage statistics for the selected disk.

The Usage tab displays one or both of the following graphs:

- **Usage Summary**: Displays a rolling time interval usage monitor 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 (see Analysis Dashboard on page 396).
- **Tier-wise Usage** (host only): Displays a pie chart divided into the percentage of storage space used by each disk tier on the host. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.

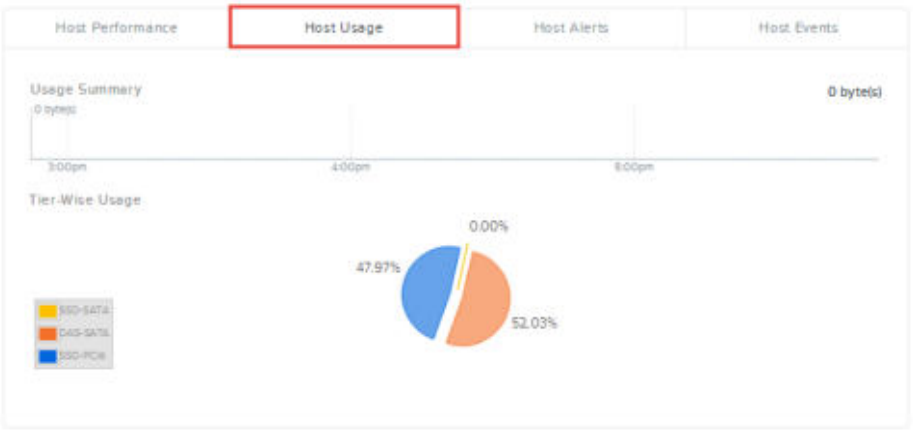

Figure 147: Usage Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about hosts, disks, and other hardware in the same form as the Alerts page (see Alerts Summary View on page 412). Click the **Unresolved X** button in the filter field to also display resolved alerts.

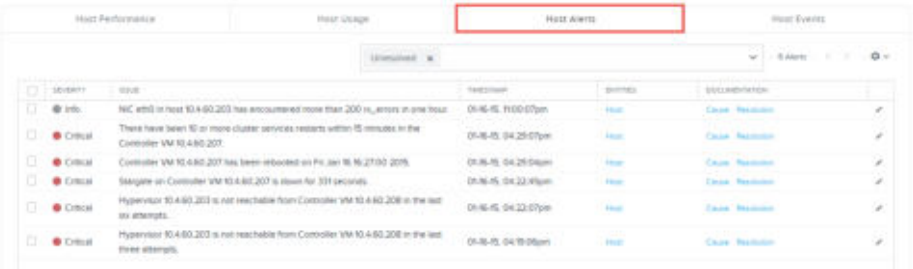

Figure 148: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about hosts, disks, and other hardware in the same form as the Events page (see Events Summary View on page 418). Click the **Include Acknowledged** button to also display acknowledged events.

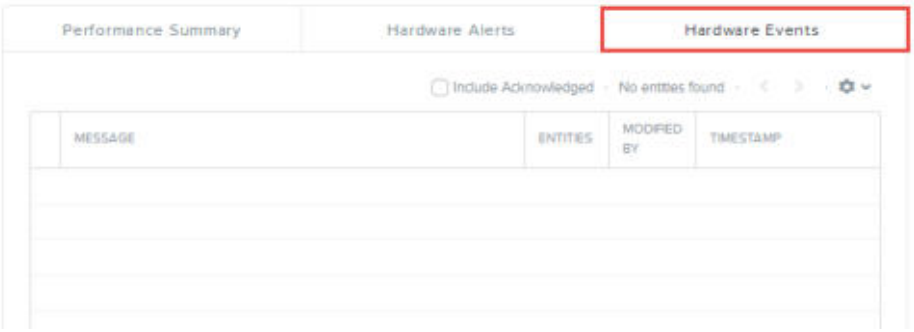

Figure 149: Events Tab

Host NICs Tab

The Host NICs tab displays information in tabular form about the host NICs used to support traffic through the virtual NICs. (This tab appears only when a host is selected and the hypervisor is ESXi.) Each line represent a host NIC, and the following information is displayed for each NIC:

- **Host NIC**. Displays the host NIC name.
- **Speed (in KBps)**. Displays the host NIC transmission speed.
- **MAC Address**. Displays the host NIC MAC address.
- **Received Packets**. Displays the number of packets received by the host NIC.
- **Transmitted Packets**. Displays the number of packets transmitted by the host NIC.
- **Dropped Rx Packets**. Displays the number of received packets dropped by the host NIC.
- **Dropped Tx Packets**. Displays the number of transmitted packets dropped by the host NIC.
- **Rx Rate**. Displays the host NIC received packets rate.
- **Tx Rate**. Displays the host NIC transmitted packets rate.
- **Network Usage Rate**. Displays the host NIC total network usage rate.

When you click a host NIC entry, a set of usage graphs about that NIC appear below the table:

- **Total Packets Received**. Displays a monitor of the total packets received over time. Place the cursor anywhere on the line to see the rate for that point in time. (This applies to all the monitors in this tab.)
- **Total Packets Transmitted**. Displays a monitor of the total packets that were transmitted.
- **Dropped Packets Received**. Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted**. Displays a monitor of transmitted packets that were dropped.
- **Error Packets Received**. Displays a monitor of error packets that were received.

Figure 150: Host NICs Tab

Hardware Table View

The Hardware Table view displays information about hosts and disks in a tabular form. Click the **Host** tab in the screen menu bar to display host information; click the **Disk** tab to display disk information; click the **Switch** tab to display switch information. The displayed information is dynamically updated to remain current.

The Hardware table view is divided into two sections:

- The top section is a table. Each row represents a single host or disk and includes basic information about that host or disk. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

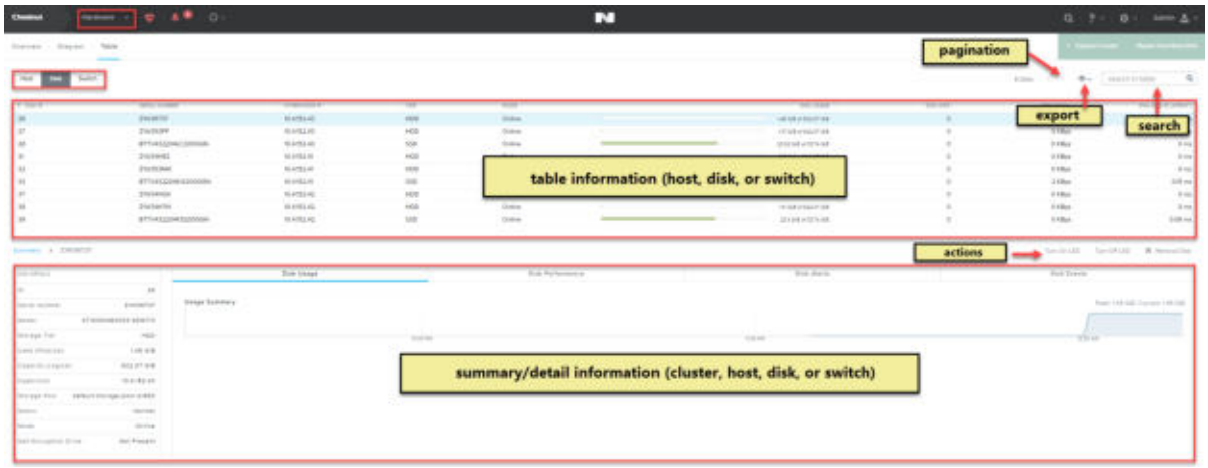

Figure 151: Hardware Table Screen

Host Tab

Clicking the **Host** tab displays information about hosts in the cluster.

- The table at the top of the screen displays information about all the hosts, and the details column (lower left) displays additional information when a host is selected in the table. The following table describes the fields in the host table and detail column.
- When a host is selected, **Summary:** ***host_name*** appears below the table, and action links appear on the right of this line:
  - Click the **Turn On LED** link to light up the host LED light on the chassis.
  - Click the **Turn Off LED** link to turn off the host LED light on the chassis.
  - Click the **Repair Host Boot Device** link to repair the boot drive of the selected host.
  - Click the **Remove Host** link to remove this host from the cluster.
- Five tabs appear that display information about the selected host (see following sections for details about each tab): **Host Performance**, **Host Usage**, **Host Alerts**, **Host Events**, **Host NICs**.

Figure 152: Hardware Table View: Host Tab

Table 33: Host Table and Detail Fields

Parameter	Description	Values
Host Table Fields (upper screen)		
Host Name	Displays the name of the host.	(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.	[ESXi|AHV|Hyper-V]
CPU Usage	Displays the percentage of CPU capacity currently being used.	0 - 100%
CPU Capacity	Displays the CPU capacity of this host.	xxx [GHz]
Memory Usage	Displays the percentage of memory capacity currently being used by this host.	0 -100%

Parameter	Description	Values
Memory Capacity	Displays the memory capacity of this host.	xxx [MiB|GiB]
Total Disk Usage	Displays the storage space used and total disk capacity of this host.	xxx [MiB|GiB|TiB] of xxx [MiB|GiB|TiB]
Disk IOPS	Displays I/O operations per second (IOPS) for this host.	[0 - unlimited]
Disk IO B/W	Displays I/O bandwidth used per second for this host.	xxx [MBps|KBps]
Disk IO Latency	Displays the average I/O latency (in milliseconds) for this host.	xxx [ms]
Host Detail Fields (lower screen)		
Host Name	Displays the name of the host.	(host name)
Hypervisor IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
Controller VM IP	Displays the IP address assigned to the Controller VM.	(IP address)
IPMI IP	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console. This field does not appear in Prism Central.	(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)
Block Serial	Displays the block serial number.	(block serial number)
Block Model	Displays the block model number.	(model series number)
Storage Capacity	Displays the total amount of storage capacity on this host.	xxx [GB|TB]
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)
Memory	Displays the total memory capacity for this host.	xxx [MB|GB]
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CPU Model	Displays the CPU model name	(CPU model name)
No. of CPU Cores	Displays the number of CPU cores on this host.	(number of CPU cores)
No. of Sockets	Displays the number of sockets.	(number of sockets)
No. of VMs	Displays the number of VMs running on this host.	(number)

Parameter	Description	Values
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on every SSD.	[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 every SSD is reserved for the oplog.	xxx [GB]
Monitored	Displays whether the host is high availability (HA) protected. A **true** value means HA is active for this host. A **false** 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 **true**. A **false** value is likely a sign of a problem situation that should be investigated.	[true|false]
Hypervisor	Displays the name and version number of the hypervisor running on this host.	(name and version #)
Datastores	Displays the names of any datastores.	(name)

Disk Tab

Clicking the **Disk** tab displays information about disks in the cluster.

- The table at the top of the screen displays information about all the disks, and the details column (lower left) displays additional information when a disk is selected in the table. The following table describes the fields in the disk table and detail column.
- When a disk is selected, **Summary:** ***disk_name*** appears below the table, and action links appear on the right of this line:
  - Click the **Turn On LED** link to light up the LED light on the disk.
  - Click the **Turn Off LED** link to turn off the LED light on the disk.
  - Click the **Remove Disk** link to remove this disk from the cluster.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Disk Usage**, **Disk Performance**, **Disk Alerts**, **Disk Events**.

Figure 153: Hardware Table View: Disk Tab

Table 34: Disk Table and Detail Fields

Parameter	Description	Values
Disk Table Fields (upper screen)		
Disk ID	Displays the disk identification number.	(ID number)
Serial Number	Displays the disk serial number.	(serial number)
Hypervisor IP	Displays the IP address assigned to the hypervisor running on 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 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]
Disk IOPS	Displays I/O operations per second (IOPS) for this disk.	[0 - unlimited]
Disk IO B/W	Displays I/O bandwidth used per second for this disk.	xxx [MBps|KBps]
Disk Avg IO Latency	Displays the average I/O latency for this disk.	xxx [ms]
Disk Detail Fields (lower screen)		
ID	Displays the disk identification number.	(ID number)
Serial Number	Displays the disk serial number.	(serial number)
Model	Displays the disk model number.	

Parameter	Description	Values
Storage 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]
Used (Physical)	Displays the amount of used space on the drive.	xxx [GB|TB]
Capacity (Logical)	Displays the total physical space on the drive.	xxx [GB|TB]
Hypervisor	Displays the IP address of the hypervisor controlling the disk.	(IP address)
Storage Pool	Displays the name of the storage pool in which the disk resides.	(name)
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
Mode	Displays whether the disk is currently online or offline.	[online|offline]
Self Encryption Drive	Displays whether this is a self encrypting drive (SED).	[present|not present]
Password Protection Mode [SED only]	Displays whether data-at-rest encryption is enabled for the cluster. When it is enabled, a key is required to access (read or write) data on the drive. This field appears only when the drive is a SED.	[protected|not protected]

Switch Tab

Clicking the **Switch** tab displays information about the physical switches used by the host NICs to support traffic through the virtual NICs. The table at the top of the screen displays information about the switches, and the lower portion of the screen displays additional information when a switch is selected in the table. You can configure any number of switches (see Configuring Network Switch Information on page 202), but only the switches that are actually being used for virtual NIC traffic appear in this table. The following table describes the fields in the switch table, in the detail column (lower left), and in the Physical Switch Interfaces tab (lower right).

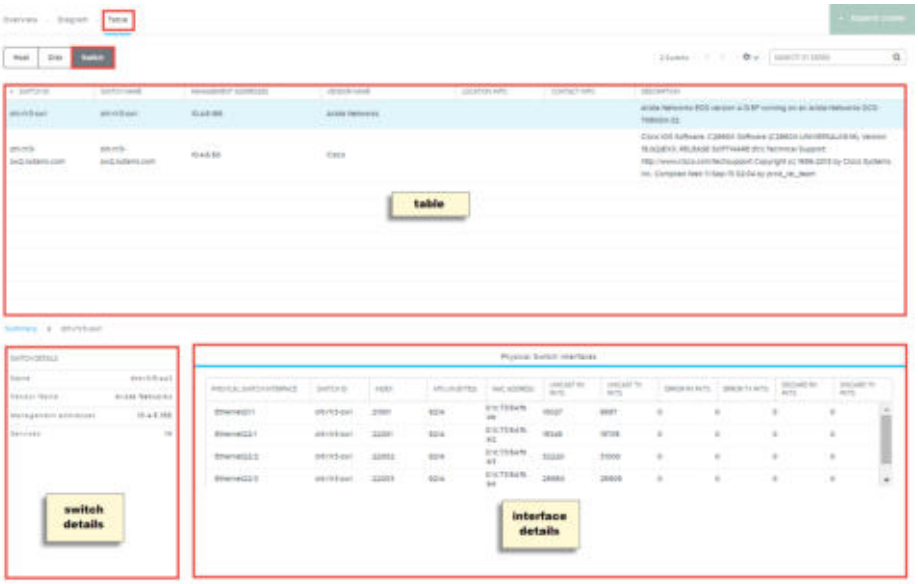

Figure 154: Hardware Table View: Switch Tab

Table 35: Switch Table and Detail Fields

Parameter	Description	Values
Switch Table Fields (upper screen)		
Switch ID	Displays the switch identification number.	(ID value)
Switch Name	Displays the switch name.	(name)
Management Addresses	Displays the switch management IP address(es).	(IP address)
Vendor Name	Displays the name of the switch vendor.	(company name)
Location Info	Displays the switch vendor location.	(company address)
Contact Info	Displays the switch vendor contact information.	(company contact)
Description	Describes the switch model and type.	(switch description)
Switch Detail Fields (lower left screen)		
Name	Displays the switch name.	(name)
Vendor Name	Displays the name of the switch vendor.	(company name)
Management Addresses	Displays the IP address(es) for the switch management ports.	(IP address)
Services	Displays the number of services being used.	(number)
Physical Switch Interfaces Fields (lower right screen)		
Physical Switch Interface	Displays the interface name.	(name)

Parameter	Description	Values
Switch ID	Displays the switch identification number.	(ID value)
Index	Displays the index value.	(number)
MTU (in bytes)	Displays the size in bytes of the largest protocol data unit (maximum transmission unit) that the layer can pass onwards.	(number)
MAC Address	Displays the interface MAC address	(mac address)
Unicast Rx Pkts	Displays the number of unicast packets received.	(number)
Unicast Tx Pkts	Displays the number of unicast packets transmitted.	(number)
Error Rx Pkts	Displays the number of received packets with an error.	(number)
Error Tx Pkts	Displays the number of transmitted packets with an error.	(number)
Discard Rx Pkts	Displays the number of received packets that were discarded.	(number)
Discard Tx Pkts	Displays the number of transmitted packets that were discarded.	(number)

When you click a physical switch interface entry, usage graphs about that interface appear below the table:

- **Unicast Packets Received**. Displays a monitor of the received unicast packets over time. Place the cursor anywhere on the line to see the value for that point in time. (This applies to all the monitors in this tab.)
- **Unicast Packets Transmitted**. Displays a monitor of the transmitted unicast packets.
- **Error Packets Received**. Displays a monitor of error packets received.
- **Dropped Packets Received**. Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted**. Displays a monitor of transmitted packets that were dropped.

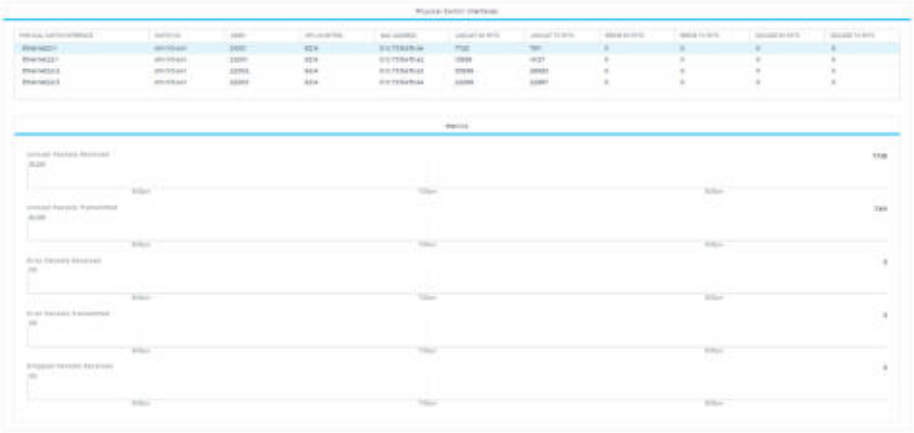

Figure 155: Physical Switch Interface Monitors

Cluster Summary Information

When a host or disk is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Hardware Summary** column (on the left) includes the following fields:
  - **Blocks**. Displays the number of blocks in the cluster.
  - **Hosts**. Displays the number of hosts in the cluster.
  - **Total Memory**. Displays the total memory capacity (GBs) in the cluster.
  - **Total CPU Capacity**. Displays the total CPU capacity (GHz) in the cluster.
  - **Disks**. Displays the number of disks in each storage tier (DAS-SATA, SSD-SATA, and SSD-PCIe) in the cluster. Tier types vary depending on the Nutanix model type.
  - **Network Switches**. Displays the number of network switches being used in the cluster.
  - **GPUs**. (AHV only) Comma-separated list of GPUs installed on the host. GPU information includes the model name and a count in parentheses if multiple GPUs of the same type are installed on the host. If the firmware on the GPU is in compute mode, the string "compute" is appended to the model name. No string is appended if the GPU is in graphics mode.

    The field is hidden if no GPUs are configured or if the hypervisor is not AHV.
- Three tabs appear that display cluster-wide information (see following sections for details about each tab): **Performance Summary**, **Hardware Alerts**, **Hardware Events**.

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label and number of graphs varies depending on what is selected in the table:

- **Performance Summary** (no host or disk selected). Displays resource performance statistics (CPU, memory, and disk) across the cluster.
- **Host Performance** (host selected). Displays resource performance statistics (CPU, memory, and disk) for the selected host.
- **Disk Performance** (disk selected). Displays disk performance statistics for the selected disk.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see Analysis Dashboard on page 396). The Performance tab includes the following graphs:

- **[Cluster-wide] CPU Usage**: Displays the percentage of CPU capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] Memory Usage**: Displays the percentage of memory capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] IOPS**: Displays I/O operations per second (IOPS) for the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Latency**: Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected host, or selected disk.

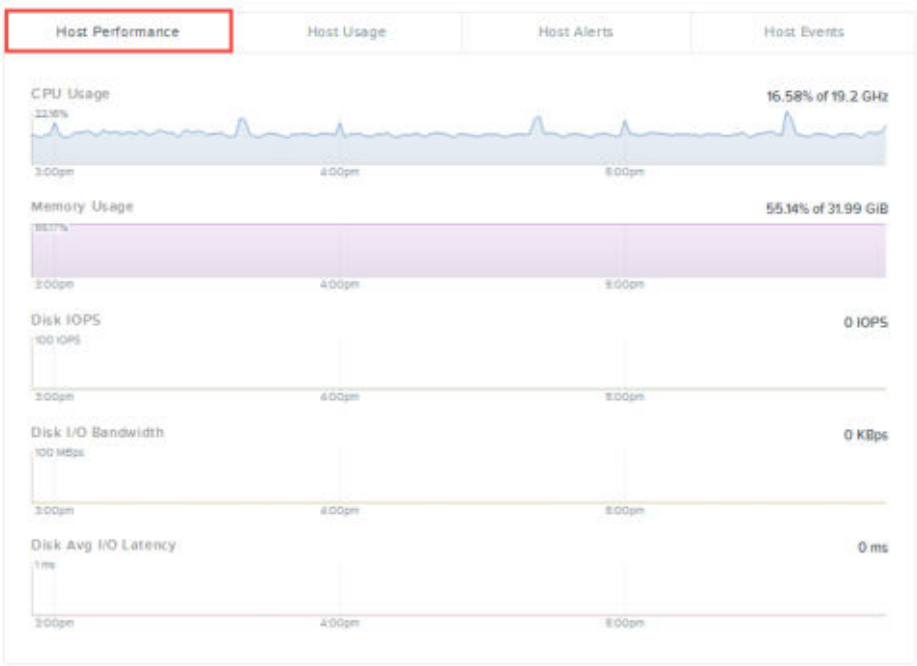

Figure 156: Hardware Table View: Performance Tab

Usage Tab

The Usage tab displays graphs of storage usage. This tab appears only when a host or disk is selected. The tab label varies depending on what is selected in the table:

- **Host Usage** (host selected). Displays usage statistics for the selected host.
- **Disk Usage** (disk selected). Displays usage statistics for the selected disk.

The Usage tab displays one or both of the following graphs:

- **Usage Summary**: Displays a rolling time interval usage monitor 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 (see Analysis Dashboard on page 396).
- **Tier-wise Usage** (host only): Displays a pie chart divided into the percentage of storage space used by each disk tier on the host. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.

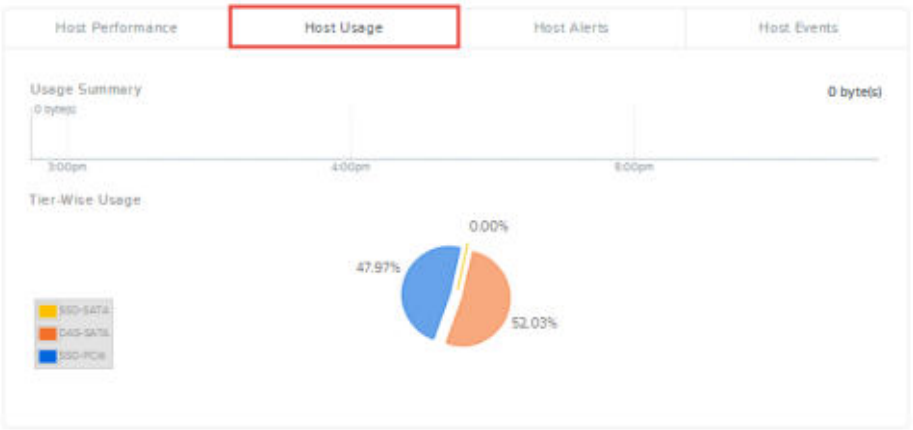

Figure 157: Hardware Table View: Usage Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about hosts, disks, and other hardware in the same form as the Alerts page (see Alerts Summary View on page 412). Click the **Unresolved X** button in the filter field to also display resolved alerts.

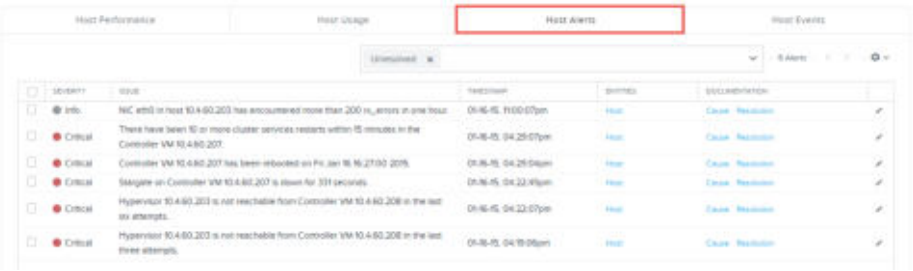

Figure 158: Hardware Table View: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about hosts, disks, and other hardware in the same form as the Events page (see Events Summary View on page 418). Click the **Include Acknowledged** button to also display acknowledged events.

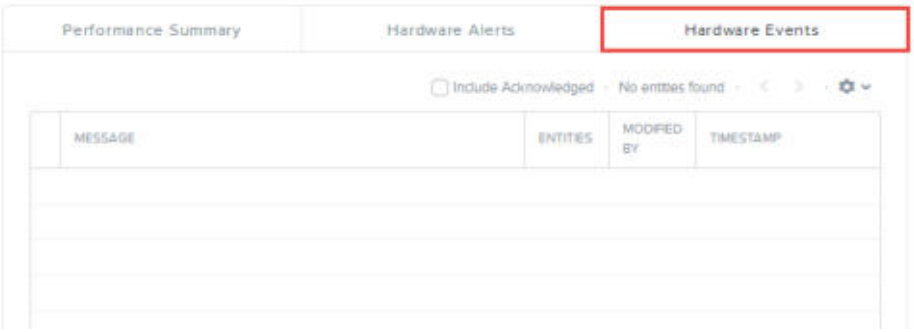

Figure 159: Hardware Table View: Events Tab

Host NICs Tab

The Host NICs tab displays information in tabular form about the host NICs used to support traffic through the virtual NICs. (This tab appears only when a host is selected.) Each line represents a host NIC, and the following information is displayed for each NIC:

- **Host NIC**. Displays the host NIC name.
- **Speed (in KBps)**. Displays the host NIC transmission speed.
- **MAC Address**. Displays the host NIC MAC address.
- **Received Packets**. Displays the number of packets received by the host NIC.
- **Transmitted Packets**. Displays the number of packets transmitted by the host NIC.
- **Dropped Rx Packets**. Displays the number of received packets dropped by the host NIC.
- **Dropped Tx Packets**. Displays the number of transmitted packets dropped by the host NIC.
- **Rx Packet Errors**. Displays the number of error packets received by the host NIC.
- **Tx Packet Errors**. Displays the number of error packets transmitted by the host NIC.

When you click a host NIC entry, a set of usage graphs about that NIC appear below the table:

- **Total Packets Received**. Displays a monitor of the total packets received by the host NIC (in KBs or MBs) over time. Place the cursor anywhere on the line to see the value for that point in time. (This applies to all the monitors in this tab.)
- **Total Packets Transmitted**. Displays a monitor of the total packets transmitted by the host NIC (in KBs or MBs) .
- **Dropped Packets Received**. Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted**. Displays a monitor of transmitted packets that were dropped.
- **Error Packets Received**. Displays a monitor for error packets received.

Figure 160: Host NICs Tab

Expanding a Cluster

A cluster is a collection of nodes. You can add new nodes to a cluster at any time after physically installing and connecting them to the network on the same subnet as the cluster. The cluster expansion process compares the AOS version on the existing and new nodes and performs any upgrades necessary for all nodes to have the same AOS version.

Before you begin

- Review the relevant sections in Prerequisites and Requirements on page 250 before attempting to add a node to the cluster. The process for adding a node varies depending on several factors. This section covers specific considerations based on your AOS, hypervisor, encryption, and hardware configuration.
- Check the Health Dashboard (see Health Dashboard on page 287). If any health checks are failing, resolve them before adding any nodes. As a final check, run NCC to ensure that the cluster is healthy.
- Allow any current expand cluster operations to complete.
- Ensure that all nodes are in the correct metadata state by checking the Hardware dashboard (see Hardware Dashboard on page 219). If any nodes show Metadata store disabled on the node or Node is removed from metadata store, enable the metadata store by clicking **Enable Metadata Store**.

About this task

To add one or more nodes to an existing cluster (you can add multiple nodes at the same time), do the following:

> Note: Steps 9-12 are for special cases (rack fault tolerance, data-at-rest encryption, Hyper-V, and Citrix Hypervisor), and they are not in proper sequence. Refer to these steps as needed during the procedure for any that apply.

Procedure

1. Do one of the following:

   » Click the gear icon in the main menu and then select **Expand Cluster** in the **Settings** page.

   » Go to the Hardware dashboard and click the **Expand Cluster** button.

2. In the Expand Cluster window, select (click the radio button for) the desired option and then click the **Next** button:

   » Select **Expand Cluster** to begin the expansion immediately (after you complete the remaining configuration steps).

   » Select **Prepare Now and Expand Later** to prepare the nodes now but delay adding them to the cluster until a later time. Preparing the nodes includes imaging the hypervisor (if needed), upgrading the AOS version (if needed), and preparing new node network configuration (if needed).

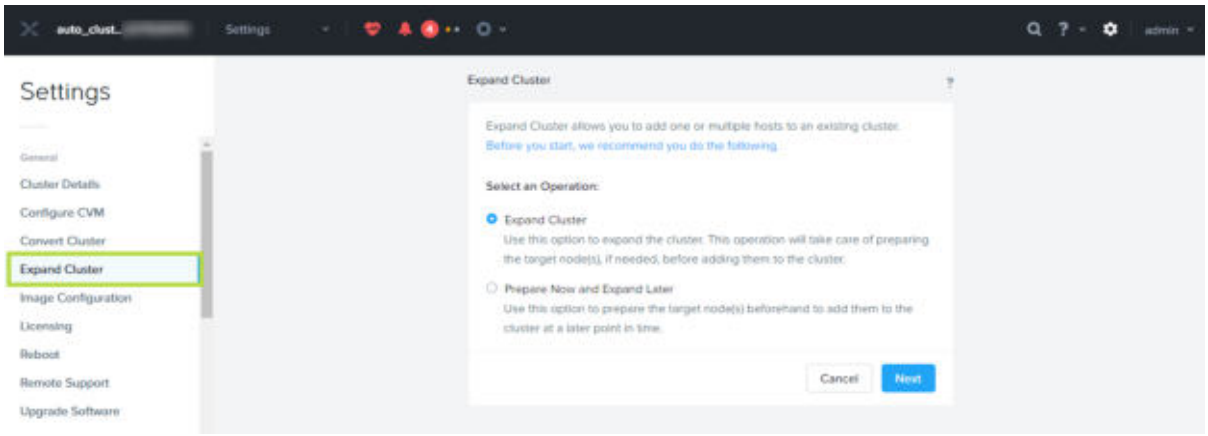

Figure 161: Expand Cluster Window

The network is searched for Nutanix nodes and then the Select Host tab appears with a graphical list of the discovered blocks and nodes. Discovered blocks are blocks with one or more unassigned factory-prepared nodes (hypervisor and Controller VM installed) residing on the same subnet as the cluster. Discovery requires that IPv6 multicast packets are allowed through the physical switch. If you see a failure message, review the requirements in Prerequisites and Requirements on page 250.

3. Do one of the following:

   » If the list includes all desired nodes, go to the next step.

   » If IPv6 is not supported for the cluster or the list does not include all desired nodes, go to the bottom of the list, click the **Discover Hosts Manually** button, and do the following to retry discovery using IPv4:

   1. In the Manual Host Discovery window, click the **+Add Host** link. A line appears in the table below the link.
   2. Enter the Controller VM IP address on that host (or the host IP address for compute-only nodes) and then click **Save**.
   3. Add more hosts as desired.
   4. When the list of hosts is complete, click the **Discover and Add Hosts** button. Manual discovery uses IPv4 to find and add the hosts to the list.

      Note: For an ESXi or Hyper-V cluster, manual host discovery requires that the target node has the same hypervisor type and version. In addition, the AOS version must be lower than or equal to that of the cluster.

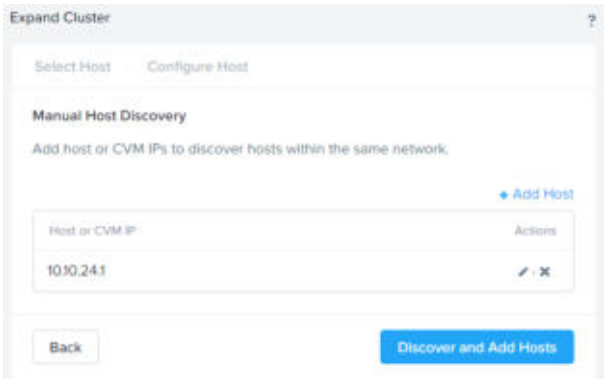

Figure 162: Manual Host Discovery

4. Select the check box for each block to be added to the cluster. All nodes within a checked block are also checked automatically; uncheck any nodes you do not want added to the cluster.

   When a block is checked, more fields appear below the block diagram. A separate line for each node (host) in the block appears under each field name.

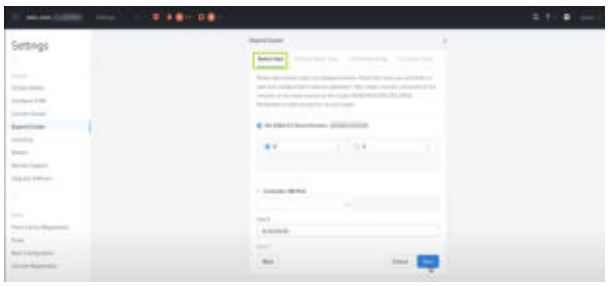

Figure 163: Select Host Tab

5. Do the following in the indicated fields for each checked block:

   a. **Host Name** (if present): Enter the name of the host.

      Enter just the host name, not the fully qualified domain name. The host name is required for Hyper-V clusters, but you can ignore the field for ESXi and AHV clusters.

   b. **Controller VM [IPv4|IPv6]**: Review the Controller VM IP address assigned to each host and do one of the following:

      > Note: For Controller VM, hypervisor, and IPMI, both IPv6 and IPv4 address fields appear to account for either protocol. When entering addresses, be sure to enter each IP address in the correct field for the relevant protocol. Usually you can ignore the IPv6 field and only update the IPv4 address if needed.

      - If the address is correct, do nothing in this field.
      - If the address is not correct, either change the incorrect address or enter a starting address on the top line (for multiple hosts). The entered address is assigned to the Controller VM of the first host, and consecutive IP addresses (sequentially from the entered address) are assigned automatically to the remaining hosts.

   c. **Hypervisor [IPv4|IPv6]**: Repeat the previous step for this field.

      This field sets the hypervisor IP addresses for all the hosts to be added.

   d. **IPMI IP [IPv4|IPv6]**: Repeat the previous step for this field.

      This field sets the IPMI port IP addresses for all the hosts to be added. An IPMI port is used for the hypervisor host console.

   e. When all the node values are correct, click the **Next** button (lower right).

   The network addresses are validated before continuing. If an issue is discovered, the problem addresses are highlighted in red.

6. In the **Choose Node Type** tab, select the node type (**HCI Node** or **Storage-only**) for each node from the pull-down list and then click the **Next** button.

   Selecting **HCI Node** means that the node will be added as a standard node in the cluster. Selecting **Storage-only** means that the node will be added as a storage-only node. Storage-only nodes have specific requirements. Review Storage-only Node Configuration on page 281 before adding a storage-only node.

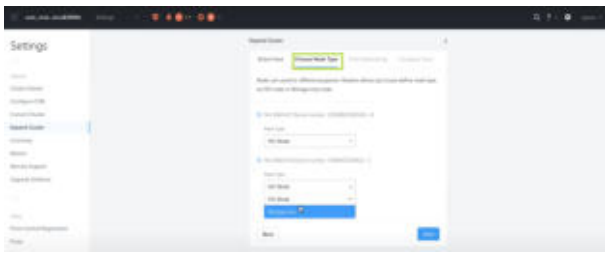

Figure 164: Choose Node Type Tab

7. In the **Host Networking** tab, configure the uplinks for each management bridge or vSwitch to be created or updated for the nodes.

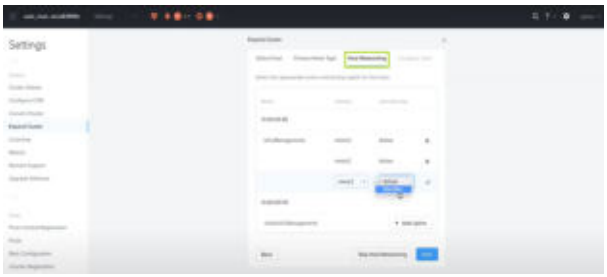

Figure 165: Networking Tab

The **Host Networking** tab displays a list of the target nodes where you can specify the status (active or standby) for each node's uplinks to the management bridge or vSwitch. The uplink list depends on the hypervisor of the node to be added, while the listed bridge or vSwitch is for the base cluster. The standard form for AHV uplinks is ethx (eth0, eth1, eth2, and so on); the standard form for ESXi uplinks is vmnicx (vmnic0, vmnic1, and so on.) For example, if the hypervisor is ESXi on the node but AHV on a three-node base cluster, each entry displays the AHV vSwitch name (br0 or br1), ESXi uplink name (vmnic0, vmnic1, or vmnic2), and status (active or standby). See Prerequisites and Requirements on page 250 for hypervisor-specific details.

> Note: To skip network configuration, click the **Skip Networking** button.

a. Click the **Add Uplink** button for the first node.

b. Select the target link from the pull-down list in the first field.

c. Select **Active** or **Standby** for the link status from the pull-down list in the second field.

d. Repeat for all the links to configure for the node.

e. Repeat steps for each node.

f. After configuring all nodes in the list, click the **Next** button.

8. In the **Configure Host** tab, specify the hypervisor image and whitelist for nodes that require imaging.

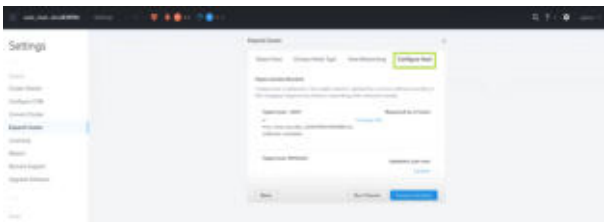

Figure 166: Configure Host Tab

» If the detected hypervisor and AOS version are the same as the cluster, no imaging is required and a message acknowledging that fact appears. Skip to the next step.

» If a hypervisor image is listed in the **Hypervisor:** *<type>* field and it is the desired one, skip to the next step. If you uploaded a hypervisor image when adding nodes previously (and

it is still compatible for imaging the new nodes), that image appears here. You can use that image or upload a different one.

» If no hypervisor image is listed or the listed one is not the desired one, click the **Choose File** (no listed file) or **Change File** (file listed already) button. In the search window, find the image file on your workstation, and then click the **Open** button (in the search window) to upload that image file.

You must provide an ISO file for imaging ESXi or Hyper-V. You can get a compatible AHV image from the **Downloads** > **AHV** page of the Nutanix Support portal (see)Accessing the Nutanix Support Portal on page 804. AOS includes a default AHV installation bundle named kvm_host_bundle_*version#*.tar.gz. If this version is appropriate, use it instead of uploading an installation bundle from the support portal.

» If a message appears that the hypervisor image is not compatible, either select (choose and upload) a hypervisor image that is compatible or update the hypervisor ISO whitelist.

The cluster includes a hypervisor ISO whitelist, which lists the approved hypervisor images. You can only use an image that is in the approved list. See "Hypervisor ISO Images" in the *Field Installation Guide* for more information about the whitelist. If your hypervisor image does not appear in the whitelist (because you created the cluster before the image was approved), you can update the whitelist as follows:

1. Download the latest hypervisor ISO whitelist from the **Downloads** > **Foundation** page of the Nutanix Support portal.
2. Click the **Update** link in the **Hypervisor ISO Whitelist** field.
3. Click the **Choose File** button (which opens a search window), find and select the whitelist file on your workstation, and then click the **Open** button.
4. Click the **Upload** button. The incompatible message disappears and you can continue provided the hypervisor image is on the uploaded whitelist.

9. [rack fault tolerance only] In the **Assign Rack** tab, select the block and the rack (from the **Assign to Rack** drop-down menu) to map the block placement in that rack. To add a rack to the list, click **Create Rack**.

   This tab appears only when rack fault tolerance is enabled. See Configuring Rack Fault Tolerance on page 19 for details.

10. [data-at-rest encryption only] In the **Encrypt Host** tab, do the following.

    This tab appears only when data-at-rest encryption is enabled with an external KMS. See “” () for more information. You can apply encryption licenses after the new nodes are added to the cluster.

    a. In the **Certificate Signing Request Information** field, click the **Generate and Download** link for each node to be added.

       Clicking the link generates a certificate signing request (CSR) named csr_for_discovered_node for the node, which you download to your workstation.

    b. Get the CSRs signed by a certificate authority (CA).

    c. Click the **select files** link for each key management server and upload the signed certificates for the nodes to be added.

    d. Click the **Next** button.

11. [Hyper-V only] Specify the credentials to join the new nodes to Active Directory and to a failover cluster.

    a. Specify the name of the Hyper-V failover cluster in the **Failover Cluster Name** text box.

    b. Specify the user name and password of the domain account that has the privileges to create a new or modify an existing computer account in the Active Directory domain. The user name must be in the DOMAIN\USERNAME format.

12. [Citrix Hypervisor only] Specify the credentials of the leader node of the Citrix Hypervisor pool.

13. When all the fields are correct, click the **Run Checks** button to verify that the nodes are ready.

    This runs a set of precheck tests on the nodes. Progress messages appear in the Expand Cluster window. You can also monitor progress from the Tasks dashboard (see View Task Status on page 747).

14. When the checks pass successfully, do one of the following:

    » Click the **Expand Cluster** button to begin the cluster expansion process. (This button appears if you selected **Expand Cluster** in step 2.)

    » Click the **Prepare Node(s)** button to begin preparing the nodes. (This button appears if you selected **Prepare Now and Expand Later** in step 2.)

    The expand cluster or node preparation process begins. As with the prechecks, progress messages appear in the Expand Cluster window, and you can also monitor progress from the Tasks dashboard. Nodes are processed (upgraded or reimaged as needed) and added in parallel. Adding nodes can take some time. Imaging a node typically takes a half hour or more depending on the hypervisor.

15. If the cluster has multiple storage pools, assign the new storage capacity to a storage pool after the nodes are added successfully:

    > Note: When the cluster has only one storage pool, you can ignore this step because the new storage is added to that storage pool automatically.

    a. Select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Pool** tabs.

    b. Select the target storage pool (upper display) and then click the **update** link.

    c. In the Update Storage Pool window, check the **Use unallocated capacity** box in the **Capacity** line and then click the **Save** button.

    This step adds all the unallocated capacity to the selected storage pool.

    d. Go to the Tasks dashboard and monitor progress until the node addition completes successfully.

    Cluster expansion is not complete until the node is added to the metadata ring.

What to do next

One or more of the following items might apply to the added nodes (see also Prerequisites and Requirements on page 250):

- Nondefault timezones are not updated on added nodes and must be reconfigured manually.

- If the Controller VM password for the cluster was changed from the default, the password on any new nodes gets updated automatically to match the cluster password.
- To check Controller VM memory compatibility after cluster expansion, run the `cvm_same_mem_level_check` and `cvm_memory_check` NCC checks (see Running Checks by Using Web Console on page 295).

  - If the `cvm_same_mem_level_check` result is FAIL, the memory size in the added nodes is not the same as the other Controller VMs in the cluster. If the memory is less than the common size, increase the memory to match.
  - If the `cvm_memory_check` result is FAIL, the Controller VM memory is less than the minimum required for the workload. Increase the memory to (at least) the minimum size.

  To increase the memory, see Increasing the Controller VM Memory Size on page 118. Also, the hypervisor-specific administration guides include the procedure to change (increase) Controller VM memory (*AHV Administration Guide*, *vSphere Administration Guide for Acropolis*, *Hyper-V Administration for Acropolis*, or *Citrix Hypervisor on Nutanix Administration Guide*). For Controller VM memory size recommendations, see the "Controller VM Memory Configurations" section of the *Acropolis Advanced Administration Guide*.

Prerequisites and Requirements

The process for adding a node varies depending on the AOS version, hypervisor host type, data-at-rest encryption status, and certain hardware configuration factors.

AOS Considerations

The following apply to any cluster:

- The expand cluster process does not support compute-only node preparation.
- The expand cluster discovery step (finding nodes to add) requires that IPv6 multicast packets are allowed through the physical switch. Therefore, the expand cluster process does not work when IPv6 is disabled in your network. If IPv6 is not enabled currently, do one of the following:

  - Enable IPv6 in your network and retry the expand cluster operation.
  - If enabling IPv6 is not an option, the expand cluster procedure allows you to enter the IP addresses of nodes to add manually and run discovery using IPv4. This requires that you have the IP addresses before starting the expand cluster operation.
- If the Controller VM memory on a new node is less than the current nodes in the cluster, the expand cluster process increases the memory on the new node to the same base value as the current nodes. The new Controller VM is upgraded to a maximum of 32 GB.

  The Controller VM is upgraded to a maximum of 28 GB for ESXi nodes with 64 GB or less of total physical memory. With total physical memory greater than 64 GB, the existing Controller VM memory is increased by 4 GB.
- A new node is reimaged automatically before being added under certain conditions. The following table describes those conditions.

Table 36: Node Imaging Criteria

Configuration	Description
Same AOS and hypervisor versions	The node is added to the cluster without reimaging it.
Same hypervisor version but different AOS version	The node is automatically reimaged before it is added. However, if the AOS version on the node is higher than the version on the cluster, you can upgrade the cluster to the higher version. If you do not upgrade the base cluster to match the node AOS version, the node is reimaged automatically to match the lower AOS version of the cluster. To upgrade your cluster, see the *Acropolis Upgrade Guide*.
Same AOS version but different hypervisor version	The node is automatically reimaged before it is added.

Note: If you are expanding a cluster (now) on which the network is segmented only by traffic type (management and backplane), see "Network Segmentation During Cluster Expansion" in the Nutanix Security Guide.

AHV Considerations

The following apply to clusters running AHV:

- If the Controller VMs in the cluster reside in a VLAN configured network, the discovery process still finds any factory-prepared nodes regardless of their current VLAN status (configured or not configured).
- Network configuration has the following restrictions and requirements:
  - You cannot migrate management from br0 to other bridges.
  - You can only have the Controller VM management interface (eth0) and hypervisor management interface deployed on the br0 bridge.

ESXi Considerations

The following apply to clusters running ESXi:

- If the ESXi root user password was changed from the default, the expand cluster operation might fail. In this case, reset the ESXi root user password to the default (see KB 1661 for default cluster credentials), and then retry the expand cluster procedure.
- While expanding a Nutanix cluster running NSX enabled ESXi hosts, add the newly imaged node to the Nutanix cluster where the host and CVM management network are configured with a standard vSwitch (see Expanding a Cluster on page 243), and then add the node in NSX manager. Otherwise, the cluster expansion operation fails with the following error.

  ```
  Failed to get VLAN tag of node <MAC Address of the node>
  ```

- If network segmentation is not enabled, expand cluster supports mixed (ESXi + storage-only) node clusters. However, when network segmentation is enabled, you cannot use expand cluster for a mixed cluster.

- Network configuration has the following restrictions and requirements:
  - You cannot configure the network when either a target node or the base cluster has LACP enabled. Prepare LACP nodes offline using a Foundation VM.
  - You cannot migrate management from vSwitch0 to another standard vSwitch.
  - Management interfaces must either be on a VSS or a DVS; a mixed setup is not supported.
  - If on VSS, the Controller VM management interface (eth0) and the hypervisor management interface must be deployed on vSwitch0 and connected to port group VM Network and Management Network, respectively.
  - If on DVS, all Controller VM management interfaces and all host management interfaces must be connected to the same distributed virtual switch and same port group. (However, the Controller VM interfaces can be connected to a different DVS port group than the host management interfaces.)
  - Segmented network interfaces (backplane, volume, DR) must be on same vSwitch type (VSS or DVS) as the management.
  - If network segmentation is enabled on the base cluster and the backplane is deployed on a separate vSwitch than management, you can create vSwitches (VSS or DVS) and prepare the required Controller VM interfaces. (Manual switch configuration is required for expand now but is integrated into the expand later work flow.)
  - If the base cluster is on DVS and you are doing an expand later, you can migrate the target nodes from the default vSwitch0 to that DVS.
- If the Controller VMs in the cluster reside in a VLAN configured network, you must first configure the new nodes in the sane VLAN before attempting to add them. Otherwise, the discovery process does not find these nodes. See the "Configuring the Cluster in a VLAN-Segmented Network" section in the *Acropolis Advanced Setup Guide* for VLAN configuration instructions.
- To expand a cluster (now) configured with DVS for Controller VM external communication, ensure that you do the following:
  - Expand DVS with the new node.
  - Make sure both the host and the CVM are configured with DVS.
  - Make sure that host to CVM and CVM to CVM communications are working.
  - Follow the cluster expansion procedure.
- After adding the new nodes, note the following:
  - The common Nutanix datastores are mounted on the new nodes by default after cluster expansion.
  - The target storage containers must be set to mount on the new hosts. You can check the mount status from the Storage dashboard (see Storage Table View on page 146). Click the **Storage Container** tab, select the target storage container, click the **Update** button, and verify that **Mount on all ESXi Hosts** (or the new hosts are checked in **Mount/Unmount on the following ESXi Hosts**) is selected in the **NFS DATASTORE** field.
  - If an added node has an older processor class than the existing nodes in the cluster, cluster downtime is required to enable EVC (enhanced vMotion compatibility) with the lower feature set as the baseline. For an indication of the processor class of a node,

see the **Block Serial** field in the Hardware dashboard (see Hardware Diagram View on page 223 or Hardware Table View on page 231). For instructions on enabling EVC, see the vSphere documentation.

> CAUTION: If you mix processor classes without enabling EVC, vMotion/live migration of VMs is not supported between processor classes. If you add the host with the newer processor class to vCenter Server before enabling EVC, cluster downtime is required to enable EVC later because all VMs (including the Controller VM) must be shut down.

- Add the new nodes to the appropriate vCenter Server cluster. If an added node has a newer processor class (for example, Haswell) than the existing nodes in the cluster (Ivy Bridge or Sandy Bridge), enable EVC with the lower feature set as the baseline before adding the node to vCenter.
- If you are adding multiple nodes to an existing EVC-enabled vCenter cluster, which requires powering off the Controller VM for each node to complete the addition, add just one node at a time and wait for data resiliency to return to "OK" before adding the next node to vCenter.

  > CAUTION: Adding multiple nodes to vCenter simultaneously can cause a cluster outage when all the Controller VMs are powered off at the same time.

- If you are adding a node to a cluster where HA is enabled with APD and VMCP is enabled, you must enable APD and APD timeout on the new host.
- If you are adding new nodes to vCenter EVC configured cluster, ensure the following requirements.
  - Enabling EVC requires ESXi version 6.0 or above.
  - All ESXi nodes should be on same version and build.

Hyper-V Considerations

The following apply to clusters running Hyper-V:

- Network preparation is not supported. To bypass this step and prepare the nodes otherwise, click the **Skip Networking** button when you get to the **Networking** tab. In addition, imaging to AHV or ESXi is not possible if the node or base cluster is Hyper-V.
- If the Controller VMs in the cluster reside in a VLAN configured network, you must first configure the nodes in the VLAN before attempting to add them. Otherwise, the discovery process does not find these nodes. See the "Configuring the Cluster in a VLAN-Segmented Network" section in the *Acropolis Advanced Setup Guide* for VLAN configuration instructions.
- After adding the new nodes, if you manage your Hyper-V cluster by using Microsoft System Center VM Manager (SCVMM), do not use the Prism web console or Microsoft Failover Cluster Manager to add the new node to the failover cluster. Instead, use SCVMM to add

the node to the Hyper-V cluster and then perform the following steps in the SCVMM user interface.

1. Open the SCVMM user interface.
2. Refresh the cluster in SCVMM. The new node is displayed under the failover cluster in the **Pending** state.
3. Right-click the node and select **Add to host cluster**.
4. Choose a run-as account that has the local administrator permissions on the new node.
5. Click **OK**. The SCVMM agent is installed on the node and file shares are registered to the new node.
6. Update the networking and other settings of the node to match your standard configuration.

For Microsoft Windows Server 2012 R2 deployments, after adding the node to the cluster in SCVMM, ensure that node disks are not added as clustered resources. Open the Failover Cluster Manager and click **Storage** > **Disks** to check.

Data-At-Rest Encryption Considerations

The following apply to clusters with data-at-rest encryption enabled (see “” ()):

- Configure data-at-rest encryption for the new nodes. The new nodes must have self-encrypting disks or software-only encryption.
- If the cluster uses an external key manager, either with self-encrypting drives or software only, then reimaging cannot be done through the expand cluster workflow. In this case, any nodes to add must already have the correct hypervisor and AOS version. Use Foundation to image the nodes (if needed) before attempting to add them to the cluster.
- Adding a node to a cluster with self-encrypting drives (SED) where the added node is running a different hypervisor is not supported. In this case, image the node to the same hypervisor by using Foundation before adding it to the SED cluster. For more information, see KB-4098.

Nutanix Clusters Considerations

The following apply to clusters hosted on a cloud platform:

- When expanding a cluster on AWS, the nodes are not added to the Cassandra ring (wait in a queue) until there are enough nodes to extend the ring. In addition, new nodes are added to the Cassandra ring only when they are do not break domain awareness. (Other services remain unaffected.)

Hardware Considerations

Note the following when it applies to the nodes you are adding:

- Ensure that you use the minimum version of Foundation required by your hardware platform. To determine whether Foundation needs an upgrade for a hardware platform, see the respective system specifications guide. If the nodes you want to include in the cluster are of different models, determine which of their minimum Foundation versions is the most recent version, and then upgrade Foundation on all the nodes to that version.
- When adding partially populated all-SSD nodes to an existing cluster, the minimum number of partially populated all-SSD nodes added must be equal to the maximum RF in the cluster.

- If you are adding a node with a different processor class to the cluster, ensure that there are no running VMs on the node and the host has the following configuration:

  - ESXi: Verify that EVC is enabled on the cluster.
  - Hyper-V: If you want to move the VMs between the nodes, ensure that you have selected the **Migrate to a physical computer with a different processor version** option for each VM by browsing to **Settings** > **Processor** > **Compatibility** in the **Action** pane of the Hyper-V Manager.

    Note: Do not shut down more than one Controller VM at the same time.

- If you expand a cluster by adding a node with older generation hardware to a cluster that was initially created with later generation hardware, power cycle (do not reboot) any guest VMs before migrating them to the added older generation node or before upgrading the cluster.

  Guest VMs are migrated during hypervisor and firmware upgrades (but not AOS upgrades).

  For example, if you are adding a node with G4 Haswell CPUs to a cluster that also has newer G5 nodes with Broadwell CPUs, you must power cycle guest VMs hosted on the G5 nodes before you can migrate the VMs to the node with G4 CPUs. Power cycling the guest VMs enables them to discover G4 processor changes.

  Power cycle guest VMs from the Prism web console VM dashboard. Do not perform a Guest Reboot; a VM power cycle is required in this case.

- If you physically add a node to a block (for example, a single node shipped from Nutanix is placed into an empty slot in an existing chassis), log on to the Controller VM for that node, and update the following parameters in the **/etc/nutanix/factory_config.json** file:

  - rackable_unit_serial: Set it to the same value as the other Controller VMs in the same block.
  - node_position: Set it to the physical location of the node in the block (A, B, C, D).

  After changing the configuration file, restart Genesis with the `genesis restart` command.

Node Maintenance

You are required to gracefully place a node into the maintenance mode or non-operational state for reasons such as making changes to the network configuration of a node, performing manual firmware upgrades or replacements, performing CVM maintenance or any other maintenance operations.

Entering and Exiting Maintenance Mode

You can only place one node at a time in maintenance mode for each cluster. When a host is in maintenance mode, the CVM is placed in maintenance mode as part of the node maintenance operation and any associated RF1 VMs are powered-off. The cluster marks the host as unschedulable so that no new VM instances are created on it. When a node is placed in the maintenance mode from the Prism web console, an attempt is made to evacuate VMs from the host. If the evacuation attempt fails, the host remains in the "entering maintenance mode" state, where it is marked unschedulable, waiting for user remediation.

When a host is placed in the maintenance mode, the CVM is placed in the maintenance mode as part of the node maintenance operation. The non-migratable VMs (for example, pinned or RF1 VMs which have affinity towards a specific node) are powered-off while live migratable or high availability (HA) VMs are moved from the original host to other hosts in the cluster. After

exiting the maintenance mode, all non-migratable guest VMs are powered on again and the live migrated VMs are automatically restored on the original host.

> Note: VMs with CPU passthrough or PCI passthrough, pinned VMs (with host affinity policies), and RF1 VMs are not migrated to other hosts in the cluster when a node undergoes maintenance. Click **View these VMs** link to view the list of VMs that cannot be live-migrated.

See Putting a Node into Maintenance Mode (AHV) on page 257 to place a node under maintenance.

You can also enter or exit a AHV host under maintenance through the CLI.

- To use the CLI method to place a AHV host in maintenance mode, see the *Putting a Node into Maintenance Mode using CLI* section in the *AHV Administration Guide*.

  > Note: Using the CLI method to place a AHV host under maintenance only places the hypervisor under maintenance mode. The CVM is up running in this method. To place the entire node under maintenance, Nutanix recommends using <UI>.

- To use the CLI method to exit a node from the maintenance mode, see the *Exiting a Node from the Maintenance Mode using CLI* section in the *AHV Administration Guide*.

Exiting a Node from Maintenance Mode

See Exiting a Node from the Maintenance Mode (AHV) on page 263 to remove a node from the maintenance mode.

Viewing a Node under Maintenance Mode

See Viewing a Node under the Maintenance Mode on page 261 to view the node under maintenance mode.

UVM Status when Node under Maintenance Mode

See UVM Status when Node under Maintenance Mode on page 265 to view the status of UVMs when a node is undergoing maintenance operations.

Best Practices and Recommendations

Nutanix strongly recommends using the **Enter Maintenance Mode** option on the Prism web console to place a node under maintenance.

Known Issues and Limitations

- The Prism web console enabled maintenance operations (enter and exit node maintenance) are currently supported only on AHV.
- Entering or exiting a node under maintenance from the CLI is not equivalent to entering or exiting the node under maintenance from the Prism Element web console. For example, placing a node under maintenance from the CLI only places the AHV host under maintenance while the CVM continues to remain powered on.
- You cannot exit the node from maintenance mode from Prism Element web console if the node is placed under maintenance mode from CLI. However, you can enter the node maintenance through the Prism Element web console and exit the node maintenance through the CLI.

Putting a Node into Maintenance Mode (AHV)

Before you begin

Check the cluster status and resiliency before putting a node under maintenance. You can also verify the status of the UVMs. See UVM Status when Node under Maintenance Mode on page 265 for more information.

About this task

Perform the following steps to put the node into maintenance mode.

Procedure

1. Login to the Prism Element web console.
2. On the home page, select **Hardware** from the drop-down menu.
3. Go to the **Table** > **Host** view.
4. Select the node which you intend to put under maintenance.

5. Click the **Enter Maintenance Mode** option.

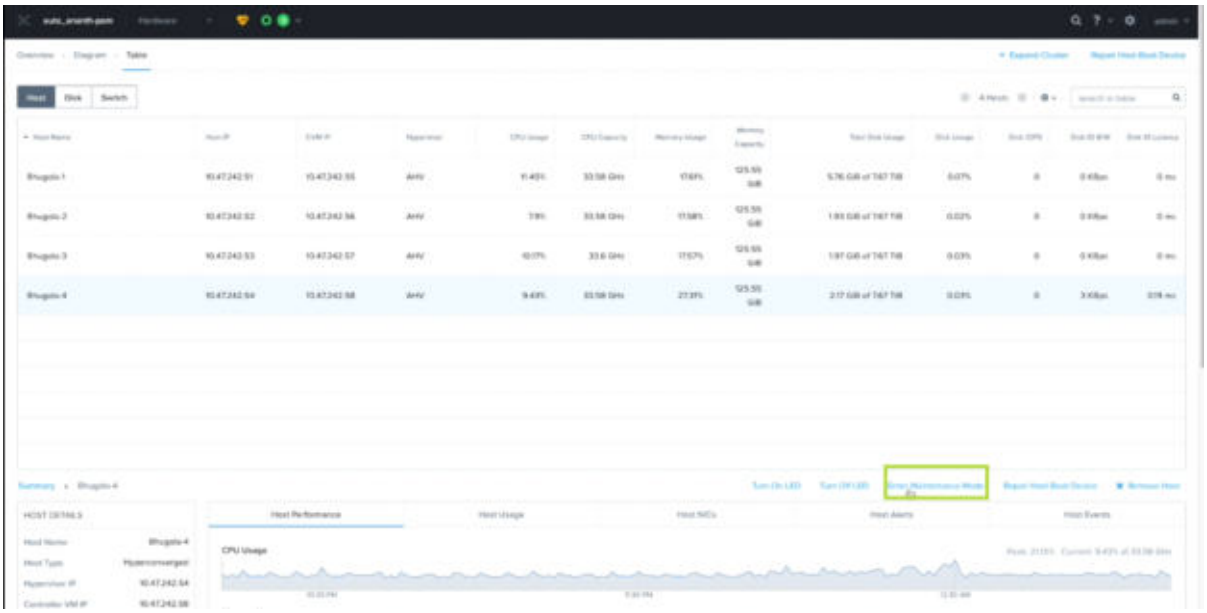

Figure 167: Enter Maintenance Mode Option

The Host Maintenance window appears with a prompt to power-off all VMs that cannot be live migrated,.

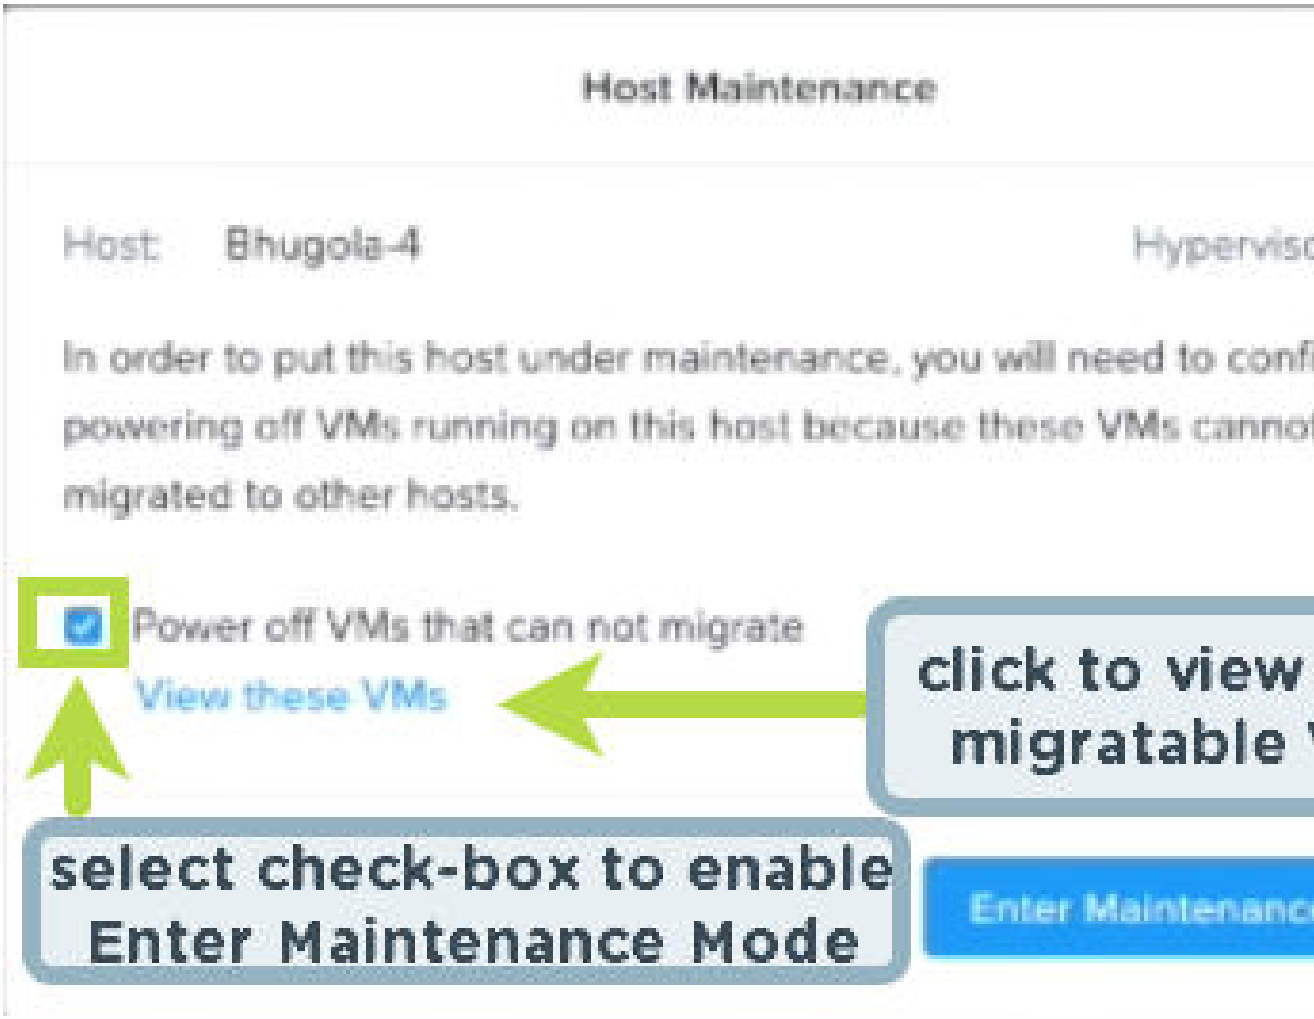

Figure 168: Host Maintenance Window (Enter Maintenance Mode Enabled)

Note: VMs with CPU passthrough, PCI passthrough, pinned VMs (with host affinity policies), and RF1 are not migrated to other hosts in the cluster when a node undergoes maintenance. Click **View these VMs** link to view the list of VMs that cannot be live-migrated.

6. Check the **Power-off VMs that can not migrate** check-box to enable the **Enter Maintenance Mode** button.

7. Click the **Enter Maintenance Mode** button.

   - A revolving icon appears as a tool tip beside the selected node and also in the Host Details view. This indicates that the host is entering the maintenance mode.
   - The revolving icon disappears and the **Exit Maintenance Mode** option is enabled after the node completely enters the maintenance mode.

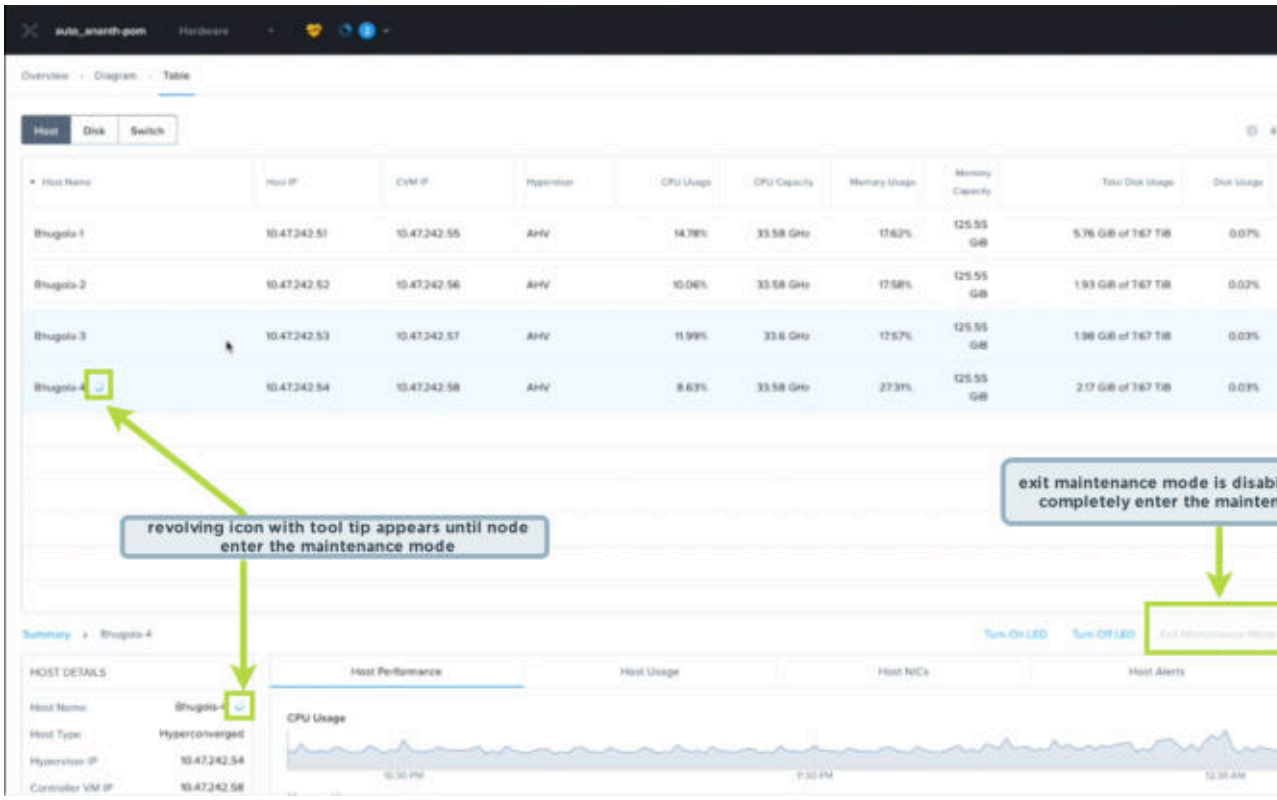

Figure 169: Enter Node Maintenance (On-going)

   - You can also monitor the progress of the node maintenance operation through the newly created Host enter maintenance and Enter maintenance mode tasks which appear in the task tray.
   - As the node enter the maintenance mode, the following high-level tasks are performed internally.

     1. The AHV host initiates entering the maintenance mode.
     2. The HA VMs are live migrated.
     3. The pinned and RF1 VMs are powered-off.
     4. The AHV host completes entering the maintenance mode.
     5. The CVM enters the maintenance mode.
     6. The CVM is shutdown.

     For more information, see UVM Status when Node under Maintenance Mode on page 265 to view the status of the UVMs.

     Note: In case of a node maintenance failure, certain rolled-back operations are performed. For example, the CVM is rebooted. But the live migrated are not restored to the original host.

What to do next

Once the maintenance activity is complete, you can perform any of the following.

- View the nodes under maintenance, see Viewing a Node under the Maintenance Mode on page 261
- View the status of the UVMs, see UVM Status when Node under Maintenance Mode on page 265
- Remove the node from the maintenance mode Exiting a Node from the Maintenance Mode (AHV) on page 263)

Viewing a Node under the Maintenance Mode

After you perform any maintenance activity, exit the node from the maintenance mode.

About this task

Perform the following steps to view a node under maintenance.

Procedure

1. Login to the Prism Element web console.
2. On the home page, select **Hardware** from the drop-down menu.
3. Go to the **Table** > **Host** view.

4. Observe the icon along with a tool tip that appears beside the node which is under maintenance. You can also view this icon in the host details view.

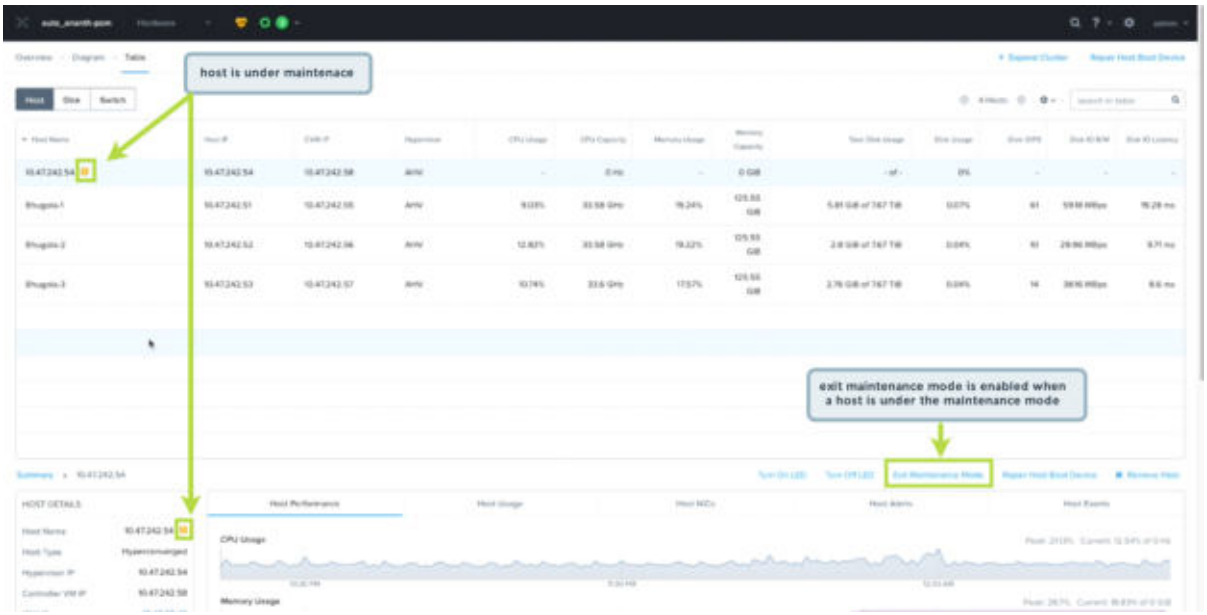

Figure 170: Node under Maintenance (Table and Host Details View)

Alternatively, you can view the node under maintenance from the **Hardware** > **Diagram** view.

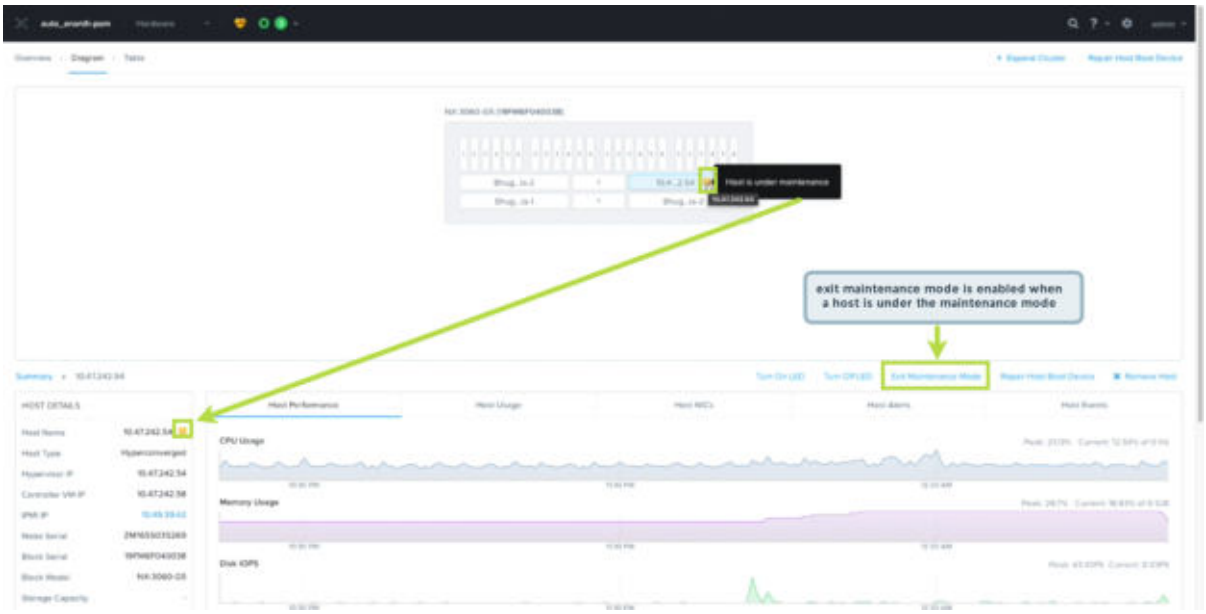

Figure 171: Node under Maintenance (Diagram and Host Details View)

What to do next

When a node is under maintenance, you can perform any of the following.

- View the status of the UVMs, see UVM Status when Node under Maintenance Mode on page 265.

- Remove the node from the maintenance mode Exiting a Node from the Maintenance Mode (AHV) on page 263)

Exiting a Node from the Maintenance Mode (AHV)

After you perform any maintenance activity, exit the node from the maintenance mode.

About this task

Procedure

1. On the Prism web console home page, select **Hardware** from the drop-down menu.
2. Go to the **Table** > **Host** view.
3. Select the node which you intend to remove from the maintenance mode.
4. Click the **Exit Maintenance Mode** option.

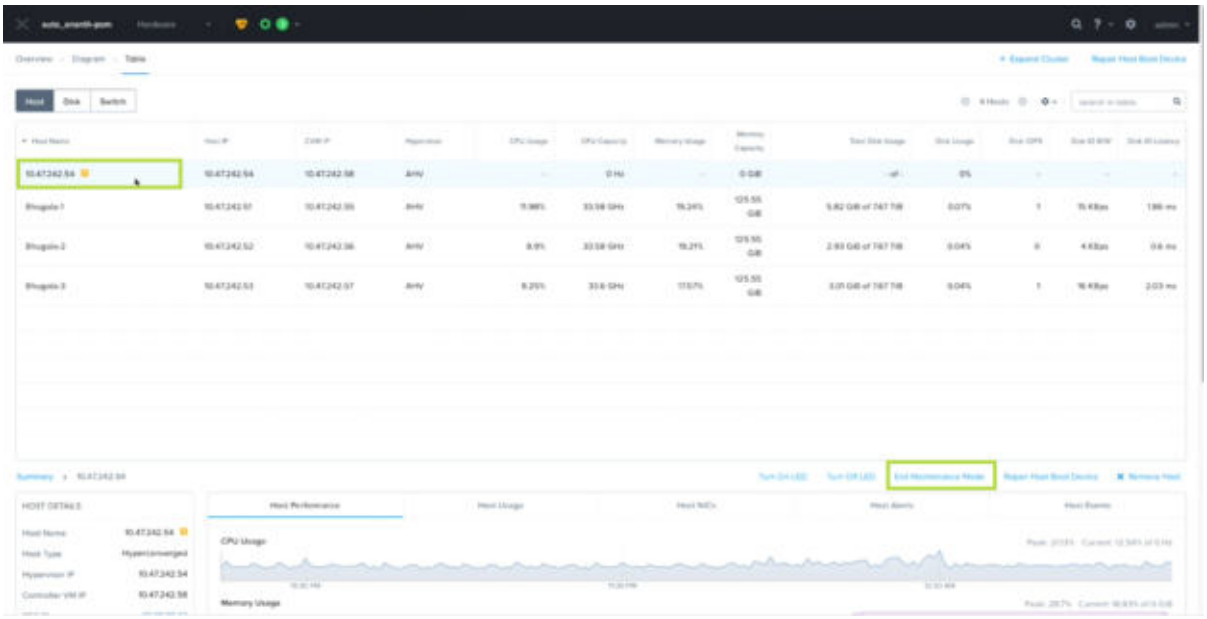

Figure 172: Exit Maintenance Mode Option

The Host Maintenance window appears.

5. Click the **Exit Maintenance Mode** button.

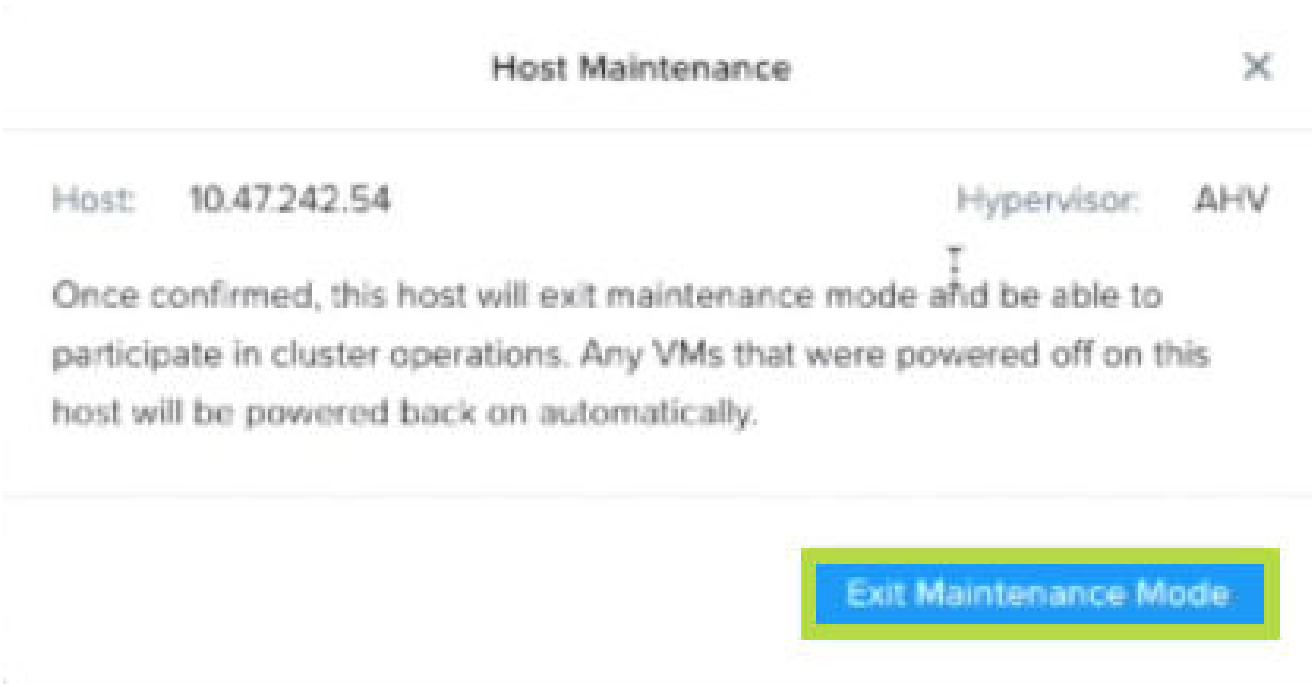

Figure 173: Host Maintenance Window (Exit Maintenance Mode)

- A revolving icon appears as a tool tip beside the selected node and also in the Host Details view. This indicates that the host is exiting the maintenance mode.
- The revolving icon disappears and the **Enter Maintenance Mode** option is enabled after the node completely exits the maintenance mode.

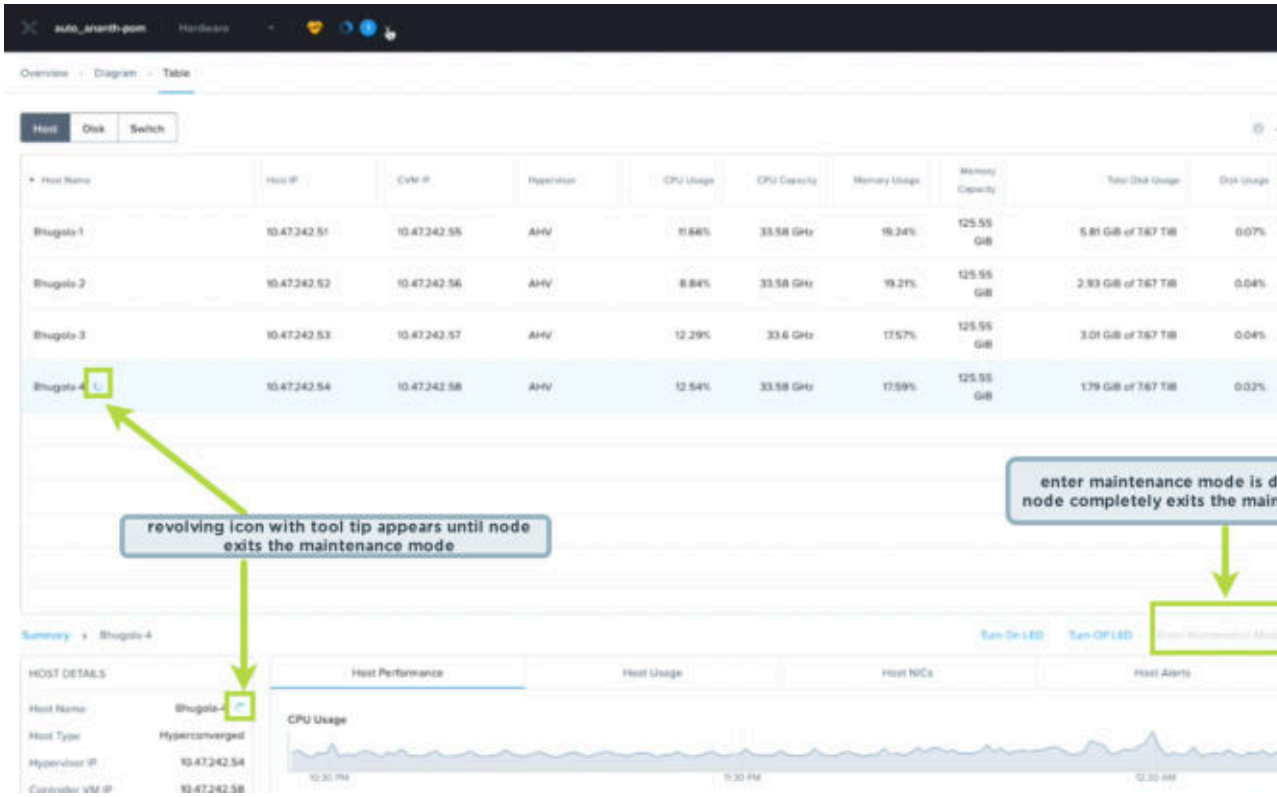

Figure 174: Exit Node Maintenance (On-going)

- You can also monitor the progress of the exit node maintenance operation through the newly created Host exit maintenance and Exit maintenance mode tasks which appear in the task tray.
- As the node exits the maintenance mode, the following high-level tasks are performed internally.

  1. The CVM is powered on.
  2. The CVM is taken out of maintenance.
  3. The host is taken out of maintenance.

  After the host exits the maintenance mode, the RF1 VMs continue to be powered on and the VMs migrate to restore host locality.

  For more information, see UVM Status when Node under Maintenance Mode on page 265 to view the status of the UVMs.

What to do next

Once a node exits the maintenance mode, you can perform any of the following.

- View the status of node under maintenance, see Viewing a Node under the Maintenance Mode on page 261.
- View the status of the UVMs, see UVM Status when Node under Maintenance Mode on page 265.

UVM Status when Node under Maintenance Mode

The following scenarios demonstrate the behavior of the UVMs when a node enters and exits a maintenance operation. The examples showcase the behavior of three VM types. high

availability (HA) VMs, pinned VMs, and RF1 VMs. The HA VMs are live migratable and can migrate across nodes if the host server goes down or reboots. The pinned VMs have the affinity towards a specific AHV node and the RF1 VMs have affinity towards a specific node or a CVM. To view the status of the UVMs, go to **VM** > **Table** > **view.**.

Scenario 1: UVMs before Node Entering Maintenance Mode

In this example, you can observe the status of the UVMs on the node prior to the node entering the maintenance mode. All the UVMs are powered-on and reside on the same host.

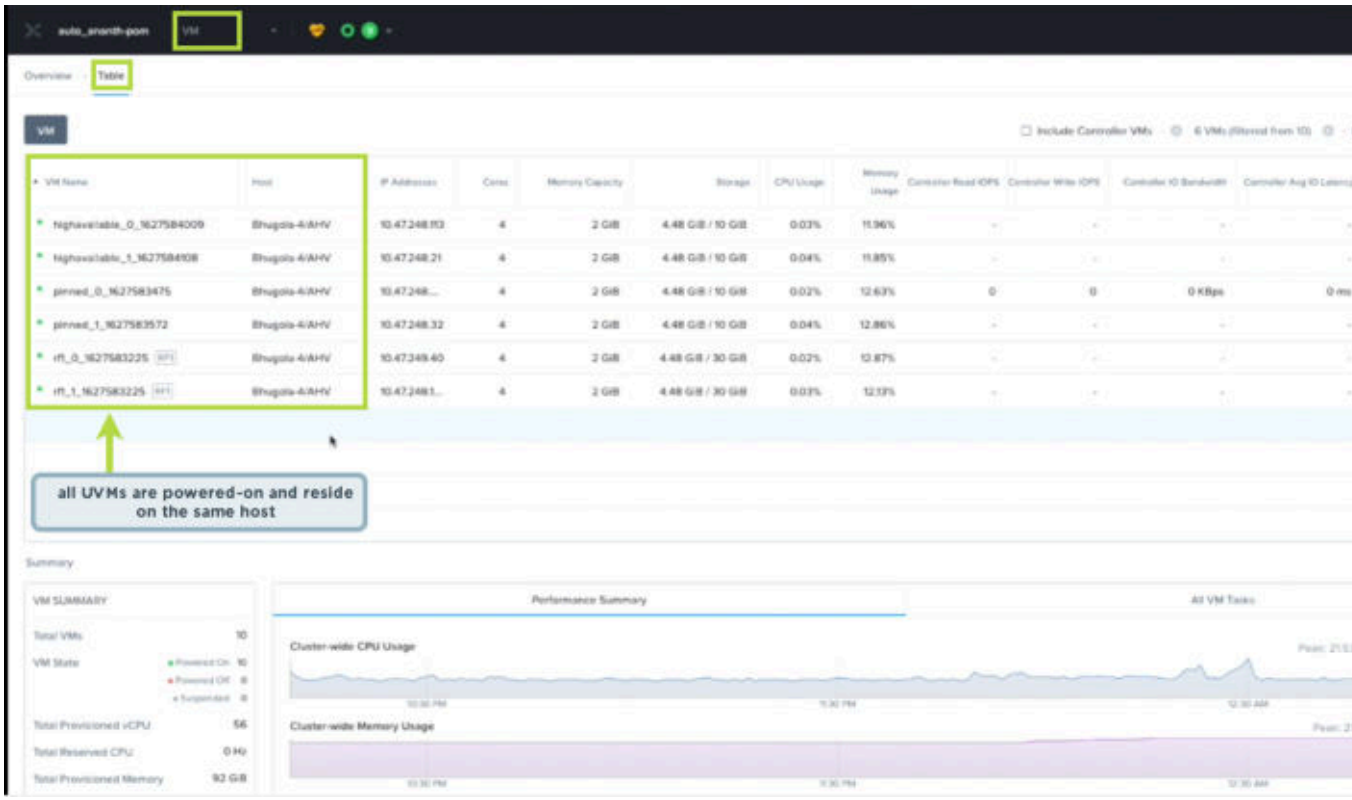

Figure 175: Example: Original State of VM and Hosts

Scenario 2: UVMs during Node Maintenance Mode

- As the node enter the maintenance mode, the following high-level tasks are performed internally.
  1. The AHV host initiates entering the maintenance mode.
  2. The HA VMs are live migrated.
  3. The pinned and RF1 VMs are powered-off.
  4. The AHV host completes entering the maintenance mode.
  5. The CVM enters the maintenance mode.
  6. The CVM is shutdown.

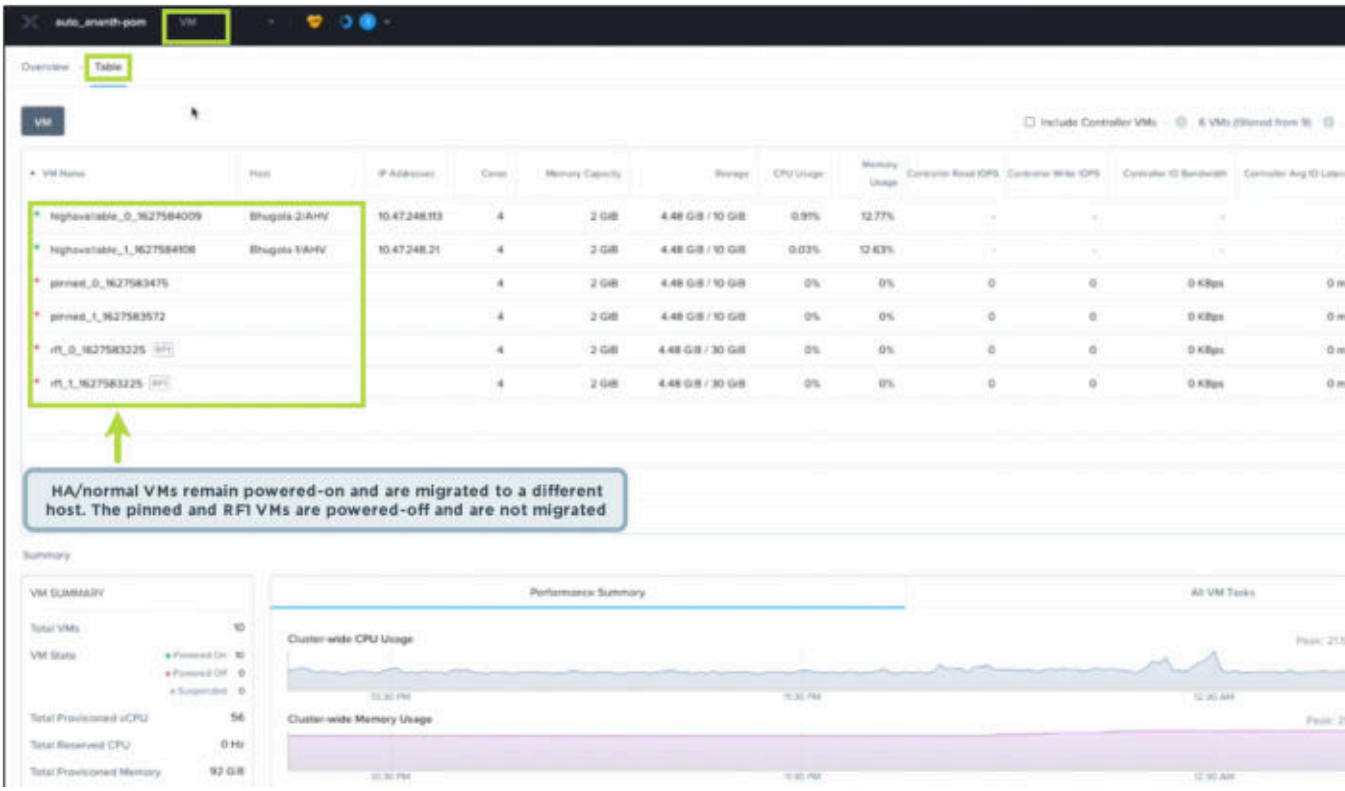

Figure 176: Example: VM and Hosts before Entering Maintenance Mode

Scenario 3: UVMs after Node Exiting Maintenance Mode

- As the node exits the maintenance mode, the following high-level tasks are performed internally.

  1. The CVM is powered on.
  2. The CVM is taken out of maintenance.
  3. The host is taken out of maintenance.

  After the host exits the maintenance mode, the RF1 VMs continue to be powered on and the VMs migrate to restore host locality.

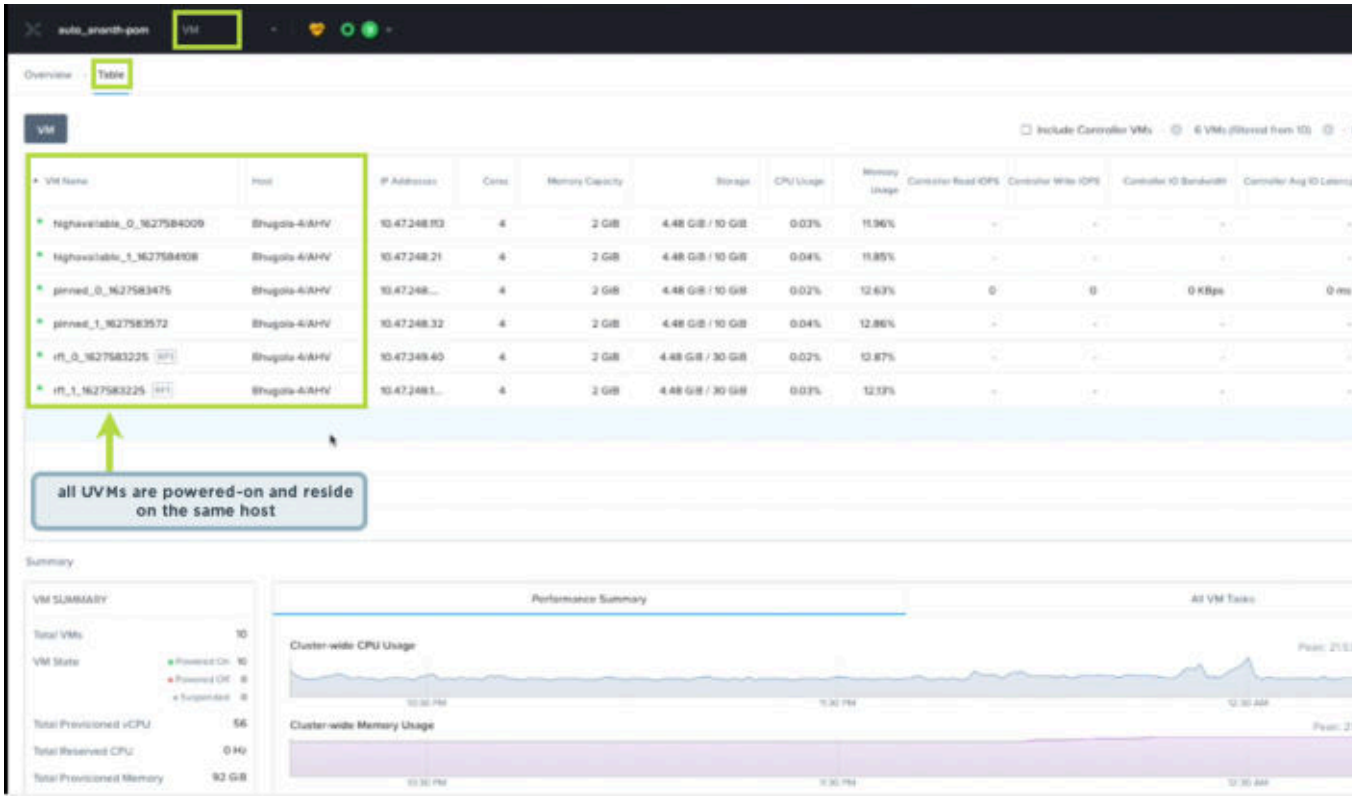

Figure 177: Example: Original State of VM and Hosts

Starting Host Boot Disk Repair

This procedure describes how to proactively install a hypervisor on a boot device that has not completely failed (for example, a SATA DOM, or an M.2 device). A master node helps clone the boot device image onto the target node. This procedure works for all hypervisors. This procedure can clone and restore a complete image before the boot device fails. If this procedure fails, perform the "ungraceful" procedure for a boot device that has completely failed. For some platforms, a single procedure serves for both proactive and failed M.2 drives.

CAUTION:

Verify that it is safe to remove the node from the cluster before you start this procedure. You must follow the hypervisor boot drive replacement procedure for your specific platform. The documents are on the Nutanix support portal under Hardware Replacement Documentation. in one of following categories.

- For Hypervisor Boot Drive Replacement, see https://portal.nutanix.com/page/documents/list?type=hardware&filterKey=Component&filterVal=Hypervisor%20Boot%20Drive
- For SATA DOM, see https://portal.nutanix.com/page/documents/list?type=hardware&filterKey=Component&filterVal=SATA%20DOM
- For M.2, see https://portal.nutanix.com/page/documents/list?type=hardware&filterKey=Component&filterVal=M2

Do not attempt to begin hypervisor boot drive repair without the appropriate document. For additional assistance, contact Nutanix support.

Modifying a Cluster

Hardware components (nodes and disks) can be removed from a cluster or reconfigured in other ways when the conditions warrant it.

Before you begin

If the Data-at-Rest Encryption is enabled then before removing a drive or node from a cluster, test the certificates again by clicking **Test all nodes** and ensure that testing is successful and the status is **Verified**. For a detailed procedure, see *Data-at-Rest Encryption*. In case of an SED drive or node, if the drive or node is not removed as recommended then the drive or node is locked.

About this task

To reconfigure one or more hardware components in a cluster, do the following:

Procedure

- In the Hardware dashboard (see Hardware Dashboard on page 219), click the **Diagram** or **Table** tab.

  The following actions can be performed from either the **Diagram** or **Table** view.

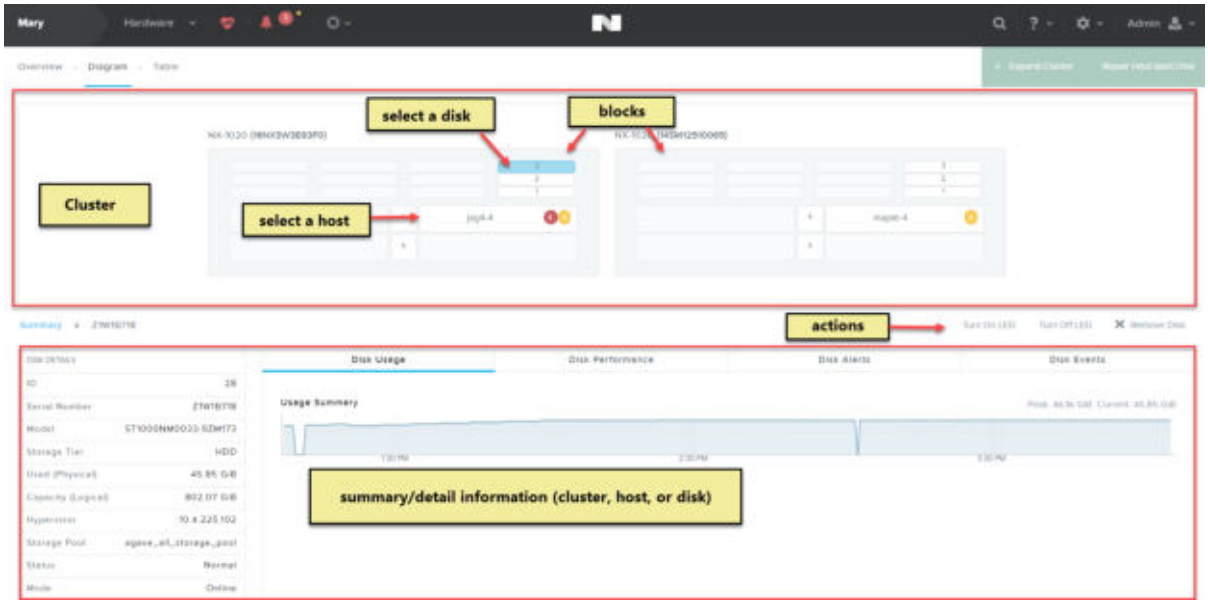

Figure 178: Cluster Hardware (Diagram View)

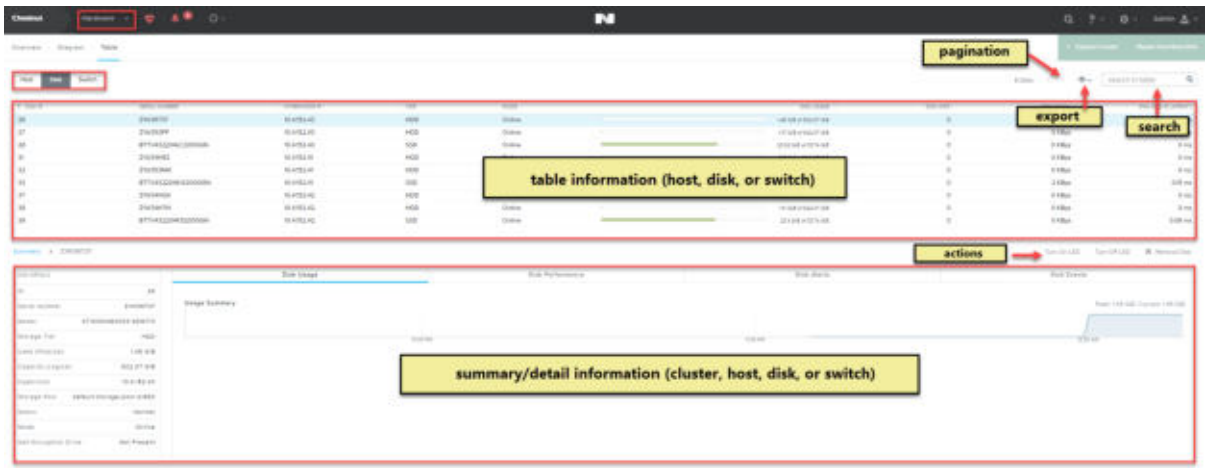

Figure 179: Cluster Hardware (Table View)

Removing a Disk

You may need to remove a disk for various reasons such as during the replacement of a failed disk.

About this task

Removing a disk takes time because the data on that disk must be migrated to other disks before it is removed from a node. Monitor the progress on the disk removal through the dashboard messages.

To remove a disk, do the following.

Procedure

1. Select the disk you want to remove in one of the following ways:
   - » Select the target disk in the diagram (Diagram view).
   - » Click the **Disk** tab and select that disk in the table (Table view).
2. Click the **Remove Disk** link on the right of the **Summary** line.
   A dialog box appears to verify the action.
3. Click the **OK** button in the confirmation dialog box.

   CAUTION: Do not physically remove a disk until that disk status indicator turns red in the diagram. The status message that appears on completion of the steps may indicate that the data migration is complete, but the disk is not ready for removal until that the disk status indicator turns red in the diagram.

Removing a Node from a Cluster

You may need to remove a node for various reasons such as:

- Replacement of a failed node
- Node is either unreachable or powered off
- Deprecate old nodes for cluster expansion

You can remove a node using the Prism web console or nCLI.

Prerequisites for Removing a Node

Consider the following before you remove a node:

- The Prism web console displays a warning message that you need to reclaim the license after you have removed the node. See the Licensing Guide or License Manager Guide for information about reclaiming or rebalancing your cluster licenses.
- Removing a node (host) takes some time because data on that node must be migrated to other nodes before it can be removed from the cluster. You can monitor progress through the dashboard messages. Removing a node implicitly removes all the disks in that node.
- (Hyper-V only) Initiating a removal of a node running Hyper-V fails if the node is running as a part of a Hyper-V failover cluster and the following message appears.

  Node *node id* is a part of a Hyper-V failover cluster *failover cluster name*. Please drain all the roles, remove the node from the failover cluster and then mark the node for removal.

  If this message is displayed in either nCLI or in web interface, as a cluster administrator, you must use the management tools provided by Microsoft such as *Failover Cluster Manager* to drain all the highly-available roles off the node. Then remove the node from the failover cluster followed by removing the node from the AOS cluster.
- (ESXi only) As a cluster administrator, you must use the management tools provided by VMware to first migrate all the VMs off the node/host, disconnect the node from vCenter, and then remove the node from the vCenter cluster. After that you can remove the node (host) from the AOS cluster.

  CAUTION: Ensure that you migrate the VMs to an ESXi cluster before removing a host or node. Verify that the cluster has enough available compute capacity before actually migrating

the VMs. If you remove a node or host without first migrating the VMs, the VMs data may be migrated without any notice resulting in loss of service.

Removing a Single Node

About this task

Review the prerequisites for removing a node listed in the *Prerequisites for Removing a Node* topic.

To remove a single node (host) from the cluster, do the following.

Procedure

1. Select the node you want to remove in one of the following ways:

   - In the Diagram page, select the target node (host).

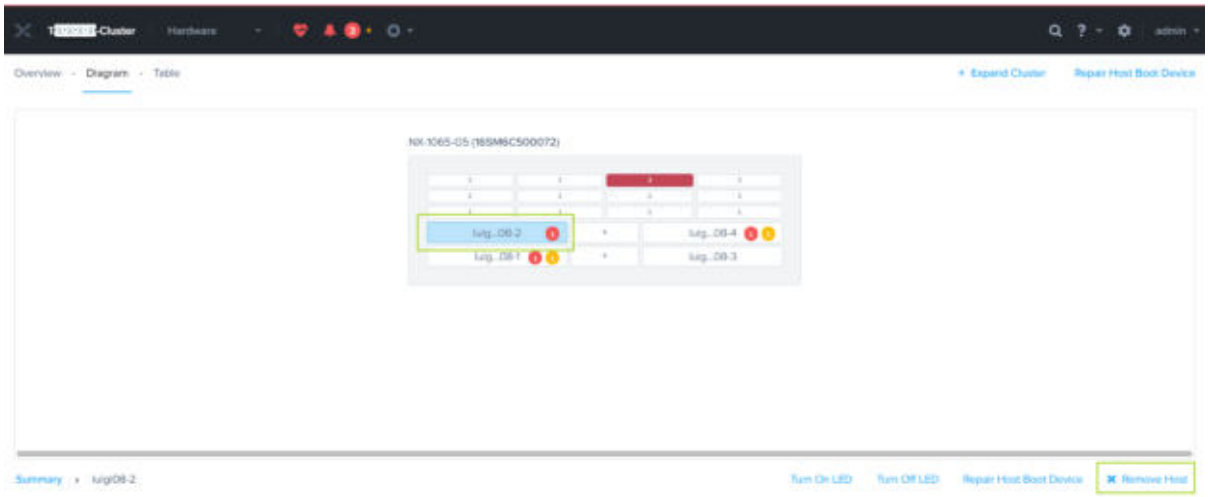

Figure 180: Diagram View

   - In the Table page, click the Host tab and select the node (host).

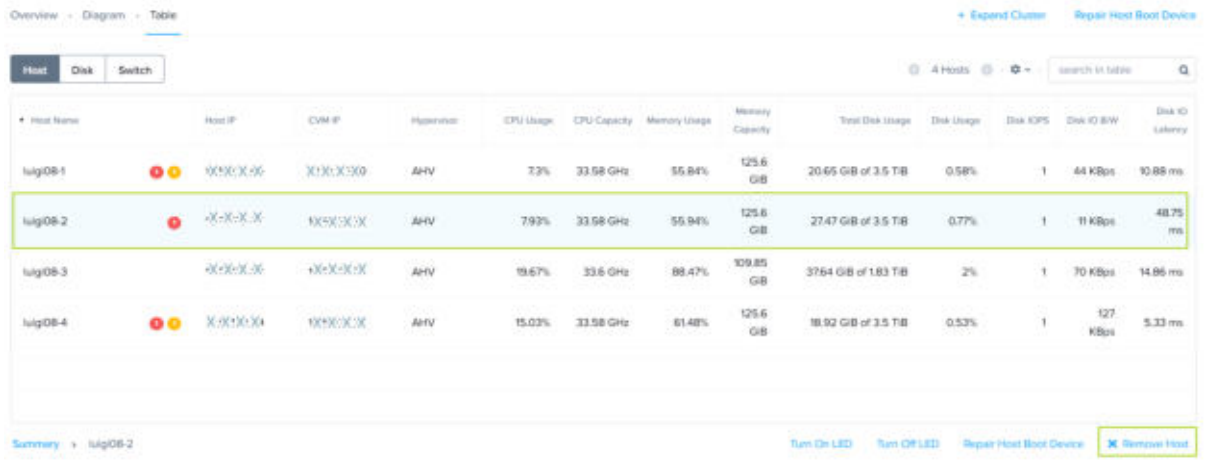

Figure 181: Table View

2. Click the **Remove Host** link on the right of the **Summary** line.

   > Note: The **Remove Host** link on the right of the **Summary** line does not appear for the removal of a host from a three-node cluster.

   A dialog box appears to verify the action.

3. Click the **OK** button in the confirmation dialog box.

4. If the node that you are trying to remove is unreachable or powered off, a notification is triggered in the Prism UI alerting you that the storage utilization for this node could not be calculated and also suggesting the possible impact of removing this node. If you still want to go ahead, you can use the force option to forcefully mark this node for removal.

What to do next

After a node is removed, it goes into a state without any configuration. You can add such a node back into the cluster through the Expanding a Cluster on page 243 workflow.

> CAUTION: After adding the removed node back into the cluster, the same cluster ID is applied to both clusters at these circumstances:
>
> - When the removed node is the first node (lowest IP address in the cluster)
> - When you reuse the removed node in another cluster
> - When the removed node again becomes the first node (lowest IP address) in the new cluster
>
> To prevent this occurrence, reimage the node (using Foundation) before adding it to the new cluster.
>
> See the *Prepare Bare-Metal Nodes for Imaging* chapter in the *Field Installation Guide* for details about how to image a node.

Remove Multiple Nodes

AOS allows you to remove multiple nodes from a cluster. The multinode removal process is sequential, that means you can request only one node at a time. However, you can mark other nodes for removal one after the other while data is being evicted in background from other nodes that are marked for removal.

Here is the multinode removal process overview:

- User requests to remove multiple nodes (serially one after the other) from the Prism web console or nCLI.
- System runs prechecks on each node to allow node removal request. The node must pass all the checks to be successfully accepted for node removal.

  For more information, see Prechecks to Allow Multiple Node Removal on page 274topic in this guide.
- System removes nodes serially from metadata ring of Cassandra.
- System rebuilds the data in parallel in the background for all the nodes that are marked for removal.

Prechecks to Allow Multiple Node Removal

When you mark multiple nodes for removal, system runs the prechecks before accepting the node for removal. Below are some of the prechecks:

- The cluster upgrade process is not in progress.
- The minimum number of nodes required by Cassandra and Zookeeper are available and are in healthy (kNormal status) state.
- No other nodes are pending to be added to the cluster.
- The cluster continues to meet the maximum RF requirements after removal of the node.
- The cluster has enough usable storage capacity to rebuild the data from the removed node.

The system does not allow multinode removal if a node fails any of the prechecks.

Limitations for Multinode Removal

The following limitations are applicable for the multinode removal process:

- System allows a maximum of four nodes to be removed simultaneously from a cluster.
- For clusters with EC (erasure coding), the system computes the longest strip size possible on the cluster based on the EC param and desired fault tolerance of the container. The system does not allow multinode removal if there are not enough viable entities of the desired fault domain (node/ rack/ru) to accommodate the largest EC strip possible on this cluster.
- In the nCLI, the multinode removal is not allowed if you skip the space usage check.

```
<ncli> host rm-start skip-space-check=true
```

Removing Multiple Nodes

About this task

> Note: This feature is supported on a cluster with a minimum of four hosts (nodes).

Review the prerequisites for removing a node listed in the Prerequisites for Removing a Node on page 271 topic.

To remove multiple nodes from the cluster, do the following:

Procedure

1. Select the node you want to remove in one of the following ways:

   - In the Diagram page, select the target node (host).

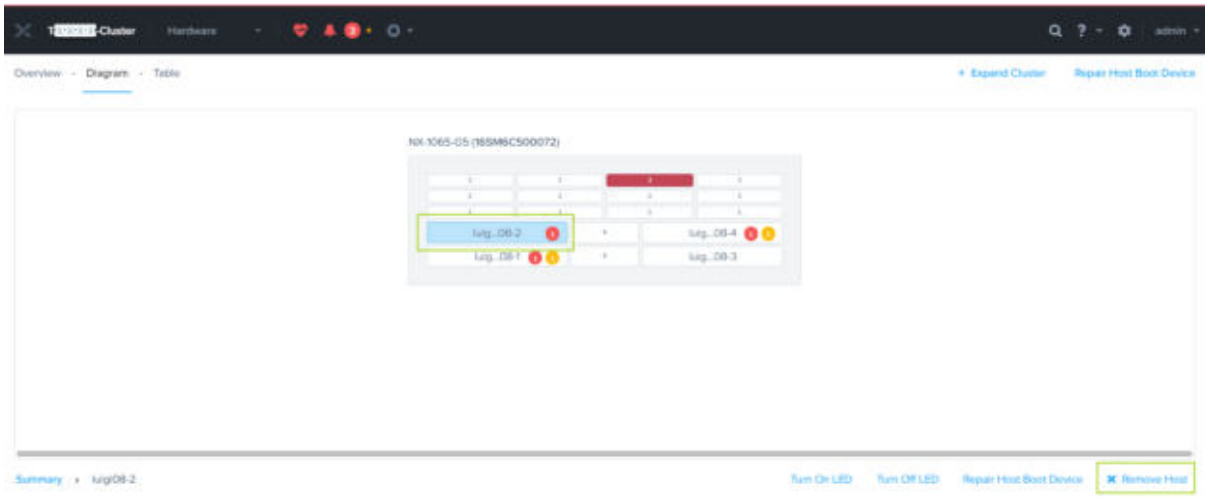

Figure 182: Diagram View

   - In the Table page, click the Host tab and select the node (host).

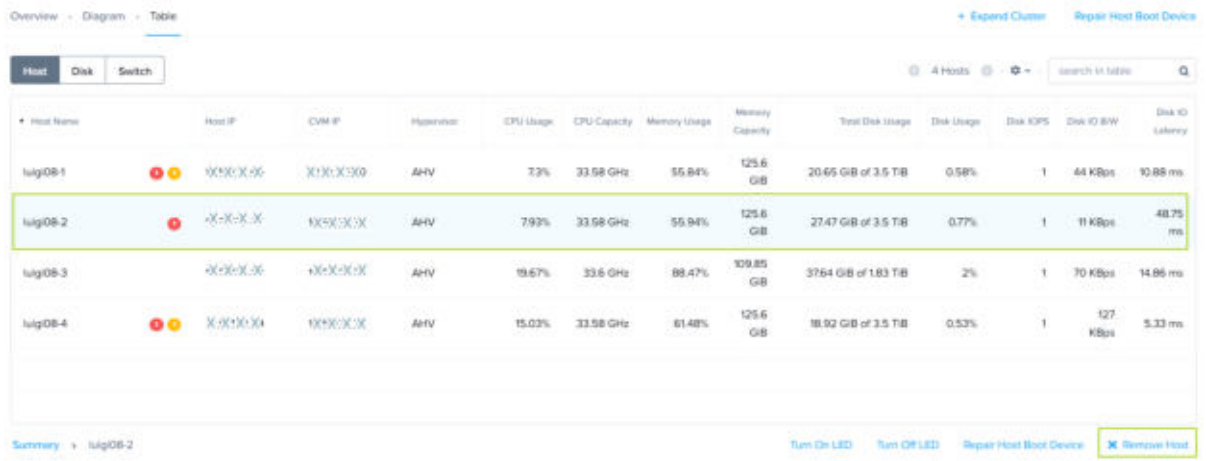

Figure 183: Table View

2. Click the **Remove Host** link on the right of the **Summary** line.

   Note: The **Remove Host** link on the right of the **Summary** line does not appear for the removal of a host from a three-node cluster.

   A dialog box appears to verify the action.

3. Click **Remove** in the confirmation dialog box.

   You can monitor the node removal progress in the **Tasks** menu.

4. To remove multiple nodes, repeat Step 1 through Step 3 for each node.

   The Tasks menu shows the progress of the node removal. The first node that was selected to be removed has two subtasks running:

   - Transferring metadata to replacement replicas
   - Extent store replication

   The subsequent nodes have one subtask running in parallel.

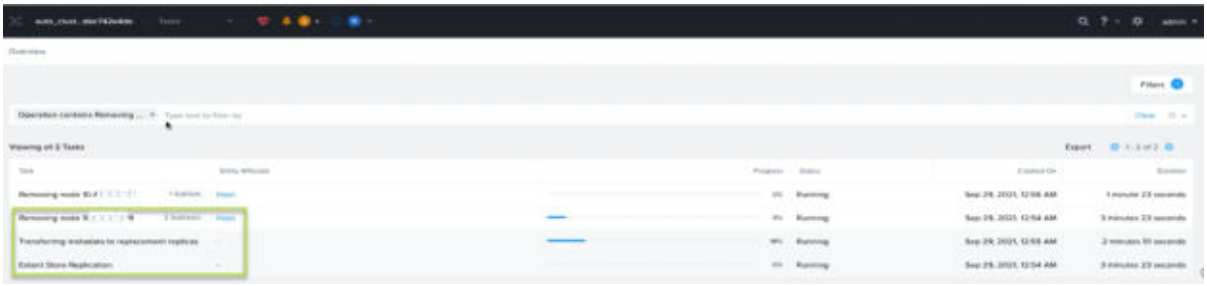

Figure 184: Tasks View

You can also monitor the data rebuild progress in the **Data Resiliency Status** ticker that appears in the Prism web console dashboard.

What to do next

After a node is removed, it goes into a state without any configuration. You can add such a node back into the cluster through the Expanding a Cluster on page 243 workflow.

CAUTION: After adding the removed node back into the cluster, the same cluster ID is applied to both clusters at these circumstances:

- When the removed node is the first node (lowest IP address in the cluster)
- When you reuse the removed node in another cluster
- When the removed node again becomes the first node (lowest IP address) in the new cluster

To prevent this occurrence, reimage the node (using Foundation) before adding it to the new cluster.

See the *Prepare Bare-Metal Nodes for Imaging* chapter in the *Field Installation Guide* for details about how to image a node.

Adding a Node

You can need to add a node or host into the metadata store in events such as the aftermath of the replacement of a failed metadata disk.

About this task

Each node includes a disk used for metadata storage, and AOS maintains a metadata store across these disks to ensure uninterrupted resiliency should a metadata disk fail. After such a failure, that node is taken out of the metadata store group and the cluster continues to operate seamlessly without it. Normally, the node is brought back into the metadata store automatically after the failed metadata disk is replaced. However, under certain (rare) circumstances this

might not happen. If the node is ready but was not added back automatically, you need to add the node back into the metadata store manually.

When the node is not added back automatically, the following alert message is displayed:

Node ready to be added to metadata store

To add a host into the metadata store, do the following.

Procedure

1. Select the node you want to remove in one of the following ways:
   - » Select the target node (host) in the diagram (Diagram view).
   - » Click the **Host** tab and select that disk in the table (Table view).
2. Click the **Enable Metadata Store** link on the right of the **Summary** line.
   The **Enable Metadata Store** link appears only when the node is not added back automatically and the alert message is displayed

Compute-Only Node Configuration (AHV Only)

A compute-only (CO) node allows you to seamlessly and efficiently expand the computing capacity (CPU and memory) of your AHV cluster. The Nutanix cluster uses the resources (CPUs and memory) of a CO node exclusively for computing purposes.

You can use a supported server or an existing hyperconverged (HC) node as a CO node. To use a node as CO, image the node as CO by using Foundation and then add that node to the cluster by using the Prism Element web console. See *Field Installation Guide* for information about how to image a node as a CO node.

> Note: If you want an existing HC node that is already a part of the cluster to work as a CO node, remove that node from the cluster, image that node as CO by using Foundation, and add that node back to the cluster. See the *Modifying a Cluster* topic in *Prism Web Console Guide* for information about how to remove a node.

Key Features of Compute-Only Node

Following are the key features of CO nodes.

- CO nodes do not have a Controller VM (CVM) and local storage.
- AOS sources the storage for vDisks associated with VMs running on CO nodes from the hyperconverged (HC) nodes in the cluster.
- You can seamlessly manage your VMs (CRUD operations, ADS, and HA) by using the Prism Element web console.
- AHV runs on the local storage media of the CO node.
- To update AHV on a cluster that contains a compute-only node, use the Life Cycle Manager. See *Life Cycle Manager Guide: LCM Updates* for more information.

Use Case of Compute-Only Node

CO nodes enable you to achieve more control and value from restrictive licenses such as Oracle. A CO node is part of a Nutanix HC cluster, and there is no CVM running on the CO node (VMs use CVMs running on the HC nodes to access disks). As a result, licensed cores on the CO node are used only for the application VMs.

Applications or databases that are licensed on a per CPU core basis require the entire node to be licensed and that also includes the cores on which the CVM runs. With CO nodes, you get a much higher ROI on the purchase of your database licenses (such as Oracle and Microsoft SQL Server) since the CVM does not consume any compute resources.

Minimum Cluster Requirements

Following are the minimum cluster requirements for compute-only nodes.

- The Nutanix cluster must be at least a three-node cluster before you add a compute-only node.

  However, Nutanix recommends that the cluster has four nodes before you add a compute-only node.

- The ratio of compute-only to hyperconverged nodes in a cluster must not exceed the following:

  1 compute-only : 2 hyperconverged

- All the hyperconverged nodes in the cluster must be all-flash nodes.
- The number of vCPUs assigned to CVMs on the hyperconverged nodes must be greater than or equal to the total number of available cores on all the compute-only nodes in the cluster. The CVM requires a minimum of 12 vCPUs. See the *CVM vCPU and vRAM Allocation* topic in *Field Installation Guide* for information about how Foundation allocates memory and vCPUs to your platform model.
- The total amount of NIC bandwidth allocated to all the hyperconverged nodes must be twice the amount of the total NIC bandwidth allocated to all the compute-only nodes in the cluster.

  Nutanix recommends you use dual 25 GbE on CO nodes and quad 25 GbE on an HC node serving storage to a CO node.

- The AHV version of the compute-only node must be the same as the other nodes in the cluster.

  When you are adding a CO node to the cluster, AOS checks if the AHV version of the node matches with the AHV version of the existing nodes in the cluster. If there is a mismatch, the add node operation fails.

See the *Expanding a Cluster* topic in the *Prism Web Console Guide* for general requirements about adding a node to a Nutanix cluster.

Restrictions

Nutanix does not support the following features or tasks on a CO node in this release:

1. Host boot disk replacement
2. Network segmentation

Supported AOS Versions

Nutanix supports compute-only nodes on AOS releases 5.11 or later.

Supported Hardware Platforms

Compute-only nodes are supported on the following hardware platforms.

- All the NX series hardware
- Dell XC Core

- Cisco UCS

Networking Configuration

To perform network tasks on a compute-only node such as creating or modifying bridges or uplink bonds or uplink load balancing, you must use the manage_ovs commands and add the **--host** flag to the **manage_ovs** commands as shown in the following example:

> Note: If you have storage-only AHV nodes in clusters with compute-only nodes being ESXI or Hyper-V, deployment of default virtual switch vs0 fails. In such cases, the Prism Element, Prism Central or CLI workflows for virtual switch management are unavailable to manage the bridges and bonds. Use the manage_ovs command options to manage the bridges and bonds.

```
nutanix@cvm$ manage_ovs --host IP_address_of_co_node --bridge_name bridge_name create_single_bridge
```

Replace *IP_address_of_co_node* with the IP address of the CO node and *bridge_name* with the name of bridge you want to create.

> Note: Run the **manage_ovs** commands for a CO from any CVM running on a hyperconverged node.

Perform the networking tasks for each CO node in the cluster individually.

See the *Host Network Management* section in the *AHV Administration Guide* for more information about networking configuration of the AHV hosts.

Adding a Compute-Only Node to an AHV Cluster

About this task

Perform the following procedure to add a compute-only node to a Nutanix cluster.

Procedure

1. Log on to the Prism Element web console.
2. Do one of the following:
   - » Click the gear icon in the main menu and select **Expand Cluster** in the **Settings** page.
   - » Go to the hardware dashboard (see the *Hardware Dashboard* topic in *Prism Web Console Guide*) and click **Expand Cluster**.

3. In the **Select Host** screen, scroll down and, under **Manual Host Discovery**, click **Discover Hosts Manually**.

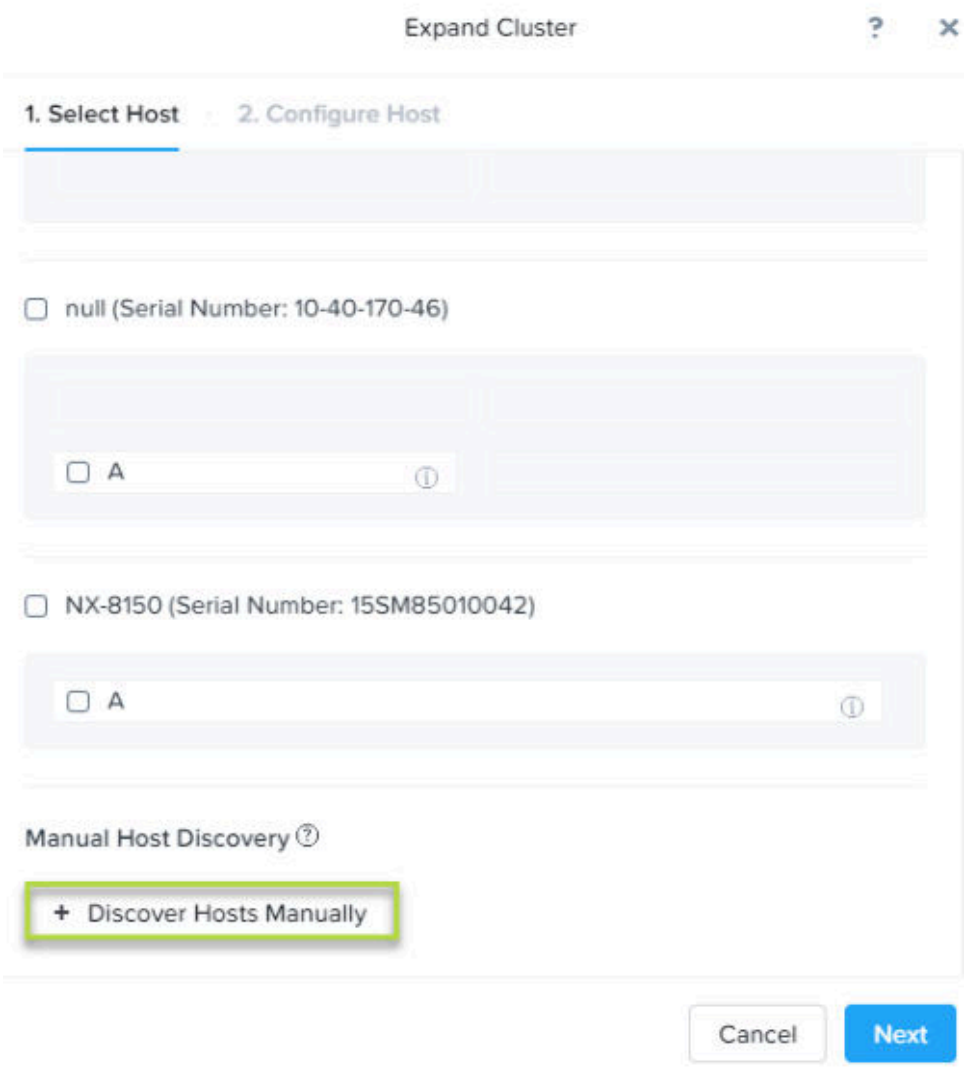

Figure 185: Discover Hosts Manually

4. Click **Add Host**.

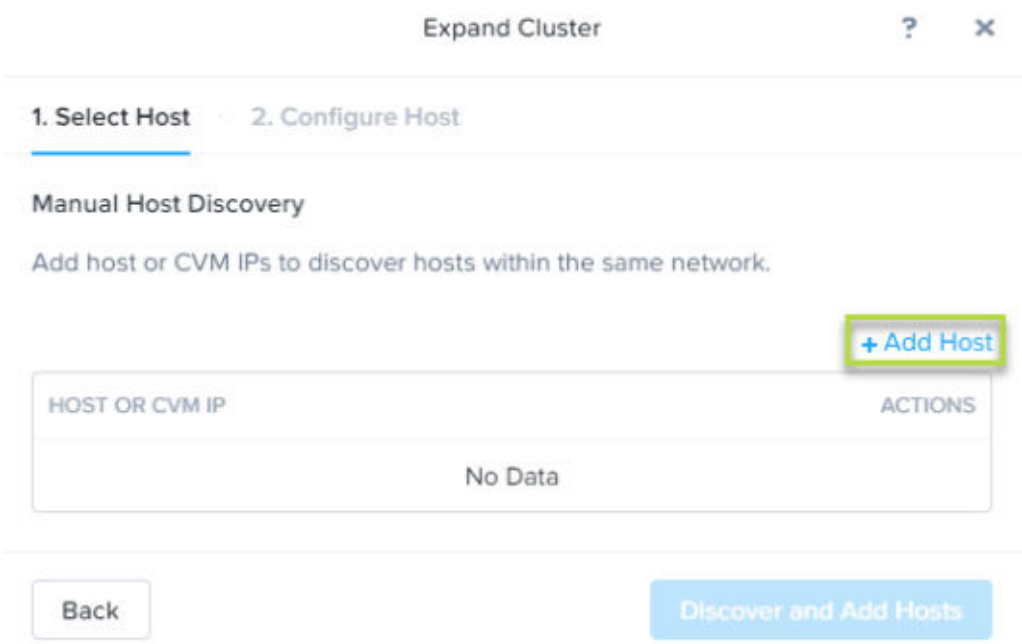

Figure 186: Add Host

5. Under **Host or CVM IP**, type the IP address of the AHV host and click **Save**.
   This node does not have a Controller VM and you must therefore provide the IP address of the AHV host.
6. Click **Discover and Add Hosts**.
   Prism Element discovers this node and the node appears in the list of nodes in the Select Host screen.
7. Select the node to display the details of the compute-only node.
8. Click **Next**.
9. In the **Configure Host** screen, click **Expand Cluster**.
   The add node process begins and Prism Element performs a set of checks before the node is added to the cluster.
   Check the progress of the operation in the **Tasks** menu of the Prism Element web console. The operation takes approximately five to seven minutes to complete.
10. Check the **Hardware Diagram** view to verify if the node is added to the cluster.
    You can identity a node as a CO node if the Prism Element web console does not display the IP address for the CVM.

Storage-only Node Configuration

A storage-only node allows you to seamlessly expand the storage capacity in your cluster. A storage-only node always runs the AHV hypervisor. If you add a storage-only node to any cluster (regardless of the hypervisor), the hypervisor on the storage-only node is always AHV. Therefore, if you want to scale up only the storage capacity in your cluster, you do not need to purchase additional hypervisor licenses.

Key Features of Storage-only Nodes

Following are the key features of storage-only nodes:

- A storage-only node includes an AHV hypervisor, a Controller VM (CVM), and memory and CPU resources enough to run only the CVM.
- A storage-only node always runs the AHV hypervisor. Therefore, if you add a storage-only node to an ESXi or Hyper-V cluster, the hypervisor on the storage-only node is always AHV.
- A storage-only node is supported on all NX hardware models. For other hardware models, check with the manufacturer if storage-only nodes are supported on that hardware.
- A storage-only node is supported on ESXi, Hyper-V, and AHV clusters.

  If you add storage-only nodes to ESXi or Hyper-V clusters, you cannot run guest VMs on the storage-only nodes.

  To ensure that no guest VMs run on a storage-only node in an AHV cluster, you can either configure AHV affinity rules or add the storage-only node as a never-schedulable node. See the *Adding a Never-Schedulable Node* topic in *AHV Administration Guide* for more information about never-schedulable nodes. You can run guest VMs on a storage-only node that is part of an AHV cluster (all AHV nodes). However, ensure that the storage-only node has enough memory and CPU resources.
- If you have storage-only AHV nodes in clusters with compute-only nodes being ESXI or Hyper-V, deployment of default virtual switch vs0 fails. In such cases, the Prism Element, Prism Central or CLI workflows for virtual switch management are unavailable to manage the bridges and bonds. Use the `manage_ovs` command options to manage the bridges and bonds.

  Note: A storage-only node is not the same as a storage-heavy node. A storage-heavy node is a regular Nutanix hyperconverged node, but with a greater storage capacity. A storage-heavy node can run any hypervisor (AHV, ESXi, or Hyper-V) and can run guest VMs.

Storage-only Node Best Practices

Following are the best practice recommendations if you plan to include storage-only nodes in your cluster.

- Ensure that your cluster has sufficient capacity if a storage-only node fails.

  Your cluster must have sufficient capacity to handle the rolling-upgrade operations during which nodes are restarted.

  To achieve optimal cluster resiliency, deploy storage-only nodes in a replication factor (RF) plus one (RF + 1) configuration. For example, if the RF of your cluster is two, deploy at least three storage-only nodes. If the RF of your cluster is three, deploy at least four storage-only nodes. The RF + 1 configuration is especially important if you plan to configure your storage-only nodes with higher storage capacity than the regular hyperconverged nodes.
- Ensure that the SSD capacity of the node is at least 10% of the HDD capacity of the node. For example, if your cluster has 20 TB HDD capacity, the SSD capacity of the node is at least 2 TB.

  This recommendation does not apply to all-flash clusters.
- Ensure that you configure all CVMs in the cluster with the same resources. For example, if the CVMs on nodes other than the storage-only nodes have 32 GB RAM, configure the CVMs on the storage-only nodes with 32 GB RAM.
- Ensure that a cluster that has all storage-only nodes has at least 3 nodes.

The following table shows the supported and unsupported configurations for a cluster with storage-only nodes.

Configuration	Supported	Description
A four-node cluster with 4 TB capacity on each node. One large capacity storage-only node of 30 TB capacity.	X	If the only storage-only node in the cluster fails, the cluster does not have the sufficient capacity of 30 TB. **Resolution**: Add the minimum number of storage-only nodes equal to the replication factor of your cluster.
A four-node cluster with 4 TB capacity on each node. Two large capacity storage-only nodes of 30 TB capacity each.	#	If one storage-only node fails, the other can handle the required capacity.
A four-node cluster with 4 TB capacity on each node. One same-capacity storage-only node as the rest of the cluster.	#	The cluster is able to handle the failure of any node in the cluster, because all nodes in the cluster are of the same capacity.
Nutanix cluster with 2 TB SSD capacity and 20 TB HDD capacity.	#	The total SSD capacity is minimum 10% of the total HDD capacity.
Nutanix cluster with 2 TB SSD capacity and 40 TB HDD capacity.	X	The total SSD capacity is not minimum 10% of the total HDD capacity.

= Supported configuration

X = Unsupported configuration

Imaging and Adding Storage-only Nodes

You can image storage-only nodes when you are creating a cluster by using Foundation. You can also expand your existing cluster to add a storage-only node.

- See the *Field Installation Guide* for instructions on imaging storage-only nodes when creating a cluster.
- See Expanding a Cluster on page 243 for instructions on adding storage-only nodes to an existing cluster.

6

NUTANIX VOLUMES

Nutanix Volumes acts as one or more targets for client Windows or Linux operating systems running on a bare metal server or as guest VMs using iSCSI initiators. You can use any storage available in any new or existing Nutanix cluster for Volumes.

See the Nutanix Volumes Guide for more information. For best practice information, Nutanix provides the Nutanix Volumes Solutions Documentation.

7

FILE SERVER MANAGEMENT

Nutanix Files (Files) lets you share files across user work stations in a centralized and protected location to eliminate the requirement for a third-party file server. The File Server dashboard displays dynamically updated information about the file servers and shares/exports in a cluster. To view the File Server dashboard, select **File Server** from the pull-down list on the left of the main menu. See the Nutanix Files Guide for more information.

8

DATA PROTECTION

Nutanix offers data protection solutions for virtual datacenters. Nutanix provides data protection functions at the VM, file, and volume group level, so VMs and data remain safe in a crash-consistent environment.

> Note: See *Nutanix DR Solutions Workflow* for more data protection solutions available from Nutanix.

- Nutanix supports several types of protection strategies including one-to-one or one-to-many replication (see *Protection Strategies* in the *Data Protection and Recovery with Prism Element*).
- You can configure asynchronous disaster recovery (see *Asynchronous Disaster Recovery Between On-Premises Datacenters* in the *Data Protection and Recovery with Prism Element*).
  - The Cloud Connect feature gives you the option to asynchronously back up data on the Amazon Web Service (AWS) cloud (see *Asynchronous Replication With Cloud Connect (On-Premises To Cloud)* in the *Data Protection and Recovery with Prism Element*) .
- You can also configure synchronous disaster recovery. There are two options:
  - For environments that support metro availability, you can configure a metro availability protection domain (see *Metro Availability (ESXi and Hyper-V 2016)* in the *Data Protection and Recovery with Prism Element*).
  - For other environments, you can configure synchronous replication (see *Synchronous Replication (ESXi and Hyper-V 2012)* in the *Data Protection and Recovery with Prism Element*).

9

HEALTH MONITORING

Nutanix provides a range of status checks to monitor the health of a cluster.

- Summary health status information for VMs, hosts, and disks appears on the home dashboard (see Home Dashboard on page 53).
- In depth health status information for VMs, hosts, and disks is available through the Health dashboard (see Health Dashboard on page 287).
- You can customize the frequency of the scheduled health checks and how frequently to run them. (see Configuring Health Checks on page 292).
- You can run NCC health checks directly from the Prism. (see Running Checks by Using Web Console on page 295).
- You can collect logs for all the nodes and components. (see Collecting Logs by Using Web Console on page 296).
- For a description of each available health check, see Alerts/Health checks on page 426.

> Note: If the Cluster Health service status is DOWN for more than 15 minutes, an alert email is sent by the AOS cluster to configured addresses and Nutanix support (if selected). In this case, no alert is generated in the Prism web console. The email is sent once per 24 hours. You can run the NCC check cluster_services_down_check to see the service status.

Health Dashboard

The Health dashboard displays dynamically updated health information about VMs, hosts, and disks in the cluster. To view the Health dashboard, select **Health** from the pull-down list on the left of the main menu.

Menu Options

The Health dashboard does not include menu options other than those available from the main menu (see Main Menu on page 46).

> Note: When you first visit the Health dashboard, a tutorial opens that takes you through a guided tour of the health analysis features. Read the message and then click the **Next** button in the text box to continue the tutorial (or click the **X** in the upper right to close the text box and end the tutorial.) You can view this tutorial at any time by selecting the **Health Tutorial** option in the user menu (see Main Menu on page 46).

Screen Details

The Health dashboard is divided into three columns:

- The left column displays tabs for each entity type (VMs, hosts, disks, storage pools, storage containers, cluster services, and [when configured] protection domains and remote sites). Each tab displays the entity total for the cluster (such as the total number of disks) and the

number in each health state. Clicking a tab expands the displayed information (see following section).

- The middle column displays more detailed information about whatever is selected in the left column.
- The right column displays summary of all the health checks. You also have an option to view individual checks from the **Checks** button (success, warning, failure, or disabled).
  - The **Summary** tab provides summarized view of all the health checks according to check status and check type.
  - The **Checks** tab provides information about individual checks. Hovering the cursor over an entry displays more information about that health check (see following figure). You can filter the checks by clicking appropriate field type and clicking **Apply**. The checks are categorized as follows.

    Filter by Status

    Passed, Failed, Warning, Error, Off, or All

    Filter by Type

    Scheduled, Not Scheduled, Event Triggered, or All

    Filter by Entity Type

    VM, Host, Disk, Storage Pool, Storage Container, Cluster Service, or All

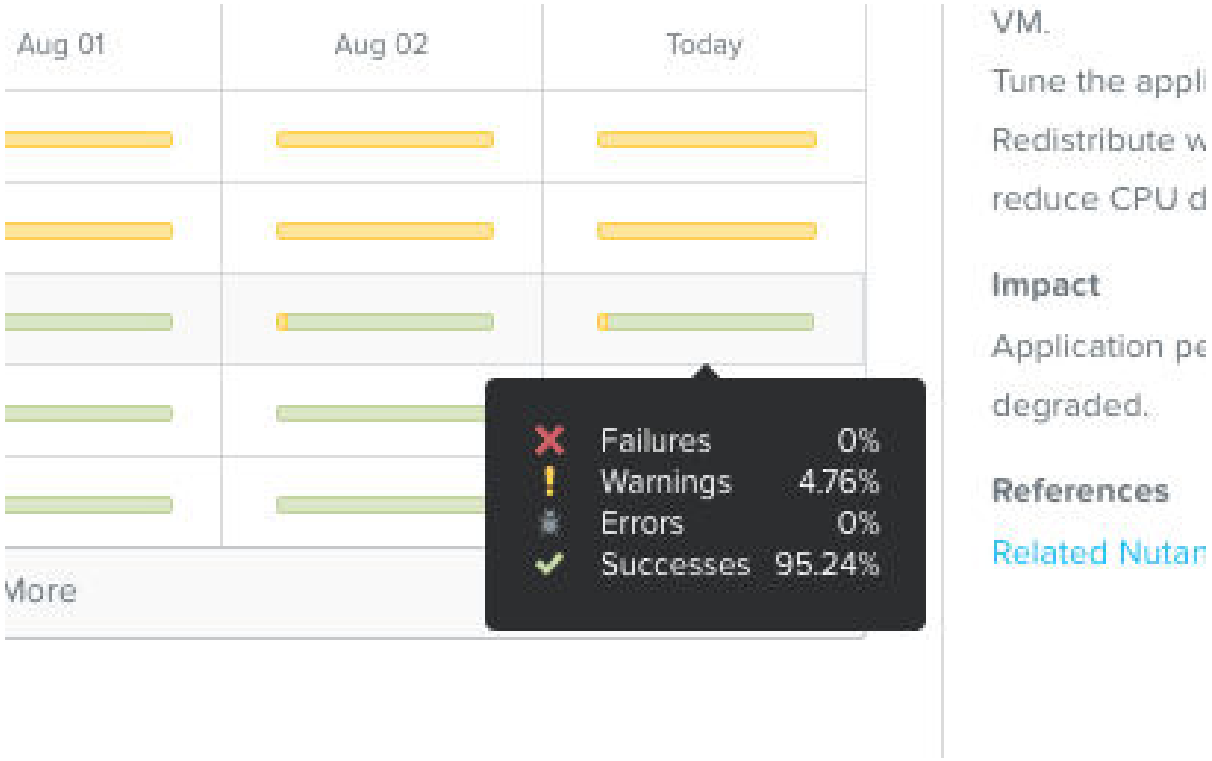

Figure 187: Hover Information

For example, if you want to see only the failed checks, filter the checks by selecting the **Failed** option. If you click on the specific check, the middle column will provide the detailed history of when the checks failed and what is the percentage of the check failure. If you click the bar, a detailed graph of the pass and fail history is displayed (as shown

below). Hovering the mouse along the graph line displays information about that point in time.

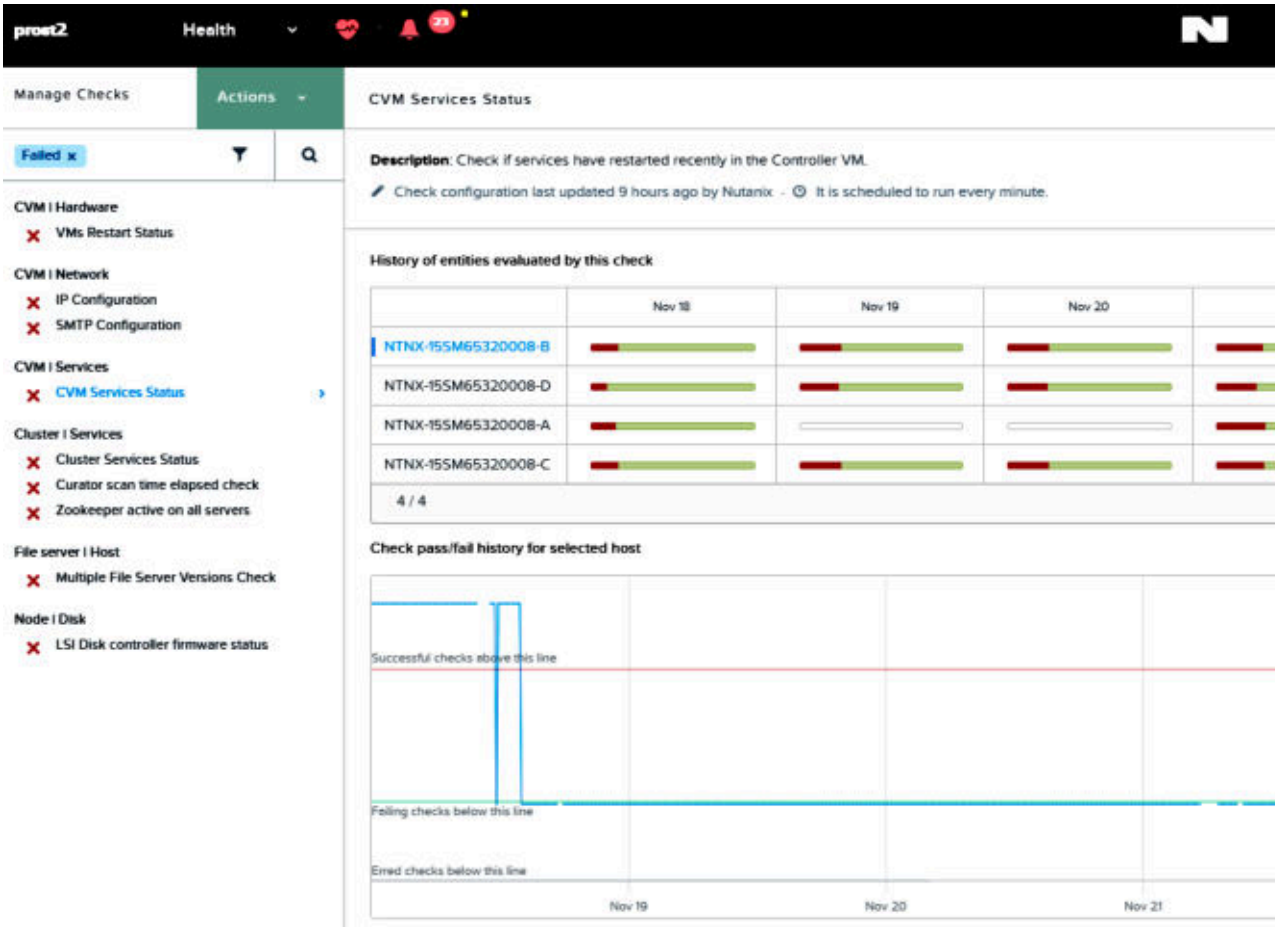

Figure 188: Filter Categorization

Note: For the checks with status as error, follow the similar process as described above to get detailed information about the errors.

You can also search for specific checks by clicking the

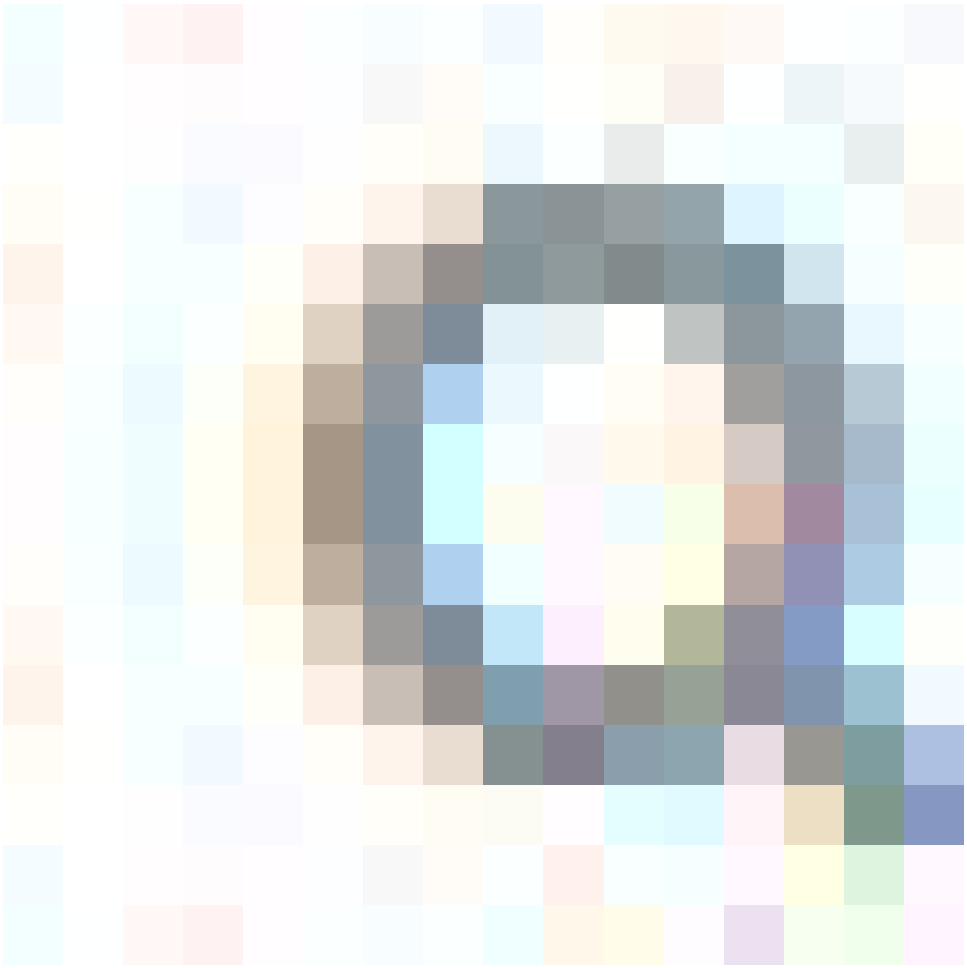

icon and then entering a string in the search box.

- The **Actions** tab provides you with an option to manage checks, run checks, and collect logs.

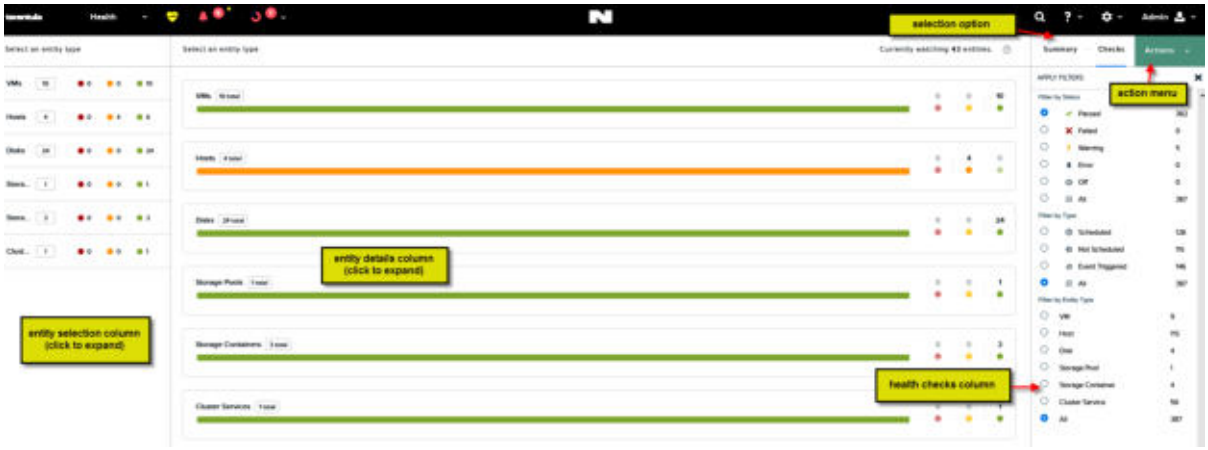

Figure 189: Health Dashboard

Focus and Filtering Options

The Health dashboard allows you to display entity health information through various views. Clicking a left column tab expands that tab to display grouping categories for that entity type (VMs, hosts, or disks). The middle section also expands with additional detail. The **Checks** tab of the right column displays the health checks that are relevant for that entity type.

Figure 190: Health Dashboard Expanded

Clicking on a grouping category displays detailed information about that grouping:

- The left column expands to display a set of grouping and filter options. The selected grouping is highlighted. You can select a different grouping by clicking on that grouping. Each grouping entry lists how many categories are in that grouping, and the middle section displays information about those categories. In the following example, the disks storage tier is selected, and there are two categories (SSD and HDD) in that grouping. By default, all entities (in this example, all disks) are included in the category information. You can narrow the included list by clicking one or more of the filters.
- The middle column displays a field for each category in the selected grouping. Each field provides details about that category. You can see additional information by hovering the cursor over a specific entry. There is a drop-down select list for filtering (same as the grouping filters) and a drop-down sort by list for ordering the information.
- The right column continues to display the health checks that are relevant for that entity type.

Figure 191: Health Dashboard Grouping Screen

The middle column provides two viewing options: a diagram view (see figure above) and a table view (see figure below). The table view provides more detailed information in tabular form. You can sort the entries by clicking a column header.

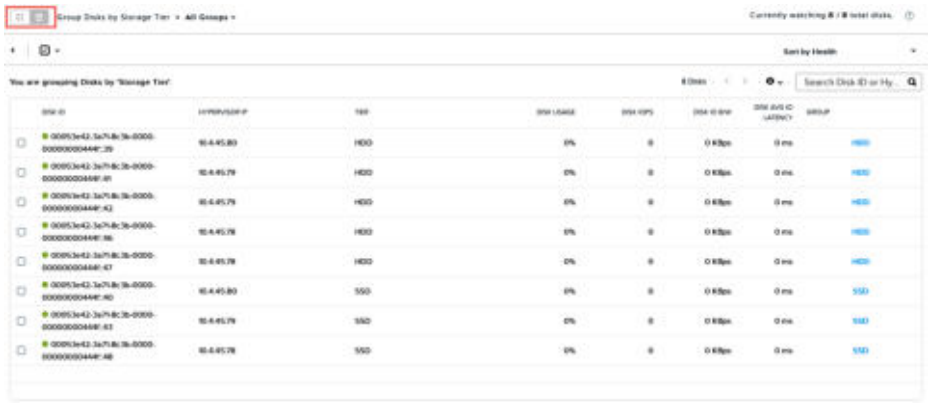

Figure 192: Health Dashboard Table View (middle column)

The middle column also includes watch list information at the top ("Currently watching *xx* entities" or "Currently watching *x* / *xx* total *entity_type*"). The Health dashboard is dynamically adjusted to reflect information about the entities in the current watch list. In this example, all disks are currently selected for the watch list ("Currently watching 18 / 18 total disks"), so the status information (middle column) and relevant health checks (right column) reflect the 18 disks. When you change the watch list to a subset of the current entity type (such as a single disk) or a different entity type (such as hosts), the status information and relevant health checks are customized accordingly for the new watch list.

Configuring Health Checks

A set of health checks are run regularly that provide a range of clusters health indicators. You can specify which checks to run and configure the schedulable checks and other parameters for each health check.

About this task

The cluster health checks cover a range of entities including AOS, hypervisor, and hardware components. A set of checks are enabled by default, but you can run, disable, or reconfigure any of the checks at any time. To reconfigure one or more health checks, do the following:

Procedure

1. In the Health dashboard (see Health Dashboard on page 287), click **Actions** > **Manage Checks** .
   The Health dashboard redisplays with information about the health checks. If you clicked on a specific health check, that check is highlighted. Either you are prompted to select a health check (first time) or the previously selected health check is highlighted.

2. Select a check.

   a. The left column lists the health checks with the selected check highlighted. Click any of the entries to select and highlight that health check.

   b. The middle column describes what this health check does, and it provides the run schedule and history across affected entities (hosts, disks, or VMs).

   c. The right column describes what failing this health check means (cause, resolution, and impact).

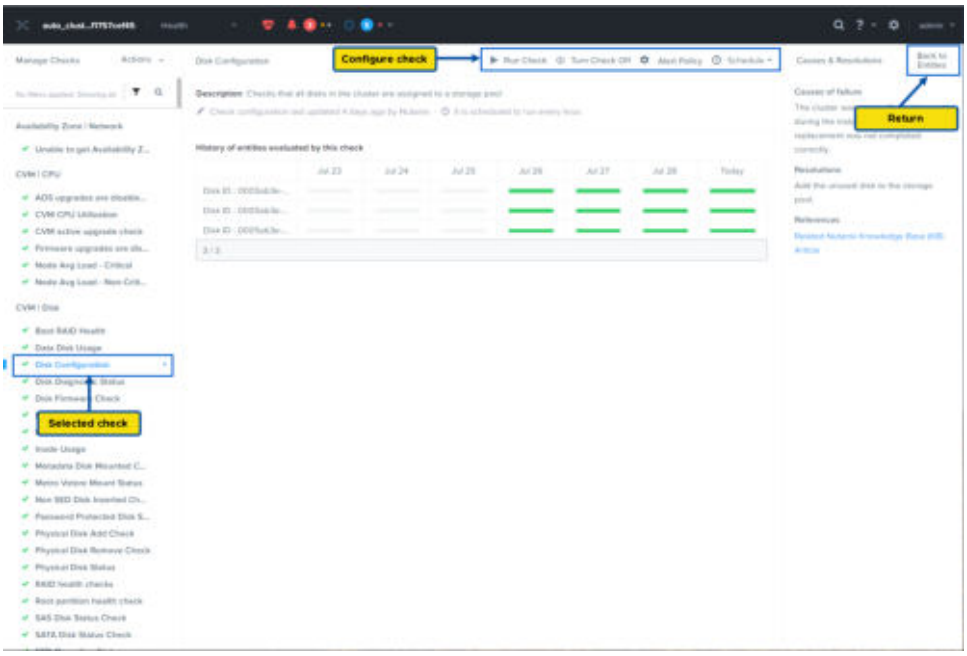

Figure 193: Health Check Configuration Screen

3. To run a particular check, click **Run Check**.

4. To turn off (or turn on) a health check, click the **Turn Check Off** (or **Turn Check On**) link at the top of the middle column and then click the **Yes** button in the dialog box.

5. To change the alert policy for those health checks that have a configurable policy, click Alert Policy at the top of the middle column, change one or more of the parameter values in the drop-down window, and then click **Save**.

   a. Depending on the alert condition (**Info**, **Critical**, and so on), unselect the condition to disable these alert condition messages.

      All the alerts are enabled by default (box checked). 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 90%) and a critical alert when a second threshold is reached (95%). 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.)

6. To change a parameter setting for those health checks that have configurable parameters, click the **Parameters** link at the top of the middle column, change one or more of the parameter values in the drop-down window, and then click the **Update** button.

   This link appears only when the health check includes configurable parameters. The configurable parameters are specific to that health check. For example, the CPU Utilization health check includes parameters to specify the host average CPU utilization threshold percentage and host peak CPU utilization threshold percentage.

Figure 194: Health Check Parameters

7. To change the schedule for running the health check, click the **Schedule** link for the schedulable checks at the top of the middle column, and select an interval from the drop-down list.

   You can choose an interval from 1 minute to 48 hours. Each check has a default interval, which varies from as short as once a minute to as long as once a day depending on the health check. The default intervals are optimal in most cases, and changing the interval is not recommended (unless requested to do so by Nutanix customer support).

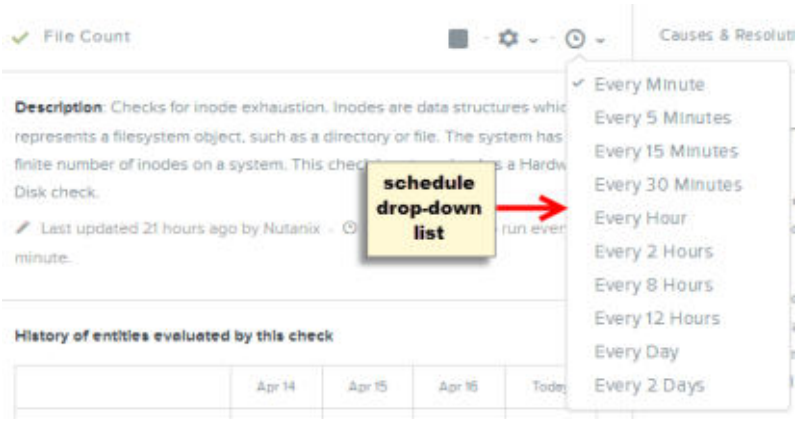

Figure 195: Health Check Schedule

Configuring NCC Frequency

Perform the following procedure to configure NCC to run automatically after the specified period of time from Prism.

Procedure

1. In the Health dashboard, from the **Actions** drop-down menu select **Set NCC Frequency**.
2. Select the configuration schedule.
   - » **Every 4 hours**: Select this option to run the NCC checks at four hours interval.
   - » **Every Day**: Select this option to run the NCC checks on a daily basis.

     Select the time of the day when you want to run the checks from **Start Time** field.
   - » **Every Week**: Select this option to run the NCC checks on a weekly basis.

     Select the day and time of the week when you want to run the checks from the **On** and **Start Time** fields. For example, if you select Sunday and Monday from the **On** field and select 3:00 p.m. from the **Start Time** field, every Sunday and Monday at 3 p.m. the NCC checks are run automatically.

   The Email address that you have configured by using Configuring Alert Emails on page 420 is also displayed. A report will be sent as an email to all the recipients.
3. Click **Save**.

Running Checks by Using Web Console

You can now run the NCC checks from the **Health** dashboard of the Prism web console. You can select to run all the checks at once, the checks that have failed or displayed some warning, or even specific checks of your choice.

About this task

> Note: If you are running checks by using web console, you will not be able to collect the logs at the same time.

Procedure

1. In the Health dashboard, from the **Actions** drop-down menu select **Run NCC Checks**.
2. Select the checks that you want to run for the cluster.
   a. **All checks**: Select this option to run all the checks at once.
   b. **Only Failed and Warning Checks**: Select this option to run only the checks that were failed or gave warning during the health check runs.
   c. **Specific Checks**: Select this option and type the check or checks name in the text box that appears that you want to run.

      This field gets auto-populated once you start typing the name of the check. All the checks that you have selected for this run are listed in the **Added Checks** box.

3. Select the **Send the cluster check report in the email** option to receive the report after the cluster check.

   To receive the email configuration ensure that you have configured email configuration for alerts. For more information, see Configuring Alert Emails on page 420.

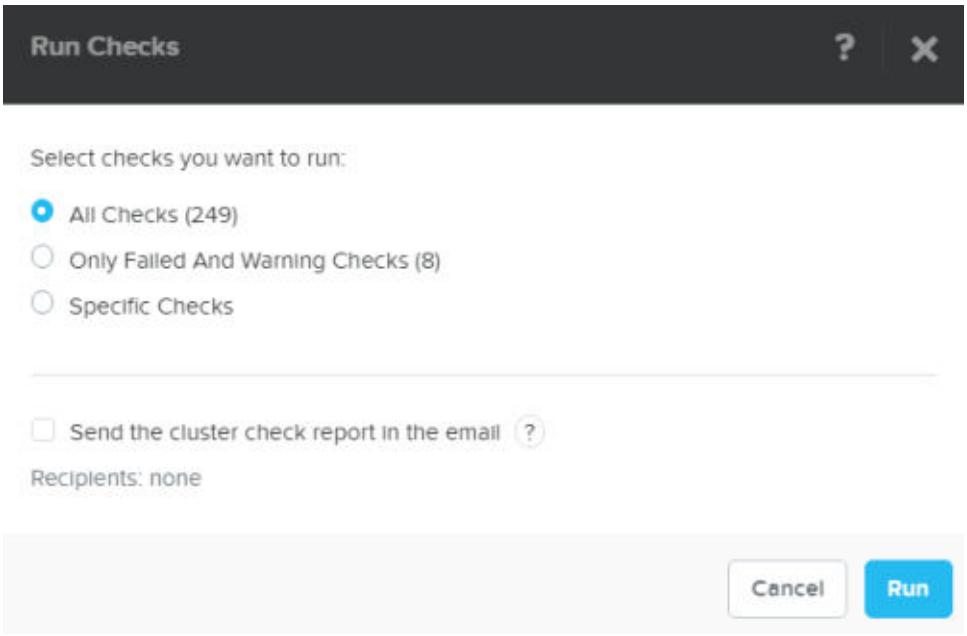

Figure 196: Run Health Checks

The status of the run (succeeded or aborted) is available in the **Tasks** dashboard. By default, all the event triggered checks are passed. Also, the **Summary** page of the **Health** dashboard will be updated with the status according to health check runs.

Collecting Logs by Using Web Console

You can collect logs for Controller VMs, file server, hardware, alerts, hypervisor, and for the system. After the task finishes, the log bundle is available for download purpose from the **Tasks** dashboard.

About this task

Log bundle includes logs and configuration information from one or more Controller VMs, configuration information for hypervisors, and information about alerts and so on. To collect the logs and download the log bundle, perform the following procedure.

Note:

- While this method works, the preferred method for collecting logs is through a CLI tool called `logbay`. See the *Nutanix Cluster Check (NCC) Guide* for information about `logbay`.
- The timestamps for all Nutanix service logs are moved to UTC (in ISO 8601:2020-01-01 T00:00:00Z) from Prism version 5.18.
- All operating system logs will not be moved to UTC, hence Nutanix recommends that you set the server local time to UTC.

Procedure

1. In the **Health** dashboard, from the **Actions** drop-down menu, select **Collect Logs**.
2. In **Node Selection**, click **+ Select Nodes**. Select the nodes for which you want to collect the logs and click **Done**.
3. Click **Next**.
4. In **Log Settings**, select one of the following:
   - » **All**. Select this option if you want to collect the logs for all the tags.
   - » **Specific (by tags)**. Select this option, click **+ Select Tags** if you want to collect the logs only for the selected tags and then click **Done**.

5. In **Output Preferences**, do the following in the indicated fields.

   - 1. **Select Duration**. Select the duration for which you want to collect the logs. You can collect the logs either in hours or days. Click the drop-down list to select the required option.
     2. **Cluster Date**. Select the date from which you want to start the log collection operation. Click the drop-down list to select either **Before** or **After** to collect logs before or after a selected date.
     3. **Cluster Time**. Select the time from when you want to start the log collection operation.
     4. **Select Destination for the collected logs**. Click the drop-down list to select the server where you want the logs to be collected.

        - **Download Locally**
        - **Nutanix Support FTP**. If you select this option, enter the case number in the **Case Number** field.
        - **Nutanix Support SFTP**. If you select this option, enter the case number in the **Case Number** field.
        - **Custom Server**. Enter server name, port, username, password, and archive path if you select this option.
     5. **Anonymize Output**. Select this option if you want to mask all the sensitive information like the IP addresses.

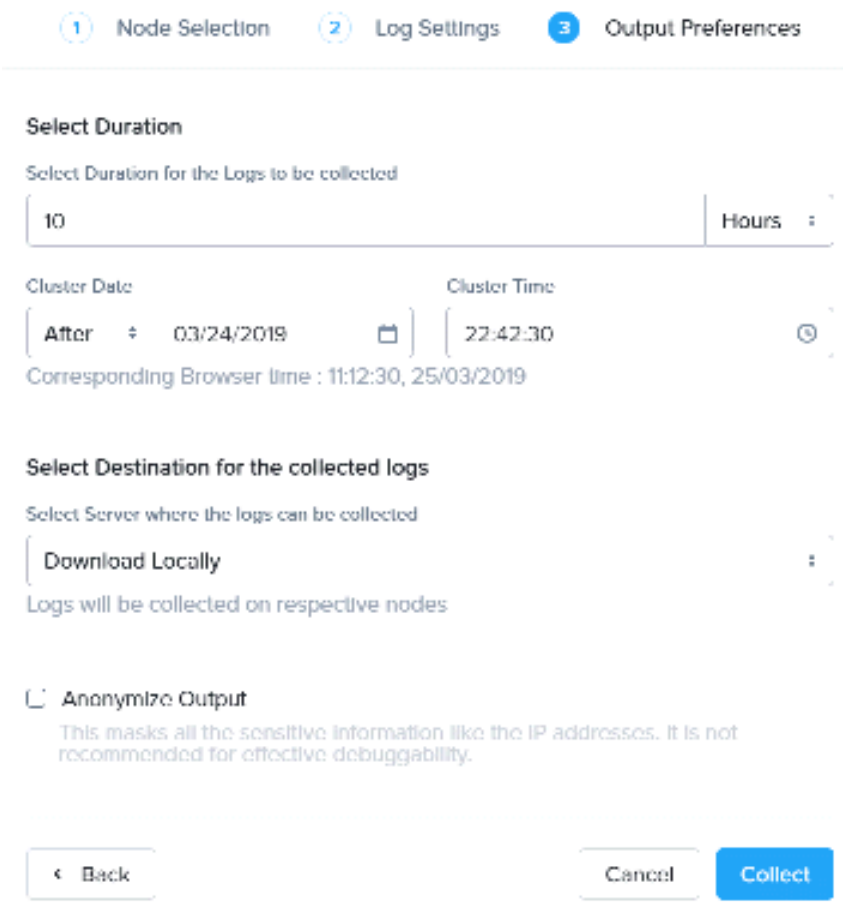

Figure 197: Collect Logs

6. To start the operation, click **Collect**.
7. After the operation completes, you can download the log bundle for the last two runs and (as needed) add it to a support case as follows:
   a. Go to the **Task** dashboard (see View Task Status on page 747), find the log bundle task entry, and click the **Succeeded** link for that task (in the Status column) to download the log bundle.

      Note: If a pop-up blocker in your browser stops the download, turn off the pop-up blocker and try again.

   b. Log in the support portal, click on the target case in the **360 View** widget on the dashboard (or click the **Create a New Case** button to create a new case), and upload the log bundle to the case (click the **Choose Files** button in the **Attach Files** section to select the file to upload).

10

VIRTUAL MACHINE MANAGEMENT

See also Recycle Bin on page 179.

Each node in a cluster includes local storage (flash and hard disk), a Controller VM, a hypervisor, and any number of host VMs running on that hypervisor.

- The web console allows you to monitor status of the VMs across the cluster (see VM Dashboard on page 300).
- In Acropolis managed clusters, the web console also allows you to do the following:
  - Create VMs (see Creating a VM (AHV) on page 316).
  - Manage VMs (see Managing a VM (AHV) on page 326).
  - Enable VM high availability (see Enabling High Availability for the Cluster on page 366).
  - Configure network connections (see Network Configuration for VM Interfaces on page 194).
- You can create and manage VMs directly from Prism Element when the hypervisor is ESXi.
  - Create VMs (see Creating a VM (ESXi) on page 347).
  - Manage VMs (see Managing a VM (ESXi) on page 350)

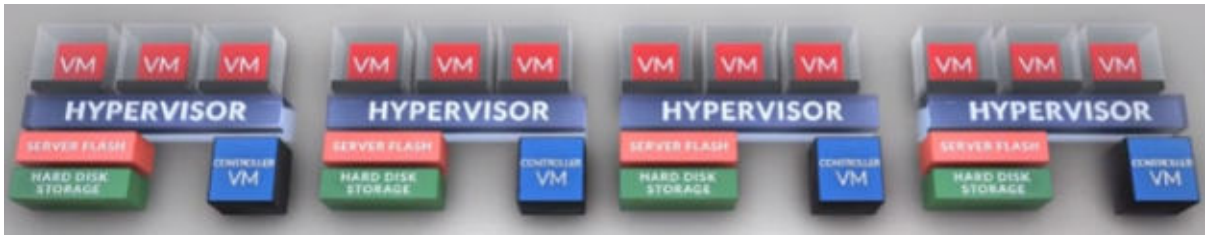

Figure 198: Node Architecture

VM Dashboard

The virtual machine (VM) dashboard displays dynamically updated information about virtual machines in the cluster. To view the VM dashboard, select **VM** from the pull-down list on the left of the main menu.

Menu Options

In addition to the main menu (see Main Menu on page 46), the VM screen includes a menu bar with the following options:

- View selector. Click the **Overview** button on the left to display the VM dashboard (see VM Overview View on page 302), or click the **Table** button to display VM information in a tabular form (see VM Table View on page 303).
- Action buttons. Click the **Create VM** button on the right to create a virtual machine (see Creating a VM (AHV) on page 316). Click the **Network Config** button to configure a network connection (see Network Configuration for VM Interfaces on page 194).

  Note: The action buttons appear only in Acropolis managed clusters.

- CVM filter. In the Table view, the Controller VMs are not listed by default. To include them in the table list, check the **Include Controller VMs** box.
- Page selector. In the Table view, VMs are listed 10 per page. When there are more than 10 VMs, left and right paging arrows appear on the right, along with the total VM count and the VM numbers for the current page.
- Export VM information. In the Table view, you can export the table containing the list of VMs and their information to a file in either CSV or JSON format by clicking the gear icon

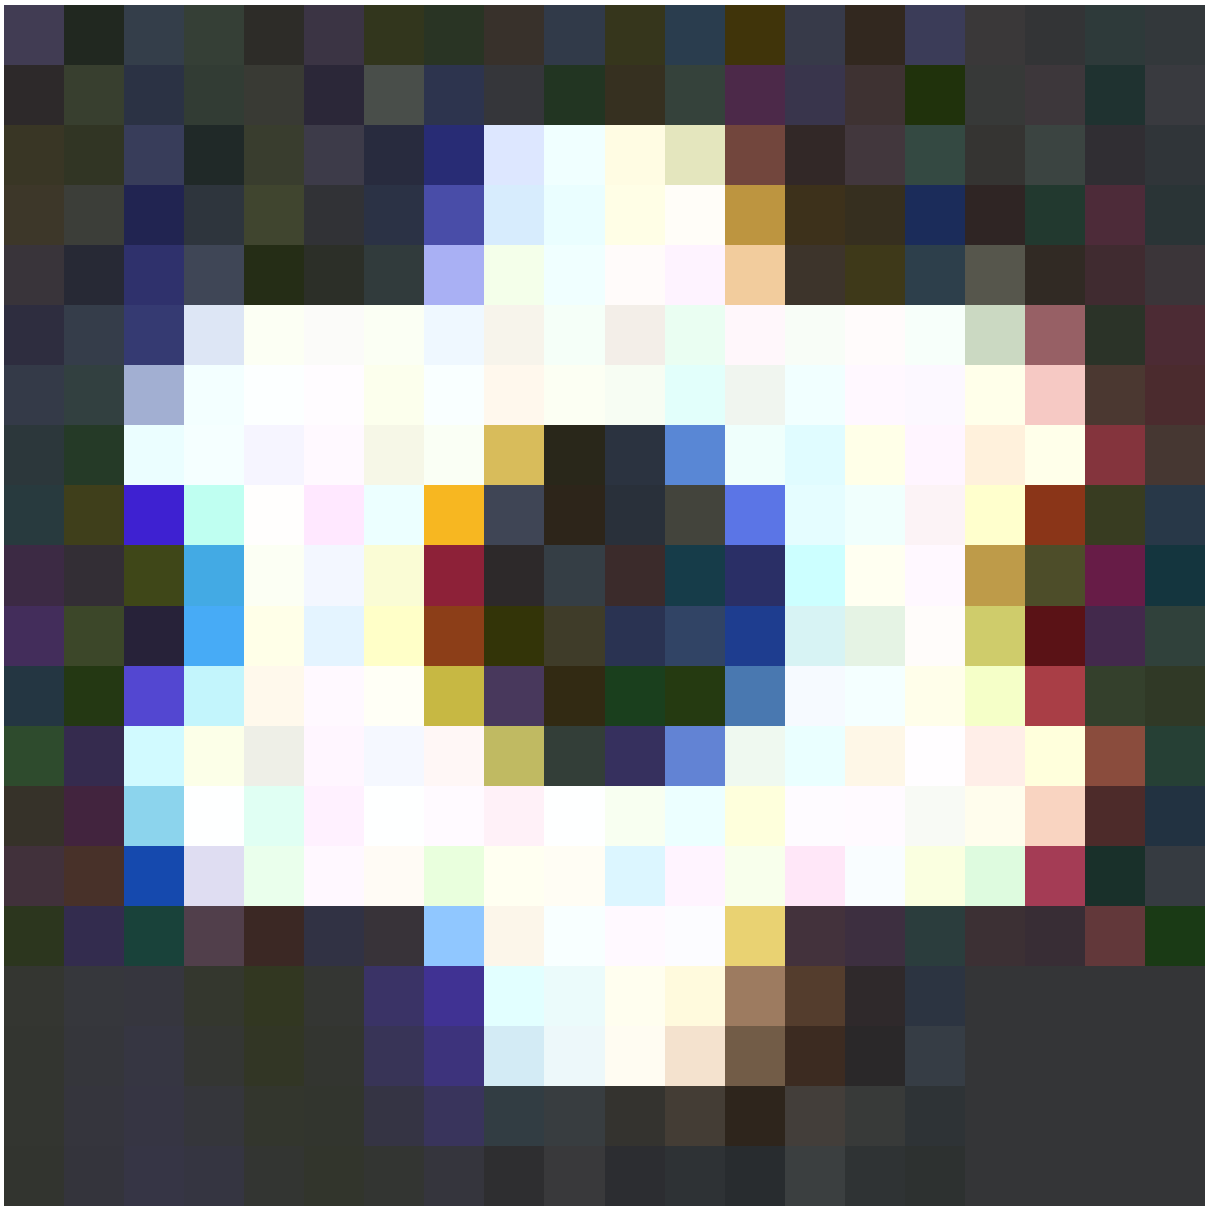

on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.

- Search box. In the Table view, you can search for entries in the table by entering a search string in the box.

Figure 199: VM Dashboard Menu

VM Overview View

The VM Overview view displays VM-specific performance and usage statistics on the left plus the most recent VM-specific alert and event messages on the right. The following figure is a sample view, and the table describes each field in this view. 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.

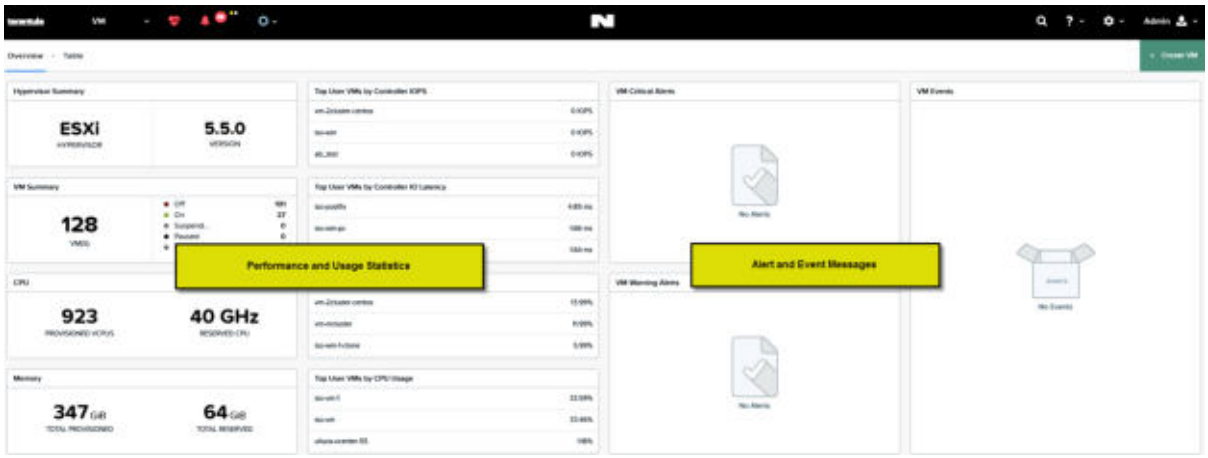

Figure 200: VM Overview View

Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived.

Table 37: VM Overview View Fields

Name	Description
Hypervisor Summary	Displays the name and version number of the hypervisor.
VM Summary	Displays the total number of VMs in the cluster broken down by on, off, and suspended states.

Name	Description
CPU	Displays the total number of provisioned virtual CPUs and the total amount of reserved CPU capacity in GHz for the VMs.
Memory	Displays the total amount of provisioned and reserved memory in GBs for the VMs.
Top User VMs by Controller IOPS	Displays I/O operations per VM for the 10 most active VMs.
Top User VMs by Controller IO Latency	Displays I/O bandwidth used per VM for the 10 most active VMs. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume.
Top User VMs by Memory Usage	Displays the percentage of reserved memory capacity used per VM for the 10 most active VMs.
Top User VMs by CPU Usage	Displays the percentage of reserved CPU capacity used per VM for the 10 most active VMs.
VM Critical Alerts	Displays the five most recent unresolved VM-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the "view all alerts" button at the bottom of the list (see Alerts Dashboard).
VM Warning Alerts	Displays the five most recent unresolved VM-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the "view all alerts" button at the bottom of the list.
VM Events	Displays the ten most recent VM-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the "view all events" button at the bottom of the list.

VM Table View

The VM Table view displays information about each VM in a tabular form. The displayed information is dynamically updated to remain current. In Acropolis managed clusters, you can both monitor and manage VMs through the VM Table view.

Table View Fields

The VM Table view is divided into two sections:

- The top section is a table. Each row represents a single VM and includes basic information about that VM. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a summary or details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

The following table describes each field in the table portion of the view. The details portion and tab contents are described in the subsequent sections.

> Note: See Understanding Displayed Statistics on page 64 for information about how the statistics are derived. VirtIO must be installed in a VM for AHV to display correct VM memory statistics.

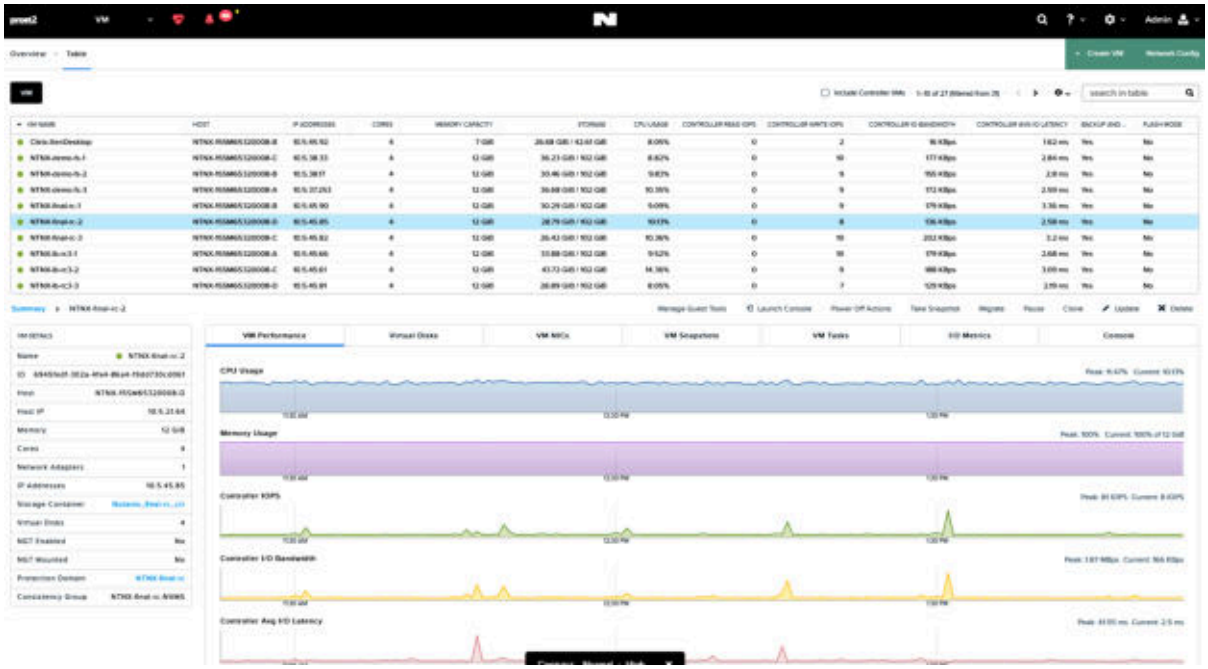

Figure 201: VM Table View

Table 38: VM Table View: Table Fields

Parameter	Description	Values
VM Name	Displays the name given the VM.	(VM name)
Host	Displays the name of the host	(Host name)
IP Addresses	Displays the IP address assigned to the VM.	(IP address)
Cores	Displays the number of CPU cores being used by this VM.	(number)
Memory Capacity	Displays the total amount of memory available to this VM.	xxx [MB | GB]
Storage	Displays the used capacity in relation to the total capacity (for example 1.9 GiB/ 5 GiB).	xxx / xxx [MiB | GiB]
CPU Usage	Displays the percentage of allocated CPU capacity currently being used by this VM.	0 - 100%
Memory Usage	Displays the percentage of allocated memory capacity currently being used by this VM.	0 -100%

Parameter	Description	Values
[Controller] Read IOPS	Displays read I/O operations per second (IOPS) for this VM. Note: In this and the following three fields, the column name includes the word Controller if the information comes from the Controller VM instead of the hypervisor. For ESXi, the information comes from the hypervisor; for Hyper-V and AHV, the information comes from the Controller VM.	(number)
[Controller] Write IOPS	Displays write I/O operations per second for this VM.	(number)
[Controller] IO Bandwidth	Displays I/O bandwidth used per second for this VM.	xxx [MBps|KBps]
[Controller] Avg IO Latency	Displays the average I/O latency for this VM.	xxx [ms]
Backup and Recovery Capable	Indicates (Yes or No) whether the VM can be protected (create backup snapshots) and recovered if needed. When the value is No, click the question mark icon for an explanation.	[Yes|No]
Flash Mode	Displays whether flash mode feature is enabled or not for the VM	[Yes|No]

VM Detail Information

When a VM is selected in the table, information about that VM appears in the lower part of the screen.

- **Summary:** ***vm_name*** appears below the table and **VM Details** fields appear in the lower left column. The following table describes the fields in this column.

- For VMs in Acropolis managed clusters, action links appear on the right of the **Summary: *vm_name*** line (see Managing a VM (AHV) on page 326 for more information about these actions):
  - Click **Manage NGT** to enable and mount Nutanix guest tools for this VM.
  - Click the **Launch Console** link to open a console window for this VM.
  - Click the **Power on** (or **Power Off Actions**) link to start (or shut down) this VM.
  - Click the **Take Snapshot** link to create a backup snapshot on demand.
  - Click the **Migrate** link to migrate the VM onto a different host.
  - Click the **Pause** (or **Resume**) link to pause (or resume) this VM.

    Note: VM pause and resume feature is not supported on AHV.
  - Click the **Clone** link to clone a copy of the VM.
  - Click the **Update** link to update the VM configuration.
  - Click the **Delete** link to delete the VM. (A VM must be powered off before it can be deleted.)

    Note: If the VM is a Prism Central VM, all Actions except **Launch Console**, **Take Snapshot**, and **Migrate** are disabled as a protective measure.
- A set of tabs appear to the right of the details section that display information about the selected VM. The set of tabs varies depending on whether the VM is an Acropolis managed VM or not. The following sections describe each tab.
  - Standard VM tabs: **VM Performance**, **Virtual Disks**, **VM NICs**, **VM Alerts**, **VM Events**, **I/O Metrics**, and **Console**.
  - Acropolis managed VM tabs: **VM Performance**, **Virtual Disks**, **VM NICs**, **VM Snapshots**, **VM Tasks**, **I/O Metrics**, and **Console**.

Figure 202: VM Table View: VM Details

Table 39: VM Detail Fields

Parameter	Description	Values
Name	Displays the name given the VM.	(VM name)

Parameter	Description	Values
Host	Displays the host name on which this VM is running.	(IP address)
Host IP	Displays the host IP address for this VM.	(IP address)
Guest OS	Displays the operating system running on this VM, such as Windows 7 or Ubuntu Linux. (This information is not available when running AHV.)	(operating system name)
Memory	Displays the amount of memory available to this VM.	xxx [MB|GB]
Reserved Memory	Displays the amount of memory reserved for this VM (by the hypervisor).	xxx [MB|GB]
Assigned Memory (Hyper-V only)	Displays the amount of dynamic memory currently assigned to the VM by the hypervisor.	xxx [MB|GB]
Cores	Displays the number of CPU cores being used by this VM.	(number)
Reserved CPU	Displays the amount of CPU power reserved for this VM (by the hypervisor).	xxx [GHz]
Disk Capacity	Displays the total disk capacity available to this VM.	xxx [GB|TB]
Network Adapters	Displays the number of network adapters available to this VM.	(# of adapter ports)
IP Addresses	Displays the IP address assigned to the VM.	(IP address)
Storage Container	Displays the name of the storage container in which the VM resides.	(storage container name)
Virtual Disks	Displays the number of virtual disks in the VM.	(number)
NGT Enabled	Displays whether NGT is enabled or not for the VM.	[Yes|No]
NGT Mounted	Displays whether NGT is mounted or not for the VM.	[Yes|No]
GPU Configuration	(AHV only) Comma-separated list of GPUs configured for the VM. GPU information includes the model name and a count in parentheses if multiple GPUs of the same type are configured for the VM. If the firmware on the GPU is in compute mode, the string `compute` is appended to the model name. The field is hidden if no GPUs are configured or if the hypervisor is not AHV.	(list of GPUs)
GPUs in Use	(AHV only) Number of GPUs in use by a VM. The field is hidden if no GPUs are configured or if the hypervisor is not AHV.	(number)

Parameter	Description	Values
VMware Guest Tools Mounted	Displays whether VMware guest tools are mounted or not on the VM	[Yes|No]
VMware Guest Tools Running Status	Displays whether VMware guest tools are running or not on the VM.	[Yes|No]

Cluster Summary Information

When a VM is not selected in the table (or when the word **Summary** is clicked), summary information across all VMs in the cluster appears in the lower part of the screen.

- The **VM Summary** fields appear in the lower left column. The following table describes the fields in this column.
- Three tabs appear that display cluster-wide information (see following sections for details about each tab): **Performance Summary**, **All VM Alerts**, **All VM Events**.

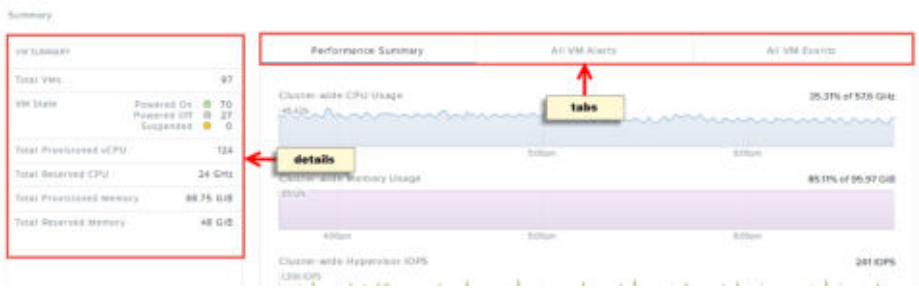

Figure 203: VM Table View: VM Summary

Table 40: VM Summary Fields

Parameter	Description	Values
Total VMs	Displays the total number of VMs in the cluster.	(number)
VM State	Displays the number of powered on, powered off, and suspended VMs in the cluster.	[number] powered on, powered off, suspended
Total Provisioned vCPU	Displays the total number of provisioned virtual CPUs in the cluster.	(number)
Total Reserved CPU	Displays the total amount of CPU power reserved for the VMs (by the hypervisor).	xxx [GHz]
Total Provisioned Memory	Displays the total amount of memory provisioned for all VMs.	xxx [GB]
Total Reserved Memory	Displays the total amount of memory reserved for all VMs (by the hypervisor).	xxx [GB]

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label varies depending on what is selected in the table:

- **Performance Summary** (no VM selected). Displays resource performance statistics (CPU, memory, and I/O) across all VMs in the cluster.
- **VM Performance** (VM selected). Displays resource performance statistics (CPU, memory, and I/O) for the selected VM.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see Analysis Dashboard on page 396). The Performance tab includes the following graphs:

- **[Cluster-wide] CPU Usage**: Displays the percentage of CPU capacity currently being used (0 - 100%) across all VMs or for the selected VM.
- **[Cluster-wide] Memory Usage**: Displays the percentage of memory capacity currently being used (0 - 100%) across all VMs or for the selected VM. (This field does not appear when the hypervisor is Hyper-V.)
- **[Cluster-wide] {Hypervisor|Controller} IOPS**: Displays I/O operations per second (IOPS) across all VMs or for the selected VM.

  Note: In this and the following two fields, the field name is either Controller or Hypervisor to indicate where the information comes from. For ESXi, the information comes from the hypervisor; for Hyper-V and AHV, the information comes from the Controller VM.

- **[Cluster-wide] {Hypervisor|Controller} I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) across all VMs or for the selected VM.
- **[Cluster-wide] {Hypervisor|Controller} Avg I/O Latency**: Displays the average I/O latency (in milliseconds) across all VMs or for the selected VM.

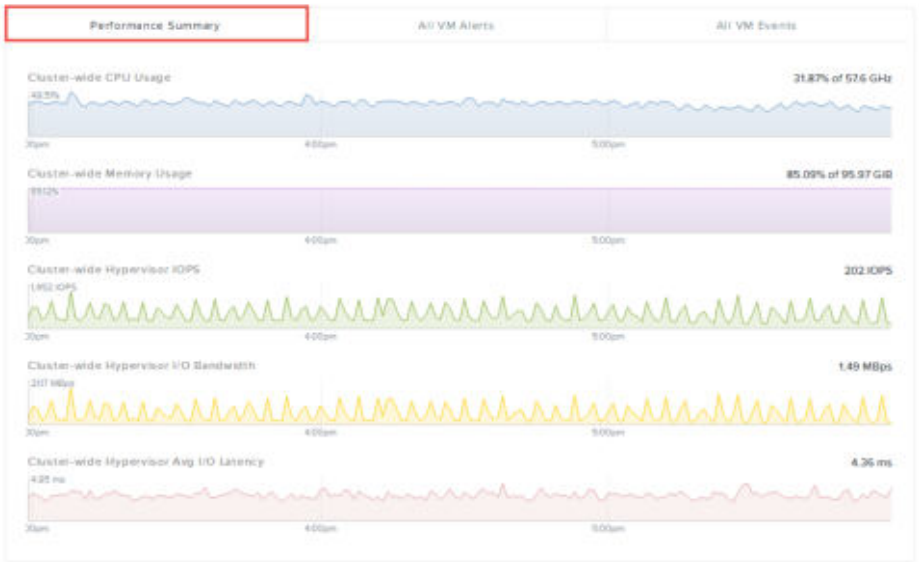

Figure 204: VM Table View: Performance Tab

Virtual Disks Tab

The Virtual Disks tab displays information in tabular form about the virtual disks in a selected VM. (This tab appears only when a VM is selected.) Each line represents a virtual disk, and the following information is displayed for each disk organized under "Default" and "Additional Stats" subtabs.

Default Tab:

> Note: Clicking on a virtual disk (line) displays subtabs for total, read, and write IOPS, I/O bandwidth, and I/O latency performance graphs for the virtual disk (see the "Performance Tab" section for more information about the graphs).

- **Virtual Disk**. Displays the virtual disk identification number.
- **Total Capacity**. Displays the total capacity of the virtual disk (in GiBs).
- **Physical Usage**. Displays the used space of the virtual disks (in GiBs).
- **Read IOPS**. Displays the read IOPS for the virtual disk.
- **Read BW**. Displays the bandwidth used by the virtual disk for read operations.
- **Read Latency**. Displays the average I/O latency for read requests to this virtual disk.
- **Write IOPS**. Displays the write IOPS for the virtual disk.
- **Write BW**. Displays the bandwidth used by the virtual disk for write operations.
- **Write Latency**. Displays the average I/O latency for write requests to this virtual disk.
- **Flash Mode**. Displays whether flash mode is enabled for the virtual disk or not.

Additional Stats Tab:

- **Total IOPS**. Displays the total (both read and write) I/O operations per second (IOPS) for the virtual disk.
- **Random IO**. Displays the percentage of I/O that is random (not sequential).
- **Read Source Cache**. Displays the amount of cache data accessed for read requests.
- **Read Source SSD**. Displays the amount of SSD data accessed for read requests.
- **Read Source HDD**. Displays the amount of HDD data accessed for read requests.
- **Read Working Size Set**. Displays the amount of data actively being read by applications in the VM that are using this virtual disk.
- **Write Working Size Set**. Displays the amount of data actively being written by applications in the VM that are using this virtual disk.
- **Union Working Size Set**. Displays the total amount of data used by the VM for either reads or writes.

Figure 205: Virtual Disks Tab

VM NICs Tab

The VM NICs tab displays information in tabular form about the virtual NICs in a selected VM. (This tab appears only when a VM is selected.) Each line represents a virtual NIC, and the following information is displayed for each NIC:

- **Virtual NIC**. Displays the virtual NIC identification number.
- **Adapter Type**. Displays the adaptor type defined for the virtual NIC.
- **MAC Address**. Displays the virtual NIC MAC address
- **IPv4 Addresses**. Displays the virtual NIC IPv4 address(es).
- **IPv6 Addresses**. Displays the virtual NIC IPv6 address(es).
- **Received Packets**. Displays the number of packets received by the virtual NIC.
- **Transmitted Packets**. Displays the number of packets transmitted by the virtual NIC.
- **Dropped Rx Packets**. Displays the number of received packets dropped by the virtual NIC.
- **Dropped Tx Packets**. Displays the number of transmitted packets dropped by the virtual NIC.

Figure 206: VM NICs Tab

When you click a virtual NIC entry, three more tabs appear below the list of virtual NICs. Clicking the **Virtual NICs Stats** tab displays the following statistics for that virtual NIC:

- **Total Packets Received**. Displays a monitor of the total packets received (in KB or MB) over time. Place the cursor anywhere on the line to see the value for that point in time. (This applies to all the monitors on this tab.)
- **Total Packets Transmitted**. Displays a monitor of the transmitted data rate.
- **Dropped Packets Received**. Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted**. Displays a monitor of transmitted packets that were dropped.
- **Error Packets Received**. Displays a monitor for error packets received.

Figure 207: VM NICs Tab: Virtual NIC Stats

Clicking the **Host NICs Stats** tab displays the following statistics for each host NIC (one per line) that is used by the selected virtual NIC to send the traffic:

- **Host NIC**. Displays the host NIC name.
- **Speed (in KBps)**. Displays the host NIC transmission speed.
- **MAC Address**. Displays the host NIC MAC address.
- **Received Packets**. Displays the number of packets received by the host NIC.
- **Transmitted Packets**. Displays the number of packets transmitted by the host NIC.
- **Dropped Rx Packets**. Displays the number of received packets dropped by the host NIC.
- **Dropped Tx Packets**. Displays the number of transmitted packets dropped by the host NIC.
- **Rx Packet Errors**. Displays the number of error packets received by the host NIC.
- **Tx Packet Errors**. Displays the number of error packets transmitted by the host NIC.

Figure 208: Virtual NICs Tab: Host NIC Stats

Clicking the **Physical Switch Interface Stats** tab displays the following statistics for each physical switch interface (one per line) used by the selected virtual NIC to send the traffic:

- **Physical Switch Interface**. Displays the switch interface name.
- **Switch ID**. Displays the switch interface ID value.
- **Index**. Displays the switch interface index number.
- **MTU (in Bytes)**. Displays the size in bytes of the largest protocol data unit (maximum transmission unit) that the layer can pass onwards.
- **MAC Address**. Displays the interface MAC address.
- **Unicast Rx Pkts**. Displays the number of unicast packets received.
- **Unicast Tx Pkts**. Displays the number of unicast packets transmitted.
- **Error Rx Pkts**. Displays the number of received packets with an error.
- **Error Tx Pkts**. Displays the number of transmitted packets with an error.
- **Discard Rx Pkts**. Displays the number of received packets that were discarded.
- **Discard Tx Pkts**. Displays the number of transmitted packets that were discarded.

Figure 209: Virtual NICs Tab: Physical Switch Interface Stats

VM Alerts Tab

The VM Alerts tab displays the unresolved alert messages about all VMs or the selected VM in the same form as the Alerts page (see Alerts Summary View on page 412). Click the **Unresolved X** button in the filter field to also display resolved alerts.

VM Events Tab

The VM Events tab displays the unacknowledged event messages about all VMs or the selected VM in the same form as the Events page (see Events Summary View on page 418). Click the **Include Acknowledged** button to also display acknowledged events.

VM Snapshots Tab (Acropolis only)

The VM Snapshots tab displays information in tabular form about backup snapshots of the VM. (This tab appears only when a VM is selected.) Each line represents 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.
- **Actions**. Displays four action links:
  - Click the **Details** link to open a window that displays the VM configuration (see Creating a VM (AHV) on page 316) plus a creation time stamp field.

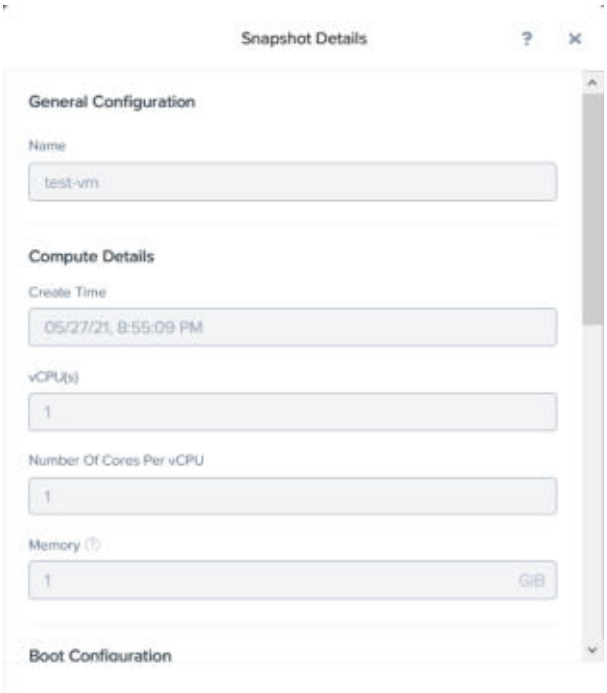

Figure 210: Snapshot Details Window

  - Click the **Clone** link to clone a VM from the snapshot (see Managing a VM (AHV) on page 326).
  - 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.

Figure 211: VM Snapshots Tab

VM Tasks Tab (Acropolis only)

The VM Tasks (selected VM) or All VM Tasks (cluster-wide) tab displays a log in tabular form of the running and completed tasks for the selected VM or across the cluster. Each line represents a task, and the following information is displayed for each task:

- **Operation**. Describes the type of task.
- **Entities**. Lists the affected VM and node.
- **Percent Complete**. Displays the run status (0-100%) of the task.
- **Progress Status**. Displays the completion status of the task (succeeded, in progress, failed).
- **Create Time**. Displays when the task began.
- **Duration**. Displays how long the task took to complete.

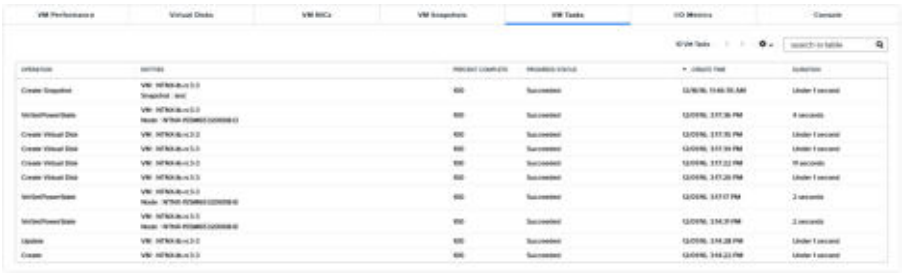

Figure 212: [All] VM Tasks Tab

I/O Metrics

The I/O Metrics tab displays information about different I/O metrics for the VM (latency and performance distribution).

- Average I/O latency for the VM. Displays a graph of average I/O latency (in milliseconds) for reads and writes over a period of time. If you hover the cursor over the graph, the read and write latency at that particular time is displayed.
- Performance distribution. Displays a bar chart for the performance distribution of read and write size (in bytes) and read and write latency (in milliseconds). This widget also provides information on the read source (HDD, SSD, or DRAM), random vs sequential read and writes in a pie chart format. This information changes according to the time that you select in the **Avg I/O Latency** widget.

Figure 213: I/O Metrics Tab

Console Tab

The Console tab displays a live console window. (This tab appears only when a VM is selected.) In addition to entering commands in the console window, you can invoke several options from this tab:

- Click the language (left-most) button and select the desired language from the pull-down list to set the language key mapping for the console keyboard.
- Click the **Send CtrlAltDel** button to send a Ctrl+Alt+Delete key signal. This is the same as pressing Ctrl+Alt+Delete from the console keyboard.
- Click the **Take Screenshot** button to take a screen shot of the console display that you can save for future reference.
- Click the **New Window** button to open a new console window. This is the same as clicking the **Launch Console** link on the Summary line.

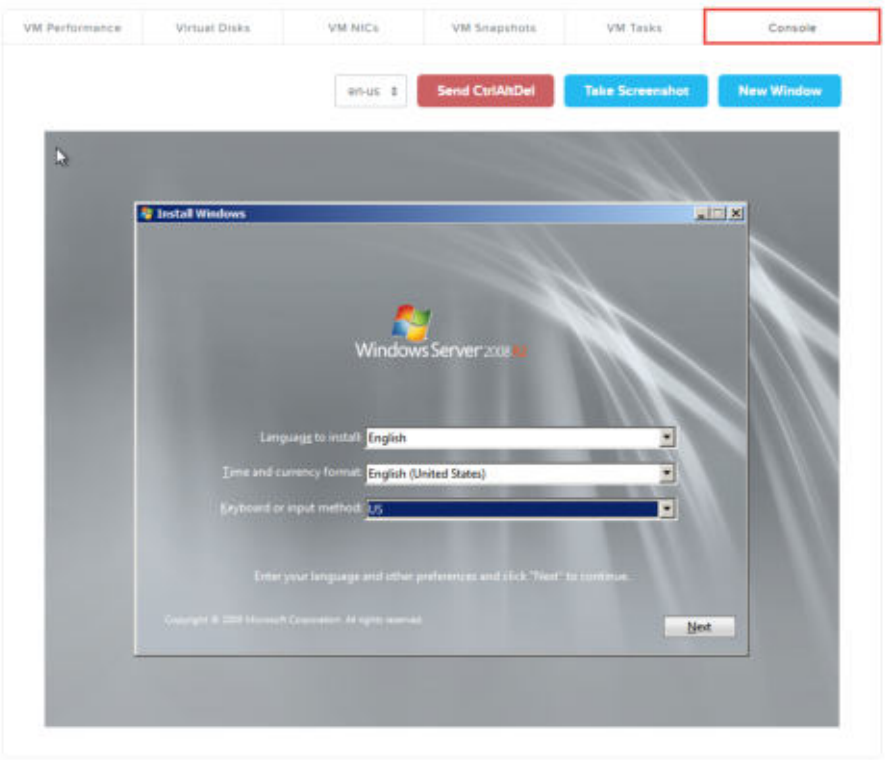

Figure 214: Console Tab

VM Management

You can create and manage VMs directly from Prism Element when the hypervisor is either ESXi or AHV. Following topics provide more information on creating and managing VM configuration on ESXi and AHV.

- ESXi
  - To create a VM, see Creating a VM (ESXi) on page 347.
  - To manage guest tools, launch VM console, power actions, clone, update, or delete operations, see Managing a VM (ESXi) on page 350.
- AHV
  - To create a VM, see Creating a VM (AHV) on page 316.
  - To manage guest tools, launch console, power actions, take snapshot, migrate, power operations, clone, update, or delete operations, see Managing a VM (AHV) on page 326.

Creating a VM (AHV)

In AHV clusters, you can create a new virtual machine (VM) through the Prism Element web console.

About this task

When creating a VM, you can configure all of its components, such as number of vCPUs and memory, but you cannot attach a volume group to the VM. Attaching a volume group is possible only when you are modifying a VM.

To create a VM, do the following:

Procedure

1. In the VM dashboard, click the **Create VM** button.

   Note: This option does not appear in clusters that do not support this feature.

   The Create VM dialog box appears.

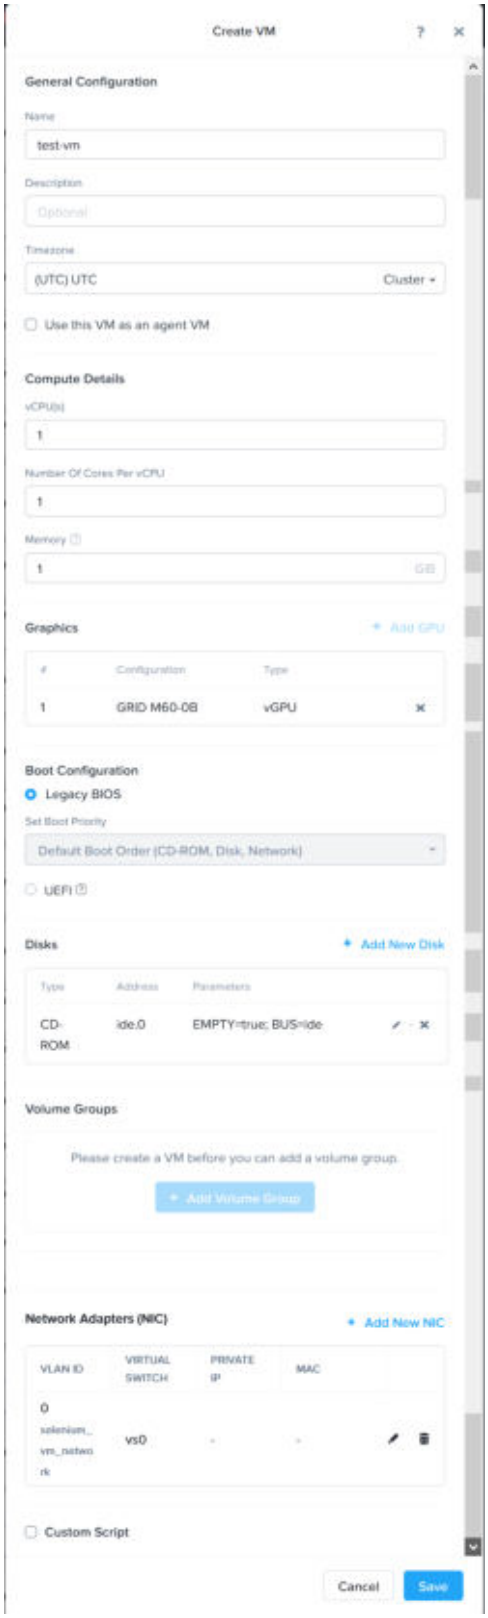

Figure 215: Create VM Dialog Box

2. 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. **Timezone**: Select the timezone that you want the VM to use. If you are creating a Linux VM, select **(UTC) UTC**.

      > Note:
      >
      > The RTC of Linux VMs must be in UTC, so select the UTC timezone if you are creating a Linux VM.
      >
      > Windows VMs preserve the RTC in the local timezone, so set up the Windows VM with the hardware clock pointing to the desired timezone.

   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 the rest of the VMs are powered on on the host) and must be powered off after the rest of the VMs are powered off (for example, during maintenance mode operations). Agent VMs are never migrated to any other host in the cluster. If an HA event occurs or the host is put in maintenance mode, agent VMs are powered off and are powered on on the same host once that host comes back to a normal state.

      If an agent VM is powered off, you can manually start that agent VM on another host and the agent VM now permanently resides on the new host. The agent VM is never migrated back to the original host. Note that you cannot migrate an agent VM to another host while the agent VM is powered on.

   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 GiB) to allocate to this VM.

3. (For GPU-enabled AHV clusters only) To configure GPU access, click **Add GPU** in the **Graphics** section, and then do the following in the Add GPU dialog box:

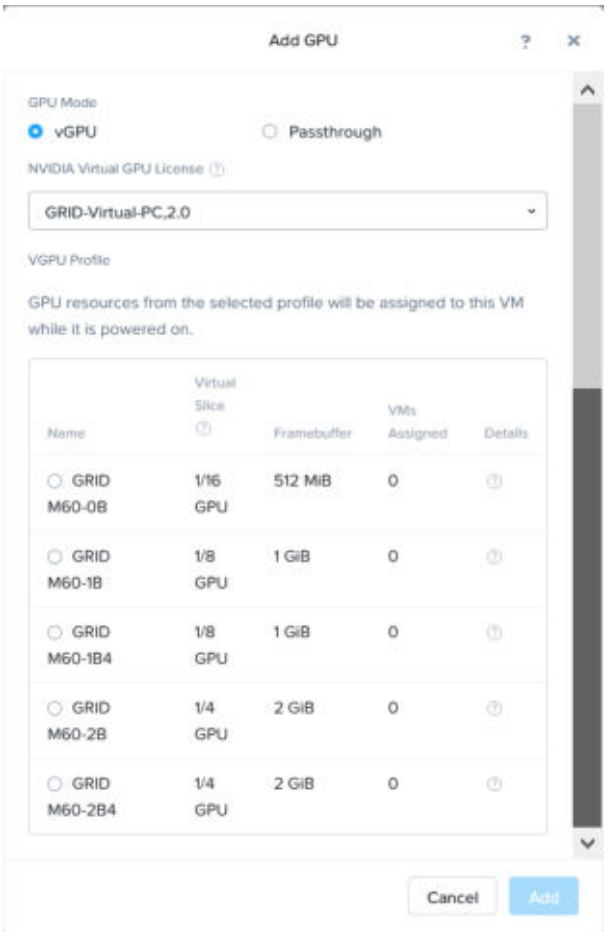

Figure 216: Add GPU Dialog Box

See *GPU and vGPU Support* in *AHV Administration Guide* for more information.

a. To configure GPU pass-through, in **GPU Mode**, click **Passthrough**, select the GPU that you want to allocate, and then click **Add**.

   If you want to allocate additional GPUs to the VM, repeat the procedure as many times as you need to. Make sure that all the allocated pass-through GPUs are on the same host. If all specified GPUs of the type that you want to allocate are in use, you can

proceed to allocate the GPU to the VM, but you cannot power on the VM until a VM that is using the specified GPU type is powered off.

See *GPU and vGPU Support* in *AHV Administration Guide* for more information.

b. To configure virtual GPU access, in **GPU Mode**, click **virtual GPU**, select a GRID license, and then select a virtual GPU profile from the list.

> Note: This option is available only if you have installed the GRID host driver on the GPU hosts in the cluster.
>
> See the NVIDIA Grid Host Driver for Nutanix AHV Installation Guide for the NVIDIA GRID host driver installation instructions.

You can assign multiple virtual GPU to a VM. A vGPU is assigned to the VM only if a vGPU is available when the VM is starting up.

Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.

> Note:
>
> Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.

After you add the first vGPU, to add multiple vGPUs, see *Adding Multiple vGPUs to the Same VM* in the *AHV Admin Guide*.

4. 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*.

   » Secure Boot is supported with AOS 5.16. The current support to Secure Boot is limited to the aCLI. For more information about Secure Boot, see *Secure Boot Support for VMs* section in the *AHV Administration Guide*. To enable Secure Boot, do the following:

   - Select UEFI.
   - Power-off the VM.
   - Log on to the aCLI and update the VM to enable Secure Boot. For more information, see *Updating a VM to Enable Secure Boot* in the *AHV Administration Guide*.

5. To attach a disk to the VM, click the **Add New Disk** button.
   The Add Disk dialog box appears.

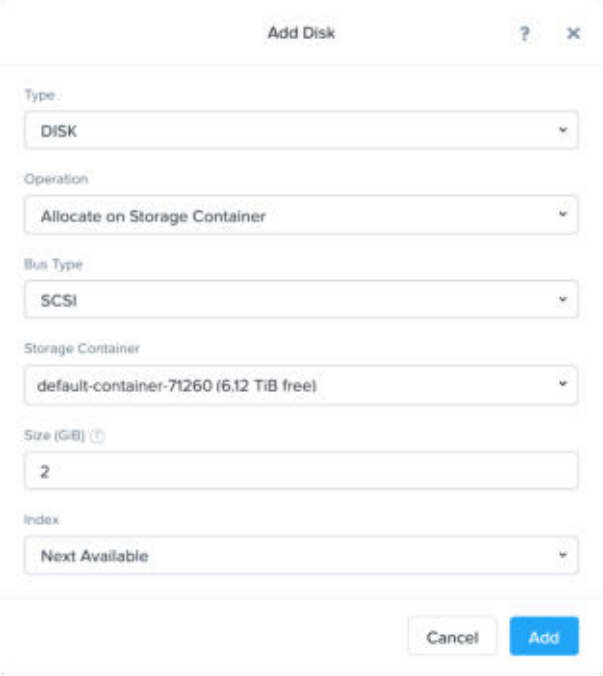

Figure 217: Add Disk Dialog Box

Do the following in the indicated fields:

a. **Type**: Select the type of storage device, **DISK** or **CD-ROM**, from the drop-down list.
   The following fields and options vary depending on whether you choose **DISK** or **CD-ROM**.

b. **Operation**: Specify the device contents from the drop-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 **Empty CD-ROM** to create a blank CD-ROM device. (This option appears only when **CD-ROM** is selected in the previous field.) A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
   - 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.
   - Select **Clone from Image Service** to copy an image that you have imported by using image service feature onto the disk. For more information on the Image Service

feature, see the *Configuring Images* section of *Prism Web Console Guide* and *Image Management* section in the *Prism Central Guide*.

c. **Bus Type**: Select the bus type from the pull-down list. The choices are **IDE**, **SCSI**, or **SATA**.

The options displayed in the **Bus Type** drop-down list varies based on the storage device Type selected in Step a.

- For device **Disk**, select from **SCSI**, **SATA**, **PCI**, or **IDE** bus type.
- For device **CD-ROM**, you can select either **IDE**,or **SATA** bus type.

Note: SCSI bus is the preferred bus type and it is used in most cases. Ensure you have installed the VirtIO drivers in the guest OS.

CAUTION: Use SATA, PCI, IDE for compatibility purpose when the guest OS does not have VirtIO drivers to support SCSI devices. This may have performance implications.

Note: For AHV 5.16 and later, you cannot use an IDE device if **Secured Boot** is enabled for **UEFI Mode** boot configuration.

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 */storage_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 (*/storage_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.

f. **Storage Container**: Select the storage container to use from the drop-down list.

This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.

g. **Size**: Enter the disk size in GiB.

h. **Index**: Displays **Next Available** by default.

i. 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.

j. Repeat this step to attach additional devices to the VM.

6. To create a network interface for the VM, click the **Add New NIC** button.
   The Create NIC dialog box appears.

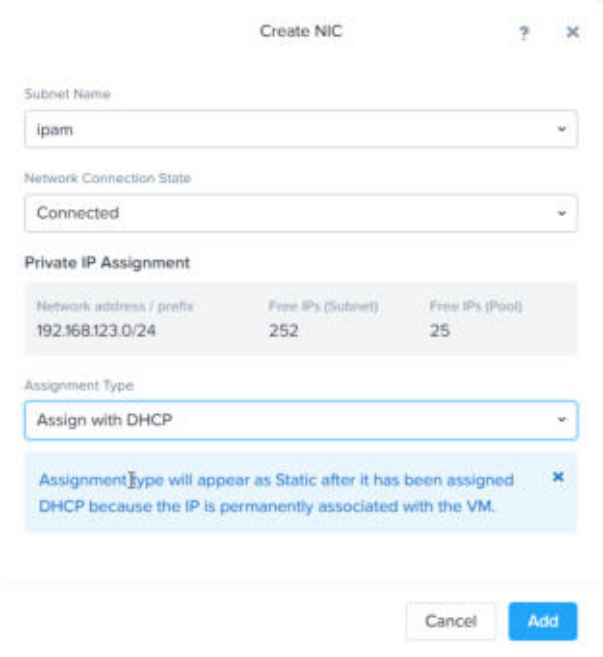

Figure 218: Create NIC Dialog Box

Do the following in the indicated fields:

a. **Subnet Name**: Select the target virtual LAN from the drop-down list.
   The list includes all defined networks (see the *Network Configuration For VM Interfaces* in the *Prism Web Console Guide*).

   > Note: Selecting IPAM enabled subnet from the drop-down list displays the **Private IP Assignment** information that provides information about the number of free IP addresses available in the subnet and in the IP pool.

b. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are *Connected* or *Disconnected*.

c. **Private IP Assignment**: This is a read-only field and displays the following:

   - **Network Address/Prefix**: The network IP address and prefix.
   - **Free IPs (Subnet)**: The number of free IP addresses in the subnet.
   - **Free IPs (Pool)**: The number of free IP addresses available in the IP pools for the subnet.

d. **Assignment Type**: This is for IPAM enabled network. Select **Assign with DHCP** to assign IP address automatically to the VM using DHCP. For more information, see *IP Address Management* in the *AHV Administration Guide*.

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.

Note: Nutanix guarantees a unique VM MAC address in a cluster. You can come across scenarios where two VM in different clusters can have the same MAC address.

Note: Acropolis leader generates MAC address for the VM on AHV. The first 24 bits of the MAC address is set to 50-6b-8d (0101 0000 0110 1101 1000 1101) and are reserved by Nutanix, the 25th bit is set to 1 (reserved by Acropolis leader), the 26th bit to 48th bits are auto generated random numbers.

7. To configure affinity policy for this VM, click **Set Affinity**.
   The Set VM Host Affinity dialog box appears.

   a. Select the host or hosts on which you want configure the affinity for this VM.

   b. Click **Save**.
      The selected host or hosts are listed. 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.

8. 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.

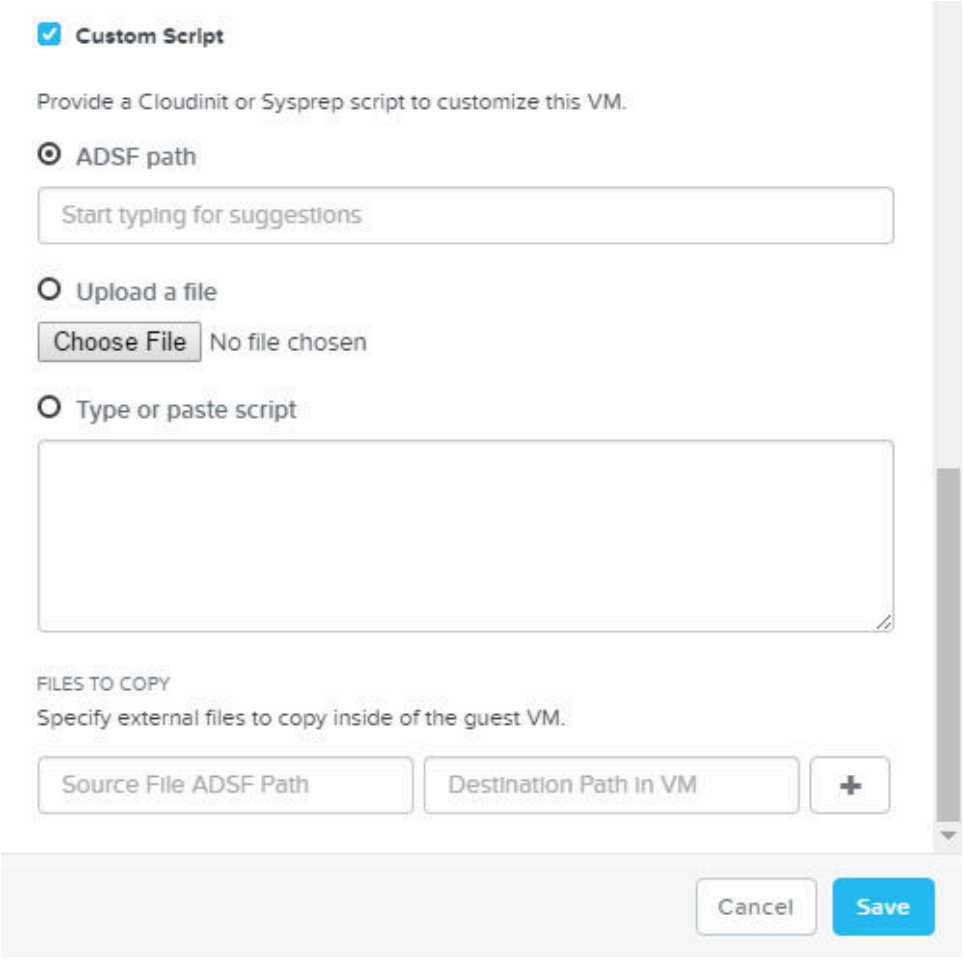

Figure 219: Create VM Dialog Box (custom script fields)

9. 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.

10. 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 absolute path to the file.

    b. In **Destination Path in VM**, enter the absolute path to the target directory and the file name.

    For example, if the source file entry is /home/my_dir/myfile.txt then the entry for the **Destination Path in VM** should be /<directory_name>/copy_desitation> i.e. /mnt/myfile.txt.

    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.

11. 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 VM table view.

Managing a VM (AHV)

You can use the web console to manage virtual machines (VMs) in Acropolis managed clusters.

About this task

After creating a VM (see Creating a VM (AHV) on page 316), you can use the web console to start or shut down the VM, launch a console window, update the VM configuration, take a snapshot, attach a volume group, migrate the VM, clone the VM, or delete the VM.

> Note: Your available options depend on the VM status, type, and permissions. Unavailable options are grayed out.

To accomplish one or more of these tasks, do the following:

Procedure

1. In the VM dashboard, click the **Table** view.

2. Select the target VM in the table (top section of screen).

   The Summary line (middle of screen) displays the VM name with a set of relevant action links on the right. You can also right-click on a VM to select a relevant action.

   The possible actions are **Manage Guest Tools**, **Launch Console**, **Power on** (or **Power off**), **Take Snapshot**, **Migrate**, **Clone**, **Update**, and **Delete**.

   > Note: VM pause and resume feature is not supported on AHV.

   The following steps describe how to perform each action.

Figure 220: VM Action Links

3. To manage guest tools as follows, click **Manage Guest Tools**.

   You can also enable NGT applications (self-service restore, Volume Snapshot Service and application-consistent snapshots) also as part of manage guest tools.

   a. Select **Enable Nutanix Guest Tools** check box to enable NGT on the selected VM.

   b. Select **Mount Nutanix Guest Tools** to mount NGT on the selected VM.

      Ensure that VM must have at least one empty IDE CD-ROM slot to attach the ISO.

      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.

   c. To enable self-service restore feature for Windows VMs, click **Self Service Restore (SSR)** check box.

      The Self-Service Restore feature is enabled of the VM. The guest VM administrator can restore the desired file or files from the VM. For more information on self-service restore

feature, see Self-Service Restore in the guide titled *Data Protection and Recovery with Prism Element*.

d. After you select **Enable Nutanix Guest Tools** check box the VSS snapshot feature is enabled by default.

   After this feature is enabled, Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent takes snapshots for VMs that support VSS.

   > Note:
   >
   > The AHV VM snapshots are not application consistent. The AHV snapshots are taken from the **VM** entity menu by selecting a VM and clicking **Take Snapshot**.
   >
   > The application consistent snapshots feature is available with Protection Domain based snapshots and Recovery Points in Prism Central. For more information, see *Implementation Guidelines for Asynchronous Disaster Recovery* section in the guide titled *Data Protection and Recovery with Prism Element*.

e. Click **Submit**.

   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. If the cloned VM is powered off, enable NGT from the UI and power on the VM. If cloned VM is powered on, enable NGT from the UI and restart the nutanix guest agent service.
   > - If you want to enable NGT on multiple VMs simultaneously, see *Enabling NGT and Mounting the NGT Installer Simultaneously on Multiple Cloned VMs* in 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.

   ```
   nutanix@cvm$ 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.

   ```
   nutanix@cvm$ ncli ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987
   ```

4. To launch a console window, click the **Launch Console** action link.

   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 console window includes four menu options (top right):

   - Clicking the **Mount ISO** button displays the following window that allows you to mount an ISO image to the VM. To mount an image, select the desired image and CD-ROM drive from the drop-down lists and then click the **Mount** button.

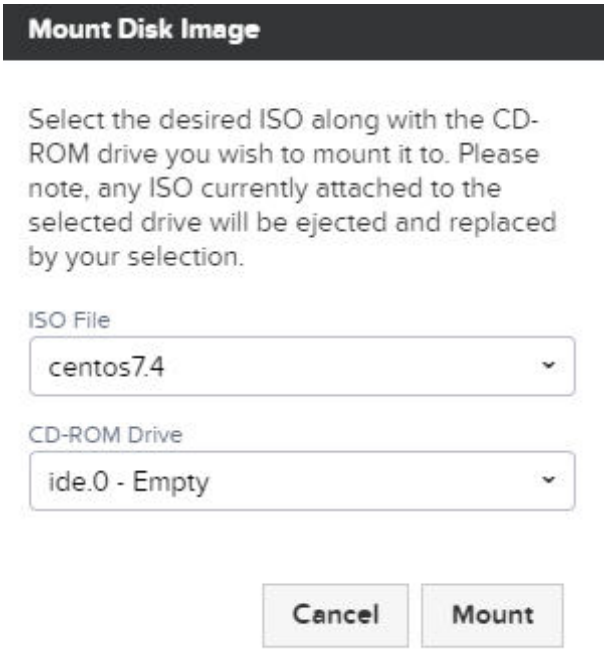

Figure 221: Mount Disk Image Window

> Note: See *Add New Disk* in Creating a VM (AHV) on page 316 section for information about how to select CD-ROM as the storage device when you intent to provide a system image from CD-ROM.

   - Clicking the **C-A-D** icon button sends a **CtrlAltDel** command to the VM.
   - Clicking the camera icon button takes a screenshot of the console window.
   - Clicking the power icon button allows you to power on/off the VM. These are the same options that you can access from the **Power On Actions** or **Power Off Actions** action link below the VM table (see next step).

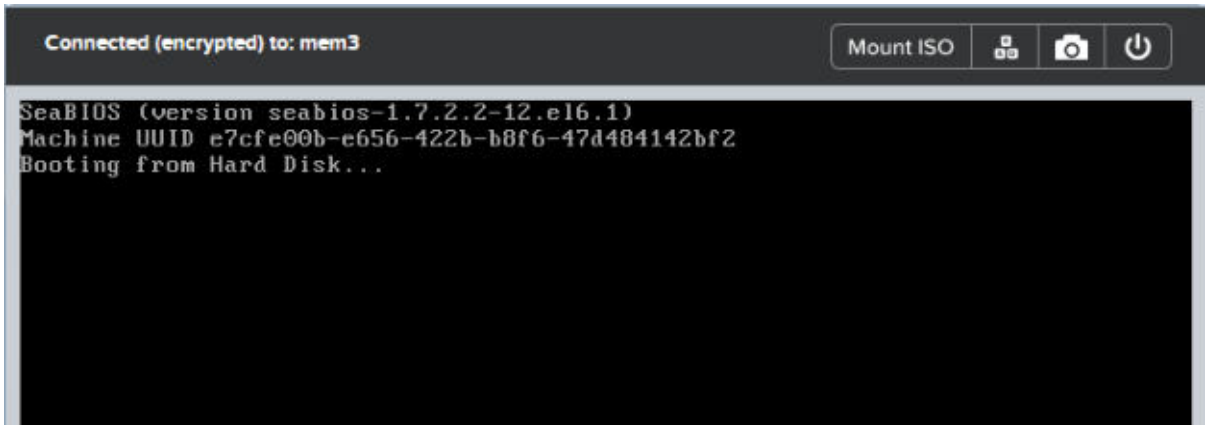

Figure 222: Virtual Network Computing (VNC) Window

5. To start or shut down the VM, click the **Power on** (or **Power off**) action link.

   Power on begins immediately. If you want to power off the VMs, you are prompted to select one of the following options:

   - **Power Off**. Hypervisor performs a hard power off action on the VM.
   - **Power Cycle**. Hypervisor performs a hard restart action on the VM.
   - **Reset**. Hypervisor performs an ACPI reset action through the BIOS on the VM.
   - **Guest Shutdown**. Operating system of the VM performs a graceful shutdown.
   - **Guest Reboot**. Operating system of the VM performs a graceful restart.

     Note: If you perform power operations such as Guest Reboot or Guest Shutdown by using the Prism Element web console or API on Windows VMs, these operations might silently fail without any error messages if at that time a screen saver is running in the Windows VM. Perform the same power operations again immediately, so that they succeed.

6. To make a snapshot of the VM, click the **Take Snapshot** action link.

   See Virtual Machine Snapshots on page 338 for more information.

7. To migrate the VM to another host, click the **Migrate** action link.

   This displays the Migrate VM dialog box. Select the target host from the drop-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.

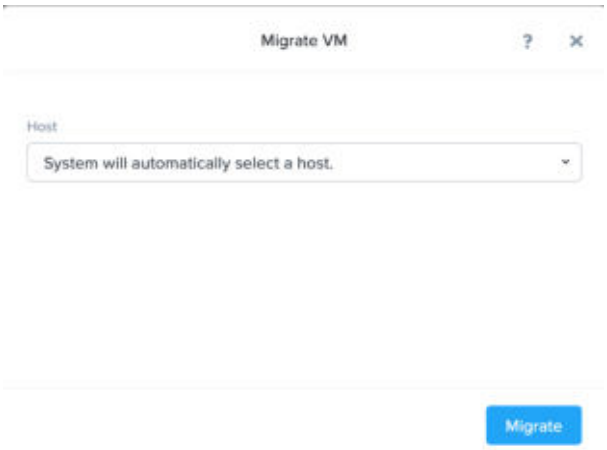

Figure 223: Migrate VM Dialog Box

> 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.

8. To clone the VM, click the **Clone** action link.

   This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits the most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

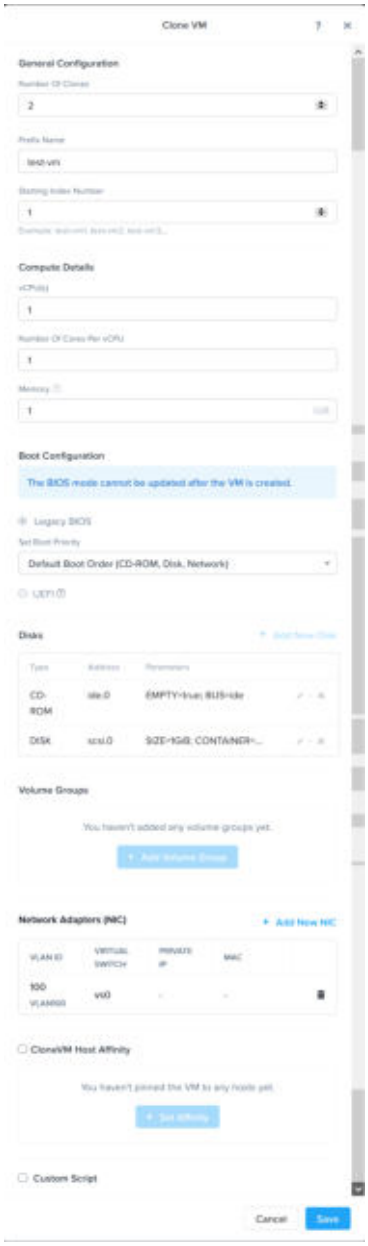

Figure 224: Clone VM Window

9. To modify the VM configuration, click the **Update** action link.

   The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Modify the configuration as needed, and then save the configuration. In addition

to modifying the configuration, you can attach a volume group to the VM and enable flash mode on the VM. If you attach a volume group to a VM that is part of a protection domain, the VM is not protected automatically. Add the VM to the same Consistency Group manually.

(For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the

specified vGPU license. However, you need to power off the VM before you perform these operations.

- Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.
- Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
- For more information on vGPU profile selection, see:
  - *Virtual GPU Types for Supported GPUs* in the *NVIDIA Virtual GPU Software User Guide* in the NVIDIA's *Virtual GPU Software Documentation* webpage, and
  - *GPU and vGPU Support* in the *AHV Administration Guide*.
- After you add the first vGPU, to add multiple vGPUs, see *Adding Multiple vGPUs to the Same VM* in the *AHV Admin Guide*.

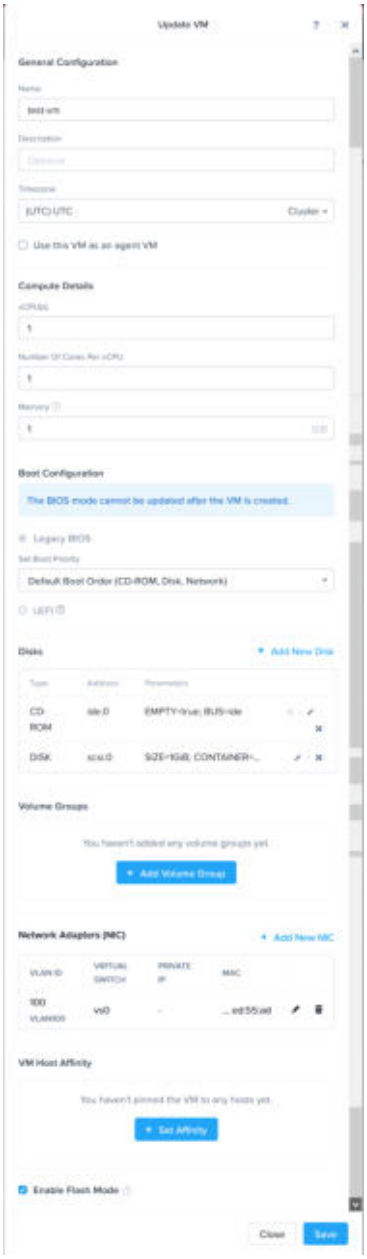

Figure 225: VM Update Dialog Box

> Note: If you delete a vDisk attached to a VM and snapshots associated with this VM exist, space associated with that vDisk is not reclaimed unless you also delete the VM snapshots.

To increase the memory allocation and the number of vCPUs on your VMs while the VMs are powered on (hot-pluggable), do the following:

a. In the **vCPUs** field, you can increase the number of vCPUs on your VMs while the VMs are powered on.

b. In the **Number of Cores Per vCPU** field, you can change the number of cores per vCPU only if the VMs are powered off.

> Note: This is not a hot-pluggable feature.

c. In the **Memory** field, you can increase the memory allocation on your VMs while the VMs are powered on.

For more information about hot-pluggable vCPUs and memory, see the *Virtual Machine Memory and CPU Hot-Plug Configurations* topic in the *AHV Administration Guide*.

To attach a volume group to the VM, do the following:

a. In the **Volume Groups** section, click **Add volume group**, and then do one of the following:

   » From the **Available Volume Groups** list, select the volume group that you want to attach to the VM.

   » Click **Create new volume group**, and then, in the **Create Volume Group** dialog box, create a volume group (see *Creating a Volume Group* in the *Prism Web Console Guide*). After you create a volume group, select it from the **Available Volume Groups** list.

   Repeat these steps until you have added all the volume groups that you want to attach to the VM.

b. Click **Add**.

10. To enable flash mode on the VM, click the **Enable Flash Mode** check box.

    » After you enable this feature on the VM, the status is updated in the VM table view. To view the status of individual virtual disks (disks that are flashed to the SSD), click the update disk icon in the **Disks** pane in the Update VM window.

    » You can disable the flash mode feature for individual virtual disks. To update the flash mode for individual virtual disks, click the update disk icon in the **Disks** pane and deselect the **Enable Flash Mode** check box.

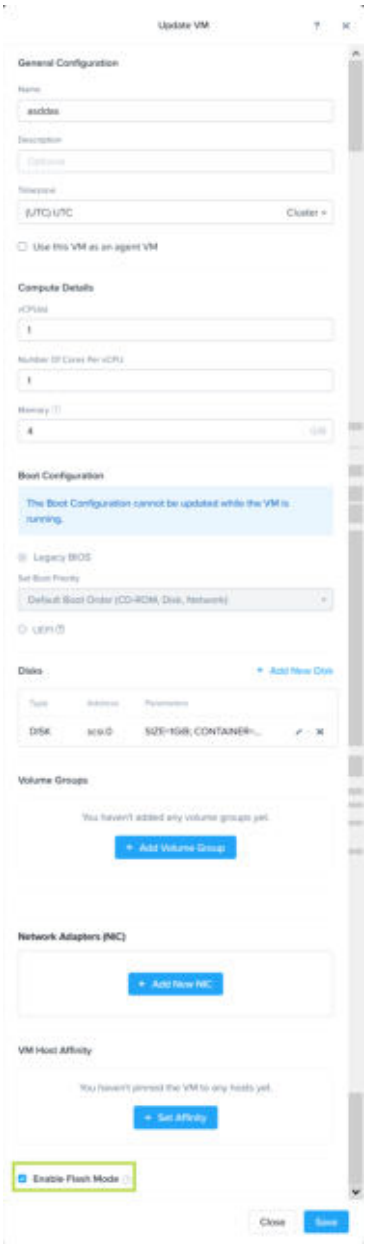

Figure 226: Update VM Resources

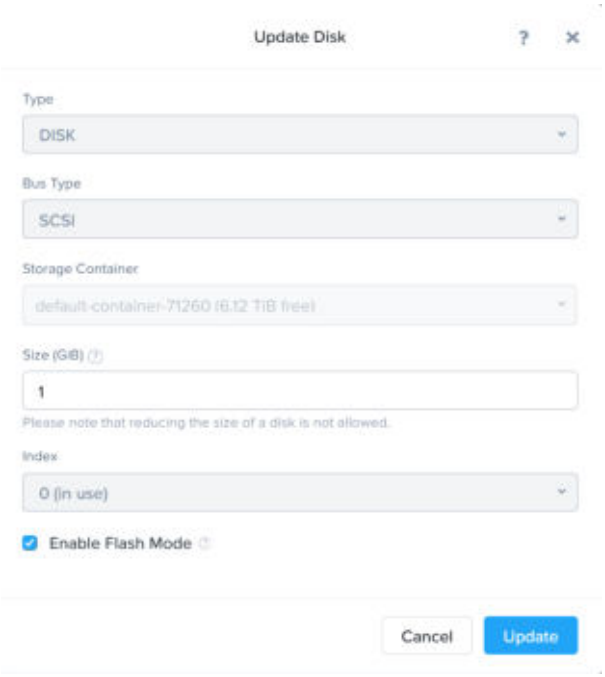

Figure 227: Update VM Resources - VM Disk Flash Mode

11. To delete the VM, click the **Delete** action link. A window prompt appears; click the **OK** button to delete the VM.
    The deleted VM disappears from the list of VMs in the table.

Virtual Machine Snapshots

You can generate snapshots of virtual machines or VMs. You can generate snapshots of VMs manually or automatically. Some of the purposes that VM snapshots serve are as follows:

- Disaster recovery
- Testing - as a safe restoration point in case something went wrong during testing.
- Migrate VMs
- Create multiple instances of a VM.

A snapshot is a copy of a state of the system - files and data - at a specific point in time. You can generate snapshots and store them locally or remotely. Snapshots are mechanism to capture the delta changes that has occurred over time. Snapshots are primarily used for data protection and disaster recovery. Snapshots are not autonomous like backup, in the sense that they depend on the underlying VM infrastructure and other snapshots to restore the VM. Snapshots consume less resources compared to a full autonomous backup. Typically, a VM snapshot captures the following:

- The state including the power state (for example, powered-on, powered-off, suspended) of the VMs.
- The data includes all the files that make up the VM. This data also includes the data from disks, memory, and other devices, such as virtual network interface cards.

VM Snapshots and Snapshots for Disaster Recovery

The VM Dashboard only allows you to generate VM snapshots manually. You cannot select VMs and schedule snapshots of the VMs using the VM dashboard. The snapshots generated manually have very limited utility.

> Note: These snapshots (stored locally) cannot be replicated to other sites.

You can schedule and generate snapshots as a part of the disaster recovery process using Nutanix DR solutions. AOS generates snapshots when you protect a VM with a protection domain using the Data Protection dashboard in Prism Web Console (see the *Data Protection and Recovery with Prism Element* guide). Similarly, Recovery Points (snapshots are called Recovery Points in Prism Central) when you protect a VM with a protection policy using Data Protection dashboard in Prism Central (see the *Leap Administration Guide*).

For example, in the Data Protection dashboard in Prism Web Console, you can create schedules to generate snapshots using various RPO schemes such as asynchronous replication with frequency intervals of 60 minutes or more, or NearSync replication with frequency intervals of as less as 20 seconds up to 15 minutes. These schemes create snapshots in addition to the ones generated by the schedules, for example, asynchronous replication schedules generate snapshots according to the configured schedule and, in addition, an extra snapshot every 6 hours. Similarly, NearSync generates snapshots according to the configured schedule and also generates one extra snapshot every hour.

Similarly, you can use the options in the Data Protection section of Prism Central to generate Recovery Points using the same RPO schemes.

Creating a VM Snapshot Manually

You can create or generate a VM snapshot manually in Prism Web Console.

About this task

Do the following:

Procedure

1. In the VM Table view, click the VM you want to take a snapshot of.

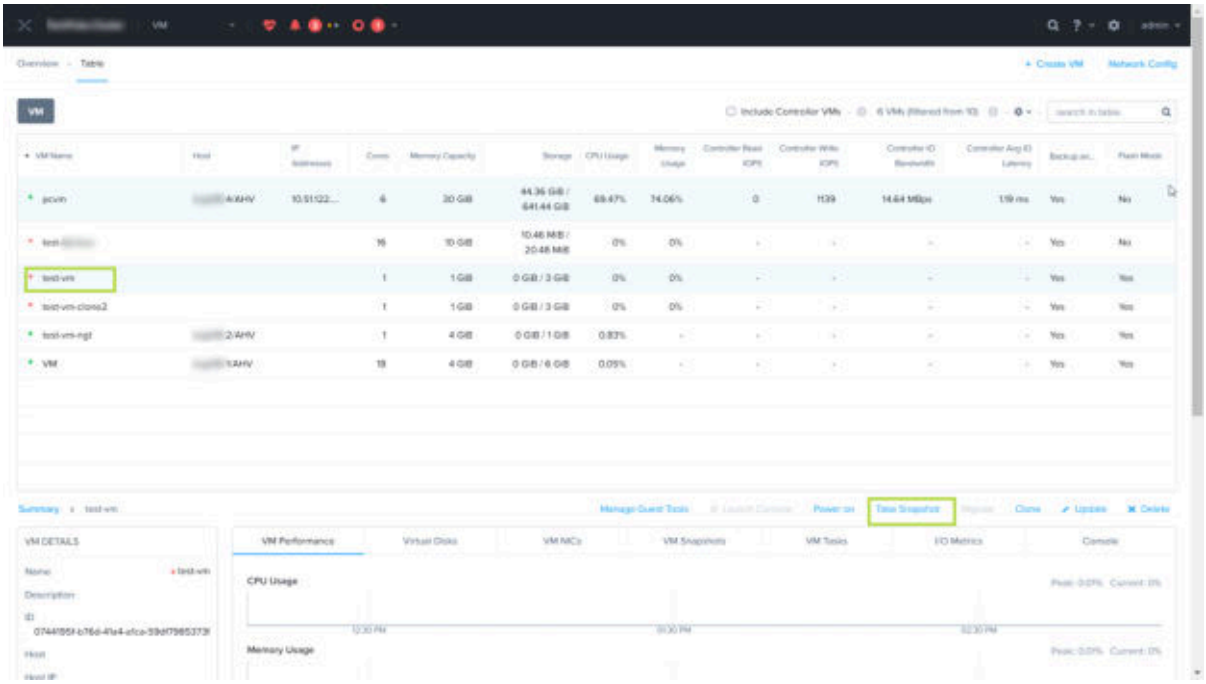

Figure 228: VM Table View

2. In the Take Snapshot dialog box, provide a name for the snapshot and click **Submit**. Ensure that the name contains only alphanumeric, dot, dash, or underscore characters.

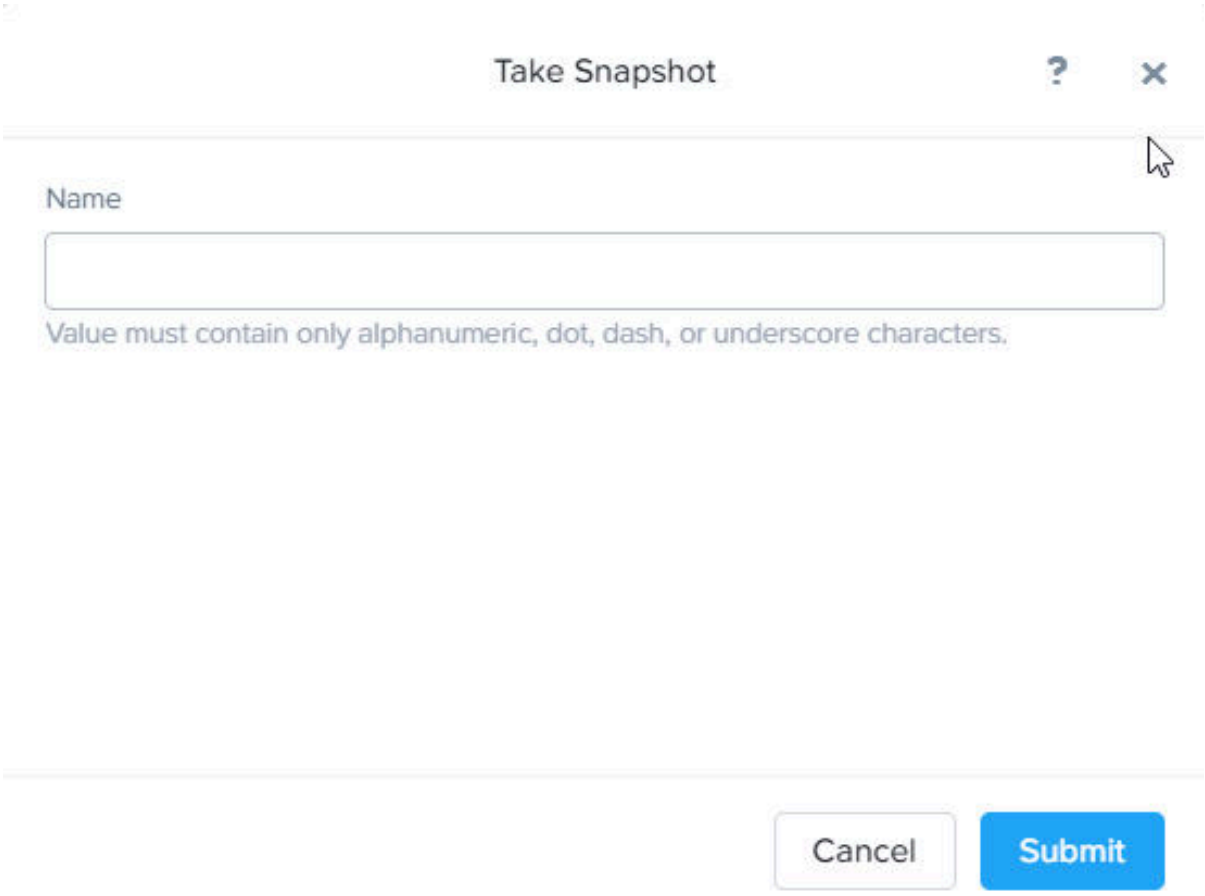

Figure 229: Take Snapshot

The snapshot is listed with the name you provided in the **VM Snapshots** tab in the **Summary** section of the VM Table view.

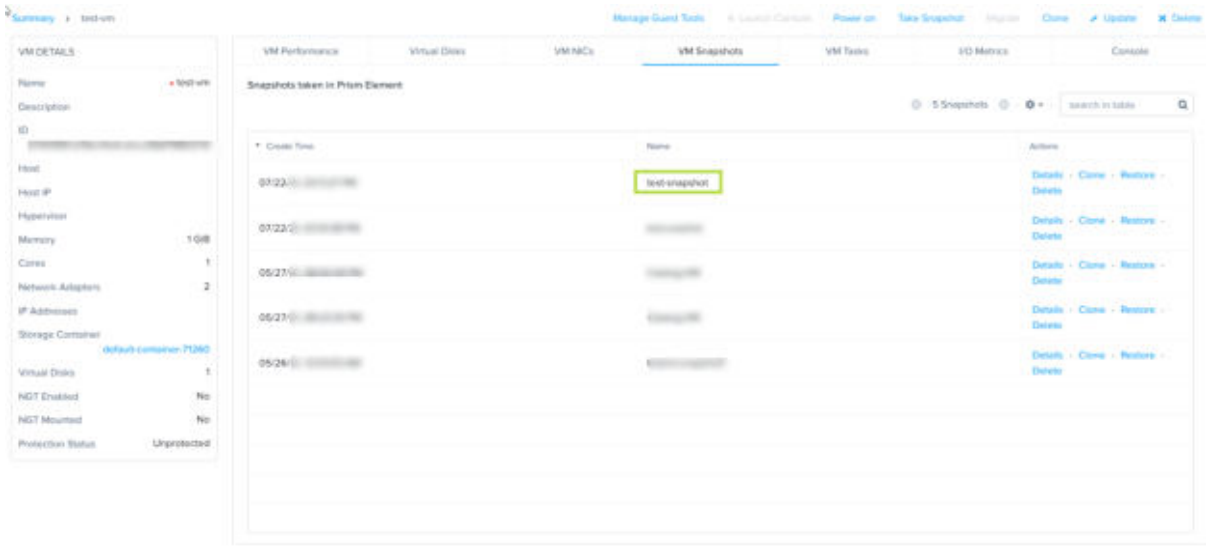

Figure 230: VM Snapshots

What to do next

You can **Delete** or **Clone** the VM snapshot. You can use the VM snapshot to **Restore** the VM to the previous state captured in the snapshot. Click **Details** to view the details of the snapshot.

Adding Multiple vGPUs to the Same VM

About this task

You can add multiple vGPUs of the same vGPU type to:

- A new VM when you create it.
- An existing VM when you update it.

> Important:
>
> Before you add multiple vGPUs to the VM, see *Multiple Virtual GPU Support* and *Restrictions for Multiple vGPU Support* in the *AHV Admin Guide*.

After you add the first vGPU, do the following on the Create VM or Update VM dialog box (the main dialog box) to add more vGPUs:

Procedure

1. Click **Add GPU**.

2. In the Add GPU dialog box, click **Add**.

   The License field is grayed out because you cannot select a different license when you add a vGPU for the same VM.

   The **VGPU Profile** is also auto-selected because you can only select the additional vGPU of the same vGPU type as indicated by the message at the top of the dialog box.

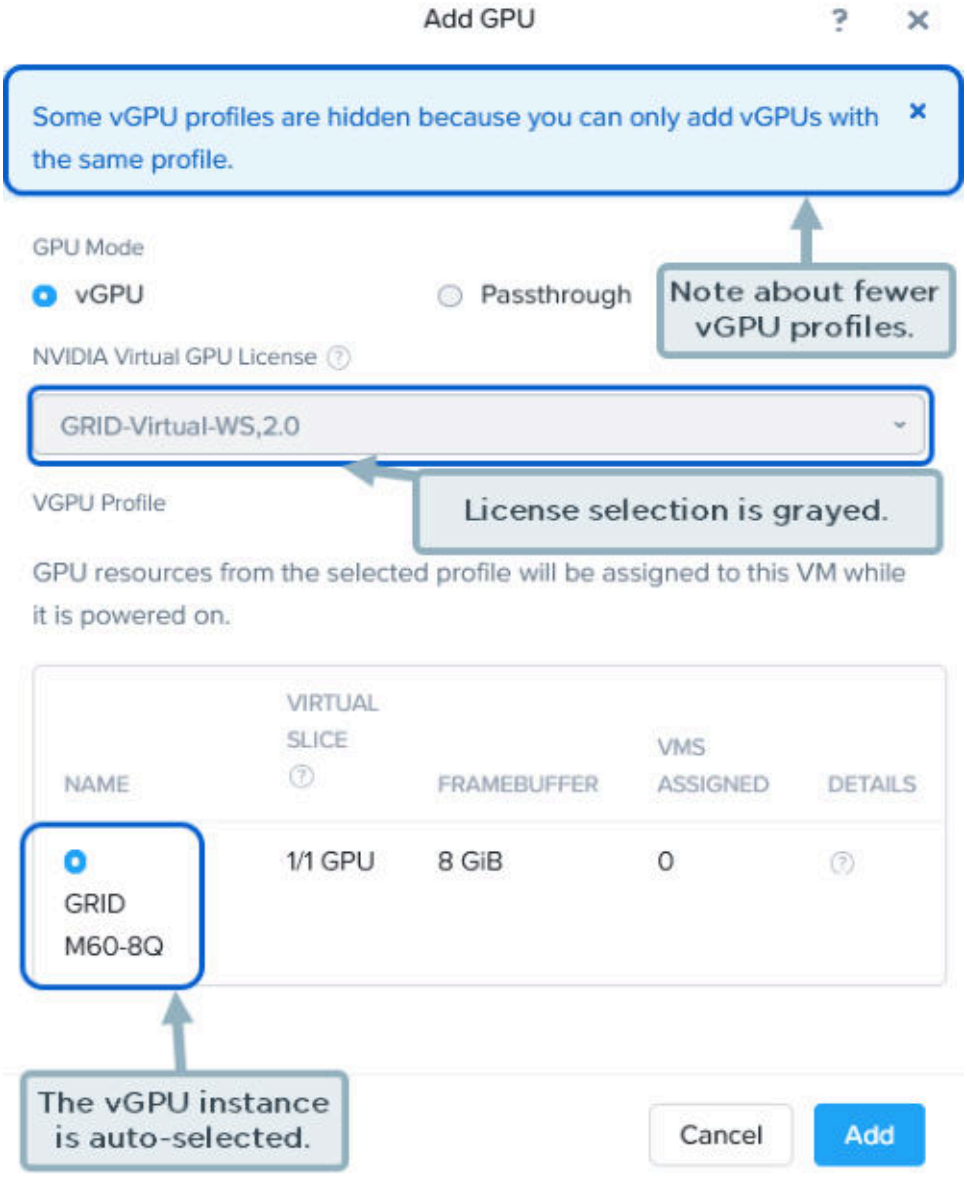

Figure 231: Add GPU for multiple vGPUs

3. In the main dialog box, you see the newly added vGPU.

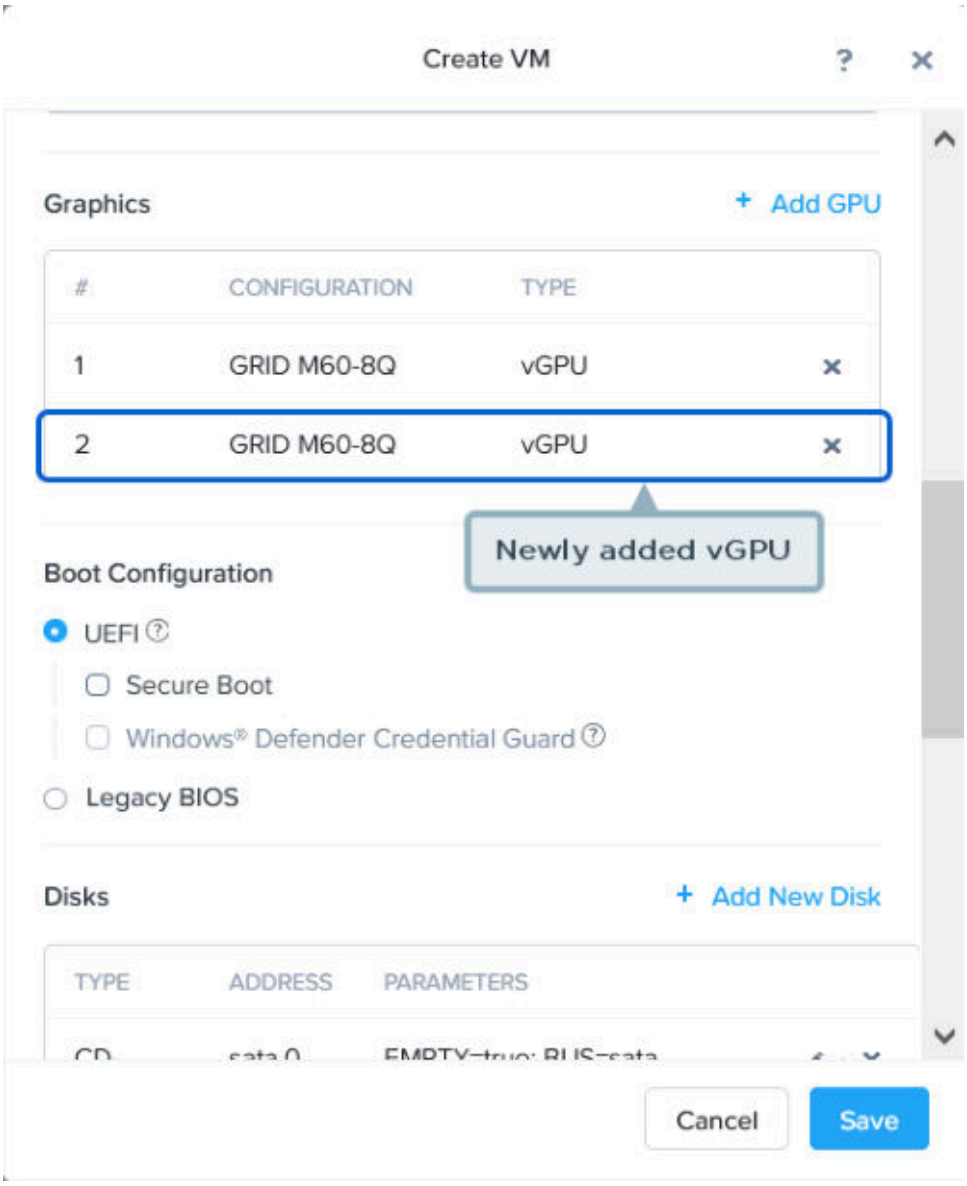

Figure 232: New vGPUs Added

4. Repeat the steps for each vGPU addition you want to make.

Migrating Live a vGPU-enabled VM Within the Cluster

You can migrate a vGPU-enabled VM only to another host in the same cluster on Prism Element. To migrate a vGPU-enabled VM to a host outside the cluster, use Prism Central. See *Migrating Live a VM with Virtual GPUs* in the *Prism Central Guide*.

Before you begin

Ensure the following:

1. The VM is not powered off.

2. The host affinity is not set for the VM.

   If the host affinity of the VM is set to only one host, you cannot live migrate the VM. However, based on the GPU resources required, if the host affinity is set to multiple hosts with the same or similar GPU resources, you can migrate the VM among the hosts with which the affinity is set.
3. You have another host in the same cluster to migrate the VM.

For limitations applicable to live migration of vGPU-enabled VMs, see *Limitations for Live Migration* in the *NVIDIA GRID Virtual GPU Support on AHV* section of the *AHV Administration Guide*.

> Note: Live migration of VMs with vGPUs is supported:
>
> - For vGPUs created with minimum NVIDIA Virtual GPU software version 10.1 (440.53).
> - Minimum AOS version that supports live migration within the same cluster is 5.18.1.
> - Minimum AOS version that supports live migration across cluster is 6.1.

About this task

To migrate the vGPU-enabled VM to another host within the same cluster on Prism Element, do the following.

Procedure

1. To migrate the VM, go to the **VM** > **Table**.
2. Select the VM you want to migrate live. Click **Migrate**.

Figure 233: VM Selection

3. In the Migrate VM dialog box, do the following:

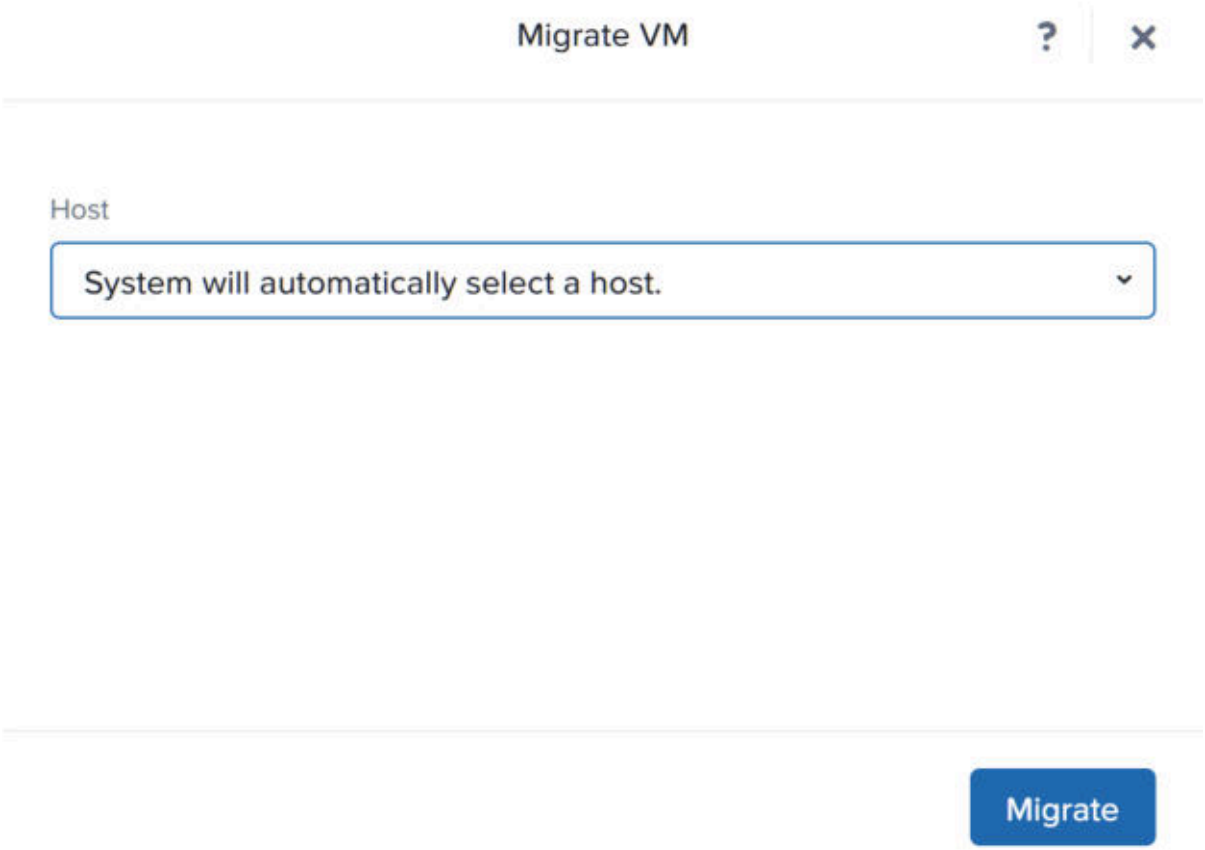

Figure 234: Migrate VM dialog box

a. In the **Host** drop-down list, do one of the following:

   » Retain the *System will automatically select a host* option if you want to migrate the VM to a host selected by the system.

   The system selects a host based on the GPU resources available with the host as appropriate for the VM to be migrated live.

   » Select the host listed in the drop-down list that you want to migrate the VM to.

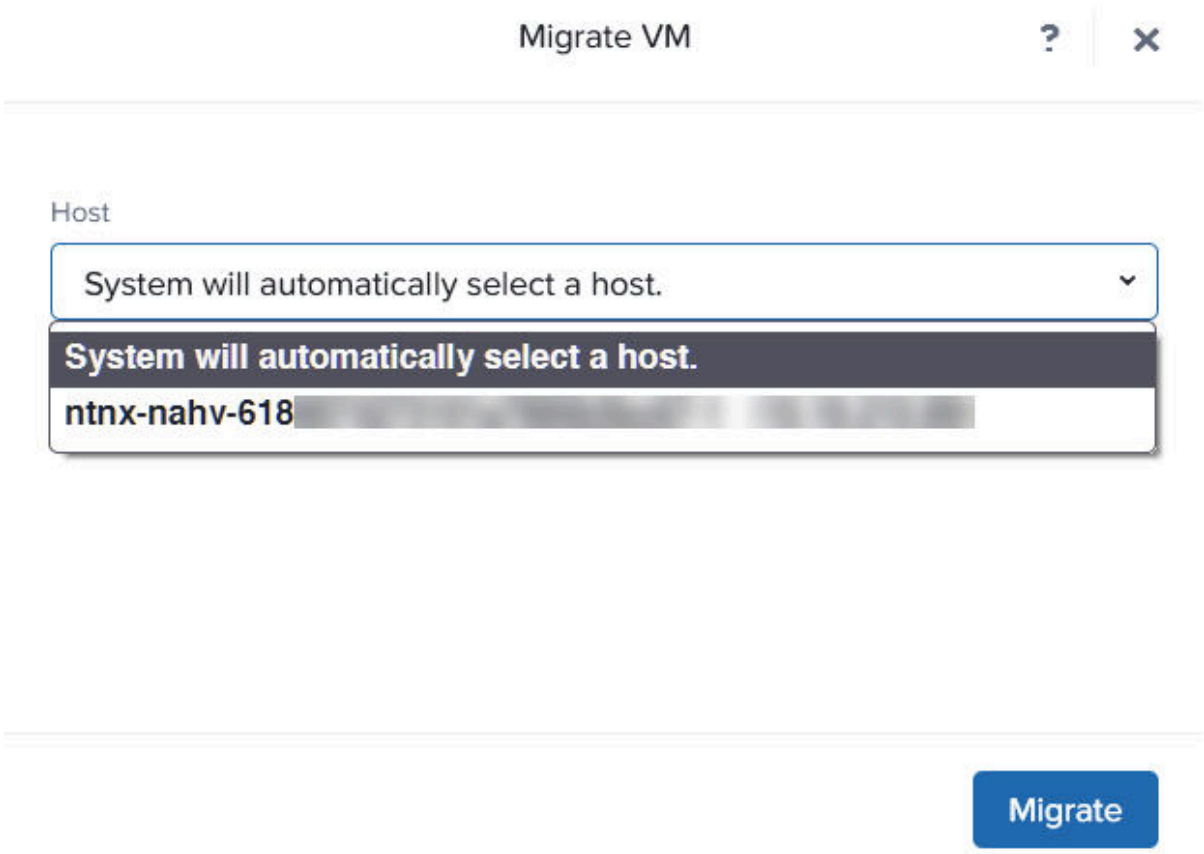

Figure 235: Host Selection

4. Click **Migrate**.

   Prism submits the task and displays the following message:

   ```
   Successfully submitted migrate operation.
   Task details
   ```

   Task details is a link to the Tasks page. Click the link to monitor the migration task on the Tasks page.

   When the migration is complete, the host name of the VM in the List view changes to the host name to which you migrated the VM.

Creating a VM (ESXi)

In ESXi clusters, you can create a new virtual machine (VM) through the web console.

Before you begin

- See the requirements and limitations section in the *VM Management through Prism Element (ESXi)* topic in the *Prism Web Console Guide* before proceeding.
- Register the vCenter Server with your cluster. For more information, see Registering a Cluster to vCenter Server on page 771.

About this task

When creating a VM, you can configure all of its components, such as number of vCPUs and memory, but you cannot attach a volume group to the VM.

To create a VM, do the following:

Procedure

1. In the VM dashboard , click the **Create VM** button.
   The Create VM dialog box appears.
2. 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 *VM Management through Prism Element (ESXi)* topic in the *Prism Web Console Guide*.
   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.

3. To attach a disk to the VM, click the **Add New Disk** button.
   The Add Disks dialog box appears.

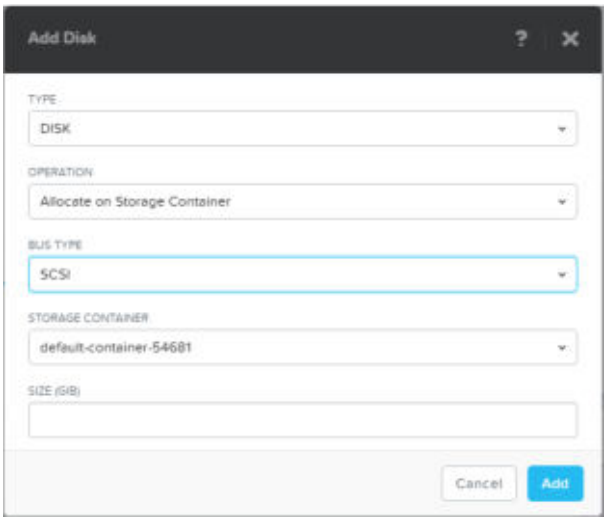

Figure 236: Add Disk Dialog Box

Do the following in the indicated fields:

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**.

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 */storage_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 (*/storage_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 are 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 more devices to the VM.

4. 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 *Network Configuration for VM Interfaces* topic in the *Prism Web Console Guide*).

   b. **Network Adapter Type**: Select the network adapter type from the pull-down list.
      See *VM Management through Prism Element (ESXi)* topic in the *Prism Web Console Guide* 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 more network interfaces for the VM.

5. 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 VM table view (see *VM Table View* topic in the *Prism Web Console Guide*).

Managing a VM (ESXi)

You can use the web console to manage virtual machines (VMs) in the ESXi clusters.

Before you begin

- See the requirements and limitations section in the *VM Management through Prism Element (ESXi)* topic in the *Prism Web Console Guide* before proceeding.
- Ensure that you have registered the vCenter Server with your cluster. For more information, see Registering a Cluster to vCenter Server on page 771.

About this task

After creating a VM, you can use the web console to manage guest tools, power operations, suspend, launch a VM console window, update the VM configuration, clone the VM, or delete the VM. To accomplish one or more of these tasks, do the following:

Note: Your available options depend on the VM status, type, and permissions. Unavailable options are unavailable.

Procedure

1. In the VM dashboard , click the **Table** view.
2. Select the target VM in the table (top section of screen).
   The summary line (middle of screen) displays the VM name with a set of relevant action links on the right. You can also right-click on a VM to select a relevant action.

   The possible actions are **Manage Guest Tools**, **Launch Console**, **Power on** (or **Power off actions**), **Suspend** (or **Resume**), **Clone**, **Update**, and **Delete**. The following steps describe how to perform each action.

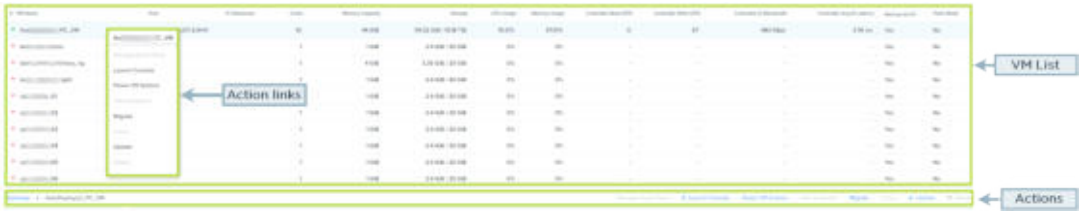

Figure 237: VM Action Links

3. To manage guest tools as follows, click **Manage Guest Tools**.
   You can also enable NGT applications (self-service restore, volume snapshot service and application-consistent snapshots) as part of manage guest tools.

   a. Select **Enable Nutanix Guest Tools** check box to enable NGT on the selected VM.

   b. Select **Mount Nutanix Guest Tools** to mount NGT on the selected VM.
      Ensure that VM has at least one empty IDE CD-ROM or SATA slot to attach the ISO.

      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.

   c. To enable self-service restore feature for Windows VMs, click **Self Service Restore (SSR)** check box.
      The self-service restore feature is enabled of the VM. The guest VM administrator can restore the desired file or files from the VM. For more information on self-service restore feature, see self-service restore in the guide titled *Data Protection and Recovery with Prism Element*.

   d. After you select **Enable Nutanix Guest Tools** check box the VSS and application-consistent snapshot feature is enabled by default.
      After this feature is enabled, Nutanix native in-guest VmQuiesced snapshot service (VSS) agent is used to take application-consistent snapshots for all the VMs that support VSS. This mechanism takes application-consistent snapshots without any VM stuns (temporary unresponsive VMs) and also enables third-party backup providers like Commvault and Rubrik to take application-consistent snapshots on Nutanix platform in a hypervisor-

agnostic manner. For more information, see Implementation Guidelines for Asynchronous Disaster Recovery in the guide titled *Data Protection and Recovery with Prism Element*.

e. To mount VMware guest tools, click **Mount VMware Guest Tools** check box.
The VMware guest tools are mounted on the VM.

> Note: You can mount both VMware guest tools and Nutanix Guest Tools at the same time on a particular VM provided the VM has sufficient empty CD-ROM slots.

f. Click **Submit**.
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. If the cloned VM is powered off, enable NGT from the UI and start the VM. If cloned VM is powered on, enable NGT from the UI and restart the Nutanix guest agent service.
> - If you want to enable NGT on multiple VMs simultaneously, see *Enabling NGT and Mounting the NGT Installer on Cloned VMs* topic in 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
```

> CAUTION: In AOS 4.6, for the powered-on Linux VMs on AHV, ensure that the NGT ISO is ejected or unmounted within the guest VM before disabling NGT by using the web console. This issue is specific for 4.6 version and does not occur from AOS 4.6.x or later releases.

> Note: If you have created the NGT ISO CD-ROMs prior to AOS 4.6 or later releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You must unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to 4.6 or later version.

4. To launch a VM console window, click the **Launch Console** action link.
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 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 Google Chrome. (Firefox typically works best.)

5. To start (or shut down) the VM, click the **Power on** (or **Power off**) action link.

   Power on begins immediately. If you want to shut down the VMs, you are prompted to select one of the following options:

   - **Power Off**. Hypervisor performs a hard shut down action on the VM.
   - **Reset**. Hypervisor performs an ACPI reset action through the BIOS on the VM.
   - **Guest Shutdown**. Operating system of the VM performs a graceful shutdown.
   - **Guest Reboot**. Operating system of the VM performs a graceful restart.

     Note: The **Guest Shutdown** and **Guest Reboot** options are available only when VMware guest tools are installed.

6. To pause (or resume) the VM, click the **Suspend** (or **Resume**) action link. This option is available only when the VM is powered on.

7. To clone the VM, click the **Clone** action link.

   This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box. A cloned VM inherits the most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For

example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- In the Clone window, you cannot update the disks.

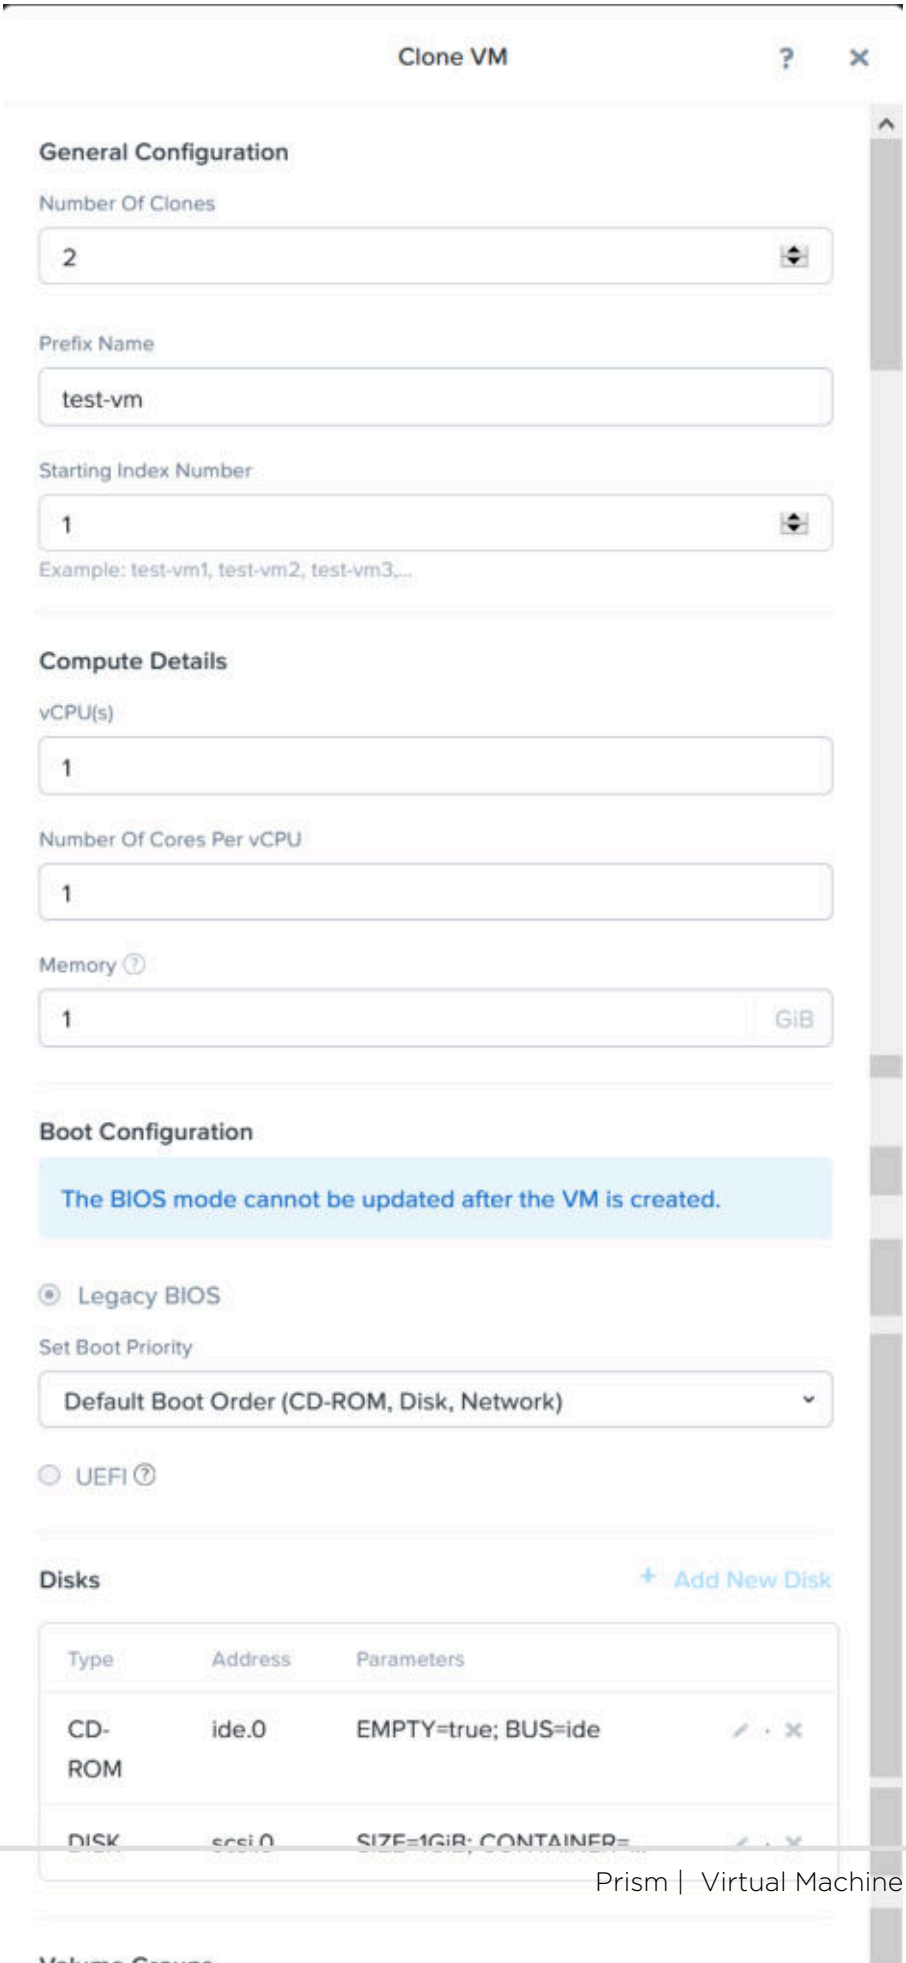
Clone VM
General Configuration
Number Of Clones
2
Prefix Name
test-vm
Starting Index Number
1
Example: test-vm1, test-vm2, test-vm3,...
Compute Details
vCPU(s)
1
Number Of Cores Per vCPU
1
Memory
1
GiB
Boot Configuration
The BIOS mode cannot be updated after the VM is created.
Legacy BIOS
Set Boot Priority
Default Boot Order (CD-ROM, Disk, Network)
UEFI
Disks
Add New Disk
Type
Address
Parameters
CD-ROM
ide.0
EMPTY=true; BUS=ide
DISK
scsi.0
SIZE=1GiB; CONTAINER=

8. To modify the VM configuration, click the **Update** action link.

   The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Modify the configuration as needed (see Creating a VM (ESXi) on page 347), and in addition you can enable Flash Mode for the VM.

   > Note: If you delete a vDisk attached to a VM and snapshots associated with this VM exist, space associated with that vDisk is not reclaimed unless you also delete the VM snapshots.

   a. Click the **Enable Flash Mode** check box.

      » After you enable this feature on the VM, the status is updated in the VM table view. To view the status of individual virtual disks (disks that are flashed to the SSD), go the **Virtual Disks** tab in the VM table view.

      » You can disable the Flash Mode feature for individual virtual disks. To update the Flash Mode for individual virtual disks, click the update disk icon in the **Disks** pane and deselect the **Enable Flash Mode** check box.

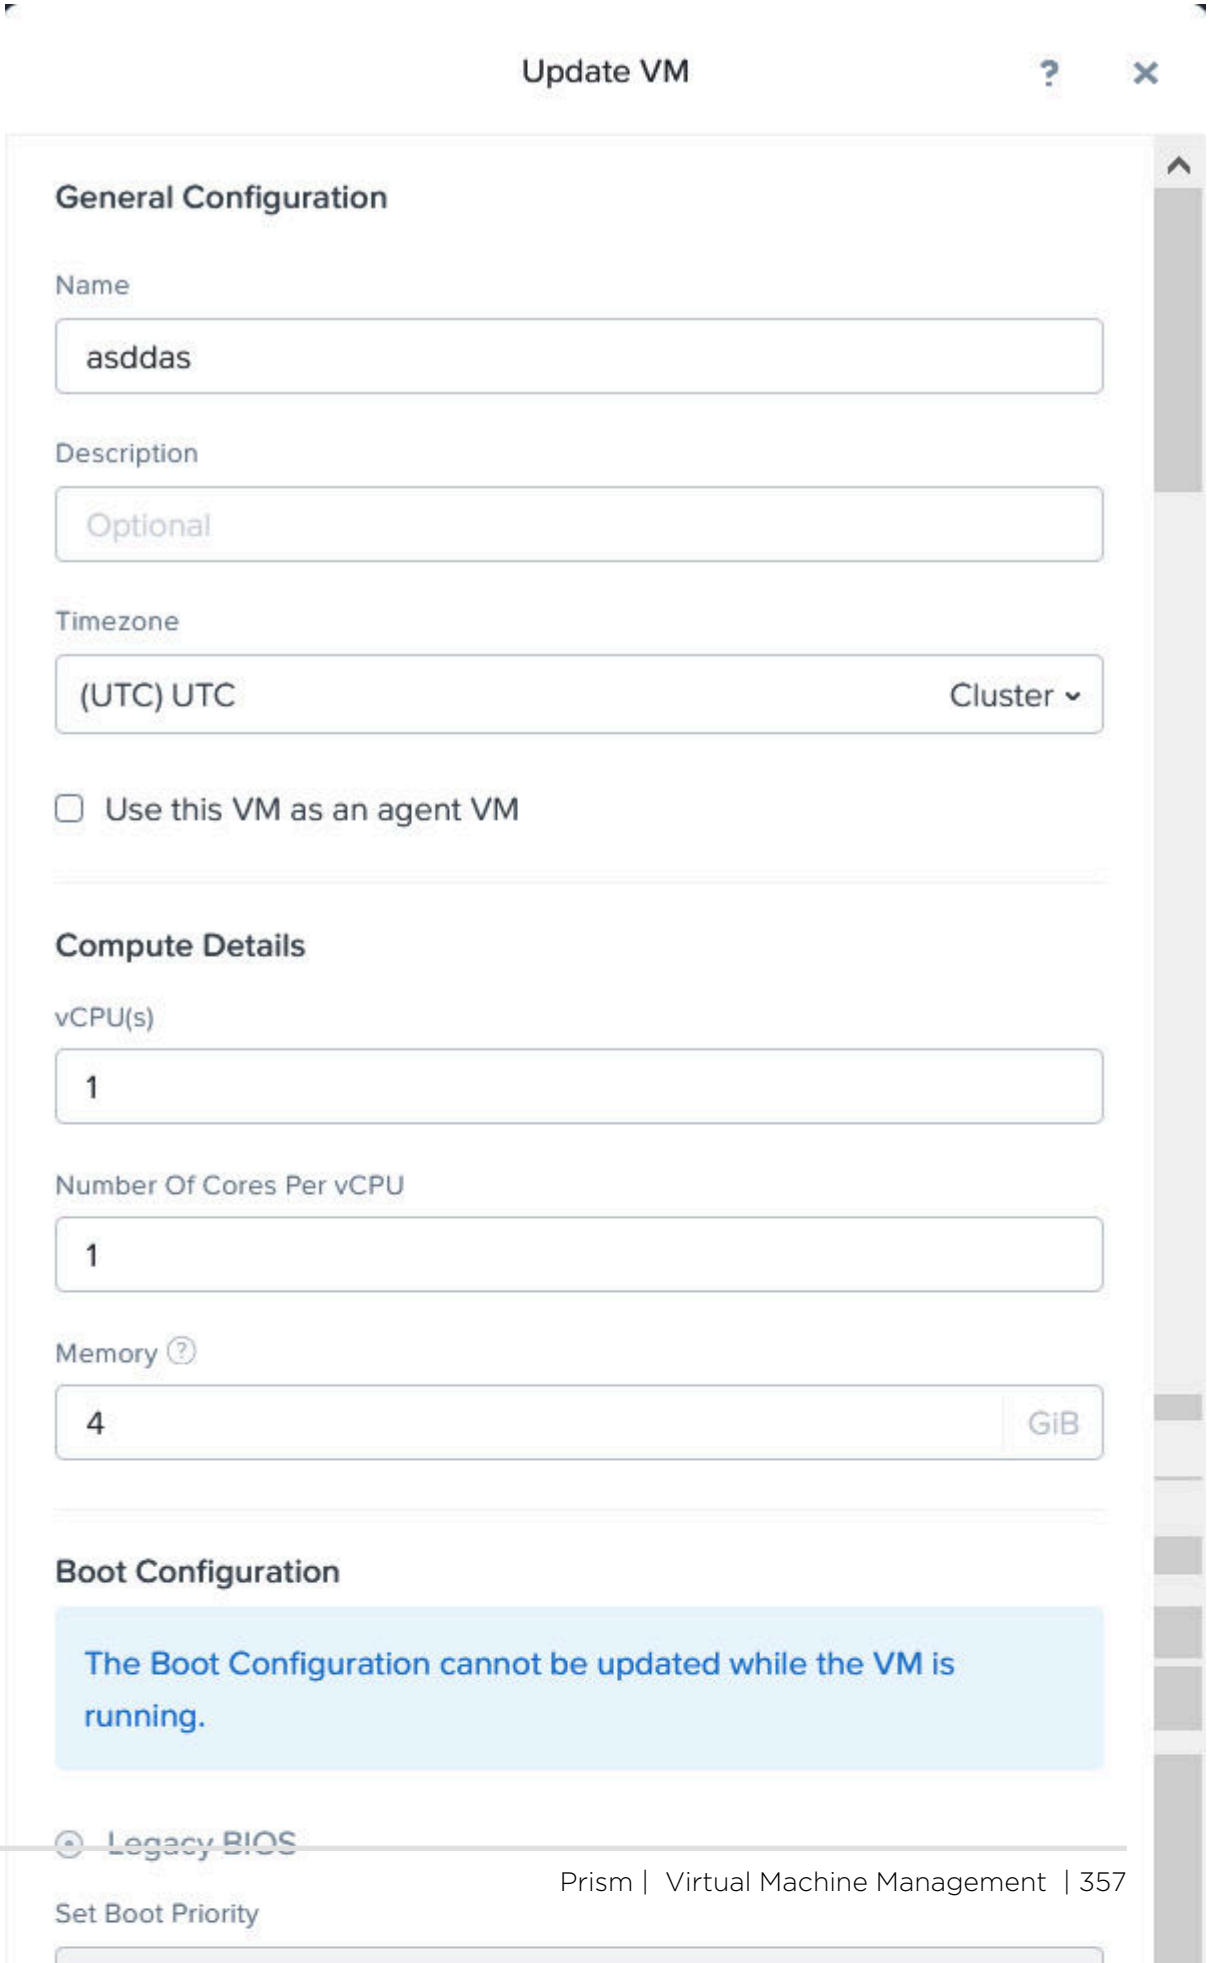
Update VM
General Configuration
Name
asddas
Description
Optional
Timezone
(UTC) UTC
Cluster
Use this VM as an agent VM
Compute Details
vCPU(s)
1
Number Of Cores Per vCPU
1
Memory
4
GiB
Boot Configuration
The Boot Configuration cannot be updated while the VM is running.
Legacy BIOS
Set Boot Priority

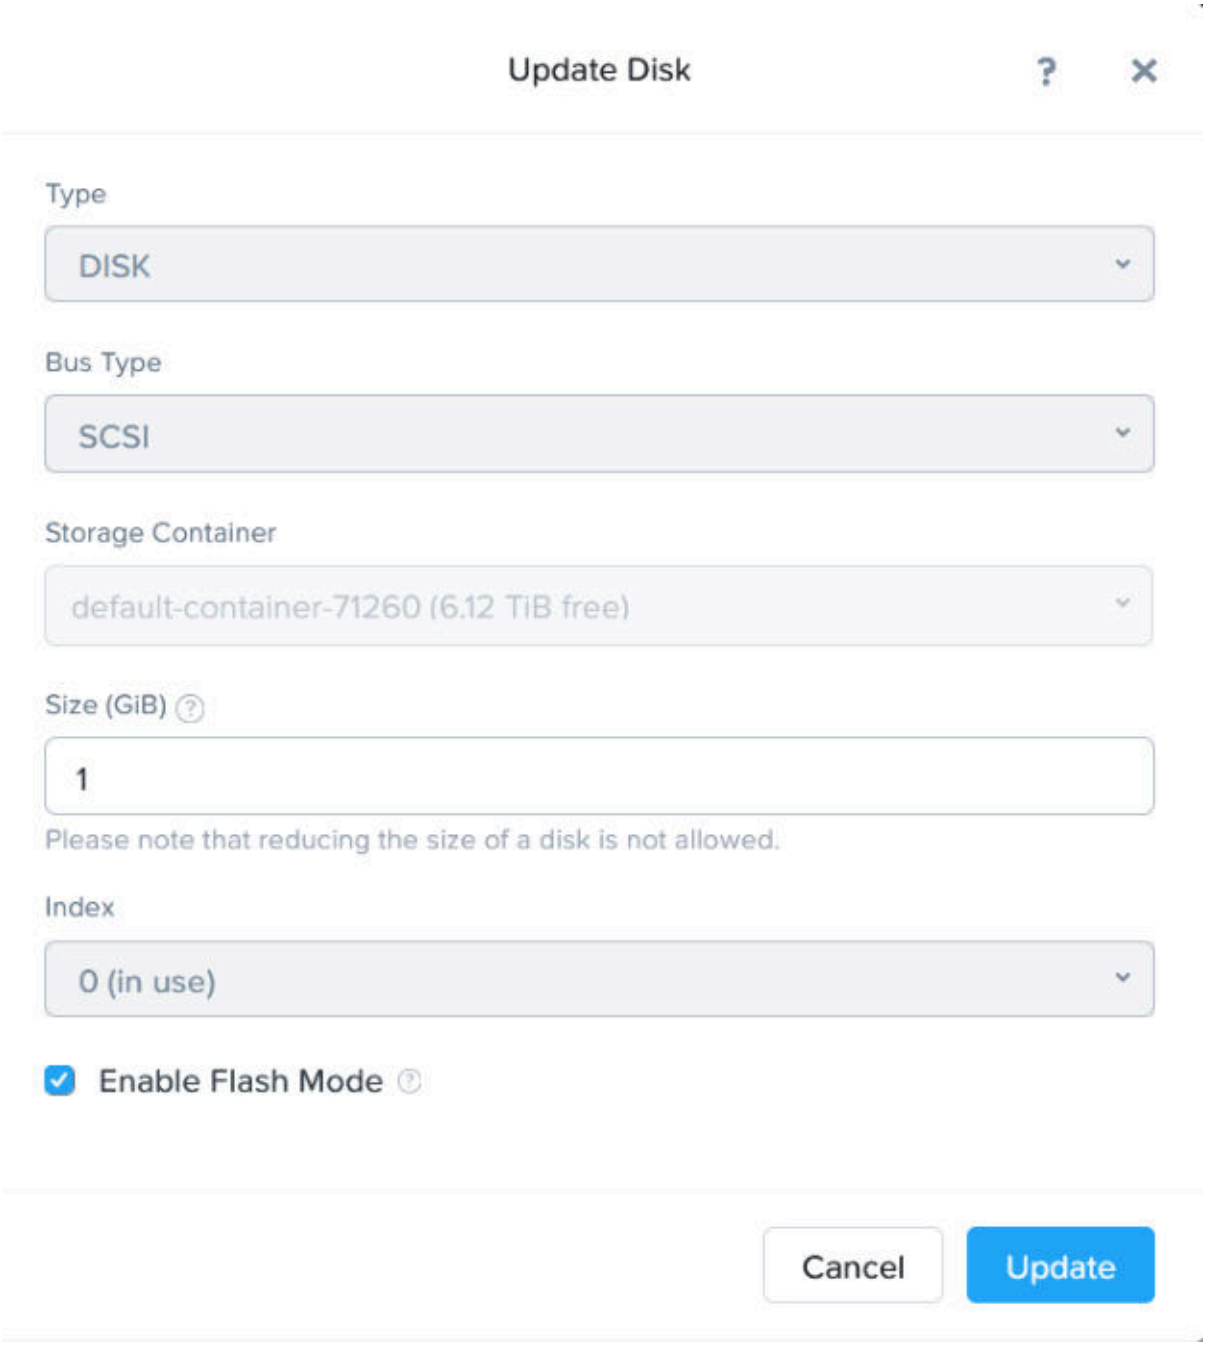

Figure 240: Update VM Resources - VM Disk Flash Mode

9. To delete the VM, click the **Delete** action link. A window prompt appears; click the **OK** button to delete the VM.
   The deleted VM disappears from the list of VMs in the table. You can also delete a VM that is already powered on.

Configuring Images

In AHV clusters, you can import and configure operating system ISO and disk image files through the web console. You can also convert previously imported files to the format that AHV uses.

About this task

The image service feature allows you to build a store of imported files that you can use to create a CD-ROM from an ISO image or an operating system. Disk from a disk image when creating a VM. The image service supports raw, vhd, vhdx, vmdk, vdi, ova, iso, and qcow2 disk formats. (The image service uses port 2007, so that port must be open. For the complete list of required ports, see Port Reference.)

Image create, update, and delete (CUD) behavior depends on whether a cluster (also known as Prism Element) is registered to Prism Central.

- Images created on a Prism Element reside on Prism Element and can be managed from Prism Element.
- For better centralized management, you can migrate images manually to Prism Central by using the image import feature in Prism Central. An image migrated to Prism Central in this way remains on Prism Element, but you can manage the image only from Prism Central. Migrated images cannot be updated from the Prism Element.
- In the case of a local image upload, with more than one PE cluster managed by Prism Central, the image state is active on that PE cluster. All other PE clusters will show the image as inactive. If you create a VM from that image, the image bits are copied to the other PE clusters. The image then appears in an active state on all managed PE clusters.

To import and configure an image file, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Image Configuration** in the **Settings** page.

   The Image Configuration window appears.

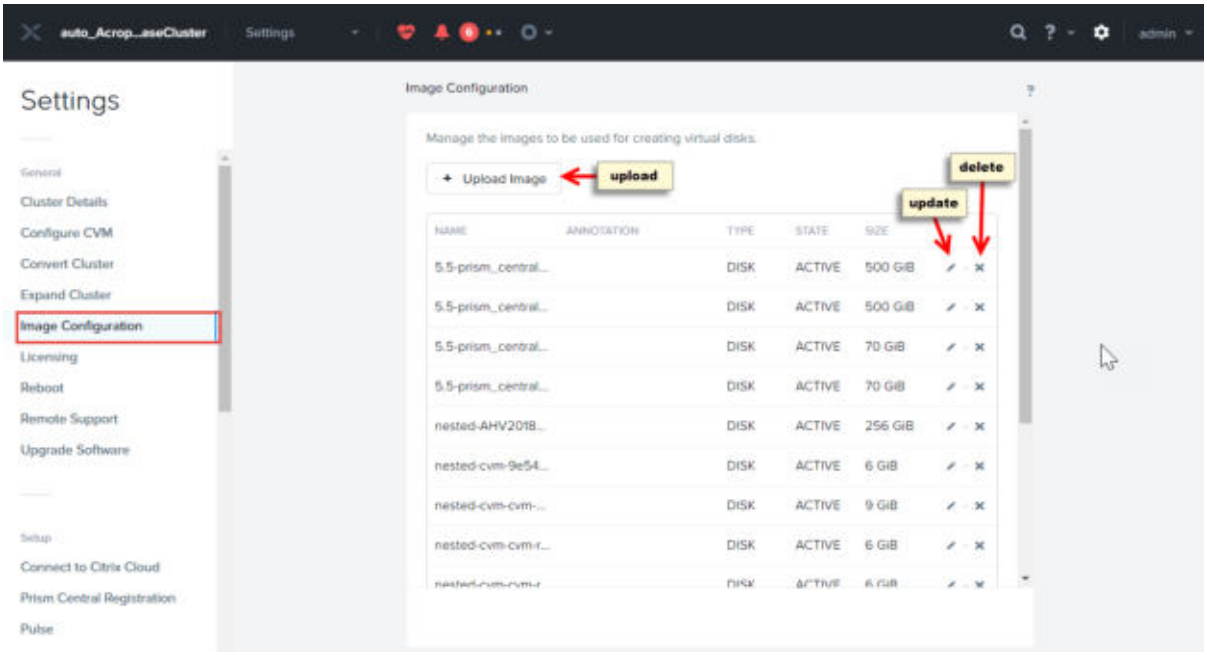

Figure 241: Image Configuration Window

2. To upload an image file to the cluster, click the **Upload Image** button.
   The Create Image window appears. Do the following in the indicated fields:

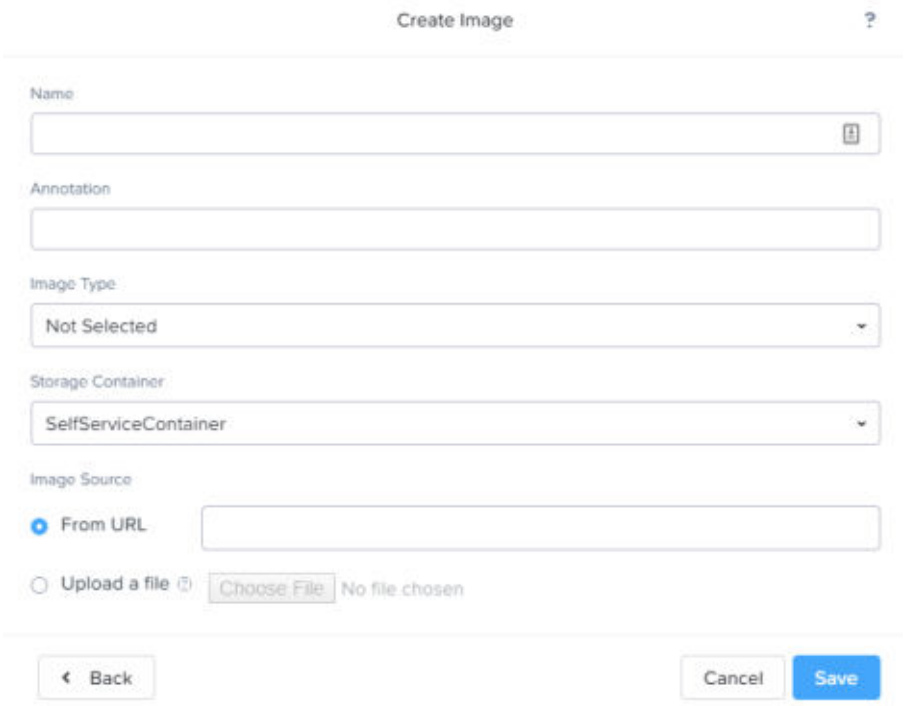

Figure 242: Create Image Window

a. **Name**: Enter a name for the image.

b. **Annotation** (optional): Enter a description for the image.

c. **Image Type** (optional): Select the image type, either **ISO** or **Disk**, from the pull-down list.

d. **Storage Container**: Select the storage container to use from the pull-down list.
   The list includes all storage containers created for this cluster. If there are no storage containers currently, a Create Storage Container link appears to create a storage container (see Creating a Storage Container on page 157).

e. **Image Source**: Do one of the following:

   - Click the **From URL** radio button to import the image from the Internet. Enter the appropriate URL address in the field using the following syntax for either NFS or HTTP. (NFS and HTTP are the only supported protocols.)

     ```
     nfs://[hostname|IP_addr]/path
     http://[hostname|IP_addr]/path
     ```

     Enter either the name of the host (*hostname*) or the host IP address (*IP_addr*) and the path to the file. If you use a *hostname*, the cluster must be configured to point at a DNS

server that can resolve that name (see Configuring Name Servers on page 751). A file uploaded through NFS must have 644 permissions.

If the image files have been copied to a container on the cluster, replace *IP_addr* with 127.0.0.1, the IP address of the loopback interface.

For example, enter nfs://127.0.0.1/*container_name*/*vm_name*/*vm_name*.vmdk or nfs://127.0.0.1/*container_name*/*vm_name*/*vm_name*-flat.vmdk.

Replace *container_name* with the name of the container where the image is placed and replace *vm_name* with the name of the VM where the image is placed.

- Click the **Upload a file** radio button to upload a file from your workstation. Click the **Choose File** button and then select the file to upload from the file search window.

f. When all the fields are correct, click the **Save** button.
The Create Image window closes and the Image Configuration window reappears with the new image appearing in the list.

3. To update the image information, click the pencil icon for that image.
The **Update Image** window appears. Update the fields as desired and then click the **Save** button.

> Note: The pencil icon is unavailable for images imported to Prism Central. Use Prism Central to manage such images.

4. To delete an image file from the store, click the X icon for that image.
The image file is deleted and that entry disappears from the list.

Virtual Machine Customization

In an Acropolis cluster, you can use Cloud-init to customize Linux VMs and the System Preparation (Sysprep) tool to customize Windows VMs.

About Cloud-Init

Cloud-init is a utility that is used to customize Linux VMs during first-boot initialization. The utility must be pre-installed in the operating system image used to create VMs. Cloud-init runs early in the boot process and configures the operating system on the basis of data that you provide (user data). You can use Cloud-init to automate tasks such as setting a host name and locale, creating users and groups, generating and adding SSH keys so that users can log in, installing packages, copying files, and bootstrapping other configuration management tools such as Chef, Puppet, and Salt. For more information about Cloud-init, see https://cloudinit.readthedocs.org/.

About Sysprep

Sysprep is a utility that prepares a Windows installation for duplication (imaging) across multiple systems. Sysprep is most often used to generalize a Windows installation. During generalization, Sysprep removes system-specific information and settings such as the security identifier (SID) and leaves installed applications untouched. You can capture an image of the generalized installation and use the image with an answer file to customize the installation of Windows on other systems. The answer file contains the information that Sysprep needs to complete an unattended installation. For more information about Sysprep and answer files, see the Microsoft Sysprep documentation.

The Customization Process in a Nutanix Cluster

You can use Cloud-init or Sysprep both when creating and when cloning VMs in a Nutanix cluster. For unattended provisioning, you can specify a user data file for Cloud-init and an answer file for Sysprep. All Cloud-init user-data formats are supported. For example, you can use the Cloud Config format, which is written in YAML, or you can provide a multi-part archive. To enable Cloud-init or Sysprep to access the script, AOS creates a temporary ISO image that includes the script and attaches the ISO image to the VM when you power on the VM.

> Note: The ISO image is mounted on bus IDE 3, so ensure that no other device is mounted on that bus.

You can also specify source paths to the files or directories that you want to copy to the VM, and you can specify the target directories for those files. This is particularly useful if you need to copy software that is needed at start time, such as software libraries and device drivers. For Linux VMs, AOS can copy files to the VM. For Windows VMs, AOS can copy files to the ISO image that it creates for the answer file.

After customizing a VM, you can copy the VDisk of the VM to Image Service for backup and duplication.

Customizing Linux Virtual Machines with Cloud-Init

Keep the user data file ready, either saved locally or uploaded to a storage container on the cluster. Alternatively, you can create or paste the script in the web console. If you want files copied to the VM during initialization, upload the files to a storage container on the cluster.

About this task

To customize a Linux VM by using Cloud-init, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials.
2. In the VM dashboard (see VM Dashboard on page 300), do one of the following:
   - » To create a VM, click **Create VM**.
   - » To clone a VM, click the VM that you want to clone, and then click **Clone**.
3. In the **Create VM** or **Clone VM** dialog box, specify a name for the VM and allocate resources such as vCPUs, memory, and storage. Select the **Custom Script** check box and specify how you want to customize the VM.

   For information about creating a VM and specifying customization options, see Creating a VM (AHV) on page 316. For information about cloning a VM, see Managing a VM (AHV) on page 326.
4. In the VM dashboard, select the VM, and then click **Power On**.
   The VM is powered on and initialized based on the directives in the user data file. To create a reference image from the VM, use Image Service. See Image Service in the VM Management chapter of the *Acropolis App Mobility Fabric Guide*.

Cloud-Init Limitations and Guidelines

- The maximum size of a VM guest customization script is 32 KB.
- Config drive version 2 data format is supported.

Customization of Windows Virtual Machines with System Preparation

To customize a Windows VM by using Sysprep, you need to perform the following tasks:

1. Create a reference image by using Sysprep.
2. Create a VM from the reference image.

You can also customize a VM when performing a fresh installation of Windows with an ISO file.

If you require unattended provisioning, keep the answer file ready, either saved locally or uploaded to a storage container on the cluster. Alternatively, you can create or paste the script in the web console. If you have files that need to be copied to the temporary ISO image, upload the files to a storage container on the cluster.

Preparing a VM to Create a Reference Image

Creating a reference image requires knowledge of Sysprep. For information about how to use Sysprep, see the Sysprep documentation on the Microsoft TechNet website.

About this task

To create a reference image, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials, and then browse to the VM dashboard (see VM Dashboard on page 300).
2. Select the VM that you want to clone, click **Launch Console**, and then log in to the VM with administrator credentials.
3. Configure Sysprep with the system cleanup action of your choice, specify whether or not you want to generalize the installation, and then choose to shut down the VM.
   - Open the command prompt as an administrator and navigate to the sysprep folder.
     ```
     cd C:\windows\system32\sysprep
     ```
   - Generalize the installation and then shut down the VM.
     ```
     C:\Windows\system32\sysprep>sysprep /generalize /shutdown /oobe
     ```
     The VM shuts down automatically.

     > Note: Make sure to shut down the VM. Restarting the VM will result in the VM losing its generalized state and in Sysprep attempting to find an answer file that has not been provided yet. For the same reasons, until you have completed this procedure, do not start the VM.
4. Create a reference image from the VM by using Image Service. See Configuring Images on page 358 for more information.

Creating a Customized Virtual Machine from a Reference Image

About this task

To use a reference image, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials, and then browse to the VM dashboard (see VM Dashboard on page 300).

2. Click **Create VM**, and then, in the **Create VM** dialog box, do the following:

   a. Specify a name for the VM and allocate resources such as vCPUs, memory, and storage.

   b. Click **Add new disk**, select the **Clone from Image Service** operation, and select the Windows reference image that you copied to Image Service.

   c. Click the **Custom Script** check box and specify how you want to customize the VM.

   For more information about creating a VM, see Creating a VM (AHV) on page 316.

3. In the VM dashboard, select the VM, and then click **Power On**.
   The VM is powered on and initialized based on the directives in the answer file. To create a reference image from the VM, use Image Service. See Image Service in the VM Management chapter of the *Acropolis App Mobility Fabric Guide*.

Customizing a Fresh Installation

You can perform a fresh installation only if you attach an empty vDisk and an installation CD-ROM to the VM. If you specify an image from Image Service or ADSF, for use as a vDisk, the VM is created from that image, and the install is no longer a fresh install.

About this task

To customize a fresh installation of Windows by using Sysprep, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials, and then browse to the VM dashboard (see VM Dashboard on page 300).

2. Click **Create VM**, specify the details that are required for installing Windows on the new VM, and then do the following:

   a. Specify a name for the VM and allocate resources such as vCPUs, memory, and storage.

   b. In the **Disks** area, click the edit button that is provided against the default CD-ROM entry. In the **Update Disk** dialog box, select the operation (**Clone from ADSF File** or **Clone from Image Service**), and then specify the image that you want to use. Click **Update**.

   c. Click **Add new disk**. Allocate space for a new disk on a storage container, and then click **Add**.

   d. Click the **Custom Script** check box and specify how you want to customize the VM.

   For more information about creating a VM, see Creating a VM (AHV) on page 316.

3. In the VM dashboard, select the VM, and then click **Power On**.
   The VM is powered on and initialized based on the directives in the answer file. To create a reference image from the VM, use Image Service. See Image Service in the VM Management chapter of the *Acropolis App Mobility Fabric Guide*.

VM High Availability in Acropolis

If you have not modified High Availability configuration from previous version of AOS releases, best effort VM availability is by default implemented in Acropolis.

If you have not enabled High Availability, in case of host failure VMs are restarted from the failed host to any available space on the other hosts in the cluster. Once the failed host joins the cluster again, VMs are migrated back to the host. This type of VM high availability is implemented without reserving any resources. Admission control is not enforced and hence there may not be sufficient capacity available to start all the VMs.

Note: Nutanix does not support VMs that are running with 100% remote storage for High Availability. The VMs must have at least one local disk that is present on the cluster.

VM high availability can be configured for host-based reservations using segment-based reservation method that is the default.

In segment-based reservation, the cluster is divided into segments to ensure that enough space is reserved for any host failure. Each segment corresponds to the largest VM that is guaranteed to be restarted in case the failure occurs. The other factor is the number of host failures that can be tolerated. Using these inputs, the scheduler implements admission control to always have enough resources reserved so that the VMs can be restarted upon failure of any host in the cluster.

The segment size ensures that the largest VM can be powered on in HA failover when cluster is fully loaded (if the cluster is fully used except the reserved segments). The number of segments that is reserved is such a way that for each host enough resources are reserved to ensure that any host failure in the cluster is tolerated. Multiple VMs may fit into a segment. If anything changes in the cluster, the reservation is computed again. The total resources reserved for segments can be more than the resources used by running VMs. This implementation guarantees successful failover even in the case of fragmentation of segments. The actual number of reserved resources depends on the current load of the cluster, but it is typically at 1 to 1.25 times the resource usage on the most loaded host.

If the host enters maintenance mode (in case of host upgrade), you might not be protected against further host failures. Maintenance mode uses reservations made for HA for migrating VMs from the host. Although you are not protected against host failure if you have reservation for HA, hypervisor upgrade occurs without any difficulty because from the perspective of a user it is the same as host failure except that the VMs are migrated (instead of restarted) and hence no runtime state is lost. The HA status goes through the same states as it goes when the host failure had occurred.

Segment-based reservation is the default method that Acropolis uses to enable VM high availability.

CAUTION: Do not change segment-based reservation method with the reserve host method.

Fault Detection for High Availability

Acropolis version 6.1 (with minimum supported AHV version 20201105.2229) onwards, the fault detection mechanism for High Availability checks for the heartbeat of the management services on the host. If any one of the services are down, then the host is marked as disconnected.

In addition to this, the fault detection mechanism also performs the following health checks for the host:

- Root file system corruption
- Read-only root file system

If any of the above-mentioned health checks are affirmative, then the host is marked as disconnected.

If the host remains in the disconnected status for 40 seconds, the VMs running on the affected host are automatically restarted (based on the resource availability) .

You can view the alerts raised for any of the above-mentioned checks in the **Activity** > **Alerts** view in Prism UI.

Enabling High Availability for the Cluster

In Acropolis managed clusters, you can enable high availability for the cluster to ensure that VMs can be migrated and restarted on another host in case of failure.

About this task

After you enable high availability for the cluster, if a host failure occurs the cluster goes through following changes.

- OK: This state implies that the cluster is protected against a host failure.
- Healing: Healing period is the time that Acropolis brings the cluster to the protected state. There are two phases to this state. The first phase occurs when the host fails. The VMs are restarted on the available host. After restarting all the VMs if there are enough resources to protect the VM, the HA status of the cluster comes back to OK state. If this does not occur, the cluster goes into critical state. The second phase occurs when the host comes back from the failure. Once the host comes back from failure, no VMs are present on the host and hence during this healing phase restore locality task occurs (VMs are migrated back). Apart from restoring the locality of the VMs, the restore locality task ensures that the cluster is back to the same state before the HA failure. Once it is finished, the HA status is back to OK state.
- Critical: If the host is down, the HA status of the cluster goes into Critical state. This happens because the cluster cannot tolerate any more host failures. You have to ensure that you bring back the host so that your cluster is protected against any further host failures.

> Note: On a less loaded cluster, it is possible for HA to go directly back to OK state if enough resources is reserved to protect another host failure. The start and migrate operations on the VMs are restricted in the Critical state because Acropolis continuously tries to ensure that the HA status is back to the OK state.

Procedure

1. Click the gear icon in the main menu and then select **Manage VM High Availability** in the **Settings** page.

   > Note: This option does not appear in clusters that do not support this feature.

   The Manage VM High Availability dialog box appears.

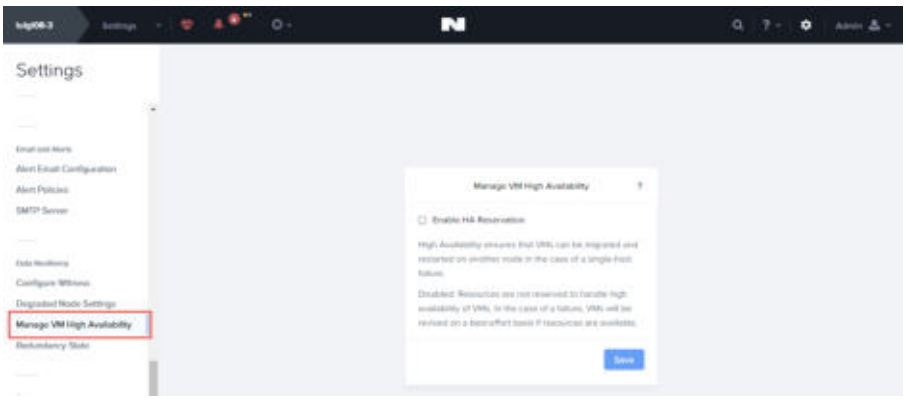

Figure 243: Manage VM High Availability Window

2. Check the **Enable HA Reservation** box and then click the **Save** button to enable.

Nutanix Guest Tools

Nutanix Guest Tools (NGT) is a software package that comes bundled with AOS. You can install NGT in a guest virtual machine (Microsoft Windows or Linux) to enable the advanced VM management features provided by Nutanix.

To use the NGT feature, you must do the following:

1. Enable the NGT feature for a VM in the Prism Element web console.

   When you enable NGT for a VM, Prism Element prepares the VM so that you can successfully install and use NGT in that VM.
2. Mount the NGT installer (ISO disk file) in a VM.
3. Install NGT in a VM.

If you install NGT in guest VMs, you can use the following advanced features (NGT applications):

File Level Restore CLI

Performs self-service file-level recovery from the VM snapshots. For more information about self-service restore, see *Self-Service Restore* in the *Data Protection and Recovery with Prism Element* guide.

Nutanix VM Mobility

Facilitates VM migration between ESXi and AHV, in-place hypervisor conversion, and cross-hypervisor disaster recovery (CHDR) features through the Nutanix VM Mobility drivers. For more information about cross-hypervisor disaster recovery, see *Nutanix Cross Hypervisor Disaster Recovery* in the *Data Protection and Recovery with Prism Element* guide. For more information about in-place hypervisor conversion, see In-Place Hypervisor Conversion on page 774.

VSS requestor and hardware provider for Windows VMs

Enables application-consistent snapshots of AHV or ESXi Windows VMs. For more information about Nutanix VSS-based snapshots for the Windows VMs, see the best practices section for taking application-consistent snapshots under *Implementation Guidelines for Asynchronous Disaster Recovery* in the *Data Protection and Recovery with Prism Element* guide.

Application-consistent snapshot for Linux VMs

Supports application-consistent snapshots for Linux VMs by running specific scripts on VM quiesce. For more information about Nutanix VSS-based snapshots for the Linux VMs, see the best practices section for taking application-consistent snapshots in *Implementation Guidelines for Asynchronous Disaster Recovery* in the *Data Protection and Recovery with Prism Element* guide.

The NGT software package includes the following:

NGT Installer

Allows you to install NGT in a guest VM.

Nutanix Guest Agent (NGA) Service

Maintains a communication channel between the Nutanix Controller VM and guest VMs.

Nutanix VirtIO Package

Includes the Nutanix VM mobility drivers that enable the following features:

- VM migration between ESXi and AHV
- In-place hypervisor conversion

- Cross-hypervisor disaster recovery (CHDR)

Nutanix Guest Tools Requirements and Limitations

Ensure that the following requirements are met, so that you can successfully use all the features of NGT.

General Requirements

- Virtual IP address must be configured for the Nutanix cluster. If you change the virtual IP address of the cluster, the reconfiguration impacts all the VMs in your cluster on which NGT is installed. See Virtual IP Address Impact on page 67 for more information.
- VMs must have at least one empty IDE CD-ROM or SATA slot to attach the ISO.
- The following ports must be open between the Controller VMs and guest VMs on which NGT is installed:
  - Controller VM: 2074. TCP port 2074 must be open in the Controller VM, so that user VMs can communicate with the Controller VM.
  - User VM: TCP port 23578 if you want to use the VSS service.

  For the complete list of required ports, see Port Reference.
- Guest VMs must be connected to a network that can be accessed by using the virtual IP address of the cluster.
- Hypervisor versions must be as follows:
  - ESXi version 5.1 or later
  - AHV version 20160215 or later

Windows Requirements

- Local VM accounts with administrative privileges are required.
- For Windows Server Edition VMs, ensure that Microsoft VSS is enabled before you install NGT.
- PowerShell path must be available in the environment for NGT installation. Otherwise, installation of NGT fails.
- For NGT installation on Windows VM, PowerShell path must be correctly set in the environment variables.
- For Windows to trust the device drivers for AHV, you must install the correct root CA in the domain. NGT installation fails if the root CA is not installed in the domain. For more information, see article KB 5488.

Linux Requirements

- For Linux operating systems, NGT is installed at **/usr/local**. Hence, ensure that the user has the write permissions for this location.
- For Linux operating systems, the lvm2 package version must be version 2.02.109 or later.

Limitations

- If a VM is connected to a volume group (VG), NGT captures the IQN of the VM and stores the information. If you change the VM IQN and you take a snapshot of the VM before the NGT refresh cycle (currently at 5 minutes) occurs, NGT will not be able to provide the auto restore capability because the snapshot operation will not be able to capture the VM-VG connection. As a workaround, restart the Nutanix Guest Agent service by running the `$ sudo service ngt_guest_agent restart` command on Linux VMs. For Windows VMs, in the command prompt, run the `net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"` command.

  > Note: By default, NGT client certificates expire every 1,000 days. Regenerate the certificates before they expire. See Regenerating NGT Certificates for Guest VMs on page 376 for instructions about how to regenerate the certificates.

See the supported operating system information for the specific NGT features to verify if the operating system is supported for the particular feature or not.

Enabling NGT and Mounting the NGT Installer in a VM

By default, the NGT feature is disabled for a VM running in a Nutanix cluster. To install and use the NGT feature in a VM, you must first enable the NGT feature (allow the installation and usage of NGT) in a VM and then mount the NGT installer (ISO disk file) in that VM by using the Prism Element web console.

About this task

When you are enabling the NGT feature and mounting the NGT installer in a VM, you must also select the NGT applications (self-service restore, volume snapshot service, and application-consistent snapshots) that you want to use in that VM.

Perform the following procedure to enable the NGT feature, mount the NGT installer, and select the NGT applications that you want to use in a VM by using the Prism Element web console.

> Note: For the ESXi clusters, you can enable NGT if you already have VMware guest tools installed on the VM.

Procedure

1. Log on to the Prism Element web console.
2. In the drop-down list of the main menu, select **VM** and select the **Table** view.

3. Select the VM for which you want to enable NGT and click **Manage Guest Tools**.

   The **Manage Guest Tools** dialog box appears. In the **Manage Guest Tools** dialog box, in addition to enabling the NGT feature and mounting the NGT installer, you must also select the NGT applications that you want to use in a VM.

   Do the following in the indicated fields.

   a. **Enable Nutanix Guest Tools**. Select this check box to enable the NGT feature in the selected VM.

   b. **Mount Nutanix Guest Tools**. Select this check box to mount the NGT installer on the selected VM.

      Note: Ensure that the VM has at least one empty IDE CD-ROM or SATA slot to attach the ISO.

   c. **Self Service Restore (SSR)**. Select this check box to enable the self-service restore feature for Windows VMs.

      For more information about the self-service restore feature, see *Self-Service Restore* in the *Data Protection and Recovery with Prism Element* guide.

   d. After you select the **Enable Nutanix Guest Tools** check box, the VSS and application-consistent snapshot applications are automatically selected.

      With these applications enabled, the Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent is used to take application-consistent snapshots for all the VMs that support VSS. This mechanism takes application-consistent snapshots without any VM stuns (temporary unresponsive VMs) and also enables third-party backup providers like CommVault and Rubrik to take application-consistent snapshots on the Nutanix platform

regardless of which hypervisor is used in the cluster. For more information, see *Conditions for Implementing Asynchronous Replication*.

e. Click **Submit**.

Prism Element now enables the NGT feature, mounts the NGT installer, and attaches a CD (ISO disk file) with the volume label NUTANIX_TOOLS to the selected VM.

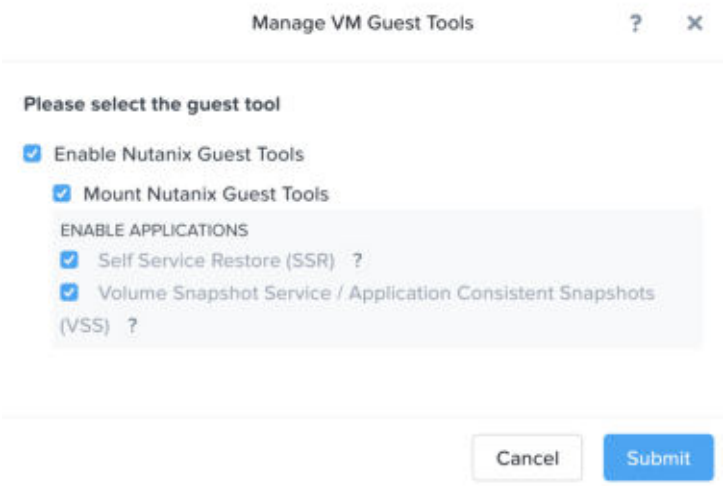

Figure 244: Manage NGT

> Note:
>
> If you clone a VM, NGT is not enabled on the cloned VM by default. If the cloned VM is powered off, enable NGT from the Prism Element web console and power on the VM. If the cloned VM is powered on, enable NGT from the Prism Element web console and restart the Nutanix Guest Agent service.
>
> Restart the Nutanix guest agent service by running the **$ sudo service ngt_guest_agent restart** command on Linux VMs. For Windows VMs, in the command prompt, run the **net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"** command.

If you eject the CD, you can manually mount the CD again by running the following nCLI command.

```
ncli> ngt mount vm-id=virtual_machine_id
```

For example, to mount the CD 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
```

NGT Installation

After you enable NGT and mount the NGT installer in a VM, you must install NGT in a VM (Windows or Linux) by logging into that VM. You can also use the silent installer package to install NGT on Windows VMs without affecting the normal VM operations.

Installing NGT on a Windows VM

The NGT software consisting of the NGT installer comes bundled with AOS. After you enable NGT and mount the NGT installer in a VM, log on to the Windows VM to install NGT.

Before you begin

You can verify if NGT is already installed on the Windows VM by monitoring the **Nutanix Guest Agent** service in the Windows Task Manager.

For more information about installing Nutanix VirtIO, used for installing Windows on a VM, see *Installing or Updgrading Nutanix VirtIO for Windows* in the *AHV Administration Guide*.

About this task

Procedure

1. Log on to the Windows VM.
2. In the Windows explorer, in the left pane, click **CD Drive NUTANIX_TOOLS**, and, in the right pane, double-click **setup** (setup.exe file).

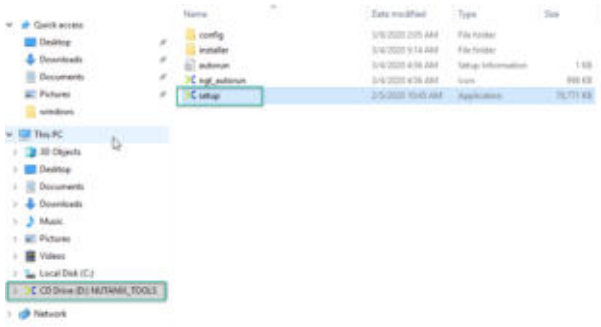

Figure 245: Start NGT Installation on Windows

As shown in the screenshot, the Nutanix icon appears next to the CD drive and setup.exe file.

> Note: If you mount the NGT installer while the VM is powered off, the Nutanix icon is not displayed after you power on the VM. You must open the CD (right-click the CD and click **Open**) and double-click the setup.exe file.

3. Accept the license agreement and follow the prompts to install NGT in the virtual machine.

> Note: The Nutanix VSS HW Provider does not get automatically registered after installation. The process of registration starts after you take the snapshot of the VM and it automatically unregisters after the operation is complete.

> CAUTION: If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and a prompt that asks for a restart of the VMs is displayed. Restart the VMs to upgrade the mobility driver version successfully.

> Note: For information on troubleshooting any NGT related issues, see KB-3741 available on the Nutanix support portal.

Installing NGT on Windows Machines (Silent Install)

The NGT software consisting of the NGT installer comes bundled with AOS. The silent installer package allows you to install NGT on Windows VMs without affecting the normal VM operations. You can also use the silent installer package to install NGT on multiple Windows VMs simultaneously. Add the commands described in this topic to a custom script to install NGT on multiple Windows VMs simultaneously.

About this task

The NGT silent installer package consists of the following components.

- Custom module that performs checks and fulfills all the prerequisites.
- Nutanix VM mobility drivers
- Python
- Microsoft Visual C++ 2008/2012 x64 redistributable setup for Windows Server operating systems
- Nutanix VSS package
- Nutanix Guest Agent service package

Before you begin

You can verify if NGT is already installed on the Windows VM by monitoring the **Nutanix Guest Agent** service in the Windows Task Manager.

About this task

Perform the following procedure to install NGT by using the NGT silent installer package.

Procedure

1. Open the command prompt and go to the drive on which the NGT installer is mounted.
2. Install NGT by using the silent installer package.

> CAUTION:
>
> If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and you must restart the VMs to upgrade the mobility driver version successfully. See the Release Notes of your AHV and AOS versions to verify if the Nutanix VM Mobility driver version is different in AHV and your VM.
>
> On AHV, depending on whether you need to restart the VMs or not, use one of the following commands to install NGT. You can use any of the following commands for ESXi, because

VMs running on ESXi are not affected due to the mismatch of the Nutanix VM Mobility driver version.

» `drive:\> setup.exe /quiet ACCEPTEULA=yes /norestart`

Use this command to ensure that VMs are not restarted after installing NGT. Note that if you use this command, mobility drivers are not updated. You must manually restart the VMs to update the mobility drivers.

» `drive:\> setup.exe /quiet ACCEPTEULA=yes`

If you use this command, VMs are restarted and drivers are updated.

By default, NGT logs are generated in **Event Viewer** and in the **%TEMP%** directory starting with *Nutanix_Guest_Tools_timestamp*. Event Viewer log files are created for all the components as part of the NGT installation.

Note: If you get any warnings after running this command, these warnings are treated as errors and are captured in the logs. If you want to ignore the warnings and complete the installation, you can use the flag IGNOREALLWARNINGS=yes. For example, drive:\> setup.exe /quiet ACCEPTEULA=yes IGNOREALLWARNINGS=yes. If you enable this flag, all the warnings are ignored and the silent installation proceeds smoothly.

3. (Optional) Install NGT if you want the logs to be generated in a location other than the **%TEMP%** directory.

   `drive:\> setup.exe /quiet ACCEPTEULA=yes -log log_path`

   Replace *log_path* with the path where you want to create the log files.

   Logs are generated in **Event Viewer** and in the path that you provide. Logs are not generated in the **%TEMP%** directory. Ensure that the path you provide has the necessary write permissions.

   CAUTION: If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and you must restart the VMs to upgrade the mobility driver version successfully. See the Release Notes of your AHV and AOS versions to verify if the Nutanix VM Mobility driver version is different in AHV and your VM.

   Note: For information on troubleshooting any NGT related issues, see KB-3741 available on the Nutanix support portal.

Installing NGT on Linux Machines

The NGT software consisting of the NGT installer comes bundled with AOS. After you enable NGT and mount the NGT installer in a VM, log on to the Linux VM to install NGT. Graphical Linux deployments auto-discover the CD-ROM and mount the CD-ROM appropriately. For non-graphical Linux deployments, use the following procedure.

Before you begin

You can run the following command to verify if NGT is already installed in the Linux guest VM.

```
[nutanix@localhost ~]$ sudo service ngt_guest_agent status
```

If NGT is installed, the command displays the following output.

```
Nutanix Guest Agent is running.
```

Procedure

1. Log on to the Linux VM.

2. Determine the directory in which the CD has been mounted.

   For example, the CD might be mounted in /dev/sr0 or /dev/sr1.

   You can use the following command to determine the directory where the CD has been mounted:

   ```
   dmesg | egrep -i --color 'cdrom|dvd|cd/rw|writer'
   ```

3. Create a temporary directory (/mnt) and mount the content of the CD in the temporary directory you created.

   ```
   $ sudo mount /dev/directory-name /mnt
   ```

   Replace *directory-name* with the directory you determined in step 2. For example, sr0 or sr1.

4. Install NGT.

   ```
   $ sudo /mnt/installer/linux/install_ngt.py
   ```

Enabling NGT and Mounting the NGT Installer on Cloned VMs

If you have cloned multiple VMs from a single VM (master VM), enable NGT and mount the NGT installer on the cloned VMs by using the master VM image.

Before you begin

Ensure the following before you perform this task:

1. Install NGT on the master VM.
2. Clone the required number of VMs from the master VM.
3. Shut down the cloned VMs.

Perform the following procedure to enable NGT and mount the NGT installer on cloned VMs.

> Note: After you perform the following procedure, you do not need to separately install NGT on the cloned VMs.

Procedure

1. For every cloned VM, log on to a Controller VM and run the following command.

   ```
   ncli> ngt mount vm-id=clone_vm_id
   ```

   Replace *clone_vm_id* with the ID of the cloned VM.

   To find the ID of the cloned VM, run `ncli> vm list name="<clone-vm-name>"` command. Note the value of the **Id** field as *clone_vm_id*.

   ```
   ncli> vm list name="<clone-vm-name>"

     Id                    : 00058a81-64bb-27af-69d2-ac1f6b602812::50132bb1-a07b-2e34-a768-d6b0ffc940e7
     Uuid                  : 5522f8cd-1c11-456c-b287-485e5e9de123
     Name                  : abcdef asdf
     VM IP Addresses       :
     Hypervisor Host Id    : 00058a81-64bb-27af-69d2-ac1f6b602812::3
     Hypervisor Host Uuid  : ae80deac-0224-43c9-9234-7a2071217e6d
     Hypervisor Host Name  : gologolo2-1
     Memory                : 1 GiB (1,073,741,824 bytes)
     Virtual CPUs          : 1
     VDisk Count           : 0
     Protection Domain     :
     Consistency Group     :
   ```

2. Power on the cloned VM.
3. (Optional) If you enable NGT and mount the NGT installer while the cloned VMs are powered on, do one of the following:
   - » Restart the Nutanix guest agent service by running the `$ sudo service ngt_guest_agent restart` command on Linux VMs. For Windows VMs, in the command prompt, run the `net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"` command.
   - » Guest power cycle the cloned VMs.

   Note: For information on troubleshooting any NGT related issues, see KB-3741 available on the Nutanix support portal.

Regenerating NGT Certificates for Guest VMs

NGT client certificates expire every 1000 days (about 3 years). This setting ensures that NGT meets the necessary security requirements.

About this task

Regenerate the NGT client certificates before they expire.

If the certificate of a guest VM is expiring in less than 45 days, the **Alerts** dashboard of the Prism Element web console displays an alert with a warning status. If the certificate is expiring in less than 7 days, the alert displays a critical status. The alert displays the name and UUID of the guest VM whose certificate is expiring. For more information, see Alerts Summary View on page 412.

If the certificate of a guest VM is expiring in less than 45 days, the **Health** dashboard displays the result of the health check in a warning state. If the certificate is expiring in less than 7 days, the health check displays critical state. The health check result displays the name and UUID of the guest VM whose certificate is expiring. For more information, see Health Dashboard on page 287

Alternatively, log on to a CVM with SSH and run the following NCC check **ngt_client_cert_expiry_check** to determine the NGT client certificates that are expiring in less than 90 days.

Perform the following procedure to regenerate the client certificates.

Procedure

1. Log on to the Prism Element web console.
2. In the **Alerts** or **Health** dashboard, determine the NGT client certificates that are expiring in less than 90 days. The dashboards display the name and UUID of the guest VMs whose certificates are expiring.

   Alternatively, log on to a CVM with SSH and run the following NCC check **ngt_client_cert_expiry_check** to determine the NGT client certificates that are expiring in less than 90 days.

   ```
   nutanix@cvm$ ncc health_checks ngt_checks ngt_client_cert_expiry_check
   ```

3. Log on to a CVM with SSH.

4. Regenerate the NGT client certificates.

   ```
   nutanix@cvm$ nutanix_guest_tools_cli refresh_vm_tools_entity [vm_uuids=string-containing vm_uuid1,vm_uuid2....]
   [threshold_days=number-of-days]
   ```

   vm_uuids and **threshold_days** are optional parameters.

   If you do not use the optional parameters, AOS regenerates the NGT client certificates of all guest VMs that are expiring in less than 45 days.

   If you want to use the **vm_uuids** parameter, replace *string-containing vm_uuid1,vm_uuid2....* with the UUID of the VM whose certificate you want to regenerate. If you want to regenerate certificates for multiple VMs, you can specify a comma-separated list of UUIDs.

   If you want to use the **threshold_days** parameter, replace *number-of-days* with the threshold value in number of days. For example, if you want to regenerate the certificates of VMs that are expiring in less than 30 days, replace *number-of-days* with 30.

5. Check the **Alerts** or **Health** dashboard in the Prism Element web console or run the NCC check mentioned in step 2 to verify if those VMs are still displayed in the alerts or NCC check.

   Note: For information on troubleshooting any NGT related issues, see KB-3741 available on the Nutanix support portal.

Nutanix Guest Agent and Controller VM Communication

After NGT you install in a VM, the Nutanix Guest Agent (NGA) service in the VM starts periodic communication with the Controller VM over SSL connections.

Each Nutanix cluster is configured as a Certificate Authority (CA). When NGT is enabled in a VM, a certificate pair is generated for that specific VM and it is embedded in an ISO that is configured for that VM. The security certificates are installed inside the VM as part of the installation process.

The NGA service running inside the VM initiates an SSL connection to the virtual IP port 2074 of the Controller VM to communicate with the Controller VM. No change in the firewall is required from the VM side. However, the VM must be able to reach port 2074 through SSL on the Controller VM. For the complete list of required ports, see Port Reference.

NGT provides three levels of security for VM communication.

- SSL certificates
- Capability-based authorization
- System or BIOS UUID of the VM must be same as seen by the hypervisor

Communication is successful only if all the three conditions are met. For example, if a VM that has NGT installed is cloned, the new VM will not be able to communicate with the Controller VM. You must separately enable NGT and then mount the NGT installer on the cloned VM. However, you do not need to separately install NGT again on the cloned VMs.

NGA also publishes information about the VM to the Controller VM. For example, guest OS type, status of VM mobility, and VSS services.

CD-ROM Eject Functionality of NGT

After NGT is installed in a VM, the CD-ROM is auto-ejected from the VM in the following two stages.

1. CD-ROM is automatically ejected from the guest VM immediately after the installation of NGT is successfully completed.

2. NGA then indicates to the Controller VM that the installation is complete, and the CD-ROM is detached from the VM through the hypervisor and destroyed. This operation depends on the periodicity of the VM and Controller VM communication and may take up to 10 minutes.

> Note:
>
> - There can be a time difference between these operations. Hence, there may be a possibility that the CD-ROM shows up as empty inside a VM, but from the NGT-Controller VM service Tools Mounted may be displayed as true. If this situation occurs, verify the hypervisor reported values.
> - If the NGT software version in the ISO is greater than the installed version inside the VM, the CD-ROM does not get automatically ejected. This helps with the upgrade of the NGT software inside the VM.

Nutanix Guest Tools Usage in Disaster Recovery

If a VM that is protected by the Nutanix data protection feature and has NGT enabled for it or vice-versa, the relevant NGT information including capabilities is added as part of the disaster recovery snapshot record.

If you migrate a VM to a remote site, restore it in place, or clone it from a snapshot, the NGT information is preserved for the restored VM. A new NGT ISO image containing only the relevant configuration information (SSL certificates, Controller VM IP address, etc.) is created for the recovered VM and the image is automatically attached to the VM.

When the VM is powered on, the NGA service running on the VM copies the relevant configuration information and detaches the CD-ROM automatically.

If you restore a VM on the remote site or retrieve the snapshot back to the local or source cluster, NGT information or functionality is lost. Therefore, you must enable NGT again for the restored VM.

NGT Metrics Collection for System Performance (Windows Perfmon)

You can monitor host specific metrics for a VM in a guest VM running the Windows OS. The supported metrics are fetched from the CVM and are published to the Microsoft Windows Performance Monitor (Perfmon) (also called System Monitor utility) on the UVM after enabling the NGT metrics collection at the UVM followed by Nutanix Guest Agent service restart.

Managing Metrics Collection in a Guest VM

This section describes how to enable and monitor the NGT metrics in a guest VM that run Microsoft Windows operating systems. Nutanix publishes VM-specific metrics in the NGT Metrics performance counter object in the Perfmon utility. You can leverage all Perfmon functionalities to view the metrics data in different ways. For example, you can view the metrics in different graph formats. The options available for graphs are Line, Histogram, and Report.

The guest VM makes an RPC call to the CVM every 30 seconds to fetch the metrics data. Once the data is fetched successfully, it internally checks if this counter is already registered in the Perfmon utility. If the counter is not registered, it first registers the counter and then publishes this data to the Perfmon utility. By default, the following metrics are collected from the CVM after enabling NGT metrics collection in the guest VM.

- hypervisor_cpu_usage_ppm. Indicates the VM processor time in percentage.
- hypervisor_cpu_ready_time_ppm. Indicates the VM stolen time in percentage.

> Important: NGT metrics collection is disabled by default in the guest VMs.

- For information on enabling NGT metrics collection in the guest VMs, see Enabling or Disabling Metrics Collection on UVM on page 379.
- For information on monitoring or viewing NGT metrics in the guest VM, see Monitoring NGT Metrics for System Performance in Guest VM on page 380.

Managing Metrics Collection from CVM

Although the NGT metrics collection is enabled by default on CVMs, you can monitor the NGT metrics in the Perfmon utility only in the guest VM.

You can modify the metrics list from the CVM by editing the ~/config/nutanix_guest_tools/ngt_metrics_info.json file to add or remove the metrics to be published on the UVMs. This change in the metrics data to be collected reflects only after the NGA service restarts on the CVM. After the guest VM fetches the metrics, it updates this data in the Perfmon utility without a manual intervention (like service restart) in the guest VM.

> Note: For example, in the ngt_metrics_info.json file you can modify the array list to add or remove metrics in the metrics info block. Any changes in the metrics list reflects after you restart the NGA service on the CVM.
>
> ```
> {
>   "metrics_info" : [
>     "hypervisor_cpu_usage_ppm",
>     "hypervisor.cpu_ready_time_ppm"
>   ]
> }
> ```

Known Issues and Limitations

- Any failure in registering a counter or in publishing the metrics data is logged in the Nutanix Guest Agent logs in the guest VM. Currently, the user is not alerted if such an event occurs.
- In Windows Performance monitor, the NGT hypervisor CPU metric is collected in Parts per Million. Since the default scale is 1, it gives values over 100%. Hence, it is recommended to adjust the scale in the **Performance Monitor Properties** > **Data** tab to .0001.

Enabling or Disabling Metrics Collection on UVM

NGT metrics collection is disabled by default in the guest VM. You can manage NGT metrics collection at the guest VM for system performance through the NGT CLI by setting the respective gflag.

About this task

Perform the following procedure to enable NGT metrics.

Procedure

1. Log on to the UVM where NGT is installed.
2. Go to **Computer** > **Program Files** > **Nutanix** > **ngtcli** directory.
3. Run ngtcli.exe.

4. Run the following command.

   ```
   ncli> ngt set-gflag ngt_metrics_collection_enabled=True
   ```

   The NGT metrics collection is enabled after you restart the **Nutanix Guest Agent** service.

   > Note: To disable the NGT metrics collection, set ngt_metrics_collection_enabled=True.

5. To restart the NGA service, perform the following.

   a. Go to windows Task Manager.

   b. Right-click and restart the **Nutanix Guest Agent** service from the list of services.

What to do next

After you successfully enable metrics collection at the guest VM, the Ngt Metrics collection capabilities got enabled string reflects in the nutanix_guest_tools.out log file.

Enabling or Disabling Metrics Collection on CVM

NGT metrics collection is enabled by default on CVMs. You can manage NGT metrics collection at CVM for system performance through the NGT CLI by setting the respective gflag.

About this task

Perform the following procedure to disable NGT metrics collection at CVM.

Procedure

1. Log on to the CVM.
2. Edit the ~/config/nutanix_guest_tools.gflags file to add an entry for the gflag, enable_ngt_metrics_info_collection_capabilities = False.

   The NGT metrics collection is enabled after you restart the Nutanix Guest Tools service.
3. Run the following command to restart the NGT service in the CVM.

   ```
   nutanix@cvm$ genesis stop nutanix_guest_tools && cluster start
   ```

   > Note: To enable the NGT metrics collection, set the gflag, enable_ngt_metrics_info_collection_capabilities=True.

Monitoring NGT Metrics for System Performance in Guest VM

Before you begin

- Verify if the guest VM has Microsoft Windows OS and NGT installed.
- Verify if the guest VM has Microsoft Windows Perfmon utility installed. You can monitor the supported NGT metrics in the guest VM only when the Perfmon utility is installed.
- Verify if the NGT metric collection is enabled in the guest VM. See Enabling or Disabling Metrics Collection on UVM on page 379 for information on how to enable NGT metrics collection from CVM.
- Ensure that the CVM timezone is in sync with the global timezone. Failing to do so does not generate the accurate metrics results.

About this task

Perform the following steps to view the NGT metrics in the guest VM.

Procedure

1. In the VM dashboard, click the **Table** view.
2. Select the target VM in the table.
3. Click **Launch Console** and login.
4. Go to **Start** > **Run**.
5. Type perfmon and click **OK**.
   This launches the Performance Monitor utility.
6. Select **Performance Monitor** under **Monitoring Tools**.

7. Click the **Add** icon to view the counters.
   The Add Counters window appears. You can observe that the NGT metrics are listed in the counter list.

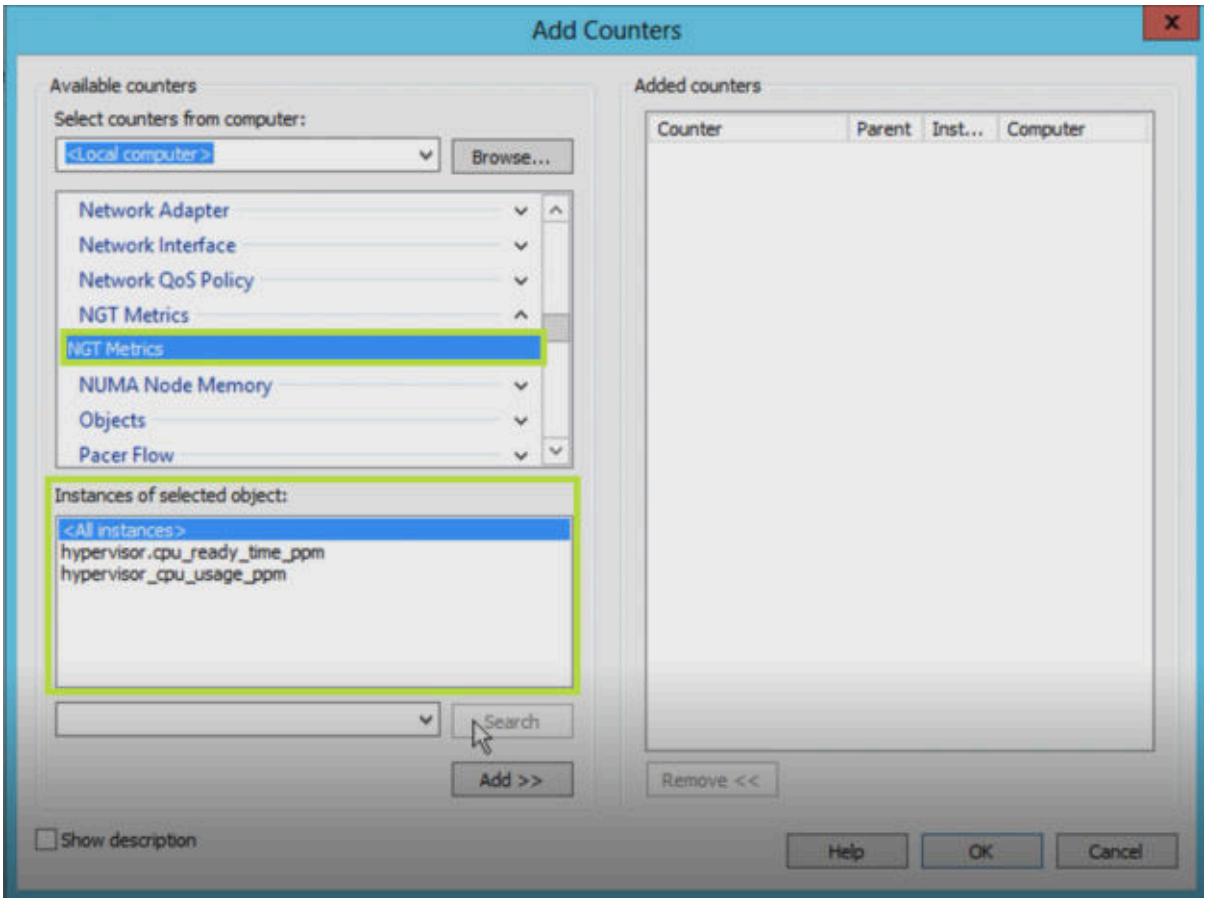

Figure 246: Performance Monitor (Add Counters)

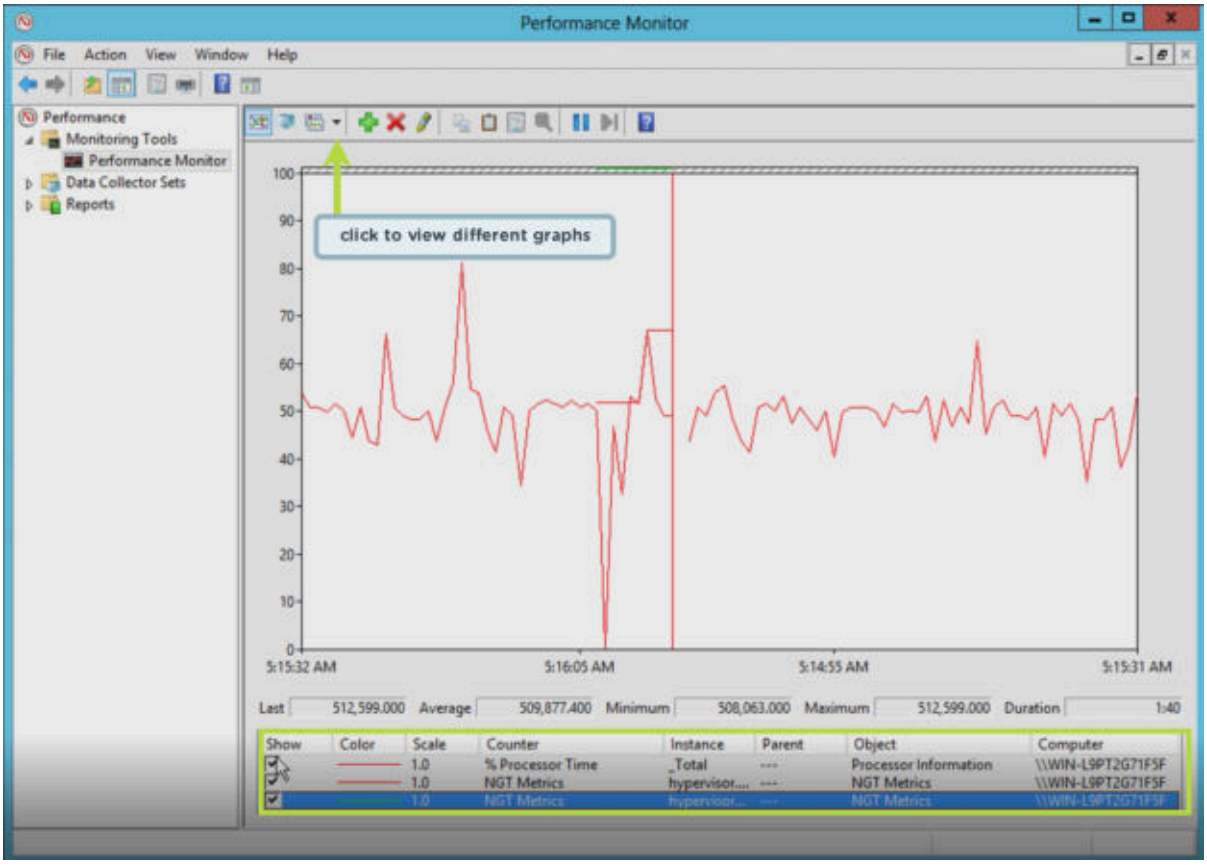

Figure 247: NGT Metrics (Graphical View)

What to do next

After metrics are collected, you can leverage all Perfmon functionalities like generating reports, viewing metrics graphically.

Upgrading NGT

If you have upgraded your AOS, you can upgrade NGT to the latest version. NGT is not automatically upgraded.

Perform the following procedure to upgrade NGT.

1. Enable NGT and mount the NGT installer. See Enabling NGT and Mounting the NGT Installer in a VM on page 369 for more information.
2. Install NGT again. See NGT Installation on page 371 for more information.

> Note: If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and a prompt that asks for a restart of the VMs is displayed. Restart the VMs to upgrade the mobility driver version successfully.

Reconfiguring NGT

If you reconfigure the cluster IP address, NGT loses connection with the CVM. You must then reconfigure NGT to establish the connection again.

About this task

Perform the following procedure to reconfigure NGT.

Procedure

1. Mount the NGT installer. For more information about mounting the NGT installer, see Enabling NGT and Mounting the NGT Installer in a VM on page 369.
2. Do one of the following.
   - » Restart the VM.
   - » If you do not want to restart the VM, restart the Nutanix Guest Agent (NGA) service. To restart the NGA service, use one of the following commands, depending on your operating system.

     On Linux VMs, run the command: `$ sudo service ngt_guest_agent restart`.

     On Windows VMs, in the command prompt, run the command: `net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"`.

     After you mount the NGT ISO and restart NGA, NGA fetches the latest configuration (new cluster IP address) mounted in the guest VM. The guest VM can now use the new IP address to communicate with the cluster.

     Note: For information on troubleshooting any NGT related issues, see KB-3741 available on the Nutanix support portal.

Disabling, Uninstalling, and Removing Nutanix Guest Tools

If you disable NGT, only the communication between the VM and NGT-Controller VM service is stopped. To completely remove NGT from the database, disable NGT by using the Prism Element web console, uninstall NGT by using the guest VM console, and then remove NGT by using nCLI.

About this task

Perform the following procedure to disable, uninstall, and remove NGT from a guest VM.

Procedure

1. Disable NGT.
   a. Log on to the Prism Element web console.
   b. Go to the **Table** view of the **VM** page.
   c. Select the VM in which you want to disable NGT.
   d. Click **Manage Guest Tools**.
   e. Clear the **Enable Nutanix Guest Tools** check box.
   f. Click **Submit**.

2. Uninstall NGT.

   For Windows VMs, log on to the Windows VM and run the following command.

   ```
   drive:\> setup.exe /uninstall /quiet ACCEPTEULA=yes /norestart
   ```

   For Linux VMs, log on to the Linux VM and run the following command:

   ```
   $ /usr/local/nutanix/ngt/python/bin/python3 /usr/local/nutanix/ngt/python/bin/uninstall_ngt.py
   ```

   For Windows VMs, log on to the Windows VM and perform the following.

   a. Go to **Control Panel** > **Programs** > **Programs and Features** page.

   b. Select the **Nutanix Guest Tools** service.

   c. Click **Uninstall**.

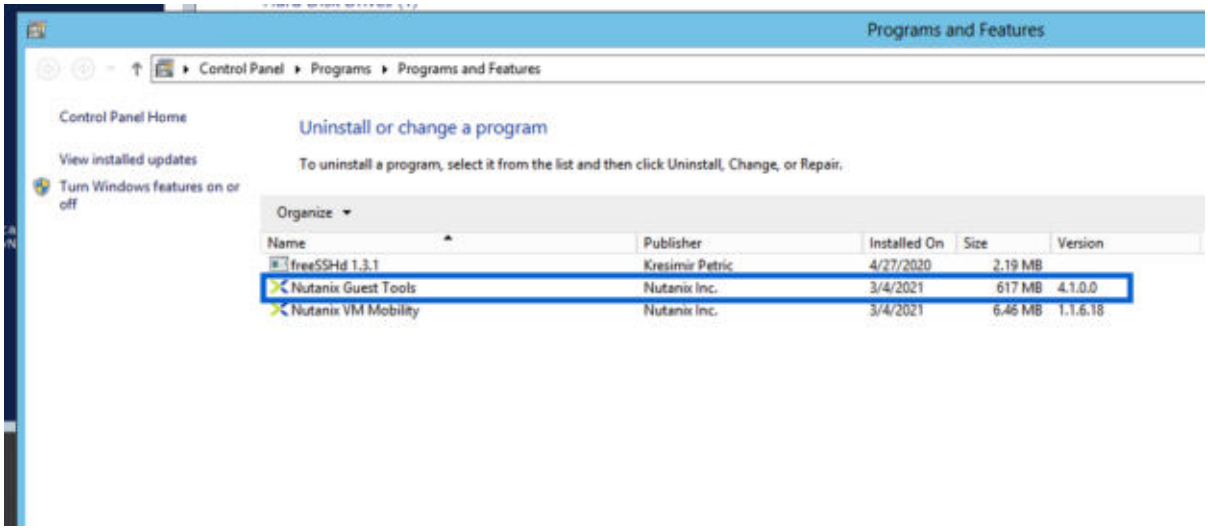

Figure 248: Uninstall NGT

The following confirmation message appears on successfully uninstalling NGT.

Figure 249: NGT Uninstall Complete

Note: To uninstall NGT using the PowerShell command, perform the following steps.

Run the following PowerShell command.

```
Get-ChildItem -Path HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall, HKLM:\SOFTWARE
\Wow6432Node\Microsoft\Windows\CurrentVersion\Uninstall | Get-ItemProperty | Where-Object {$_.DisplayName -
match "Nutanix Guest Tools" } | Select-Object -Property QuietUninstallString | Out-String -Width 4096
```

Result: An output string appears on the console. For example,

```
"C:\ProgramData\Package Cache\{3cfa83ac-a36f-49f1-ad23-3a51c0e6964a}\NutanixGuestTools.exe" /uninstall /quiet
```

To the generated output string, append /norestart, and run the updated command from the command prompt. For example,

```
"C:\ProgramData\Package Cache\{3cfa83ac-a36f-49f1-ad23-3a51c0e6964a}\NutanixGuestTools.exe" /uninstall /quiet /
norestart
```

Note: If you have an NGT ISO mounted, run the following command to uninstall NGT.

```
drive:\> setup.exe /uninstall /quiet ACCEPTEULA=yes /norestart
```

3. Remove NGT completely.

   a. Log on to the Controller VM with SSH.

   b. Determine the ID of the VM from which you want to remove NGT.

```
nutanix@cvm$ ncli ngt list

  VM Id                    : 0005b6a7-6bcc-03f2-0000-0000000097fe::e4a1216a-a287-43c2-bbbf-28951e1bf615
  VM Name                  : win 2012
  NGT Enabled              : true
  Tools ISO Mounted        : false
  Vss Snapshot             : false
  File Level Restore       : false
  Communication Link Active : true
```

   In the VM ID column, the text after "::" is the ID of the VM as shown in this example.

   c. Run the following command.

```
nutanix@cvm$ ncli ngt delete vm-id=virtual_machine_id
```

   Replace *virtual_machine_id* with the ID of the VM you determined in step b.

   > Note: For information on troubleshooting any NGT related issues, see KB-3741 available on the Nutanix support portal.

Affinity Policies Defined in Prism Element

In Prism Element, you can define scheduling policies for virtual machines on an AHV cluster at a VM level. By defining these policies, you can control the placement of a virtual machine on specific hosts within a cluster.

You can define two types of affinity policies in Prism Element.

VM-Host Affinity Policy

The VM-host affinity policy controls the placement of a VM. 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.
> - VMs configured with host affinity settings retain these settings if the VM is migrated to a new cluster. Remove the VM-host affinity policies applied to a VM that you want to migrate to another cluster, as the UUID of the host is retained by the VM and it does not allow VM restart on the destination cluster. When you attempt to protect such VMs, it is successful. However, some disaster recovery operations like migration fail and attempts to power on these VMs also fail.

You can define the VM-host affinity policies by using Prism Element during the VM create or update operation. For more information, see *Creating a VM (AHV)* in *Prism Web Console Guide* or *AHV Administration Guide*.

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 the 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 388.
> - 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 though if a host is removed from a cluster, the host UUID is not removed from the host-affinity list for a VM.

Configuring VM-VM Anti-Affinity Policy

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.

About this task

> Note: Currently, the VM-VM affinity policy is not supported.

Perform the following procedure to configure the VM-VM anti-affinity policy.

Procedure

1. Log on to the Controller VM with SSH session.
2. Create a group.

   ```
   nutanix@cvm$ 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.

   ```
   nutanix@cvm$ 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.

   ```
   nutanix@cvm$ 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 on to the Controller VM with SSH session.
2. Remove the VM-VM anti-affinity policy.

   ```
   nutanix@cvm$ 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).

Connect to Citrix Cloud

The **Connect to Citrix Cloud** feature leverages the automated installer functionality of the Citrix Cloud connector to establish a secure communication channel between Nutanix and Citrix Cloud. This feature provides single-click integration of the on-premise Nutanix clusters as a resource location with the Citrix Cloud environment.

> Note: This feature is supported on AHV only.

This feature helps you configure the Citrix Cloud integration settings in the following way:

1. Establishes the connection to the Citrix Cloud workspace.
2. Configures Nutanix cluster as resource location in the Citrix cloud.
3. Configures the Citrix Cloud connector VM.
4. Registers the Citrix Cloud connector VM to the Active Directory (AD) domain on the Citrix cloud.

Once the integration is complete, VDIs can be created using the *Nutanix AHV MCS Plug-in for Citrix XenDesktop* 1.1.1.0 or later. The AHV MCS Plug-in is designed to create and manage VDIs in a Nutanix Acropolis infrastructure environment. See the *Nutanix AHV MCS Plug-in for Citrix XenDesktop 7.9 or later install guide and release notes* for details.

Thus, to begin deploying your VMs and applications, you must perform the following:

- Configure the Citrix cloud integration settings using the **Connect to Citrix Cloud** feature.
- Install the *Nutanix CWA Plug-In for Citrix Cloud Connector*.

Refer to the *Citrix documentation* for details on Citrix Cloud connector.

Connecting to the Citrix Cloud

To integrate with the Citrix Cloud workspace, perform the following procedure.

Before you begin

Ensure that you have created a Sysprep (System Preparation) VM with Windows 2012 R2 server as the base image and the VM should be available on the Nutanix cluster.

> CAUTION: The Sysprep VM must be in the powered-off state. If you power on the VM, you will lose the Sysprep state and the configuration will fail.

See *Microsoft documentation* on *Sysprep (Generalize) a Windows installation* for more information.

Procedure

1. Log on to the Prism Web Console using your Nutanix administrator credentials.

   > Note: The procedure does not work if Prism Element is launched from Prism Central web console.

2. Click the gear icon and select **Connect to Citrix Cloud**. Click the gear icon in the main menu and then select **Connect to Citrix Cloud** in the **Settings** page.

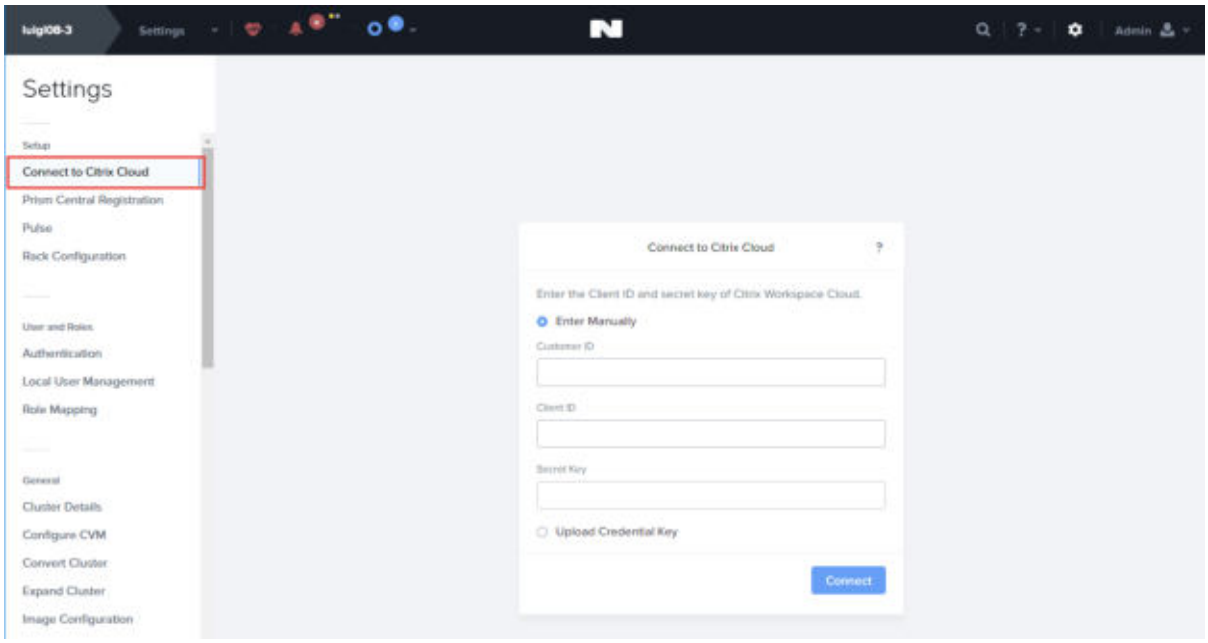

Figure 250: Connect to Citrix Cloud

3. Authenticate your connection to the Citrix cloud by using one of the following methods.

   - **Enter Manually**

     1. Enter your **Customer ID** for the Citrix Cloud.
     2. Enter your secure **Client ID** for the Citrix Cloud.
     3. Enter the downloaded secure client **Secret Key**.

        > Note: You can find the **Customer ID**, **Client ID**, and **Secret Key** from the *API Access* page in the Citrix Cloud console.

     4. Click **Connect**.

   - **Upload Credential Key**

     1. Click **Upload Key File** to browse and select the key file in the CSV format.

        You can create or download the key file from the *API Access* page in the Citrix Cloud console (within *Identity and Access Management*). This key file is used for the Citrix Cloud connector installation.
     2. Enter your **Customer ID** for the Citrix Cloud.
     3. Click **Connect**.

4. Select the **Resource Location** from the drop-down menu.

   The drop-down menu displays the resource locations that are created on the Citrix Cloud. If you have not previously created any resource locations, enter a name for the resource location that you want to create in the **Resource location** field.

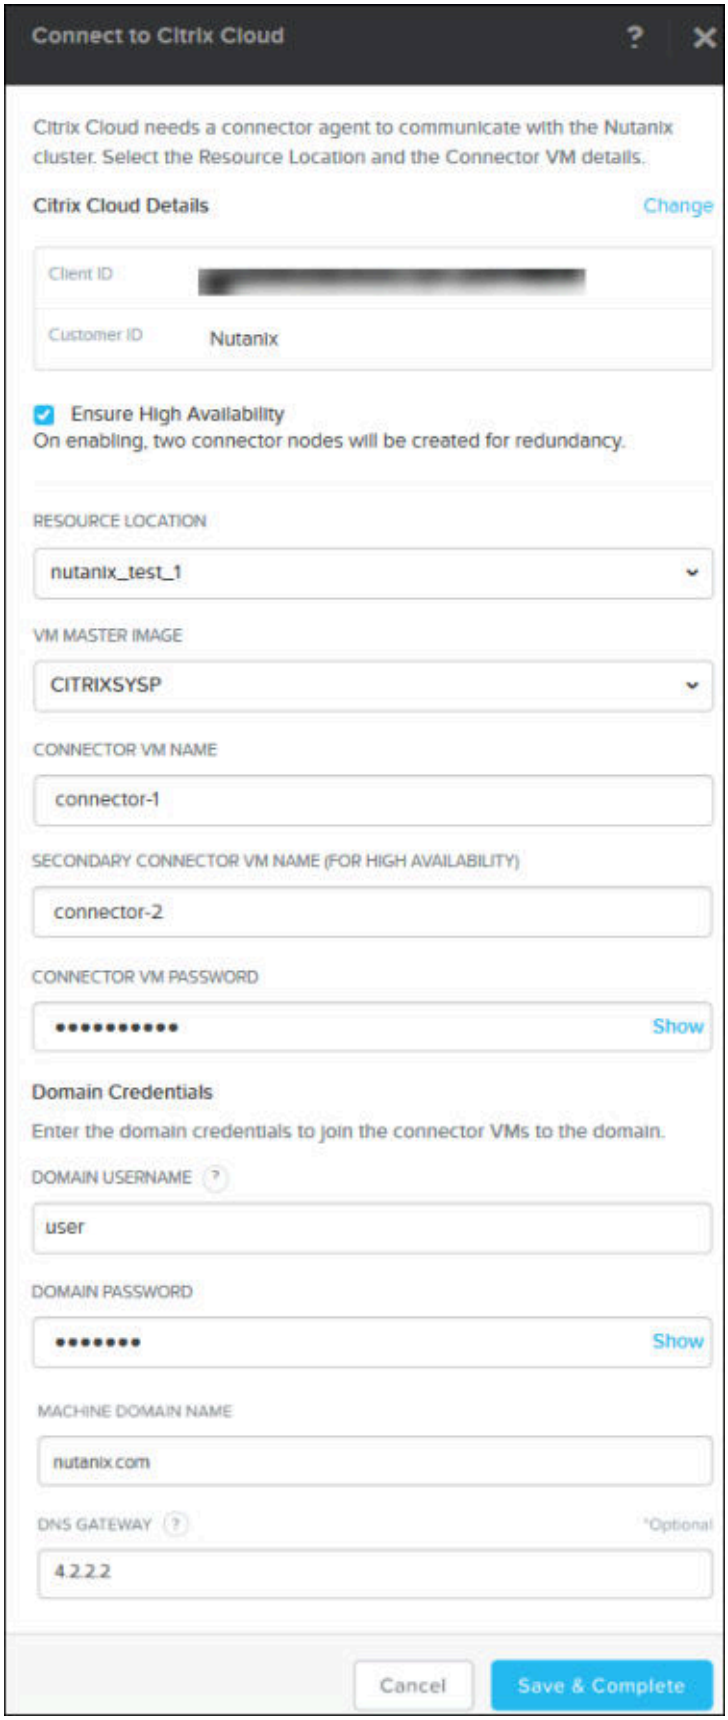

Figure 251: Configuring Resource Location and Connector VM

5. (Optional) Review the **Citrix Cloud** details. If any change is required, click the **Change** hyperlink to edit the connection details.
6. (Optional) Select the **High Availability** check-box to enable or disable high availability for the connector nodes.
   By default, high availability is enabled. On enabling high availability, two connector nodes are created for the redundancy purpose.

   > Note: Citrix Cloud recommends that you install two connectors for redundancy and high availability.

7. In the **VM Master Image** search box, start typing the initial letters of the previously created Sysprep VM image and select the auto-completed option.
8. Enter the **Connector VM Name**.
9. If high availability is enabled, enter the **Secondary Connector VM Name (for high availability)**.
10. Enter the **Connector VM Password**.
11. Enter the **Domain Credentials** to join your enterprise domain to the resource location.
    a. Enter the **Domain Username** in the format USERDOMAIN\USERNAME.
    b. Enter the **Domain Password**.
    c. Enter the **Machine Domain Name**.
    d. Enter the **DNS Gateway**.
12. Click **Save and Continue**.
    The connector VM or VMs are configured successfully.

What to do next

- Verify the connection status, see Viewing Citrix Connection Status on page 393.
- Verify if the connector VM or VMs are created from the **VM Dashboard**.
- Wait for the state of the connector VM or VMs to change from powered off to powered on and then click **Launch Console**. The VM preparation process starts.
- Once the VM starts, install the *Nutanix CWA Plug-In for Citrix Cloud Connector* to start the application deployment.

Viewing Citrix Connection Status

You can view the current status of your Citrix Cloud connection. This helps to verify the registration of the connector VM to the Citrix cloud.

To view details of the Citrix connection, click the gear icon and select **Connect to Citrix Cloud**. If the connector VM is configured, the connector VM name and the status are listed under Connector VMs. It might take some time for the connection status to refresh.

The Connection Status shows the following statuses:

- Connected - The connector VM is successfully registered with Citrix cloud.
- Connection in progress - The connector VM is registration with Citrix cloud is in progress.

- Not Connected - The connector VM is not registered with the Citrix cloud or the Citrix registration has failed. In this case, wait for the connection status to refresh, it may take several minutes depending on your environment. Otherwise, click **Delete Connection** to delete the Citrix cloud connection and configure it again. The **Delete Connection** option deletes the connector VM and unregisters the account from the Citrix Cloud.

  Note: For XenServer, the **Delete Connection** option does not delete the connector VMs. You must delete the connector VMs manually.

Guest VM Cluster Configuration (AHV Only)

In AHV clusters, you can create guest VM clusters by either directly attaching volume groups to guests VMs or by using iSCSI. Follow this document if you want to create guest VM clusters by directly attaching volume groups to guests VMs.

See the Nutanix Volumes Guide if you want to create guest VM clusters by using iSCSI.

Creating a Guest VM Cluster by Directly Attaching a Volume Group (AHV Only)

In AHV clusters, you can create a guest VM cluster by directly attaching a volume group to guest VMs. After you attach a volume group to guest VMs, vDisks appear as SCSI devices to the guest operating system and you do not need to set up any in-guest connections when you are creating a guest cluster. If you directly attach volume groups to guest VMs, you can seamlessly share vDisks across VMs in the guest cluster.

About this task

You can directly attach a volume group to guest VMs to create the following guest clusters:

- Microsoft Failover Cluster (MSFT)
- Red Hat Enterprise Linux (RHEL) Cluster

  Note: To create an MSFT cluster, ensure that the minimum Nutanix VirtIO version installed on the guest VM is 1.1.4.

You cannot create a guest VM cluster by directly attaching a volume group to guest VMs in ESXi and Hyper-V clusters.

Perform the following sequence of tasks to create a guest cluster by directly attaching a volume group to guest VMs.

Procedure

1. Log on to the Prism Element web console.
2. Create the number of VMs you want in the guest cluster.

   See Creating a VM (AHV) for instructions about how to create a VM in an AHV cluster.
3. Create a volume group for the guest cluster.

   See Creating a Volume Group for instructions about how to create a volume group.
4. Attach the volume group you created to each VM in the guest cluster.

   See Managing a VM (AHV) for instructions about how to attach a volume group to a VM.

5. Create a guest cluster.

   See the documentation of your guest operating system for instructions about how to create a guest cluster.

11

PERFORMANCE MONITORING

Nutanix provides several mechanisms to maximize performance in the cluster. The converged Distributed Storage Fabric (DSF) architecture is designed to service the VM requests locally on each node whenever possible. Each node employs data tiers so that frequently accessed ("hot") data is retained in memory or solid state disk (SSD) storage while seldom accessed ("cold") data is moved to hard disk drive (HDD) storage. Each Controller VM has an in-memory read cache to access highly requested data directly from memory.

- The web console allows you to monitor and analyze performance across the cluster (see Analysis Dashboard on page 396).

Analysis Dashboard

The Analysis dashboard allows you to create charts that can monitor dynamically a variety of performance measures. To view the Analysis dashboard, select **Analysis** from the pull-down list on the left of the main menu.

Menu Options

The Analysis dashboard does not include menu options other than those available from the main menu (see Main Menu on page 46).

Analysis Screen Details

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 a single metric for one or more entities. 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 & Events 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 3 hours to a month) 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 and events. Any alerts and events 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 64 for information about how the metrics are measured.

Figure 252: Analysis Dashboard

Table 41: 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	Allows you to create a chart that tracks a single metric for one or more entities (see Creating a Metric Chart on page 399).
New Entity Chart	Allows you to create a chart that tracks one or more metrics for a single entity (see Creating an Entity Chart on page 398).
(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, WTD [week to date], 1 month) 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 & Events Monitor	Displays a monitor of alert and event messages that were generated during the time interval. Alerts and events 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 = orange, informational alert = blue, event = gray).
(defined chart monitors)	Displays monitors for any enabled (checked) charts. In the figure above, three charts are enabled (memory usage, CPU/memory, and disk IOPS). You can export the chart data by clicking on the chart header. 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 (see Alerts Dashboard). Clicking a message causes the monitor line to move to the time when that alert occurred.
Events	Displays the event messages that occurred during the time interval. Clicking a message causes the monitor line to move to the time when that event occurred.

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 (see Analysis Dashboard on page 396), click **New** > **New Entity Chart**.
   The New Entity Chart dialog box appears.
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 include host, disk, storage pool, storage container, virtual machine, volume group, 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.

      Note: If you are creating this chart for Prism Central, the list spans the registered clusters. Otherwise, the list is limited to the current cluster.

   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 400.

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.

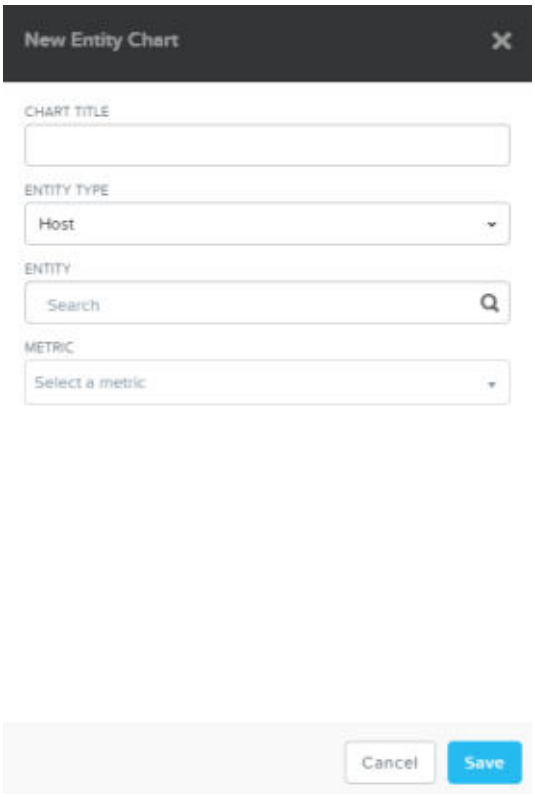

Figure 253: New Entity Chart Window

Creating a Metric Chart

About this task

A metric chart monitors the performance of a single metric for one or more entities. To create a metric chart definition, do the following:

Procedure

1. In the Analysis dashboard (see Analysis Dashboard on page 396), click **New** > **New Metric Chart**.
   The New Metric Chart dialog box appears.

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 400.

   c. **Entity Type**: Select an entity type from the pull-down list. (Repeat to include additional entities.)
   The entity types include host and cluster.

   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. (Repeat to include additional names.)

   > Note: If you are creating this chart for Prism Central the list spans the registered clusters. Otherwise, the list is limited to the current cluster.

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.

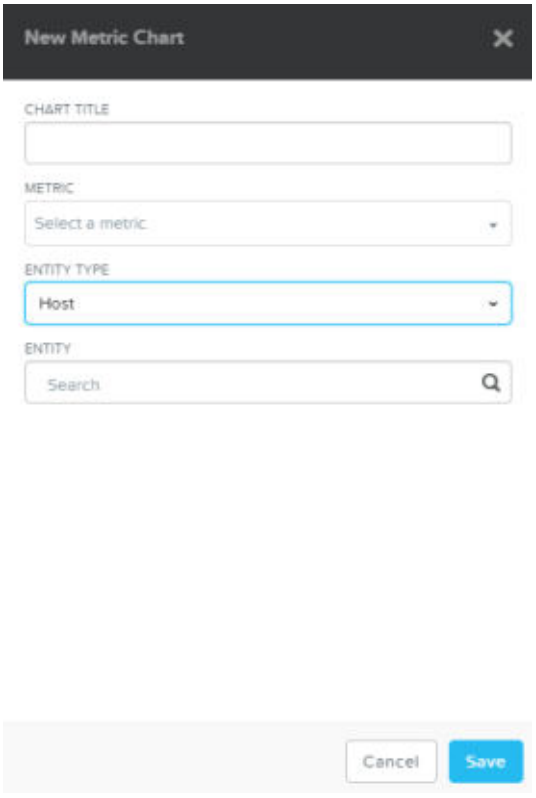

Figure 254: New Metric Chart Window

Chart Metrics

These metrics can be added to charts.

> Note: The mapping between a metric and an entity type is hypervisor dependent.

Metric	Entity Type(s)	Description
Content Cache Hit Rate (%)	• Host • Cluster	Content cache hits over all lookups. ID: CONTENT_CACHE_HIT_PPM
Content Cache Hits	• Host • Cluster	Number of hits on the content cache. ID: CONTENT_CACHE_NUM_HITS
Content Cache Logical Memory Usage	• Host • Cluster	Logical memory (in bytes) used to cache data without deduplication. ID: CONTENT_CACHE_LOGICAL_MEMORY_USAGE_BYTES
Content Cache Logical SSD Usage	• Host • Cluster	Logical SSD memory (in bytes) used to cache data without deduplication. ID: CONTENT_CACHE_LOGICAL_SSD_USAGE_BYTES
Content Cache Lookups	• Host • Cluster	Number of lookups on the content cache. ID: CONTENT_CACHE_NUM_LOOKUPS
Content Cache Memory Saved	• Host • Cluster	Memory (in bytes) saved due to content cache deduplication. ID: CONTENT_CACHE_SAVED_MEMORY_USAGE_BYTES
Content Cache Physical Memory Usage	• Host • Cluster	Real memory (in bytes) used to cache data by the content cache. ID: CONTENT_CACHE_PHYSICAL_MEMORY_USAGE_BYTES
Content Cache Reference Count	• Host • Cluster	Average number of content cache references. ID: CONTENT_CACHE_NUM_DEDUP_REF_COUNT_PPH
Content Cache SSD Usage	• Host • Cluster	Real SSD usage (in bytes) used to cache data by the content cache. ID: CONTENT_CACHE_PHYSICAL_SSD_USAGE_BYTES
Content Cache SSD Usage Saved	• Host • Cluster	SSD usage (in bytes) saved due to content cache deduplication. ID: CONTENT_CACHE_SAVED_SSD_USAGE_BYTES
Deduplication Fingerprints Cleared	• Host • Cluster	Number of written bytes for which fingerprints have been cleared. ID: DEDUP_FINGERPRINT_CLEARED_BYTES
Deduplication Fingerprints Written	• Host • Cluster	Number of written bytes for which fingerprints have been added. ID: DEDUP_FINGERPRINT_ADDED_BYTES

Metric	Entity Type(s)	Description
Disk I/O Bandwidth	• Host • Cluster • Disk • Storage Pool	Data transferred per second in KB/second from disk. ID: STATS_BANDWIDTH
Disk I/O Bandwidth - Read	• Host • Cluster • Disk • Storage Pool	Read data transferred per second in KB/second from disk. ID: STATS_READ_BANDWIDTH
Disk I/O Bandwidth - Write	• Host • Cluster • Disk • Storage Pool	Write data transferred per second in KB/second from disk. ID: STATS_WRITE_BANDWIDTH
Disk I/O Latency	• Host • Cluster • Disk • Storage Pool	I/O latency in milliseconds from disk. ID: STATS_AVG_IO_LATENCY
Disk IOPS	• Host • Cluster • Disk • Storage Pool	Input/Output operations per second from disk. ID: STATS_NUM_IOPS
Disk IOPS - Read	• Host • Cluster • Disk • Storage Pool	Input/Output read operations per second from disk. ID: STATS_NUM_READ_IOPS
Disk IOPS - Write	• Host • Cluster • Disk • Storage Pool	Input/Output write operations per second from disk. ID: STATS_NUM_WRITE_IOPS

Metric	Entity Type(s)	Description
GPU Framebuffer Usage	Virtual Machine	Framebuffer usage in percentage. ID: FRAMEBUFFER_USAGE_PPM Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU Usage	Virtual Machine	GPU compute usage in percentage. ID: GPU_USAGE_PPM Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU video decoder Usage	Virtual Machine	GPU video decoder usage in percentage. ID: DECODER_USAGE_PPM Note: The Virtual Machine entity is applicable for AHV hypervisor.
GPU video encoder usage	Virtual Machine	GPU video encoder usage in percentage ID: ENCODER_USAGE_PPM Note: The Virtual Machine entity is applicable to AHV hypervisor.
Hypervisor CPU Ready Time (%)	Virtual Machine	Percentage of time that the virtual machine was ready, but could not get scheduled to run. ID: STATS_HYP_CPU_READY_TIME
Hypervisor CPU Usage (%)	• Host • Cluster • Virtual Machine	Percent of CPU used by the hypervisor. ID: STATS_HYP_CPU_USAGE
Hypervisor I/O Bandwidth	• Host • Cluster • Virtual Machine	Data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_BANDWIDTH Note: Virtual Machine entity is applicable for ESXi. hypervisor
Hypervisor I/O Bandwidth - Read	• Host • Cluster • Virtual Machine	Read data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_READ_BANDWIDTH Note: Virtual Machine entity is applicable for ESXi hypervisor.

Metric	Entity Type(s)	Description
Hypervisor I/O Bandwidth - Write	• Host • Cluster • Virtual Machine	Write data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_WRITE_BANDWIDTH Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency	• Host • Cluster • Virtual Machine	I/O latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_IO_LATENCY Note: Virtual Machine entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency - Read	• Host • Cluster	I/O read latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_READ_IO_LATENCY Note: Cluster entity is applicable only for AHV hypervisor.
Hypervisor I/O Latency - Write	• Host • Cluster	I/O write latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_WRITE_IO_LATENCY Note: Cluster entity is applicable only for AHV hypervisor.
Hypervisor IOPS	• Host • Cluster • Virtual Machine	Input/Output operations per second from Hypervisor. ID: STATS_HYP_NUM_IOPS Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor IOPS - Read	• Host • Cluster • Virtual Machine	Input/Output read operations per second from Hypervisor. ID: STATS_HYP_NUM_READ_IOPS Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.

Metric	Entity Type(s)	Description
Hypervisor IOPS - Write	• Host • Cluster • Virtual Machine	Input/Output write operations per second from Hypervisor. ID: STATS_HYP_NUM_WRITE_IOPS Note: • Cluster entity is applicable for AHV hypervisor. • Virtual Machine entity is applicable for ESXi.
Hypervisor Memory Usage (%)	• Host • Cluster • Virtual Machine	Percent of memory used by the hypervisor. ID: STATS_HYP_MEMORY_USAGE
Logical Usage	Storage Container	Logical usage of storage (physical usage divided by replication factor). ID: STATS_UNTRANSFORMED_USAGE
Memory Usage(%)	• Virtual Machine	Percent of memory used by the hypervisor. Consolidated guest memory usage in percentage. ID: MEMORY_USAGE_PPM
Network Rx Bytes	Virtual Machine	Write data transferred in KB/second from the Storage Controller. ID: HYPERVISOR_NUM_RECEIVED_BYTES Note: Virtual Machine entity is applicable for AHV hypervisor.
Network Tx Bytes	Virtual Machine	Write data transferred per second in KB/second. ID: HYPERVISOR_NUM_TRANSMITTED_BYTES Note: Virtual Machine entity is applicable for AHV hypervisor.
Physical Usage	• Host • Cluster • Disk • Storage Pool • Storage Container	Actual usage of storage. ID: STATS_TRANSFORMED_USAGE

Metric	Entity Type(s)	Description
Read IOPS (%)	• Host • Cluster • Disk • Storage Pool	Percent of IOPS that are reads. ID: STATS_READ_IO_PPM
Replication Bandwidth - Received	• Cluster • Remote Site • Protection Domain • Replication Link	Replication data received per second in KB/second ID: STATS_REP_BW_RECEIVED
Replication Bandwidth - Transmitted	• Cluster • Remote Site • Protection Domain • Replication Link	Replication data transferred per second in KB/second ID: STATS_REP_BW_TRANSFERRED
Replication Bytes - Received	• Cluster • Remote Site • Protection Domain • Replication Link	Number of bytes received. ID: STATS_REP_NUM_RECEIVED_BYTES Note: Cluster entity is applicable for AHV hypervisor.
Replication Bytes - Total Received	Replication Link	Total number of bytes received. ID: STATS_REP_TOT_RECEIVED_BYTES
Replication Bytes - Total Transmitted	Replication Link	Total number of bytes transmitted. ID: STATS_REP_TOT_TRANSMITTED_BYTES
Replication Bytes - Transmitted	• Cluster • Remote Site • Protection Domain • Replication Link	Number of bytes transmitted. ID: STATS_REP_NUM_TRANSMITTED_BYTES Note: Cluster entity is applicable for AHV hypervisor.

Metric	Entity Type(s)	Description
Storage Controller Bandwidth	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_BANDWIDTH
Storage Controller Bandwidth - Read	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Read data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_READ_BANDWIDTH
Storage Controller Bandwidth - Write	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Write data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_WRITE_BANDWIDTH
Storage Controller IOPS	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_IOPS Note: The Host entity is applicable for AHV hypervisor only.

Metric	Entity Type(s)	Description
Storage Controller IOPS - Read	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output read operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_READ_IOPS Note: The Host entity is applicable for AHV hypervisor.
Storage Controller IOPS - Read (%)	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Percent of Storage Controller IOPS that are reads. ID: STATS_CONTROLLER_READ_IO_PPM Note: The Host entity is applicable for AHV hypervisor.
Storage Controller IOPS - Write	• Host • Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Input/Output write operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_WRITE_IOPS
Storage Controller IOPS - Write (%)	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Percent of Storage Controller IOPS that are writes. ID: STATS_CONTROLLER_WRITE_IO_PPM

Metric	Entity Type(s)	Description
Storage Controller Latency	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	I/O latency in milliseconds from the Storage Controller. ID: STATS_CONTROLLER_AVG_IO_LATENCY Note: • Custer entity is applicable only for AHV hypervisor. • Virtual Machine entity is applicable for AHV and Hyper-V hypervisors.
Storage Controller Latency - Read	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Storage Controller read latency in milliseconds. ID: STATS_CONTROLLER_AVG_READ_IO_LATENCY Note: Custer entity is applicable for AHV hypervisor.
Storage Controller Latency - Write	• Cluster • Storage Container • Virtual Machine • Volume Group • Virtual Disk	Storage Controller write latency in milliseconds. ID: STATS_CONTROLLER_AVG_WRITE_IO_LATENCY Note: Custer entity is applicable for AHV hypervisor.
Storage container own usage	Storage Container	Storage container's own usage + Reserved (not used). ID: NEW_CONTAINER_OWN_USAGE_LOGICAL
Swap In Rate	Virtual Machine	Rate of data being swapped in. ID: STATS_HYP_SWAP_IN_RATE Note: Virtual Machine entity is applicable for ESXi and Hyper-V hypervisors.
Swap Out Rate	Virtual Machine	Rate of data being swapped out. ID: STATS_HYP_SWAP_OUT_RATE Note: Virtual Machine entity is applicable for ESXi and Hyper-V hypervisors.

Metric	Entity Type(s)	Description
Virtual NIC bytes received packets with error.	Virtual Machine	Virtual NIC bytes received packets with error. STATS_NETWORK_ERROR_RECEIVED_PACKETS Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC bytes received rate.	Virtual Machine	Virtual NIC bytes received rate in kbps. STATS_NETWORK_RECEIVED_RATE Note: Virtual Machine entity is applicable only for ESXi hypervisor.
Virtual NIC bytes transmitted rate.	Virtual Machine	Virtual NIC bytes transmitted rate in kbps. STATS_NETWORK_TRANSMITTED_RATE Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC dropped transmitted packets.	Virtual Machine	Number of dropped transmitted packets by the Virtual NIC. STATS_NETWORK_DROPPED_TRANSMITTED_PACKETS Note: Virtual Machine entity is applicable for ESXi hypervisor.
Virtual NIC receive packets dropped.	Virtual Machine	Number of receive packets dropped by the Virtual NIC. STATS_NETWORK_DROPPED_RECEIVED_PACKETS Note: Virtual Machine entity is applicable for ESXi hypervisor.
Write IOPS (%)	• Host • Cluster • Disk • Storage Pool	Percent of IOPS that are writes. ID: STATS_WRITE_IO_PPM

Exporting Performance Data

This topic describes how to view and export the performance chart.

About this task

Performance data is stored for three months in the UI. You can export the performance data results in CSV or JSON format.

Note: Performance data is available for 3 months.

Procedure

1. Log into the Prism Web Console.
2. Navigate to the Analysis Dashboard. Click **Home** > **Analysis**.
   The Analysis page displays.
3. Set the range to **1 Month**. The **1 Month** range shows the data in monthly segments.

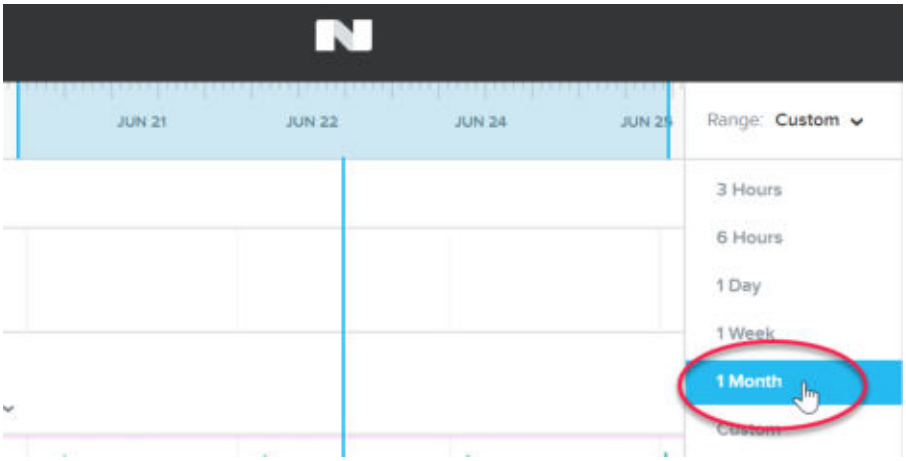

Figure 255: Set Performance Data Range

4. Export the performance data into a CSV or JSON file. Click the drop-down arrow next to the cluster chart you want to export.

Figure 256: Export Performance Data

12

ALERT AND EVENT MONITORING

The Prism web console provides several mechanisms to monitor and configure cluster alerts and events.

- The Alerts Dashboard on page 412 allow you to view and monitor alert and event activity in the cluster.
  - The Alerts Summary View on page 412 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.
  - The Events Summary View on page 418 displays a list of events that you filter in various ways. You can also drill down for detailed information for an event.
- You can configure rules for who should receive email alerts and customize the alert messages (see Configuring Alert Emails on page 420).
- You can customize the list of events that generate an alert (see Configuring Alert Policies on page 424).

When alerting is enabled, some potentially high impact events also generate a case automatically with Nutanix customer support. Events that might generate a Nutanix support case include hardware problems such as clock, power supply, or disk failures; monitoring problems such as Stargate outages or Curator scan failures; and data processing problems such as disabled compression or deduplication.

> Note: These views reflect that Prism 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 Dashboard

The Alerts dashboard displays alert and event messages. To view the Alerts dashboard, select **Alerts** from the pull-down list on the far left of the main menu.

The Alerts dashboard provides two viewing modes.

- Click the **Alerts** tab (on the left) to display the alert messages (see Alerts Summary View on page 412).
- Click the **Events** tab to display the event messages (see Events Summary View on page 418).

Alerts Summary View

The Alerts summary view displays a list of alert messages. 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.

> Note: See Alerts/Health checks on page 426 for a list of alert messages and corrective actions.

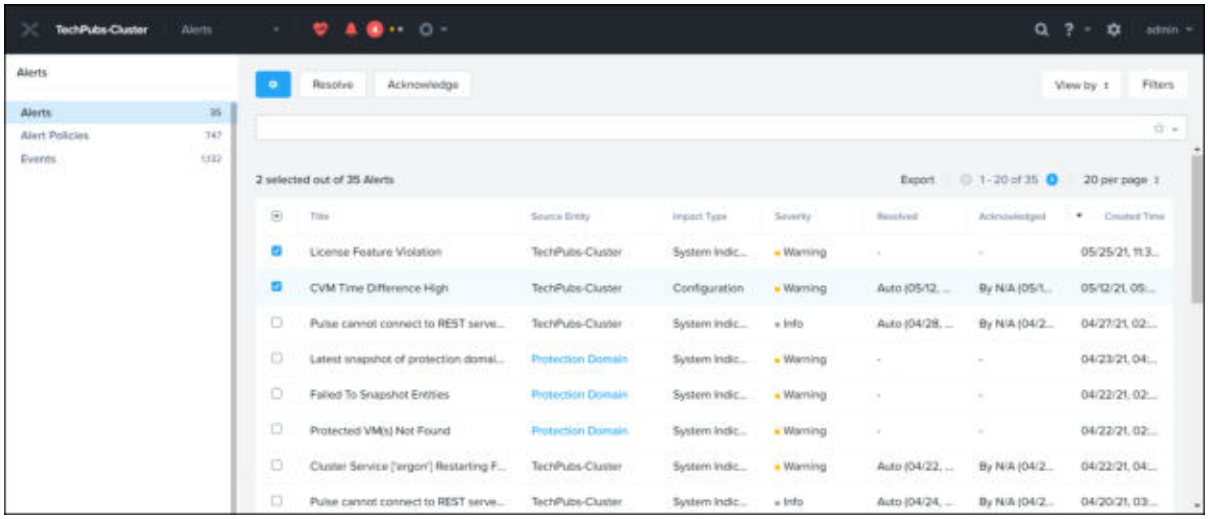

Figure 257: Alerts View

- To order the alerts chronologically, click the **Created 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 **Export**. The maximum number of alerts you can export is 1000.
- To display the details page for an alert, click the alert title (see Alert Details on page 415).
- To enable alert emails and specify email addresses to which alerts should be sent, click the **Email Configuration** button (see Configuring Alert Emails on page 420). This button turns into a gear icon when you select one or more alerts in the table.
- To configure alert policies, click **Alert Policies** to specify what events should generate an alert (see Configuring Alert Policies on page 424).
- 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.
  - 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.
  - After you select the first page of alerts, click **Select All Alerts** to select 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 42: Alerts View Fields

Parameter	Description	Values
(selection box)	Click this box to select the message for acknowledgement or resolution.	n/a
Title	Displays the title of the alert. Click the title to displays the details page for that alert (see Alert Details on page 415).	(alert title)
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 (Info)** 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 424). • 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 424). • 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)

Alert Details

Clicking on an alert message in the summary view (see Alerts Summary View on page 412) 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.

Figure 258: Alert Details Screen

- 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 43: 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)

Parameter	Description	Values
Severity	Displays the severity level of this condition. There are three levels: **Critical** An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. **Warning** An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. **Informational (Info)** An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational
Created Time	Displays the date and time when the alert first occurred.	(time and date)
Last Occurred	Displays the date and time for the most recent occurrence of the alert.	(time and date)
Impact Type	Displays the impact type to which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Policy	Displays the name of the alert policy.	(alert policy name)
Status (for manually resolved alerts)	Displays the current alert status: acknowledged, resolved, or neither (blank).	(blank), Acknowledged, Resolved
Status (for auto resolved alerts)	Displays the current alert status: acknowledged, resolved, auto resolved, or neither (blank).	(blank), Acknowledged, Auto Resolved

Parameter	Description	Values
Acknowledged By	Indicates whether the alert has been acknowledged. • A blank (dash) value means the alert is not acknowledged. • An "N/A *date_time*)" entry indicates this alert was acknowledged automatically at the specified date and time. This requires that auto alert is enabled for this alert type (see Configuring Alert Policies on page 424). • 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*

Events Summary View

The Events summary view displays a list of event messages. Event messages describe cluster actions such as adding a storage container 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.

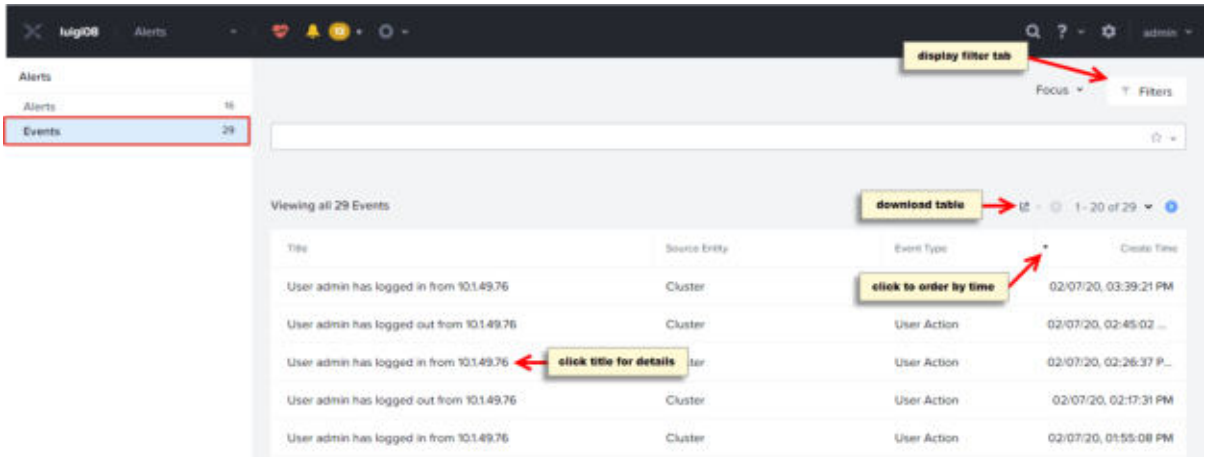

Figure 259: Events View

- To display the details page for an event, click the event title (see Event Details on page 419).

- To order the events chronologically, click the **Create Time** column header, which toggles the event list latest-to-earliest or earliest-to-latest.
- 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
  - Time Range: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*

Table 44: 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 419).	(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 type of event. Three types are defined, and each event is assigned to one of these types.	Behavioral Anomaly, System Action, User Action
Create Time	Displays the date and time when the event occurred.	(time and date)

Event Details

Clicking on an event message in the summary view (see Events Summary View on page 418) 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.

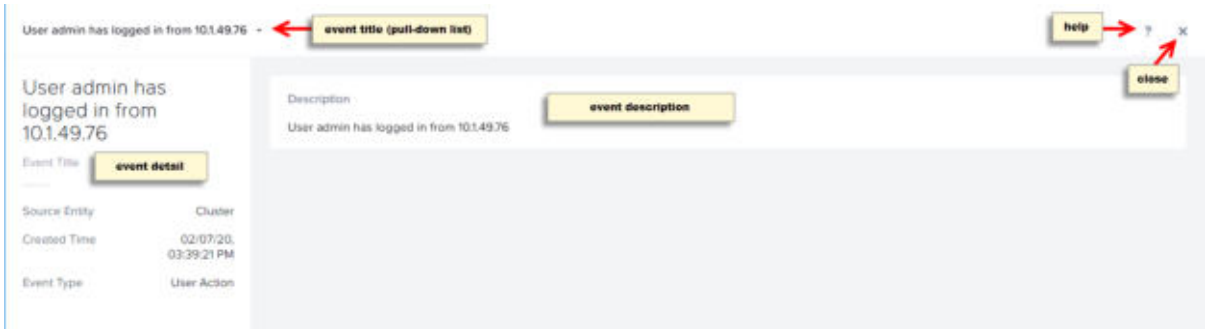

Figure 260: Event Details Screen

- 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 45: 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)
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

E-mail notification of alerts is enabled by default and sends alert messages automatically to Nutanix customer support through customer-opened ports 80 or 8443. For the complete list of required ports, see Port Reference.To customize who should receive the alert e-mails (or to disable e-mail notification), do the following:

Procedure

1. Do one of the following:

   » Click the gear icon in the main menu and then select **Alert Email Configuration** in the **Settings** page.

   » Go to the Alerts dashboard, click the **Configure** button, and select **Email Configuration** from the pull-down menu.

   The Alert Email Configuration window appears.

2. Click the **Settings** tab and do the following:

   a. Select the **Email Every Alert** box to send an email whenever the event occurs.

   b. Select 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.

   c. Select the **Nutanix Support** box to send alert notification to Nutanix customer support. Note that an alert notification to the Nutanix customer support is sent only if you have selected the **Email Every Alert** option.

   d. To send alert notifications to others, enter their email addresses in a comma separated list in the field below the **Nutanix Support** box.

   e. 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, default Nutanix tunnel is used for mail service. For more information on configuring the SMTP server settings window (see Configuring an SMTP Server on page 754).

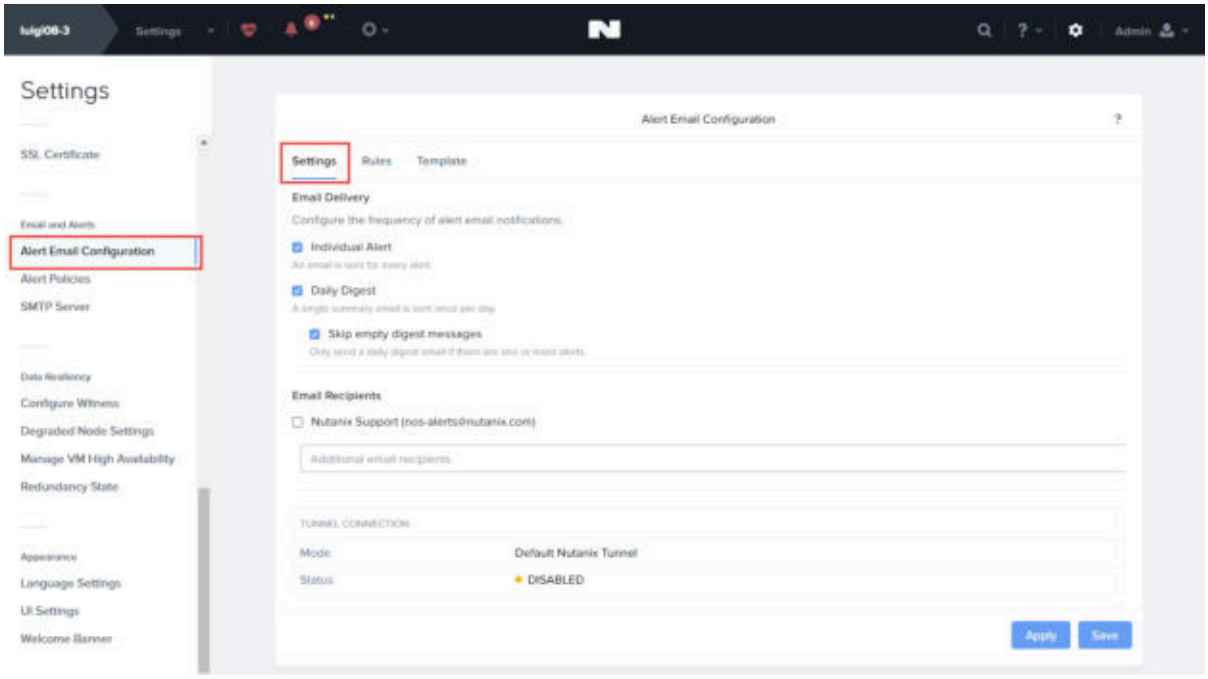

Figure 261: Alert Email Configuration: Settings Tab

3. To create a custom alert email rule, click the **Rules** tab and the **New Rule** button, and then configure the rule as follows:

   a. Specify the conditions for generating the alert:

   - **Alert Severity**: Select one or more of the severities from the pull-down list (**Critical**, **Warning**, **Info**, **All**).
   - **Impact Type**: Select one or more of the categories from the pull-down list (**Availability**, **Capacity**, **Configuration**, **Performance**, **System Indicator**, **All**).
   - **Cluster**: Select one or more of the clusters from the pull-down list (*cluster_name*, All).
   - **Alert Contains**: Enter a key phrase or word that should generate an email notification whenever the alert contains that phrase.

     For example, if you want to receive an email notification for when an alert contains the phrase "NIC down", then enter NIC down.

     In the following example, if you want to receive an email notification for A1029 alert, the **Alert Contains** field must contain phrases that effectively matches phrases from the *Name* and *Alert message* fields.

     Table 46: Example: Zeus Config Mismatch [6208] [A1029]

Name	Hypervisor IP Address Configuration
Description	Checks that hypervisor IP address is in sync with zeus configuration
Alert message	Hypervisor IP address on Controller VM *svm_ip_address* was updated from *zeus_ip_address* to *invalid_ip_address* without following the Nutanix IP Reconfiguration procedure.
Cause	The IP address configured in the cluster does not match the actual setting of the hypervisor.
Impact	The hypervisor host is unreachable from other Controller VMs in the cluster.

Resolution	Follow the IP address change procedure in the Nutanix documentation.

b. Specify who should receive the alert email by entering recipient email addresses as a comma separated list in the **Send Email To** field.

c. Click the **Save Rule** button to save the rule.

d. Repeat these steps to create more custom rules.

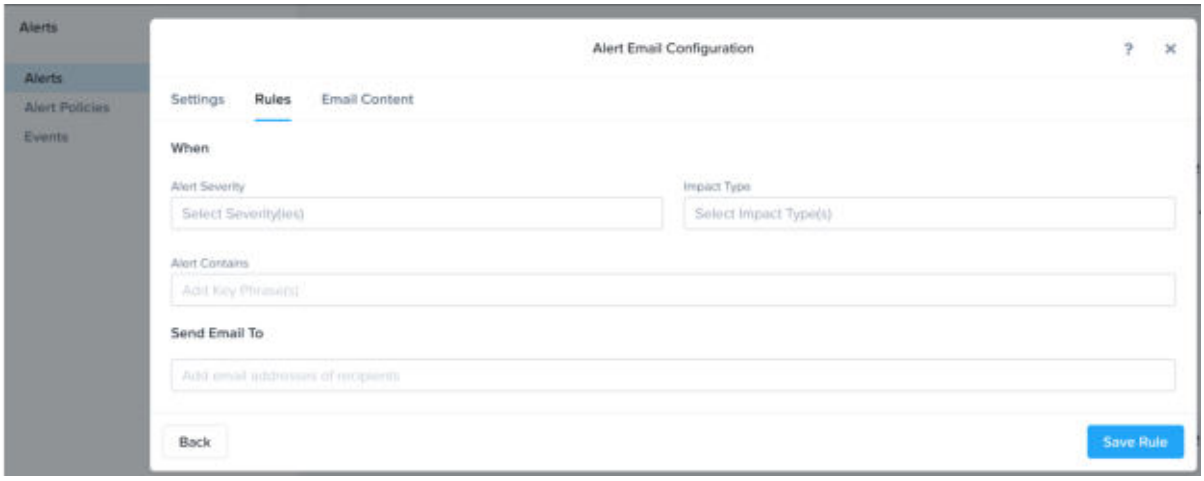

Figure 262: Alert Email Configuration Window: Rules tab

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).

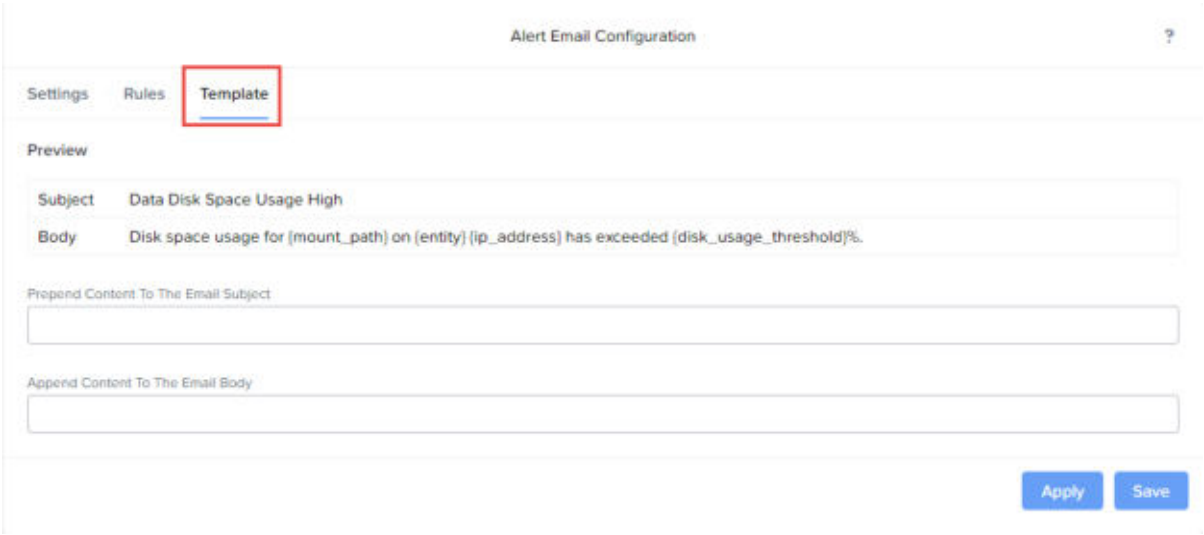

Figure 263: Alert Email Configuration: Settings Tab

Configuring Alert Policies

About this task

The system comes with default alert policies, but you can customize which events should trigger an alert. To customize the alert policies, do the following:

Procedure

1. Do one of the following:

   » Click the gear icon in the main menu and then select **Alert Policies** in the **Settings** page.

   » Go to the Alerts dashboard, click the **Configure** button, and select **Alert Policies** from the pull-down menu.

   The Alert Policies page appears. This page displays the list of events that can trigger an alert. Each line includes the ID number (Axxxx), title, impact type (system indicator, availability, configuration, and so on), entity type (cluster, host, protection domain, and so on), rule setting (default or overridden), enabled and auto-resolve status (enabled or disabled), and

the last time this alert policy was updated. Enter a string in the search box at the top or use the scroll bar on the right to search through the list.

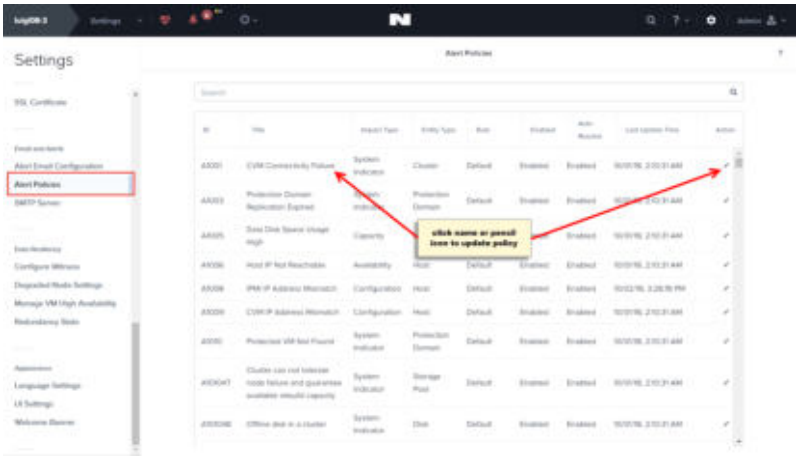

Figure 264: Alert Policies Page

2. To modify the alert policy for an event, click the title or pencil icon for that event to open the Update Policy window, and do the following in the indicated fields:

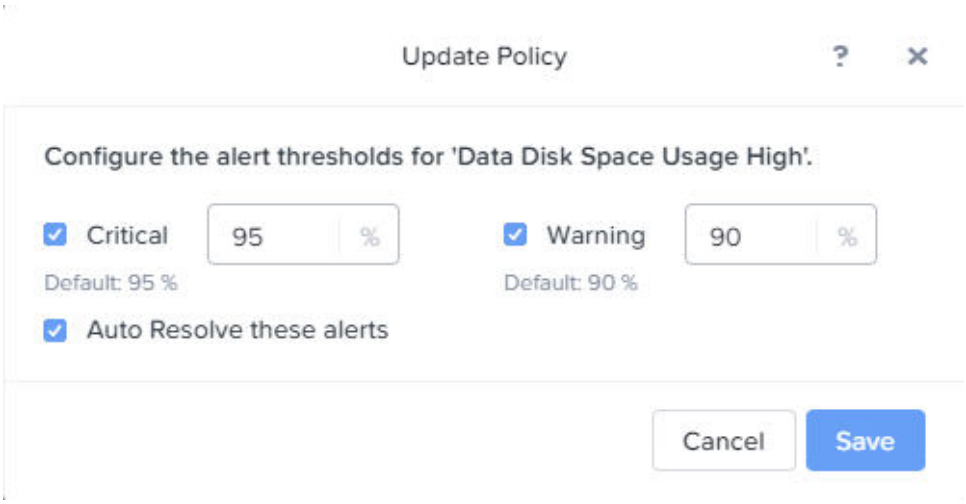

Figure 265: Update Policy Window

a. **Critical**, **Warning**, and/or **Info** severity fields: Uncheck (or check) the box next to the severity to disable (or re-enable) this as an alert.

   All the alerts are enabled by default (box checked). 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 90%) and a critical alert when a second threshold

is reached (95%). 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. Click the **Save** button to save the changes and return to the Alert Policies page.

3. Repeat the previous step for any additional alert policies that you want to update.

Alerts/Health checks

Availability Zone

Table 47: Unable to Retrieve the Availability Zone Endpoint [130173] [A130173]

Name	Unable to get Availability Zone Endpoint
Description	Unable to retrieve the endpoint information for the Availability Zone
Alert message	Unable to retrieve the endpoint information for the Availability Zone *availability_zone_physical_name*
Cause	Prism Central, Prism Element, or Xi - Availability Zone may be unreachable
Impact	VMs may not be fully protected
Resolution	Check if the Prism Central, Prism Element, or Xi - Availability Zone are reachable

Table 48: Unable to communicate with the Data Center Manager [130175] [A130175]

Name	Failure to communicate with DCM
Description	Unable to retrieve the endpoint information for the Availability Zone from the Data Center Manager
Alert message	Unable to retrieve the endpoint information for the Availability Zone *availability_zone_name* from the Data Center Manager
Cause	Data Center Manager may be unreachable or has provided an invalid response
Impact	VMs may not be fully protected
Resolution	Check if Data Center Mananger is reachable from the Xi - Availability Zone

Cluster

Table 49: High vDisk count in the cluster [1032] [A1182]

Name	vDisk Count Check
Description	Checks for high vDisk count in the cluster.
Alert message	Number of *vdisk_type* in the cluster is above the threshold (*vdisk_count*/*vdisk_threshold*)
Cause	Aggressive replication/snapshot schedules may generate large number of vDisks on the remote site
Impact	The cluster may become unstable due to increased load on core Nutanix services
Resolution	Immediately contact Nutanix support for assistance
KB Article	1624
Severity	Warning

Table 50: Cluster contains an unsupported mix of nodes with different storage combinations [1054] [A1191]

Name	All-flash Node Intermixed Check
Description	Checks that all-flash nodes do not coexist with non-all-flash nodes in a cluster
Alert message	*alert_msg*
Cause	All-flash nodes and non-all-flash nodes are put into the same cluster
Impact	The cluster may experience performance degradation or instability.
Resolution	Remove either all-flash nodes or non-all-flash nodes from the cluster
KB Article	3551
Severity	Warning

Table 51: Cluster In Read-Only Mode [1058] [A1195]

Name	Cluster In Read-Only Mode
Description	Single-node cluster is in read-only mode.
Alert message	Single node cluster with CVM *service_vm_external_ip* is running in read-only mode due to *reason*.
Cause	One of the metadata SSDs has failed.
Impact	Native backups are disabled.
Resolution	Replace the failed SSD. Refer to the Nutanix documentation for instructions.
KB Article	8156
Severity	Warning

Table 52: CVM Time Difference High [3025] [A1017]

Name	Time Drift Check

Description	Checks whether CVM times match.
Alert message	**Time differs between Controller VMs.**
Cause	Timedrift between CVMs is high. NTP connectivity may be lost.
Impact	The cluster may experience downtime.
Resolution	Check NTP/network configuration or connectivity.
Time Drift Threshold	Warning Threshold: 3 seconds
KB Article	1512
Severity	Warning

Table 53: Upgrade bundle available [3031] [A3031]

Name	Upgrade Bundle Available
Description	Bundle available for upgrade.
Alert message	**Bundle *bundle_name* is available for *upgrade_type* upgrade.**
Cause	Newer bundle available for upgrade
Impact	Cluster is running on older version.
Resolution	Upgrade the cluster with newer bundle. Refer to the Nutanix documentation for instructions.
Severity	Info

Table 54: Cluster services are down [3032] [A3032]

Name	Cluster Services Down Check
Description	Check that no services are down in the cluster.
Alert message	**Cluster Service: *service_name* is down on the Controller VM *ip_address*.**
Cause	One or more services in the cluster are down.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to serve I/O requests.
Resolution	[If the reason for the service(s) being down is unknown, run "ncc health_checks run_all", perform log collection covering the period of this alert timestamp, and provide the results to Nutanix Support via case attachment for further investigation. If due to planned maintenance or triage, please ensure the service is started again using "cluster start" as soon as possible. Report any reoccurring issues.]
KB Article	1518
Severity	Critical

Table 55: Alert Manager Service Check [3033]

Name	Alert Manager Service Check
Description	Check if Alert Manager service is available.
Cause	This check usually indicates that the Alert Manager service is not healthy, but there could be other causes.
Impact	Cluster issues may stay unnoticed.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.
KB Article	3377

Table 56: Cluster Service Restarting Frequently [3034] [A3034]

Name	Cluster Services Status
Description	Check if services have restarted recently across the cluster.
Alert message	There have been multiple service restarts of *services* across all Controller VM(s). Latest crash of these services have occurred at *timestamps* respectively.
Cause	This alert indicates that one or more services in the cluster were restarted.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.
KB Article	3378
Severity	Warning

Table 57: File system inconsistencies are detected. [3038] [A3038]

Name	EXT4-fs Error Check
Description	Captures EXT4-fs error messages
Alert message	EXT4 file system errors are detected on CVM: *svm_ip*: *alert_msg*
Cause	File system inconsistencies are present on the node.
Impact	Inability for the CVM to boot or for the upgrade pre-checks to run.
Resolution	Look for any problems in the file system. Review KB 8514.
KB Article	8514
Severity	Critical

Table 58: Baseline configuration changes [3039] [A3039]

Name	AIDE baseline configuration changes check

Description	Checks the log files of aide for changes to baseline configs if aide is enabled
Alert message	*message*
Cause	Baseline configurations changed
Impact	Aide execution may fail depending on the files that were changed
Resolution	Check changed files and revert their changes if needed
KB Article	10461
Severity	Warning

Table 59: Found marker files [3040] [A3040]

Name	Marker Files Check
Description	This check determines if there are Marker files present on any of the CVMs of a cluster
Alert message	Found marker files: *alert_msg*
Cause	Presence of Marker files detected in one of the CVM(s).
Impact	The presence of marker files may result in either CVM services getting impacted during a service restart or may cause data loss if multiple CVMs have marker files.
Resolution	Please remove the marker files from the CVM, Refer KB 10524 for more details.
KB Article	10524
Severity	Critical

Table 60: Host IP Not Reachable [3065] [A1006]

Name	Host IP Pingable
Description	Checks if the local CVM can ping all host IP addresses.
Alert message	Hypervisor *target_ip* is not reachable from Controller VM *source_ip* in the last 6 attempts.
Cause	The hypervisor host is down or there is a network connectivity issue.
Impact	If the host and CVM are down, cluster compute and storage capacity are reduced and VM data may be under-replicated.
Resolution	Ensure that the hypervisor host is running and that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	2484
Severity	Critical

Table 61: NIC Link Down [3067] [A1082]

Name	NIC Link Status
Description	Checks whether any NIC is down.

Alert message	Link on NIC *nic_name* of host *host_ip* is down. NIC description: *nic_description*.
Cause	NIC Link is physically disconnected from the upstream Switch.
Cause	Upstream TOR switch port connected to the CVM has been disabled
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	Ensure that the cable between the node and the switch is plugged in properly on both ends.
Resolution	Ensure that the TOR switch port connected to the CVM is enabled and in forwarding state.
KB Article	2480
Severity	Warning

Table 62: Remote Support Enabled [6206] [A1051]

Name	Remote Support Status
Description	Checks if remote support tunnel to Nutanix HQ is enabled on this cluster.
Alert message	Remote support tunnel to Nutanix HQ is enabled on this cluster.
Cause	Nutanix support staff are able to access the cluster to assist with any issue.
Impact	Nutanix support staff are able to access the cluster to assist with any issue.
Resolution	No action is necessary.
KB Article	2466
Severity	Info

Table 63: CVM IP Address Mismatch [6207] [A1009]

Name	CVM IP Address Configuration
Description	Checks that CVM IP address is in sync with zeus configuration
Alert message	IP address of Controller VM *zeus_ip_address* has been updated to *invalid_ip_address*. The Controller VM will not be part of the cluster once the change comes into effect, unless zeus configuration is updated.
Cause	A CVM has been assigned an IP address that does not match the cluster's configuration.
Impact	The misconfigured CVM may not be a functional member of the cluster, reducing storage capacity. VM data may be under-replicated.
Resolution	Follow the IP address change procedure in the Nutanix documentation.
KB Article	2465
Severity	Critical

Table 64: Zeus Config Mismatch [6208] [A1029]

Name	Hypervisor IP Address Configuration
Description	Checks if the hypervisor IP address is in sync with zeus configuration.
Alert message	Hypervisor IP address on Controller VM *svm_ip_address* was updated from *zeus_ip_address* to *invalid_ip_address* without following the Nutanix IP Reconfiguration procedure.
Cause	The IP address configured in the cluster does not match the actual setting of the hypervisor.
Impact	The hypervisor host is unreachable from other Controller VMs in the cluster.
Resolution	Follow the IP address change procedure in the Nutanix documentation.
KB Article	2465
Severity	Critical

Table 65: CVM Subnet Configuration [6212]

Name	CVM Subnet Configuration
Description	Checks that subnet is in sync with zeus configuration
Cause	The subnet configured in the cluster does not match the actual subnet
Impact	The Controller VM will not be part of the cluster, unless the configuration is updated
Resolution	Follow the IP address change procedure in the Nutanix documentation.
KB Article	2465

Table 66: CVM using default password [6213] [A6213]

Name	CVM using default password
Description	CVM is using default password
Alert message	Please update the default password on CVM for user(s): *users* to harden the security on the clusters. Nutanix highly recommends not using default passwords for clusters.
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration - Using default passwords could expose your system to unauthorized access.
Resolution	Ensure that all CVMs are using a custom password for improved security.
KB Article	6153
Severity	Critical

Table 67: Host using default password [6214] [A6214]

Name	Host using default password

Description	Host is using default password
Alert message	Please update the default password on host for user(s): *users* to harden the security on the clusters. Nutanix highly recommends not using default passwords for clusters.
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration - Using default passwords could expose your system to unauthorized access.
Resolution	Ensure that all hosts are using a custom password for improved security.
KB Article	6153
Severity	Critical

Table 68: Multiple Vcenters discovered [6415] [A6415]

Name	Multiple vCenter Servers Discovered
Description	Multiple vCenter servers discovered
Alert message	Discovered more than one vCenter servers : '*vcenter_ips*'
Cause	Nodes in the current Nutanix cluster are managed by multiple vCenters.
Impact	You will not be able to perform VM Management related operations.
Resolution	(A) Make sure all the nodes of this Nutanix cluster are moved to any one vCenter. (B) Contact Nutanix support.
Severity	Warning

Table 69: Common port group between ESXi hosts is absent [6416] [A6416]

Name	No common port group among ESXi hosts
Description	Common port group between ESXi hosts is not present
Alert message	'*message*'
Cause	ESXi hosts does not have any common port group.
Impact	You will not be able to perform network related operations.
Resolution	Ensure that all the nodes of this cluster have at least one common port group or all the nodes are in a single datacenter. If you are still not able to resolve the issue, contact Nutanix support.
Severity	Warning

Table 70: Number of datastores configured is higher than ESXiMaxConnPerIP. [6447] [A6447]

Name	Check ESXi stale worlds
Description	Checks if the number of datastores configured in the ESXI host is higher than ESXi MaxConnPerIP advanced setting.

Alert message	Number of datastores configured in the ESXi host is *no_of_datastore* and it is higher than ESXi MaxConnPerIP advanced setting value #MaxConnPerIP@. This may cause stale VM worlds leading to degraded IO performance on the host *host_ip*.
Cause	Number of datastores configured in the ESXi host is higher than ESXi MaxConnPerIP advanced setting. This may cause stale VM worlds leading to degraded IO performance on the host.
Impact	Degraded hypervisor IO performance
Resolution	Review KB-7833 to increase MaxConnPerIP parameters.
KB Article	7833
Severity	Warning

Table 71: ESXi SIOC status check [6448]

Name	ESXi SIOC status check
Description	Checks if the ESX Storage IO Controller (SIOC) feature is disabled
Cause	ESX SIOC feature may be enabled
Impact	ESX SIOC feature should be disabled on all Nutanix datastores
Resolution	ESX SIOC feature needs to be disabled
KB Article	3358

Table 72: vCenter Server Not Registered [6450] [A6450]

Name	vCenter Registration Check
Description	Check if vCenter server is registered.
Alert message	vCenter Server for this vSphere cluster is not registered in the Prism.
Cause	vCenter Sever is not registered in the Prism.
Impact	vSphere VM management operations on Prism will be disabled.
Impact	Some VMs might be missed from the snapshots created during upgrade.
Resolution	Register the vCenter server for this vSphere cluster in the Prism settings
KB Article	3815
Severity	Critical

Table 73: CVM CPU Resource settings [6451]

Name	CVM CPU Resource settings
Description	Checks CVM CPU Resource settings
Cause	Underprovisioned CVM CPU Resources
Impact	Underprovisioned CVM CPU Resources may lead to multiple CVMs being in a degraded state.

Resolution	Increase CVM CPU resource settings to default
KB Article	10287

Table 74: IPMI IP Address Mismatch [15013] [A1008]

Name	IPMI IP Address Configuration
Description	Checks if the IPMI IP address is in sync with Zeus configuration.
Alert message	IPMI IP address on Controller VM *svm_ip_address* was updated from *zeus_ip_address* to *invalid_ip_address* without following the Nutanix IP Reconfiguration procedure.
Cause	The IP address configured in the cluster does not match the actual setting of the IPMI interface.
Impact	The host is unreachable through the IPMI interface.
Resolution	Follow the IP address change procedure in the Nutanix documentation.
KB Article	2465
Severity	Critical

Table 75: IPMI sensor values check [15021]

Name	IPMI sensor values check
Description	Check that all IPMI sensor values are within threshold
Cause	IPMI sensors are not within threshold.
Impact	Cannot monitor hardware health.
Resolution	Review KB article 1524.
KB Article	1524

Table 76: IPMI SEL assertions check [15023]

Name	IPMI SEL assertions check
Description	Check for IPMI SEL assertions in the past 24 hours
Cause	IPMI events logged in the last 24 hours.
Impact	Host might have restarted.
Resolution	Refer to KB 1590.
KB Article	1590

Table 77: IPMI SEL Log Power Failure [15030] [A15030]

Name	IPMI SEL power failure check
Description	Check for IPMI SEL Power failure in the past 24 hours

Alert message	Power failure Over-Current-Protection engaged - host *ip_address*: *alert_msg* do not remove and add back power - Immediately Contact Nutanix Support.
Cause	The power supply has detected an Over-Current condition for this block and its protection mechanism has activated.
Impact	Block will automatically shut down as part of the Over-Current protection mechanism.
Resolution	Do not replace or reseat the power supply. Immediately call Nutanix technical support for assistance.
KB Article	1590
Severity	Critical

Table 78: IPMI using default password [15037] [A15037]

Name	IPMI using default password
Description	IPMI is using default password
Alert message	Please update the default password on IPMI devices for user(s): *users* to harden the security on the clusters. Nutanix highly recommends not using default passwords for clusters.
Cause	IPMI access credentials have not changed from the shipped default.
Impact	Insecure configuration - Using default passwords could expose your system to unauthorized access.
Resolution	Ensure that all IPMI devices are using a custom password for improved security.
KB Article	6153
Severity	Critical

Table 79: Power supplies of different types detected on a node [15038] [A15038]

Name	Mixed Power Supply check
Description	Check for mixed power supplies on chassis
Alert message	*per_block_msg*
Cause	Power supplies of different types are installed on a chassis
Impact	Mixed power supplies is not a tested combination and may result in unanticipated results.
Resolution	Replace the power supply to match PSUs defined for the NX platform.
KB Article	6282
Severity	Info

Table 80: Number of orphaned egroups is over the recommended threshold. [20013] [A20013]

Name	Garbage egroups check

Description	Check that the number of orphaned egroups are below the recommended threshold.
Alert message	Orphaned egroups are high
Cause	Number of orphaned groups is over the recommended threshold.
Impact	Space is underutilized.
Resolution	Review KB 1574.
KB Article	1574
Severity	Critical

Table 81: Snapshot chain height exceeds threshold [20014] [A20014]

Name	Snapshot chain height check
Description	Check that snapshot chain height is less than 25
Alert message	Snapshot chain height has exceeded the threshold limit.
Cause	Snapshot chain height exceeds threshold of 25.
Impact	Some VMs might experience increased latencies.
Resolution	Review KB 1732.
KB Article	1732
Severity	Warning

Table 82: Oplog Episode Count Check [20015] [A20015]

Name	Oplog episodes check
Description	Check that oplog episode count is within threshold
Alert message	Oplog episode count is high for some virtual disks for a sustained period of time.
Cause	Oplog episode count is high for one or more vDisks.
Impact	Oplog draining to extent store might get slower or stalled, impacting random writes latency.
Resolution	One or more VMs have an unusual burst of write IO, or there might be data inconsistency in oplog episode files. Review KB 1541.
KB Article	1541
Severity	Critical
Severity	Warning

Table 83: Automatic Deduplication disabled check [20016]

Name	Automatic Deduplication disabled check
Description	Check that deduplication is not automatically disabled.

Cause	Deduplication on the Storage Container is disabled by Stargate service.
Impact	Space savings may not be observed on the Storage Container.
Resolution	Review KB 1730.
KB Article	1730

Table 84: Compression disabled check [20017]

Name	Compression disabled check
Description	Check whether compressions is automatically disabled.
Cause	Metadata usage has exceeded the 65 percent default threshold.
Impact	Compression is disabled
Resolution	Contact Nutanix Support to investigate the cause and manually re-enable compression.
KB Article	2725

Table 85: List VMs with QoS attributes [20025]

Name	List VMs with QoS attributes
Description	Check if there are VMs with QoS attributes
Cause	User set QoS attributes on the VMs using APIs
Impact	Increase in VM latency also contributing to increase in cluster latency
Resolution	Not Applicable
KB Article	9952

Table 86: RDMA Enabled [20026] [A20026]

Name	RDMA Enabled Check
Description	Check if RDMA is enabled in specific AOS version
Alert message	RDMA is enabled on AOS versions lower or equal to 5.10.8, 5.11.2, 5.16 which can result in high network latency
Cause	RDMA is enabled on the cluster
Impact	High Network Latency
Resolution	Upgrade AOS to a newer STS or LTS version
KB Article	10350
Severity	Warning

Table 87: List Storage Containers being converted to AES [20027]

Name	List Storage Containers being converted to AES
Description	Check if there are any Storage Containers which are currently being converted to AES
Cause	Automatically enabled on non-AES Storage Containers which meet the eligibility criteria for conversion to AES
Impact	Minor performance dip during the conversion
Resolution	Not Applicable
KB Article	10668

Table 88: Discrepancy between usage is over the recommended threshold [20029] [A20029]

Name	Usage discrepancy check
Description	Check that the discrepancy between the total container utilization and total transformed usage are below the recommended threshold
Alert message	Usage discrepancy is high: *usage_discrepancy_placeholder* GB
Cause	Discrepancy between usage is over the recommended threshold
Impact	Storage usage increasing faster than normal
Resolution	Review KB 10459
KB Article	10459
Severity	Critical

Table 89: Cassandra service status check [21009]

Name	Cassandra service status check
Description	Check that all Cassandra nodes are online and and functioning normally
Cause	Cassandra service is not running in one or more hosts.
Impact	Cluster resilience might degrade.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.
KB Article	1547

Table 90: Cassandra service crashed [21010] [A21010]

Name	Cassandra service restarts check
Description	Check if Cassandra service is crashing continuously.
Alert message	Cassandra service has recently crashed. Contact Nutanix support for assistance.

Cause	Cassandra service has restarted frequently in the last 30 minutes.
Impact	Data resiliency is compromised.
Resolution	Contact Nutanix support.
KB Article	1592
Severity	Critical

Table 91: Cassandra service is running out of memory [21011] [A21011]

Name	Cassandra memory usage
Description	Check if the Cassandra service is running out of memory.
Alert message	Cassandra service is running out of memory. Contact Nutanix support for assistance.
Cause	Cassandra has crossed memory threshold more than 5 times in the last 3 minutes.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.
KB Article	1576
Severity	Critical

Table 92: Cassandra Keyspace/Column family check [21013]

Name	Cassandra Keyspace/Column family check
Description	Check that there are no undefined Keyspace or CF in Cassandra
Cause	The Cassandra keyspace/column families are no longer used in this AOS version.
Impact	Some metadata usage by deprecated Cassandra keyspace/column-family.
Resolution	For more information, see KB 4274. If cleanup is required, contact Support.
KB Article	4274

Table 93: Cassandra Waiting For Disk Replacement [21014] [A21014]

Name	Cassandra Waiting For Disk Replacement
Description	Metadata disk not replaced for Disk Replace Op
Alert message	Disk *disk_id* for node *service_vm_external_ip* not replaced. Node performing disk replace: *ip_address_doing_ring_change*
Cause	Disk not replaced for Disk Replace Op.
Impact	Disk Replacement operation will not proceed to completion.

Resolution	Ensure that a replacement disk is present on the node undergoing disk replace. If it is and the alert persists, contact Nutanix support.
Severity	Critical

Table 94: Multiple Cassandra nodes have similar tokens [21015] [A21015]

Name	Cassandra Similar Token check
Description	Check for similar tokens in Cassandra
Alert message	Multiple Cassandra nodes have similar tokens
Cause	Multiple Cassandra nodes have similar tokens
Impact	Cluster may not tolerate a Cassandra failure.
Resolution	Refer to KB 4301.
KB Article	4301
Severity	Critical

Table 95: Domain Fault Tolerance is low for Metadata [21016] [A21016]

Name	Cassandra Domain Fault Tolerance Warning
Description	Domain fault tolerance is low for metadata
Alert message	*domain* fault tolerance is reduced from desired fault tolerance(*desired_fault_tolerance*) to *current_fault_tolerance* for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Cassandra-domain-aware fixer has been triggered to re-arrange metadata and improve the domain fault tolerance. No action is required from your side.
Severity	Info

Table 96: Domain Fault Tolerance is low for Metadata [21017] [A21017]

Name	Cassandra Domain Fault Tolerance Error
Description	Domain fault tolerance is low for metadata
Alert message	*domain* fault tolerance is reduced from desired fault tolerance(*desired_fault_tolerance*) to *current_fault_tolerance* for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Please run Cassandra-domain-aware fixer operation using the nCLI command 'cluster cassandra-domain-aware-fixer-start'.

Severity	Warning

Table 97: Domain Fault Tolerance is low for Metadata [21018] [A21018]

Name	Cassandra Domain Fault Tolerance Critical
Description	Domain fault tolerance is low for metadata
Alert message	*domain* fault tolerance is reduced from desired fault tolerance(*desired_fault_tolerance*) to *current_fault_tolerance* for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Please add more nodes to reduce the skew in number of nodes in each domain. It is preferred to maintain almost equal number of nodes in all the domains to achieve best domain fault tolerance.
Severity	Warning

Table 98: Node detached from metadata ring [21019] [A21019]

Name	Metadata Drive Detached With Fixer Check
Description	Node detached from metadata ring
Alert message	Metadata drive on CVM *ip_address* is now detached from ring due to *reason*.
Cause	Node was removed from the metadata store as part of the ring fixer operation.
Impact	Cluster performance may be degraded till node gets added back.
Resolution	No action required. Node will get added back to the metadata store automatically.
Severity	Info

Table 99: One or more Cassandra nodes have invalid tokens [21020] [A21020]

Name	Cassandra Invalid Token check
Description	Check for short tokens in Cassandra
Alert message	Nodes *node_list* have shorter tokens than recommended
Cause	One or more Cassandra nodes have short tokens
Impact	Cluster may not work properly after some time.
Resolution	Refer to KB 6656.
KB Article	6656
Severity	Critical

Table 100: Health warnings detected in metadata service [21021] [A21021]

Name	Cassandra SSTable Health Warning check
Description	Check for cassandra SStables with health warnings
Alert message	Health warnings detected in metadata service
Cause	Cassandra SSTable health warning
Impact	Data resiliency is critical. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Refer to KB 7077.
KB Article	7077
Severity	Critical

Table 101: Metadata service restarting frequently due to long GC pauses. [21022] [A21022]

Name	Cassandra Restarting Due to Long GC Pauses
Description	Metadata service restarting frequently due to long GC pauses.
Alert message	Metadata service on node *ip_address* is restarting frequently due to long garbage collection pauses.
Cause	Metadata Service is undergoing long garbage collection pauses. This could possibly be due to cluster overload.
Impact	Cluster performance may be significantly degraded.
Resolution	Engage Nutanix Support to investigate why the Metadata service is significantly overloaded.
Severity	Warning

Table 102: SSD tier size is small on one or more nodes. [21023] [A21023]

Name	Cassandra SSD Size check
Description	Check if SSD disks are sized appropriately relative to HDD disk size.
Alert message	SSD on node(s) *node_list* is not sized for the expected metadata growth. Projected metadata growth is going beyond 50% of SSD capacity. Increase SSD disk capacity by at least *expected_ssd_size_list* on *node_list* respectively
Cause	SSD tier size is small to hold the metadata when cluster reaches maximum capacity
Impact	Performance may be degraded
Resolution	SSD tier size of specified nodes has to be increased. For more information, see KB-8842 or engage Support
KB Article	8842
Severity	Warning

Table 103: ESXi RAM disk root usage [101041]

Name	ESXi RAM disk root usage
Description	Check that RAM disk root directory usage below 85%
Cause	ESXi ramdisk root usage exceeds 85%.
Impact	Host might be inaccessible if ramdisk is full.
Resolution	Refer to KB 1846.
KB Article	1846

Table 104: Discoverable disks check [101042]

Name	Discoverable disks check
Description	Check if disks are discoverable
Cause	HyperV utilities in the host are unable to detect controllers and disks.
Impact	NA
Resolution	The host may need a reboot. Review KB 2712 for mitigation steps.
KB Article	2712

Table 105: Duplicate disk id check [101044]

Name	Duplicate disk id check
Description	Check for duplicate disk ids
Cause	Multiple disks with same disk id were found.
Impact	This can interfere with cluster functionality.
Resolution	Review KB 1876.
KB Article	1876

Table 106: Container available unreserved space check [101045]

Name	Container available unreserved space check
Description	Checks if available unreserved space on any Storage Container is below 10 percent of the maximum capacity.
Cause	Unreserved available space is below the threshold.
Impact	Virtual Machines might go into a stunned state if the available unreserved capacity drops to zero.
Resolution	Increase the available storage capacity on the cluster by either moving or removing virtual machines from the cluster or adding more nodes to the cluster.
KB Article	2047

Table 107: Cluster does not have enough free space to tolerate a node failure [101047] [A101047]

Name	High disk space usage
Description	Check if the storage pool has sufficient free space to maintain the desired replication factor in the event of a node failure, based on highest replication factor container requirements
Alert message	Cluster cannot tolerate *num_node_failure_limit* node failure. Current usage of *used_bytes* bytes exceeds max limit of *usable_bytes* bytes.
Cause	Inadequate free space on the cluster
Impact	Cluster can not tolerate single node(RF2) or two node(RF3) failure and guarantee available storage to provide data redundancy.
Resolution	Clean out unused data from guest VMs, unneeded reservation at container level, or VMs from the container to be below the threshold.
KB Article	1863
Severity	Critical

Table 108: Offline disk in a cluster [101048] [A101048]

Name	Disk online check
Description	Check for the offline disks
Alert message	Disk mounted at *mount_path* on Controller VM *cvm_ip* is marked offline.
Cause	A disk may have failed or was manually removed.
Impact	Cluster performance may be degraded.
Resolution	Please refer to KB 1536 for more information.
KB Article	1536
Severity	Warning

Table 109: Incomplete disk removal [101049]

Name	Incomplete disk removal
Description	Check for incomplete disk removal
Cause	Cluster Health detected an incomplete disk removal operation.
Impact	Data integrity might be compromised.
Resolution	If the disk removal has not been manually triggered contact Nutanix Support.
KB Article	1545

Table 110: SED protection consistency check [101050]

Name	SED protection consistency check

Description	Checks that all the drives in the system can authenticate to their internal security agents
Cause	(A) The encrypted disks may have tentative passwords. (B) The disks can not authenticate themselves with the configured passwords. (C) Protection is disabled but the disk has a password configured.
Impact	Automatic mounting of disks on boot may have issues or take longer than normal
Resolution	Contact Nutanix Support.
KB Article	2042

Table 111: Automatic disabling of Deduplication [101051]

Name	Automatic disabling of Deduplication
Description	Check that On disk deduplication is not automatically disabled
Cause	Deduplication was disabled due to high metadata size, CVM RAM or SSD requirements are not met.
Impact	Deduplication is disabled.
Resolution	Ensure the RAM and SSD requirement are met and enable deduplication from prism or CLI.
KB Article	1851

Table 112: Cluster In Override Mode [101057] [A101057]

Name	Cluster In Override Mode
Description	Single-node cluster is in override mode.
Alert message	Single node cluster with CVM *service_vm_external_ip* is running in override mode due to *reason*.
Cause	One of the metadata SSDs has failed.
Impact	Cluster is not disk fault tolerant.
Resolution	Replace the failed SSD. Refer to the Nutanix documentation for instructions.
KB Article	8132
Severity	Warning

Table 113: SED Disk Configuration Misconfigured [101080] [A101080]

Name	Disk configuration v/s intent consistency check
Description	Check if SED disk configuration is consistent with the intent
Alert message	Either the cluster had SED protection enabled and the disks did not have a password set OR the cluster did not have SED protection enabled and disks had password set OR found tentative passwords.
Cause	Inconsistency in disk's zeus configuration and the intent

Impact	Data unavailability requiring SRE intervention.
Resolution	This issue is usually seen when the KMS is switched off when the disks are still in protected state. Please check whether the KMS is on or off. If the issue persists, contact Nutanix Support.
KB Article	11322
Severity	Critical
Severity	Warning

Table 114: Reset Current Passwords for admin and nutanix SSH accounts on {cluster_name} [101081] [A101081]

Name	Password Strength Status
Description	High-strength password policy has been enabled, please reset the password.
Alert message	High-strength password was recently enabled, please reset current passwords for admin and nutanix SSH accounts on *cluster_name* - *cluster_uuid*.
Cause	High-strength password policy has been enabled.
Impact	High-strength password policy has been enabled but low-strength passwords may still be in use on the cluster.
Resolution	Please reset the current passwords. Follow the Security Guide's Hardening Instructions on Nutanix Support Portal for the password requirements.
Severity	Warning

Table 115: Host FQDN resolution [103069]

Name	Host FQDN resolution
Description	Check FQDN resolution of host IPs
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve "host FQDN to IP" or "IP to host FQDN".
Resolution	Check if hostname to IP and IP to hostname resolution is working.
KB Article	1709

Table 116: NTP server FQDN resolution [103070]

Name	NTP server FQDN resolution
Description	Check NTP Server FQDN resolution
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve NTP server FQDN.
Resolution	Check if ntp server name resolution is working.
KB Article	1709

Table 117: SMTP server FQDN resolution [103071]

Name	SMTP server FQDN resolution
Description	Check SMTP server FQDN resolution
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve SMTP server FQDN.
Resolution	Check if SMTP server name resolution is working.
KB Article	1709

Table 118: NSC(Nutanix Service Center) server FQDN resolution [103072]

Name	NSC(Nutanix Service Center) server FQDN resolution
Description	Check NSC server FQDN resolution
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve NSC(Nutanix Service Center) server FQDN.
Resolution	Check if cluster has Internet access and correct name server setting exists.
KB Article	1709

Table 119: CVM Not Uplinked to Active 10Gbps Link [103073] [A103073]

Name	CVM 10GBe uplink check
Description	Check that CVM uplinked to active 10Gbps link
Alert message	**CVM is not uplinked to active 10Gbps link. Please check your Network configurations**
Cause	CVM is not uplinked to active 10Gbps link.
Impact	Performance may degrade.
Resolution	Verify bonds do not contain a 1GbE link. Please refer to KB 1584.
KB Article	1584
Severity	Warning

Table 120: Storage routed to alternate CVM check [103074]

Name	Storage routed to alternate CVM check
Description	Check if storage is routed to an alternate CVM
Cause	Stargate traffic reroute is engaged on host of the cluster.
Impact	Due to stargate traffic reroute, the data resilience and performance may get impacted
Resolution	Check the following cases: 1. CVM is completely down. 2. Stargate is not running or CVM is partly down. 3. Stargate is crashing.
KB Article	1497

Table 121: NIC driver and firmware version check [103075]

Name	NIC driver and firmware version check
Description	Check if Quad port NIC driver version is above minimum recommended version
Cause	Quad-port nic card is not running minimum recommended driver or firmware version.
Impact	NIC is not on the latest driver / firmware version.
Resolution	Upgrade to the latest version of driver and firmware version for the Quad port NIC card.
KB Article	2715

Table 122: VM IDE bus check [103079]

Name	VM IDE bus check
Description	Check UVMs with IDE disk
Cause	Detected VMs running in the cluster on IDE bus.
Impact	VMs Performance degradation.
Resolution	Review KB 2998.
KB Article	2998

Table 123: CVM memory check [103080]

Name	CVM memory check
Description	Check CVM Memory
Cause	CVM memory is less than the recommended threshold with the current feature set.
Impact	Cluster services may crash and fail to restart on CVMs with less than recommended memory.
Resolution	Increase the memory of the Controller VM as per platform requirements listed in KB1513.
KB Article	1513

Table 124: Same timezone check [103085] [A103085]

Name	CVM same timezone check
Description	Check that all CVMs are in the same timezone
Alert message	*vm_type*s are not in the same timezone.
Cause	Some of the Controller VMs in the cluster are not in the same timezone.

Impact	It is recommended to have all the CVMs in the same timezone to avoid potential issues.
Resolution	Make sure all Controller VMs are in the same time zone. If they are not, then set the timezone using KB 1050
KB Article	2044
Severity	Warning

Table 125: Cluster NCC version check [103086]

Name	Cluster NCC version check
Description	Check that NCC version is consistent on the cluster
Cause	NCC version is different across the CVMs.
Impact	Some NCC checks can not be executed on all CVMs.
Resolution	Upgrade NCC to the desired version to keep it consistent across CVMs.
KB Article	3640

Table 126: Duplicate CVM IP address detected [103089] [A103089]

Name	Duplicate CVM IP check
Description	Check for duplicate CVM IP address
Alert message	**Duplicate CVM IP address detected. Associated IP address is** *cvm_ip*
Cause	One (or more) Nutanix cluster Controller VM external(eth0) interfaces conflicts with other hosts on the same network.
Impact	CVM unavailable from network
Resolution	Please refer to KB 3025.
KB Article	3025
Severity	Critical

Table 127: Duplicate Host IP address detected [103091] [A103091]

Name	Duplicate Hypervisor IP check
Description	Check for duplicate Hypervisor IP address
Alert message	**Duplicate Host IP address detected. Associated IP address is** *host_ip*
Cause	One or more of your Nutanix cluster hypervisor interfaces conflicts with other hosts on the same network.
Impact	Hypervisor host connectivity can become unstable or unavailable, leading to performance impact, redundancy concerns and and potential downtime.
Resolution	Reconfigure the conflicting hypervisor interfaces to have different IP addresses or VLAN settings

KB Article	5578
Severity	Critical

Table 128: Host same timezone check [103092]

Name	Host same timezone check
Description	Check that all hosts are in the same timezone
Cause	Not all hosts in the cluster are in the same timezone.
Impact	Ensure all hosts in the cluster are in the same timezone to help ensure all guest VMs share the same timezone and log messages are timestamped consistently.
Resolution	Ensure that all hosts are in the same timezone by following the instructions in KB 1050
KB Article	2044

Table 129: CVM NIC Link Down [103094] [A103094]

Name	CVM NIC Link Status
Description	Checks whether any CVM NIC is down.
Alert message	Link on NIC *nic_name* of cvm *svm_ip* is down.
Cause	The rdma0 NIC is disconnected from the switch, or the switch port is failing.
Impact	Cluster performance may be significantly degraded. When rdma0 port is down, the CVM will switch to TCP mode via host vSwitch, and the RDMA mode of entire cluster is turned off. Critical large storage io service will suffer in non-RDMA mode.
Resolution	Ensure that the rdma0 NIC is connected to the switch and that the switch port is functioning properly.
KB Article	2480
Severity	Warning

Table 130: Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM [103095] [A103095]

Name	RDMA network switch and NIC PFC config check
Description	Checks if network switch DCBX and NIC PFC are correctly configured and enabled.
Alert message	Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM *cvm_ip*: *alert_msg*
Cause	Network switch DCBX is disabled or corresponding NIC PFC bit is not enabled.
Impact	RDMA traffic between CVMs is not enabled, which degrades cluster performance.

Resolution	Enable DCBX on network switch and configure PFC bit on RDMA NICs.
KB Article	5925
Severity	Warning

Table 131: NDP Check [103100]

Name	NDP Check
Description	Checks if the NDP(Nutanix Discovery Protocol) works in the cvm external network
Cause	IPv6 may be disabled at the CVM, or an upstream network switch may be blocking IPv6 multicast or has IPv6 MLD snooping enabled.
Impact	Node discovery will fail
Resolution	NDP sends an IPv6 UDP payload to multicast address ff02::1 DstPort 13000. Confirm IPv6 is enabled in your network and ensure multicast propagation is enabled on listeners (Controller VMs in the cluster) and senders (any node to be added) to avoid issues when expanding cluster.
KB Article	5660

Table 132: DVS ESXi compatibility check [103102]

Name	DVS ESXi compatibility check
Description	Check if DVS is being used on an affected ESXi version - ESXi 6.0GA,U1,U2,U3 and 6.5GA.
Cause	ESXi on the cluster is not of the recommended version - ESXi 6.0GA,U1,U2,U3 and 6.5GA.
Impact	CVMs might be disconnected after reboot and risk a hung upgrade.
Resolution	Upgrade ESXi on the cluster to ESXi 6.5 U1 or later. Please refer to KB 8149. If issue persists contact Nutanix Support
KB Article	8149

Table 133: LSI Disk controller firmware status [106018]

Name	LSI Disk controller firmware status
Description	Check LSI Firmware Revision
Cause	LSI Disk Controller is running a superseded version of firmware.
Impact	LSI Disk Controller firmware issues can result in a node to hang and disk instability.
Resolution	Update the LSI Controller Firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.
KB Article	3035

Table 134: SATA Controller status is bad. [106020] [A106020]

Name	SATA controller
Description	Check SATA controller ports
Alert message	*controller_type controller_instance* on host *ip_address* has the following problems: *alert_msg*
Cause	The SATA controller is experiencing errors.
Impact	The system may experience issues while accessing physical storage within the node.
Resolution	Review KB 2312. If required Contact Nutanix support.
KB Article	2312
Severity	Critical
Severity	Warning

Table 135: Intel SSD S3610 on {ip_address} has configuration problems. [106021] [A106021]

Name	Intel SSD Configuration
Description	Check if node has supported configuration for Intel SSD S3610.
Alert message	Intel SSD S3610 on host *ip_address* has configuration problems: *alert_msg*.
Cause	Intel SSD S3610 is installed on an NX-1020 or NX-9040 node. Intel SSD S3610 is installed on a cluster whose AOS version is before 4.1.1.4. Intel SSD S3610 is not properly down-formatted.
Impact	Intel SSD S3610 drive performance and write endurance are degraded.
Resolution	Remove Intel SSD S3610 from the NX-1020 or NX-9040 nodes. Upgrade AOS to at least version 4.1.1.4. Contact Nutanix support to replace with a new drive. Refer to the Nutanix documentation for instructions.
KB Article	2309
Severity	Critical

Table 136: Hardware FRU information correctness [106025]

Name	Hardware FRU information correctness
Description	Check that FRU fields are programmed to correct value
Cause	Hardware information may be incorrectly configured.
Impact	No functional Impact.
Resolution	Contact Nutanix support to correctly configure the FRU fields.
KB Article	3022

Table 137: DIMMs of different manufacturers in one memory channel [106028] [A106028]

Name	DIMMs Interoperability Check

Description	DIMMs Interoperability Check
Alert message	Memory channels on host *ip_address* have DIMMs of multiple manufacturers: *err_msg*
Cause	DIMMs of different manufacturers are in the same memory channel.
Impact	Unsupported configuration.
Resolution	Ensure that all DIMMs in a memory channel are of the same manufacturer. Verify capacity, model, and manufacturer of DIMMs.
KB Article	3651
Severity	Warning

Table 138: DIMMs have invalid Part Number. [106038] [A106038]

Name	DIMM Part Number Check
Description	DIMMs Part Number Check
Alert message	DIMMs on host *ip_address* have invalid Part Number: *err_msg*
Cause	DIMMs have invalid Part Number.
Impact	Unsupported configuration.
Resolution	Replace DIMMs that have invalid Part Number. Ensure that all DIMMs have the valid Part Number M393A4K40BB2-CTD.
KB Article	3651
Severity	Warning

Table 139: DIMM Configuration is Wrong [106041] [A106041]

Name	DIMM Config Check
Description	DIMM Configuration Check
Alert message	*dimm_id* on host *ip_address* has the following problem: *alert_msg*
Cause	The DIMM configuration is unsupported. Either number of CPUs or the number of DIMMs, or capacity of the DIMMS are not supported.
Impact	Unsupported DIMM configuration. Performance may be degraded.
Resolution	Ensure that DIMM configuration is correct.
KB Article	3651
Severity	Warning

Table 140: Samsung 863 or 863a on {ip_address} has configuration problems. [106046] [A106046]

Name	Samsung 863(a) Configuration
Description	Check if node has supported configuration for Samsung 863 or 863a.

Alert message	Samsung 863 or 863a on host *ip_address* has configuration problems: *alert_msg*.
Cause	Samsung 863 or 863a is not properly down-formatted.
Impact	Samsung 863 or 863a drive performance and write endurance are degraded.
Resolution	Contact Nutanix support to replace with a new drive. Refer to the Nutanix documentation for instructions.
KB Article	2309
Severity	Critical

Table 141: Supported virtual hardware version check [106047]

Name	Supported virtual hardware version check
Description	Supported virtual hardware version check
Cause	One or more CVM is using unsupported virtual hardware version
Impact	Virtual hardware version lesser then 7 might cause performance issue
Resolution	Use the recommended minimum Virtual machine hardware version
KB Article	1719

Table 142: USB boot device missing on node. [106048] [A106048]

Name	Klas USB boot device connection check
Description	USB boot device connection check on Klas platform
Alert message	USB boot device is not detected on node *ip_address* with Klas platform. *alert_msg*
Cause	USB boot device is not present on Klas platform.
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	Replug USB boot device again into Klas node and restart the node. If restart fails, contact Klas for USB replacement.
KB Article	3195
Severity	Critical

Table 143: Secure boot feature status [106065] [A106065]

Name	Secure Boot Check
Description	Check if the secure boot configuration is consistent across all nodes in the cluster.
Alert message	Secure boot is not enabled on following hosts: *host_ips*
Cause	Secure boot is not consistently configured across all the nodes in the cluster.
Impact	Cluster is only partially protected by the host secure boot.

Resolution	Enable secure boot on all hosts in the cluster to ensure intended cluster security.
KB Article	8193
Severity	Info

Table 144: ESXi version compatibility [106415]

Name	ESXi version compatibility
Description	Check ESX version compatibility with AOS
Cause	ESXi on the cluster is not of the recommended version.
Impact	Identified ESXi version on the cluster is either incompatible or has issues with CBT/VADP APIs.
Resolution	Upgrade ESXi on the cluster. Please refer to KB 1729 for more information.
KB Article	1729

Table 145: ESXi Driver compatibility check [106418]

Name	ESXi Driver compatibility check
Description	Check the compatibility of current installed ixgbe/ixgben driver versions with the current ESX version.
Cause	Current installed Intel ixgbe/ixgben driver version is not recommended on current ESX version.
Impact	Nutanix has seen instability in Network with the older driver version on later ESX versions.
Resolution	Nutanix recommends to keep the driver on the latest qualified version based on Intel 82599 and x540 10Gigabit Ethernet Controllers.
KB Article	1702

Table 146: ESXi services check [106419]

Name	ESXi services check
Description	Check if ESXi services are up on the host
Cause	One of ssh, ntpd, shell, vpxa, lbtd, DCUI, sfcbd services are not running.
Impact	Cluster connectivity, cluster synchronization, adding host to cluster, accessing console may be impacted.
Resolution	Check following services on the ESX host: SSH, ntpd, vpxa, lbtd, DCUI, sfcbd services.
KB Article	2051

Table 147: ESXi ProductLocker Configuration [106420]

Name	ESXi ProductLocker Configuration
Description	Check ESXi ProductLocker settings
Cause	Host contains a ProductLocker location that is not properly configured
Impact	CVMs appear to be stunned or frozen periodically on an ESXi cluster, potentially causing instability of cluster services and VMs. This issue can affect multiple CVMs at once, causing complete cluster outage
Resolution	Manually configure the ProductLocker location.Review KB 10204.
KB Article	10204

Table 148: VAAI plugin is not {alert_info} [106421] [A106421]

Name	ESXi VAAI plugin installed
Description	Check if the latest version of VAAI Plugin is installed
Alert message	VAAI plugin on host *host_ip* is not *alert_info*. *version_info*
Cause	Incorrect ESXi VAAI plugin version for AOS version.
Impact	Suboptimal storage performance due to lack of supported VAAI primitives (VMware KB 1021976).
Resolution	Review KB 1868.
KB Article	1868
Severity	Info

Table 149: hostd Service Not Running [106427] [A106427]

Name	VMware hostd service check
Description	Check if hostd access is available
Alert message	hostd service is not running.
Cause	ESXi hostd service may not be running or is unresponsive.
Impact	If hostd service is not running, AOS upgrade will fail.
Resolution	Please refer to KB 1849 for more information. If unsure, please contact Nutanix Support.
KB Article	1849
Severity	Warning

Table 150: CVM Port Group Renamed [106428] [A106428]

Name	CVM port group renamed
Description	Check that CVM port group name has not changed

Alert message	CVM port group has been renamed on the ESXi hosts
Cause	CVM port group has been renamed on the ESXi hosts.
Impact	AOS upgrade may fail.
Resolution	Refer to KB 1904
KB Article	1904
Severity	Warning

Table 151: ESXi Scratch Configuration [106429]

Name	ESXi Scratch Configuration
Description	Check ESXi Scratch settings
Cause	Host contains a scratch location that is not properly configured or the locker.conf cannot be read.
Impact	If scratch space is not available, temporary data is stored on a RAM disk with limited space, causing issues and log files not being updated.
Resolution	Manually configure the scratch location. Please refer to KB 2046 for more information.
KB Article	2046

Table 152: ESXi configured VMK check [106430]

Name	ESXi configured VMK check
Description	Check that correct vmk is chosen
Cause	Host Management IP in esx.conf does not match the host ip in zeus config.
Impact	A cluster upgrade or node addition operation might fail
Impact	The cluster will choose the wrong ESXi interface as management server IP address when a node is added.
Resolution	Review KB 2049 and KB 2522
KB Article	2049

Table 153: VMKNICs subnets check [106442]

Name	VMKNICs subnets check
Description	Check if vmknics have different subnets
Cause	vmknics have ip address configured in the same subnet
Impact	vmknics in the same ip subnet on the same esxi host is unsupported.
Resolution	Correct the IP addressing in the network. Please refer to KB 2722 for more information.
KB Article	2722

Table 154: Invalid APD configuration [106445] [A106445]

Name	ESXi APD handling check
Description	Check if APD handling is as recommended.
Alert message	*alert_msg*
Cause	APD handling is not configured to the recommended value.
Impact	If APD is not enabled, VMs could potentially end up in a critical state during a failure event.
Resolution	Refer to KB 3328 for the steps to enable or disable APD handling.
KB Article	3328
Severity	Warning

Table 155: ESXi NFS heartbeat timeout check [106446]

Name	ESXi NFS heartbeat timeout check
Description	Check if NFS heartbeat timeout is set to recommended value.
Cause	NFS heartbeat parameter is NOT set to 30 seconds.
Impact	This parameter essentially decides the amount of time ESXI waits before aborting a heartbeat request. If set incorrectly, it may cause ESXI to perceive the datastore status incorrectly.
Resolution	Refer to KB 3328 for steps to disable APD handling.
KB Article	3328

Table 156: ESXi AHCI Driver version check [106447]

Name	ESXi AHCI Driver version check
Description	Checks if VMW AHCI driver has the required version and accordingly prompt to enable or disable the driver.
Cause	The VMW AHCI driver has satadom performance issues on ESXi 6.5, And has been fixed in ESXi 6.5U1. Disablement of this driver on ESXi 6.7 may break ESXi upgrades
Impact	Esxi 6.5 with AHCI driver version less than or equal to 1.0.0-34 installed.
Resolution	Disable the AHCI driver if driver version is below 1.0.0-34 and enable if it is above 1.0.0-34.
KB Article	4532

Table 157: ESXi Storage access from Host [106449]

Name	ESXi Storage access from Host
Description	Check if storage is accessible from the host

Cause	Storage is not properly configured on the host.
Impact	All storage I/O might go to a single node.
Resolution	Review KB 2263.
KB Article	2263

Table 158: Incorrect Kerberos Setup [106450] [A106450]

Name	Hyper-V Kerberos setup check
Description	Check if Kerberos is set up correctly
Alert message	Kerberos is not set up correctly.
Cause	Kerberos is not setup correctly.
Impact	Kerberos authentication is impacted.
Resolution	Check Kerberos setup. Please refer to KB 3030 for more information.
KB Article	3030
Severity	Warning

Table 159: Hyper-V VMQ Check [106451]

Name	Hyper-V VMQ Check
Description	Check if VMQ is disabled on host
Cause	VMQ is enabled on Hyper-V host.
Impact	Enabling VMQ on NICs with a faulty driver can cause VMs to fail during live migration.
Resolution	Nutanix recommends that you disable VMQ. To do so, find the VMQ status: Get-Netadaptervmq and for the component for which VMQ is enabled, disable VMQ. Disable-NetAdapterVmq -name <component>.
KB Article	1866

Table 160: Hyper-V Host pending reboots check [106452]

Name	Hyper-V Host pending reboots check
Description	Check that host has no pending reboots
Cause	Host has pending reboot.
Impact	Hyper-v Hosts pending reboots after Windows cluster aware or other updates/hyperv controller utilities fail to show any disks attached even though the disks are physically attached.
Resolution	Restart the Hyper-v Host.
KB Article	2713

Table 161: Orphaned VSS copies are detected. [106453] [A106453]

Name	Orphaned shadow copies
Description	Check if any orphan VSS copies are present.
Alert message	Orphaned VSS Copies Detected
Cause	Orphaned VSS copies are detected.
Impact	Performance degradation
Resolution	Please delete orphaned VSS copies. See KB 2798 for more information.
KB Article	2798
Severity	Warning

Table 162: Domain join failure [106454] [A106454]

Name	Domain join check
Description	Check if host is joined to a domain
Alert message	Node(s) is/are not joined to a domain.
Cause	Node(s) is/are not joined to a domain.
Impact	The host will not be manageable using Hyper-V Manager or SCVMM.
Resolution	Join the cluster and nodes to a domain. Please refer to KB 2231 for more information.
KB Article	2231
Severity	Critical

Table 163: Hyper-V services check [106455]

Name	Hyper-V services check
Description	Check if essential services are running on the host.
Cause	NutanixHAService/DiskMonitoerService/NutanixHostAgent not running properly.
Impact	Cluster cannot do disk monitoring and storage traffic re-routing if the Controller VM is unreachable or down.
Resolution	Verify status of Nutanix services. Review KB2242 for detailed information.
KB Article	2242

Table 164: Window update check [106456]

Name	Window update check
Description	Check if automatic Windows updates are disabled on the host
Cause	Windows update mode is set to automatic.

Impact	Hyper-V hosts can reboot in an uncontrolled manner because of automatic windows updates.
Resolution	Set the update mode to Manual. Refer to KB article 2253
KB Article	2253

Table 165: Failover cluster creation failure [106457] [A106457]

Name	MS Failover cluster check
Description	Check if the Hyper-V failover cluster is properly configured
Alert message	Hyper-V failover cluster is not configured with host disk drives.
Cause	Hyper-V failover cluster is not configured with host disk drives.
Impact	It can cause outage as it will prevent disks from being available to CVMs after CVM and Host reboot.
Resolution	Refer to KB 2282.
KB Article	2282
Severity	Critical

Table 166: Required Windows features check [106458]

Name	Required Windows features check
Description	Check if the required Windows features are installed and the banned Windows features are not installed
Cause	Required windows roles are not enabled on the host.
Impact	Recommended features are not enabled.
Resolution	Review KB 2244.
KB Article	2244

Table 167: Network is not properly configured on the host [106459] [A106459]

Name	Host networking check
Description	Check if networking is properly configured on the host
Alert message	Network is not properly configured on the host
Cause	The Nutanix-specific virtual network switches are not correctly configured.
Impact	Nutanix services will not come up in the Controller VM.
Resolution	Create the missing vmswitches/network adapters.
KB Article	2283
Severity	Warning

Table 168: Hyper-V host power scheme [106460]

Name	Hyper-V host power scheme
Description	Check if high performance profile is enabled on the host
Cause	Power scheme on the host is not configured to High Performance.
Impact	Hosts can experience increased latency and performance might be degraded.
Resolution	Review KB 2217.
KB Article	2217

Table 169: CVM configuration check [106462]

Name	CVM configuration check
Description	Check if the CVM is properly configured
Cause	A problem with the CVM configuration has been detected.
Impact	CVM services may not start after reboot.
Resolution	Review KB 2283 and 2219.
KB Article	2219

Table 170: Unable to access storage from the host. [106463] [A106463]

Name	HyperV Storage access from Host
Description	Check if storage is accessible from the host
Alert message	Unable to access storage from the host. Ensure no external storage is connected.
Cause	Storage is not properly configured on the host.
Impact	All storage I/O might go to a single node.
Resolution	Review KB 2263.
KB Article	2263
Severity	Warning

Table 171: WinRM service check [106464]

Name	WinRM service check
Description	Check if the winrm service on the host is listening on port 5985 by connecting to the host using WSMan.
Cause	Either password is not updated in the cluster configuration or WinRM service not running in one or more Hyper-V hosts.
Impact	The CVM cannot connect to the host to manage it if wsman connection is failing or WinRM service not running in one or more Hyper-V hosts.

Resolution	(A) Ensure that the WinRM service is running and listening on the port 5986. (B) Update the correct password in the cluster configuration if wsman connection is failing. (C) Review KB 2243 for more details.
KB Article	2243

Table 172: Crash dumps check [106465]

Name	Crash dumps check
Description	Check crash dumps on the host
Cause	Dedicated kernel crash dump is not enabled or kernel crash dump exists on the host.
Impact	Host kernel crash dump may not work correctly.
Resolution	Please refer to KB 2249 for more information.
KB Article	2249

Table 173: User VMs status [106468]

Name	User VMs status
Description	Check if the user virtual machines are properly configured
Cause	User VMs are in critical state.
Impact	User VMs in critical state can not be accessible.
Resolution	Review KB 2259.
KB Article	2259

Table 174: CVM Panics check [106469]

Name	CVM Panics check
Description	Check if any CVMs have panicked.
Cause	At least one CVM panicked recently.
Impact	Storage may have been unavailable while the CVM was restarting.
Resolution	Follow the instructions in KB 4379.
KB Article	4379

Table 175: Host Missing Critical Windows Updates [106470] [A106470]

Name	Critical Windows updates check
Description	Check if required windows updates are installed
Alert message	Host is missing critical Windows Updates.

Cause	Host is missing critical Windows Updates.
Impact	Host might experience problems regarding storage / live migration.
Resolution	Review KB 3011 for more information.
KB Article	3011
Severity	Warning

Table 176: HyperV VSS file limit check [106473]

Name	HyperV VSS file limit check
Description	Check containers file count within limit.
Cause	VSS can fail on containers with higher file count.
Impact	VSS may fail on the container.
Resolution	Reduce the number of VMs and associated files on the container below the file limit.
KB Article	4384

Table 177: Hyper-V Disk Pass-Through Check [106474]

Name	Hyper-V Disk Pass-Through Check
Description	Verify if the disks are pass-through to the Controller VM in Hyper-V 2016.
Cause	All the storage controllers are not pass-through to the Controller VM.
Impact	Misconfigured physical drives are not be accessible.
Resolution	Run Repair-CVM command to fix the Controller VM configuration.
KB Article	4907

Table 178: Active Directory DC(s) running on cluster [106476] [A106476]

Name	DC running on Nutanix Cluster Check
Description	Check if all active directory domain controllers are running on Nutanix cluster
Alert message	Active Directory DC(s) running on cluster which will prevent failover cluster from starting after Nutanix Cluster starts and stops
Cause	Active Directory Domain Contoller(s) configured on the UVMs in the cluster
Impact	Failover cluster will fail to start after the Nutanix cluster starts and stops(due to power failure for example), as it depends on availability of domain controllers which will result in User VM downtime
Resolution	There should be at least one domain controller running outside Nutanix cluster, that can be used by failover cluster in case locally running DC servers are failed
KB Article	5178

Severity	Warning

Table 179: DNS server(s) running on cluster [106477] [A106477]

Name	DNS on Cluster Check
Description	Check if all DNS servers are running on Nutanix cluster
Alert message	DNS server(s) running on cluster which will prevent failover cluster from starting after Nutanix Cluster starts and stops.
Cause	DNS servers configured on the UVMs in the cluster
Impact	Failover cluster will fail to start after the Nutanix cluster starts and stops(due to power failure for example), as it depends on availability of dns servers which will result in User VM downtime
Resolution	There should be at least one DNS server running outside Nutanix cluster, that can be used by failover cluster in case locally running DNS servers are failed
KB Article	5178
Severity	Warning

Table 180: This is a non supported configuration. [106478] [A106478]

Name	External Storage Connected on Host Check
Description	Check if external datastore(s) is connected to a host
Alert message	*alert_msg*
Cause	Non-supported configuration.
Impact	Non-supported configuration
Resolution	Unmount/Disconnect external datastores from host(s). For more details review KB 2263.
KB Article	2263
Severity	Warning

Table 181: Remote Site Time Sync [110001]

Name	Remote Site Time Sync
Description	Checks whether time drift between remote site is within limit.
Cause	Time drift between source and remote cluster is high.
Impact	Remote replication may fail to complete as expected.
Resolution	Use NTP to sync time
KB Article	1848

Table 182: Remote site connectivity not normal. [110002] [A110002]

Name	Remote Site Connectivity
Description	Check if connectivity to remote sites is normal.
Alert message	Connectivity to remote site *remote_name* is not normal.
Cause	Remote site may not have the local site configured as a remote.
Cause	Remote site might be stale (destroyed and re-created).
Cause	Cluster services at remote site might not be running.
Cause	Network connectivity to remote site might have been lost.
Cause	Firewall/VPN on source cluster, remote cluster, or both may prevent connection establishment.
Impact	Remote replication may fail to complete as expected.
Resolution	Configure the local site as remote on the remote site.
Resolution	Remove and add back the remote site on the local cluster.
Resolution	Ensure that cluster services at remote site are running.
Resolution	Check network connectivity to remote site.
Resolution	Check firewall/VPN on source cluster and remote cluster to allow them to connect to each other.
KB Article	1596
Severity	Warning

Table 183: Remote Site MTU settings [110005]

Name	Remote Site MTU settings
Description	Check MTU setting for connecting to remote cluster
Cause	MTU size mismatch on CVMs between remote sites.
Impact	Degraded replication performance may be observed if MTU size is not consistent between remote sites.
Resolution	Review KB 1949.
KB Article	1949

Table 184: Remote Site configuration [110006]

Name	Remote Site configuration
Description	Check remote site configuration on local and remote sites
Cause	Remote site is not properly configured.
Impact	Nutanix snapshot replication might fail to the remote site.
Resolution	Review KB 3335.
KB Article	3335

Table 185: Remote Site virtual external IP(VIP) [110007]

Name	Remote Site virtual external IP(VIP)
Description	Check if the remote site configuration on the remote sites have virtual IP address of this cluster
Cause	Remote cluster may not be using the virtual IP address (VIP) of the local cluster for the remote site connectivity.
Cause	Remote Cluster may not be using the virtual IP address of the DR network segmentation of the local cluster for the remote site connectivity.
Impact	Shutting down of the CVMs on the remote site may result in the replication problems.
Impact	Partial or incorrect DR network segmentation configuration.
Resolution	Add the virtual IP address (VIP) of the local cluster to the remote site configuration of the remote cluster. Add the virtual IP address (VIP) of the remote cluster to the remote site configuration of the local cluster.
Resolution	If the local cluster has DR network segmentation enabled, add the virtual IP address (VIP) of the DR network segmentation and the network segmented external subnet of the local cluster to the remote site configuration of the local cluster.
KB Article	3024

Table 186: AWS cloud instance not configured properly. [110008] [A110008]

Name	AWS Instance Check
Description	Verify that the AWS cloud instances have recommended configuration.
Alert message	**AWS cloud instance at remote *remote_name* not configured properly.**
Cause	AWS instance does not have at least two 190GB+ disks.
Impact	Curator scans may fail.
Resolution	Make sure that AWS instance has at least two 190GB+ disks.
KB Article	3814
Severity	Warning

Table 187: Old generation AWS instance configured. [110009] [A110009]

Name	AWS Instance Type Check
Description	Checks whether old generation AWS instance is configured for remote site.
Alert message	**Old generation AWS instance is configured for remote site *remote_name*. Follow KB 3861 to change the instance type to m3.2xlarge and cloud CVM configuration for better RPO/RTO.**
Cause	Present AWS instance has slow replication performance.
Impact	Present AWS instance has slow replication performance.
Resolution	Follow the KB 3861 to change the instance type to m3.2xlarge and cloud CVM configuration.

KB Article	4416
Severity	Info

Table 188: AOS version of cloud remote site is less than source cluster. [110010] [A110010]

Name	Cloud Remote Version Check
Description	Check AOS version for cloud remote site.
Alert message	AOS version of cloud remote site *remote_name*: (*remote_nos_version*) is less than that of source cluster(*src_nos_version*).
Cause	AOS version of cloud remote site is lower than that of source cluster.
Impact	Performance improvements made in new AOS version will not be available for cloud CVM.
Resolution	Upgrade cloud remote site.
KB Article	4246
Severity	Warning

Table 189: Cloud cluster does not have all recommended gflags set. [110012] [A110012]

Name	Cloud Gflags Check
Description	Check if cloud remote sites have recommended gflags set.
Alert message	Cloud cluster does not have all recommended gflags set.
Cause	Cluster may have been upgraded to a newer version.
Impact	Some of the improvements made in new release will not be enabled.
Resolution	Set recommended gflags. The cloud_connect_fix script can be run.
KB Article	4249
Severity	Critical

Table 190: Egroup count on cloud disk is higher than the recommended threshold. [110013] [A110013]

Name	Cloud Egroup Count check
Description	Check if egroup count on cloud remote is within threshold.
Alert message	Egroup count on cloud disk is higher than the recommended threshold.
Cause	Large amount of data may have been stored on cloud appliance.
Impact	Slow performance can be seen.
Resolution	Spread the protection domains that you want to replicate among multiple cloud remote sites.
KB Article	4247

Severity	Warning

Table 191: Incorrect Cluster Information in the Remote Site {remote_name} [110014] [A110014]

Name	Incorrect Cluster Information in Remote Site
Description	Check if the cluster information in remote site is correct.
Alert message	Replications to this remote site may fail because the cluster information is incorrect in the remote site *remote_name*.
Cause	Incorrect IP addresses may have been found in remote site configuration.
Cause	Remote cluster may have been destroyed and recreated.
Impact	Replications to this remote site may fail.
Resolution	Remove incorrect IP addresses and add correct IP addresses in the remote site configuration.
Resolution	Delete the remote site and add it again. You need to configure the schedules again.
KB Article	4505
Severity	Warning

Table 192: Metadata disk usage is higher than the specified limit. [110015] [A110015]

Name	Cloud Metadata Disk Usage Check
Description	Check if the metadata disk usage on cloud Controller VM is within the specified limit.
Alert message	Metadata disk usage is higher than the specified limit: *limit*.
Cause	Large amount of data may have been stored on cloud appliance.
Impact	Slow performance can be seen.
Resolution	(A) Spread the protection domains that you want to replicate among multiple cloud remote sites. (B) Follow the KB 4302 to add a new metadata disk to existing cloud Controller VM.
KB Article	4504
Severity	Warning

Table 193: Cloud disk usage is nearing the threshold. [110016] [A110016]

Name	Cloud Disk Usage Check
Description	Check if the cloud disk usage is within the threshold limit.
Alert message	Cloud disk usage: *usage_bytes* TB, is nearing the threshold: *threshold_bytes* TB. Cloud remote site will be turned to read-only once the threshold is crossed.
Cause	Large amount of data may have been stored on the cloud appliance.

Impact	Slow performance can be seen.
Resolution	(A) Spread the protection domains that you want to replicate among multiple cloud remote sites. (B) Reduce retention of snapshots on the cloud remote site.
KB Article	4708
Severity	Warning

Table 194: Azure cloud Controller VM has smaller disks. [110017] [A110017]

Name	Azure Disk Size Check
Description	Verify that the Azure cloud Controller VM has recommended configuration.
Alert message	**Azure cloud Controller VM has smaller than the two 256GB+ recommended disks. Contact Nutanix support for further assistance.**
Cause	Azure cloud Controller VM does not have minimum of two 256GB+ disks.
Impact	Slow performance can be seen.
Impact	Curator scans may fail.
Resolution	Contact Nutanix support so that the Azure cloud Controller VM has at least two 256GB+ disks.
KB Article	4763
Severity	Warning

Table 195: Remote site in same vCenter datacenter. [110018] [A110018]

Name	Remote Datacenter Check
Description	Check if any remote site is added under the same datacenter in vCenter.
Alert message	**Hosts of current cluster and Async DR remote site '*remote_name*' are registered under the same datacenter in vCenter.**
Cause	Source and remote site of an Async DR protection domain are under the same datacenter in the vCenter.
Impact	Migration to this remote site may fail.
Resolution	(A) Register source and remote sites with different datacenter in vCenter. (B) If source and remote site are in same vCenter datacenter, make sure that source and remote container names are different.
KB Article	3296
Severity	Warning

Table 196: Remote site is Unreachable [110020] [A110020]

Name	DR Service Reachability Check
Description	Check if the DR services are reachable.

Alert message	Remote site is unreachable because *message*.
Cause	Connectivity issues between the two sites.
Cause	Firewall rule to whitelist the Xi Leap subnet is not applied to the upstream network.
Impact	DR workflows will not work.
Resolution	Resolve the connectivity with the remote site. If the issue persists, contact Nutanix Support.
Resolution	Whitelist the IP address of one or more nodes belonging to the Xi Leap subnet in the firewall.
KB Article	6642
Severity	Warning

Table 197: Cloud Subnet not found or deleted [110023] [A110023]

Name	Cloud Subnet Not Found
Description	Cloud subnet not found or deleted
Alert message	Cloud Subnet for CIDR *cidr* not found: *message*.
Cause	Cloud subnet might be deleted while it is in use by Nutanix Cluster
Impact	Networking operations may be impacted on the affected node(s) in the cluster
Resolution	Create Cloud Subnet
KB Article	10013
Severity	Critical

Table 198: Default UVM Security Group Not Found [110024] [A110024]

Name	Default UVM Security Group Not Found
Description	Cannot find default AWS security group for all UVMs
Alert message	Cannot find default AWS UVMs Security Group for all UVMs: *message*
Cause	AWS Default UVM Security Group might be deleted while it is in use by Nutanix Cluster
Impact	Networking operations may be impacted on the affected node(s) in the cluster
Resolution	Ensure AWS security group with tags as "Key: nutanix:clusters:cluster-uuid Value: <cluster uuid> Key: nutanix:clusters:networks Value: "all-uvms"" exists. If issue still persists, please contact Nutanix Support
KB Article	10019
Severity	Critical

Table 199: Active Backup Schedule Check [110200]

Name	Active Backup Schedule Check
Description	Checks if a backup schedule exists for protection domain protecting some entities.
Cause	No backup schedule exists for protection domain protecting some entities.
Cause	Backup schedules exist for protection domain not protecting any entity.
Cause	If schedule was intentionally deleted, an expired snapshot(s) could still consume storage space.
Impact	No backups will be made.
Impact	Any expired snapshots (if existing), would still consume storage space.
Resolution	Delete backup schedule(s) for protection domain not protecting any entity. Delete the protection domain as well once the schedule is deleted.
Resolution	Configure a backup schedule for protection domain protecting some entities.
Resolution	If the schedule(s) was intentionally deleted to reclaim storage space, delete this protection domain also to reclaim the space taken up by snapshots.
KB Article	1910

Table 200: Storage Container(s) is/are not mounted on all nodes. [110201] [A110201]

Name	Storage Container Mount Configuration
Description	Verify that the Storage Container(s) is/are mounted on all nodes.
Alert message	Storage container(s) *unmounted_containers* for protection domain *pd_name* is/are not mounted on all nodes.
Cause	Storage Container(s) is/are not mounted on all nodes.
Impact	Registration of VM(s) may fail during restoration from backup.
Resolution	Mount the Storage Container(s) on all nodes.
KB Article	1888
Severity	Warning

Table 201: Local Backup Configuration [110204]

Name	Local Backup Configuration
Description	Verify if recovery is possible from a local backup.
Cause	Backup schedule is not configured properly.
Cause	Backup is not happening as expected.
Impact	Recovery from local backup may not be possible.
Resolution	Verify backup schedule configuration.
Resolution	Verify if backups are happening without error.
KB Article	1879

Table 202: Cluster Certificate Expiring [110215] [A1115]

Name	SED CA Certificate
Description	Check if CA certificates are about to expire.
Alert message	Cluster CA certificate *certificate_name* expires on *expiration_date*.
Cause	Certificates have defined expiration dates.
Impact	If the cluster's certificate expires, it will be unable to verify trust of server certificates.
Resolution	Get a new signed certificate.
KB Article	2493
Severity	Critical
Severity	Warning
Severity	Info

Table 203: Cloud remote check [110220]

Name	Cloud remote check
Description	Check backup and on-wire compression settings for cloud remote sites
Cause	The remote site for cloud site is not configured with recommended settings.
Impact	Not having recommended setting may lead to slow replication speeds.
Resolution	Contact Nutanix support.
KB Article	4245

Table 204: Cloud cluster does not have recommended configuration locally [110221] [A110221]

Name	Cloud local check
Description	Check that cloud cluster has recommended configuration locally.
Alert message	Cloud cluster does not have recommended configuration locally. Verify that cloud cluster has recommended configuration locally and then proceed.
Cause	(A) AWS cluster does not have exactly two storage pools (backup-sp and DoNotUse-sp). (B) Storage Pool backup-sp does not have exactly one disk with CLOUD storage tier. (C) Some remote sites configured on cloud do not have onwire compression enabled.
Impact	Snapshots may get replicated to non cloud disks leading to slow performance and other issues.
Impact	Disabling on wire compression may lead to slow replication speeds.
Resolution	Contact Nutanix support.
KB Article	4248
Severity	Critical

Table 205: Standby Protection Domains not in sync with primary [110222] [A110222]

Name	Metro Availability Secondary PD sync check
Description	Check if standby Metro Protection Domain is in sync with Active site
Alert message	Standby Metro Protection domain is no longer in sync with Active Protection domain on primary.
Cause	Link between primary and secondary is not healthy.
Impact	Standby protection domain is no more in sync.
Resolution	Ensure Metro is configured as per guidelines in the Prism Web Console guide on Nutanix Portal
KB Article	3032
Severity	Critical

Table 206: Checkpoint snapshot on Metro configured Protection Domain [110224]

Name	Checkpoint snapshot on Metro configured Protection Domain
Description	Check that last auto checkpoint snapshot was taken in the last 6 hours
Cause	Automatic checkpoint snapshot might not be created or created with delay.
Impact	Time for data replication and resync will be longer when Metro Availability is re-enabled.
Resolution	Create a schedule manually for checkpoint snapshot.
KB Article	3031

Table 207: ESX VM Virtual Hardware Version Compatible [110227]

Name	ESX VM Virtual Hardware Version Compatible
Description	Check Virtual hardware version compatibility of protected VMs
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote site.
Impact	VMs will not be registered on the remote site hosts upon PD failover.
Resolution	Upgrade the hypervisor version on the remote site to support virtual hardware version for the VM.
KB Article	2037

Table 208: Duplicate VM names [110228]

Name	Duplicate VM names
Description	Check for duplicate VM names
Cause	Duplicate named VMs are found.
Impact	Duplicate named VMs will not be protected by protection domains.

Resolution	Review KB 1907.
KB Article	1907

Table 209: Protection Domains File share [110229]

Name	Protection Domains File share
Description	Check for files/VMs protected by multiple protection domains
Cause	Some files/VMs are being protected by multiple protection domains.
Impact	Migrating such a PD would remove the VMs from the inventory.
Resolution	Review KB 2034.
KB Article	2034

Table 210: Snapshot file location check [110233]

Name	Snapshot file location check
Description	Check snapshot file location
Cause	VM protected by Nutanix Protection Domain has a trailing backslash in the snapshot file location.
Impact	Nutanix snapshots might be skipped for these VMs if they are protected by Nutanix Protection Domains.
Resolution	Review 2069.
KB Article	4382

Table 211: Detected VMs with non local data [110237]

Name	Detected VMs with non local data
Description	Metro Availability VMs are accessing data from local cluster.
Cause	Detected VMs polling data from remote location.
Impact	Performance degradation
Resolution	Review KB 2093.
KB Article	2093

Table 212: Files in a stretched VMs should be in the same Storage Container [110241]

Name	Files in a stretched VMs should be in the same Storage Container
Description	Check that VMs are consistently stretched
Cause	VM that is protected by Metro protection domain has some of its files in a different Storage Container from the Metro protected Storage Container.

Cause	In case of Hyper-V hosts, the VMs on stretch container may not be using Metro name.
Impact	VMs might not be properly recovered on the remote site.
Resolution	If the VM resides on a Metro enabled Storage Container, move all its associated files to the same Storage Container.
Resolution	Use the Metro name for all the VM components to avail the benefits of Metro features.
KB Article	2094

Table 213: Vstore VM Files Consistency Check [110253]

Name	Vstore VM Files Consistency Check
Description	Check that VMs in a Vstore are consistently protected
Cause	VM that is protected by Vstore protection domain has some of its files in a different Storage Container from the Vstore protected Storage Container.
Impact	VMs might not be properly recovered on the remote site.
Resolution	If the VM resides on a Vstore protection enabled Storage Container, move all its associated files to the same Storage Container.
KB Article	3819

Table 214: Internal Consistency Groups Check [110254]

Name	Internal Consistency Groups Check
Description	Check if a PD contains internal consistency groups
Cause	PDs may have some leftover internal consistency groups because metro availability enablement was aborted.
Impact	Snapshots of a PD in this state cannot be used to perform a restore operation.
Resolution	Review KB.
KB Article	4415

Table 215: Linked Clones in high frequency snapshot schedule [110256]

Name	Linked Clones in high frequency snapshot schedule
Description	Verify if the protection domain contains any linked clones that are configured with a high frequency snapshot schedule.
Cause	The protection domain might have a high frequency snapshot schedule configured with linked clones.
Impact	Snapshots of the protection domain in this state might lose data.
Resolution	Reduce the snapshot frequency or move the linked clones out of the protection domain.

KB Article	4906

Table 216: SSD Snapshot reserve space check [110257]

Name	SSD Snapshot reserve space check
Description	Check if the snapshot reserve space on SSD is filling up.
Cause	Snapshots' retention policy is high.
Cause	Available SSD space is low.
Cause	Replication link is too slow.
Impact	High frequency snapshotting cannot be sustained for some/all Protection Domains
Resolution	Reduce the number of entities that are configured for high-frequency snapshots.System will automatically try to reclaim SSD space by falling back to a lower snapshot frequency.
Resolution	Increase the SSD capacity in the cluster.
Resolution	Check the network throughput and bandwidth to ensure no bottlenecks are causing slower replications.
KB Article	4904

Table 217: Aggressive break replication timeout. [110258] [A110258]

Name	Metro aggressive break replication timeout check
Description	Check if break replication timeout of metro protection domain is more than the recommended limit.
Alert message	Break replication timeout of metro protection domain is below the recommended minimum of 15 seconds.
Cause	Break replication timeout of a metro protection domain is below recommended limit.
Cause	Metro availability is disabled.
Impact	Performance degradation.
Resolution	Set break replication timeout of the affected protection domain to recommended value.
Resolution	Enable metro availability.
KB Article	4951
Severity	Warning

Table 218: Incomplete Metro setup. [110261] [A110261]

Name	Metro HyperV 2016 metro cluster name setup check
Description	Metro setup for HyperV 2016 is incomplete.

Alert message	Metro setup for HyperV 2016 is incomplete.
Cause	Metro cluster name is not setup.
Cause	Different Metro cluster name on two clusters.
Cause	AD or DNS entry is not created for Metro cluster name.
Cause	Password on two clusters may not match.
Impact	Unsupported Metro feature.
Resolution	Please use setup_hyperv.py to add Metro cluster name.
Resolution	Use a common Metro cluster name across clusters to setup Metro.
Resolution	Please use setup_hyperv.py script to setup Metro.
Resolution	Please provide correct password to setup_hyperv.py while setting up Metro.
KB Article	5159
Severity	Critical
Severity	Warning

Table 219: NearSync Replication of Protection Domain has not progressed. [110264] [A110264]

Name	Nearsync Replication Stuck
Description	Nearsync replication of the protection domain has not progressed.
Alert message	NearSync replication of the protection domain *protection_domain_name* to the remote site *remote_name* for the snapshot *snapshot_id* has not made any progress in the last *time_string*.
Cause	Network outage, saturated network or unavailability of services like Stargate, Cerebro on the remote site. Snapshots of the protection domain will not be available on the remote site.
Cause	SSD snapshot reserve on the remote site does not have enough free disk space.
Impact	Protection domain may transition out of NearSync.
Impact	Snapshots of the protection domain will not be available on the remote site.
Resolution	Check the network connection and status of the remote site.
Resolution	Ensure that the remote site has enough disk space available in the SSD snapshot reserve.
KB Article	10249
Severity	Warning

Table 220: Latest snapshot of protection domain is missing entities [110265] [A110265]

Name	Latest snapshot missing entities check
Description	Check that all the VMs/VGs/Files in the protection domain at the time of snapshot creation are protected by the snapshot.
Alert message	Latest snapshot of protection domain *pd_name* is missing entities *missing_entities*

Cause	Snapshot doesn't include all the VMs/VGs/Files protected by that protection domain at the time of snapshot creation.
Impact	Some VMs/VGs/Files are not protected.
Resolution	If missing entities do not exist on the cluster, remove them from the protection domain else if entities are valid, create a out of band snapshot of the protection domain. If the issue still persists, contact Nutanix Support
KB Article	6258
Severity	Warning

Table 221: Maximum Entities in a Consistency Group Check [110266]

Name	Maximum Entities in a Consistency Group Check
Description	Check if each CG in a PD does not contain more than 20 entities
Cause	Various
Impact	Snapshots of a PD in this state cannot be used to perform a restore operation.
Resolution	Review KB.
KB Article	6247

Table 222: Invalid break replication timeout. [110269] [A110269]

Name	Metro invalid break replication timeout check
Description	Check if the break replication timeout of metro protection domain is zero or invalid.
Alert message	Break replication timeout of metro protection domain is invalid, please refer the KB to resolve this issue.
Cause	Break replication timeout of the metro protection domain is invalid state.
Impact	Metro relationship
Resolution	Set break replication timeout of the affected protection domain to the recommended value.
KB Article	7424
Severity	Warning

Table 223: Inconsistent file groups in the System [110271] [A110271]

Name	Inconsistent file groups check
Description	Inconsistent file groups possibly consuming unwarranted resources
Alert message	Possible unwarranted resource consumption because of inconsistent file groups
Cause	Possibly unsupported Metro transition.
Impact	Storage scarcity

Resolution	Please break and re-enable Metro relationship. Contact Nutanix support if that does not work.
KB Article	8607
Severity	Warning

Table 224: Stale App Consistent Snapshot Metadata Chunks Check [110274]

Name	Stale App Consistent Snapshot Metadata Chunks Check
Description	Stale app consistent snapshot metadata possibly consuming unwarranted resources
Cause	Unable to garbage collect app consistent snapshot metadata chunk entities probably due to Insights service unavailability.
Impact	IDF Higher Space Utilization
Resolution	Contact Nutanix Support to clean up stale entries. Refer KB 11244.
KB Article	11244

Table 225: Validation failed for open ports for required services for Synchronous Replication [110454] [A110454]

Name	Ports Open Check
Description	Checks if the ports for the services Acropolis(2030), Anduril(2036), Cerebro(2020), Ergon(2090), NGT(2073), Stargate(2009) which are required for Synchronous Replication are open
Alert message	*error_msg*
Cause	One or more services required for Synchronous Replication cannot be reached
Impact	Synchronous Replication will not work.
Resolution	Check firewall rules on the CVM and in your environment and reach out to Nutanix Support for any service related issues
KB Article	7365
Severity	Warning

Table 226: VM config updates disabled [110455]

Name	VM config updates disabled
Description	Checks for VMs that have configuration updates disabled
Cause	Update operations on a VM might fail after a failed migration
Impact	VM updates are blocked for the affected VMs
Resolution	Depending on the failure, a manual VM cleanup may be needed. Please refer to KB for detailed guidance
KB Article	9902

Table 227: E-mail alerts check [111000]

Name	E-mail alerts check
Description	Check email alerts
Cause	E-mail alerts disabled
Impact	Alerts will not be sent through email, asup cases might not be created.
Resolution	Review SMTP configuration. Review KB 1586.
KB Article	1586

Table 228: E-mail alerts contacts configuration [111001]

Name	E-mail alerts contacts configuration
Description	Check email contacts
Cause	E-mail alerts do not have contacts configured.
Impact	Alerts are enabled, but no contacts are configured.
Resolution	Configure emails contacts where alerts should be sent.
KB Article	1586

Table 229: Cassandra nodes up [111002]

Name	Cassandra nodes up
Description	Check that all Cassandra nodes are running
Cause	One or more Cassandra nodes are down which may be due to several issues including network issue, disk related issue, CPU or resource crunch.
Impact	Meta data load is inconsistent.
Resolution	Cassandra will repair itself. If other Cassandra checks are failing, contact Nutanix support.
KB Article	1369

Table 230: Cassandra tokens consistent [111003]

Name	Cassandra tokens consistent
Description	Check that all Cassandra nodes are consistent
Cause	The tokens may be inconsistent because of manual intervention or because of bad configuration.
Impact	Meta data load is inconsistent.
Resolution	Follow instructions described in KB 1369.
KB Article	1369

Table 231: Cassandra metadata balanced across CVMs [111004]

Name	Cassandra metadata balanced across CVMs
Description	Check that Cassandra metadata balanced across CVMs
Cause	Metadata is not distributed evenly among nodes in the cluster.
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	Correct imbalance by using the nCLI command 'cluster cassandra-token-range-skew-fixer-start'.
KB Article	1342

Table 232: Zookeeper fault tolerance check [111007]

Name	Zookeeper fault tolerance check
Description	Checks if Zookeeper services have the desired fault tolerance level
Cause	Zookeeper and/or Zookeeper monitor services are down on one or more nodes
Impact	The cluster may not be able to continue operating in the case of a node or block failure.
Resolution	Contact Nutanix support
KB Article	1587

Table 233: Time since last Curator scan is beyond threshold [111008] [A111008]

Name	Curator scan time elapsed check
Description	Check if time since last Curator Scan is beyond threshold.
Alert message	Time since last Curator scan is beyond threshold. Contact Nutanix support for assistance.
Cause	Background cluster maintenance tasks are not occurring.
Impact	Cluster operation may be affected.
Resolution	Contact Nutanix Support.
KB Article	1668
Severity	Critical

Table 234: Zookeeper Not Active on All CVMs [111009] [A111009]

Name	Zookeeper active on all CVMs
Description	Checks if number of active zookeepers is equal to the total number of zookeeper CVMs
Alert message	Zookeeper is not active on all CVMs.

Cause	Zookeeper service is not running on all CVMs.
Impact	Cluster resiliency is affected.
Resolution	Contact Nutanix support.
KB Article	1587
Severity	Warning

Table 235: Zookeeper nodes distributed in multi-block cluster [111010]

Name	Zookeeper nodes distributed in multi-block cluster
Description	Checks if zookeeper nodes are distributed in a multi block cluster
Cause	Zookeeper service has nodes running in the same block.
Impact	Cluster is not block aware.
Resolution	Review KB 1587.
KB Article	1587

Table 236: Zookeeper alias is incorrectly configured in the cluster. [111011] [A111011]

Name	Zookeeper Alias Check
Description	Check Zookeeper alias information configuration
Alert message	*msg*
Cause	ZK aliases in one or more sources do not match with the Zookeeper server config maintained by Zookeeper.
Impact	Nutanix utilities or Nutanix cluster services may not work as expected.
Resolution	Ensure the ZK aliases are consistent across all locations.
KB Article	1700
Severity	Critical

Table 237: GPU drivers installed [111015]

Name	GPU drivers installed
Description	Check if drivers are installed
Cause	GPU driver may be corrupt or missing or not functioning properly.
Impact	GPU card (NVIDIA) may not function properly.
Resolution	Confirm that GPU is installed and re-install or upgrade the driver. Review KB 2714.
KB Article	2714

Table 238: CVM passwordless SSH to host [111019] [A111019]

Name	Host passwordless SSH
Description	Check passwordless SSH into local hypervisor
Alert message	CVM *svm_ip* is not able to passwordless SSH to its local hypervisor
Cause	SSH disabled on hypervisor
Impact	Node will be removed from the cluster, leading to reduced capacity
Resolution	Enable SSH on hypervisor. Please refer to KB 1501 for more information.
KB Article	1501
Severity	Critical

Table 239: Non default gflags check [111020]

Name	Non default gflags check
Description	Check for non-default settings and warn if values not consistent across CVMs
Cause	One or more settings have non-default values or non-default settings do not have the same value on all CVMs.
Impact	Non-default settings will be overwritten during upgrade.
Impact	Cluster services such as stargate may fail.
Resolution	Follow the instructions in KB 1530.
KB Article	1530

Table 240: Node is in degraded state [111021] [A1305]

Name	Degraded Node check
Description	Check if any node is in degraded state
Alert message	Node *svm_ip* *svm_id* has been marked as being in a degraded state *dnd_ts*
Cause	Unstable environment such as partially bad disks, network bandwidth reduction, network packet drops, soft lockups, hardware issues(such as unreliable DIMM with ECC errors).
Impact	Overall cluster and user VM performance may be adversely affected.
Resolution	Refer to KB4206.
KB Article	4206
Severity	Critical

Table 241: Name server configuration [111022]

Name	Name server configuration
Description	Check for DNS Server configuration

Cause	Name server is not configured on CVMs and hypervisor hosts or is not able to resolve queries.
Impact	Name resolution will not work.
Resolution	Verify if Name server is configured on CVMs and on hypervisor hosts and if it can resolves queries. Review KB 3005 for more details.
KB Article	3005

Table 242: HTTP proxy check [111023]

Name	HTTP proxy check
Description	Check if HTTP proxy is working
Cause	(A) No proxy configured. (B) Unable to connect to proxy on the port.
Impact	Cluster may not have access to Internet-based service such as 1-click auto and manual downloads, Pulse, LCM inventory, etc.
Resolution	Check the proxy server settings (network, port, user credentials, etc.)
KB Article	1708

Table 243: LDAP configuration [111024]

Name	LDAP configuration
Description	Check LDAP configuration.
Cause	LDAP not correctly configured in the cluster.
Impact	Directory users might not be able to log properly.
Resolution	Review KB 2997.
KB Article	2997

Table 244: Auto support check fails [111025] [A111025]

Name	Pulse configuration
Description	Check for auto support configuration
Alert message	Pulse is not configured. Connection Status: *auto_support_check_connection_status*, Error Message: *auto_support_check_message*
Cause	Pulse is not configured.
Impact	Pulse emails will not be sent.
Resolution	Verify if Pulse is configured, if emails can be sent via SMTP server or connections to nsc01.nutanix.net/nsc02.nutanix.net hosts are successful.Review KB 1585 for more details.
KB Article	1585
Severity	Info

Table 245: Snapshots space utilization status [111026]

Name	Snapshots space utilization status
Description	Check that space occupied by snapshots is below the threshold
Cause	Snapshots disk space utilization exceeds the threshold.
Impact	Snapshots occupies most of the disk space on the cluster.
Resolution	Review KB 1853.
KB Article	1853

Table 246: Virtual IP Check [111027] [A111027]

Name	Virtual IP check
Description	Check if virtual IP is configured and reachable.
Alert message	*vm_type* **Virtual IP is configured but unreachable.**
Cause	Cluster virtual IP is not configured.
Cause	Cluster services are down or cluster is not started yet.
Impact	Nutanix features that use virtual IP address might be adversely affected.
Impact	
Resolution	Configure a valid virtual IP for the cluster.
Resolution	Verify that all cluster services are up.
KB Article	3708
Severity	Critical

Table 247: Remote syslog server not reachable [111028] [A111028]

Name	Remote syslog server check
Description	Check remote syslog server connectivity
Alert message	**Remote syslog server is not reachable.**
Cause	Remote syslog server is not reachable.
Impact	remote syslog will not be receiving log output from cluster.
Resolution	Check for network problem with rsyslog server or upstream network.
KB Article	2043
Severity	Warning

Table 248: Storage Pool SSD tier usage [111031]

Name	Storage Pool SSD tier usage
Description	Check that storage pool SSD tier usage under ILM threshold

Cause	Storage Pool SSD utilization consistently above 75%.
Impact	Performance degradation.
Resolution	Review knowledge article 2882 for more information.
KB Article	2882

Table 249: Cluster version check [111032]

Name	Cluster version check
Description	Check that cluster runs GA version of AOS
Cause	Non GA version of AOS detected.
Impact	Cluster is on Non GA version.
Resolution	Review KB 2720.
KB Article	2720

Table 250: Hardware configuration validation [111033]

Name	Hardware configuration validation
Description	Check that /etc/nutanix/factory_config.json is valid
Cause	factory_config.json does not contain valid configuration.
Impact	Service degradation.
Resolution	Review KB 3001.
KB Article	3001

Table 251: Advanced Vector Extensions (AVX) enabled [111035]

Name	Advanced Vector Extensions (AVX) enabled
Description	Checks AVX enabled on CPUs.
Cause	Intel's Advanced Vector Extensions (AVX) is disabled on the node.
Impact	If AVX is disabled, upgrade will fail.
Resolution	Enable AVX on the node. Review KB 8148.
KB Article	8148

Table 252: Hypervisor version check [111037]

Name	Hypervisor version check
Description	Check whether the same hypervisor version running on all cluster nodes.
Cause	Hypervisor version is not same across all hosts.

Impact	It is recommended to keep the hypervisor version same across all the hosts in the cluster.
Resolution	Keep hypervisor versions same across the hosts in the cluster.
KB Article	3559

Table 253: Zookeeper connections check [111038]

Name	Zookeeper connections check
Description	Check that Zookeeper is correctly configured for a cluster
Cause	Zookeeper is not configured for the number of services running per node
Impact	Some services may experience issues in connecting to Zookeeper
Resolution	Review KB 1587.
KB Article	1587

Table 254: Check that cluster virtual IP address is part of cluster external subnet [111044] [A111044]

Name	Check that cluster virtual IP address is part of cluster external subnet
Description	Check that cluster virtual IP address is part of cluster external subnet
Alert message	Cluster virtual IP address *vip* is in a different subnet from the cluster external subnet *external_subnet*
Cause	Cluster virtual IP address is in a different subnet from the cluster external subnet
Impact	Data protection in between remote sites will be affected.
Resolution	Ensure that cluster virtual IP address is part of cluster external subnet
KB Article	3817
Severity	Critical

Table 255: Mellanox NIC Speed is not 10GbE and 40GbE or both 10GbE and 40GbE NICs are installed on one node [111047] [A111047]

Name	Mellanox NIC Status check
Description	Checks if Mellanox NICs are down or if any Mellanox NIC has speed other than 10GbE or 40GbE. Checks if both 10GbE and 40GbE Mellanox NICs are installed on one node.
Alert message	Mellanox NIC on host *host_ip* has problem: *alert_msg*
Cause	The NIC is disconnected from the switch, or the switch port is not functioning correctly, or both 10GbE and 40GbE Mellanox NICs are installed on one node.

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	Ensure that all the Mellanox NIC installed on one node run on the same speed.
KB Article	4114
Severity	Warning

Table 256: Prism Central using default password [111050] [A111050]

Name	PC using default password
Description	PC is using default password
Alert message	Please update the default password on PC for user(s): *users* to harden the security on the clusters. Nutanix highly recommends not using default passwords for clusters.
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration - Using default passwords could expose your system to unauthorized access.
Resolution	Ensure that Prism Central is using a custom password for improved security.
KB Article	6153
Severity	Critical

Table 257: Mellanox NIC not Installed or with Wrong Type on host machine [111051] [A111051]

Name	Mellanox NIC Mixed Family check
Description	Checks if Mellanox NICs belong to the same Mellanox NIC family.
Alert message	Mellanox NIC installed on host : *host_ip* belong to different families.
Cause	Mellanox NICs are not from the same Mellanox NIC family.
Impact	Mixing Mellanox NICs from different families may cause degraded cluster performance.
Resolution	Replace Mellanox NICs to make sure they belong to same Mellanox NIC family.
KB Article	4329
Severity	Warning

Table 258: Notifications dropped check [111052]

Name	Notifications dropped check
Description	Some alert notifications were dropped
Cause	One or more cluster services are not healthy.
Impact	Some of the alerts will not be raised. Cluster issues may stay unnoticed.
Resolution	Contact Nutanix support.

KB Article	4905

Table 259: Cluster SSD configuration compatibility check for high frequency snapshotting [111058]

Name	Cluster SSD configuration compatibility check for high frequency snapshotting
Description	Check if the SSD configuration of the cluster is compatible with high frequency snapshotting.
Cause	Cluster does not have an SSD configuration recommended for high frequency snapshotting. (All hybrid node SSDs must be greater than 1.2TB)
Impact	High frequency snapshotting may not be sustained for some or all of the protection domains.
Impact	High latency may be seen in the user VMs.
Resolution	(A) Use a cluster with one of the recommended SSD configurations. (B) Replace the smaller SSDs in the hybrid nodes with SSD size greater than 1.2TB. (C) Use a lower snapshot frequency for entities in this cluster.
KB Article	5225

Table 260: Alert E-mail Failure [111066] [A111066]

Name	Alert E-mail Failure
Description	Failed to send alert Emails
Alert message	Unable to send alert Email from the Controller VM *service_vm_external_ip* due to *reason*.
Cause	Issue with the Email client
Impact	Alerts notification Emails will not be sent.
Resolution	Contact Nutanix Support.
KB Article	9937
Severity	Critical

Table 261: The frequency of CPU {cpu_id} on host {host_ip} is extremely low [111068] [A111068]

Name	Host CPU Frequency Check
Description	Checks the frequencies of CPUs on a node.
Alert message	CPU *cpu_id* on host *host_ip* has extremely low /high frequency: *alert_msg*
Cause	Either a power supply or a fan has failed. A failed power supply reduces the power available to the CPU, and a failed fan leads to overheating of the CPU. Each condition causes the frequency of the CPU to drop below the frequency that is required for node operations.
Impact	When CPU frequency is extremely low, the node may fail and impact workload.

Resolution	Replace the failed power supply or fans. Refer to the Nutanix documentation for instructions.
KB Article	5542
Severity	Critical

Table 262: AOS Version EOL [111070] [A111070]

Name	AOS EOL Check
Description	Check if the cluster nodes are running an AOS Version that will no longer be supported
Alert message	The AOS Version #current_AOS_version@ installed on this cluster is end of life (EOL). Upgrade AOS immediately.
Cause	End of Life (EOL) AOS Version is installed on the cluster nodes
Impact	Support may not be able to resolve issues since this AOS version will no longer be maintained
Resolution	Upgrade to a supported AOS version immediately. The AOS EOL Schedule can be found on the Nutanix support portal in KB2720: https://portal.nutanix.com/kb/2720.To determine the AOS upgrade path, whether it be long-term supported (LTS) or short-term supported (STS) refer to https://portal.nutanix.com/#/page/upgradePaths.The Nutanix SW EOL policy can be found at https://www.nutanix.com/support-services/product-support/policies/
KB Article	2720
Severity	Info

Table 263: The {storage_type} capacity of the cluster's Node(s) exceeds maximum allowed threshold. [111073] [A111073]

Name	Maximum supported HDD capacity per node for AOS version/Workloads
Description	Check to verify HDD/Total capacity per node can be supported by AOS version for cluster workload.
Alert message	The *storage_type* capacity of the Node(s) *node_list* exceeds maximum allowed threshold of *threshold_limits*TiB *reason*
Cause	HDD/Total capacity per node exceeds maximum allowed threshold for the current AOS version.
Cause	HDD/Total capacity per node exceeds maximum allowed threshold for the workload running on the cluster.
Impact	VMs might experience increased latencies.
Resolution	If the node's HDD/Total capacity exceeds the maximum allowed threshold for AOS versions pre-5.11, Nutanix recommends upgrading the cluster to AOS 5.15 (LTS) or later. Please refer to KB 7196 for additional details.

Resolution	Nutanix does not support HDD/Total capacity for the nodes that exceed the maximum allowed threshold limit for Objects/Files/Mine workload. Please refer to KB 7196 for additional details.
KB Article	7196
Severity	Warning

Table 264: Current memory configuration on the nodes does not support the {storage_type} capacity. [111074] [A111074]

Name	Maximum supported HDD/Total capacity per node for memory size check
Description	Check to verify whether the HDD/Total capacity of the node is supported by current memory configured on the node.
Alert message	Current memory configuration *available_memory_list* does not support the *storage_type* capacity on the node(s): *node_list*
Cause	Configured memory does not support the HDD/Total capacity installed on the nodes.
Impact	VMs might experience increased latencies.
Resolution	Current memory configured on the node(s) does not support the HDD/Total capacity installed. Please consider upgrading memory configuration on the node(s).Please refer to KB 7196 for additional details.
KB Article	7196
Severity	Warning

Table 265: SW Encryption Keys from {kms_name} are unavailable [111075] [A111075]

Name	SW Encryption Keys Unavailable
Description	Check to see if the KMS actually has the expected passwords for the container encryption keys.
Alert message	SW encryption keys for one or more containers are not available from external key manager *kms_name*.
Cause	If a 3rd party KMS gets restored from a backup or other incorrect management operation, it may have missing keys.
Impact	Cluster will have data unavailability if rebooted.
Resolution	Contact Nutanix and 3rd party key manager support.
KB Article	8223
Severity	Critical

Table 266: Cluster's local KMS is running in unsafe mode [111078] [A111078]

Name	Local Key Manager running in unsafe mode
Description	Local Key Manager is configured on a cluster with less than 3 nodes.

Alert message	Cluster's local KMS is configured on a cluster with less than 3 nodes.
Cause	Local Key Manager needs at least 3 nodes in a cluster to securely store its master key.
Impact	Local Key Manager's master key is not stored securely.
Resolution	Register to Prism Central and switch to Prism Central based key manager or switch to an External Key Manager.
KB Article	8223
Severity	Warning

Table 267: High number of tasks in the cluster [111082]

Name	High number of tasks in the cluster
Description	High number of tasks in the cluster
Cause	High number of tasks in the cluster
Impact	Operation on entities in the cluster might not progress.
Resolution	Contact Nutanix support for help.
KB Article	9866

Table 268: Current configuration of the cluster cannot support RPO under 6 hours [111083] [A111083]

Name	Minimum dense node cluster configuration check to support RPO under 6 hours
Description	Check to verify if the dense node cluster configuration can support RPO under 6 hours.
Alert message	The node(s) *node_list* do not support RPO under 6 hours. *message*
Cause	HDD/Total Capacity of the Node is greater than the allowed threshold for PDs with RPO less than 6 hours
Cause	The CVM configuration of the Node cannot support the current RPO of under 6 hours.
Impact	Unsupported configuration. Cluster performance may be significantly degraded.
Resolution	If the current capacity of the node cannot support the RPO, either the RPO must be increased to 6 hours or hardware configuration needs to be updated. Please refer to KB 7196 for additional details.
Resolution	If the CVM configuration of the node cannot support the RPO, either the RPO must be increased to 6 hours or hardware configuration needs to be updated. Please refer to KB 7196 for additional details.
KB Article	7196
Severity	Warning

Table 269: RF1 Container Check [111085]

Name	RF1 Container Check
Description	List all RF1 containers and active VMs in the cluster
Cause	Found RF1 containers and VMs in the cluster
Impact	RF1 is enabled in the cluster. VMs that have vDisks or VGs on RF1 containers have single copy of data in the cluster. Those VMs will experience data un-availability if the affinity host experiences failure.
Resolution	Not Applicable
KB Article	10663

Table 270: Remote syslog server forwarding failed. [111087] [A111087]

Name	Remote syslog forwarding failure check
Description	Check remote syslog server log forwarding.
Alert message	Remote syslog server forwarding failures detected on CVM: *svm_ip*
Cause	Remote syslog server forwarding failed.
Impact	remote syslog will not be receiving log output from cluster.
Resolution	Check for forwarding failure causes with rsyslog server. Refer KB 10735
KB Article	10735
Severity	Warning

Table 271: Cannot Remove Password Protected Disk(s) [130003] [A1106]

Name	Password Protected Disk Status
Description	Cannot remove password protected disk(s).
Alert message	Cannot remove password protected disks *disk_id_list* from the configuration as they cannot be cleaned.
Cause	The drives are offline, the drives are removed, or the node is down.
Impact	The disks cannot be removed from the configuration.
Resolution	Ensure the disks are accessible.
Severity	Warning

Table 272: Automatic Addition Of Node To Metadata Ring Disabled [130004] [A1079]

Name	Node Auto Add Check
Description	Automatic addition of the node to the metadata ring is disabled.
Alert message	Automatic addition of node to metadata ring has been disabled on CVM *service_vm_id* with IP address *ip_address*.

Cause	This node has been removed from the metadata store after being automatically added recently. Automatic addition of this node to the metadata store has now been disabled.
Impact	Cluster performance may be significantly degraded.
Impact	In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Automatic addition will be re-enabled when the node is manually added to the metadata store. If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. Otherwise, contact Nutanix support.
Severity	Warning

Table 273: Node Detached From Metadata Ring [130005] [A1055]

Name	Node Detached From Ring Check
Description	The metadata drive is detached from the ring.
Alert message	Node *ip_address* is now detached from metadata ring due to *reason*.
Cause	Either a metadata drive has failed, the node was down for an extended period of time, or an unexpected subsystem fault was encountered, so the node was removed from the metadata store.
Impact	Cluster performance may be significantly degraded.
Impact	In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	(A) If the metadata drive has failed, replace the metadata drive as soon as possible. Refer to the Nutanix documentation for instructions. (B) If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. For assistance, contact Nutanix support.
Severity	Warning

Table 274: Metadata Dynamic Ring Change Operation Stuck [130006] [A1117]

Name	Metadata DynRingChangeOp Status
Description	Dynamic Ring Change operation is not making progress.
Alert message	*operation* for node *ip_address* is not progressing. Node performing the operation: *ip_address_doing_ring_change*.
Cause	Node addition or removal is not making progress. A node in the cluster is unhealthy.
Impact	Node addition and node removal operations will not proceed to completion.
Resolution	If the alert persists, contact Nutanix Support.
KB Article	8195
Severity	Critical

Table 275: Metadata Dynamic Ring Change Operation Too Slow [130007] [A1116]

Name	Metadata DynRingChangeOp Slow Check
Description	Dynamic Ring Change operation too slow.
Alert message	*operation* for node *ip_address* is progressing slowly. Total elapsed time: *elapsed_time_mins* min(s). Node performing the operation: *ip_address_doing_ring_change*.
Cause	A node in the cluster is unhealthy, or the cluster contains an extraordinarily large amount of data.
Impact	Cluster performance may be degraded.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.
Severity	Warning

Table 276: Node Forwarding Metadata Requests [130008] [A1037]

Name	Node In Forwarding Mode Check
Description	Node Forwarding Metadata Requests
Alert message	Metadata service on node *ip_address* is running in forwarding mode due to *reason*.
Cause	Either a metadata drive has failed, node removal has been initiated, or an unexpected subsystem fault has been encountered.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
KB Article	8266
Severity	Critical

Table 277: Large Metadata Size Detected [130009] [A1119]

Name	Metadata Size
Description	Large metadata size detected.
Alert message	Node *ip_address* contains a large amount of metadata.
Cause	A node in the cluster contains a large amount of metadata and has exceeded thresholds.
Impact	Cluster performance may be degraded.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.
Severity	Warning

Table 278: Node Marked To Be Auto Added To Metadata Ring [130010] [A1078]

Name	Node Marked To Be Auto Added Check
Description	Node Marked To Be Auto Added To Metadata Ring
Alert message	Node *service_vm_id* with IP address *ip_address* is marked to be added to metadata ring. *reason*.
Cause	Node was detached from metadata ring earlier. It is now stable and will be added back to the metadata store.
Impact	The metadata service will start to use the node.
Resolution	No action is necessary.
Severity	Info

Table 279: Node Marked To Be Detached From Metadata Ring [130011] [A1054]

Name	Node Marked To Be Detached Check
Description	The node is marked to be detached from the metadata ring.
Alert message	Node *ip_address* is marked to be detached from metadata ring due to *reason*.
Cause	Either a metadata drive has failed, the node was down for an extended period of time, or an unexpected subsystem fault was encountered, so the node is marked to be removed from the metadata store.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	(A) If the metadata drive has failed, replace the metadata drive as soon as possible. Refer to the Nutanix documentation for instructions. (B) If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. For assistance, contact Nutanix support.
KB Article	8408
Severity	Critical

Table 280: Metadata Ring Imbalance [130012] [A1072]

Name	Metadata Imbalance Check
Description	Cassandra metadata imbalance.
Alert message	Metadata ring imbalance for *num_partitions* partitions. The partitions are *partition_list*.
Cause	One or more nodes have a disproportionately larger token range size. This imbalance may cause performance bottlenecks on the node(s) affected.
Impact	Node performance may be significantly degraded.
Resolution	Execute cassandra ring skew fix operation using the NCLI command 'cluster cassandra-token-range-skew-fixer-start'. UVMs' performance might be affected for a small amount of time.
Severity	Warning

Table 281: Curator Job Running Too Long [130015] [A1120]

Name	Curator Job Status
Description	The Curator job has been running for too long.
Alert message	Curator job *name* with id *execution_id* has been running for a long time i.e. *elapsed_time_secs* seconds.
Cause	Various
Impact	Background cluster maintenance tasks might get affected in the future. The root cause should be addressed as soon as possible.
Resolution	Contact Nutanix support if this issue persists.
Severity	Info

Table 282: Curator Scan Failure [130016] [A1081]

Name	Curator Scan Status
Description	Curator scan failure.
Alert message	Curator *scan_type* scans have repeatedly failed to complete.
Cause	Various
Impact	Background cluster maintenance tasks are not occurring. While cluster operation will not be immediately affected, the root cause should be addressed as soon as possible.
Resolution	Contact Nutanix support.
KB Article	3786
Severity	Critical

Table 283: Duplicate IP Address Detected [130017] [A1096]

Name	IP Configuration
Description	Duplicate IP address is detected.
Alert message	Detected IP address conflict for cluster virtual IP address *duplicate_ip*. IP found on NIC with MAC address *mac_address*.
Cause	The cluster virtual IP address may be configured on another host.
Impact	The cluster is not accessible through the virtual IP address. Other services that use the virtual IP address, such as remote replication or SCCM, may be affected.
Resolution	Either configure a different cluster virtual IP address or unconfigure the IP address from the other host.
Severity	Warning

Table 284: Duplicate Remote Cluster ID [130018] [A1038]

Name	Duplicate Remote Cluster ID Check
Description	Duplicate remote cluster ID.
Alert message	Remote site already exists with the name '*conflicting_remote_name*'.
Cause	Two remote sites with different names or different IP addresses have same cluster ID. This can happen in two cases: (a) A remote cluster is added twice under two different names (through different IP addresses) or (b) Two clusters have the same cluster ID.
Impact	Protected data is not replicated to the remote site.
Resolution	In case (a) remove the duplicate remote site. In case (b) verify that the both clusters have the same cluster ID and contact Nutanix support.
Severity	Warning

Table 285: IP Address Not Hosted [130023] [A1097]

Name	Virtual IP Configuration
Description	IP address is not hosted.
Alert message	Unable to host virtual IP address *ip_address*.
Cause	The network configuration is incorrect.
Impact	The cluster is not accessible through the virtual IP address. Other services that use the virtual IP address, such as remote replication or SCCM, may be affected.
Resolution	Contact Nutanix support.
Severity	Critical

Table 286: SMTP Error [130060] [A1080]

Name	SMTP Configuration
Description	Failed to send Pulse and/or alert emails to Nutanix support with the current configuration.
Alert message	Failed to send email from Controller VM *ip_address* due to error: *error*.
Cause	SMTP server may not be configured.
Cause	The SMTP server is not reachable from the Controller VM.
Cause	Firewall is preventing emails from being forwarded through the SMTP server.
Impact	Nutanix support cannot proactively monitor cluster health and offer assistance before issues become acute.
Resolution	Configure the SMTP server.
Resolution	SMTP server should be reachable.
Resolution	Ensure that the firewalls allow the cluster to send email messages through the Nutanix SMTP servers (nsc01.nutanix.net and nsc02.nutanix.net) on port 80.

KB Article	1638
Severity	Warning

Table 287: Remote Site Latency Is High. [130075] [A1168]

Name	RemoteSiteHighLatency
Description	Latency to a Metro remote site is high.
Alert message	Latency to remote site '*remote_name*' is more than '*remote_site_high_latency_threshold_ms*' ms , for '*sustained_high_latency_secs*' seconds due to *reason*.
Cause	Network issues or CVM resource constraints
Impact	Replication will fail.
Resolution	Check if remote site is reachable and resolve any network issues. Monitor CVM CPU and memory usage
Severity	Warning

Table 288: Possible degraded Node [130087] [A130087]

Name	Node Degradation Status
Description	Services on one of the nodes in the cluster are in a degraded state and possibly not making progress.
Alert message	Possible degraded node *ip_address* with Controller VM ID *service_vm_id* reported by component *component_name*.
Cause	Various
Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix support.
KB Article	3827
Severity	Critical

Table 289: CPS Deployment Evaluation Mode [130093] [A130093]

Name	CPS Deployment Evaluation Mode
Description	Checks if the CPS deployment is running in evaluation mode.
Alert message	CPS deployment of cluster with cluster id *cluster_id* is running in evaluation mode. Please apply a license by following the CPS deployment guide.
Cause	CPS is running in evaluation mode.
Impact	The deployment cannot be in evaluation and the user has to apply a license.
Resolution	Apply a license to the deployment.
Severity	Warning

Table 290: Metro connectivity lost [130119] [A130119]

Name	Metro Connectivity Lost
Description	Metro connectivity is lost
Alert message	Metro connectivity from local site '*local_stargate_handler*' to remote site '*remote_name*' is lost due to '*reason*'
Cause	Remote site(s) not reachable
Impact	Metro connection of the remote site is unavailable
Resolution	Check remote site(s) connectivity and re-enable Metro if necessary
Severity	Warning

Table 291: Alert raised on cloud remote site {remote_name}: {alert_message} [130135] [A130135]

Name	Cloud Remote Alert
Description	Alert generated on cloud remote site.
Alert message	Alert raised on cloud remote site *remote_name*: *alert_message*
Cause	Various
Impact	Various
Resolution	Resolve the issue stated in the alert. If you cannot resolve the issue, please contact Nutanix support.
Severity	Warning

Table 292: Snapshot Reserve on SSD is Full. [130139] [A130139]

Name	LWS Store Full
Description	SSD reserved for the Nearsync snapshots is full.
Alert message	The SSD space that is reserved for the snapshots is full.
Cause	Snapshot reserve on the SSD is full or has high IO workload.
Impact	RPO cannot be met and will be temporarily transitioned to 60 minutes.
Resolution	Reduce the retention time of the snapshots.
Severity	Warning

Table 293: Two node cluster state changed [130151] [A130151]

Name	Change in state of two node cluster
Description	Change in state of two node cluster
Alert message	Two node cluster state changed: *message*
Cause	There is a change in availability of a node (or network between the nodes).

Impact	If the cluster is now in stand-alone mode, it cannot handle the remaining nodes failure.
Resolution	Two node cluster will auto-heal when node (or network between the nodes) is available again.
KB Article	8284
Severity	Critical

Table 294: Two node cluster changed state to stand-alone mode. [130152] [A130152]

Name	Two node cluster changed state to stand-alone mode
Description	One of the nodes in a two node cluster is currently unavailable, the cluster is now operating in stand-alone mode.
Alert message	**Two node cluster state change:** *message*
Cause	There was an unplanned event that changed the availability of a node (or network between the nodes).
Impact	The cluster is not node fault tolerant until the cluster auto-heals.
Resolution	Two node cluster will auto-heal when node (or network between the nodes) is available again.
KB Article	8285
Severity	Warning

Table 295: Two node cluster state change to stand-alone mode [130153] [A130153]

Name	Two node cluster changed now in stand-alone mode
Description	Change in state of two node cluster
Alert message	**Two node cluster state change:** *message*
Cause	There is a change in availability of a node (or network between the nodes). This event was planned.
Impact	The cluster is not node fault tolerant until the cluster auto-heals.
Resolution	Two node cluster will auto-heal when the planned event is completed successfully.
Severity	Warning

Table 296: State changed for two node cluster [130154] [A130154]

Name	Two node cluster changed state back to a two-node mode.
Description	Change in state of two node cluster
Alert message	**Two node cluster state changed:** *message*
Cause	A previously unavailable node or service is now available again.

Impact	No impact.
Resolution	No resolution necessary.
Severity	Info

Table 297: Witness is unreachable from node. [130155] [A130155]

Name	Witness VM is unreachable from a node in two node cluster
Description	Unable to reach the Witness VM over the network, the cluster may not be able to handle a node failure.
Alert message	The node cannot issue witness APIs to witness: *message*
Cause	Node is unable to issue APIs to witness VM.
Impact	The cluster may not be able to handle node failure until the witness is reachable again.
Resolution	Check the witness VM to verify it is up and verify the cluster nodes can communication with the Witness VM over the network.
Severity	Warning

Table 298: Drive Removal Stuck [130156] [A130156]

Name	Drive Removal Stuck
Description	Drive Removal Stuck
Alert message	Removal of drive *disk_id* is stuck.
Cause	Various
Impact	Drive cannot be removed from the cluster.
Resolution	Contact Nutanix support.
Severity	Warning

Table 299: Unable to retrieve the endpoint information for the Availability Zone. [130173] [A130173]

Name	Unable to get Availability Zone Endpoint
Description	Unable to retrieve the endpoint information for the Availability Zone
Alert message	Unable to retrieve the endpoint information for the Availability Zone *availability_zone_physical_name*
Cause	Prism Central, Prism Element, or Xi - Availability Zone may be unreachable
Impact	VMs may not be fully protected
Resolution	Check if the Prism Central, Prism Element, or Xi - Availability Zone are reachable. If the problem persists, contact Nutanix support.
KB Article	8274

Severity	Critical

Table 300: Unable to communicate with the Data Center Manager [130175] [A130175]

Name	Failure to communicate with DCM
Description	Unable to retrieve the endpoint information for the Availability Zone from the Data Center Manager
Alert message	Unable to retrieve the endpoint information for the Availability Zone *availability_zone_name* from the Data Center Manager
Cause	Data Center Manager may be unreachable or has provided an invalid response
Impact	VMs may not be fully protected
Resolution	Check if Data Center Manager is reachable from the Xi - Availability Zone
KB Article	6901
Severity	Critical

Table 301: Data At Rest Encryption key backup warning [130184] [A130184]

Name	Data At Rest Encryption key backup warning.
Description	New encryption keys not backed up.
Alert message	Data At Rest Encryption key backup warning: *message*
Cause	New encryption keys have been created since the last backup was downloaded.
Impact	If the encryption keys are lost, the data in the storage container cannot be accessed.
Resolution	Download the latest copy of the cluster's key backup.
Severity	Warning

Table 302: Local key manager master key rotation warning. [130185] [A130185]

Name	Local key manager master key rotation warning.
Description	Unable to rotate local key manager's master key.
Alert message	Local key manager master key rotation warning. Last failure reason: *failure_message*
Cause	One or more nodes in the cluster are unavailable.
Impact	If the cluster experiences any more node failures, the data on the cluster may be inaccessible.
Resolution	Ensure Mantle service is UP on all nodes. If the warning persists, contact Nutanix support.
Severity	Warning

Table 303: Remote Availability Zone Latency is High [130187] [A130187]

Name	Latency to Remote Availability Zone is High
Description	Network latency to the remote Availability Zone is high.
Alert message	Network latency to the remote Availability Zone '*availability_zone_physical_name*' exceeds the threshold limit.
Cause	Network latency to the remote Availability Zone exceeds the threshold limit.
Impact	Replications will fail
Resolution	Check if the remote Availability Zone is reachable over the network.
Severity	Warning

Table 304: Cluster Join to Domain Failure [130189] [A1154]

Name	ClusterJoinToDomainFailure
Description	Failed to join AOS cluster to domain.
Alert message	Failed to join AOS cluster to domain *domain_name*: *reason*
Cause	Cluster could not be joined to domain because of incorrect login credentials or other causes.
Impact	Workflows such as backup service (volume shadow copy service) and Kerberos for SMB will not work.
Resolution	Please ensure that the specified domain account credentials are valid and that the virtual computer object representing the AOS cluster does not already exist.
KB Article	8080
Severity	Critical

Table 305: Background Encryption Stuck [130194] [A130194]

Name	Background Encryption Stuck
Description	Background Encryption Stuck
Alert message	Encryption of container *container_id* is not progressing. Contact Nutanix support for further assistance.
Cause	Oplog entries, extent groups, NFS entries or other components are not being encrypted.
Impact	Encryption is not progressing.
Resolution	Contact Nutanix support.
Severity	Warning

Table 306: Failed to reserve host memory for Atlas networking [130202] [A130202]

Name	Failed to reserve host memory for Atlas networking
Description	Failed to reserve host memory for Atlas networking.
Alert message	Failed to reserve host memory for Atlas networking
Cause	Failed to reserve host memory for Atlas networking.
Impact	If hypervisor memory is exhausted, random VMs may be powered off to reclaim memory.
Resolution	Power off VMs to reduce host memory consumption.
KB Article	8525
Severity	Critical

Table 307: Cluster is susceptible to CopyUpBlockMap Issue [130336] [A130336]

Name	Stale snapshot chain view due to CopyUpBlockMap issue
Description	Check whether cluster is susceptible to CopyUpBlockMap issue
Alert message	This cluster is potentially affected by a rare data integrity issue. For corrective actions, please review Field Advisory #0075 under the following support portal link: https://portal.nutanix.com/page/documents/field-advisories/list/
Cause	CopyUpBlockMap caused stale snapshot chain view.
Impact	Multiple stargate may fatal due to stale snapshot chain view which can cause outage.
Resolution	Review Field Advisory #0075 under the following support portal link: https://portal.nutanix.com/page/documents/field-advisories/list/
KB Article	8999
Severity	Critical

Table 308: Possible Service Degradation. [130338] [A130338]

Name	Service instances receiving bad degraded node scores
Description	Service instances on a few nodes of the cluster are possibly in a bad state.
Alert message	Possible degradation of *service_name* instances on the Controller VM IDs *service_vm_id_list*.
Cause	Various
Impact	Cluster performance may be significantly degraded.
Resolution	Follow KB9390 to troubleshoot reasons for the service receiving bad scores.
KB Article	9390
Severity	Warning

Table 309: Cross-container disk migration task is paused [130340] [A130340]

Name	Cross-container disk migration task is paused
Description	Cross-container disk migration task is paused.
Alert message	Migration of VM disk '*vdisk_name*' (VM '*vm_name*') from source storage container '*src_container_name*' to destination storage container '*dst_container_name*' is paused
Cause	Cross-container disk migration task is paused due to lack of free space in the destination storage container.
Impact	Cross-container disk migration cannot be completed.
Resolution	Free up additional space in the destination storage container or abort the parent task.
KB Article	10248
Severity	Warning

Table 310: Zookeeper server unresponsive. [130344] [A130344]

Name	Zookeeper server unresponsive for a long time
Description	Zookeeper server might be unresponsive on a CVM for a long time.
Alert message	Zookeeper server on the Controller VM ID *service_vm_id* is unresponsive since *death_timestamp_secs*.
Cause	Zookeeper server on one of the nodes has been unresponsive for a long time.
Impact	Cluster performance may be significantly degraded.
Resolution	Follow KB1587 to troubleshoot reasons for Zookeeper server being unresponsive.
KB Article	1587
Severity	Warning

Table 311: Detected iSCSI discovery or login through cluster external IP address [130346] [A130346]

Name	Invalid iSCSI portal
Description	Detected iSCSI discovery or login through cluster external IP address
Alert message	Detected iSCSI *access_type* through cluster external IP address
Cause	The iSCSI discovery or login portal is incorrectly configured
Impact	The iSCSI service won't be highly available and the performance would be affected
Resolution	Use the cluster external data services IP address for iSCSI discovery
Severity	Warning

Table 312: MEM-NVMe Tiering Cluster Constraints Check Failure [130360] [A130360]

Name	MEM-NVMe tiering disabled
Description	The cluster requirements for enabling the MEM-NVMe storage tier feature are not met. MEM-NVMe storage tiering will be disabled.
Alert message	Cluster fails to meet the MEM-NVMe tiering constraints due to reason: *reason*
Cause	Failing to meet the MEM-NVMe tiering constraints results in MEM-NVMe tiering being disabled.
Impact	All the MEM-NVMe disks will be treated as SSD-PCIe disks.
Resolution	Configure the cluster to meet the MEM-NVMe tiering constraints. All the nodes in the cluster should have at least one MEM-NVMe disk and the cluster should not contain any {tiers_not_allowed} disk.
KB Article	11650
Severity	Warning

Table 313: Pulse cannot connect to REST server endpoint [140001] [A140001]

Name	REST endpoint connection status
Description	Check if Pulse can connect to REST server endpoint https://insights.nutanix.com:443.
Alert message	Pulse cannot connect to REST server endpoint https://insights.nutanix.com:443. Connection Status: *connection_status*, Pulse Enabled: *enable*, Error Message: *message*
Cause	Pulse cannot connect to REST server endpoint https://insights.nutanix.com:443.
Impact	Data driven serviceability and customer support cannot be performed.
Resolution	Ensure that the REST server endpoint https://insights.nutanix.com:443 is reachable from Pulse.
KB Article	5490
Severity	Warning

Table 314: Metering server endpoint connection unavailable. [140003] [A140003]

Name	Metering endpoint connection status
Description	Checks if the connection to the metering server endpoint is available and if the metering stats can be successfully sent to server endpoint.
Alert message	Cannot publish metering stats to server endpoint on a metering enabled cluster.
Cause	Unable to verify metering enablement state and/or unable to establish a connection to the metering server endpoint.
Impact	Metering and billing stats cannot be collected from this cluster violating the terms of service.
Resolution	Ensure that the metering server endpoint is resolvable and reachable from cluster and that the network connection is stable.

KB Article	10651
Severity	Critical

Table 315: Pulse is not enabled on cluster. [140004] [A140004]

Name	Pulse enablement check
Description	Checks if the pulse is disabled on cluster.
Alert message	**Pulse is not enabled on cluster**
Cause	Pulse is not enabled on cluster.
Impact	Telemetry data cannot be sent thus impacting Nutanix's ability to provide proactive support
Resolution	Ensure that pulse is enabled on the cluster to allow Nutanix to provide proactive support.
KB Article	10743
Severity	Info

Table 316: OVS Service Restart [150002] [A150002]

Name	OVS Service Restarted
Description	OVS service has unexpectedly restarted on one or more AHV hosts.
Alert message	Node *uuid_node*: detected unexpected restart of ovs-vswitchd process on host *host_ip* from old_pid *old_pid* to new_pid *new_pid*. OVS service configuration was restored after restart. Guest VMs might have experienced temporary network disconnect.
Cause	OVS service was restarted.
Impact	Guest VMs may temporarily lose network connectivity.
Resolution	Bridges and flow rules on hosts have been rebuilt. No user action is required.
KB Article	6151
Severity	Warning

Table 317: Flow visualization statistics collector service restart detected [150003] [A150003]

Name	Flow visualization statistics collector service monitor
Description	Check if flow visualization statistics collector service was restarted.
Alert message	Flow visualization statistics collector service was restarted *max_restart_count* times in last *time_interval_min* minutes on host *host_ip*.
Cause	Flow visualization statistics collector service was restarted 10 times in the last 15 minutes.
Impact	Flow visualization may not show real time data.
Resolution	Review KB 8911.

KB Article	8911
Severity	Warning

Table 318: Inconsistent Virtual Switch State Detected [150004] [A150004]

Name	Inconsistent virtual switch state was detected
Description	Virtual switch state is inconsistent across nodes in the cluster.
Alert message	Virtual switch *vs_name* is in an inconsistent state for the *host_info*. Reason: *error_message*
Cause	A virtual switch is not present on at least one node in the cluster.
Impact	VMs using this virtual switch may lose network connectivity when migrated to the affected node.
Resolution	(A) Review virtual switch configuration and ensure every host (storage-only hosts are optional) in the cluster is included. (B) Ensure all the hosts and their physical NICs used in the virtual switch configuration are present.
Severity	Critical
Severity	Warning

Table 319: Default Virtual Switch Error [150005] [A150005]

Name	Default virtual switch error
Description	Default virtual switch has error in the cluster.
Alert message	Default virtual switch cannot be created on the *host_info*. Reason: *error_message*
Cause	Default virtual switch could not be created on a node in the cluster.
Impact	Networking operations may be impacted on the affected node(s) in the cluster.
Resolution	Check Acropolis service logs and refer to KB 9233 to identify the problem and steps to restore the default virtual switch in the cluster.
KB Article	9233
Severity	Critical

Table 320: IGMP Snooping needs to be re-configured on newly migrated virtual switch. [150008] [A150008]

Name	IGMP Snooping needs to be re-configured on newly migrated virtual switch
Description	Notifies when IGMP Snooping has been configured on an OVS bridge that's migrated to a virtual switch
Alert message	IGMP Snooping was found to be configured on bridge *bridge_name*, when it was migrated to virtual switch *vs_name* with UUID *vs_uuid*. This configuration won't be carried over to the virtual switch *vs_name*, please re-configure it.

Cause	IGMP Snooping was previously enabled on the bridge before its migration to a virtual switch.
Impact	Existing IGMP Snooping configuration on the OVS bridge is not reflected in the migrated virtual switch configuration. Any further update of virtual switch will reconfigure IGMP Snooping based on its configuration in virtual switch.
Resolution	Please re-configure IGMP Snooping on the bridge using the migrated virtual switch.
Severity	Warning

Table 321: Inconsistent SPAN Session State Detected [150009] [A150009]

Name	Inconsistent SPAN session state was detected
Description	SPAN session state is inconsistent on one of the nodes in the cluster.
Alert message	SPAN session *session_name* is in an inconsistent state for the *host_info*. Reason: *message*
Cause	A SPAN session is in an inconsistent state on one of the nodes in the cluster.
Impact	The SPAN session mirroring traffic from the host NIC(s) to a destination VM NIC is no longer active.
Resolution	Ensure that the source ports and/or destination ports for the SPAN session are connected and in active state.
KB Article	10440
Severity	Warning

Table 322: IPFIX exporter cannot be enabled on the cluster [150010] [A150010]

Name	IPFIX exporter cannot be enabled on the cluster
Description	Cannot enable IPFIX exporter on the cluster
Alert message	IPFIX exporter *exporter_name* cannot be enabled on cluster with UUID *cluster_uuid*. Reason: *cause*
Cause	IPFIX exporter could not be enabled on the cluster.
Impact	Network connection details cannot be exported via IPFIX from this cluster.
Resolution	Check the software version(s) of Prism Element and AHV and upgrade them to the minimum required version(s). Make sure the connection between Prism Central and Prism Element is healthy.
Severity	Warning

Table 323: File server unreachable check [160008] [A160008]

Name	File Server Network Reachable
Description	Checks if the file server is reachable.
Alert message	File server *file_server_name* is unreachable.

Cause	File server network is unreachable.
Impact	File server operations will not work.
Resolution	Refer to KB article 3691. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Warning

Table 324: File server down check [160009] [A160009]

Name	File Server Reachable
Description	Checks if the file server is down.
Alert message	File server *file_server_name* is down.
Cause	All the file server VMs are powered off.
Impact	File server operations will not work.
Resolution	Refer to KB article 3691. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Critical

Table 325: File server multiple VMs on single node check [160010] [A160010]

Name	File Server VM Status
Description	Checks if multiple file server VMs are running on a single node.
Alert message	File server *file_server_name* has multiple VMs on a single node.
Cause	File server VMs are running on a single node.
Impact	Single node failure may bring down the file server.
Resolution	Refer to KB article 3027. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3027
Severity	Warning

Table 326: File server services down check [160011] [A160011]

Name	File Server Status
Description	Checks if all the file server services are running.
Alert message	File server *file_server_name* services (*service_names*) are down.
Cause	One or more services are not running on File server VMs.
Impact	File server operations will not work.

Resolution	Refer to KB article 3691. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Critical

Table 327: File Server PD activates on multiple sites [160022] [A160022]

Name	File Server PD Active On Multiple Sites
Description	Checks that File Server PD Active On Multiple Sites
Alert message	File server *file_server_name* protection domain *protection_domain_name* activates on multiple site(s) *remote_site_names*
Cause	File Server PD activates on multiple sites
Impact	File server operations might not work
Resolution	Refer to KB article 4380. Contact Nutanix support if issue still persists or assistance needed.
KB Article	4380
Severity	Warning

Table 328: File server invalid snapshot [160027] [A160027]

Name	File Server Invalid Snapshot Warning
Description	Checks that File server is in healthy state to take snapshot
Alert message	File server *file_server_name* is not healthy and leads to invalid snapshot
Cause	File Server NVMs might deleted
Impact	File server operations will not work and creates the bad snapshots
Resolution	Refer to KB article 4312. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Warning

Table 329: File server entities not protected [160029] [A160029]

Name	File Server Entities Not Protected
Description	Checks all File server entities are protected
Alert message	Some of File server *file_server_name* entities are not protected
Cause	Some of the File Server entities are not protected
Impact	File server migration might fails and creates the bad snapshots
Resolution	Refer to KB article 4312. Contact Nutanix support if issue still persists or assistance needed.

KB Article	4312
Severity	Warning

Table 330: Multiple File server versions are present in the cluster. [160030] [A160030]

Name	Multiple File Server Versions Check
Description	Checks whether multiple File server versions are running in the cluster
Alert message	File Servers are on different versions. Please upgrade all File Servers to the same version.
Cause	Multiple File server(s) running on different versions.
Impact	Feature availability is not consistent across file servers.
Resolution	Refer to KB article 4297 for more details. Contact Nutanix support if assistance is needed.
KB Article	4297
Severity	Info

Table 331: File server upgrade task hung for too long [160031] [A160031]

Name	File Server Upgrade Task Stuck Check
Description	Check File server task is hung for a long time.
Alert message	File server *file_server_name* upgrade task hung for too long
Cause	Various. Check alert message for details.
Impact	FileServer upgrade task could fail.
Resolution	Refer to KB article 4314. Contact Nutanix support if issue still persists or assistance needed.
KB Article	4314
Severity	Warning

Table 332: File Server PD enabled on non-compatible Remote Site [160033] [A160033]

Name	Remote Site Not File Server Capable
Description	Checks File server remote sites are file server capable
Alert message	File server *file_server_name* protection domain *protection_domain_name* enabled on non-compatible Remote site *remote_site_name*
Cause	Remote site is not File server capable
Impact	File server migration/failover will fail
Resolution	Refer to KB article 4380. Contact Nutanix support if issue still persists or assistance needed.
KB Article	4380

Severity	Warning

Table 333: Remote Site AOS not compatible with File Server version [160056] [A160056]

Name	Remote Site AOS not compatible with File Server version
Description	Checks File server remote sites AOS are file server version compatible
Alert message	File server *file_server_name* protection domain *protection_domain_name* enabled on non-compatible Remote site *remote_site_name*: *reason*
Cause	Remote site AOS is not compatible with File server version
Impact	File server migration/failover will fail
Resolution	Refer to KB article 4380. Contact Nutanix support if issue still persists or assistance needed.
KB Article	4380
Severity	Critical

Table 334: Analytics VM one or more component failure. [160064] [A160064]

Name	Analytics VM one or more component failure.
Description	One or more components of the File Analytics VM {avm_ip} are not functioning properly or have failed.
Alert message	One or more components of the File Analytics VM *avm_ip* are not functioning properly or have failed.
Cause	One or more File Analytics VM components have failed.
Impact	File Server analytics will not perform appropriately.
Resolution	Contact Nutanix Support.
KB Article	8239
Severity	Critical

Table 335: Analytics VM high CPU usage. [160065] [A160065]

Name	Analytics VM high CPU usage.
Description	File Analytics VM {avm_ip} has a high CPU usage of {usage_percent}%
Alert message	File Analytics VM *avm_ip* has a high CPU usage of *usage_percent*%
Cause	File Analytics VM is performing too many CPU intensive operations.
Impact	File Analytics VM's performance will be impacted.
Resolution	Monitor which component has a high CPU usage and reduce the processing load.
Severity	Warning

Table 336: Analytics VM high disk usage. [160066] [A160066]

Name	Analytics VM high disk usage.
Description	File Analytics VM {avm_ip} has a high disk usage of {usage_percent}%
Alert message	File Analytics VM *avm_ip* has a high disk usage of *usage_percent*%
Cause	Extensive disk usage on the File Analytics VM
Impact	File Analytics VM's performance will be impacted.
Resolution	Add more capacity to the File Analytics VM or delete some data.
Severity	Warning

Table 337: Analytics VM low memory available. [160067] [A160067]

Name	Analytics VM low memory available.
Description	The memory available on File Analytics VM {avm_ip} is {memory_available} bytes which is low.
Alert message	The percent memory available on File Analytics VM *avm_ip* is *memory_available* which is low
Cause	Excessive memory usage on the File Analytics VM
Impact	File Analytics VM's performance will be impacted.
Resolution	Add more RAM
Severity	Warning

Table 338: File server version check [160070]

Name	File server version check
Description	Check that file server runs GA version of Nutanix Files
Cause	Non-GA version of file server detected.
Impact	File server is on non-GA version.
Resolution	Refer to KB article 6601. Contact Nutanix support if issue still persists or assistance needed.
KB Article	6601

Table 339: File server VM Hardware Clock Timezone not Supported [160071] [A160071]

Name	File server VM Hardware Clock Timezone not Supported
Description	Checks that file server VM hardware clock timezone is UTC
Alert message	File server *file_server_name* : File server VM *vm_name_list* hardware clock timezone are not UTC.
Cause	File server VM hardware clock timezone is set to unsupported value
Impact	Services on file server VM may fail due to time drift

Resolution	Refer to KB article 6559. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Warning

Table 340: Nutanix Files Version EOL [160073] [A160073]

Name	File server version EOL check
Description	File server(s) running versions which have reached end-of-life: {fs_eol_msg}
Alert message	File server(s) running versions which have reached end-of-life. *fs_eol_msg*
Cause	End-of-life Nutanix Files version running on cluster.
Impact	Support may not be able to resolve issues since this Nutanix Files version is no longer maintained.
Resolution	Refer to KB article 6601. Contact Nutanix support if issue still persists or assistance needed.
KB Article	6601
Severity	Info

Table 341: File Server upgrade task stuck [160075] [A160075]

Name	File Server Upgrade Task Stuck
Description	File Server Upgrade Task Stuck.
Alert message	Upgrade of file server *file_server_name* is stuck due to *reason*.
Cause	Check alert message for details
Impact	File server has not been upgraded.
Resolution	Refer to KB article 8438 for more details. Contact Nutanix support if issue still persists or assistance needed. Once the issue is resolved, execute 'afs infra.resume_fs_upgrade' on cvm to resume Nutanix Files upgrade.
KB Article	8438
Severity	Critical

Table 342: Files cluster HA takeover failure [160076] [A160076]

Name	File server HA Takeover Failure
Description	Files cluster HA takeover failed
Alert message	Files cluster *file_server_name* HA takeover has failed multiple times on FSVM *fsvm_internal_ip* due to *message*.
Cause	Files cluster HA takeover process failed for multiple attempts
Impact	Files cluster may be unavailable.

Resolution	Refer to KB article 8419 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8419
Severity	Critical

Table 343: File server container deduplication check [160077] [A160077]

Name	File Server Container Deduplication Enabled
Description	Nutanix Files container has deduplication policy enabled which is not recommended
Alert message	File server(s) with deduplication policy enabled on storage container: *fs_dedup_msg*
Cause	On-disk deduplication policy is enabled on Nutanix Files container which is not recommended
Impact	Enabling deduplication policy on file server container may impact storage performance.
Resolution	Refer to KB article 6944. Contact Nutanix support if issue still persists or assistance needed.
KB Article	6944
Severity	Info

Table 344: File server VM down check [160078] [A160078]

Name	File Server VM Reachable
Description	Checks if all the file server VMs are reachable.
Alert message	File server *file_server_name*'s VMs (*svm_names*) are down
Cause	One or more VMs are down on file server.
Impact	File server operations will not work.
Resolution	Refer to KB article 3691. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Warning

Table 345: File Server Clone Grace Period Check [160079] [A160079]

Name	File Server Clone Grace Period Check
Description	Check for licensing of cloned file servers
Alert message	Cloned file server(s) in grace period: *fs_clone_msg*
Cause	One or more cloned file servers are within grace period.
Impact	Cloned file servers will be considered for licensing after end of grace period.

Resolution	Refer to KB article 7865. Contact Nutanix support if issue still persists or assistance needed.
KB Article	7865
Severity	Info

Table 346: File Server Alert [160080] [A160080]

Name	File Server Alert
Description	Files server alert
Alert message	*file_server_alert_msg*
Cause	Check alert summary
Impact	Check alert summary
Resolution	Check alert summary
Severity	Critical
Severity	Warning
Severity	Info

Table 347: File Server Manager Upgrade Failed. [160085] [A160085]

Name	File Server Manager Upgrade Failed.
Description	The File Server Manager upgrade task failed on the cluster.
Alert message	*reason*
Cause	Please check the alert message for details.
Impact	The File Server Manager upgrade failed on the cluster; File Server operations may not work as expected.
Resolution	Please refer to KB 8342 at the Nutanix Portal.
Severity	Warning

Table 348: File Server Container has Unexpected Files [160087] [A160087]

Name	File Server Container Unexpected Files
Description	Nutanix Files container has unexpected files which is not recommended.
Alert message	File server container(s) with unexpected files present: *unexpected_files_msg*
Cause	Unexpected files are present on Nutanix Files container.
Impact	File Server delete operation can fail. Nutanix Files container deletion will also result in deletion of any resident non fileserver VM disks or images.
Resolution	Move any non fileserver VMs out of Nutanix Files container. Refer to KB article 6944. Contact Nutanix support if issue still persists or assistance needed.

KB Article	6944
Severity	Warning

Table 349: File Server Container Managed by Nutanix [160088] [A160088]

Name	File Server Container Nutanix Managed Check
Description	File Server container should not be Nutanix Managed.
Alert message	File Server container(s) which are managed by Nutanix: *fs_nutanix_managed_msg*
Cause	File Server has the flag isNutanixManaged set as True.
Impact	File Server container stats will not be collected.
Resolution	isNutanixManaged flag should be set as false for the File Server Container. Refer to KB article 6944. Contact Nutanix support if issue still persists or assistance needed.
KB Article	6944
Severity	Info

Table 350: Nutanix Files Incompatible AOS/FSM Version [160089] [A160089]

Name	File Server Version Compatibility Check
Description	File server(s) running 4.0.0 versions and above whose compatibility versions failed: {fs_comp_msg}
Alert message	File server(s) running on incompatible AOS or FSM version *fs_comp_msg*
Cause	Older AOS or FSM version running with 4.0.0 or above Files version.
Impact	Infra operations is not supported for 4.0.0 file servers running older AOS or FSM versions.
Resolution	Please upgrade AOS to version 5.15.7 or later or upgrade FSM to version 2.1.0 or the latest.
KB Article	6601
Severity	Warning

Table 351: File Server task hung for too long [160090] [A160090]

Name	File Server Task Stuck Check
Description	File Server Task is in a hung state and system couldn't recover it in last 24 hours.
Alert message	File Server Task hung for too long
Cause	File Server Task running on the Controller VM and/or File Server VM is in a hung state & system is not able to recover from it. It may be due to a software issue or a task conflict.
Impact	File Server hung tasks may cause a degraded system state.

Resolution	Refer to KB article 9072. Contact Nutanix support if issue still persists or assistance needed.
KB Article	9072
Severity	Warning

Table 352: Cassandra status is not normal on few nodes of file server [160092] [A160092]

Name	File Server Cassandra Status Failure
Description	Cassandra status is not normal on few nodes of file server
Alert message	*message*
Cause	See details in the alert message.
Impact	Some file server features may not be functioning
Resolution	Contact Nutanix support.
Severity	Critical

Table 353: Inconsistent state found in File Server [160093] [A160093]

Name	File Server Invalid IDF State
Description	One or more inconsistent entities have been found. This may cause upgrades to fail.
Alert message	One or more inconsistent entities have been found. This may cause upgrades to fail.
Cause	IDF is in an inconsistent state.
Impact	File server upgrades may fail.
Resolution	Contact Nutanix support.
Severity	Critical

Table 354: Post upgrade handler failed [160094] [A160094]

Name	File Server Master One Time Upgrade Failure
Description	Post upgrade handler failed.
Alert message	*message*.
Cause	Master One Time Upgrade failed.
Impact	File server upgrade failed. Some file server features may not be functioning.
Resolution	Contact Nutanix support.
Severity	Critical

Table 355: File Server SMB Share is unavailable [160095] [A160095]

Name	File Server Multi-Protocol Backward Compatibility Failure
Description	File Server SMB Share is unavailable.
Alert message	*message*
Cause	Multi-protocol share backward compatibility failed.
Impact	File Server SMB Share is unavailable.
Resolution	Please retry enabling multi-protocol access, if the issue persists, contact Nutanix support.
Severity	Critical

Table 356: Auto Snapshot Failed [160096] [A160096]

Name	File Server Auto Snapshot Failure
Description	File server auto snapshot failed.
Alert message	*message*
Cause	Auto snapshot was running for a long time. System is overloaded.
Impact	File server auto snapshot failed.
Resolution	Contact Nutanix support.
Severity	Critical

Table 357: Some files are possibly corrupted [160097] [A160097]

Name	File Server Corrupted Files
Description	Some files are possibly corrupted.
Alert message	*message*
Cause	Some files are possibly corrupted for file server.
Impact	Affected data may not be readable until fixed.
Resolution	Contact Nutanix support.
Severity	Critical

Table 358: Some File server share(s) are unavailable [160098] [A160098]

Name	File Server Shares Unavailable
Description	Some File server share(s) are unavailable.
Alert message	*message*
Cause	See details in the alert message.
Impact	Some shares may be unavailable.

Resolution	Contact Nutanix support.
Severity	Critical

Table 359: File server rebalance is aborted [160099] [A160099]

Name	File Server Rebalance Aborted
Description	File server rebalance is aborted
Alert message	*reason_and_resolution_msg*
Cause	See details in the alert message.
Impact	Rebalance is not performed.
Resolution	See details in the alert message. Contact Nutanix support if necessary.
KB Article	7035
Severity	Critical
Severity	Warning

Table 360: File Server VM Boot ISO Disconnected [160100] [A160100]

Name	File Server VM Boot ISO Disconnected
Description	Checks if the File Server VM boot ISO is connected or not
Alert message	File server *file_server_name* : File server VM *vm_name_list* boot ISO are disconnected.
Cause	File Server VM boot ISO is disconnected
Impact	File Server VM will not come up after reboot
Resolution	Refer to KB article 7887. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3691
Severity	Critical

Table 361: File server share is inaccessible or read-only [160102] [A160102]

Name	File Server Share Access Changed
Description	Share has become readonly or inaccessible.
Alert message	share *share_name* on File Server *file_server_name* is *message*
Cause	File analytics has set a filter to set share readonly or inaccessible.
Impact	Share may be inaccessible or read-only.
Resolution	Check File Analytics page for available actions. Contact Nutanix support if problem persists.
Severity	Critical
Severity	Warning

Table 362: File Server internal network is not on the same subnet as Controller VM [160104] [A160104]

Name	File Server Cvm Subnet Check
Description	File Server internal network and Controller VM are in different Subnet
Alert message	File Server *file_server_name* internal network *nvm_subnet* is not on the same subnet as Controller VM *cvm_ip* subnet *cvm_subnet*
Cause	File Server internal network and the Controller VM are in different subnets. This may cause performance issues.
Impact	File Server performance might have degraded.
Resolution	Re-ip the file server's internal network with IP as the same subnet as the Controller VM. Refer to KB article 9845. Contact Nutanix support if issue still persists or assistance needed.
KB Article	9845
Severity	Warning

Table 363: File server protection domain schedule suspended [160115] [A160115]

Name	File Server Protection Domain Schedule Suspended
Description	Checks all File server protection domain schedule is active
Alert message	Some of File server *file_server_name* protection domain *protection_domain_name* schedule is suspended
Cause	Some of the File Server protection domain schedule is suspended
Impact	File server migration might fail and create unhealthy snapshots
Resolution	Refer to KB article 4312. Contact Nutanix support if issue still persists or assistance needed.
KB Article	4312
Severity	Warning

Table 364: File server metadata scan is aborted [160122] [A160122]

Name	File Server Metadata Scan Aborted
Description	File server metadata scan is aborted
Alert message	*message*
Cause	See details in the alert message.
Impact	Metadata scan is not completed.
Resolution	See details in the alert message.
Severity	Warning

Table 365: File server external network not functional. [160123] [A160123]

Name	File Server External IP Failure
Description	One or more of the external IP(s) may not be operational or reachable.
Alert message	*message*
Cause	Most probable cause is either network unreachability on the external interface or core Nutanix files cluster services are down/unstable. See details in the alert message.
Impact	Shares may become unavailable.
Resolution	Ensure that external network is reachable, please refer to KB10434 for more details. Contact Nutanix support if the issue still persists.
KB Article	10434
Severity	Critical

Table 366: File server persistent store volume-group unavailable [160124] [A160124]

Name	File Server Persist Store VG Unavailable
Description	One or more of persist store Volume Group(s) may not be operational.
Alert message	*message*
Cause	Most probable cause can be underneath storage issues on AOS or Files cluster virtual IP connectivity problems. See details in the alert message.
Impact	Some Persist Store Volume Group(s) may be unavailable which can lead to issues with Continuously Available shares.
Resolution	Execute Cluster Health check to ensure there are no underlying AOS issues, please refer to KB10435 for more details. Contact Nutanix support if the issue still persists.
KB Article	10435
Severity	Critical

Table 367: File server Inode usage is high [160125] [A160125]

Name	File Server Inode Usage High
Description	The number of free Inodes in the cluster is getting low. Inode contains information about data files attributes.
Alert message	*message*
Cause	There can be multiple causes including overall high storage utilization. See details in the alert message.
Impact	The system may experience unavailability.
Resolution	Execute Cluster Health check to ensure there are no underlying storage high utilization issues, please refer to KB10436 for more details. Contact Nutanix support if the issue still persists.
KB Article	10436

Severity	Critical

Table 368: File server iSCSI service is not healthy [160126] [A160126]

Name	File Server iSCSI Service Unavailable
Description	iSCSI service is not in an healthy state on the cluster and can lead to storage unavailability.
Alert message	*message*
Cause	iSCSI service is in a hung state or not in a functional state.
Impact	Shares may become unavailable.
Resolution	Execute Cluster Health check to ensure there are no underlying cluster issues, please refer to KB10437 for more details. Contact Nutanix support if the issue still persists.
KB Article	10437
Severity	Critical

Table 369: File Server Protection Domain to Container Mapping Incorrect [160132] [A160132]

Name	File Server Protection Domain to Container Mapping Incorrect
Description	File server(s) protection domain are mapped to incorrect container.
Alert message	File server protection domains have incorrect container mapping: *fs_fail_msg*
Cause	Missing or invalid protection domain container mapping.
Impact	Files data protection capability might be affected.
Resolution	Correct the protection domain container mapping. Refer to KB article 5146. Contact Nutanix support if issue still persists or assistance needed.
KB Article	5146
Severity	Critical

Table 370: File Server is not reachable at required ports from the platform. [160133] [A160133]

Name	File Server Cvm Port Check
Description	File Server is not reachable at required ports from the platform.
Alert message	Unreachable File server(s): *fs_port_msg*
Cause	Either File server is not reachable from the platform OR File Server ports are blocked by firewall.
Impact	File Server services are impacted.
Resolution	Check file server status. Ensure the File Server Port Requirements are met entirely, Please refer to the Nutanix Files User Guide Files Deployment section for port requirements. Refer to KB article 9845. Contact Nutanix support if issue still persists or assistance needed.

KB Article	9845
Severity	Critical

Table 371: Platform is not reachable at required ports from the fileserver VM. [160134] [A160134]

Name	File Server Cvm Port Unreachable
Description	Platform is not reachable at required ports from the fileserver VM.
Alert message	*message*
Cause	Either Platform services are down OR Platform ports are blocked by firewall.
Impact	File Server services are impacted.
Resolution	Check Platform status. Ensure the File Server Port Requirements are met entirely, Refer Nutanix Files User Guide. Refer to KB article 11022. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11022
Severity	Critical

Table 372: File server slog partitioning failed. [160140] [A160140]

Name	File Server Slog Partition Failure
Description	File server slog partition failure for volume group.
Alert message	*message*
Cause	Device partition creation failed.
Impact	Share write performance is impacted.
Resolution	Refer to KB article 11400. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11400
Severity	Critical

Table 373: File server slog update task failed. [160141] [A160141]

Name	File Server Slog Update Task Failure
Description	File server slog update task failure for volume groups.
Alert message	*message*
Cause	Storage update with slog device failed.
Impact	File server is in heterogenous state.
Resolution	Refer to KB article 11412. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11412

Severity	Critical

Table 374: IDF Source to Remote Sync Status [200001] [A200001]

Name	IDF DB to DB sync heartbeat check
Description	Checks if IDF replica is receiving requests from IDF master.
Alert message	IDF data on this cluster is not in sync with the IDF data on cluster:*cluster_name*
Cause	Cluster network connectivity or cluster services such as insights server, insights uploader, insights receiver, Aplos, or Prism gateway could be down.
Impact	IDF data that is replicated from master cluster is not up to date.
Resolution	Ensure that cluster network connectivity is up and all cluster services are up.
KB Article	5582
Severity	Warning

Table 375: Dynamic scheduling failure. [200201] [A200201]

Name	Acropolis Dynamic Scheduler Status
Description	One or more nodes have resource contention. This imbalance can cause performance bottlenecks on the node(s) affected.
Alert message	Dynamic scheduling failure. *reason*
Cause	Cluster may have insufficient CPU or Controller VM resources.
Cause	Cluster may have insufficient resources to satisfy VM group affinity policies or VM host affinity policies.
Impact	Cluster performance may be significantly degraded.
Resolution	Shut down unneeded VMs to free cluster resources.
Resolution	Expand the cluster to add resources.
Severity	Warning

Table 376: Policy not applicable to any Host. [200305] [A200305]

Name	No viable Hosts to place VMs in PE
Description	Policy as defined does not apply to any of the Hosts.
Alert message	Rule *rule_uuid* applies to no Host in cluster *cluster_uuid*
Cause	Addition of a new policy(s) or update to existing policy(s) or Attaching/ Detaching categories to Hosts.
Impact	VMs will not have any Hosts to run on.
Resolution	Make appropriate updates to affinity policies or Host categories in the affected PE.
Severity	Warning

Table 377: VM forcibly powered off. [200401] [A200401]

Name	VM forcibly powered off
Description	We have detected that VM {vm_name} (vm_id: {vm_id}) was running on multiple hosts. To avoid data corruption this VM was forcibly powered off on host {host_address}. {reason}
Alert message	VM named *vm_name* with id *vm_id* has been forcibly powered off on host *host_address*
Cause	Unexpected restart of Acropolis service during VM migration.
Cause	If VM does not have local disks, but using network-attached storage, it will be forcibly powered off to prevent data corruption on network-attached storage.
Impact	If the VM does not have any disks on Nutanix storage and is only using network-attached storage, data corruption on network-attached storage is possible.
Resolution	Please collect the log bundle and engage Nutanix support.
Resolution	Check the integrity of data located on network-attached storage that is in use by this VM.
Severity	Warning

Table 378: Power cycle VMs before performing upgrade or migrate operation [200404] [A200404]

Name	Power cycle VMs before performing upgrade or migrate operation
Description	Power cycle all the VMs before you perform the upgrade or migrate operation.
Alert message	Power cycle the VMs before you perform the upgrade or migration operation
Cause	A node with a lower CPU feature set has been added to the cluster.
Impact	You will not be able to migrate the existing VMs to newly added node(s). Operations that require VM migration, including but not limited to hypervisor upgrade, will fail.
Resolution	Ensure that all the VMs in the cluster are power cycled.
Severity	Warning

Table 379: Report Generation Failure [200501] [A200501]

Name	Report Generation Failure
Description	Report Generation failed. Database may be down or cluster may have insufficient storage.
Alert message	Report generation failed for report *report_name* scheduled for *timestamp*. *reason*
Cause	Cluster may have insufficient storage.
Impact	Some of the reports can be missed.
Resolution	(A) Free the space by deleting older reports that have been generated. (B) Manually run the report instance scanner utility.
Severity	Warning

Table 380: Send Report Through E-mail Failure [200502] [A200502]

Name	Send Report Through E-mail Failure
Description	Sending of report through E-mail failed. SMTP might be down or report might have been deleted.
Alert message	Sending of report through E-mail failed for report *report_name*. *reason*
Cause	(A) SMTP server might be unreachable or might not have been configured. (B) Generated report may have been deleted.
Impact	Some of the report emails can be missed.
Resolution	Check the SMTP configuration and ensure that its reachable.
Severity	Warning

Table 381: Report Quota Scan Failure [200503] [A200503]

Name	Report Quota Scan Failure
Description	Report Quota scan failed. Cassandra might be unreachable.
Alert message	Report quota scan failed for the report *report_name* at *timestamp*. *reason*
Cause	Cassandra may not be reachable for some time.
Impact	There might be some stale reports on the cluster that are taking unnecessary space on the cluster.
Resolution	Manually run the report instance scanner utility.
Severity	Warning

Table 382: VPC VMs detected in Flow Security Policy [200608] [A200608]

Name	VPC VMs detected in Flow Security Policy
Description	VMs that have NICs in one or more VPCs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for these VMs.
Alert message	*num_vms* VPC VM(s) detected in Flow Security Policy *rule_name* (*network_security_rule_uuid*). VMs: *vms*. Flow Security Policies will not be enforced for the affected VMs.
Cause	VMs with NICs in VPC subnets have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the affected VMs from VPC subnets for Flow Security Policies to be enforced on these VMs.
Severity	Warning

Table 383: VMs with kDirectNics detected in Flow Security Policy [200609] [A200609]

Name	VMs with kDirectNics detected in Flow Security Policy

Description	VMs that have one or more kDirectNICs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for the traffic through these kDirectNics.
Alert message	*num_vms* VM(s) with kDirectNics detected in Flow Security Policy *rule_name* (*network_security_rule_uuid*). VMs: *vms*. Flow Security Policies will not be enforced for the traffic through these kDirectNics.
Cause	VMs with kDirectNics have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the kDirectNics from the vms for Flow Security Policies to be enforced.
Severity	Warning

Table 384: Availability Zone Connection Failure [200801] [A200801]

Name	Availability Zone Connection Failure
Description	The remote availability zone {az_url} is unreachable.
Alert message	The remote availability zone *az_url* is unreachable.
Cause	(A) Services on remote site might be down. (B) Connection might be missing on remote site. (C) Access tokens might have expired.
Impact	Disaster recovery operations might fail due to unreachability of remote site.
Resolution	Unpair from the remote availability zone, if already paired. Pair with the remote availability zone again.
KB Article	8252
Severity	Critical

Table 385: PE-PC Connection Failure [200802] [A200802]

Name	PE-PC Connection Failure
Description	The remote {remote_vm_type} {cluster_ip} is unreachable.
Alert message	The remote *remote_vm_type* *cluster_ip* is unreachable.
Cause	The name servers might be missing.
Cause	Services on PC/PE might be down.
Cause	Access tokens might have expired.
Impact	Management tasks for various entities, including VM, might fail due to unreachability of remote {remote_vm_type}.
Resolution	Ensure that the appropriate name servers are configured.
Resolution	Check that no services are crashing on the PC or PE.
Resolution	Contact Nutanix Support to troubleshoot this issue further.
Severity	Warning

Table 386: VM Recovery Storage Container(s) is/are not mounted on all hosts. [300431] [A300431]

Name	VM Recovery Storage Container Mount Check
Description	Verify that the Storage Container(s) is/are mounted on all the hosts.
Alert message	Container(s) *container* part of recovery point of entities, is/are not mounted on all the hosts.
Cause	Storage Container(s) is/are not mounted on all the hosts.
Impact	Registration of ESX VM(s) will fail during data recovery workflows.
Resolution	Mount the Storage Container(s) on all the hosts.
KB Article	7911
Severity	Warning

Table 387: Entity Sync Failure [500101] [A500101]

Name	Entity Sync failed
Description	Entity Sync failed
Alert message	*message*
Cause	An update made to the entity locally might not have synced to the connected Availability Zone(s).
Cause	One or more internal services might be down or not working as expected.
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Resolution	Contact Nutanix support.
Severity	Warning

Table 388: Entity Sync Failure for the Protection Policy [500102] [A500102]

Name	Entity Sync failed for Protection Policy
Description	Entity Sync failed for the Protection Policy
Alert message	*message*
Cause	An update made to the Protection Policy locally might not have synced to the connected Availability Zone(s).
Cause	One or more internal services might be down or not working as expected.
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Resolution	Contact Nutanix support.
Severity	Warning

Table 389: Entity Sync Failure for the Recovery Plan [500103] [A500103]

Name	Entity Sync failed for Recovery Plan
Description	Entity Sync failed for the Recovery Plan
Alert message	*message*
Cause	An update made to the Recovery Plan locally might not have synced to the connected Availability Zone(s).
Cause	Protection Policy configuration for the entities specified in the Recovery Plan may be different across connected Availability Zones.
Cause	One or more internal services might be down or not working as expected.
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue persists, contact Nutanix support.
Resolution	Ensure that the Protection Policy configuration is same across connected Availability Zones for the entities specified in the Recovery Plan.
Resolution	Contact Nutanix support.
Severity	Warning

Table 390: Entity Sync Failure for the Availability Zone [500104] [A500104]

Name	Entity Sync failed for Availability Zone
Description	Entity Sync failed for the Availability Zone
Alert message	*message*
Cause	Unable to send request to the Availability Zone.
Cause	Request timed out.
Impact	Data Protection and Recovery might be impacted.
Resolution	Ensure Availability Zone is reachable and do a 'Forced Entity Sync' to ensure that all the entities are in sync. If the issue still persists, contact Nutanix support.
Resolution	Ensure there are no sync related tasks running under the 'Tasks' page on all the connected Availability Zone(s) and then invoke 'Forced Entity Sync' from the chosen site.
KB Article	10241
Severity	Critical

Table 391: Translated Addresses Retrieval Failure in NGT [600102] [A600102]

Name	Unable to get translated addresses in NGT
Description	NGT is unable to retrieve the XAT address translations.
Alert message	Unable to retrieve translated addresses in NGT because the Prism Element IDF service may be unreachable.

Cause	(A) Prism Element IDF may be unreachable. (B) Translation information is not available in IDF.
Impact	NGT features may not work on the User VMs.
Resolution	Check if the Prism Element IDF service is reachable.
Severity	Warning

Table 392: Data provider collector is in crashloop [650000] [A650000]

Name	Data provider collector is in crashloop
Description	Data provider collector is restarting frequently
Alert message	*data_provider_name* data collector with uuid *metrics_data_provider_instance_uuid* is in crashloop
Cause	Collector heartbeat may not be reaching Prism Central
Impact	Users may notice missing stats for some of the endpoints that are configured through monitoring integration feature.
Resolution	Manually attempt to start mercury service on all Prism Central VMs. if situation persists, or any assistance is needed please contact Nutanix Support.
KB Article	9258
Severity	Warning

Table 393: Data provider collector has violated sizing limits [650001] [A650001]

Name	Data provider collector has sizing violation
Description	Data provider collector metrics usage has violated sizing limits
Alert message	*data_provider_name* data collector has reached *threshold*% of the sizing limits.
Cause	Total no. of instances or Metric collection has reached the supported limit by Data Provider Collector
Impact	Users may not be able to configure new instances through monitoring integration feature.
Resolution	Delete some existing instance to configure new instances if situation persists, or any assistance is needed please contact Nutanix Support.
KB Article	9473
Severity	Critical
Severity	Warning

Table 394: Cluster is not healthy [802004] [A802004]

Name	Failure to configure a cluster for Advanced Networking
Description	Failure to configure a cluster for Advanced Networking.
Alert message	Cluster (*cluster_uuid*) failed to be configured for Atlas Networking

Cause	Cluster cannot accept Advanced Networking Configuration.
Impact	Ability to make network-related configurations may be impacted.
Resolution	(A) Verify health of cluster. (B) Verify cluster is reachable from Prism Central. (C) Verify cluster AOS and AHV version are compatible with Advanced Networking.
Severity	Warning

Table 395: Security Planning is disabled [803001] [A803001]

Name	Security Planning is disabled
Description	Security Planning is disabled due to de-registration of Prism Element from Prism Central.
Alert message	Security Planning is disabled. Reason: *details*
Cause	The cluster that is hosting the Epoch data collector was de-registered from Prism Central.
Impact	The Security Planning feature has been automatically disabled.
Resolution	Enable Security Planning from the Prism Central UI to deploy a new instance of the Epoch data collector.
Severity	Warning

Table 396: ID Firewall lost connectivity to domain controller [803003] [A803003]

Name	ID Firewall is unable to connect to one or more domain controllers
Description	ID Firewall is unable to connect to one or more domain controllers used for scraping identities.
Alert message	Domain controller *domain_controller* on domain *domain* is not reachable and accepting LDAP or WMI connections. Details: *details*.
Cause	The domain controller may not be reachable via the network.
Cause	The domain controller may not be accepting LDAP/WMI connections due to permissions issues.
Cause	The domain controller may not be running or fully booted up.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until access to the domain controller is restored.
Resolution	Check network connectivity from Prism Central to the domain controller. Please refer article KB-10219 for more details.
Resolution	Check that the domain controller is fully up and accepting WMI and LDAP connections. Please refer article KB-10219 for more details.
Resolution	Check that the service account registered for ID Firewall is active and has permissions for both WMI and LDAP access, and that the correct password is entered in Prism Central. Please refer article KB-10219 for more details.
KB Article	10219

Severity	Warning

Table 397: ID Firewall unable to locate mapped Active Directory object [803004] [A803004]

Name	ID Firewall mapped object not found
Description	ID Firewall was unable to find a mapped object in Active Directory.
Alert message	The mapped *object_type* *object_dn* with GUID *object_guid* was not found in *domain*.
Cause	The mapped object in Active Directory which is referenced for ID Firewall could not be found because it was deleted from Active Directory.
Impact	Flow VDI security policies written around the mapped object will not function as expected.
Resolution	The mapping should be checked, and if the mapped object does not exist, or if it has been deleted and re-created, the mapping must also be deleted and possibly re-created. Security policies should be updated to use the new mapping or not reference the mapping at all if the object was permanently deleted.
Severity	Warning

Table 398: ID Firewall did not recover state after reconnecting to a temporarily unreachable domain controller [803005] [A803005]

Name	ID Firewall state recovery failed
Description	ID Firewall was unable to recover state after reconnecting to a temporarily unreachable domain controller.
Alert message	ID Firewall was unable to recover state from domain controller *domain_controller* on domain *domain*. Details: *details*.
Cause	Recovery failed because either too much time has passed since the domain controller was last reachable, or the event log on the domain controller has rolled over and the events necessary to recover are no longer present.
Impact	The state enforced by ID Firewall may be inconsistent, and therefore applied policies may not be properly enforced until action is taken.
Resolution	In order to ensure security policies are properly enforced, all active Nutanix VDI VM users should log out and log back into their VMs. Please refer KB-10220 for further details.
KB Article	10220
Severity	Warning

Table 399: Epoch Data Collector upgrade available [803006] [A803006]

Name	Epoch Data Collector needs upgrade
Description	Checks whether there is an upgrade available for the Epoch Data Collector.
Alert message	*message*

Cause	A new version of the Epoch Data Collector is available and not yet installed.
Impact	Some new functionality or bugfixes in the Epoch Data Collector may not be available.
Resolution	Go to Epoch Data Collector Enablement page to upgrade the Epoch Data Collector.
Severity	Info

Table 400: ID Firewall service account is invalid [803007] [A803007]

Name	Service account used for ID Firewall is invalid
Description	ID Firewall has detected that the service account configured is not valid.
Alert message	ID Firewall service account *account*@*domain* is invalid. Details: *details*.
Cause	The service account password may have been changed.
Cause	The service account may have been deleted.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until the service account is valid again.
Resolution	Check that the configured service account password is up-to-date and update it in Prism if needed.
Resolution	Check that the configured service account exists, and update it in Prism if needed.
Severity	Warning

Table 401: Robo witness configured status [804010]

Name	Robo witness configured status
Description	Checks if the robo two node cluster has witness configured
Cause	Robo cluster has no witness configured
Impact	Cluster will not tolerate any failure.
Resolution	Configure witness by using the NCLI command 'cluster add-witness'
KB Article	8562

Table 402: Robo cluster witness current state [804011]

Name	Robo cluster witness current state
Description	Checks the current leader node and cluster operation mode for the robo cluster
Cause	Robo cluster may have been through cluster transitions and recommendations may be needed based on the current state
Impact	Robo cluster may not tolerate another failure depending on the current cluster state

Resolution	Ensure that the robo cluster is in Normal mode with no leadership
KB Article	8563

Table 403: Robo cluster read only state [804012]

Name	Robo cluster read only state
Description	Check if the robo cluster current state is Read Only.
Cause	The cluster is running in single node mode and has only one functional SSD
Impact	Data resiliency is degraded, and user VM writes are blocked.
Resolution	Make sure there are at least 2 working SSDs, and let the cluster transition out of read-only mode
KB Article	8564

Table 404: Robo cluster witness sync with witness VM [804013]

Name	Robo cluster witness sync with witness VM
Description	Check if the witness state in robo cluster and witness VM are in sync.
Cause	Network unavailability during robo cluster transition could have caused witness to be out of sync with the witness VM
Impact	Robo cluster cannot handle all future failures, transitions and planned outages.
Resolution	Contact Nutanix support for resolution
KB Article	8565

Table 405: Robo cluster nodes ping to witness VM [804014]

Name	Robo cluster nodes ping to witness VM
Description	Check that both nodes on robo cluster are able to ping the witness VM configured to the cluster.
Cause	Network failures either at robo cluster or at the witness VM will cause the pings to fail.
Impact	Robo cluster may not handle any failures ,transitions or planned outages.
Resolution	Look for any network issues at robo cluster or the witness VM site and resolve the network issues for the pings to succeed.
KB Article	8566

Table 406: Robo cluster nodes ping to each other [804015]

Name	Robo cluster nodes ping to each other
Description	Check that both nodes on robo cluster are able to ping each other.

Cause	Network failures/outages will cause pings to fail to between the nodes
Impact	Robo cluster will transition to single node
Resolution	Look for any network issues at robo cluster site, resolve the network issues for the pings to succeed and cluster to transition to Normal state with no leadership
KB Article	8567

Controller VM

Table 407: CVM Boot Raid Degraded [1019] [A1107]

Name	Boot RAID Health
Description	Checks the health of boot RAID volumes.
Alert message	Boot RAID volumes on CVM *service_vm_external_ip* are degraded.
Cause	RAID volume may be in an abnormal state due to power loss or temporary or permanent drive failure.
Impact	This might be due to a boot drive failure. The CVM might become inoperable if the issue is not resolved.
Resolution	Verify boot disk status from the web console Hardware page and replace the disk if required.
KB Article	2463
Severity	Critical

Table 408: Abnormal Host Boot RAID State [1020] [A1105]

Name	Host Boot RAID State
Description	Checks the health of boot RAID volumes.
Alert message	Hypervisor RAID volume '*volume_label*' is in abnormal state '*volume_state*' on host *host_ip*.
Cause	RAID volume may be in an abnormal state due to power loss or temporary or permanent drive failure.
Impact	This might be due to a boot drive failure. The CVM might become inoperable if issue is not resolved.
Resolution	Use remote console to verify the RAID status to decide if replacement of drive is required.
KB Article	2463
Severity	Critical

Table 409: CVM or Prism Central VM RAM Usage High [3023] [A1056]

Name	CVM Memory Usage

Description	Check that CVM or Prism Central VM memory usage is not high.
Alert message	Main memory usage in *vm_type ip_address* is high, *available_memory_kb* KB is free.
Cause	The RAM usage on the Controller VM or Prism Central VM has been high.
Impact	Cluster performance may be significantly degraded.
Resolution	Check the memory utilization of Prism Central VM or Controller VM. If abnormal behavior is seen, please collect logs and contact Nutanix Support.
CVM Available Memory Threshold	Critical Threshold: 589824 KB
KB Article	2473
Severity	Critical

Table 410: CVM/PCVM time not synchronized with external servers. [3026] [A3026]

Name	CVM/PCVM NTP Time Synchronized
Description	Checks that the {vm_type} is synchronizing time with an NTP server.
Alert message	*alert_msg*
Cause	External NTP servers are not configured or are not reachable
Impact	Workflows involving Kerberos may fail if the time difference between the Controller VM and the NTP server is greater than 5 minutes.
Resolution	Please verify that the external NTP servers are configured and are reachable.
KB Article	4519
Severity	Warning

Table 411: CVM Rebooted [3028] [A1024]

Name	CVM Rebooted Check
Description	Checks if the CVM is not rebooted recently.
Alert message	*vm_type ip_address* has been rebooted on *reboot_timestamp_str*.
Cause	CVM is rebooted.
Impact	During the time the Controller VM is down, cluster compute and storage capacity are reduced.
Resolution	Check CVM status. Ignore if intentional.
KB Article	2474
Severity	Critical

Table 412: CVM Service(s) Restarting Frequently [3029] [A1032]

Name	CVM Services Status

Description	Check if services have restarted recently on the CVM.
Alert message	One or more cluster services have restarted within 15 minutes in the Controller VM *ip_address*, the total restarts being >= 10.The following are the service(s) that have restarted : *failed_components*.
Cause	Faulty behavior in one or more service(s), leading to frequent crashes,in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs more than once, please collect CVM logs and contact Nutanix support.
KB Article	2472
Severity	Critical

Table 413: {services} facing segmentation fault [3036] [A3036]

Name	CPP Services Segmentation fault Check
Description	Check if services faced a segmentation fault recently in PCVM/CVM
Alert message	More than 5 segmentation fault errors happened in the last 15 minutes in the PCVM/CVM *ip_address*. Services: *services*
Cause	Faulty behavior in one or more service(s).
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this occurs more than 1 time, please refer to KB 2472 for more information. If unsure, please contact Nutanix Support.
KB Article	2472
Severity	Critical

Table 414: CVM Python Service(s) Restarting Frequently [3037] [A3037]

Name	Python Services Fatal Check
Description	Check if python services have crashed/restarted recently in PCVM/CVM
Alert message	One or more cluster services have restarted within 15 minutes in the PCVM/CVM *ip_address*, the total restarts being >= 10. The following are the service(s) that have restarted : *failed_components*.
Cause	Faulty behavior in one or more service(s) leading to frequent crashes in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs more than once, please contact Nutanix support.
KB Article	2472

Severity	Warning

Table 415: Baseline configuration changes [3039] [A3039]

Name	AIDE baseline configuration changes check
Description	Checks the log files of aide for changes to baseline configs if aide is enabled
Alert message	*message*
Cause	Baseline configurations changed
Impact	Aide execution may fail depending on the files that were changed
Resolution	Check changed files and revert their changes if needed
KB Article	10461
Severity	Warning

Table 416: CVM CPU Utilization [3041]

Name	CVM CPU Utilization
Description	Checks whether CPU usage is high on CVM.
Cause	CPU usage is high.
Impact	High I/O latency may be experienced by some workloads.
Resolution	Reduce CPU-intensive processes.
KB Article	2733

Table 417: Multiple cpu_unblock processes running [3042] [A3042]

Name	cpu_unblock Process Check
Description	Check that there are no stale cpu_unblock processes running.
Alert message	Multiple cpu_unblock processes are running on *svm_ip*.
Cause	Zookeeper restarting frequently on cluster node.
Impact	Cluster performance may be significantly degraded.
Resolution	Kill all cpu_unblock processes and restart cluster services on node.
KB Article	5581
Severity	Critical
Severity	Warning

Table 418: Jumbo Frames Enabled for NIC {nic_name} on {service_vm_external_ip} [3063] [A1095]

Name	NIC MTU Configuration
Description	Check MTU of the CVM network interfaces.
Alert message	Controller VM *service_vm_external_ip* has NIC *nic_name* with MTU set to *mtu* instead of *desired_mtu*.
Cause	MTU is not set correctly on eth0/eth1
Impact	If NIC MTU is not set correctly on the CVM, services may be unstable or the system may not perform optimally.
Resolution	Set MTU to the right value
KB Article	2486
Severity	Critical

Table 419: CVM Connectivity Failure [3064] [A1001]

Name	CVM Passwordless Connectivity
Description	Check that CVM has passwordless connection to each other.
Alert message	Controller VM *target_ip* is not reachable from Controller VM *source_ip* in the last 3 attempts.
Cause	Controller VM is not running or is in a hung state.
Cause	Controller VM does not have the correct network configuration.
Cause	Upstream TOR switch has configuration or status change on the port connected to the CVM.
Cause	Virtual switch in the Hypervisor is not working correctly.
Cause	SSH Keys are set up incorrectly.
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Ensure the CVM is powered on and running. If the CVM is in a hung state, reference KB 1526 for further guidance.
Resolution	Correct the network configuration settings of the CVM.
Resolution	Ensure that the port connected to the CVM is up and in a forwarding state.
Resolution	Contact Nutanix Support.
Resolution	Contact Nutanix Support.
KB Article	1526
Severity	Critical

Table 420: MTU configuration across Controller VMs is not consistent. [3068] [A3068]

Name	NIC MTU Ping Check

Description	Check connectivity between Controller VMs by pinging with configured MTUs.
Alert message	Failed to ping other CVMs with MTU set to *mtu* on Controller VM *service_vm_external_ip*.
Cause	MTU is not consistent across Controller VMs and host external interfaces, or upstream network.
Impact	An inconsistent MTU configuration can lead to network instability across Controller VMs and hosts, which in certain conditions can lead to storage availability issues.
Resolution	Set MTU consistently across Controller VM and host external interfaces. Ensure upstream physical network MTU supports the configuration.
KB Article	2486
Severity	Critical

Table 421: NIC MTU Uniformity Check [3069]

Name	NIC MTU Uniformity Check
Description	Check if the MTU setting is the same on the matching external interface across every CVM in the cluster.
Cause	MTU is not uniform across the CVMs' management interface
Impact	If MTU is not the same across Controller VM interfaces which are in the same layer 2 broadcast domain or VLAN, some packets could be dropped which could cause the storage to be unavailable.
Resolution	Set MTU on the matching external interface uniformly across the CVMs in the cluster
KB Article	6442

Table 422: CVM NIC Speed Low [6008] [A1058]

Name	TenGig Adaptor Status
Description	Checks whether CVM is uplinked to 10GbE NIC.
Alert message	*alert_msg*
Cause	CVM not uplinked to 10GbE NIC.
Cause	1GbE NIC is part of the bond.
Cause	If hypervisor is ESXi, it may be in lockdown.
Impact	The Controller VM is not configured to use the 10GbE NIC or is configured to share load with a slower NIC.
Resolution	Check network configuration. Please refer to KB 1584.
Resolution	Check network configuration.
Resolution	If hypervisor is ESXi, check for Lockdown Mode enabled and disable the configuration.
KB Article	1584

Severity	Warning

Table 423: CVM Host Subnet Mismatch [6202] [A1048]

Name	CVM Subnet
Description	Checks that host and cvm share the same subnet.
Alert message	Controller VM *svm_ip* with network address *svm_subnet* is in a different network than the Hypervisor *hypervisor_ip*, which is in the network *hypervisor_subnet*.
Cause	The Controller VM and the hypervisor are not on the same subnet.
Impact	IO Redirection in the event of CVM failure will not work as expected, leading to guest VM unavailability.
Resolution	Reconfigure the cluster. Either move the Controller VMs to the same subnet as the hypervisor hosts or move the hypervisor hosts to the same subnet as the Controller VMs.
KB Article	1515
Severity	Critical

Table 424: Node Avg Load - Non-Critical [6511]

Name	Node Avg Load - Non-Critical
Description	Checks whether the system load average is high.
Cause	I/O workload on the cluster is high.
Impact	High I/O latencies may be experienced by some workloads.
Resolution	Redistribute VMs to reduce load on the affected node. Tune the applications to reduce CPU demand. Reduce the number of VMs on the host.
KB Article	3469

Table 425: HDD IO Latency [6514]

Name	HDD IO Latency
Description	Verify whether the average time for I/O requests to the HDDs managed by the Controller VM is not high.
Cause	I/O workload on the cluster is high.
Impact	High I/O latencies may be experienced by some workloads.
Resolution	Redistribute VMs to reduce load on the affected Controller VM.
KB Article	1637

Table 426: CPU Average Load Critically High on Controller VM. [6516] [A6516]

Name	Node Avg Load - Critical
Description	Checks whether the system load average is critically high.
Alert message	Average CPU load on Controller VM *ip_address* is critically high (above critical average load threshold value *critical_high_value*)
Cause	Critically high load on the system.
Impact	Low performance may be experienced due to cluster-wide I/O latency.
Resolution	Contact Nutanix Support for assistance.
KB Article	3469
Severity	Critical

Table 427: CPU Average Load critically high on CVM. [6518] [A6518]

Name	Node Avg Load - Very High
Description	Checks whether the system load average is critically high.
Alert message	*alert_msg*
Cause	Critically high load on the system for quite a long time
Impact	Low performance may be experienced due to cluster-wide I/O latency.
Resolution	Reboot the CVM experiencing high load average.
CVM Load average very critical threshold	Critical Threshold: 100000
KB Article	3469
Severity	Critical

Table 428: Stargate's memory fragmentation check [20023]

Name	Stargate's memory fragmentation check
Description	Checks if stargate's memory fragmentation should be below certain threshold.
Cause	If memory fragmentation is high, this will result in lower cache hits and hence degrade performance.
Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix Support for diagnosis.
KB Article	6735

Table 429: High oplog index memory usage [20031] [A20031]

Name	Oplog index memory usage check

Description	Check if the oplog index memory usage is approaching its limit on CVM.
Alert message	Oplog index memory usage is very high compared to the amount of CVM memory, consider increasing the CVM memory to see improved performance.
Cause	Oplog index memory usage high.
Impact	Makes the oplog draining more aggressive, thereby increasing the load on the system.
Resolution	Increase the CVM memory.
KB Article	10714
Severity	Warning

Table 430: Single active root partition check [101063]

Name	Single active root partition check
Description	Verifies that there only one active root partition in the CVM.
Cause	The storage partitioning on the CVM may have been modified manually after installation. A normal CVM should have one active root partition.
Impact	AOS upgrade will fail.
Resolution	Contact Nutanix Support.
KB Article	6459

Table 431: CVM Boot Raid Degraded [101064] [A101064]

Name	Root partition health check
Description	Checks the health of boot RAID volumes.
Alert message	Boot RAID volumes on CVM *service_vm_external_ip* are degraded.
Cause	RAID volume may be in an abnormal state due to power loss or temporary or permanent drive failure.
Impact	This might be due to a boot drive failure. The CVM might become inoperable if the issue is not resolved.
Resolution	Verify the boot disk status from the PRISM web console and contact Nutanix Support if you suspect a device failure.
KB Article	6459
Severity	Critical

Table 432: RAID health checks [101065]

Name	RAID health checks
Description	Checks the health of all the RAID devices in the CVM.
Cause	One or more storage devices on the node are failing.

Impact	AOS upgrade will fail.
Resolution	Contact Nutanix Support.
KB Article	6459

Table 433: Core dumps are enabled on this CVM or Prism Central VM. [101069] [A101069]

Name	Saltstate coreoff check
Description	Checks if saltstate is correctly set to default coreoff state to ensure services do not produce core dumps unnecessarily.
Alert message	Core dumps are enabled on the *vm_type* with IP *service_vm_external_ip*.
Cause	Core dumps are enabled for services running on this CVM or Prism Central VM.
Impact	Core dumps are useful for isolated troubleshooting but leaving enabled may unnecessarily fill the node /home system partition, and this can affect node stability and availability.
Resolution	Follow KB7586 to validate and fix coreoff saltstate. Resolution may involve graceful node reboot, so ensure the cluster is otherwise healthy and resiliency status is OK prior to this action.
KB Article	7586
Severity	Critical

Table 434: SED Keys Unavailable [101071] [A101071]

Name	SED Keys Unavailable
Description	Check to see if the KMS actually has the expected passwords for the SED disks.
Alert message	SED keys for one or more disks are not available from external key manager *kms_name*.
Cause	If a 3rd party KMS gets restored from a backup or other incorrect management operation, it may have missing keys.
Impact	Cluster will have data unavailability if rebooted.
Resolution	Contact Nutanix and 3rd party key manager support.
KB Article	8223
Severity	Critical

Table 435: Kernel crash dump found. [101073] [A101073]

Name	Kernel crash dump check
Description	Checks if Kernel crash dump logs exists on a node.
Alert message	*alert_msg*
Cause	Kernel dump logs are present on the node.

Impact	Neglecting old kernel crash dump files may unnecessarily fill the system partition, in turn affecting node stability.
Resolution	Please refer to KB 11831 to validate core dump.
KB Article	11831
Severity	Warning

Table 436: SED Disk Configuration Misconfigured [101080] [A101080]

Name	Disk configuration v/s intent consistency check
Description	Check if SED disk configuration is consistent with the intent
Alert message	Either the cluster had SED protection enabled and the disks did not have a password set OR the cluster did not have SED protection enabled and disks had password set OR found tentative passwords.
Cause	Inconsistency in disk's zeus configuration and the intent
Impact	Data unavailability requiring SRE intervention.
Resolution	This issue is usually seen when the KMS is switched off when the disks are still in protected state. Please check whether the KMS is on or off. If the issue persists, contact Nutanix Support.
KB Article	11322
Severity	Critical
Severity	Warning

Table 437: Reset Current Passwords for admin and nutanix SSH accounts on {cluster_name} [101081] [A101081]

Name	Password Strength Status
Description	High-strength password policy has been enabled, please reset the password.
Alert message	High-strength password was recently enabled, please reset current passwords for admin and nutanix SSH accounts on *cluster_name* - *cluster_uuid*.
Cause	High-strength password policy has been enabled.
Impact	High-strength password policy has been enabled but low-strength passwords may still be in use on the cluster.
Resolution	Please reset the current passwords. Follow the Security Guide's Hardening Instructions on Nutanix Support Portal for the password requirements.
Severity	Warning

Table 438: Host time not synchronized with any external servers. [103090] [A103090]

Name	Host time synchronized

Description	Checks to ensure that the host is configured and synchronizing time with an NTP server.
Alert message	*alert_msg*
Cause	External NTP servers are not configured or are not reachable.
Impact	Logs may have different timestamps in the host and the CVMs/PCVM. Host may not work as expected.
Resolution	Please verify that the NTP servers are configured and are reachable from the host.
KB Article	4519
Severity	Warning

Table 439: Validate network configuration for whitespace [103096]

Name	Validate network configuration for whitespace
Description	Check if Network configuration files (ifcfg-ethX and /etc/hosts) are malformed, with whitespaces inserted.
Cause	Network configuration files are malformed, with whitespaces inserted.
Impact	Genesis may crash, causing services to restart.
Resolution	Follow the instructions in KB 5796.
KB Article	5796

Table 440: Validate network configuration for empty files [103097]

Name	Validate network configuration for empty files
Description	Check if there are empty network configuration files (ifcfg-ethX and /etc/hosts).
Cause	Network configuration files are malformed, with empty files created.
Impact	Genesis may crash, causing services to restart.
Resolution	Follow the instructions in KB 5796.
KB Article	5796

Table 441: CVM TCP checksum error rate check [103104]

Name	CVM TCP checksum error rate check
Description	Check if corrupted packets with compromised checksums are reaching the CVM and not dropped by the host NICs.
Cause	Hardware failure or NIC firmware/driver compatibility issue.
Impact	Cluster performance may be significantly degraded.

Resolution	Check for any hardware failures or issues with firmware/driver/cables. Please refer to KB 9623 for more details.
KB Article	9623

Table 442: Disk firmware needs upgrade [106059] [A106059]

Name	Disk Firmware Check
Description	Checks if any of the disks passthrough to CVM, requires firmware upgrade or not
Alert message	*alert_msg*
Cause	Current disk firmware is not recommended by Nutanix
Impact	Disk may not function properly, resulting in data unavailability.
Resolution	Upgrade Disk Firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.
KB Article	7069
Severity	Warning

Table 443: Controller VM Certificate Expiring [110214] [A1114]

Name	SED Node Certificate
Description	Check if node certificates are about to expire.
Alert message	Certificate for key management server *key_management_server_name* for Controller VM *service_vm_external_ip* expires on *expiration_date*.
Cause	Certificates have defined expiration dates.
Impact	If the node's certificate expires then it will be unable to authenticate with external servers.
Resolution	Get a new signed certificate.
KB Article	2492
Severity	Critical
Severity	Warning
Severity	Info

Table 444: Node in Maintenance Mode [130001] [A1013]

Name	Node Maintenance Mode Status
Description	Node in Maintenance Mode.
Alert message	Controller VM *ip_address* is put in maintenance mode due to *reason*.
Cause	Node removal has been initiated.

Impact	The node will be removed from the cluster.
Resolution	No action is necessary.
Severity	Info

Table 445: Kerberos Clock Skew Failure [130024] [A1083]

Name	Kerberos Clock Skew Status
Description	Kerberos clock skew failure.
Alert message	Kerberos authentication failure occurred on the Controller VM *ip_address* due to clock skew between the Controller VM and the Domain Controller (reason: *reason*).
Cause	There is considerable clock skew between the Controller VM and the Domain Controller.
Impact	The Controller VM cannot manage the host, which may lead to node instability.
Resolution	Ensure that the time on the Controller VM is synchronized with the time of the host. It can be accomplished by configuring the NTP server properly.
KB Article	8036
Severity	Warning

Table 446: Metadata Volume Snapshot Persistent Failure [130026] [A1087]

Name	Metadata Volume Snapshot Persistent
Description	Metadata Volume Snapshot persistent failure.
Alert message	The last *failure_count* metadata volume snapshots have failed. *reason*
Cause	The EBS snapshot service is unavailable, possibly because AWS could not be reached.
Impact	If the cloud instance fails, backups of the Guest VMs may be unavailable.
Resolution	Please check connectivity to cloud instance and reach out to Nutanix Support for assistance.
KB Article	8413
Severity	Warning

Table 447: Metadata Volume Snapshot Timeout [130027] [A1088]

Name	Metadata Volume Snapshot Status
Description	Metadata Volume Snapshot timeout failure.
Alert message	Metadata volume snapshot could not be created in *duration_mins* minutes.
Cause	Metadata volume snapshot frequency is too high.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.

Resolution	Contact Nutanix support.
Severity	Warning

Table 448: Physical Disk Added To Slot [130034] [A1102]

Name	Physical Disk Add Check
Description	Physical Disk added to slot.
Alert message	Disk with serial *disk_serial* and model *disk_model* was added in drive bay *disk_location* on Controller VM *service_vm_external_ip*. Action Required: *action*. Visit Prism hardware page for further actions.
Cause	A drive was inserted into an empty slot.
Impact	The drive may not be available to the cluster until further configuration is completed.
Resolution	Ensure that the drive is added correctly from the Hardware page.
KB Article	9288
Severity	Critical
Severity	Warning
Severity	Info

Table 449: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	Drive *disk_id* with serial *disk_serial* and model *disk_model* in drive bay *disk_location* on Controller VM *service_vm_external_ip* has failed. Node serial is *node_serial_number* and Node position *node_position*. #rf1_custom_message@
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 450: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical drive is removed from the slot.
Alert message	Drive with serial *disk_serial* and model *disk_model* was removed from drive bay *disk_location* on Controller VM *service_vm_external_ip*. Node serial is *node_serial_number* and Node position *node_position*. #rf1_custom_message@

Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the drive will start.
Resolution	Remove the drive from the slot and reseat it. Contact the Hardware Vendor if issue persists.
KB Article	8031
Severity	Critical

Table 451: Self-encrypting Drive Operation Failure [130051] [A1111]

Name	SED Operation Status
Description	Self-encrypting drive operation failure.
Alert message	Self-encrypted Drive operation *operation* has failed for disk *disk_id* with serial *disk_serial* on node *service_vm_external_ip*.
Cause	A self-encrypting drive operation could not be performed.
Impact	The desired action could not be performed.
Resolution	Retest key management server configuration to ensure connectivity and certificates are valid then retry the command.
Severity	Warning

Table 452: Stargate Temporarily Down [130054] [A1030]

Name	Stargate Status
Description	Stargate Temporarily Down.
Alert message	Stargate on Controller VM *ip_address* is down for *downtime* seconds. #rf1_custom_message@
Cause	Zookeeper service not running.
Cause	Cassandra service not running.
Cause	Stargate keeps timing out during restart.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If CVM node running the zeus service was powered off, power it on. Otherwise contact Nutanix support.
Resolution	Contact Nutanix support.
Resolution	Contact Nutanix support.
KB Article	3784
Severity	Critical
Severity	Warning

Table 453: Unsupported Configuration For Redundancy Factor 3 [130055] [A1092]

Name	FT2 Configuration
Description	Unsupported Configuration For Redundancy Factor 3.
Alert message	Controller VM *service_vm_id* with IP address *service_vm_external_ip* has *actual_ram_size_gb*GB RAM which does not meet the configuration requirement of *min_ram_required_gb*GB RAM to support 'Redundancy Factor 3' feature.
Cause	To support 'Redundancy Factor 3' feature all controller VMs in the cluster must meet the minimum requirements for RAM.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Increase RAM on the Controller VM to meet the minimum requirement. Contact Nutanix support for assistance.
Severity	Warning

Table 454: Stargate Responsive [130069]

Name	Stargate Responsive
Description	Check if any Stargate process is still running after being killed.
Cause	Stargate got into an unresponsive state.
Impact	Nutanix datastore may become inaccessible.
Resolution	Restart the Stargate process.
KB Article	3527

Table 455: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	Drive *disk_id* with serial *disk_serial* in drive bay *disk_location* on Controller VM *service_vm_external_ip* has failed diagnostic test. This drive is failing or has failed. #rf1_custom_message@
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 456: File Server VM High CPU Utilization [160060]

Name	File Server VM High CPU Utilization

Description	Checks whether CPU usage is high on file server VM.
Cause	CPU usage of one or all of file server VMs is high.
Impact	Some workloads may encounter high I/O latency.
Resolution	Refer to KB article 5446. Contact Nutanix support if issue still persists or assistance needed.
KB Article	5446

Table 457: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed
Description	Programming a Flow rule failed.
Alert message	Flow policy operation failed for *rule_name*. *reason*
Cause	Microsegmentation service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.
Resolution	Check PC microsegmentation service, PE acropolis service, PC to PE connection, PE to AHV connection.
Severity	Warning

Table 458: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed
Description	Flow failure event.
Alert message	Flow operation failed on *remote_uuid*. *reason*
Cause	Microsegmentation service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microsegmentation service, Check PE acropolis service, AHV microsegmentation module.
Severity	Warning

Table 459: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	Flow mode change failed on *remote_uuid*. *reason*

Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect cvm_logs and contact Nutanix Support
Severity	Warning

Table 460: VPC VMs detected in Flow Security Policy [200608] [A200608]

Name	VPC VMs detected in Flow Security Policy
Description	VMs that have NICs in one or more VPCs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for these VMs.
Alert message	*num_vms* VPC VM(s) detected in Flow Security Policy *rule_name* (*network_security_rule_uuid*). VMs: *vms*. Flow Security Policies will not be enforced for the affected VMs.
Cause	VMs with NICs in VPC subnets have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the affected VMs from VPC subnets for Flow Security Policies to be enforced on these VMs.
Severity	Warning

Table 461: VMs with kDirectNics detected in Flow Security Policy [200609] [A200609]

Name	VMs with kDirectNics detected in Flow Security Policy
Description	VMs that have one or more kDirectNICs have been detected that are also part of a Flow Security Policy. This Security Policy will not be enforced for the traffic through these kDirectNics.
Alert message	*num_vms* VM(s) with kDirectNics detected in Flow Security Policy *rule_name* (*network_security_rule_uuid*). VMs: *vms*. Flow Security Policies will not be enforced for the traffic through these kDirectNics.
Cause	VMs with kDirectNics have been placed into Flow Security Policies.
Impact	Flow Security Policies will not be enforced for the affected VMs.
Resolution	Remove the kDirectNics from the vms for Flow Security Policies to be enforced.
Severity	Warning

DR

Table 462: VM is protected in multiple PDs. [110217] [A110217]

Name	Protection Domains sharing VMs
Description	VM is protected in multiple Protection Domains.
Alert message	VM *vm_name* is protected in multiple protection domains *pd_names*.
Cause	VM is protected in multiple Protection Domains.
Impact	Recovery of the VM will fail.
Resolution	Verify that the VM is not protected in multiple Protection Domains.
KB Article	3337
Severity	Warning

Table 463: NGT installation required. [110242] [A110242]

Name	Cross Hypervisor NGT Installation Check
Description	Some VMs in protection domain replicating to cross-hypervisor remote site do not have NGT installed.
Alert message	Some VMs *vm_list* in cross-hypervisor protection domain *pd_name* do not have NGT installed.
Cause	Some VMs in protection domain replicating to cross-hypervisor remote site do not have NGT installed.
Impact	Recovery of VMs on the remote site might fail.
Resolution	Verify that NGT is installed on all the VMs of the protection domain that you want to recover on the remote site.
KB Article	3567
Severity	Warning

Table 464: Too many files in the Consistency Group. [110243] [A110243]

Name	Metro/Vstore Consistency Group File Count Check
Description	Checks if too many files are being protected by a single consistency group of any Metro/Vstore protection domain.
Alert message	The consistency group *cg_name* (part of protection domain *pd_name*) protects *num_files* files, which *message* the currently set threshold of *threshold_num_files*. Crossing the threshold of *threshold_num_files* files will cause RPO loss.
Cause	Too many files are being protected by a single consistency group of a Metro/Vstore protection domain.
Impact	Performance on Metro-protected or Vstore-protected Storage Container may be degraded.
Resolution	(A) Delete some VMs/files from the consistency group, or move to another Storage Container and protect the new Storage Container. (B) Change alert threshold from the alert policies section. (C) If needed, contact Nutanix Support for assistance.

Vstore consistency group file count threshold exceeded	Warning Threshold: 1200 files Info Threshold: 960 files
KB Article	3820
Severity	Warning
Severity	Info

Table 465: Too many files in the Protection Domain. [110244] [A110244]

Name	Metro/Vstore Protection Domain File Count Check
Description	Checks if too many files are being protected by any Metro/Vstore protection domain.
Alert message	The protection domain *pd_name* protects *num_files* files, which exceeds the currently set threshold of *threshold_num_files*.
Cause	Too many files are being protected by a Metro/Vstore protection domain.
Impact	Performance on Metro-protected or Vstore-protected Storage Container may be degraded.
Resolution	(A) Delete some VMs/files from the protection domain, or move to another Storage Container and protect the new Storage Container. (B) Change alert threshold from the alert policies section. (C) If needed, contact Nutanix Support for assistance.
Vstore protection domain file count threshold	Warning Threshold: 3600 files
KB Article	3821
Severity	Warning

Table 466: Found old clones on cluster. [110245] [A110245]

Name	Clone Age Check
Description	Check for any clones restored from protection domain snapshots that are too old.
Alert message	Found *count* clones related to protection domain *pd_name* which are older than the currently configured threshold of *clone_age_threshold* days.
Cause	Clones were not removed automatically.
Cause	Possible abrupt exit of backup operation.
Impact	Storage performance may be degraded.

Resolution	Remove any clones restored from protection domain snapshots that are too old.
Resolution	Please contact backup application provider to remove the clone.
Age threshold for clones restored from protection domains.	Warning Threshold: 7 days
KB Article	4298
Severity	Warning

Table 467: Too many clones on cluster. [110246] [A110246]

Name	Clone Count Check
Description	Check if there are too many clones restored from protection domain snapshots.
Alert message	Found *count* clones related to protection domain *pd_name* which is more than the currently configured threshold of *clone_count_threshold*.
Cause	Clones were not removed automatically.
Impact	Storage performance may be degraded.
Resolution	Remove any clones restored from protection domain snapshots that are too old.
Count threshold for clones restored from protection domains.	Warning Threshold: 3 clones
KB Article	4298
Severity	Warning

Table 468: Protecting VMs that are using shared VHDX disks is unsupported. [110247] [A110247]

Name	Check VHDX Disks
Description	Check if VMs with shared VHDX disks are not part of any protection domain.
Alert message	Protecting VMs that are using shared VHDX disks is unsupported.
Cause	VMs with shared VHDX disks are protected, which is an unsupported configuration.
Impact	Snapshotting shared VHDX disks that are connected to protected VMs may fail.
Resolution	Unprotect VMs that use shared VHDX disks.

KB Article	3816
Severity	Warning

Table 469: Symlinks found on metro/vstore protected container. [110248] [A110248]

Name	Metro Vstore Symlinks Check
Description	Check for symlinks in metro/vstore protection domain.
Alert message	Symlinks found on metro/vstore protected container *ctr_name*, which might lead to cluster instability.
Cause	Symlinks present in metro/vstore protection domain.
Impact	Snapshot operations on the metro/vstore protected container may fail, leading to replication failures.
Resolution	Remove symlinks from metro/vstore protected container.
KB Article	4276
Severity	Warning

Table 470: Aged third-party backup snapshots present [110250] [A110250]

Name	Aged Third-party Backup Snapshot Check
Description	Check for aged third-party backup snapshots.
Alert message	*pd_name* has *num_snapshot* aged third-party backup snapshot(s) and they may unnecessarily consume storage space in the cluster.
Cause	Third-party backup snapshots are present in the cluster longer than the configured threshold.
Impact	Aged snapshots may unnecessarily consume storage space in the cluster.
Resolution	Contact Nutanix support.
Age threshold for third-party backup snapshots.	Warning Threshold: 7 days
KB Article	4273
Severity	Warning

Table 471: Replications are scheduled on entities from the storage container(s) that have deduplication enabled. [110251] [A110251]

Name	Replication of deduped entities to non-compliant remote sites
Description	Replications are scheduled on entities from the storage container(s) that have deduplication enabled

Alert message	Protection domain *pd_name* with entities from the storage container(s) *dedup_ctrs* that have deduplication enabled are being replicated to single-node backup or cloud remote site(s) *remote_names*.
Cause	Protection domain with entities from the storage container(s) that have deduplication enabled are being replicated to a single-node backup or cloud remote site.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable deduplication on the storage container(s), or use a multi-node AOS cluster as a replication target.
KB Article	4343
Severity	Warning

Table 472: Protection Domain activation may fail as conflicting files exist. [110255] [A110255]

Name	Protection Domain File Conflict Check
Description	Check if any existing file turns into a conflicting state after activating an inactive Protection Domain.
Alert message	Files *conflicting_files* conflict with activation of Protection Domain. All entities may not be recovered.
Cause	This might occur if some entities were unprotected from the Protection Domain before making it inactive or some files were created outside the scope of the protection domain.
Impact	Some entities may not be recovered after you activate the Protection Domain.
Resolution	Before activating this Protection Domain, please identify the entities (VMs or VGs) referring the conflicting files and delete the VMs/VGs after making sure that they are not in use. Do not delete disks directly. Additionally, in case of ESX, please remove any other VM files associated with stale entities. If the issue persists, contact Nutanix support.
KB Article	4592
Severity	Critical

Table 473: Remote Replication is Lagging for Protection Domain Snapshot [110260] [A110260]

Name	PD Snapshot Lagging Remote Replication
Description	Remote replication of the protection domain snapshot is lagging.
Alert message	Snapshot '*snapshot_uuid*' of the protection domain '*protection_domain_name*' has replication to remote '*remote_name*' that is lagging by *lagging_minutes* minutes.
Cause	Replication of the protection domain snapshot to the remote cluster is too slow.
Impact	Recovery point objective for the protection domain will be affected.
Resolution	(A) Ensure that the number of entities are within the supported limit. (B) Ensure that replication bandwidth to remote cluster is unaffected.

KB Article	10197
Severity	Warning
Severity	Info

Table 474: Old Entity-Centric Third-Party Backup Snapshots Present [110263] [A110263]

Name	Aged Entity-Centric Third-party Backup Snapshot Check
Description	Check for the old entity-centric third-party backup snapshots.
Alert message	VM *vm_name* has *num_snapshot* old third-party backup snapshot(s) and may unnecessarily consume storage space in the cluster.
Cause	Entity-centric third-party backup snapshots are present in the cluster longer than the configured threshold.
Impact	Old snapshots may unnecessarily consume storage space in the cluster.
Resolution	The Recovery Point for the specified VM needs to be deleted from Recoverable Entities in Prism Central. If the issue persists, contact Nutanix support.
Age threshold for third-party backup snapshots.	Warning Threshold: 7 days
KB Article	4273
Severity	Warning

Table 475: NearSync Replication of Protection Domain has not progressed. [110264] [A110264]

Name	Nearsync Replication Stuck
Description	Nearsync replication of the protection domain has not progressed.
Alert message	NearSync replication of the protection domain *protection_domain_name* to the remote site *remote_name* for the snapshot *snapshot_id* has not made any progress in the last *time_string*.
Cause	Network outage, saturated network or unavailability of services like Stargate, Cerebro on the remote site. Snapshots of the protection domain will not be available on the remote site.
Cause	SSD snapshot reserve on the remote site does not have enough free disk space.
Impact	Protection domain may transition out of NearSync.
Impact	Snapshots of the protection domain will not be available on the remote site.
Resolution	Check the network connection and status of the remote site.
Resolution	Ensure that the remote site has enough disk space available in the SSD snapshot reserve.
KB Article	10249
Severity	Warning

Table 476: Protected VM is not capable of backup and recovery. [110268] [A110268]

Name	Protected VM CBR Incapable
Description	Some of the protected VMs are not capable of backup and recovery.
Alert message	Protected VM *vm_name* is not capable of backup and recovery. Reason: *cbr_incapable_reason*
Cause	Not all protected VMs have a valid configuration for backup and recovery
Impact	RPO will be affected.
Resolution	Review and update the VM configuration. If the configuration of the VM cannot be changed then unprotect the VMs.
KB Article	7635
Severity	Warning

Table 477: Stale VM present in DR system [110457] [A110457]

Name	Stale VM present in the DR system
Description	A stale VM in the DR protection conflicts with the live VM
Alert message	A stale VM *stale_vm_name* with ID(*stale_vm_hs_id*, *vm_ha_uuid*) against VM *vm_name* with ID(*vm_id*, *vm_ha_uuid*). Please unprotect the stale VM entry *protection_information*.
Cause	An in-place restore workflow might have been performed on the VM by backup software resulting in newly created VM with the same Nutanix ID
Impact	Some or all DR workflows on this VM might fail
Resolution	Please unprotect the stale VM and if needed, protect the restored VM. If the issue still persists contact Nutanix Support
KB Article	10545
Severity	Critical

Table 478: Cerebro Stats Collector Failed. [111059] [A111059]

Name	Cerebro stats collector
Description	Collect Cerebro stats and publish to CFS.
Alert message	Cerebro stats collection has failed.
Cause	(A) Snapshot chain length has exceeded the threshold limit. (B) Garbage collection is not able to keep up with the configured RPO.
Impact	Cluster is not able to keep up with the configured RPO
Resolution	Lower the frequency of the snapshot schedule.
KB Article	6200
Severity	Warning

Table 479: Secure Tunnel To Remote Site Down [130000] [A1090]

Name	Remote Site Tunnel Status
Description	Secure Tunnel To Remote Site Down.
Alert message	Secure tunnel to remote site *remote_name* is down.
Cause	Incorrect remote site configuration or network connectivity issue.
Impact	Replication to remote site will fail.
Resolution	Check if the IP address specified in the remote site is reachable.
Severity	Critical
Severity	Warning
Severity	Info

Table 480: Metro Availability Is Disabled [130002] [A1124]

Name	Metro Connectivity
Description	Metro availability is disabled.
Alert message	Metro availability for the protection domain '*protection_domain_name*' to the remote site '*remote_name*' is disabled because of '*reason*'.
Cause	Remote site unreachable.
Impact	Metro availability operation is disabled.
Resolution	Check network connectivity and health of the remote cluster. If issue persists reach out to Nutanix Support
KB Article	8228
Severity	Critical
Severity	Warning

Table 481: Duplicate Remote Cluster ID [130018] [A1038]

Name	Duplicate Remote Cluster ID Check
Description	Duplicate remote cluster ID.
Alert message	Remote site already exists with the name '*conflicting_remote_name*'.
Cause	Two remote sites with different names or different IP addresses have same cluster ID. This can happen in two cases: (a) A remote cluster is added twice under two different names (through different IP addresses) or (b) Two clusters have the same cluster ID.
Impact	Protected data is not replicated to the remote site.
Resolution	In case (a) remove the duplicate remote site. In case (b) verify that the both clusters have the same cluster ID and contact Nutanix support.
Severity	Warning

Table 482: Entities Restored But Unprotected [130019] [A1134]

Name	Restored Entities Protected
Description	Entities are restored but unprotected.
Alert message	Some entities became unprotected after restoring *protection_domain_name*. Count: *num_entities*. Entity names: *entity_names*. Reason: *reason*.
Cause	Some other protected entities conflict with the entities being recovered.
Impact	Some entities became unprotected after restoring the snapshot.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 483: Entities Skipped During Restore [130020] [A1132]

Name	Entities Restored Check
Description	Entities were skipped during restore.
Alert message	Some entities skipped during restore of *protection_domain_name*. Count: *num_entities*. Entity names: *entity_names*. Reason: *reason*.
Cause	Existing files conflict with the files to be recovered.
Impact	Not all entities could be restored from the snapshot.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	8415
Severity	Critical

Table 484: Entity conflict in Consistency Groups [130022] [A1136]

Name	Consistency Group Configuration
Description	Entity conflict in consistency group(s)
Alert message	Failed to create a snapshot of protection domain *protection_domain_name* as some files in consistency group *consistency_group_name* overlap other consistency groups.
Cause	VMs with common files are not in the same Consistency Group.
Impact	Snapshot will fail.
Resolution	Put VMs with common files in the same Consistency Group.
Severity	Warning

Table 485: Metadata Volume Snapshot Persistent Failure [130026] [A1087]

Name	Metadata Volume Snapshot Persistent

Description	Metadata Volume Snapshot persistent failure.
Alert message	The last *failure_count* metadata volume snapshots have failed. *reason*
Cause	The EBS snapshot service is unavailable, possibly because AWS could not be reached.
Impact	If the cloud instance fails, backups of the Guest VMs may be unavailable.
Resolution	Please check connectivity to cloud instance and reach out to Nutanix Support for assistance.
KB Article	8413
Severity	Warning

Table 486: Metadata Volume Snapshot Timeout [130027] [A1088]

Name	Metadata Volume Snapshot Status
Description	Metadata Volume Snapshot timeout failure.
Alert message	Metadata volume snapshot could not be created in *duration_mins* minutes.
Cause	Metadata volume snapshot frequency is too high.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.
Severity	Warning

Table 487: Metro Availability Configuration Failed [130028] [A1123]

Name	Metro Availability
Description	Metro availability start failed.
Alert message	Failed to *operation* for the protection domain '*protection_domain_name*' to the remote site '*remote_name*'. Reason: '*reason*'.
Cause	Check the alert message for the reason of failure
Impact	Metro availability operation could not be started.
Resolution	Resolve the issue as stated in the alert message and retry the Metro operation. If the issue persists contact Nutanix Support
KB Article	10250
Severity	Critical

Table 488: Stale NFS Mount [130029] [A1135]

Name	Metro Vstore Mount Status
Description	Stale NFS mount.
Alert message	Unable to process NFS requests for vstore id *vstore_id* due to *reason*.

Cause	Stale NFS file handle.
Impact	Unable to process NFS requests.
Resolution	Unmount the datastore from all hosts and remount again.
Severity	Warning

Table 489: Compliance Failure of the Protected VMs. [130037] [A1109]

Name	Protected VM CBR Capability
Description	Protected VM is not Nutanix backup and recovery compliant.
Alert message	*message*
Cause	Protected VM may have been modified to include files that are not on a Nutanix Storage Container.
Impact	Any data associated with the VM may not be backed up or replicated to the remote site.
Resolution	Remove or detach files from the VM that are not on the Nutanix Storage Containers and make the VM Nutanix backup and recovery compliant.
Severity	Warning

Table 490: Protected VM Not Found [130038] [A1010]

Name	Protected VM Not Found
Description	Protected VM not found.
Alert message	Unable to locate the VM '*vm_name* (ID *vm_id*)' protected by protection domain '*protection_domain_name*'.
Cause	The protected VM cannot be found and may have been deleted.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove the VM from the protection domain.
Severity	Info

Table 491: Protection Domain Activation Failed [130040] [A1043]

Name	PD Active
Description	Protection domain activation failed.
Alert message	Unable to make protection domain '*protection_domain_name*' active on remote site '*remote_name*' due to '*reason*'.
Cause	Various
Impact	Protected VMs could not be started during failover to a remote site.

Resolution	Resolve the stated reason for the failure. If issue persists, contact Nutanix support.
KB Article	8468
Severity	Critical

Table 492: Protection Domain Activation or Migration Failure [130041] [A1060]

Name	PD Change Mode Status
Description	Protection domain activation or migration failed.
Alert message	Protection domain *protection_domain_name* activate/deactivate failed with the error : *reason*
Cause	(A) Protection domain cannot be activated or migrated. (B) Protection domain with same name might be active on remote site. (C) Remote sites might not be configured correctly, please check the remote sites on source and target clusters. (D) Protection Domain has one or more missing VMs/VGs.
Impact	Protected VMs could not be started during failover to a remote site.
Impact	Some of the VMs might have been powered off. User has to power on those VMs.
Resolution	Resolve the stated reason for the failure. If issue persists, please contact Nutanix Support.
KB Article	8021
Severity	Critical

Table 493: Registration Of One Or More VMs Failed [130043] [A1093]

Name	PD VM Registration Status
Description	Registration of one or more VMs failed.
Alert message	Failed to register one or more VMs of the protection domain *protection_domain_name*.
Cause	VMs in the snapshot have invalid VM configuration or the local hypervisor reported an error.
Impact	VMs restored from the snapshot will not be unavailable.
Resolution	Check the existence of the VM in the snapshot and check any errors reported in the hypervisor management software. If you cannot resolve the error, contact Nutanix support.
KB Article	8005
Severity	Critical

Table 494: Protection Domain Replication Expired [130045] [A1003]

Name	PD Replication Expiry Status
Description	Protection domain replication expired.

Alert message	Protection domain *protection_domain_name* replication to the remote site *remote_name* has expired before it is started.
Cause	Replication is taking too long to complete before the snapshots expire.
Impact	The snapshot that was scheduled to be replicated could not be replicated. Data that is expected to be protected may not be protected.
Resolution	Review replication schedules taking into account bandwidth and overall load on systems. Confirm retention time on replicated snapshots. If issue persists, please contact Nutanix Support.
Severity	Warning

Table 495: Protection Domain Replication Failure [130046] [A1015]

Name	PD Replication Status
Description	Protection Domain Replication Failure
Alert message	Protection domain *protection_domain_name* replication to remote site *remote_name* failed. *reason*.
Cause	Various
Impact	Snapshots of protected VMs were not replicated to the remote site.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 496: Skipped Replication Of Snapshot For Protection Domain [130047] [A1113]

Name	PD Replication Skipped Status
Description	Skipped replication of the snapshot.
Alert message	Replication was skipped for protection domain *protection_domain_name* of the latest snapshot *snapshot_name* to remote site *remote_name*. *reason* There could be other snapshots that may have been skipped.
Cause	Snapshot replication was skipped because a newer snapshot is available.
Impact	Snapshot is not present on the remote site.
Resolution	Make sure that enough bandwidth is available for replication. If the issue persists, consider reducing the frequency of replication. Perform a manual resync of the skipped snapshot, if required.
KB Article	4282
Severity	Warning

Table 497: Failed to Receive Snapshot for the Protection Domain [130048] [A1127]

Name	Protection Domain Receive Snapshot Failure
Description	Protection Domain Receive Snapshot Failure.

Alert message	Failed to receive snapshot for the protection domain '*protection_domain_name*' from the remote site *remote_name*. *reason*
Cause	Incorrect configuration on DR site or Network Error.
Impact	The snapshot receive operation could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 498: Protection Domain Snapshot Failure [130049] [A1064]

Name	PD Snapshot Status
Description	Protection Domain Snapshot Failure.
Alert message	Protection domain *protection_domain_name* snapshot '*snapshot_id*' failed. *reason*.
Cause	Protection domain has VMs being protected by other vstore(s) or in case of a cluster where Metro Availability is configured, Metro protection domain has more entities than supported.
Impact	A requested snapshot of guest VMs and files in the protection domain did not succeed.
Resolution	Unprotect VMs from other vstore(s) before snapshotting the concerned protection domain. In case of a cluster where Metro Availability is configured, make sure number of entities are within the supported limit.
KB Article	8038
Severity	Critical
Severity	Warning

Table 499: Remote Site Is Unhealthy [130050] [A1125]

Name	Remote Site Health
Description	Remote site is unhealthy.
Alert message	Remote '*remote_name*' is unhealthy due to reason '*reason*'.
Cause	Various
Impact	Replication to remote site may fail.
Resolution	Check if the IP address specified in the remote site is reachable.
Severity	Warning

Table 500: Snapshot Partially Crash Consistent [130052] [A1110]

Name	Snapshot Crash Consistent
Description	Snapshot partially crash consistent.

Alert message	Failed to create an application-consistent snapshot for one or more VMs in snapshot *snapshot_name* of protection domain *protection_domain_name*. A crash-consistent snapshot has been created for these VMs instead.
Cause	VSS or hypervisor error
Impact	Recovery may take longer and involve manual steps.
Resolution	No action is necessary if this issue is intermittent. If this error persists, contact Nutanix support.
Severity	Warning

Table 501: VM Action Error [130057] [A1033]

Name	VM Action Status
Description	VM action error.
Alert message	Failed to *action* VM with name '*vm_name* and internal ID (ID *vm_id*)' due to *reason*.
Cause	A VM could not be restored because of a hypervisor error, or could not be deleted because it is still in use.
Impact	The requested VM action (restore or delete) could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10658
Severity	Warning

Table 502: VM Virtual Hardware Version Not Compatible [130058] [A1133]

Name	VM Virtual Hardware Version Compatible
Description	Virtual hardware version of the VM is not compatible.
Alert message	Virtual hardware version of the VM *vm_name* is not supported by any of the nodes on the remote site *remote_name*. Protection domain: *protection_domain_name*. Reason: *reason*.
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote site.
Impact	VM may not register properly on the remote site on restore/clone.
Resolution	Upgrade the hypervisor version on the remote site to support virtual hardware version for the VM.
Severity	Warning

Table 503: Missing Network Mapping for the Remote Site. [130062] [A1156]

Name	Remote Site Network Configuration
Description	Network mapping not specified for the remote site.
Alert message	No network mapping specified for the remote site *remote_name*.

Cause	Network mapping not specified for the remote site.
Impact	When the VM is restored or cloned on the remote site, networking configuration may not be recovered.
Resolution	Specify network mapping for the remote site.
Severity	Warning

Table 504: Remote Site {remote_name} Network Mapping Invalid [130063] [A1157]

Name	Remote Site Network Mapping Configuration
Description	Invalid Network Mapping Specified.
Alert message	Invalid network mapping specified for remote site *remote_name*: *reason*.
Cause	Either the source or destination network configuration is not present in the network mapping specified on the remote site, or the networks specified in the network mapping do not exist.
Cause	The network associated with one or more VM(s) does not have the corresponding recovery network mapping in the remote site.
Impact	When the VM is restored or cloned on the remote site, networking configuration may not be recovered.
Resolution	Delete existing network mappings and create the network mapping with the associated networks present on the source and the destination cluster for the remote site.
Resolution	Update the network mapping of the remote site with all the network(s) associated with the VM(s).
KB Article	10243
Severity	Warning

Table 505: Protection Domain Snapshot Operation Skipped [130066] [A1147]

Name	PD Snapshot Skipped Status
Description	Skipped snapshot operation.
Alert message	Replication was skipped for protection domain *protection_domain_name* to remote site *remote_name*. *reason*
Cause	Snapshot operation was skipped.
Impact	Snapshot is not present on both the local and remote sites.
Resolution	Wait for the completion of the upgrade process before initiating snapshot schedules.
Severity	Info

Table 506: Volume Group Action Error [130071] [A1163]

Name	Volume Group Action Status
Description	Volume group action error.
Alert message	Failed to *action* volume group with name '*vg_name*' and internal ID '*vg_uuid*' because *reason* in protection domain '*protection_domain_name*'
Cause	A volume group could not be restored or could not be deleted because it is still in use.
Impact	The requested volume group action (restore or delete) could not be completed.
Resolution	Detach the volume group from VMs and external initiators before recovery. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10120
Severity	Critical

Table 507: VSS Snapshot Failed [130073] [A130073]

Name	VSS Snapshot Status
Description	VSS snapshot failed.
Alert message	VSS snapshot failed for the VM(s) *vm_names* protected by the *protection_domain_name* in the snapshot *snapshot_id* because *reason*.
Cause	Guest is not able to quiesce VM due to internal error.
Impact	Crash consistent snapshot is taken instead of application consistent.
Resolution	Review the logs in the guest VM and take necessary actions. If issue persists contact Nutanix Support.
Severity	Warning

Table 508: Nutanix Guest Tools Not Installed [130074] [A130074]

Name	NGT Configuration
Description	Nutanix Guest Tools is not Installed.
Alert message	VSS is enabled but Nutanix Guest Tools are not installed on the VM(s) *vm_names* protected by *protection_domain_name*.
Cause	VSS is enabled but Nutanix Guest Tools are not installed on the VM(s).
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Install Nutanix Guest Tools on the VM(s).
Severity	Warning

Table 509: Protection Domain might have symlinks. [130077] [A130077]

Name	Snapshot Symlink Check
Description	Protection domain snapshot has symlinks.
Alert message	Protection Domain *protection_domain_name* might have symlinks that have been skipped in the snapshot.
Cause	The protection domain might have symlinks that have been skipped in the snapshot.
Impact	The symlinks in the protection domain have been skipped in the snapshot.
Resolution	Delete the symlinks or protect the entities individually.
Severity	Warning

Table 510: Entity Restore Aborted [130078] [A130078]

Name	EntityRestoreAbort
Description	Entity restore aborted.
Alert message	Restoring VMs failed for the snapshot *snapshot_id* protected by the *protection_domain_name* because *reason*.
Cause	Check the alert message for the reason of failure.
Impact	Snapshot expiry is removed to prevent data loss, and you should remove it manually.
Resolution	Restart the restore process. Snapshot expiry is removed to allow subsequent restore operations. Snapshot will need to be deleted manually to free up the space if required. If this error persists, contact Nutanix support.
Severity	Warning

Table 511: Remote Site Is Incompatible For Replication [130079] [A130079]

Name	Remote Site Snapshot Replication Status
Description	Skipped replication of the snapshot.
Alert message	Replication was skipped for protection domain *protection_domain_name* to remote site *remote_name*. *reason*
Cause	Snapshot replication was skipped because the remote is not fully upgraded.
Impact	Snapshot is not present on the remote site.
Resolution	(A) Wait for the completion of the upgrade process on the remote before initiating snapshot schedules. (B) Manually replicate the snapshot from Prism UI after the remote site upgrade process completes or wait for next scheduled snapshot to replicate data.
Severity	Info

Table 512: Nutanix Guest Tools Agent Is Not Reachable On The VM [130081] [A130081]

Name	VSS VM Reachable
Description	VSS on VM(s) is not reachable.
Alert message	VM(s) *vm_names* Guest Agent Service is not reachable, protected by *protection_domain_name*.
Cause	The communication link to the VMs Nutanix Guest Agent service seems to be down.
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Please check and restart Nutanix Guest Agent service inside the VM.
Severity	Warning

Table 513: Protected VM(s) Not Found [130083] [A130083]

Name	Protected VMs Not Found
Description	Protected VM(s) Not Found.
Alert message	Unable to locate VM(s) *vm_names* protected by protection domain '*protection_domain_name*'.
Cause	Protected VM(s) cannot be found and may have been deleted.
Impact	Any data associated with the VMs may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Verify that the VM(s) are registered on the host. If the VM(s) were unregistered or deleted, remove VM(s) from the protection domain. If the CVM on the host the VM(s) are registered on was shutdown or rebooting, retry the snapshot operation.
Severity	Warning

Table 514: Protected Volume Groups Not Found [130084] [A130084]

Name	Protected Volume Groups Not Found
Description	Protected volume groups not found.
Alert message	Unable to locate volume group(s) *vg_names* protected by protection domain '*protection_domain_name*'.
Cause	Protected volume group(s) cannot be found and may have been deleted.
Impact	Any data associated with the volume group(s) may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Remove volume group(s) from the protection domain.
Severity	Warning

Table 515: VSS Software or (pre_freeze/post_thaw) Scripts Not Installed [130085] [A130085]

Name	VSS Scripts Not Installed
Description	VSS software or pre_freeze/post_thaw scripts are not installed
Alert message	VSS is enabled but VSS software or pre_freeze/post_thaw scripts are not installed on the guest VM(s) *vm_names* protected by *protection_domain_name*.
Cause	VSS is enabled but VSS software or pre_freeze/post_thaw scripts are not installed on the VM(s).
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Install VSS software or (pre_freeze/post_thaw) scripts on the guest VM(s).
Severity	Warning

Table 516: VStore Snapshot Status [130086] [A130086]

Name	VStore Snapshot Status
Description	vStore snapshot status.
Alert message	Snapshot status for vstore *vstore_name*: *reason*.
Cause	Various
Impact	The requested VStore snapshot action could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 517: Failed To Snapshot Entities [130088] [A130088]

Name	Failed To Snapshot Entities
Description	Failed to snapshot entities.
Alert message	Failed to snapshot the entities *entity_names* in the protection domain *protection_domain_name*, snapshot *snapshot_id*.
Cause	Conflicting hypervisor tasks might be running on the entities.
Impact	The entities are not captured in the snapshot.
Resolution	Retry the operation after the hypervisor tasks are complete.
KB Article	8485
Severity	Warning

Table 518: Associated entities are not protected together. [130090] [A130090]

Name	Related Entity Protection Status
Description	Protection status of a related entity.

Alert message	*error_message*.
Cause	Related entity is not protected in the same protection domain.
Impact	Related VM/Volume Group will not be snapshotted and recovered.
Resolution	Protect the related entity in the same consistency group.
Severity	Warning

Table 519: Remote Site Operation Mode ReadOnly [130092] [A1189]

Name	Remote Site Operation Mode ReadOnly
Description	Operation Mode of Remote Site changed to kReadOnly.
Alert message	Remote site *remote_name* moved to read-only state, outgoing replications to this remote will fail.
Cause	SSD failure on remote site in case of single node or two-node cluster.
Cause	Cloud virtual appliance ran out of memory in case of cloud connect.
Cause	Data stored on cloud appliance exceeded 20TB.
Impact	All outgoing replications to remote site will be aborted.
Resolution	Replace the failed SSD on the remote site.
Resolution	Delete some remote snapshots to reduce memory or data usage in case of cloud connect and wait for some time.
Resolution	Increase the cloud CVM size. Refer to the Nutanix documentation for instructions.
KB Article	8082
Severity	Warning

Table 520: VM renamed on restoration [130094] [A130094]

Name	VM Renamed On Conversion
Description	VM Renamed On Restoration
Alert message	VM '*vm_name* (ID *vm_id*)' will be registered as '*new_vm_name*' during VM restore to avoid duplicate VM names.
Cause	Directory with same name as VM name exists on Storage Container.
Impact	The restored VM will be registered under the new name.
Resolution	No action is necessary.
Severity	Info

Table 521: Failed to reconfigure Nutanix Guest Tools during VMs recovery for protection domain [130095] [A130095]

Name	Failed To Recover NGT Information

Description	Failed to enable Nutanix Guest Tools during VMs recovery for protection domain.
Alert message	Failed to reconfigure Nutanix Guest Tools for VMs in protection domain *protection_domain_name* due to *reason*.
Cause	Nutanix Guest Tools service might be down.
Impact	VM
Resolution	Enable and mount Nutanix Guest Tools, and restart Nutanix Guest Agent service within the VM, for VMs in protection domain. If you cannot resolve the error, contact Nutanix support.
KB Article	10079
Severity	Warning

Table 522: Failed to reconfigure Nutanix Guest Tools for the recovered VM [130096] [A130096]

Name	Failed To Recover NGT Information for VM
Description	Failed to reconfigure Nutanix Guest Tools for a VM in protection domain.
Alert message	Failed to reconfigure Nutanix Guest Tools for VM *vm_name* in protection domain *protection_domain_name*.
Cause	Virtual IP address of the cluster might not have been configured.
Cause	Error in generating certificate for VM.
Cause	Guest VM information could not be retrieved.
Impact	VM
Resolution	Configure the virtual IP address of the cluster (if not configured).
Resolution	Enable and Mount Nutanix Guest Tools on the failed VM, and restart the Nutanix Guest Agent service within the VM.
Resolution	Resolve any issues on the Guest VM level to make sure NGT can communicate to the VM.
KB Article	10077
Severity	Warning

Table 523: Failed to reconfigure Nutanix Guest Tools for a VM in protection domain [130097] [A130097]

Name	Failed To Mount NGT ISO On Recovery of VM
Description	Failed to mount ISO image as part of Nutanix Guest Tools reconfiguration for a VM in protection domain
Alert message	Failed to reconfigure Nutanix Guest Tools for VM *vm_name* in protection domain *protection_domain_name* due to ISO mount failure.
Cause	Virtual IP address of the cluster might not have been configured.
Cause	Guest VM information could not be retrieved.

Cause	NGT ISO could not be mounted.
Impact	VM
Resolution	Configure the virtual IP address of the cluster (if not configured).
Resolution	Mount Nutanix Guest Tools ISO on the failed VM, and restart Nutanix Guest Agent service within the VM.
Resolution	Make sure there is at least one free CD-ROM slot on the Guest VM. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 524: External iSCSI Attachments Not Snapshotted [130099] [A130099]

Name	External iSCSI Attachments Not Snapshotted
Description	External iSCSI Attachments Not Snapshotted.
Alert message	*error_message*. Volume Group external iSCSI attachments cannot be snapshotted for these VMs.
Cause	VMs have duplicate IQNs.
Cause	VMs have configured iSCSI target IP addresses that do not belong to the local Nutanix cluster.
Cause	IQN could not be resolved to VM ID.
Impact	Volume group external iSCSI attachments cannot be snapshotted and restored for these VMs.
Resolution	Make sure all VMs in the protection domain have different IQNs.
Resolution	Remove iSCSI target IP addresses that do not belong to the local Nutanix cluster from the VMs.
Resolution	Install Nutanix Guest Tools on all the VMs using volume group external iSCSI attachments.
Severity	Warning

Table 525: Volume Group Attachments Not Restored [130101] [A130101]

Name	Volume Group Attachments Not Restored
Description	Volume Group Attachments Not Restored.
Alert message	*error_message* for the protection domain '*protection_domain_name*'.
Cause	Acropolis failed to attach volume groups to the VMs.
Impact	VMs will not be able to access volume group iSCSI disks.
Resolution	Manually attach volume groups to the VMs.
Severity	Warning

Table 526: Self service restore operation failed [130102] [A130102]

Name	Self service restore operation Failed
Description	Self service restore operation failed
Alert message	Unable to '*operation*' disk from VM '*vm_name* (ID *vm_id*)' due to '*reason*'.
Cause	(A) Snapshot disk could not be recovered. (B) Recovered disk could not be attached to the VM. (C) Self-service restore disk could not be detached from VM.
Impact	In the case of attachment failure, the user will not be able to restore files from snapshot.
Impact	In the case of detachment failure, the guest VM will have an extraneous disk attached.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 527: Witness is not configured [130114] [A130114]

Name	Witness Not Configured
Description	Witness is not configured
Alert message	Failed to *operation* for the protection domain '*protection_domain_name*' to the remote site '*remote_name*'. Reason: Witness is not configured or a different Witness VM is configured on the remote site
Cause	Witness is not configured
Impact	Metro availability operation failed
Resolution	Configure Witness
KB Article	10119
Severity	Critical

Table 528: Witness is not reachable [130115] [A130115]

Name	Witness Not Reachable
Description	Witness is not reachable
Alert message	Failed to *operation* since *unreachable_since* for the protection domain '*protection_domain_name*' to the remote site '*remote_name*'. Reason: Witness is not reachable or the Witness credentials are incorrect
Cause	Witness or the network is down.
Cause	Witness credentials are incorrect.
Impact	Metro availability failures will not be automatically handled
Resolution	Check if Witness is up on the network.
Resolution	Check Witness credentials.

KB Article	5483
Severity	Critical

Table 529: Metro Availability Is Promoted [130116] [A130116]

Name	Automatic Promote Metro Availability
Description	Metro availability is promoted.
Alert message	Metro availability for the protection domain '*protection_domain_name*' to the remote site '*remote_name*' is promoted because of '*reason*'.
Cause	Remote site unreachable.
Impact	Metro availability is promoted on standby cluster and storage is now inaccessible in the active side.
Resolution	(A) Check if cluster service is healthy at remote cluster. (B) Ensure VMs are running first at the promoted side. In the case of storage only outage in order to restore VM availability follow the procedure in the Prism Web Console Guide, section: <Recovering from a Storage-only Failure on the Primary Site>. After that verify the previously active cluster is healthy before re-enabling metro availability.
KB Article	8279
Severity	Critical
Severity	Warning

Table 530: Failed to update Metro Availability failure handling [130117] [A130117]

Name	Updating Metro Failure Handling Failed
Description	Error in updating failure-handling on Metro Availability protection domain.
Alert message	Error in updating Metro Availability failure handling for the protection domain '*protection_domain_name*' to the remote site '*remote_name*'. Reason: '*reason*'.
Cause	Failed to update Metro Availability Failure Handling due to possible network error
Impact	Metro Availability failure-handling could not be updated.
Resolution	Resolve network issues. If the issue persists contact Nutanix Support
KB Article	8480
Severity	Critical

Table 531: Failed to update Metro Availability failure handling on the remote site [130118] [A130118]

Name	Updating Metro Failure Handling Remote Failed
Description	Error in updating failure handling on the remote Metro Availability protection domain

Alert message	Error in updating Metro Availability failure handling on the remote site '*remote_name*' for the protection domain '*protection_domain_name*'. Reason: '*reason*'.
Cause	Failed to update Metro Availability Failure Handling due to possible network error
Impact	Failure handling on this Metro Availability protection domain might be affected
Resolution	Resolve any network related issues. If issue persists contact Nutanix Support
KB Article	8481
Severity	Critical
Severity	Warning

Table 532: Authentication failed in Witness [130123] [A130123]

Name	Authentication Failed in Witness
Description	Authentication failed in Witness
Alert message	Failed to *operation* for the protection domain '*protection_domain_name*' to the remote site '*remote_name*'. Reason: Authentication failure in Witness
Cause	Witness configured with incorrect credentials
Impact	Metro availability failures will not be automatically handled
Resolution	Configure Witness with correct credentials
KB Article	4376
Severity	Critical

Table 533: Metro Availability Is Promoted [130125] [A130125]

Name	Manual Promote Metro Availability
Description	Metro availability is promoted.
Alert message	Metro availability for the protection domain '*protection_domain_name*' to the remote site '*remote_name*' is promoted.
Cause	An administrator promoted metro availability.
Impact	Metro availability is promoted on standby site.
Resolution	Re-enable to get the benefit of Metro protection domain.
Severity	Info

Table 534: Metro Availability Is Disabled [130126] [A130126]

Name	Manual Break Metro Availability
Description	Metro availability is disabled.

Alert message	Metro availability for the protection domain '*protection_domain_name*' to the remote site '*remote_name*' is disabled.
Cause	An administrator disabled metro availability.
Impact	Metro availability operation is disabled.
Resolution	Re-enable Metro on the protection domain to get the benefit of Metro availability
Severity	Info

Table 535: VMs in the standby site of a Metro Availability protection domain are running at suboptimal performance. [130129] [A130129]

Name	Metro Protection domain VMs running at Sub-optimal performance
Description	VMs in the standby site of a Metro Availability protection domain are running at suboptimal performance.
Alert message	Warning: *num_of_vms* VMs are running on the standby site '*remote_name*' of protection domain '*protection_domain_name*'. However, all the I/O operations are being forwarded to this site. This can significantly affect the performance of VMs. It is recommended to move VMs to this site or update the DRS rules. The affected VMs list is: *vm_names*
Cause	VMs are hosted at the standby site of the Metro Availability protection domain.
Impact	Experience of decreased performance
Resolution	Host the VMs at the active site of the Metro Availability protection domain.
Severity	Warning

Table 536: Protection domain contains more than the specified VMs [130130] [A130130]

Name	Protection Domain VM Count Check
Description	Protection domain contains multiple VMs.
Alert message	It is not recommended to protect more than one VM in a protection domain *protection_domain_name* if you are replicating to the backup only remote site *remote_name*.
Cause	Replicating a protection domain with multiple VMs to a backup only remote site.
Impact	Selective retrieval of any VM will take more time as the snapshot will capture all the VMs in the protection domain.
Impact	Replications may be skipped because of a large snapshot size.
Resolution	Protect each VM in a separate protection domain.
Severity	Info

Table 537: Snapshot contains entities from the storage container that have deduplication enabled. [130131] [A130131]

Name	Replication Of Deduped Entity
Description	Snapshot contains entities from the storage container that have deduplication enabled.
Alert message	Do not replicate protection domain *protection_domain_name* comprising entities from the storage container that have deduplication enabled to a *remote_site_type* remote site *remote_name*.
Cause	Replicating a snapshot with entities from the storage container that have deduplication enabled to a single node backup or cloud remote site.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable dedup on the Storage Containers, or use a multi-node AOS cluster as a replication target.
Severity	Warning

Table 538: Cloud Appliance Deployment Failed [130133] [A130133]

Name	Cloud Appliance Deployment Failed
Description	Cloud Appliance Deployment Failed
Alert message	Failed to configure cloud appliance for remote site *remote_site_name*. *reason*.
Cause	Resource quota for buckets/instances/storage accounts exhausted for cloud account.
Cause	AWS credentials do not have permissions to perform S3/SG/EC2 related operations.
Cause	Cloud CVM deployed was not able to reach the Internet.
Impact	Cloud appliance could not be configured as remote site.
Resolution	Make sure that account has enough quota to create new buckets/instances.
Resolution	Make sure that AWS IAM user has permissions to perform S3/SG/EC2 operations.
Resolution	Make sure that the subnet in which cloud appliance is being deployed allows Internet connectivity.
Severity	Warning

Table 539: vStore is being replicated to backup only remote site. [130134] [A130134]

Name	Vstore Replication To Backup Only Remote
Description	vStore is being replicated to backup only remote site.
Alert message	vStore protection domain *protection_domain_name* is being replicated to backup only remote site *remote_name*.
Cause	vStore is being replicated to backup only remote site.

Impact	Low RTO can be seen as selective retrieve of VMs is not supported and snapshot of whole protection domain, having all the VMs in it, has to be retrieved.
Resolution	One VM per protection domain is the recommended configuration to achieve higher RTO.
Severity	Info

Table 540: Replication of protection domain {protection_domain_name} has not progressed. [130137] [A130137]

Name	Protection Domain Replication Stuck
Description	Replication of protection domain has not progressed.
Alert message	Replication of protection domain *protection_domain_name* to remote site *remote_name* for the snapshot *snapshot_id* has not made any progress in last one hour.
Cause	Network outage with remote.
Cause	Remote cluster does not have enough free disk space available.
Cause	Single-node cluster might be in read-only mode, due to one of the metadata SSDs failure.
Impact	Snapshots of protection domain will not be available on the remote site.
Resolution	Check network connection with remote.
Resolution	Make sure that the remote cluster has enough disk space available.
Resolution	Replace the failed SSD.
Severity	Warning

Table 541: Snapshot Reserve on SSD is Full. [130139] [A130139]

Name	LWS Store Full
Description	SSD reserved for the Nearsync snapshots is full.
Alert message	The SSD space that is reserved for the snapshots is full.
Cause	Snapshot reserve on the SSD is full or has high IO workload.
Impact	RPO cannot be met and will be temporarily transitioned to 60 minutes.
Resolution	Reduce the retention time of the snapshots.
Severity	Warning

Table 542: Snapshot Reserve Allocation Failure [130140] [A130140]

Name	LWS store allocation too long
Description	Failed to allocate the snapshot reserve.
Alert message	Failed to allocate the snapshot reserve in *duration_mins* minutes.

Cause	Snapshot reservation timed out.
Impact	Protection domain will not be able to meet less than 15 minutes RPO.
Resolution	Check SSD performance from Prism Web console and if issue persists contact Nutanix support
Severity	Warning

Table 543: Failed to allocate snapshot reserve on the remote site [130141] [A130141]

Name	LWS store allocation in remote too long
Description	Failed to allocate snapshot reserve on the remote site.
Alert message	Failed to allocate snapshot reserve on the remote site '*remote_name*' in *duration_mins* minutes.
Cause	Snapshot reserve allocation on the remote site timed out.
Impact	Failed to meet less than 15 minute RPO for the remote site.
Resolution	Contact Nutanix support
Severity	Warning

Table 544: Protection Domain transitioning to lower frequency snapshotting. [130142] [A130142]

Name	Protection Domain transitioning to lower frequency snapshot schedule
Description	Protection Domain transitioning to lower frequency snapshot schedule.
Alert message	Protection Domain '*protection_domain_name*' has degraded to '*rpo_string*' RPO. *reason*.
Cause	High RPO cannot be sustained due to system constraints.
Cause	New entity has been protected in the protection domain which needs to be seeded to the remote.
Impact	All ongoing high frequency replications will be aborted and RPO for data protection will be affected.
Resolution	Protection Domain will automatically attempt to transition to a higher RPO after system constraints are resolved.
Resolution	Protection Domain will automatically attempt to transition to a higher RPO after seeding for newly added entities completes.
Severity	Warning
Severity	Info

Table 545: Protection Domain falling back to lower frequency replications to remote. [130143] [A130143]

Name	Protection Domain fallback to lower frequency replications to remote
Description	Protection Domain fallback to lower frequency replications to remote.

Alert message	Protection Domain '*protection_domain_name*' has degraded to '*rpo_string*' RPO to remote '*remote_name*'. *reason*.
Cause	High frequency replications to remote site cannot be sustained due to system constraints.
Impact	Recovery Points available on the remote for Data Protection will be affected
Resolution	Protection Domain will automatically attempt to transition to higher frequency replications when resources are available.
Severity	Warning

Table 546: Protection Domain Transition to Higher Frequency Snapshotting. [130144] [A130144]

Name	Protection Domain transitioning to higher frequency snapshot schedule
Description	Protection domain transitioning back to the configured RPO.
Alert message	Protection Domain '*protection_domain_name*' has been upgraded to *rpo_string* RPO.
Cause	High frequency snapshot schedule has been enabled.
Impact	RPO for Data Protection will be met.
Resolution	Reduce snapshot schedule frequency or delete the schedule if high frequency snapshots are not needed.
Severity	Info

Table 547: Snapshot queued for replications to remote site [130145] [A130145]

Name	Protection Domain Snapshots Queued for Replication to Remote Site
Description	Protection domain {protection_domain_name} has one or more snapshots queued for replication to remote site {remote_name}
Alert message	Protection domain *protection_domain_name* has one or more snapshots queued for replication to remote site *remote_name*
Cause	Various
Impact	One or more snapshots cannot be replicated to the remote site to meet the configured recovery point objective
Resolution	Maintain enough cross-cluster network bandwidth and throttle primary IO
Severity	Warning

Table 548: Protection Domain {protection_domain_name} Seeding Process Pending to the Remote Site {remote_name} [130146] [A130146]

Name	Protection Domain Seeding Process Pending to Remote Site
Description	Protection Domain {protection_domain_name} will maintain {rpo_string} RPO until seeding process completes to the remote_site {remote_name}

Alert message	The Protection Domain *protection_domain_name* will maintain *rpo_string* RPO until seeding process completes to the remote_site *remote_name*. The seeding process gets completed when snapshot reserve has been allocated on the local and remote site
Cause	Various
Impact	The Protection Domain {protection_domain_name} will maintain {rpo_string} RPO until seeding process completes to the remote_site {remote_name}
Resolution	Maintain enough network bandwidth to help the seeding transfer to succeed. Protection domain will wait for the seeding transfer and then automatically enable the configured RPO.
Severity	Warning

Table 549: VM registration failed [130147] [A130147]

Name	VM Registration Failure
Description	VM Registration Failure
Alert message	Registration of virtual machine, *vm_name*, failed because *reason*
Cause	This might happen for several reasons. Check the alert message for more details.
Impact	VMs cannot be registered on the cluster.
Resolution	Verify the existence of the VM in the snapshot and check whether any errors are reported in the hypervisor management software. Check the VM version and the maximum VM version supported by the host. If you cannot resolve the error, contact Nutanix support
Severity	Critical
Severity	Warning

Table 550: Snapshot Creation Delayed [130148] [A130148]

Name	Protection Domain Snapshots Delayed
Description	Snapshot creation for the protection domain {protection_domain_name} has been delayed
Alert message	Snapshot creation for the protection domain *protection_domain_name* has been delayed.
Cause	Many entities might be scheduled for the snapshot creation at the same time.
Impact	One or more snapshots cannot be created and hence cannot be replicated to the remote sites to meet the configured recovery point objective
Resolution	Stagger the schedules for different protected entities. Additionally, you can throttle other administrative operations in the cluster.
Severity	Warning

Table 551: VM Recovery Point Creation Failed. [130157] [A130157]

Name	VM Recovery Point Failed
Description	VM Recovery Point creation failed.
Alert message	Failed to capture the Recovery Point for VM '*vm_name*'.
Cause	Nutanix Data Protection service error.
Impact	VM Recovery Point will not be available to recover the VM.
Resolution	Retry the VM Recovery Point operation. If issue persists please contact Nutanix Support.
Severity	Warning

Table 552: VM Protection Failed. [130158] [A130158]

Name	VM Protection Failed
Description	VM protection failed.
Alert message	Failed to protect VM: '*vm_name*', because of '*reason*'.
Cause	Nutanix Data Protection service error.
Impact	VM cannot be recovered in case of disaster.
Resolution	VM protection will get retriggered. In case if this situation persists, please contact nutanix support.
KB Article	8281
Severity	Warning

Table 553: Protection Rule Conflict Occurred. [130161] [A130161]

Name	Protection Rule Conflicts
Description	Unable to protect VM due to conflicting protection rules.
Alert message	Failed to apply the protection rule on the VM '*vm_name*', because of '*reason*'.
Cause	One or more protection rules protecting the same entity.
Impact	VM cannot be recovered if disaster occurs.
Resolution	Resolve all the conflicts between the protection rules or modify the VM to ensure that it is protected under one protection rule.
KB Article	8275
Severity	Warning

Table 554: Recovery Point Replication Skipped. [130162] [A130162]

Name	Recovery Point Replication Skipped
Description	Recovery Point Replication Skipped.

Alert message	Skipped replication of Recovery point created at '*recovery_point_create_time*' UTC of the VM: '*vm_name*' because of '*reason*'.
Cause	Latest recovery point is in the queue for the replication.
Impact	One of the recovery point replication will be skipped.
Resolution	Trigger the replication of the snapshot manually.
Severity	Warning

Table 555: Failed to find the Target cluster for Replication [130163] [A130163]

Name	Fetch Replication Target Failure
Description	Cannot find a target cluster for replication on the Remote Availability Zone.
Alert message	Cannot find the replication target on the Recovery Location '*availability_zone_physical_name*'.
Cause	None of the attached clusters support recovery.
Cause	Virtual IP is not set on the remote cluster.
Cause	Target cluster could not be found for replication due to some internal error.
Impact	Replication of recovery points to this remote Availability Zone will be affected.
Resolution	Ensure that at least one of the clusters, attached to the Remote Availability Zone, supports recovery.
Resolution	Configure the virtual IP address on the remote cluster.
Resolution	Retry the operation. Contact Nutanix support if the issue persists.
Severity	Warning

Table 556: VM Replication Failure [130164] [A130164]

Name	VM Replication Failure
Description	VM Replication Failure
Alert message	Replication of VM Recovery Point created at *recovery_point_create_time* UTC for VM *vm_name* to the Recovery Location *availability_zone_physical_name* failed.
Cause	(A) Remote cluster is in read only state. (B) Remote Availability Zone does not support PowerPC architecture. (C) Replication of Recovery Point with dedup entity to a Recovery Availability Zone without dedup support is not allowed. (D) Remote Availability Zone does not support Volume Groups. (E) Replication of Recovery Point with VMs having attached File Level Restore (FLR) disks to a Recovery Availability Zone without FLR support is not allowed.
Impact	VM Recovery Point was not replicated to the Recovery Location.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 557: Application-Consistent Recovery Point Failed [130165] [A130165]

Name	VSS Recovery Point Status
Description	Application-consistent Recovery Point failed.
Alert message	Application-Consistent Recovery Point created at *recovery_point_create_time* UTC for the VM *vm_name* failed because *reason*.
Cause	Failed to quiesce the VM.
Impact	Crash-consistent Recovery Point is taken instead of application-consistent.
Resolution	Review the logs in the guest VM and take necessary actions. If issue persists contact Nutanix Support.
Severity	Warning

Table 558: Recovery Point expired prior to start of replication. [130166] [A130166]

Name	VM Replication Expiry Status
Description	Recovery Point expired prior to start of replication.
Alert message	Replication of Recovery Point created at *recovery_point_create_time* UTC for VM *vm_name* to the recovery location *az_name* expired prior to start of replication.
Cause	Previous Recovery Point replication took longer than RPO.
Impact	The Recovery Point that was scheduled to be replicated could not be replicated. Data that is expected to be protected may not be protected.
Resolution	Review the RPO and retention time of the Recovery Point based on available bandwidth.
Severity	Warning

Table 559: VM Recovery Point Replication Time Exceeded the RPO [130167] [A130167]

Name	VM Recovery Point Replication Time Exceeds RPO
Description	VM Recovery Point Replication Time Exceeded the RPO Limit.
Alert message	Replication time of the snapshot created at '*recovery_point_create_time*' UTC for VM *vm_name* exceeded the RPO limits.
Cause	Latest replication took more time than the RPO limit.
Impact	Replications for the next snapshots will be delayed.
Resolution	Please check your network.
Severity	Warning

Table 560: NGT on VM is not reachable [130168] [A130168]

Name	NGT is unreachable
Description	NGT on the VM is not reachable.

Alert message	Communication link to Nutanix Guest Tools on the VM *vm_name* is not working.
Cause	Communication link to Nutanix Guest Tools on the VM is not working.
Impact	Crash consistent snapshot will be taken instead of application consistent snapshot and static IP address preservation/mapping will fail.
Resolution	Restart the Nutanix Guest Tools service in the VM.
Severity	Warning

Table 561: VM Replication has not progressed. [130169] [A130169]

Name	VM Recovery Point Replication Stuck
Description	Replication of the VM has not progressed.
Alert message	Replication of the Recovery Point created at *recovery_point_create_time* UTC for VM *vm_name* to Recovery Location *az_name* has not made any progress in last one hour.
Cause	Network Connectivity issues between the Source and the recovery location.
Cause	Recovery Location might not have enough free disk space available.
Cause	Single-node cluster might be in read-only mode, due to an SSD failure.
Impact	VM Recovery Point will not be available at the Recovery Location.
Resolution	Check network connection between the Source and the Recovery Location.
Resolution	Make sure that the Recovery Location has enough disk space available.
Resolution	Replace the failed SSD.
Severity	Warning

Table 562: Nutanix Guest Tools Not Installed on the VM [130170] [A130170]

Name	VM NGT Configuration
Description	Nutanix Guest Tools not installed.
Alert message	Nutanix VSS is enabled but Nutanix Guest Tools is not installed on the VM *vm_name*.
Cause	Nutanix VSS is enabled but Nutanix Guest Tools is not installed on the VM.
Impact	Application Consistent Recovery Points will fail. Instead Crash consistent recovery point will be taken.
Resolution	Install Nutanix Guest Tools on the guest VM.
Severity	Warning

Table 563: Nutanix VSS provider or (pre_freeze/post_thaw) Scripts Not Installed. [130171] [A130171]

Name	VSS Scripts Not Installed On VM
Description	Nutanix VSS provider or pre_freeze/post_thaw Scripts Not Installed.

Alert message	VSS is enabled, but Nutanix VSS provider or pre_freeze/post_thaw scripts are not installed on the VM *vm_name*.
Cause	VSS is enabled, but Nutanix VSS provider or pre_freeze/post_thaw scripts are not installed on the VM.
Impact	Crash-consistent Recovery Point will be taken instead of application-consistent Recovery Point.
Resolution	Install Nutanix VSS provider or (pre_freeze/post_thaw) scripts on the VM.
Severity	Warning

Table 564: Degraded VM Recovery Point. [130172] [A130172]

Name	Partial VM Recovery Point
Description	Partial VM Recovery Point.
Alert message	Recovery Point for VM *vm_name* failed to capture associated policies and categories because *reason*.
Cause	Recovery Point for VM failed to capture associated policies and categories.
Impact	VM will be restored without associated policies and categories.
Resolution	Check the connection between PC and the PE, then retry the Recovery Point operation.
Severity	Info

Table 565: Recovery Point Replication Failed [130174] [A130174]

Name	Recovery Point Replication Failed
Description	Recovery Point replication failed.
Alert message	Failed to replicate recovery point created at: '*recovery_point_create_time*' UTC of the VM: '*vm_name*' to the recovery location: *availability_zone_physical_name*.
Cause	Network Connectivity issues between the Primary and the Recovery Availability Zone.
Cause	Data Protection and Replication service is not working as expected. The service could be down.
Cause	VM migration process is in progress.
Cause	If NearSync is enabled, the container name on remote cluster should match the one on source cluster hosting the VM.
Cause	Virtual IP address has not been configured on the remote cluster.
Cause	Remote clusters are unhealthy.
Cause	Recovery point was created for an entity having datasource backed virtual disks.

Cause	(A) Remote cluster is in read only state. (B) Remote Availability Zone does not support PowerPC architecture. (C) Replication of Recovery Point with dedup entity to a Recovery Availability Zone without dedup support is not allowed. (D) Remote Availability Zone does not support Volume Groups. (E) Replication of Recovery Point with VMs having attached File Level Restore (FLR) disks to a Recovery Availability Zone without FLR support is not allowed.
Cause	Replication to the remote site have been suspended for an entity by the user through CLI.
Impact	Recovery Point will not be replicated to the recovery location.
Resolution	Check network connection between the Primary and the Recovery Availability Zone.
Resolution	Contact Nutanix support.
Resolution	Retry the Recovery Point replication operation after the migration is complete.
Resolution	Create container(s) on which VM is running with the same name on the remote cluster.
Resolution	Configure the Virtual IP address and then retry the Recovery Point replication operation.
Resolution	For a manually taken Recovery Point, retry replication again. For a scheduled Recovery Point replication, ensure all the emote clusters are healthy. Then wait for the scheduled Recovery Point replication.
Resolution	Retry the Recovery Point replication operation after storage migration for the datasource backed virtual disks is complete.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Resolution	Remove the suspended replication timer associated with the entity using the CLI.
KB Article	8460
Severity	Warning

Table 566: Invalid Availability Zone details configured in Protection Policy. [130177] [A130177]

Name	Availability Zone Validation Failed
Description	Availability Zone referenced in the Protection Policy is not Valid.
Alert message	Protection Policy *protection_rule_name* has invalid Availability Zone configured.
Cause	Invalid Availability Zone configured in the Protection Policy.
Impact	Protection Policy will not be processed.
Resolution	Update the Availability Zone in Protection Policy.
Severity	Warning

Table 567: Entity Unprotection Failed [130178] [A130178]

Name	Entity Unprotection Failed
Description	Unable to unprotect the VM because of some internal error.
Alert message	Failed to unprotect the VM: '*vm_name*', because of '*reason*'.
Cause	Data Protection service may be down.
Impact	VM will remain protected.
Resolution	If this issue persists for a long time, contact Nutanix support for further assistance.
Severity	Warning

Table 568: Data Protection tasks are not progressing. [130179] [A130179]

Name	Cerebro tasks Stuck
Description	Replication tasks are not progressing.
Alert message	Multiple tasks are not making progress on PE:*cluster_name* because of long running replication tasks.
Cause	Replication tasks are running for long.
Impact	Tasks can stuck forever.
Resolution	Please reach out to Nutanix Support to help abort the long running tasks.
Severity	Warning

Table 569: Recovery Location is in Read Only mode [130180] [A130180]

Name	Entity Centric Remote Site Operation Mode ReadOnly
Description	Operation Mode of the Recovery Location changed to Read Only mode.
Alert message	Availability zone: *availability_zone_physical_name* is in read-only state. Replications to this site will fail.
Cause	Single node or two-node clusters: Possible SSD failure on the Recovery Location.
Impact	Replications to this Recovery Location will be aborted.
Resolution	Single node or two-node clusters: Replace the failed SSD on the recovery location.
KB Article	8324
Severity	Warning

Table 570: Snapshot has Entities from the Storage Container that have Deduplication Enabled [130181] [A130181]

Name	Replication Of Deduped Entity

Description	Snapshot contains entities from the storage container that have deduplication enabled.
Alert message	Do not replicate the recovery point comprising entities from the storage container that have deduplication enabled to the recovery location : *availability_zone_physical_name*.
Cause	A snapshot with entities from the storage container that have deduplication enabled was replicated to a single node backup or cloud recovery location.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable the deduplication on the storage containers or use a multi-node AOS cluster as a replication target.
Severity	Warning

Table 571: Nutanix Guest Tools not Upgraded [130182] [A130182]

Name	NGT needs upgrade
Description	Nutanix Guest Tools not upgraded.
Alert message	Nutanix Guest Tools are not upgraded to version *version* or higher on the guest VM(s) *vm_names* protected by *protection_domain_name*.
Cause	VSS is enabled, but Nutanix Guest Tools are not upgraded on the VM(s).
Impact	Crash-consistent snapshot will be taken instead of application consistent snapshot.
Resolution	Install the latest version of Nutanix Guest Tools available on the cluster on the guest VM(s).
Severity	Warning

Table 572: Protected VM(s) May Not Be Recoverable [130183] [A130183]

Name	Protected VMs May Not Be Recoverable
Description	Protected VM(s) May Not Be Recoverable
Alert message	VM(s) *vm_names* protected by protection domain '*protection_domain_name*' are running from VMWare snapshots. Recovery of the VMs from the PD snapshot '*snapshot_id*' may not succeed. Manual intervention is required to boot up the VMs.
Cause	Protected VM(s) are running from VMWare snapshots and may not be recoverable.
Impact	Recovery of the VMs from the PD snapshot may not succeed
Resolution	Ensure that the VMs are not running from VMWare snapshot at the time of scheduled Protection Domain snapshot. Refer Nutanix KB: 7023 for further details.
Severity	Warning

Table 573: VM Protection Failed [130186] [A130186]

Name	VM Protection Failed
Description	Failed to protect the VM by using the defined Protection Policy.
Alert message	Failed to apply protection policy '*protection_rule_name*' on the VM '*vm_name*' because '*reason*'
Cause	VM is already protected in a legacy protection domain.
Cause	Cluster hosting the VM does not support VM in multiple schedules or replication to more than one target site.
Cause	(A) VM has volume group attachments. (B) VM is CBR incapable. (C) Source or the target cluster does not support replication within the same Availability Zone.
Impact	VM will not be protected.
Resolution	Unprotect the VM from the Protection Domain to protect it using Entity Centric Workflow (Leap).
Resolution	Upgrade cluster to version 5.19 or above.
Resolution	Resolve the stated reason for the failure.
KB Article	8472
Severity	Warning

Table 574: VM Recovery May Fail [130188] [A130188]

Name	VM Recovery May Fail
Description	VM may not start on the recovery location.
Alert message	VM *vm_name* may not start on the recovery location '*availability_zone_name*'
Cause	Primary and Remote locations are on different hypervisors and VM does not have NGT information.
Cause	VM uses vNUMA. Cannot recover vNUMA VMs on Xi.
Impact	VM will not start on the recovery location.
Resolution	Install NGT on the VM (if not already installed).
Resolution	Don't replicate VMs which use vNUMA to Xi.
Severity	Warning

Table 575: Virtual IP Address not Configured on the Cluster. [130191] [A130191]

Name	Virtual IP not set on Cluster
Description	Virtual IP address is not configured on the cluster.
Alert message	Virtual IP address is not configured on the cluster *cluster_name*. Protection and replication of VMs will be affected.
Cause	Virtual IP address has not been configured on the cluster.
Impact	Protection and replication of VMs will be affected.

Resolution	Configure the virtual IP address on the cluster.
Severity	Warning

Table 576: Full Replication Started for the VM [130193] [A130193]

Name	VM Recovery Point Full Replication
Description	A full replication was triggered for the VM.
Alert message	Replication for the VM *vm_name* recovery point with the create time *recovery_point_create_time* to recovery location *az_name* will be a full replication.
Cause	A reference recovery point could not be found.
Impact	Replication will take longer to finish.
Resolution	No action is necessary.
Severity	Info

Table 577: Protected VM Not Recoverable [130195] [A130195]

Name	VM Not Recoverable
Description	Protected VM may not be recoverable.
Alert message	VM *vm_name* is running from the VMware snapshots. Recovery of the VM from the Recovery Point created at *recovery_point_create_time* UTC may not succeed. Manual intervention is required to boot up the VM.
Cause	Protected VM is running from the VMware snapshot and may not be recoverable.
Impact	Recovery of the VM from the Recovery Point may not succeed.
Resolution	Ensure that the VM is not running from the VMware snapshot at the time of scheduled Recovery Point.
Severity	Warning

Table 578: VM Virtual Hardware Version Incompatible [130196] [A130196]

Name	VM Virtual Hardware Version Incompatible
Description	VM Virtual Hardware Version Incompatible
Alert message	Virtual hardware version of the VM *vm_name* is not supported on the availability zone *availability_zone_name*.
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote Availability Zone.
Impact	VM may not register properly on the remote Availability Zone on restore/ clone.

Resolution	Upgrade the hypervisor version on the remote Availability Zone to support the virtual hardware version of the VM.
Severity	Warning

Table 579: VM Restoration Failed. [130197] [A130197]

Name	VM Restoration Failed
Description	VM Restoration Failed.
Alert message	Failed to restore VM: '*vm_name*' due to *reason*.
Cause	Snapshot is not present.
Cause	VM Restoration failed due to Internal error.
Impact	VM will not be restored.
Resolution	Retry the VM Restoration operation with different snapshot.
Resolution	Retry the VM Restoration operation.
KB Article	8487
Severity	Critical

Table 580: VM Migration Failed [130199] [A130199]

Name	VM Migration Failed
Description	VM Migration Failed.
Alert message	Failed to migrate VM: '*vm_name*' as the user may have manually powered ON the VM on the source after migration process powered it off or due to an internal error.
Cause	User may have manually powered ON the VM on the source after migration process powered it off.
Impact	VM will not be migrated.
Resolution	Retry the VM Migration operation. If this situation persists, contact Nutanix support.
KB Article	8520
Severity	Critical

Table 581: VSS Snapshot is not supported for some VMs. [130200] [A130200]

Name	VSS Snapshot Not Supported
Description	VSS snapshot is not supported for some VMs.
Alert message	VSS snapshot is not supported for the following VMs protected by Protection Domain '*protection_domain_name*'. VMs list: '*vm_names*'
Cause	Some VMs in the Protected Domain have unsupported configuration.

Impact	Hypervisor based application consistent snapshot is taken on ESX. Crash consistent snapshot is taken for other hypervisors.
Resolution	For more details, please review the alerts generated at Data Protection page and fix the unsupported configuration.
Severity	Warning

Table 582: Recovery Location is not in Good State [130203] [A130203]

Name	Recovery Location Not In Good State
Description	Recovery location is not in good state.
Alert message	Recovery location: *availability_zone_physical_name* is not in good state. Replications to this site will fail.
Cause	If Recovery Availability Zone is different from the Primary Availability Zone, check Network connectivity issues between the two Availability Zones.
Cause	Data Protection and Replication service is not working as expected. The service could be down.
Cause	The Cluster(s) attached to the affected Availability Zone may have been destroyed and recreated.
Impact	Replications to this recovery location will fail.
Resolution	Check network connection between the Primary and the Recovery Availability Zone.
Resolution	Contact Nutanix support.
Resolution	Disconnect and reconnect the Recovery Availability Zone.
Severity	Warning

Table 583: Entity Transition to a Higher Frequency Snapshot Schedule. [130204] [A130204]

Name	Entity transitioning to higher frequency snapshot schedule
Description	Entity is being transitioned to a higher frequency snapshot schedule.
Alert message	Entity '*vm_name*' has been upgraded to *rpo_string* RPO.
Cause	High frequency snapshot schedule has been enabled and higher RPO can be sustained as the system is stable.
Impact	RPO for Data Protection will be affected.
Resolution	If high-frequency snapshots are not required, reduce the snapshot schedule frequency in the Protection Rule or remove the VM from the Protection Rule with high frequency schedule.
Severity	Info

Table 584: Entity is being transitioned to a lower frequency snapshot schedule. [130205] [A130205]

Name	Entity transitioning to lower frequency snapshot schedule
Description	Entity is being transitioned to a lower frequency snapshot schedule.
Alert message	Entity *vm_name* has degraded to *rpo_string* RPO. *reason*.
Cause	High RPO cannot be sustained due to system constraints.
Impact	All ongoing high frequency replications will be aborted and RPO for data protection will be affected.
Resolution	Entity will automatically attempt to transition to a higher RPO after system constraints are resolved.
Severity	Warning

Table 585: Subnet configured for Remote Site mis-matches with remote cluster subnet [130206] [A130206]

Name	Remote Subnet Mismatch With Provided Value
Description	Subnet configured for Remote Site mis-matches with remote cluster subnet.
Alert message	Remote *remote_name* is configured with the subnet *user_specified_subnet*, which is not matching with the subnet *remote_returned_subnet* of the remote cluster.
Cause	Subnet specified in the remote site is not matching with the subnet of the remote cluster.
Impact	Improper firewall and routing settings.
Resolution	Reconfigure the remote subnet parameter.
KB Article	10276
Severity	Warning

Table 586: Hosting of Virtual IP of the Network Segmentation DR Service Failed. [130207] [A130207]

Name	Failed Hosting Network Segmented VIP
Description	Hosting of virtual IP of the DR service network segmentation failed.
Alert message	Virtual IP *virtual_ip* of the DR service network segmentation could not be hosted on the network interface *network_interface*.
Cause	Network configuration might have been changed or has not been configured properly.
Impact	Replication of the protected entities will be affected.
Resolution	Reconfigure the DR service network configuration.
KB Article	8258
Severity	Critical

Table 587: Snapshot Replication to Remote Site is Lagging. [130210] [A130210]

Name	Recovery Point Lagging Remote Replication
Description	Replication of the snapshot to the remote cluster is lagging.
Alert message	Replication of the recovery point for the VM '*vm_name*'to the remote site '*remote_az*' is lagging by '*lagging_minutes*' minutes.
Cause	Clusters are not sized properly according to the workloads that are running on them.
Cause	Recovery Point Replication requires more bandwidth than the allocated bandwidth.
Impact	Recovery point objective for the protection domain will be affected.
Resolution	Make sure that the clusters are sized properly according to the workload.
Resolution	Maintain sufficient network bandwidth for successful seeding transfer.
Severity	Warning
Severity	Info

Table 588: VM has been Recovered at an Alternate Path [130211] [A130211]

Name	VM Recovered at Alternate Path
Description	VM has been Recovered at an Alternate Path
Alert message	The following files of the VM *old_vm_name* conflict with the files present on the target: *conflicting_file_paths*. The VM *new_vm_name* has been restored at an alternate path by adding the prefix '*path_prefix*' after the container name to its file paths.
Cause	Existing VM files conflict with the VM files that you are recovering.
Impact	The VM with the conflicting files has been restored at an alternate path.
Resolution	The VM containing the conflicting files has been recovered at an alternate path. Clean up the conflicting stale files.
Severity	Warning

Table 589: Cluster does not support Synchronous Replication. [130212] [A130212]

Name	Cluster does not support Synchronous Replication
Description	Cluster does not support Synchronous Replication.
Alert message	Protection Rule: '*protection_rule_name*' is processed but following issue is observed - '*reason*'.
Cause	The cluster given in the Protection Rule does not support Synchronous Replication.
Impact	Protection Rule might not be processed.
Resolution	Please use a cluster that supports Synchronous Replication.
Severity	Critical

Table 590: Failed to transition to a higher frequency snapshot schedule. [130213] [A130213]

Name	Failed to transition to higher frequency snapshot schedule
Description	Protection Domain failed to transition to a higher frequency snapshot schedule.
Alert message	Failed to enable nearsync for *protection_domain_name*. *reason*.
Cause	Low RPO cannot be sustained due to system constraints.
Impact	Failed to enable high frequency replications and configured RPO for data protection will be affected.
Resolution	Protection Domain will automatically attempt to transition to a high frequency snapshot schedule after system constraints are resolved.
Severity	Warning

Table 591: Nearsync Replication is lagging for Protection Domain snapshot [130214] [A130214]

Name	Nearsync Replication is Lagging
Description	Replication of the nearsync protection domain snapshot is lagging.
Alert message	Nearsync replication for *protection_domain_name* to remote *remote_name* is lagging by *replication_lagging_seconds* seconds.
Cause	Replication of the protection domain nearsync snapshot is slow.
Impact	Configured RPO to the remote might get affected.
Resolution	Ensure that the number of entities are within the supported limit.
Severity	Warning

Table 592: VM SyncRep Container not found [130335] [A130335]

Name	VM SyncRep Container not found
Description	VM SyncRep container is not found at the target cluster
Alert message	Failed to replicate the VM: '*vm_name*' synchronously since '*details*'.
Cause	Container does not exist or it is marked to be removed at the target cluster.
Impact	VM will not be synchronously replicated.
Resolution	Retry the VM SyncRep operation after creating the missing container at target cluster. If this situation persists, contact Nutanix support.
Severity	Critical

Table 593: Cluster does not support the planned failover of synchronously replicated entities. [130339] [A130339]

Name	The cluster does not support the planned failover of synchronously replicated entities

Description	Cluster does not support the planned failover of synchronously replicated entities.
Alert message	Synchronous Replication is enabled for the entity *vm_name* but cluster *cluster_name* does not support planned failover.
Cause	The cluster does not support the planned failover of synchronously replicated entities.
Impact	Planned failover of the entity will fail.
Resolution	Please upgrade the source and target clusters to a version later than 5.17.1. It is recommended to have source and target clusters on the same version.
Severity	Warning

Table 594: VM Protection Might Fail Post Recovery [130341] [A130341]

Name	VM Protection Might Fail Post Recovery
Description	VM Protection Might Fail Post Recovery
Alert message	After recovery to remote availability zone *remote_availability_zone_name*, VM protection might fail because remote cluster *remote_cluster_name* does not support multisite protection.
Cause	The remote cluster to which replication is happening does not support multisite protection.
Impact	After recovery to the specified remote site, VM protection might fail.
Resolution	Upgrade the remote Prism Element to AOS 5.19 or above, which supports multisite protection.
Severity	Warning

Table 595: Remote Site network connectivity unstable [130342] [A130342]

Name	Remote Site Connectivity Unstable
Description	Remote Site network connectivity is unstable
Alert message	Network connectivity to remote site '*remote_name*' is unstable
Cause	Connection to remote site is unstable
Impact	VMs in Metro Protection Domain can experience intermittent I/O failures due to OS disk timeouts
Resolution	Check network connectivity between local and remote sites and manually break metro if VM I/O is stalled
KB Article	9713
Severity	Warning

Table 596: Volume Group Recovery Point Creation Failed. [130343] [A130343]

Name	Volume Group Recovery Point Failed

Description	Volume Group Recovery Point creation failed.
Alert message	Failed to capture the Recovery Point for Volume Group '*volume_group_name*'.
Cause	Nutanix Data Protection service error occurred.
Impact	Volume Group Recovery Point will not be available to recover the Volume Group.
Resolution	Retry the Volume Group Recovery Point operation. If issue persists please contact Nutanix Support.
Severity	Warning

Table 597: Cluster does not support the live migration of synchronously replicated entities. [130345] [A130345]

Name	The cluster does not support the live migration of synchronously replicated entities
Description	Cluster does not support the live migration of synchronously replicated entities.
Alert message	Synchronous Replication is enabled for the entity *vm_name* but cluster *cluster_name* does not support live migration.
Cause	The cluster does not support the live migration of synchronously replicated entities.
Impact	Live migration of the entity will fail.
Resolution	Please upgrade the source and target clusters to version 5.19 or later. It is recommended to have source and target clusters on the same version.
Severity	Warning

Table 598: Consistency Group Conflict Occurred. [130349] [A130349]

Name	Consistency Group VM Conflicts
Description	Unable to protect VM due to conflicting protection policy and Consistency Group configurations.
Alert message	Failed to protect the VM *vm_name* belonging to the Consistency Group *consistency_group_name*, because of *reason*.
Cause	Multiple protection policies are applied on some of the members present in the Consistency Group.
Cause	No protection policy is applied to protect some of the members present in the Consistency Group.
Cause	Members in the same Consistency Group have different protection policies that can be applied.
Cause	VM belongs to a Consistency Group that contains VM(s) or Volume Group(s) which have been deleted.
Impact	VM cannot be recovered if disaster occurs.

Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that the members in a Consistency Group are protected by the same protection policy.
Resolution	Remove the stale/deleted VM(s) or Volume Group(s) from the Consistency Group.
KB Article	10571
Severity	Warning

Table 599: Consistency Group Conflict Occurred. [130350] [A130350]

Name	Consistency Group VG Conflicts
Description	Unable to protect Volume Group due to conflicting protection policy and Consistency Group configurations.
Alert message	Failed to protect the Volume Group *volume_group_config_name* belonging to the Consistency Group *consistency_group_name*, because of *reason*.
Cause	Multiple protection policies are applied on some of the members present in the Consistency Group.
Cause	No protection policy is applied to protect some of the members present in the Consistency Group.
Cause	Members in the same Consistency Group have different protection policies that can be applied.
Cause	Volume Group belongs to a Consistency Group that contains VM(s) or Volume Group(s) which have been deleted.
Impact	Volume Group cannot be recovered if disaster occurs.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that exactly one protection policy is applied to protect the defined members in a Consistency Group.
Resolution	Ensure that the members in a Consistency Group are protected by the same protection policy.
Resolution	Remove the stale/deleted VM(s) or Volume Group(s) from the Consistency Group.
KB Article	10572
Severity	Warning

Table 600: Volume Group Protection Failed. [130351] [A130351]

Name	Volume Group Protection Failed
Description	Volume Group protection failed.

Alert message	Failed to protect Volume Group: '*volume_group_name*', because of '*reason*'.
Cause	Nutanix Data Protection service error occurred.
Impact	Volume Group cannot be recovered in case of disaster.
Resolution	Volume Group protection will get retriggered. If issue persists, please contact nutanix support.
KB Article	10560
Severity	Warning

Table 601: Volume Group Protection Might Fail Post Recovery [130352] [A130352]

Name	Volume Group Protection Might Fail Post Recovery
Description	Volume Group Protection Might Fail Post Recovery
Alert message	After recovery to remote availability zone *remote_availability_zone_name*, Volume Group protection might fail because remote cluster *remote_cluster_name* does not support multisite protection.
Cause	The remote cluster to which replication is happening does not support multisite protection.
Impact	After recovery to the specified remote site, Volume Group protection might fail.
Resolution	Upgrade the remote Prism Element to AOS 5.19 or above, which supports multisite protection.
KB Article	10573
Severity	Warning

Table 602: Volume Group Recovery Point Replication Failed. [130355] [A130355]

Name	Volume Group Recovery Point Replication Failed
Description	Volume Group Recovery Point replication failed.
Alert message	Failed to replicate recovery point created at: *recovery_point_create_time* UTC of the volume group: *volume_group_name* to the recovery location: *availability_zone_physical_name*.
Cause	Network Connectivity issues between the Primary and the Recovery Availability Zone.
Cause	Data Protection and Replication service is not working as expected. The service could be down.
Cause	Volume Group migration process is in progress.
Cause	Virtual IP address has not been configured on the remote cluster.
Cause	Remote clusters are unhealthy.
Impact	Volume Group Recovery Point will not be replicated to the recovery location.
Resolution	Check network connection between the Primary and the Recovery Availability Zone.

Resolution	Contact Nutanix support.
Resolution	Retry the Recovery Point replication operation after the migration is complete.
Resolution	Configure the Virtual IP address and then retry the Recovery Point replication operation.
Resolution	For a manually initiated Volume Group Recovery Point replication, retry again. For a scheduled Volume Group Recovery Point replication, ensure all the emote clusters are healthy. Then wait for the next scheduled Recovery Point replication.
KB Article	10754
Severity	Warning

Table 603: Volume Group Protection Failed [130357] [A130357]

Name	Volume Group Protection Failed On PC
Description	Volume Group protection failed.
Alert message	Failed to protect Volume Group: '*volume_group_config_name*', because of '*reason*'.
Cause	Data Protection service error occurred.
Cause	Volume Group is already protected in a legacy protection domain.
Cause	The Protection Policy applied on Volume Group has a NearSync or a Synchronous snapshot schedule applied on it.
Impact	Volume Group cannot be recovered in case of disaster.
Resolution	Volume Group protection will get retriggered. If issue persists, please contact Nutanix Support.
Resolution	Unprotect the Volume Group from the Protection Domain to protect it using Leap.
Resolution	NearSync and Synchronous protection is not supported on Volume Groups. Therefore, ensure that the Protection Policy protecting the Volume Groups does not have NearSync or Synchronous snapshot schedules
KB Article	11308
Severity	Warning

Table 604: Consistency Group Contains Stale Members [130358] [A130358]

Name	Consistency Group Contains Stale Members
Description	Consistency Group contains one or more members that have been deleted.
Alert message	Consistency Group '*consistency_group_config_name*' contains stale members.
Cause	One or more VMs or Volume Groups belonging to the Consistency Group may have been deleted.
Impact	The VMs and Volume Groups part of the Consistency Group will not be protected.
Resolution	Remove the deleted members from the Consistency Group configuration.

KB Article	11413
Severity	Warning

Table 605: Synchronous Replication is not supported on the VM [130359] [A130359]

Name	Synchronous Replication is not supported on the VM
Description	Synchronous Replication is not supported on the VM
Alert message	Synchronous replication is not supported on the VM *vm_name*. Reason: *message*
Cause	VM has disks that are backed by a datasource external to the AOS cluster.
Impact	VM will not be synchronously replicated until the seeding from the external datasource is complete.
Resolution	Enable background storage migration on the disk backed by external datasource to complete the seeding.
KB Article	11344
Severity	Warning

Table 606: Partial Volume Group Recovery Point. [130361] [A130361]

Name	Partial Volume Group Recovery Point
Description	Partial Volume Group Recovery Point.
Alert message	Recovery Point for Volume Group *volume_group_name* failed to capture associated policies and categories because *reason*.
Cause	Recovery Point for Volume Group failed to capture associated policies and categories.
Impact	Volume Group will be restored without associated policies and categories.
Resolution	Check the connection between PC and the PE, then retry the Recovery Point operation.
KB Article	11744
Severity	Info

Table 607: Volume Group Recovery Point Replication Time Exceeded the RPO [130362] [A130362]

Name	Volume Group Recovery Point Replication Time Exceeds RPO
Description	Volume Group Recovery Point Replication Time Exceeded the RPO Limit.
Alert message	Replication time of the snapshot created at '*recovery_point_create_time*' UTC for Volume Group *volume_group_name* exceeded the RPO limits.
Cause	Latest replication took more time than the RPO limit.
Impact	Replications for the next snapshots will be delayed.

Resolution	Please check your network.
KB Article	11745
Severity	Warning

Table 608: Recovery Plan Execution Failure [300401] [A300401]

Name	Recovery Plan Execution Failure
Description	Execution of the Recovery Plan failed.
Alert message	*operation_type* failed for the Recovery Plan *recovery_plan_name*. One or more entities in this Recovery Plan might have failed to be started/recovered.
Cause	One or more entities were not recovered as part of the Recovery Plan execution.
Impact	One or more entities or applications in this Recovery Plan might not have started completely.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.
KB Article	8278
Severity	Critical

Table 609: Validation Warnings found during Recovery Plan Execution [300402] [A300402]

Name	Recovery Plan Execution Failure due to Validation Warnings
Description	Execution of the Recovery Plan failed.
Alert message	*operation_type* for Recovery plan *recovery_plan_name* resulted in validation warnings. You can choose to execute the Recovery Plan anyway or you can fix the issues.
Cause	The Recovery Plan validation resulted in warnings.
Impact	The Recovery Plan execution could not be started due to warnings in the validation report.
Resolution	Review the Recovery Plan validation report for detailed information.
Severity	Warning

Table 610: Recovery Plan Validation Failed With Errors [300408] [A300408]

Name	Validation Failed For Recovery Plan with Errors
Description	Recovery Plan validation failed with errors.
Alert message	Validation failed with errors for the Recovery Plan *recovery_plan_name*.
Cause	There might be various reasons for this issue.
Impact	Entities recovery might fail.

Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.
KB Article	8276
Severity	Critical

Table 611: Recovery Plan Validation Failed with Warnings [300409] [A300409]

Name	Validation Failed For Recovery Plan with Warnings
Description	Recovery Plan validation failed with warnings.
Alert message	Validation failed with warnings for the Recovery Plan *recovery_plan_name*.
Cause	There might be various reasons for this issue.
Impact	Entities recovery might fail.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.
Severity	Warning

Table 612: Planned Failover or Unplanned Failover operations will fail. [300411] [A300411]

Name	Delete the failed over entities from either the Primary or the Recovery Location
Description	Delete the failed over entities on either the primary or the recovery location.
Alert message	Unplanned Failover for the Recovery Plan *recovery_plan_name* has completed successfully. Validate and clean up the entities from the site where it is not required, for new unplanned or planned failover operations for this recovery plan to succeed.
Cause	Unplanned failover for the Recovery Plan completed successfully.
Impact	Planned Failover and Unplanned Failover operations are not possible without cleanup.
Resolution	Post unplanned failover entities are now active on the recovery Availability Zone. For synced entities, delete the entities on the primary Availability Zone to avoid performance issues. For entities protected by asynchronous replication, performance will not be affected and deletion can happen on either Availability Zone as required.
Severity	Warning

Table 613: Floating IPs Allocation Failed [300420] [A300420]

Name	Floating IPs Allocation Failure.
Description	Floating IPs allocation failed.
Alert message	Failed to allocate *num_failed_floating_ips* floating IPs, required for the dynamic allocation, for *operation_type* on the Recovery Plan *recovery_plan_name*.
Cause	The system may not have enough floating IP addresses.

Impact	VMs will come up without floating IPs.
Resolution	Try to manually assign pre-allocated floating IPs for the vNICs.
Severity	Warning

Table 614: Floating IPs Deallocation Failed [300421] [A300421]

Name	Floating IPs Deallocation Failure.
Description	Floating IPs deallocation failed.
Alert message	Failed to deallocate *num_failed_floating_ips* floating IPs, which were allocated dynamically but are not associated with any vNIC, for *operation_type* on the Recovery Plan *recovery_plan_name*.
Cause	The floating IPs might be in use.
Impact	Floating IPs will remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.
Severity	Warning

Table 615: Recovery Plan Execution Failure due to Validation Errors [300422] [A300422]

Name	Recovery Plan Execution Failure due to Validation Errors
Description	Execution of the Recovery Plan failed.
Alert message	The *operation_type* for the recovery plan *recovery_plan_name* could not be started due to errors in validation.
Cause	The Recovery Plan validation failed.
Impact	The Recovery Plan execution could not be started.
Resolution	Review the Recovery Plan validation report for detailed information.
KB Article	8277
Severity	Critical

Table 616: Floating IPs Deallocation Failed after Failback from Xi. [300427] [A300427]

Name	Unassociated Floating IPs Deallocation Failure after Failback from Xi.
Description	Failed to deallocate unassociated floating IPs after failback from Xi that were earlier dynamically created.
Alert message	Failed to deallocate unassociated floating IPs *failed_floating_ips* after failback from Xi that were earlier dynamically created, for *operation_type* on the Recovery Plan *recovery_plan_name*.
Cause	The floating IPs might be in use.
Impact	Floating IPs will continue to remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.
Severity	Warning

Entity Sync

Table 617: Entity Sync Failure [500101] [A500101]

Name	Entity Sync failed
Description	Entity Sync failed
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.

Guest VM

Table 618: Datastore VM Count High [3030] [A1170]

Name	Datastore VM Count Check
Description	Checks for high VM count on datastores.
Alert message	The number of VMs on Datastore: *datastore* is *vm_count*
Cause	High number of VMs on at least one datastore
Impact	HA service might be affected as VMware HA protects only 2048 VMs per datastore.
Resolution	Move some VMs to another datastore
VMs per Datastore Threshold	Warning Threshold: 1945 VMs per datastore Critical Threshold: 2048 VMs per datastore
KB Article	3012
Severity	Critical
Severity	Warning

Table 619: VM Transmit Packet Drop Check [3060]

Name	VM Transmit Packet Drop Check
Description	Checks whether transmitted packet drop rate is high.
Cause	Transmitted packet drop rate is high.
Impact	Application performance on the VM may be degraded.
Resolution	Check network configuration/connections.
KB Article	2733

Table 620: VM Receive Packet Drop Check [3061]

Name	VM Receive Packet Drop Check
Description	Checks whether received packet drop rate is high.
Cause	Received packet drop rate is high.
Impact	Application performance on the VM may be degraded.
Resolution	Check network configuration/connections.
KB Article	2733

Table 621: Orphan VM Snapshot Check [3200]

Name	Orphan VM Snapshot Check
Description	Checks for VM snapshots whose associated VMs are removed.
Cause	The VM associated with a snapshot has been removed.
Impact	Cluster storage resources are consumed by unneeded data.
Resolution	Delete the snapshots based on the VM that has been removed.
KB Article	3752

Table 622: VM has non-ASCII name. [106472] [A106472]

Name	Non-ASCII VM Names
Description	Check if any VMs have non-ASCII names.
Alert message	VM *vm_name* has non-ASCII name, which may cause snapshotting problems.
Cause	Some VMs have non-ASCII names.
Impact	Snapshot operations may fail.
Resolution	Update the VM names to consist only of ASCII characters.
KB Article	4255
Severity	Critical

Table 623: VM config updates disabled [110455]

Name	VM config updates disabled
Description	Checks for VMs that have configuration updates disabled
Cause	Update operations on a VM might fail after a failed migration
Impact	VM updates are blocked for the affected VMs
Resolution	Depending on the failure, a manual VM cleanup may be needed. Please refer to KB for detailed guidance
KB Article	9902

Table 624: Node Failure [130030] [A1137]

Name	Node Status
Description	Node failure.
Alert message	Host *hypervisor_address* appears to have failed. High Availability is restarting VMs on *failover_host_info*.
Cause	Host is not accessible.
Impact	VMs will restart on another host.
Resolution	Check for network connectivity or a hardware failure.
Severity	Critical
Severity	Warning

Table 625: Protected VM Not Found [130038] [A1010]

Name	Protected VM Not Found
Description	Protected VM not found.
Alert message	Unable to locate the VM '*vm_name* (ID *vm_id*)' protected by protection domain '*protection_domain_name*'.
Cause	The protected VM cannot be found and may have been deleted.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove the VM from the protection domain.
Severity	Info

Table 626: VM Action Error [130057] [A1033]

Name	VM Action Status
Description	VM action error.
Alert message	Failed to *action* VM with name '*vm_name* and internal ID (ID *vm_id*)' due to *reason*.
Cause	A VM could not be restored because of a hypervisor error, or could not be deleted because it is still in use.
Impact	The requested VM action (restore or delete) could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10658
Severity	Warning

Table 627: Failure To Restart VMs For HA Event [130064] [A1145]

Name	VMs Restart Status
Description	Failure to restart VMs for an HA event.
Alert message	Failed to restart one or more VMs that were running on failed host *hypervisor_address*.
Cause	Not enough memory or CPU resources within the cluster.
Impact	Some VMs will remain powered off.
Resolution	Shut down unneeded VMs to free cluster resources; expand the cluster to add resources; or wait until the failed host recovers. Once resources are available, power on the VMs.
KB Article	8199
Severity	Critical
Severity	Warning

Table 628: HA Host Evacuation Failure [130065] [A1146]

Name	Host Evacuation Status
Description	Failure to evacuate host while entering maintenance mode or reserving host for HA.
Alert message	Failed to evacuate some VMs from host *hypervisor_address* while *reason*.
Cause	Not enough memory/CPU resources within the cluster.
Impact	The host will not enter requested state.
Resolution	Shut down unneeded VMs to free cluster resources or expand the cluster to add resources.
Severity	Warning

Table 629: HA Healing Failure [130067] [A1155]

Name	VM HA Healing Status
Description	HA healing failure.
Alert message	Could not restore VM High Availability.
Cause	Not enough memory/CPU resources within the cluster.
Impact	VMs are no longer protected against host failure.
Resolution	Shut down unneeded VMs to free cluster resources, expand the cluster to add resources, or disable protection against host failures. VMs cannot be powered on or migrated until the system has sufficient resources to protect against the configured number of host failures.
Severity	Warning

Table 630: VM High Availability Failure [130068] [A130068]

Name	VM HA Status
Description	High availability failure.
Alert message	VMs are no longer protected against host failure. Reason: *reason*.
Cause	Not enough memory/CPU resources within the cluster.
Impact	VMs are no longer protected against host failure.
Resolution	Shut down unneeded VMs to free cluster resources, expand the cluster to add resources, or disable protection against host failures. VMs cannot be powered on or migrated until the system has sufficient resources to protect against the configured number of host failures.
Severity	Warning

Table 631: iSCSI Configuration Failed [130100] [A130100]

Name	iSCSI Configuration Failed
Description	iSCSI Configuration Failed
Alert message	Failed to re-configure iSCSI settings on the recovered VM '*vm_name*'. *reason*.
Cause	Nutanix Guest Tools failed to execute some iSCSI commands on the guest VM.
Impact	iSCSI disks may become unavailable on the guest VM.
Resolution	(A) If IQN and iSCSI target IP addresses of the VM have been updated by Nutanix Guest Agent, Discover and connect to new targets after rebooting the VM. (B) Manually configure iSCSI settings on the guest VM. (C) Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10118
Severity	Critical

Table 632: Failed to Mount Nutanix Guest Tools [130103] [A130103]

Name	NGT Mount Failure
Description	Nutanix Guest Tools mount may have failed.
Alert message	Nutanix Guest Tools mount failed for *vm_name*. *reason*.
Cause	CDROM slot is not free. This can be because of hypervisor-related error.
Impact	Nutanix Guest Tools will not be mounted on the VM.
Resolution	Check the logs on the CVM. If required, also check the VM logs for any error information to fix the issue.
Severity	Warning

Table 633: VSS Snapshot is not supported for the VM. [130105] [A130105]

Name	VSS Snapshot Not Supported
Description	VSS snapshot is not supported for the VM.
Alert message	VSS snapshot is not supported for the VM '*vm_name*', because *reason*.
Cause	VM has unsupported configuration.
Impact	Hypervisor based application consistent snapshot is taken on ESX. Crash consistent snapshot is taken for other hypervisors.
Resolution	Please look at the alert message and fix the invalid configuration.
Severity	Warning

Table 634: VSS Snapshot Aborted [130106] [A130106]

Name	VSS Snapshot Aborted
Description	VSS snapshot is aborted by the Guest VM.
Alert message	VSS snapshot is aborted for the VM '*vm_name* (ID *vm_id*)', because *reason*.
Cause	Guest VM aborted VSS snapshot operation.
Impact	Crash consistent snapshot is taken.
Resolution	Please look at the alert message for details of the failure and take necessary actions. If issue persists contact Nutanix Support.
Severity	Warning

Table 635: Recovered VM Disk Configuration Update Failed [130107] [A130107]

Name	Disk Configuration Update Failed
Description	Disk Configuration Update Failed
Alert message	Failed to make some disks online on the recovered VM '*vm_name*'. *reason*.
Cause	Nutanix Guest Tools failed to automatically bring the disks online.
Impact	Disks may become offline on the recovered VM.
Resolution	Manually bring the disks online on the recovered VM. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	8208
Severity	Critical

Table 636: Execution of the PostThaw Script Failed [130113] [A130113]

Name	PostThaw Script Execution Failed
Description	Execution of the PostThaw Script Failed

Alert message	Failed to execute the post_thaw script during the creation of the application consistent snapshot with uuid '*snapshot_uuid*' for the VM '*vm_name*'. Error : *error_message*.
Cause	Guest VM failed to execute the post_thaw script during the creation of the application consistent snapshot.
Impact	Some of the applications that got stopped by the pre_freeze script may not start after application consistent snapshot is created for the VM.
Resolution	Manually run the post_thaw script on the guest VM. Fix the cause of the script failure to avoid any further execution failures. If you cannot resolve the error, contact Nutanix support.
KB Article	8116
Severity	Critical

Table 637: Failed to take the application-consistent snapshot for the VM. [130127] [A130127]

Name	Application Consistent Snapshot Skipped
Description	Failed to take the application-consistent snapshot for the VM.
Alert message	Failed to take the application-consistent snapshot for the VM '*vm_name*', because *reason*.
Cause	VMware tools is not installed on the VM.
Cause	VMware tools is not running on the VM.
Cause	Removal of temporary hypervisor-based snapshot of the VM failed.
Impact	Crash consistent snapshot is taken for the VM.
Resolution	Install the VMware tools on the VM.
Resolution	Start the VMware tools on the VM.
Resolution	Manually remove temporary hypervisor-based snapshot of the VM.
Severity	Warning

Table 638: VM Registration Warning [130132] [A130132]

Name	VM Registration Warning
Description	VM Registration caused warning
Alert message	VM with name *vm_name* has been recovered successfully, however the following issue was observed - *reason*
Cause	Processor features may not be compatible.
Impact	VM registered successfully during the disaster recovery but with warnings.
Resolution	Ensure that the recovered VM is working as expected.
Severity	Warning

Table 639: Guest Power Operation through NGT Failed [130149] [A130149]

Name	VM Guest Power Op Failed
Description	Guest Power Operation Failed
Alert message	Failed to perform operation '*operation*'on VM '*vm_name*'. *reason*.
Cause	Power operation failed in guest operating system.
Impact	Desired power operation may not be completed.
Resolution	(A) Manually shutdown the VM by logging into the VM and running the appropriate shutdown command. (B) Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.
KB Article	10242
Severity	Warning

Table 640: Removal of the Temporary Hypervisor Snapshot Failed for the VM [130150] [A130150]

Name	Temporary Hypervisor Snapshot Cleanup Failed
Description	Removal of the temporary hypervisor snapshot that got created while taking the host-based application-consistent snapshot has failed for the VM.
Alert message	Removal of the temporary hypervisor snapshot that got created while taking the Nutanix host-based application-consistent snapshot has failed for the VM *vm_name*.
Cause	VM might be experiencing high I/O when host-based application-consistent snapshot cleanup is attempted.
Impact	Future application-consistent snapshots might fail.
Resolution	Wait until the merging of the snapshots is complete and then verify that there are no temporary Nutanix host-based snapshots before taking an application-consistent snapshot for the VM.
Severity	Warning

Table 641: Nutanix Guest Tools Failed to Initiate VM Reboot. [130190] [A130190]

Name	NGT VM Reboot Failure
Description	Nutanix Guest Tools failed to initiate reboot on the VM.
Alert message	Nutanix Guest Tools failed to initiate the VM reboot after NGT install/upgrade on the VM '*vm_name*' due to *reason*.
Cause	Failed to initiate reboot on the VM. Please check the alert message for exact cause of the failure.
Impact	VM is not rebooted after NGT install/upgrade. VM Mobility drivers may not be updated.
Resolution	Resolve the stated reason for the failure in the alert message. If the issue persists contact Nutanix Support.
Severity	Warning

Table 642: Conflicting NGT policies [130192] [A130192]

Name	NGT policy conflicts
Description	There are conflicting NGT policies for the VM.
Alert message	**Conflicting NGT policies found for the VM '*vm_uuid*'.**
Cause	More than one NGT policy is applicable for this VM. Please resolve the conflicts.
Impact	VM is not rebooted after NGT install/upgrade. Scheduled reboot of the VMs will be impacted.
Resolution	Check if more than 1 NGT policy exists for each VM and remove such policies via the Prism Central.
KB Article	10694
Severity	Warning

Table 643: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	**Failed to evacuate VMs off of node *node_uuid*: *error***
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.
KB Article	8471
Severity	Critical

Table 644: NGT CD-ROM not Unmounted on the VM [130334] [A130334]

Name	NGT CD-ROM Not Unmounted
Description	NGT CD-ROM not unmounted on the VM.
Alert message	**NGT CD-ROM not unmounted on the VM '*vm_name*' with uuid *vm_id***
Cause	NGT CD-ROM not unmounted on the VM because NGT version of the VM is incompatible with the NGT version of the cluster.
Impact	NGT workflow will not be impacted, however it is advised to unmount the stale ISO.
Resolution	Unmount the NGT ISO on the VM.
Severity	Info

Table 645: Node Removal is Stuck [130337] [A130337]

Name	Node Removal Stuck
Description	Cannot proceed with the node removal as the node is not in maintenance mode.
Alert message	Cannot proceed with the node removal due to: *error*
Cause	Insufficient cluster resources are available for the VMs to migrate off the host that is marked for removal.
Cause	Exit Maintenance Mode action triggered manually on the node that is marked for removal, and hence the node removal task cannot proceed.
Cause	The node that is being removed is either the source or the destination for VMs that are being live migrated to/from another cluster, and hence the node removal task cannot proceed.
Impact	The node marked for removal cannot be removed from the cluster until this condition is resolved.
Resolution	Ensure enough resources are available on other hosts for the VMs to evacuate from the host being removed. Consider manually shutting down unnecessary VMs and reviewing the HA settings. Put the affected node back in Maintenance Mode now.
Resolution	Put the affected node back in Maintenance Mode.
Resolution	Wait for cross cluster live migration of VM(s) to finish.
KB Article	8471
Severity	Warning

Table 646: VM is in Paused state [130348] [A130348]

Name	VM is in Paused state
Description	VM found in paused state.
Alert message	VM *vm_name* with ID *vm_uuid* has been found in paused state
Cause	Failed migration or some hypervisor issue.
Impact	VM in paused state most likely means downtime and it's in an irrecoverable state.
Resolution	Power cycle the VM to get it out of this state.
KB Article	10834
Severity	Warning

Table 647: File Analytics VM component failure [160081] [A160081]

Name	File Analytics VM component failure
Description	One or more components of the File Analytics VM {avm_ip} are not functioning properly or have failed.

Alert message	One or more components of the File Analytics VM *avm_ip* are not functioning properly or have failed.
Cause	One or more File Analytics VM components have failed.
Impact	File Server analytics will not perform appropriately.
Resolution	Contact Nutanix Support.
Severity	Critical

Table 648: File Analytics VM high CPU usage. [160082] [A160082]

Name	File Analytics VM high CPU usage.
Description	The File Analytics VM {avm_ip} has a high CPU usage of {usage_percent}%
Alert message	The File Analytics VM *avm_ip* has a high CPU usage of *usage_percent*%
Cause	The File Analytics VM is performing too many CPU intensive operations.
Impact	The File Analytics VM's performance will be impacted.
Resolution	Monitor which component has a high CPU usage and reduce the processing load.
Severity	Warning

Table 649: File Analytics VM high disk usage. [160083] [A160083]

Name	File Analytics VM high disk usage.
Description	The File Analytics VM {avm_ip} has a high disk usage of {usage_percent}%
Alert message	The File Analytics VM *avm_ip* has a high disk usage of *usage_percent*%
Cause	Extensive disk usage on the File Analytics VM
Impact	The File Analytics VM's performance will be impacted.
Resolution	Add more capacity to the File Analytics VM or delete some data.
Severity	Warning

Table 650: File Analytics VM low memory available. [160084] [A160084]

Name	File Analytics VM low memory available.
Description	The memory available on File Analytics VM {avm_ip} is {memory_available} bytes which is low.
Alert message	The percent memory available on File Analytics VM *avm_ip* is *memory_available* which is low
Cause	Excessive memory usage on the File Analytics VM
Impact	The File Analytics VM's performance will be impacted.
Resolution	Add more RAM
Severity	Warning

Table 651: vNUMA VM Pinning Failure. [200402] [A200402]

Name	vNUMA VM Pinning Failure
Description	Cluster was unable to provide requested vNUMA pinning for VM.
Alert message	VM '*vm_name*' with uuid *vm_id* is running with a non-compliant vNUMA pinning following *operation*.
Cause	Insufficient memory and/or CPU resources available.
Impact	VM performance may degrade.
Resolution	Migrate VM to a host which has adequate CPU and memory resources. Please refer to KB-10223.
KB Article	10223
Severity	Warning

Table 652: VM stored on RF1 container failed to start after host/CVM maintenance. [200405] [A200405]

Name	VM stored on RF1 container failed to start after host/CVM maintenance
Description	VM stored on RF1 container failed to start after host/CVM maintenance.
Alert message	VM *vm_name* (UUID: *vm_uuid*) failed to start with the following error: *message*
Cause	Insufficient resources available to start the VM.
Impact	VM which was running before the host/CVM maintenance will stay powered off until manual intervention.
Resolution	Free up resources and manually start the VM. Refer to KB-10804.
KB Article	10804
Severity	Critical

Table 653: NGT Update Available [600101] [A600101]

Name	NGT update available
Description	Checks if VM have latest NGT version supported by cluster.
Alert message	It is recommended that NGT on the VM *vm_name* with uuid *vm_id* should be upgraded to the latest version supported by the cluster.NGT update contains bug fixes and improvements, which will improve the overall product experience.
Cause	NGT version on VM is not latest
Impact	NGT update may contain bug fixes and improvements. This may not be available with the NGT version that is currently installed.
Resolution	Upgrade the NGT version on VM to the latest version supported by cluster.
KB Article	5487
Severity	Info

Table 654: NGT client certificates expiring [600103] [A600103]

Name	NGT client certs expiring
Description	NGT: Client certificates present on the guest VMs will expire soon
Alert message	Client certificates for VM *vm_name* with uuid *vm_id* are expiring in *expiry_time* days and need to be regenerated.Upon Certificate expiry, the CVM-Guest VM communication will be broken.
Cause	NGT Client certificates have definite expiry period of 1000 days based on ISO standards.
Impact	If a guest VM's NGT client certificate is expired, it will not be able to communicate with the CVM and will result in unavailability.
Resolution	NGT Client certificates need to be regenerated on the guest VMs. Refer to KB 10075 for further details
KB Article	10075
Severity	Critical
Severity	Warning
Severity	Info

Hardware

Table 655: SATA DOM has failed. [1025] [A1165]

Name	SATADOM Status
Description	Checks that host SATA DOM is functioning.
Alert message	SATA DOM on host *ip_address* has failed.
Cause	SATA DOM has lost power connection.
Cause	SATA DOM has failed.
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	Check the SATA DOM physical connection.
Resolution	Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	1850
Severity	Critical

Table 656: SATA DOM not reachable. [1030] [A1177]

Name	SATADOM Connection Status
Description	Checks that host SATA DOM is reachable.
Alert message	SATA DOM on host *ip_address* cannot be reached. The SATA DOM disconnection is determined by using sata dom information data: *data*

Cause	SATA DOM has lost power connection.
Cause	SATA DOM is not installed.
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	Check if SATA DOM is installed.
Resolution	Check the SATA DOM physical connection. Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	1850
Severity	Critical

Table 657: SATA DOM has worn out. [1031] [A1180]

Name	SATADOM Wearout Status
Description	Checks the wearout of SATA DOM via SMART data.
Alert message	SATA DOM with serial: *serial* on host *ip_address* has PE cycles above 4500 or PE cycles above 3000 and daily PE cycles above 15.
Cause	SATA DOM has been in use for a long time.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	3060
Severity	Critical

Table 658: NVMe Drive has errors. [1059] [A1196]

Name	NVMe Status Check
Description	Check that host NVMe drive is functioning properly.
Alert message	NVMe Drive on host *ip_address* has errors: *alert_msg*
Cause	NVMe Drive is damaged or worn out.
Impact	Cluster performance may be degraded. The integrity of data from NVMe drive may be harmed.
Resolution	Replace the problematic NVMe drive as soon as possible.
KB Article	3751
Severity	Warning

Table 659: NIC Flaps detected on the host. [3066] [A3066]

Name	NIC Flapping Check
Description	Check that all NICs have no flapping.
Alert message	NIC *nic_name* on host *host_ip* has flapped *nic_flap_count* times in one hour.
Cause	A physical networking component is failing.
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	Check all physical networking components (including switch, cables and NIC) and replace the failed component.
NIC Flap Threshold	Warning Threshold: 5 flaps
KB Article	2479
Severity	Warning

Table 660: NIC Error Rate Check [6011]

Name	NIC Error Rate Check
Description	Checks whether the NICs on hypervisor are experiencing errors.
Cause	NIC may be failing.
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	Check NIC health and if required, replace the NIC as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	1381

Table 661: NIC RX_CRC Error Rate High [6013] [A6013]

Name	NIC CRC Error Rate Check
Description	Checks whether the NICs on hypervisor are experiencing rx_crc errors.
Alert message	NIC *nic_name* in host *host_ip* has encountered many *error_type*.
Cause	NIC or its peripherals may be failing.
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	Replace the NIC, cable, SFP+ adapter, or switch port. Refer to the Nutanix documentation for instructions.
RX CRC Error Percentage Threshold	Critical Threshold: 1 %

KB Article	1381
Severity	Critical

Table 662: NIC RX_Missed Error Rate High [6014] [A6014]

Name	NIC Missed Error Rate Check
Description	Checks whether the NICs on hypervisor are experiencing rx_missed errors.
Alert message	NIC *nic_name* in host *host_ip* has encountered many *error_type*.
Cause	RX buffer size or transient condition.
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	RX buffer size might be too small to absorb current inbound traffic. This can be a transient condition. If the alert reappears frequently, contact Nutanix support for instructions on how to increase RX buffer.
RX Missed Error Percentage Threshold	Warning Threshold: 1 %
KB Article	1381
Severity	Warning

Table 663: Node NIC Link Speed Check [6015]

Name	Node NIC Link Speed Check
Description	Check if negotiated speed on any interface is less than 1 Gbps.
Cause	Speed negotiated between NIC and Switch is less than 1 Gbps.
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	Check the switch configuration and connection between NIC and switch.
KB Article	10171

Table 664: AHV Host NIC MTU Configuration [6016]

Name	AHV Host NIC MTU Configuration
Description	Check MTU of the AHV host network interfaces.
Cause	Network switch is dropping packets with configured MTU size.
Impact	If NIC MTU is not correctly configured on the AHV host or switch, VM live migration or other operations may fail.

Resolution	Review network switch configuration.
KB Article	11043

Table 665: Checks SSD latency threshold [6515] [A6515]

Name	SSD IO Latency
Description	Checks whether the system load average is high.
Alert message	SSD *ssd_list* for *svm_ip* are having await time greater than *threshold*
Cause	There may be too many VMs on one host.
Impact	High I/O latencies may be experienced by some workloads.
Resolution	Redistribute VMs to reduce load on the affected Controller VM.
SSD IO Latency	Warning Threshold: 50 seconds
KB Article	1582
Severity	Warning

Table 666: Fan Speed Low [15006] [A1045]

Name	Fan Speed Low Check
Description	Checks if the fan speed is not low.
Alert message	*fan_id* has stopped or its speed is low on Controller VM *ip_address*.
Cause	A fan has failed.
Impact	The device may fail or be permanently damaged.
Resolution	Replace the fan as soon as possible. Refer to the Nutanix documentation for instructions.
Lower threshold for fan speed in RPM	Critical Threshold: 500 rpm
KB Article	1524
Severity	Critical

Table 667: Fan Speed High [15007] [A1020]

Name	Fan Speed High Check
Description	Checks if the fan speed is not too high.
Alert message	Speed of *fan_id* exceeded threshold on Controller VM *ip_address*. Current sensor value is *fan_rpm*
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.

Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.
Upper threshold for fan speed in RPM	Critical Threshold: 25000 rpm
KB Article	1524
Severity	Critical

Table 668: Reduction in total memory size detected [15008] [A15008]

Name	DIMM Size Check
Description	Check if current available memory has gone below the installed size.
Alert message	*alert_msg*
Cause	A DIMM has failed, or DIMMs have been removed from the node.
Impact	Memory capacity on the node is less than expected.
Resolution	Determine if a failed DIMM needs to be replaced.
KB Article	9261
Severity	Critical

Table 669: CPU Temperature High [15010] [A1049]

Name	CPU Temperature
Description	Check that cpu temperature is not too high.
Alert message	Temperature of *cpu_id* exceeded threshold on Controller VM *ip_address*. Current sensor value is *temperature*C. *fan_speed*
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.
KB Article	1524
Severity	Critical

Table 670: RAM Temperature High [15011] [A1022]

Name	DIMM Temperature
Description	Check that DIMM temperature is not high.
Alert message	Temperature of *dimm_id* exceeded threshold on Controller VM *ip_address*. Current sensor value is *temperature*C

Cause	DIMM temperature is high.
Impact	The device may fail or be permanently damaged.
Resolution	Check RAM temperature.
KB Article	1524
Severity	Critical

Table 671: System Temperature High [15012] [A1012]

Name	System Temperature
Description	Checks if the system temperature is not high.
Alert message	System temperature exceeded threshold on Controller VM *ip_address*. Current sensor value *temperature*C
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.
KB Article	1524
Severity	Critical

Table 672: GPU Temperature is not within threshold [15014] [A1070]

Name	GPU Temperature
Description	Checks if GPU temperature is beyond threshold
Alert message	*alert_msg*
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.
KB Article	1524
Severity	Critical

Table 673: Hardware Clock Failure [15015] [A1059]

Name	Hardware Clock Status
Description	Check if the hardware clock has failed.
Alert message	Hardware clock in host *host_ip* has failed.
Cause	The RTC clock on the host has failed or the RTC battery has died.
Impact	The host may fail after restart and remain unavailable until it is replaced.

Resolution	Replace the node. Follow the instructions mentioned in KB 2462.
KB Article	2462
Severity	Critical

Table 674: IPMI Error [15016] [A1050]

Name	IPMI SDR Status
Description	Checks whether IPMI SDR has failures.
Alert message	Controller VM *ip_address* is unable to fetch IPMI SDR repository.
Cause	If the Hypervisor is ESXi or AHV, there is a possibility of process memory being exhausted on the host due to serviceability scripts running on the CVM.
Cause	The cluster did not correctly fail back after a Controller VM was temporarily down.
Cause	The IPMI interface is down.
Cause	There is a network connectivity issue.
Impact	Without IPMI access, we are not able to do Server management/ Troublshooting remotely, for example access DCUI or remote power action etc.
Resolution	Connect to the Hypervisor host of the CVM mentioned in the alert. Contact Nutanix Support.
Resolution	Ensure that the hypervisor password configured in the cluster matches the actual hypervisor password.
Resolution	Ensure that the IPMI interface is functioning.
Resolution	Ensure that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	1524
Severity	Critical

Table 675: Mismatch detected in count of GPUs [15017] [A1071]

Name	GPU Status
Description	Check that GPU have not faulted.
Alert message	*alert_msg*
Cause	The device is failing.
Impact	Graphics acceleration is not available to guest VMs.
Resolution	Replace the device. Refer to the Nutanix documentation for instructions.
KB Article	1524
Severity	Critical

Table 676: Power Supply Status Down/Unretrievable [15018] [A1046]

Name	Power Supply Status
Description	Check that Power supply has no errors.
Alert message	*per_psu_msg*
Cause	1. The power supply has failed or status not retrievable because BMC can not communicate with PSU properly2. BMC version may not be up to date and/or able to communicate with PSU for accurate status
Impact	The affected block does not have power supply redundancy.
Resolution	1. Restore the power supply as soon as possible.Refer to the Nutanix documentation for instructions.2. If PSU info unavailable, upgrade to the latest BMC firmware, and apply factory default after BMC upgrade. If it does not solve the issue, Refer to Nutanix KB 7386
KB Article	7386
Severity	Critical
Severity	Warning

Table 677: High number of correctable ECC errors found. [15019] [A1187]

Name	IPMI SEL CECC Check
Description	Check for number of CECC errors in the IPMI SEL.
Alert message	*dimm_errors per_dimm_msg*
Cause	A memory module in the node is failing. DIMM CECC errors / RAS event found. Installed BIOS version is P[X]42.300 or newer.
Cause	A memory module in the node is failing. DIMM CECC POST errors found. Installed BIOS version is P[X]42.300 or newer.
Cause	A memory module in the node is failing. Installed BIOS version is P[X]41.002 / 42.000
Cause	A memory module in the node is failing. Installed BIOS version below P[X]41.002 in G6/G7 series.
Cause	A memory module in the node is failing. All installed BIOS versions in G4/G5 series.
Impact	The node may stop running thereby reducing the cluster capacity.If BIOS version is P[X]42.300 or newer, and CECC POST error is found, the DIMM capacity will reduce to half
Resolution	Reboot the node as soon as possible to trigger DIMM repair. Refer to the Nutanix documentation (KB-11342) for further instructions.
Resolution	Contact Nutanix Support and replace the failing memory module. Refer to the Nutanix documentation (KB-11342) for further instructions.
Resolution	Ensure that you upgrade the BMC and BIOS to the latest version using LCM in 24 hours. New BIOS versions have improved DIMM error handling ability that reduces the need for aggressive proactive DIMM replacement. If not possible contact Nutanix Support to replace the failing DIMM. Refer to the Nutanix documentation (KB-11342) for further instructions.

Resolution	Contact Nutanix Support to replace the failing DIMM. Ensure that you upgrade the BMC and BIOS to the latest version using LCM. New BIOS versions have improved DIMM error handling ability that reduces the need for aggressive proactive DIMM replacement. Refer to the Nutanix documentation (KB-11342) for further instructions.
Resolution	Contact Nutanix Support to replace the failing DIMM. Refer to the Nutanix documentation (KB-11342) for further instructions.
KB Article	11342
Severity	Critical

Table 678: CPU Temperature Reading Error [15028] [A15028]

Name	CPU Temperature Fetch
Description	Check that CPU temperature can be fetched.
Alert message	Failed to fetch temperature of *cpu_id* on Controller VM *ip_address*
Cause	CPU temperature sensor may be unstable or BMC firmware may be reporting temperature inaccurately.
Impact	CPU temperature cannot be monitored.
Resolution	Check BMC SEL log. If CPU temperature in SEL log is 0, check CPU temperature sensor. If CPU temperature in SEL is not 0, ensure that the BMC firmware version is correct.
KB Article	1524
Severity	Critical

Table 679: CPU Temperature Low [15031] [A15031]

Name	CPU Temperature Low
Description	Check that cpu temperature is not too low.
Alert message	Temperature of *cpu_id* is below threshold on Controller VM *ip_address*. Current sensor value is *temperature*C
Cause	BMC or the sensor on CPU may have problem
Impact	The device may fail or be permanently damaged.
Resolution	Do BMC factory default. If the problem still persists, replace the CPU.
KB Article	1524
Severity	Critical

Table 680: RAM Temperature Low [15032] [A15032]

Name	DIMM Temperature Low
Description	Check that DIMM temperature is not low.

Alert message	Temperature of *dimm_id* is below threshold on Controller VM *ip_address*. Current sensor value is *temperature*C
Cause	BMC or the sensor on DIMM may have problem
Impact	The device may fail or be permanently damaged.
Resolution	Do BMC factory default. If the problem still persists, replace the chassis.
KB Article	1524
Severity	Critical

Table 681: System Temperature Low [15033] [A15033]

Name	System Temperature Low
Description	Check that system temperature is not low
Alert message	System temperature is below threshold on Controller VM *ip_address*. Current sensor value is *temperature*C
Cause	BMC or the sensor on chassis may have problem
Impact	The device may fail or be permanently damaged.
Resolution	Do BMC factory default. If the problem still persists, replace the chassis.
KB Article	1524
Severity	Critical

Table 682: System Temperature Reading Error [15034] [A15034]

Name	System Temperature Fetch
Description	Check that System temperature can be fetched.
Alert message	Failed to fetch system temperature on Controller VM *ip_address*
Cause	System temperature sensor may be unstable or BMC firmware may be reporting temperature inaccurately.
Impact	System temperature cannot be monitored.
Resolution	Check BMC SEL log. If System temperature in SEL log is 0, check System temperature sensor. If System temperature in SEL is not 0, ensure that the BMC firmware version is correct.
KB Article	1524
Severity	Critical

Table 683: High number of uncorrectable errors-UECC found. [15039] [A15039]

Name	IPMI SEL UECC Check
Description	Check for number of UECC errors for last one day in the IPMI SEL.
Alert message	DIMM UECC errors found for *dimm_id*(Serial:*serial*) by *vendor* on node *host_ip* during *event*

Cause	A memory module in the node has failed.
Impact	The node may stop running, reducing cluster capacity.
Resolution	Contact Nutanix Support to replace the failed DIMM. Ensure that you upgrade the BMC and BIOS to the latest version using LCM.New BIOS versions have improved DIMM error handling ability that reduces the need for proactive DIMM replacement.
KB Article	11342
Severity	Critical

Table 684: DIMM HPPR failure event found. [15040] [A15040]

Name	DIMM HPPR Check
Description	Check for DIMM HPPR failure event in the IPMI SEL.
Alert message	DIMM HPPR failure event found for *dimm_id*(Serial:*serial*) by *vendor* on node *host_ip*
Cause	A memory module in the node is failing.
Impact	The node may stop running, reducing cluster capacity.
Resolution	Contact Nutanix Support to replace the failed memory module as soon as possible.
KB Article	11342
Severity	Critical

Table 685: Hypervisor boot drive RAID is in an unhealthy state. [101058] [A101058]

Name	Hypervisor boot drive RAID status
Description	Checks the status of Dell Marvell BOSS hypervisor boot drive.
Alert message	Hypervisor boot drive RAID on *ip_address* is in *raid_state* state.
Cause	Hypervisor boot drive RAID is not functioning properly.
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 Support to help restore the hypervisor boot drive RAID.
KB Article	4816
Severity	Critical

Table 686: Free block count of SATA DOM has gone below threshold. [101062] [A101062]

Name	SATADOM Free Block Count Check
Description	Checks the free block count of SATADOM-SL 3IE3.
Alert message	*alert_msg*

Cause	SATADOM-SL 3IE3 free block count is less than the threshold
Impact	As FBC approaches zero the host boot disk performance may be impacted.
Resolution	If no performance issues, then nothing needs to be done. If there are performance issues, either replace the SATADOM or free up some space on SATADOM.Please make sure SATADOM is on the latest firmware and if issue still persists then contact Nutanix support for further assistance.
KB Article	4395
Severity	Info

Table 687: Bond uplink VLAN config check [103106]

Name	Bond uplink VLAN config check
Description	Checks if all uplinks in the bonds are configured with same set of VLANs.
Cause	Inconsistent VLAN configuration detected on the uplinks.
Impact	Hosts, CVM, or user VMs can lose network connectivity if active uplink changes. In balance-slb and balance-tcp bond modes, optimal performance might not be achieved.
Resolution	Check for VLAN configuration on switch ports where uplinks are connected.
KB Article	8185

Table 688: Bond uplink connectivity check [103107]

Name	Bond uplink connectivity check
Description	Checks if at least two uplinks in the bond are connected.
Cause	No cable is connected to the host uplink port. Cable, host NIC or switchport may be faulty. Switchport is in DOWN state.
Impact	If less than two uplinks in the bond are connected to network then the redundancy of network connectivity will be lost
Resolution	If the bond has two or more uplinks added, make sure that at least two of these uplinks are connected. Also check if cables and network hardware are healthy and switchport status is UP.
KB Article	8185

Table 689: SAS Connectivity Not Normal [106019] [A106019]

Name	SAS Connectivity
Description	Check SAS connectivity
Alert message	SAS cables missing or not properly connected on host.
Cause	SAS cables are not properly connected.
Impact	Cluster performance may be significantly degraded.

Resolution	Ensure that SAS cables are properly connected.
KB Article	2311
Severity	Warning

Table 690: CPUs of different types or models are in the same chassis. [106026] [A1190]

Name	CPU type on chassis check
Description	Checks whether CPUs within a chassis are of the same type.
Alert message	Chassis *chassis_id* has CPU configuration problem: *alert_msg*. This problem may affect the performance and stability of the nodes in the chassis. Chassis cpu info: *chassis_cpu_info*
Cause	A node of a different type was added to the chassis.
Impact	Having nodes of different CPU types in the same chassis is not supported.
Resolution	Remove the node that has a different CPU type from the chassis. Before moving a node to a chassis, make sure its CPU type is the same as that of other nodes in the chassis.
KB Article	3288
Severity	Critical

Table 691: SATADOM-ML 3SE has high wear. [106037] [A106037]

Name	SATADOM-ML 3SE Wearout Status
Description	Checks the wearout of SATADOM-ML 3SE via SMART data.
Alert message	SATADOM-ML 3SE with serial: *serial* on host *ip_address* has the following problem: *alert_msg*
Cause	SATADOM-ML 3SE Remaining Device Life too short.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATADOM-ML 3SE as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	3060
Severity	Warning

Table 692: Host boot disk requires attention [106042] [A106042]

Name	M.2 Boot Disk change check
Description	Check the change of host boot disk
Alert message	Host boot disk on host *ip_address* has the following problems: *err_msg*
Cause	Active host boot disk failed and needs to switch to spare host boot disk to boot the host.

Impact	Host boot disk may lose redundancy.
Resolution	Replace the failed host boot disk. Refer to the Nutanix documentation for instructions.
KB Article	4874
Severity	Warning

Table 693: Hypervisor boot drive RAID status check. [106066] [A106066]

Name	ILO boot drive RAID Status Check
Description	Checks the status of HPE Integrated Lights-Out(iLO) boot drive.
Alert message	*alert_msg*
Cause	Hypervisor RAID Boot Device is not reachable or is faulty
Impact	Cluster performance may be significantly degraded. In case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check the status of the RAID device. Please refer KB 9064 to resolve this issue. If problem persists, please engage Nutanix Support.
KB Article	9064
Severity	Warning

Table 694: Host Boot Disk Guest VM Check [106067]

Name	Host Boot Disk Guest VM Check
Description	Checks that no guest VM is installed on Host Boot Disk.
Cause	A guest VM is installed on Host Boot Disk or it is incorrectly configured.
Impact	Degradation of the Host Boot Disk will be accelerated, leading to unavailability of the node.
Resolution	Remove guest VMs from host Boot Disk or reconfigure guest VMs on the host machine.
KB Article	9963

Table 695: Disk power modification required [106070] [A106070]

Name	Disk Power Check
Description	Check for PM1643 power on Lenovo node
Alert message	*alert_msg*
Cause	Power difference due to restrictions in prior generation systems.
Impact	Performance will be degraded
Resolution	Follow instructions given in the KB-10162

KB Article	10162
Severity	Warning

Table 696: Disk Cabling is incorrect. [106082] [A106082]

Name	HPE HBA Cabling Check
Description	Checks the cabling of all the disks are in correct order as per layout
Alert message	*alert_msg*
Cause	Cabling of one or more disks is incorrect
Impact	All disks might not be utilized to full efficiency
Resolution	Need to connect all disks as per correct cabling configuration
KB Article	11310
Severity	Warning

Table 697: One or more HBAs are connected at incorrect PCIe slot [106083] [A106083]

Name	HPE HBA NUMA Check
Description	Checks if all HBAs are connected to the right PCIe slot.
Alert message	*alert_msg*
Cause	One or more HBAs are connected at incorrect PCIe slot.
Impact	Disks might not be utilized to full efficiency.
Resolution	Please connect all HBAs to the correct PCIe slot.
KB Article	11310
Severity	Warning

Table 698: M10 GPU Configuration incorrect on the Node [111040] [A111040]

Name	M10 GPU check
Description	Check number of M10 GPUs on a node
Alert message	M10 GPU configuration has problem on node *ip_address*: *err_msg*
Cause	The number of M10 GPU cards on a node is not 1 or 2 or 3 or 4.
Impact	If the number of M10 GPU cards is not 1 or 2 or 3, the node cannot reboot. If the number of M10 GPU on a card is not 4, the VMs that the missing M10 GPU is assigned to may not run.
Resolution	Add or remove M10 GPU cards so that either 1 or 2 or 3 are installed on the node. Replace M10 GPU card that does not have 4 M10 GPUs.
KB Article	3682
KB Article	3682

Severity	Warning

Table 699: Node Failure [130030] [A1137]

Name	Node Status
Description	Node failure.
Alert message	Host *hypervisor_address* appears to have failed. High Availability is restarting VMs on *failover_host_info*.
Cause	Host is not accessible.
Impact	VMs will restart on another host.
Resolution	Check for network connectivity or a hardware failure.
Severity	Critical
Severity	Warning

Table 700: Non Self Encryption Drive Disk Inserted [130031] [A1122]

Name	Non SED Disk Inserted Check
Description	Non-self-encrypting drive is inserted.
Alert message	Non encrypted disk with serial *disk_serial* was added in drive bay *disk_location* on Controller VM *service_vm_external_ip*. It is not usable because the rest of the cluster is protected using encrypted drives.
Cause	A cluster with self-encrypting drives has a non-self encrypting drive installed.
Impact	The new drive was not mounted because the cluster cannot be fully secured without all drives supporting encryption.
Resolution	Ensure all installed drives are of self-encrypting type. If more help is needed call Nutanix support.
KB Article	8417
Severity	Critical

Table 701: Disk Failed/Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Drive is marked offline.
Alert message	Disk *disk_position* with id *disk_id* on node *node_position* of block *block_position* has been marked offline due to IO errors. Serial number of the disk is *disk_serial* in host *host_ip* of block *block_serial*.
Cause	A drive on the node has failed and was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Replace the failed drive. Refer to the Nutanix documentation for instructions.

KB Article	8453
Severity	Critical

Table 702: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	Drive *disk_id* with serial *disk_serial* and model *disk_model* in drive bay *disk_location* on Controller VM *service_vm_external_ip* has failed. Node serial is *node_serial_number* and Node position *node_position*. #rf1_custom_message@
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 703: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical drive is removed from the slot.
Alert message	Drive with serial *disk_serial* and model *disk_model* was removed from drive bay *disk_location* on Controller VM *service_vm_external_ip*. Node serial is *node_serial_number* and Node position *node_position*. #rf1_custom_message@
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the drive will start.
Resolution	Remove the drive from the slot and reseat it. Contact the Hardware Vendor if issue persists.
KB Article	8031
Severity	Critical

Table 704: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	Drive *disk_id* with serial *disk_serial* in drive bay *disk_location* on Controller VM *service_vm_external_ip* has failed diagnostic test. This drive is failing or has failed. #rf1_custom_message@
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.

Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 705: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	Failed to evacuate VMs off of node *node_uuid*: *error*
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.
KB Article	8471
Severity	Critical

License

Table 706: Nutanix Objects License Expiry [14008] [A14008]

Name	ObjectLicenseExpiry
Description	Check if Nutanix Objects license is about to expire.
Alert message	This *cluster_type* has expiring or expired Nutanix Objects license(s). Detailed license expiry info: *license_expiry_msg* *license_warning_msg*
Cause	The Nutanix Objects license is set to expire or has expired.
Impact	The Nutanix Objects instance will be out of licensing compliance.
Resolution	Contact Nutanix to extend the expiration of the Object license, or apply a new Object license.
KB Article	2469
Severity	Warning

Table 707: Files License Invalid [160086] [A160086]

Name	File Server License Usage Check
Description	Check for Files license capacity violation
Alert message	Cluster does not have enough Files license capacity. *fs_license_msg*
Cause	The Files license capacity for the cluster is non-compliant.
Impact	Cluster does not have enough Files license capacity.
Resolution	Apply for a new license. Increase Files license capacity. Refer to KB article 7865. Contact Nutanix support if issue still persists or assistance needed.

KB Article	7865
Severity	Warning

Node

Table 708: SATA DOM has failed. [1025] [A1165]

Name	SATADOM Status
Description	Checks that host SATA DOM is functioning.
Alert message	SATA DOM on host *ip_address* has failed.
Cause	SATA DOM has lost power connection.
Cause	SATA DOM has failed.
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	Check the SATA DOM physical connection.
Resolution	Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	1850
Severity	Critical

Table 709: SATA DOM not reachable. [1030] [A1177]

Name	SATADOM Connection Status
Description	Checks that host SATA DOM is reachable.
Alert message	SATA DOM on host *ip_address* cannot be reached. The SATA DOM disconnection is determined by using sata dom information data: *data*
Cause	SATA DOM has lost power connection.
Cause	SATA DOM is not installed.
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	Check if SATA DOM is installed.
Resolution	Check the SATA DOM physical connection. Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	1850
Severity	Critical

Table 710: SATA DOM has worn out. [1031] [A1180]

Name	SATADOM Wearout Status
Description	Checks the wearout of SATA DOM via SMART data.
Alert message	SATA DOM with serial: *serial* on host *ip_address* has PE cycles above 4500 or PE cycles above 3000 and daily PE cycles above 15.
Cause	SATA DOM has been in use for a long time.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	3060
Severity	Critical

Table 711: NVMe Drive has errors. [1059] [A1196]

Name	NVMe Status Check
Description	Check that host NVMe drive is functioning properly.
Alert message	NVMe Drive on host *ip_address* has errors: *alert_msg*
Cause	NVMe Drive is damaged or worn out.
Impact	Cluster performance may be degraded. The integrity of data from NVMe drive may be harmed.
Resolution	Replace the problematic NVMe drive as soon as possible.
KB Article	3751
Severity	Warning

Table 712: Baseline configuration changes [3039] [A3039]

Name	AIDE baseline configuration changes check
Description	Checks the log files of aide for changes to baseline configs if aide is enabled
Alert message	*message*
Cause	Baseline configurations changed
Impact	Aide execution may fail depending on the files that were changed
Resolution	Check changed files and revert their changes if needed
KB Article	10461
Severity	Warning

Table 713: Latency between CVMs is high [6005] [A6005]

Name	Inter-CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	Latency between *vm_type*s is higher than *ping_threshold* ms.
Cause	Malfunctioning or misconfigured network infrastructure
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.
Ping Time Threshold	Critical Threshold: 15 ms
KB Article	3690
Severity	Critical

Table 714: CVM memory reservation is incorrectly configured. [6211] [A6211]

Name	CVM Memory Pinned Check
Description	Verify CVM memory reservation and pinning.
Alert message	CVM memory reservation is incorrectly configured on CVM with IP *service_vm_external_ip*
Cause	CVM memory reservation is configured incorrectly.
Impact	This can potentially result in cluster instability.
Resolution	Please contact Nutanix Support to verify memory resolution.
KB Article	1735
Severity	Warning

Table 715: Host CPU Utilization [6409]

Name	Host CPU Utilization
Description	Checks whether cpu contention is high on host.
Cause	High CPU utilization on host.
Impact	I/O performance may be degraded.
Resolution	Reduce CPU intensive VMs or reduce the total number of VMs.
KB Article	2797

Table 716: Storage tier skew on a node [20028]

Name	Storage tier skew on a node
Description	Checks for skew in storage tiers of the node
Cause	Disks in the tier are not the same size.

Impact	Unsupported configuration, impacting cluster performance
Resolution	Replace disk with the same size as the other disks in the tier on the node.
KB Article	10718

Table 717: Metadata disk(s) not mounted on CVM [101055] [A101055]

Name	Metadata Disk Mounted Check
Description	Check that all metadata disks are mounted.
Alert message	Metadata disk(s) on Controller VM *svm_ip_address* are not mounted: *unmounted_disks*
Cause	Metadata disk(s) not mounted on CVM.
Impact	Cluster performance may be significantly degraded. If this condition continues, there is a potential for data corruption and/or loss.
Resolution	Contact Nutanix Support to mount the metadata disk.
KB Article	4494
Severity	Critical
Severity	Warning

Table 718: Hypervisor boot drive RAID is in an unhealthy state. [101058] [A101058]

Name	Hypervisor boot drive RAID status
Description	Checks the status of Dell Marvell BOSS hypervisor boot drive.
Alert message	Hypervisor boot drive RAID on *ip_address* is in *raid_state* state.
Cause	Hypervisor boot drive RAID is not functioning properly.
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 Support to help restore the hypervisor boot drive RAID.
KB Article	4816
Severity	Critical

Table 719: Free block count of SATA DOM has gone below threshold. [101062] [A101062]

Name	SATADOM Free Block Count Check
Description	Checks the free block count of SATADOM-SL 3IE3.
Alert message	*alert_msg*
Cause	SATADOM-SL 3IE3 free block count is less than the threshold
Impact	As FBC approaches zero the host boot disk performance may be impacted.

Resolution	If no performance issues, then nothing needs to be done. If there are performance issues, either replace the SATADOM or free up some space on SATADOM.Please make sure SATADOM is on the latest firmware and if issue still persists then contact Nutanix support for further assistance.
KB Article	4395
Severity	Info

Table 720: CVM dvPortGroup Binding Check [103098]

Name	CVM dvPortGroup Binding Check
Description	Check for CVM dVportgroup binding
Cause	One (or more) Nutanix cluster Controller VM NIC is connected to a dVportgroup with ephemeral binding.
Impact	CVM may lose network connectivity on reboot
Resolution	Change the Dvportgroup binding to static
KB Article	3130

Table 721: Unsupported SFP is installed on host machine [103099] [A103099]

Name	Unsupported SFP message Check
Description	Check if an unsupported SFP is plugged in
Alert message	Unsupported SFP is installed on host machine *host_ip*: *alert_msg*
Cause	An unsupported SFP is plugged in.
Impact	NICs may not be recognized and network communication is affected.
Resolution	Remove the unsupported SFP
KB Article	6212
Severity	Warning

Table 722: Static Port Channel configuration found on the cluster [103108] [A103108]

Name	Static Port Channel Check
Description	Checks that static port channel is configured on host
Alert message	Detected utilization of static port-channel configuration by CVM *svm_ip* management interfaces
Cause	CVM management interfaces are utilising Static Port Channel configuration.
Impact	If upstream switch is using load sharing of traffic on a static port channel, phoenix based LCM firmware upgrades may fail.
Resolution	Remove Static Port Channel configuration.
KB Article	11072
Severity	Warning

Table 723: CPUs of different types or models are in the same chassis. [106026] [A1190]

Name	CPU type on chassis check
Description	Checks whether CPUs within a chassis are of the same type.
Alert message	Chassis *chassis_id* has CPU configuration problem: *alert_msg*. This problem may affect the performance and stability of the nodes in the chassis. Chassis cpu info: *chassis_cpu_info*
Cause	A node of a different type was added to the chassis.
Impact	Having nodes of different CPU types in the same chassis is not supported.
Resolution	Remove the node that has a different CPU type from the chassis. Before moving a node to a chassis, make sure its CPU type is the same as that of other nodes in the chassis.
KB Article	3288
Severity	Critical

Table 724: SATADOM-ML 3SE has high wear. [106037] [A106037]

Name	SATADOM-ML 3SE Wearout Status
Description	Checks the wearout of SATADOM-ML 3SE via SMART data.
Alert message	SATADOM-ML 3SE with serial: *serial* on host *ip_address* has the following problem: *alert_msg*
Cause	SATADOM-ML 3SE Remaining Device Life too short.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATADOM-ML 3SE as soon as possible. Refer to the Nutanix documentation for instructions.
KB Article	3060
Severity	Warning

Table 725: Host boot disk requires attention [106042] [A106042]

Name	M.2 Boot Disk change check
Description	Check the change of host boot disk
Alert message	Host boot disk on host *ip_address* has the following problems: *err_msg*
Cause	Active host boot disk failed and needs to switch to spare host boot disk to boot the host.
Impact	Host boot disk may lose redundancy.
Resolution	Replace the failed host boot disk. Refer to the Nutanix documentation for instructions.
KB Article	4874
Severity	Warning

Table 726: Boot drive is in degraded state [106053] [A106053]

Name	Host Hardware RAID status Check
Description	Check status of hardware raid and alert user if raid is degraded
Alert message	*alert_msg*
Cause	Boot drive mentioned in the alert message may be offline and requires attention
Impact	Node will operate without host boot drive redundancy but will be at risk of host unavailability if further hardware raid issues are seen
Resolution	Boot drive mentioned in alert message requires attention and triage the root cause of it being unavailable
KB Article	6627
Severity	Critical

Table 727: Temperature of M2 disk is out of range [106054] [A106054]

Name	Host Hardware RAID M2 Temperature Check
Description	Checks temperature of each drive and alert if it is not within 0C and 75C
Alert message	*alert_msg*
Cause	Ambient temperature in the affected host environment may be too high for effective cooling, or airflow through host may be suboptimal
Impact	Continued operation above temperature thresholds may damage the host boot drives and lead to unplanned downtime
Resolution	Review the physical environment of the affected host and ensure air-conditioning systems are functioning and correctly calibrated and confirm the absence of any airflow obstructions around vents and the host itself
KB Article	6627
Severity	Warning

Table 728: M2 disk returned UECC errors [106055] [A106055]

Name	Host Hardware RAID M2 UECC Check
Description	Checks if disk has returned UECC errors
Alert message	*alert_msg*
Cause	Host boot disk has UECC errors
Impact	Continued operation with UECC errors, may lead to unplanned downtime
Resolution	Review the UECC errors of each host boot disk and triage accordingly
KB Article	6627
Severity	Warning

Table 729: M2 disk has worn out [106056] [A106056]

Name	Host Hardware RAID M2 wearout Check
Description	Checks wearout level of the disk.
Alert message	*alert_msg*
Cause	Disk is too old, or disk has gone through too many write cycles.
Impact	Slow Performance.
Resolution	Disk needs to be replaced.
KB Article	6627
Severity	Warning

Table 730: Firmware of raid M2 disk needs to be upgraded. [106057] [A106057]

Name	Host Hardware RAID M2 Firmware Check
Description	Checks if hardware raid M2 disk firmware needs to be upgraded
Alert message	*alert_msg*
Cause	Hardware raid drive is running outdated firmware and requires update
Impact	Host boot drive may not have up-to-date performance and stability features of more recent firmware
Resolution	Review KB6627 for instructions on upgrading the host boot drive firmware
KB Article	6627
Severity	Warning

Table 731: RAID card BIOS or Firmware or boot loader needs to be updated. [106058] [A106058]

Name	Host Hardware RAID Card Firmware Check
Description	Checks if RAID card BIOS, Firmware or boot loader needs to be updated.
Alert message	*alert_msg*
Cause	Hardware raid controller is running outdated BIOS or Firmware or Bootloader version. Please refer to the alert message
Impact	Hardware RAID controller may not have up-to-date performance and stability features of more recent firmware.
Resolution	Review KB6627 for instructions on upgrading BIOS or Firmware or Bootloader of hardware RAID controller
KB Article	6627
Severity	Warning

Table 732: Temperature of RAID card is above threshold [106060] [A106060]

Name	Host Hardware Raid Board Temperature Check
Description	Checks temperature of RAID card and alert if it is more than 60C
Alert message	*alert_msg*
Cause	Ambient temperature in the affected host environment may be too high for effective cooling, or airflow through host may be suboptimal
Impact	Continued operation above temperature thresholds may damage the host boot drives and lead to unplanned downtime
Resolution	Review the physical environment of the affected host and ensure air-conditioning systems are functioning and correctly calibrated and confirm the absence of any airflow obstructions around vents and the host itself
KB Article	6627
Severity	Warning

Table 733: SMART parameters of disk are out of range [106061] [A106061]

Name	SATA Disk Status Check
Description	Disk status check for SATA drives
Alert message	*alert_msg*
Cause	One or more S.M.A.R.T. parameters of the disk may be out of range
Impact	Cluster performance may be significantly degraded or node may become unavailable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.Available storage space is reduced If the drive gets marked offline.
Resolution	Based on which S.M.A.R.T. parameters is out of range, refer to the Nutanix documentation for instructions.
KB Article	8094
Severity	Critical
Severity	Warning
Severity	Info

Table 734: SMART parameters of disk are out of range [106062] [A106062]

Name	SAS Disk Status Check
Description	Disk status check for SAS drives
Alert message	*alert_msg*
Cause	One or more S.M.A.R.T. parameters of the disk may be out of range
Impact	Cluster performance may be significantly degraded or node may become unavailable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.Available storage space is reduced If the drive gets marked offline.

Resolution	Based on which S.M.A.R.T. parameters is out of range, refer to the Nutanix documentation for instructions.
KB Article	8094
Severity	Critical
Severity	Warning
Severity	Info

Table 735: SMART parameters of disk(host) are out of range [106063] [A106063]

Name	Host Boot Disk Status Check
Description	Disk status check for host boot disks
Alert message	*per_disk_msg*
Cause	One or more S.M.A.R.T. parameters of the disk may be out of range
Impact	Cluster performance may be significantly degraded or node may become unavailable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.Available storage space is reduced If the drive gets marked offline.
Resolution	Based on which S.M.A.R.T. parameters is out of range, refer to the Nutanix documentation for instructions.
KB Article	8094
Severity	Critical
Severity	Warning
Severity	Info

Table 736: Disk firmware needs upgrade [106064] [A106064]

Name	Host Boot Disk Firmware Check
Description	Checks if any of the host boot disks, requires firmware upgrade or not
Alert message	*per_disk_msg*
Cause	Current host boot disk firmware is not recommended by Nutanix
Impact	Host boot Disk may not function properly, resulting in node unavailability.
Resolution	Upgrade host boot disk firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.
KB Article	7069
Severity	Warning

Table 737: Hypervisor boot drive RAID status check. [106066] [A106066]

Name	ILO boot drive RAID Status Check
Description	Checks the status of HPE Integrated Lights-Out(iLO) boot drive.
Alert message	*alert_msg*
Cause	Hypervisor RAID Boot Device is not reachable or is faulty
Impact	Cluster performance may be significantly degraded. In case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check the status of the RAID device. Please refer KB 9064 to resolve this issue. If problem persists, please engage Nutanix Support.
KB Article	9064
Severity	Warning

Table 738: Host Boot Disk Guest VM Check [106067]

Name	Host Boot Disk Guest VM Check
Description	Checks that no guest VM is installed on Host Boot Disk.
Cause	A guest VM is installed on Host Boot Disk or it is incorrectly configured.
Impact	Degradation of the Host Boot Disk will be accelerated, leading to unavailability of the node.
Resolution	Remove guest VMs from host Boot Disk or reconfigure guest VMs on the host machine.
KB Article	9963

Table 739: Hardware difference is noticed [106068] [A106068]

Name	Hardware differences alert
Description	Any change in any parameter of the hardware will raise this alert
Alert message	*alert_msg*
Cause	Hardware configuration is changed
Impact	NA
Resolution	NA
Severity	Info

Table 740: VMD Driver status [106069] [A106069]

Name	VMD Driver Disablement Check
Description	Check if VMD driver is present on ESXi hosts.
Alert message	VMD driver is installed on following hosts: *ip_address*

Cause	VMD enabled node on ESXi is susceptible to PSOD on reboot/shutdown.
Impact	VMD enabled node on ESXi is susceptible to PSOD on reboot/shutdown
Resolution	Uninstall VMD vib from nodes where it is installed.
KB Article	10053
Severity	Warning

Table 741: Host Boot RAID volume is in degraded state. [106071] [A106071]

Name	Host NVMe RAID status Check
Description	Check the status of NVMe RAID if degraded
Alert message	*alert_msg*
Cause	One of the M.2s connected to the RAID card is offline or has issues. Check alert message for the drive and port details
Impact	The node will be up and running with one M.2. But has a risk of going down in absence of redundancy.
Resolution	Replace the failing M.2. Check for the port and drive details in the alert message. Or Contact Nutanix Support.
KB Article	10355
Severity	Critical

Table 742: Temperature of M2 disk is out of range [106072] [A106072]

Name	Host NVMe RAID M2 Temperature Check
Description	Checks temperature of each drive and alert if it is not within 0C and 75C
Alert message	*alert_msg*
Cause	Ambient temperature in the affected host environment may be too high for effective cooling, or airflow through host may be suboptimal
Impact	Continued operation above temperature thresholds may damage the host boot drives and lead to unplanned downtime
Resolution	Check for any airflow obstructions physically near the node and the vents. Make sure air conditioning systems are functional.
KB Article	10355
Severity	Warning

Table 743: M2 disk returned UECC errors [106073] [A106073]

Name	Host NVMe RAID M2 UECC Check
Description	Checks if disk has returned UECC errors
Alert message	*alert_msg*

Cause	Both M.2s connected to the RAID card has UECC errors.
Impact	Continued operation with UECC errors, may lead to unplanned downtime
Resolution	It is recommended to replace one of the drives having UECC.
KB Article	10355
Severity	Warning

Table 744: M2 disk has worn out [106074] [A106074]

Name	Host NVMe RAID M2 wearout Check
Description	Checks wear-out level of the M.2.
Alert message	*alert_msg*
Cause	M.2 has worn out and reaching its specification limit.
Impact	Slow Performance.
Resolution	Disk needs to be replaced.
KB Article	10355
Severity	Warning

Table 745: Firmware of RAID M2 disk needs to be upgraded. [106075] [A106075]

Name	Host NVMe RAID M2 Firmware Check
Description	Checks if M2 disk firmware needs to be upgraded
Alert message	*alert_msg*
Cause	M.2 drive is running outdated firmware and requires update
Impact	M.2 drive may not have up-to-date performance and stability features of more recent firmware
Resolution	Review KB10355 for instructions on upgrading the host boot drive firmware
KB Article	10355
Severity	Warning

Table 746: RAID card BIOS or Firmware or boot loader needs to be updated. [106076] [A106076]

Name	Host NVMe RAID Card Firmware Check
Description	Checks if RAID card BIOS, Firmware needs to be updated.
Alert message	*alert_msg*
Cause	Hardware RAID controller is running outdated BIOS or Firmware version. Please refer to the alert message

Impact	Hardware RAID controller may not have up-to-date performance and stability features of more recent firmware.
Resolution	Review KB10355 for instructions on upgrading BIOS or Firmware or Bootloader of hardware RAID controller
KB Article	10355
Severity	Warning

Table 747: Temperature of RAID card is above threshold [106077] [A106077]

Name	Host NVMe RAID Board Temperature Check
Description	Checks temperature of RAID card and alert if it is more than 95C
Alert message	*alert_msg*
Cause	Ambient temperature in the affected host environment may be too high for effective cooling, or airflow through host may be suboptimal
Impact	Continued operation above temperature thresholds may damage the host boot drives and lead to unplanned downtime
Resolution	Review the physical environment of the affected host and ensure air-conditioning systems are functioning and correctly calibrated and confirm the absence of any airflow obstructions around vents and the host itself
KB Article	10355
Severity	Warning

Table 748: Host Hardware RAID AOC Link Speed [106078]

Name	Host Hardware RAID AOC Link Speed
Description	NVMe RAID controller PCIe speed is changed
Cause	RAID card link speed dropped.
Impact	Reduce link speed may impact performance.
Resolution	Probable RAID card link issue. Reboot the node. If issue persists, replace RAID card.
KB Article	10355

Table 749: Host Hardware RAID AOC Link Width [106079]

Name	Host Hardware RAID AOC Link Width
Description	Checks if NVMe RAID controller link width is changed
Cause	RAID card link width dropped.
Impact	Reduce link width may impact performance.
Resolution	Probable RAID card link issue. Reboot the node. If issue persists, replace RAID card

KB Article	10355

Table 750: Host NVMe RAID M2 Link Speed [106080]

Name	Host NVMe RAID M2 Link Speed
Description	Checks if disk link speed is dropped
Cause	Host NVMe M2 disk link speed is less than 8 Gbps
Impact	Reduce link speed may impact performance
Resolution	Probable RAID card M2 link issue or M2 issue connected to the same port. Reboot the node. If issue persists, replace M2. If issue still persists after replacing M2, replace RAID card
KB Article	10355

Table 751: Host NVMe RAID M2 Link Width [106081]

Name	Host NVMe RAID M2 Link Width
Description	Checks if disk link width is changed
Cause	Host NVMe M2 disk link width is less than 4x
Impact	Reduce link width may impact performance.
Resolution	Probably RAID card M2 port link issue or M2 issue connected to the same port. Reboot the node. If issue persists, replace M2. If issue still persists after replacing M2, replace RAID card.
KB Article	10355

Table 752: Disk Cabling is incorrect. [106082] [A106082]

Name	HPE HBA Cabling Check
Description	Checks the cabling of all the disks are in correct order as per layout
Alert message	*alert_msg*
Cause	Cabling of one or more disks is incorrect
Impact	All disks might not be utilized to full efficiency
Resolution	Need to connect all disks as per correct cabling configuration
KB Article	11310
Severity	Warning

Table 753: One or more HBAs are connected at incorrect PCIe slot [106083] [A106083]

Name	HPE HBA NUMA Check

Description	Checks if all HBAs are connected to the right PCIe slot.
Alert message	*alert_msg*
Cause	One or more HBAs are connected at incorrect PCIe slot.
Impact	Disks might not be utilized to full efficiency.
Resolution	Please connect all HBAs to the correct PCIe slot.
KB Article	11310
Severity	Warning

Table 754: HyperThreading Enabled Check [106443]

Name	HyperThreading Enabled Check
Description	Checks if Hyperthreading is enabled.
Cause	Hyperthreading is not enabled on host(s).
Impact	Performance.
Resolution	Enable Hyperthreading on host.
KB Article	11983

Table 755: BIOS-BMC upgrade required [111039] [A111039]

Name	BMC BIOS version check
Description	Check BMC/Bios version of all nodes
Alert message	*alert_msg*
Cause	Current BIOS/BMC version is not recommended for this node due to known issue
Impact	Current BIOS/BMC version may impact node behavior, skew sensor data, or lack enhanced features and security fixes.
Resolution	Upgrade BMC and BIOS version. Perform inventory from Life Cycle Manager to check and update to the latest BMC/BIOS version. Check KB 2896 and 2905 for alternative update option.
KB Article	3565
Severity	Critical

Table 756: Mellanox NIC Driver version check [111045]

Name	Mellanox NIC Driver version check
Description	Check if Mellanox NIC driver at a recommended version
Cause	Mellanox NIC driver is not at a recommended version
Impact	Current driver version may not support all functionality. In some cases using a non-recommended driver version may lead to network and service instability.

Resolution	Check LCM for an updated Mellanox NIC driver or otherwise refer to KB4289 for further detail
KB Article	4289

Table 757: AHV Crash File Check [111053]

Name	AHV Crash File Check
Description	Identify recent AHV crash file(s) on host.
Cause	Found kernel crash file(s) on AHV host(s).
Impact	Kernel issue may affect hypervisor functionalities.
Resolution	Notify Nutanix support to investigate the kernel issues.
KB Article	4866

Table 758: AHV Read-Only Filesystem Check [111055]

Name	AHV Read-Only Filesystem Check
Description	Check if any filesystem is read-only.
Cause	Found read-only filesystem(s) on AHV host(s).
Impact	The host may become unable to service I/O requests.
Resolution	Follow the instructions in the associated KB article to make the filesystems readable and writable.
KB Article	4897

Table 759: AHV CVM startup dependency check [111056]

Name	AHV CVM startup dependency check
Description	Check if AHV CVM is configured to auto start.
Cause	Missing auto start file(s) on AHV host(s).
Impact	AHV CVM might not be started properly.
Resolution	Notify Nutanix support to investigate the issue.
KB Article	4898

Table 760: Detected incompatible AHV version [111061] [A111061]

Name	AHV Version Check
Description	Check if compatible AHV version is installed on host.
Alert message	Upgrade the version of AHV on the host to *desired_ver* *or_later*.
Cause	The installed AHV version is not compatible with the current AOS version.

Impact	AHV is not running with the full set of features.
Resolution	Upgrade the version of AHV on the host to a version which is compatible with the current AOS version.
KB Article	5286
Severity	Warning

Table 761: AHV Time Zone Check [111071]

Name	AHV Time Zone Check
Description	Check whether the host's time zone is UTC, and warn that upgrading will change it.
Cause	The AHV host's time zone was changed/set to something other than UTC.
Impact	Non-default or inconsistent AHV host time zone may impact log collection and event correlation.
Resolution	Refer to KB6834 to change host time zone to UTC. AHV hypervisor upgrade will automatically reset time zone back to UTC.
KB Article	6834

Table 762: AHV Filesystem Integrity Check [111077]

Name	AHV Filesystem Integrity Check
Description	Check if EXT4 filesystem has errors.
Cause	Found corrupted EXT4 filesystem(s) on the AHV host(s).
Impact	The host may not be able to serve any filesystem I/O requests.This may cause VMs to crash, upgrades to fail, failure in CRUD operations, performance degradation, cluster instability among others.
Resolution	Follow the instructions in the associated KB article to recover from the filesystem corruption.
KB Article	9271

Table 763: Conntrack connection limit reached check [111079]

Name	Conntrack connection limit reached check
Description	Checks that the number of conntrack connections is below configured limit.
Cause	Number of connections in the conntrack table have exceeded the threshold.
Impact	User VMs may lose network access if limit is reached.
Resolution	Consider stopping some flows or migrate VMs. For higher number of allowed connections in conntrack table, upgrade the cluster to AOS 5.18 or higher. Refer to KB 9544 for more details.
KB Article	9544

Table 764: AHV bridge config check [111080]

Name	AHV bridge config check
Description	Checks if the load balancing mode is balance-tcp when LACP is configured, on all the uplink bridges on the host.
Cause	balance-slb and active-backup bond modes are not supported with LACP.
Impact	Use of other load balancing mode may impact network connectivity.
Resolution	Please change the bond mode to balance-tcp.
KB Article	9675

Table 765: AHV host is not schedulable for a long time [111084] [A111084]

Name	Node is not schedulable
Description	Node not schedulable for long period of time.
Alert message	AHV host *host_ip* is not schedulable for at least 6 hours.
Cause	Host is in maintenance mode or not reachable.
Impact	Cluster running with reduced host resources and may impact HA, ADS VM live migration and maintenance activities.
Resolution	Confirm the host is online and ensure no ongoing host maintenance tasks. Check KB 10128 if any issues exiting maintenance mode.
KB Article	10128
Severity	Info

Table 766: File integrity check found modifications. [111086] [A111086]

Name	AHV file integrity check
Description	Weekly file integrity check to detect any modifications in baseline configuration of host.
Alert message	File integrity tool has found modifications on host *host_ip*.
Cause	Detected unexpected modifications in baseline configuration of host.
Impact	Baseline configurations of host might have been changed in an unauthorized manner.
Resolution	Refer KB10414 to review the detected changes.
KB Article	10414
Severity	Warning

Table 767: Forwarding logs to remote server is failing. [111088] [A111088]

Name	AHV remote log forwarding check
Description	Check log forwarding to remote server.

Alert message	Forwarding logs to remote server is failing on host *host_ip*.
Cause	Forwarding logs to remote server is failing.
Impact	Remote server will not be receiving log output from cluster.
Resolution	Check failure causes for forwarding the logs to remote server. Refer KB10811.
KB Article	10811
Severity	Warning

Table 768: SSH keys missing [111089] [A111089]

Name	Node Missing SSH Keys check
Description	Check if /home/nutanix/ssh_keys directory contents are missing.
Alert message	*alert_msg*
Cause	Some of SSH keys are missing from /home/nutanix/ssh_keys directory.
Impact	Missing SSH keys can jeopardise ongoing infrastructure workflows in a silent manner such as node un-configuration.
Resolution	Please refer to KB 10736 for more information.
KB Article	10736
Severity	Critical

Table 769: Kernel mitigation configuration doesn't match current state. [111090] [A111090]

Name	Kernel mitigations check
Description	Check if spectre/meltdown mitigation is synced up with mitigations value in /proc/cmdline.
Alert message	*alert_msg*
Cause	Spectre/meltdown configuration doesn't match current state.
Impact	Uncertain configuration.
Resolution	Reboot CVM. Please refer to KB 10938 for more information.
KB Article	10938
Severity	Critical

Table 770: Disk utilization for audit partition is more than 75%. [111092] [A111092]

Name	AHV audit disk space check
Description	Checks if the disk usage for audit partition is higher than 75%.
Alert message	On host *host_ip*, *alert_msg*
Cause	Disk utilization for audit partition is more than 75%.
Impact	There may not be enough disk space left for new audit records.

Resolution	Refer KB11680 to address the disk utilization over 75%.
KB Article	11680
Severity	Warning

Table 771: AHV RPM database integrity check [111093]

Name	AHV RPM database integrity check
Description	Checks if RPM database is corrupt.
Cause	Found inconsistent RPM database(s) on AHV host(s).
Impact	AHV host upgrades may fail.
Resolution	Source of inconsistencies needs to be identified and fixed. RPM DB may need repair or host may need to be re-imaged. Refer to KB 10397 for details.
KB Article	10397

Table 772: Certificate Creation Error [130013] [A1112]

Name	Certificate Creation Status
Description	Certificate signing request creation failed.
Alert message	Failed to create a certificate signing request on node *service_vm_external_ip* in the cluster.
Cause	Either the OpenSSL library had a failure or a non-supported subject field, subject value, or subject alternative name was provided.
Impact	Without a CSR then data at rest encryption cannot be enabled.
Resolution	Ensure that the certificate signing request fields are valid and retry the command, or contact Nutanix support if the OpenSSL configuration file is missing or corrupted.
Severity	Info

Table 773: Agent VM Restoration Failure [130091] [A130091]

Name	Agent VM Restoration
Description	Failure to Restore Agent VM
Alert message	Failed to restore agent VM with UUID *vm_uuid*.
Cause	Not enough memory/CPU resources on this node.
Impact	VMs on this host may not function correctly because of the missing agent VM.
Resolution	Shut down unneeded VMs to free cluster resources or expand the cluster to add resources.
KB Article	8010
Severity	Critical

Table 774: VM Group Snapshot and Current State Mismatch [130110] [A130110]

Name	VM group snapshot and current mismatch
Description	Check that snapshot to restore matches current VM group
Alert message	VM group *vm_group_uuid* snapshot and current state mismatch: *message*
Cause	VM group changed after snapshot was created.
Impact	VMs might be scheduled to run on the same host instead of a different host.
Resolution	Recreate VM snapshot to ensure up-to-date VM group information.
Severity	Critical
Severity	Warning

Table 775: Protection Domain is in a Decoupled State [130112] [A130112]

Name	Protection Domain Decoupled Status
Description	Protection domain is in a decoupled state.
Alert message	Protection domain '*protection_domain_name*' protecting the Storage Container '*container_name*' is in a decoupled state. *impact*
Cause	Standby site has been promoted to the primary site.
Impact	Metro availability for the protection domain on the primary is unavailable.
Resolution	Re-enable metro availability from the remote site.
Severity	Warning

Table 776: Host Network Uplink Configuration Failed. [130160] [A130160]

Name	ClusterVswitch Configuration Update Failed
Description	Failed to apply host uplink configuration on the host(s).
Alert message	Host network uplink configuration update on Host(s):'*node_uuid*' failed due to the reason: *cause*
Cause	Host uplink configuration failed to apply on the host(s).
Impact	Host OVS configuration may not be in sync with cluster vswitch configuration.
Resolution	Use manage_ovs to configure the host individually after correcting errors.
Severity	Warning

Table 777: Failed to configure host for Atlas networking [130201] [A130201]

Name	Failed to configure host for Atlas networking
Description	Failed to configure host for Atlas networking.
Alert message	Failed to configure host *hypervisor_address* for Atlas networking: *message*
Cause	Failed to configure host for Atlas networking.

Impact	VMs on Atlas networks cannot run on this host.
Resolution	Check the hypervisor version, and upgrade if necessary.
KB Article	8526
Severity	Critical

Table 778: Virtual Switch configuration deployment failed. [150006] [A150006]

Name	Virtual Switch configuration deployment failed
Description	Notifies when a validated Virtual Switch configuration is unable to be applied to host network infrastructure
Alert message	Virtual Switch *vs_name* configuration deployment on Host(s):'*node_uuid*' failed due to the reason: *cause*
Cause	Unable to communicate with one or more host(s) to apply the configuration change.
Cause	Lost connectivity while applying the Virtual Switch configuration.
Cause	OVS did not accept or apply the requested configuration change.
Cause	Unexpected error with service(s) required for Virtual Switch configuration change (Acropolis, Genesis, IDF, Prism Gateway).
Impact	Unable to apply Virtual Switch configuration on the host(s).
Resolution	Ensure all hosts are up and reachable by all CVMs.
Resolution	Ensure the external network is properly configured to retain connectivity using the applied Virtual Switch configuration.
Resolution	Review the alert message and affected host OVS log files for further details.
Resolution	Review the alert message and related service Acropolis and Genesis log files for further details.
KB Article	9406
Severity	Critical
Severity	Warning

Table 779: Check to track bonds with no existing network device(s) [150007]

Name	Check to track bonds with no existing network device(s)
Description	Health check to track bonds with no existing network devices.
Cause	There are some interface(s) in the OVS bond for which no network device(s) exist.
Impact	If some interfaces don't have network devices configured, OVS will try to reconfigure the bonds which can cause race condition during deletion.
Resolution	Remove the interfaces having no network devices using OVS commands.
KB Article	10089

Table 780: IPFIX exporter host update failed [150011] [A150011]

Name	IPFIX exporter host update failed
Description	IPFIX exporter update failed on the host in the cluster.
Alert message	IPFIX exporter update failed on host *hypervisor_address* because of error: *cause*
Cause	IPFIX exporter update failed on the host in the cluster.
Impact	IPFIX exporter on this host in the cluster may not be working as expected.
Resolution	Check the error message and follow the appropriate steps to fix the IPFIX exporter state on the host.
Severity	Warning

Table 781: OVN Connection Unhealthy [150012] [A150012]

Name	OVN Connection Unhealthy
Description	Hypervisor node is disconnected from the Advanced Networking Controller.
Alert message	Node *uuid_node* is disconnected from the Advanced Networking Controller
Cause	The Chassis' ovn-controller does not have a Healthy connection to the Advanced Networking Controller.
Impact	Ability to make network-related configurations may be impacted.
Impact	Overlay networking may be disconnected on this Hypervisor node.
Resolution	Verify that the Hypervisor node can connect to the Prism Central hosting the Advanced Networking Controller. Verify the health of the Advanced Networking Controller.
KB Article	11129
Severity	Critical

Table 782: Conntrack Stats Collector Check [200604]

Name	Conntrack Stats Collector Check
Description	Checks if Flow records are being transferred to Prism Central from hypervisor
Cause	Conntrack stats collector on node hypervisor is offline, unconfigured or blocked by firewall. Kafka is offline on Prism Central.
Impact	Network Flow traffic for cluster node will not be visible in Prism Central
Resolution	Ensure that the firewall allows connection between registered Prism Central and AHV Hypervisor of the cluster. Ensure Kafka is running on Prism Central. Restart Conntrack stats collector.
KB Article	5765

Table 783: Flow policy-hit logging is not functional [200605] [A200605]

Name	Conntrack Mode Check
Description	Checks whether AHV is logging network traffic hitting Flow security policies.
Alert message	Network traffic hitting Flow security policies is not being logged.
Cause	Prism Central, Prism Element & AHV versions may be incompatible.
Cause	Memory requirements may not be met for policy-hit logging to be enabled.
Impact	Network traffic hitting security policies will not be logged when Flow is enabled on the cluster.
Resolution	Ensure AOS version is at least 5.11 on Prism Central and on Prism Element clusters managed by it.
Resolution	Increase memory available on AHV and upgrade to the latest AHV version.
KB Article	7098
Severity	Warning

Table 784: Openflow Table Getting Full [200607] [A200607]

Name	Openflow Table Full Check
Description	Checks that number of OVS rules is not reaching maximum number of rules
Alert message	Number of OVS rules is *num_flow_table_rules* for bridge *bridge_name* on AHV host *host_ip*. This is more than *usage_msg*% of maximum allowed number of rules.
Cause	Openflow table is getting full
Impact	When number of OVS rules will reach maximum, some VMs may lose network connectivity.
Resolution	Consider turning off unnecessary user VMs or remove unused vNICs from running user VMs on AHV host. For higher number of allowed OVS rules, ensure AOS version at least 5.10.5 and above on cluster.
KB Article	7563
Severity	Critical
Severity	Warning

Table 785: SSL Certificate is expiring soon [200705] [A200705]

Name	Check SSL Cert Expiry
Description	Notify if SSL Certificate is expiring soon
Alert message	SSL Certificate is expiring in the next 30 days. Please regenerate
Cause	SSL Certificate is expiring soon
Impact	SSL Certificates are expired and Cannot be accessed.
Resolution	Regenerate the ssl certificate.
KB Article	11493

Severity	Critical
Severity	Warning

Network

Table 786: CVM Connectivity Failure [3064] [A1001]

Name	CVM Passwordless Connectivity
Description	Check that CVM has passwordless connection to each other.
Alert message	Controller VM *target_ip* is not reachable from Controller VM *source_ip* in the last 3 attempts.
Cause	Controller VM is not running or is in a hung state.
Cause	Controller VM does not have the correct network configuration.
Cause	Upstream TOR switch has configuration or status change on the port connected to the CVM.
Cause	Virtual switch in the Hypervisor is not working correctly.
Cause	SSH Keys are set up incorrectly.
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Ensure the CVM is powered on and running. If the CVM is in a hung state, reference KB 1526 for further guidance.
Resolution	Correct the network configuration settings of the CVM.
Resolution	Ensure that the port connected to the CVM is up and in a forwarding state.
Resolution	Contact Nutanix Support.
Resolution	Contact Nutanix Support.
KB Article	1526
Severity	Critical

Table 787: Host IP Not Reachable [3065] [A1006]

Name	Host IP Pingable
Description	Checks if the local CVM can ping all host IP addresses.
Alert message	Hypervisor *target_ip* is not reachable from Controller VM *source_ip* in the last 6 attempts.
Cause	The hypervisor host is down or there is a network connectivity issue.
Impact	If the host and CVM are down, cluster compute and storage capacity are reduced and VM data may be under-replicated.
Resolution	Ensure that the hypervisor host is running and that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	2484
Severity	Critical

Table 788: NIC Link Down [3067] [A1082]

Name	NIC Link Status
Description	Checks whether any NIC is down.
Alert message	Link on NIC *nic_name* of host *host_ip* is down. NIC description: *nic_description*.
Cause	NIC Link is physically disconnected from the upstream Switch.
Cause	Upstream TOR switch port connected to the CVM has been disabled
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	Ensure that the cable between the node and the switch is plugged in properly on both ends.
Resolution	Ensure that the TOR switch port connected to the CVM is enabled and in forwarding state.
KB Article	2480
Severity	Warning

Table 789: Latency between CVMs is high [6005] [A6005]

Name	Inter-CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	Latency between *vm_type*s is higher than *ping_threshold* ms.
Cause	Malfunctioning or misconfigured network infrastructure
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.
Ping Time Threshold	Critical Threshold: 15 ms
KB Article	3690
Severity	Critical

Table 790: CVM NIC Speed Low [6008] [A1058]

Name	TenGig Adaptor Status
Description	Checks whether CVM is uplinked to 10GbE NIC.
Alert message	*alert_msg*
Cause	CVM not uplinked to 10GbE NIC.
Cause	1GbE NIC is part of the bond.
Cause	If hypervisor is ESXi, it may be in lockdown.
Impact	The Controller VM is not configured to use the 10GbE NIC or is configured to share load with a slower NIC.

Resolution	Check network configuration. Please refer to KB 1584.
Resolution	Check network configuration.
Resolution	If hypervisor is ESXi, check for Lockdown Mode enabled and disable the configuration.
KB Article	1584
Severity	Warning

Table 791: CVM Host Subnet Mismatch [6202] [A1048]

Name	CVM Subnet
Description	Checks that host and cvm share the same subnet.
Alert message	Controller VM *svm_ip* with network address *svm_subnet* is in a different network than the Hypervisor *hypervisor_ip*, which is in the network *hypervisor_subnet*.
Cause	The Controller VM and the hypervisor are not on the same subnet.
Impact	IO Redirection in the event of CVM failure will not work as expected, leading to guest VM unavailability.
Resolution	Reconfigure the cluster. Either move the Controller VMs to the same subnet as the hypervisor hosts or move the hypervisor hosts to the same subnet as the Controller VMs.
KB Article	1515
Severity	Critical

Table 792: Transmit Packet Drop Check [6404]

Name	Transmit Packet Drop Check
Description	Checks whether the host transmitted packets drop rate is high.
Cause	Transmitted packet drop rate is high.
Impact	Cluster performance may be degraded.
Resolution	Check network hardware/configuration.
KB Article	3706

Table 793: Receive Packet Drop Check [6405]

Name	Receive Packet Drop Check
Description	Checks whether the received packets drop rate is high.
Cause	Received packet drop rate is high.
Impact	Cluster performance may be degraded.
Resolution	Check network hardware/configuration.
KB Article	2883

Table 794: Unable to connect to vCenter. [6449] [A6449]

Name	vCenter Connection Check
Description	Check if vCenter connection is established.
Alert message	**vCenter connection is not established.** *cause_msg*, **Check KB 3815 for detailed information.**
Cause	vCenter server is unreachable on port 80.
Impact	Certain VM operations might fail.
Resolution	Confirm that vCenter server is reachable on port 80.
KB Article	3815
Severity	Critical

Table 795: CVM NIC Link Down [103094] [A103094]

Name	CVM NIC Link Status
Description	Checks whether any CVM NIC is down.
Alert message	**Link on NIC** *nic_name* **of cvm** *svm_ip* **is down.**
Cause	The rdma0 NIC is disconnected from the switch, or the switch port is failing.
Impact	Cluster performance may be significantly degraded. When rdma0 port is down, the CVM will switch to TCP mode via host vSwitch, and the RDMA mode of entire cluster is turned off. Critical large storage io service will suffer in non-RDMA mode.
Resolution	Ensure that the rdma0 NIC is connected to the switch and that the switch port is functioning properly.
KB Article	2480
Severity	Warning

Table 796: Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM [103095] [A103095]

Name	RDMA network switch and NIC PFC config check
Description	Checks if network switch DCBX and NIC PFC are correctly configured and enabled.
Alert message	**Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM** *cvm_ip*: *alert_msg*
Cause	Network switch DCBX is disabled or corresponding NIC PFC bit is not enabled.
Impact	RDMA traffic between CVMs is not enabled, which degrades cluster performance.
Resolution	Enable DCBX on network switch and configure PFC bit on RDMA NICs.
KB Article	5925
Severity	Warning

Table 797: Inconsistent Bridge/vSwitch configuration [103101] [A103101]

Name	NS Config Consistency Check
Description	Checks if the host bridge/vSwitch configuration matches NS config in zookeeper.
Alert message	Bridge/vSwitch *switch_name* on host *host_ip* has different configuration compared to other hosts or zeus configuration, which can result in service degradation.
Cause	In case the customer modifies configuration during cluster lifetime, but does not restart genesis, this will result in inconsistent configuration.
Impact	Potentially will break configured services (ABS/DR).
Resolution	Please refer to KB 8018.
KB Article	8018
Severity	Warning

Table 798: IPv6 Config check [103103]

Name	IPv6 Config check
Description	Check if there are any manual configuration of IPv6 on CVMs.
Cause	IPv6 has been configured manually on CVM interfaces.
Impact	Firewall rules configured using ip6tables would be erased.
Resolution	Remove the manual configuration and use manage_ipv6 CLI for IPv6 configuration. Follow KB-9399 for more information.
KB Article	9399

Table 799: CVM TCP checksum error rate check [103104]

Name	CVM TCP checksum error rate check
Description	Check if corrupted packets with compromised checksums are reaching the CVM and not dropped by the host NICs.
Cause	Hardware failure or NIC firmware/driver compatibility issue.
Impact	Cluster performance may be significantly degraded.
Resolution	Check for any hardware failures or issues with firmware/driver/cables. Please refer to KB 9623 for more details.
KB Article	9623

Table 800: CVM Interface Configuration Files Check [103105]

Name	CVM Interface Configuration Files Check
Description	Checks if the network interface configuration file for eth0 is not missing or malformed.

Cause	Manual update of the configuration file.
Impact	Affects cluster services and upgrades.
Resolution	Fix the malformed interface configuration file /etc/sysconfig/network-scripts/ifcfg-eth0.
KB Article	8378

Table 801: Bond uplink VLAN config check [103106]

Name	Bond uplink VLAN config check
Description	Checks if all uplinks in the bonds are configured with same set of VLANs.
Cause	Inconsistent VLAN configuration detected on the uplinks.
Impact	Hosts, CVM, or user VMs can lose network connectivity if active uplink changes. In balance-slb and balance-tcp bond modes, optimal performance might not be achieved.
Resolution	Check for VLAN configuration on switch ports where uplinks are connected.
KB Article	8185

Table 802: Bond uplink connectivity check [103107]

Name	Bond uplink connectivity check
Description	Checks if at least two uplinks in the bond are connected.
Cause	No cable is connected to the host uplink port. Cable, host NIC or switchport may be faulty. Switchport is in DOWN state.
Impact	If less than two uplinks in the bond are connected to network then the redundancy of network connectivity will be lost
Resolution	If the bond has two or more uplinks added, make sure that at least two of these uplinks are connected. Also check if cables and network hardware are healthy and switchport status is UP.
KB Article	8185

Table 803: Cloud Subnet Not Deleted [110025] [A110025]

Name	Cloud Subnet Not Deleted
Description	Cloud subnet not deleted during AOS network delete
Alert message	Cloud subnet for CIDR *cidr* not deleted: *message*
Cause	Cloud subnet could not be deleted as it has other resources(cloud network interfaces) that were created through cloud console or are managed by another cluster.
Impact	Subnets not in use may not get deleted, require deletion from cloud console

Resolution	If the subnet has some other NICs associated then you cannot delete it. Otherwise try cleaning up from cloud console.
KB Article	10795
Severity	Info

Table 804: Key Management Server Unavailable [110213] [A1099]

Name	Key Management Server Availability
Description	Check Key Management Server is available.
Alert message	Cannot connect to key management server *key_management_server_name* from Controller VM *service_vm_external_ip*.
Cause	The node cannot connect to a key management server either because of hardware failures, certificate validation failure, or network issues.
Impact	Secure clusters may not be able to unlock their drives on boot or change passwords.
Resolution	Ensure the key management server and nodes in a cluster have network connectivity to each other and all certificates are valid.
KB Article	3656
Severity	Critical

Table 805: Mellanox NIC Speed is not 10GbE and 40GbE or both 10GbE and 40GbE NICs are installed on one node [111047] [A111047]

Name	Mellanox NIC Status check
Description	Checks if Mellanox NICs are down or if any Mellanox NIC has speed other than 10GbE or 40GbE. Checks if both 10GbE and 40GbE Mellanox NICs are installed on one node.
Alert message	Mellanox NIC on host *host_ip* has problem: *alert_msg*
Cause	The NIC is disconnected from the switch, or the switch port is not functioning correctly, or both 10GbE and 40GbE Mellanox NICs are installed on one node.
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	Ensure that all the Mellanox NIC installed on one node run on the same speed.
KB Article	4114
Severity	Warning

Table 806: Mellanox NIC not Installed or with Wrong Type on host machine [111051] [A111051]

Name	Mellanox NIC Mixed Family check
Description	Checks if Mellanox NICs belong to the same Mellanox NIC family.

Alert message	**Mellanox NIC installed on host : *host_ip* belong to different families.**
Cause	Mellanox NICs are not from the same Mellanox NIC family.
Impact	Mixing Mellanox NICs from different families may cause degraded cluster performance.
Resolution	Replace Mellanox NICs to make sure they belong to same Mellanox NIC family.
KB Article	4329
Severity	Warning

Table 807: OVA Upload Interrupted [111076] [A111076]

Name	OVA Upload Interrupted
Description	OVA upload interrupted
Alert message	**Upload interrupted for OVA *ova_name***
Cause	Network might have been disconnected.
Cause	Prism UI may have been closed.
Cause	OVA file might not be accessible.
Impact	OVA will not be available for use.
Resolution	Validate stable network connection.
Resolution	Ensure Prism UI is open and logged in.
Resolution	Ensure OVA file is accessible.
Severity	Info

Table 808: Failed to configure host for Atlas networking [130201] [A130201]

Name	Failed to configure host for Atlas networking
Description	Failed to configure host for Atlas networking.
Alert message	**Failed to configure host *hypervisor_address* for Atlas networking: *message***
Cause	Failed to configure host for Atlas networking.
Impact	VMs on Atlas networks cannot run on this host.
Resolution	Check the hypervisor version, and upgrade if necessary.
KB Article	8526
Severity	Critical

Table 809: Failed to reserve host memory for Atlas networking [130202] [A130202]

Name	Failed to reserve host memory for Atlas networking
Description	Failed to reserve host memory for Atlas networking.

Alert message	**Failed to reserve host memory for Atlas networking**
Cause	Failed to reserve host memory for Atlas networking.
Impact	If hypervisor memory is exhausted, random VMs may be powered off to reclaim memory.
Resolution	Power off VMs to reduce host memory consumption.
KB Article	8525
Severity	Critical

Table 810: IPFIX exporter host update failed [150011] [A150011]

Name	IPFIX exporter host update failed
Description	IPFIX exporter update failed on the host in the cluster.
Alert message	**IPFIX exporter update failed on host** *hypervisor_address* **because of error:** *cause*
Cause	IPFIX exporter update failed on the host in the cluster.
Impact	IPFIX exporter on this host in the cluster may not be working as expected.
Resolution	Check the error message and follow the appropriate steps to fix the IPFIX exporter state on the host.
Severity	Warning

Table 811: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed
Description	Programming a Flow rule failed.
Alert message	**Flow policy operation failed for** *rule_name*. *reason*
Cause	Microsegmentation service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.
Resolution	Check PC microsegmentation service, PE acropolis service, PC to PE connection, PE to AHV connection.
Severity	Warning

Table 812: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed
Description	Flow failure event.
Alert message	**Flow operation failed on** *remote_uuid*. *reason*

Cause	Microsegmentation service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microsegmentation service, Check PE acropolis service, AHV microsegmentation module.
Severity	Warning

Table 813: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	Flow mode change failed on *remote_uuid*. *reason*
Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect cvm_logs and contact Nutanix Support
Severity	Warning

Table 814: Maximum VPN BGP route limit reached [801001] [A801001]

Name	Max VPN eBGP route limit reached
Description	On-prem VPN gateway is advertising more than the maximum number of routes accepted by the VPN gateway over eBGP.
Alert message	A total of *num_routes* routes received from the on-prem VPN gateway for connection "*vpn_connection_name*". Only *max_num_routes* routes will be processed.
Cause	On-prem VPN gateway is advertising more than the maximum number of routes accepted by the VPN gateway over eBGP.
Impact	Connectivity between AZs via the VPN connection might be impacted.
Resolution	Consider aggregating more specific prefixes into larger prefixes to reduce the number of routes advertised.
KB Article	8469
Severity	Critical

Table 815: VPN IPSEC tunnel between VPN endpoints is down [801002] [A801002]

Name	VPN IPSEC Tunnel down
Description	The IPSEC tunnel between the VPN endpoints is down.
Alert message	VPN IPSEC tunnel for VPN connection "*vpn_connection_name*" is down.
Cause	The IPSEC tunnel between the VPN endpoints is down.
Impact	Connectivity between AZs via VPN connection will be impacted.
Resolution	Follow the on-prem VPN device vendor's troubleshooting steps. If you suspect that the problem is not with the on-prem VPN device, please contact Nutanix support.
KB Article	8466
Severity	Critical

Table 816: eBGP session between VPN endpoints is down [801003] [A801003]

Name	eBGP session between VPN gateway peers down
Description	The eBGP session between VPN endpoints is down.
Alert message	eBGP session for VPN connection "*vpn_connection_name*" is down.
Cause	The eBGP session between VPN endpoints is down.
Impact	VPN connectivity between the VPN endpoints will be impacted because routes cannot be exchanged.
Resolution	Follow the on-prem VPN device vendor's troubleshooting steps. If you suspect that the problem is not with the on-prem VPN device, please contact Nutanix support.
KB Article	8465
Severity	Critical

Table 817: Invalid routes received by VPN connection "{vpn_connection_name}". [801004] [A801004]

Name	Invalid routes received from on-prem VPN
Description	Invalid routes received from on-prem VPN gateway.
Alert message	The following invalid routes received by VPN connection "*vpn_connection_name*" will be rejected: *message*.
Cause	Invalid routes received from on-prem VPN gateway.
Impact	VPN connectivity may be impacted.
Resolution	Remove the invalid routes being advertised from on-prem VPN gateway over eBGP.
Severity	Warning

Table 818: Deletion of the Layer-2 subnet extension failed on the remote site. [801101] [A801101]

Name	Layer-2 subnet extension deletion failure
Description	Deletion of the Layer-2 subnet extension failed on the remote site.
Alert message	Deletion of the Layer-2 subnet extension "#layer2_stretch_name@" failed on the remote site.
Cause	Deletion of the remote Layer-2 subnet extension failed (possibly due to AZ connection issues).
Impact	The remote site's Prism Central will show the remote subnet as being extended to the local site, even though the subnet extension has been deleted locally.
Resolution	Delete the Layer-2 subnet extension on the remote site by logging into the remote site's Prism Central and navigating to the Subnets page.
KB Article	10310
Severity	Critical

Table 819: ANC version does not support Layer-2 subnet extension. [801102] [A801102]

Name	ANC version does not support Layer-2 subnet extension
Description	ANC version does not support Layer-2 subnet extension.
Alert message	ANC version does not support Layer-2 subnet extension.
Cause	ANC version does not support Layer-2 subnet extension.
Impact	ARP messages (and other unknown unicast messages) from peer AZ will be dropped, causing reachability issues between VMs on the stretched subnet across AZs.
Resolution	Upgrade the ANC to the right version that has support for Layer-2 subnet extension.
KB Article	10388
Severity	Critical

Table 820: VPN gateway version does not support Layer-2 subnet extension. [801103] [A801103]

Name	VPN gateway version does not support Layer-2 subnet extension
Description	VPN gateway version does not support Layer-2 subnet extension.
Alert message	VPN gateway version does not support Layer-2 subnet extension "#layer2_stretch_name@".
Cause	VPN gateway version does not support Layer-2 subnet extension.
Impact	ARP messages (and other unknown unicast messages) from peer AZ will be dropped, causing reachability issues between VMs on the stretched subnet across AZs.
Resolution	Upgrade the VPN gateway to the right version that has support for Layer-2 subnet extension.
KB Article	10389

Severity	Critical

Table 821: VPN connection associated with Layer-2 subnet extension not found. [801104] [A801104]

Name	VPN connection associated with Layer-2 subnet extension not found
Description	VPN connection associated with Layer-2 subnet extension not found.
Alert message	VPN connection associated with Layer-2 subnet extension "#layer2_stretch_name@" not found.
Cause	VPN connection associated with Layer-2 subnet extension was deleted.
Impact	VMs across AZs on the stretched subnet cannot communicate with each other.
Resolution	Delete the Layer-2 subnet extension and create it again when VPN connection is restored.
KB Article	10390
Severity	Critical

Table 822: Peer AZ involved in Layer-2 subnet extension is not reachable. [801105] [A801105]

Name	Peer AZ involved in Layer-2 subnet extension is not reachable
Description	Peer AZ involved in Layer-2 subnet extension is not reachable.
Alert message	Peer AZ involved in Layer-2 subnet extension "#layer2_stretch_name@" is not reachable.
Cause	The cause could either be VPN connection between the AZs is down, or Atlas service in the peer AZ is down (or not responsive).
Impact	VMs in the subnet extended across AZs may be unable to communicate with each other.
Resolution	Fix the underlying issue of VPN connection being down, or Atlas service in peer AZ being down.
KB Article	10391
Severity	Critical

Table 823: Subnet involved in Layer-2 subnet extension not found. [801106] [A801106]

Name	Subnet involved in Layer-2 subnet extension not found
Description	Subnet involved in Layer-2 subnet extension not found.
Alert message	Subnet involved in Layer-2 subnet extension "#layer2_stretch_name@" not found.
Cause	Subnet involved in Layer-2 subnet extension was deleted.
Impact	If subnet is not present, no vNICs can be created using this network UUID.
Resolution	Delete the Layer-2 subnet extension, and create it again after the subnet has been restored.

KB Article	10392
Severity	Critical

Table 824: CIDR of subnets involved in Layer-2 subnet extension don't match. [801107] [A801107]

Name	CIDR of subnets involved in Layer-2 subnet extension don't match
Description	CIDR of subnets involved in Layer-2 subnet extension don't match.
Alert message	CIDR of subnets involved in Layer-2 subnet extension "#layer2_stretch_name@" don't match.
Cause	CIDR of subnets involved in Layer-2 subnet extension don't match.
Impact	Some VMs may not be able to communicate with other VMs on the extended subnet.
Resolution	Modify one of the subnets to fix the CIDR mismatch.
KB Article	10393
Severity	Critical

Table 825: DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses. [801108] [A801108]

Name	DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses
Description	DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses.
Alert message	DHCP of subnets involved in Layer-2 subnet extension "#layer2_stretch_name@" have overlapping IP addresses or include VPN interface IP addresses.
Cause	DHCP of subnets involved in Layer-2 subnet extension have overlapping IP addresses or include VPN interface IP addresses.
Impact	Some VMs may not be able to communicate with other VMs on the extended subnet.
Resolution	Modify the subnets involved in the Layer-2 subnet extension to ensure DHCP pools do not overlap, and do not contain the local or remote VPN interface IP.
KB Article	10394
Severity	Critical

Table 826: Local VPN interface IP involved in Layer-2 subnet extension is in use in peer AZ. [801109] [A801109]

Name	Local VPN interface IP involved in Layer-2 subnet extension is in use in peer AZ
Description	Local VPN interface IP involved in Layer-2 subnet extension is in use in peer AZ.

Alert message	Local VPN interface IP involved in Layer-2 subnet extension "#layer2_stretch_name@" is in use in peer AZ.
Cause	A vNIC in peer AZ was assigned the same IP address as VPN interface IP used in the local AZ.
Impact	Some VMs in the subnets involved in Layer-2 subnet extension will be unable to communicate with other VMs in peer AZ.
Resolution	Resolve the IP conflict by ensuring the VPN interface IP is not used for any user VMs in any of AZs involved in the Layer-2 subnet extension.
KB Article	10395
Severity	Critical

Table 827: Remote VPN interface IP involved in Layer-2 subnet extension is in use in this AZ. [801110] [A801110]

Name	Remote VPN interface IP involved in Layer-2 subnet extension is in use in this AZ
Description	Remote VPN interface IP involved in Layer-2 subnet extension is in use in this AZ.
Alert message	Remote VPN interface IP involved in Layer-2 subnet extension "#layer2_stretch_name@" is in use in this AZ.
Cause	A vNIC in this AZ was assigned the same IP address as VPN interface IP used in the remote AZ.
Impact	Some VMs in the subnets involved in Layer-2 subnet extension will be unable to communicate with other VMs in peer AZ.
Resolution	Resolve the IP conflict by ensuring the VPN interface IP is not used for any user VMs.
KB Article	10396
Severity	Critical

Table 828: Some IP address(es) are common across the subnets involved in the Layer-2 subnet extension. [801111] [A801111]

Name	Some IP address(es) are common across the subnets involved in the Layer-2 subnet extension
Description	Some IP address(es) are common across the subnets involved in the Layer-2 subnet extension
Alert message	Some IP address(es) are common across subnets involved in Layer-2 subnet extension "#layer2_stretch_name@". Detailed error text: *error_detail*
Cause	Some UVMs are allocated IP addresses that are in use in the peer AZ involved in Layer-2 subnet extension.
Impact	Some VMs in the subnets involved in Layer-2 subnet extension will be unable to communicate with other VMs in peer AZ.

Resolution	Resolve the IP conflict by ensuring the IP addresses allocated to UVMs are unique across the AZs involved in the Layer-2 subnet extension.
KB Article	10604
Severity	Critical

Table 829: Advanced Networking Controller is not Healthy [802001] [A802001]

Name	An Advanced Networking Controller service did not respond
Description	Atlas could not connect to an Advanced Networking Controller service.
Alert message	The service at *url* did not respond
Cause	The Advanced Networking Controller is not Healthy.
Impact	Ability to make network-related configurations may be impacted.
Resolution	Verify the health of the Advanced Networking Controller.
KB Article	9366
Severity	Critical

Table 830: Unable to resolve an Advanced Networking Controller service's DNS name [802002] [A802002]

Name	An Advanced Networking Controller service's DNS name was not resolvable
Description	Atlas was unable to resolve an Advanced Networking Controller service's DNS name.
Alert message	The IP corresponding to *url* could not be resolved with the following DNS servers: *dns_servers*
Cause	PC nameserver configuration is incorrect.
Impact	Ability to make network-related configurations may be impacted.
Resolution	Verify nameserver configuration on PC.
KB Article	9365
Severity	Critical

Table 831: Unable to apply Virtual Private Cloud Routing Policy. [802003] [A802003]

Name	A Virtual Private Cloud Routing Policy of type reroute is inactive
Description	Atlas was unable to apply a VPC Routing Policy of type reroute.
Alert message	Routing Policy with priority *priority* and type reroute in VPC *vpc_name* is not applied as VM/ service with reroute IP: *reroute_ip* is down.
Cause	VM(s) with reroute IP(s) does not exist or is powered off.
Impact	Traffic in VPC will not hit the inactive Routing Policy.
Resolution	(A) Create VM with reroute IP in the given VPC. (B) Power on VM with reroute IP in the given VPC.

KB Article	9484
Severity	Warning

Table 832: ID Firewall lost connectivity to domain controller [803003] [A803003]

Name	ID Firewall is unable to connect to one or more domain controllers
Description	ID Firewall is unable to connect to one or more domain controllers used for scraping identities.
Alert message	Domain controller *domain_controller* on domain *domain* is not reachable and accepting LDAP or WMI connections. Details: *details*.
Cause	The domain controller may not be reachable via the network.
Cause	The domain controller may not be accepting LDAP/WMI connections due to permissions issues.
Cause	The domain controller may not be running or fully booted up.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until access to the domain controller is restored.
Resolution	Check network connectivity from Prism Central to the domain controller. Please refer article KB-10219 for more details.
Resolution	Check that the domain controller is fully up and accepting WMI and LDAP connections. Please refer article KB-10219 for more details.
Resolution	Check that the service account registered for ID Firewall is active and has permissions for both WMI and LDAP access, and that the correct password is entered in Prism Central. Please refer article KB-10219 for more details.
KB Article	10219
Severity	Warning

Table 833: ID Firewall did not recover state after reconnecting to a temporarily unreachable domain controller [803005] [A803005]

Name	ID Firewall state recovery failed
Description	ID Firewall was unable to recover state after reconnecting to a temporarily unreachable domain controller.
Alert message	ID Firewall was unable to recover state from domain controller *domain_controller* on domain *domain*. Details: *details*.
Cause	Recovery failed because either too much time has passed since the domain controller was last reachable, or the event log on the domain controller has rolled over and the events necessary to recover are no longer present.
Impact	The state enforced by ID Firewall may be inconsistent, and therefore applied policies may not be properly enforced until action is taken.
Resolution	In order to ensure security policies are properly enforced, all active Nutanix VDI VM users should log out and log back into their VMs. Please refer KB-10220 for further details.

KB Article	10220
Severity	Warning

Table 834: ID Firewall service account is invalid [803007] [A803007]

Name	Service account used for ID Firewall is invalid
Description	ID Firewall has detected that the service account configured is not valid.
Alert message	ID Firewall service account *account*@*domain* is invalid. Details: *details*.
Cause	The service account password may have been changed.
Cause	The service account may have been deleted.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until the service account is valid again.
Resolution	Check that the configured service account password is up-to-date and update it in Prism if needed.
Resolution	Check that the configured service account exists, and update it in Prism if needed.
Severity	Warning

Nutanix Files

Table 835: File Server Space Usage High [160000] [A160000]

Name	File Server Space Usage
Description	The usage of file server space is high.
Alert message	File Server space usage for *file_server_name* is at *usage_pct*%.
Cause	File Server Storage utilization is high.
Impact	If storage usage keeps growing, the shares will be marked as read-only.
Resolution	Expand File Server storage size. Ask users to delete unused files. Refer to KB article 8475 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8475
Severity	Warning

Table 836: File Server Space Usage Critical. [160001] [A160001]

Name	File Server Space Usage Critical
Description	File server space usage critical.
Alert message	File Server space usage for *file_server_name* is at *usage_pct*%.
Cause	File Server storage utilization has reached a critical value.

Impact	The shares will be marked as read-only until there is more free space on the File Server.
Resolution	Expand File Server storage size. Refer to KB article 8475 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8475
Severity	Critical

Table 837: File Server Unreachable. [160002] [A160002]

Name	File Server Unavailable Check
Description	File server is unreachable.
Alert message	File server *file_server_name* is unreachable.
Cause	FSVMs are not reachable.
Impact	No operations can be executed against the File Server.
Resolution	From Controller VM run 'ncc health_checks fileserver_checks run_all' and resolve the reported issues. If you are unable to resolve the issue reported by NCC, contact Nutanix Support for assistance.
Severity	Warning

Table 838: File Server storage is not available. [160003] [A160003]

Name	File Server Storage Status
Description	File Server storage is not available.
Alert message	Storage for File Server *file_server_name* is unavailable.
Cause	File Server storage is unavailable due to network connectivity issues between FSVM and CVM.
Impact	Read and write operations on the File Server will fail.
Resolution	From Controller VM run 'ncc health_checks run_all' and resolve the reported issues. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 839: File Server scale-out failed. [160004] [A160004]

Name	File Server Scale-out Status
Description	File Server scale-out failed.
Alert message	File Server *file_server_name* scale-out failed because of *reason*
Cause	See details in the alert message.
Impact	Unable to scale-out the file server.

Resolution	From Controller VM run 'ncc health_checks run_all' and resolve the reported issues. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 840: File server could not join the AD Domain [160005] [A160005]

Name	File Server Join Domain Status
Description	File server could not join the AD domain.
Alert message	File Server *file_server_name* could not join the Domain because *reason*.
Cause	See details in the alert message. Error could be due to one of the following reasons: Domain credentials are invalid. Correct Port for connecting to AD is not being used. File Server AD Object already exists in the domain.
Impact	Domain operations can not be performed on the File Server as it has not joined the domain
Resolution	The resolution is to identify the reason from one of the possible causes in the cause list. Please refer Nutanix Files Guide 'Pre-requisite' section for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 841: Node Failed To Join Domain [160006] [A160006]

Name	File Server Node Join Domain Status
Description	The node could not join the domain.
Alert message	The node could not join the domain for file server *file_server_name* as *reason*
Cause	DNS is not correctly configured or is not reachable. Domain Controllers are not reachable or the reachable Domain controller is having issues. Domain credentials are invalid. Correct Port for connecting to AD is not being used. File Server AD Object already exists in the domain.
Impact	The node could not join the domain.
Resolution	The resolution is to identify the reason from one of the possible causes in the cause list. Please refer Nutanix Files Guide 'Pre-requisite' section for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 842: File Server Time Difference High [160007] [A160007]

Name	FSVM Time Drift Status
Description	The time drift between the FSVMs is beyond the acceptable range

Alert message	For *file_server_name*, time drift between the FSVMs (*lower_time_ip* and *higher_time_ip*) is more than the acceptable value of *time_difference_limit_secs*
Cause	NTP is not configured correctly or the NTP service is not functioning.
Impact	The Nutanix Files cluster may become inaccessible
Resolution	Check that the NTP service is running and is reachable from the FSVMs Refer KB article 8135 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8135
Severity	Warning

Table 843: File Server Storage Cleanup Failed [160012] [A160012]

Name	File Server Storage Cleanup Failure
Description	Failed to clean up storage for the file server.
Alert message	Storage *message* for the file server *file_server_name* is not cleaned up.
Cause	Acropolis service may be down on the cluster.
Impact	File server storage will not be released.
Resolution	Contact Nutanix support if issue still persists or assistance needed.
Severity	Critical

Table 844: File server cannot connect with AD server [160013] [A160013]

Name	File Server AD Connectivity Failure
Description	File server cannot connect with AD server with configured information
Alert message	File server *file_server_name* cannot connect with AD server
Cause	The machine account credentials may have been changed. There is a network connectivity issue for the AD server.
Impact	Domain operations cannot be performed for the file server. Client connectivity with file-server is disabled.
Resolution	Check that the AD server is reachable. Refer to KB article 8122 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8122
Severity	Critical

Table 845: File Server performance optimization recommended [160015] [A160015]

Name	File Server Performance Optimization Recommended
Description	File server has a recommendation to optimize performance by using scale-up, scale-out or rebalance.

Alert message	A recommendation is available to optimize the performance on one or more nodes of file server *file_server_name*.
Cause	File server has one or more nodes under extensive load.
Impact	File server performance may be impacted and new client connections may be refused.
Resolution	Run 'Performance optimization' for the specified file server. Please refer Nutanix Files 'System limits' section in the release notes for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	7035
Severity	Warning
Severity	Info

Table 846: User Quota Assignment Failed [160016] [A160016]

Name	File Server Quota allocation failed for user
Description	Failed to assign the specified quota to the user
Alert message	Failed to apply quota for user: *user_name* on share: *share_name* of File Server *file_server_name*
Cause	The user has used more space on the File Server than the specified quota value.
Impact	User has no quota limit
Resolution	Notify user to reduce their usage below the quota value. Notify user to increase the quota allocation for the user.
Severity	Warning

Table 847: Share utilization reached configured limit [160017] [A160017]

Name	Share Utilization Reached Configured Limit
Description	Share is no longer writable.
Alert message	Utilization on *share_name* on file server *file_server_name* has reached the configured limit
Cause	Share utilization reached its configured limit
Impact	Writes to the share will fail until one of the actions is taken.
Resolution	Expand File Server storage size. Refer to KB article 8475 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8475
Severity	Critical

Table 848: File Server failed to get updated CVM IP address. [160018] [A160018]

Name	File Server CVM IP update failed

Description	File server Controller VM IP update failed.
Alert message	The file server *file_server_name* could not get updated CVM IP address.
Cause	AOS cluster is not able to contact file server.
Impact	File server is not reachable
Resolution	Ensure correct Cluster IP Address Configuration procedure is followed. Refer to KB article 2290 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	2290
Severity	Critical

Table 849: Appropriate Site Not Found in Active Directory [160019] [A160019]

Name	File Server Site Not Found
Description	Unable to determine an appropriate site in Active Directory
Alert message	File server *file_server_name*'s client network is not mapped to a site on the Active Directory.
Cause	File server client network is not part of any site in Active Directory.
Impact	File server may take a long time to join the domain, depending on which domain controller is selected.
Resolution	Add the client network of file server to a local site in Active Directory for optimal performance.
Severity	Warning

Table 850: File server DNS Updates Pending [160020] [A160020]

Name	File Server DNS Updates Pending
Description	DNS updates are pending after a file server operation.
Alert message	File server *file_server_name*'s DNS entries need to be programmed on nameservers. *message*
Cause	Failed to add or remove DNS entries
Impact	The file server cannot be used by clients by its name, but can be used with the IP addresses.
Resolution	Correct the required DNS entries under 'DNS' on the file server page. Refer KB article 3722 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	3722
Severity	Warning
Severity	Info

Table 851: File Server activation failed [160021] [A160021]

Name	File Server Activation Failed
Description	File Server Activation Failed
Alert message	File server *file_server_name* activation failed due to *reason*
Cause	Check alert message for details
Impact	File Server is not usable.
Resolution	Try to activate file server again, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 852: Failed to set VM-to-VM anti-affinity rule [160023] [A160023]

Name	Failed To Set VM-to-VM Anti Affinity Rule
Description	Failed to set VM-to-VM anti affinity rule
Alert message	Failed to set VM-to-VM anti-affinity rule for *file_server_name*. *reason_and_resolution_msg*
Cause	More than one FSVM of a File Server entity is may be deployed hosted on a single physical host.
Impact	More than one FSVM of a File Server entry is hosted on a single physical host, which will cause data unavailability in case that host fails.
Resolution	Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 853: Discovery of iSCSI targets failed. [160025] [A160025]

Name	File Server iSCSI Discovery Failure
Description	Failed to discover iSCSI targets on the CVM during the discovery process.
Alert message	Discovery of iSCSI targets failed for file server *file_server_name*
Cause	The 'external_data_services_ip' or the CVM IP addresses are not reachable from the FSVM.
Impact	Share operations and File server HA will fail.
Resolution	Please refer article 8216 for more details on Changing Data Services IP Address. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8216
Severity	Critical
Severity	Warning

Table 854: File Server upgrade failed [160026] [A160026]

Name	File Server Upgrade Failed

Description	File Server Upgrade Failed.
Alert message	Upgrade of file server *file_server_name* failed due to *reason*.
Cause	Check alert message for details
Impact	File server has not been upgraded.
Resolution	From Controller VM run 'ncc health_checks fileserver_checks run_all' and resolve the reported issues. If you are unable to resolve the issue reported by NCC, contact Nutanix Support for assistance.
Severity	Warning

Table 855: Incompatible File Server activated [160028] [A160028]

Name	Incompatible File Server Activation
Description	Activating an incompatible File Server
Alert message	Incompatible File server *file_server_name* is activated
Cause	File server is activated which is incompatible with current AOS version
Impact	Next AOS upgrade will fail.
Resolution	Upgrade file server to the latest. Please refer Software Product Interoperability page on portal for reference. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 856: File Server in heterogeneous state. [160032] [A160032]

Name	File Server In Heterogeneous State
Description	File server in heterogeneous state. Nodes do not match in their CPU or memory configuration.
Alert message	File server *file_server_name* is in heterogeneous state, as File Server VM configuration update operation failed due to *reason*
Cause	Check alert message for details.
Impact	Performance of the file server is less than optimal.
Resolution	Get the largest configured CPU and memory values among VMs of reported File Server in VM page. Free up resources on hypervisor hosts and retry File Server VM configuration update with these values in File Server page. Contact Nutanix support if issue still persists or assistance needed.
Severity	Critical

Table 857: Failed to correct File Server data and meta data inconsistencies [160034] [A160034]

Name	Failed To Run File Server Metadata Fixer Successfully

Description	Failed to Run File Server Metadata Fixer tool successfully
Alert message	Failed to run File Server *file_server_name* Metadata fixer task successfully. *reason_and_resolution_msg*
Cause	Metadata fixer task might have timed out
Impact	Some of the top level directories would be unavailable on file server
Resolution	Run 'afs share.fix_tld_metadata' command form one of the FSVM shell to fix the metadata inconsistency. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 858: File server share deletion failed [160035] [A160035]

Name	File Server Share Deletion Failed
Description	Failed to delete share.
Alert message	Failed to delete share *share_name* on File Server *file_server_name*. *message*
Cause	Multiple causes possible. Most likely, a failure in destroying file systems from FSVMs.
Impact	Unable to clean up share storage space.
Resolution	Please retry Share delete operation. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 859: File server compatibility check skipped [160036] [A160036]

Name	Skipped File Server Compatibility Check
Description	File server compatibility check skipped
Alert message	File server compatibility check skipped
Cause	File server compatibility check was skipped during AOS upgrade
Impact	Some management functions of the file server may not be available.
Resolution	Upgrade file server to the latest. Please refer Software Product Interoperability page on portal for reference. Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 860: Failed to add one or more file server administrator users or groups [160038] [A160038]

Name	Failed to add one or more file server admin users or groups
Description	Failed to add one or more users or groups as file server administrators
Alert message	File server *file_server_name*: *failure_msg* Update file server administrators on the file server page

Cause	One or more users or groups could not be resolved on Active Directory.
Impact	File server administrator users or groups may not function as administrators
Resolution	Update file server administrators on the file server page
Severity	Info

Table 861: Maximum connections limit about to reach on a file server VM [160039] [A160039]

Name	Maximum connections limit reached on a file server VM
Description	Maximum connections limit is about to reach on a file server VM
Alert message	File server VM *fsvm_name* on file server *file_server_name* has reached *conn_percentage*connections limit.
Cause	File server has one or more nodes under extensive load.
Impact	File server performance may be impacted and new client connections may be refused.
Resolution	Run 'Performance optimization' for the specified file server. Refer to KB article 7371 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	7371
KB Article	7371
Severity	Critical
Severity	Warning

Table 862: File Server Clone failed [160040] [A160040]

Name	File Server Clone Failed
Description	File Server Clone Failed
Alert message	File server *file_server_name* clone failed due to *reason*
Cause	Check alert message for details
Impact	Clone File Server not available.
Resolution	Check alert message, Retry clone operation after rectifying, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 863: File Server rename failed [160041] [A160041]

Name	File Server Rename Failed
Description	File Server Rename Failed
Alert message	File server rename from *file_server_name* to *file_server_new_name* failed due to *reason*
Cause	Check alert message for details

Impact	File server could not be renamed.
Resolution	Check alert message, Retry rename operation after rectifying, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 864: File Server AntiVirus - ICAP Server Down [160042] [A160042]

Name	File Server AntiVirus - ICAP Server Down
Description	ICAP server is not responding to scan requests.
Alert message	ICAP server *scan_server_host* is not responding from file server *file_server_name*.
Cause	Failed to reach ICAP server from FSVM.
Impact	File server antivirus scanning is degraded.
Resolution	Check ICAP server and connectivity between FSVMs to ICAP server. Refer to KB article 8414 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8414
Severity	Warning

Table 865: File Server AntiVirus - All ICAP Servers Down [160043] [A160043]

Name	File Server AntiVirus - All ICAP Servers Down
Description	All configured ICAP servers are not responding to scan requests.
Alert message	All configured ICAP servers are not responding from file server *file_server_name*.
Cause	Failed to reach all configured ICAP server from FSVM.
Impact	File server antivirus scanning is not possible. Access is decided by scan policy.
Resolution	Check ICAP server and connectivity between FSVMs to ICAP server. Refer to KB article 8414 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8414
Severity	Critical

Table 866: File Server network change failed [160044] [A160044]

Name	File Server Network Change Failed
Description	File Server Network Change Failed
Alert message	File server *file_server_name* network change failed due to *reason*
Cause	Check alert message for details
Impact	File Server is not usable.

Resolution	Try to change file server network again, and if the failure persists, then contact Nutanix support.
Severity	Warning

Table 867: File Server AntiVirus - Scan Queue Full on FSVM [160045] [A160045]

Name	File Server AntiVirus - Scan Queue Full on FSVM
Description	Antivirus Scan Queue is Full on FSVM.
Alert message	Antivirus Scan Queue is Full on FSVM *nvm_name* of file server *file_server_name*.
Cause	ICAP Daemon is not able to keep up with in coming requests.
Impact	File server antivirus scanning is slow, clients may observe high latency.
Resolution	Check ICAP server status. Check number of ICAP servers configured. Check workload on FSVM.
Severity	Critical

Table 868: File Server AntiVirus - Scan Queue Piling Up on FSVM [160046] [A160046]

Name	File Server AntiVirus - Scan Queue Piling Up on FSVM
Description	Antivirus Scan Queue is Piling Up on FSVM.
Alert message	Antivirus Scan Queue is Piling Up on FSVM *nvm_name* of file server *file_server_name*.
Cause	Incoming scan request rate is higher than the scan processing rate.
Impact	File server antivirus scanning is slow, clients may observe high latency.
Resolution	Check ICAP server status. Check number of ICAP servers configured. Check workload on FSVM.
Severity	Warning

Table 869: File Server AntiVirus - Excessive Quarantined / Unquarantined Files [160047] [A160047]

Name	File Server AntiVirus - Excessive Quarantined / Unquarantined Files
Description	Excessive Quarantined / Unquarantined Files on File Server.
Alert message	File Server *file_server_name* has a lot of quarantined / unquarantined files.
Cause	Number of quarantined / unquarantined files on file server is approaching limit.
Impact	File server will stop scanning the files once the limit is reached. Default action specified in the scan policy will be applied.
Resolution	Delete unwanted quarantined / unquarantined files, check ICAP server for false positives or check system for virus activity.
Severity	Warning

Table 870: File Server AntiVirus - Quarantined / Unquarantined Files Limit Reached [160048] [A160048]

Name	File Server AntiVirus - Quarantined / Unquarantined Files Limit Reached
Description	Quarantined / Unquarantined Files Limit is Reached on File Server.
Alert message	File Server *file_server_name* has reached quarantined / unquarantined files limit.
Cause	Number of quarantined / unquarantined files on file server reached limit.
Impact	File server stopped scanning files. Default action in the scan policy is applied.
Resolution	Delete unwanted quarantined / unquarantined files, check ICAP server for false positives or check system for virus activity. Refer to KB article 8416 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8416
Severity	Critical

Table 871: File Server time is out of sync with the Active Directory [160049] [A160049]

Name	File Server time is out of sync with the Active Directory
Description	File server time is out of sync with the Active Directory domain controllers.
Alert message	Time drift between the file server VMs and the Active Directory is at *time_drift_secs* seconds on file server *file_server_name*. If the drift becomes more than 300 seconds, the clients may get disconnected from the fileserver.
Cause	NTP is not configured correctly or the NTP service is not functioning.
Impact	The Nutanix Files cluster may become inaccessible.
Resolution	Check that the NTP service is running and the NTP servers are reachable from the file server VMs. Refer to KB article 8404 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8404
Severity	Warning

Table 872: File Server User Management Configuration Failed [160050] [A160050]

Name	File Server User Management Configuration Failed
Description	Given user management options could not be configured for the file server.
Alert message	The values provided for File server *file_server_name*'s user management options are incorrect.
Cause	Failed to configure user management with given options
Impact	The file server may have restricted or insecure access from various clients.
Resolution	Update correct entries against fields in the User Management tab. Refer to KB article 9123 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	9123

Severity	Warning

Table 873: /home partition usage on a file server VM higher than threshold [160051] [A160051]

Name	File Server Disk Usage Warning
Description	Space consumption under /home partition on file server VM is higher than the set threshold
Alert message	/home partition usage on File server VM *fsvm_name* on file server *file_server_name* at *usage_pct* (greater than threshold *threshold_pct*).
Cause	File server has one or more nodes under extensive load.
Impact	Deleting a share from a File server fails.
Resolution	Remove unwanted data or files from under /home for the specified file server. Refer to KB article 8478 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8478
Severity	Critical
Severity	Warning

Table 874: File server DNS records cannot be refreshed [160052] [A160052]

Name	File Server DNS - Unable to refresh DNS records
Description	File server DNS records periodic refresh failed
Alert message	File server *file_server_name* DNS records cannot be updated
Cause	DNS entries may not be created with Nutanix Files computer account due to insufficient permissions.
Impact	If scavenging is active on the DNS servers, then Nutanix Files DNS records may get removed. This can cause clients to not connect with Nutanix Files using its DNS hostname.
Resolution	Create static DNS entries for the fileserver manually using 'DNS' button. Update required permissions for the Nutanix Files computer account. Refer to KB article 8405 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8405
Severity	Critical

Table 875: File Server Share Backup diff path translation failed [160053] [A160053]

Name	File Server Backup Diff Path Translation Warning
Description	Backup diff found inodes whose path translation failed
Alert message	Backup diff found inodes whose path translation failed *file_server_name*, error information: *failure_msg*

Cause	File server share may have hardlinks.
Impact	Some files may not have been backed up.
Resolution	Refer to KB article 8874 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8874
Severity	Warning

Table 876: File Server Partner Server Connectivity Down [160054] [A160054]

Name	File Server Partner Server Connectivity Down
Description	Partner server is not responding to file notifications.
Alert message	Partner server *partner_server_host* is not responding from file server *file_server_name*.
Cause	Failed to reach partner server from File Server VM.
Impact	File server stopped notifying file operation events.
Resolution	Ensure that the partner server is functioning and that there is connectivity between File Server VMs and the partner server.
Severity	Critical

Table 877: File Server DNS resolver IP connectivity Failure [160055] [A160055]

Name	File Server DNS resolver IP connectivity Failure
Description	File server cannot resolve its domain name using configured DNS resolver IP addresses.
Alert message	File server *file_server_name* cannot resolve its domain name using configured DNS resolver IP addresses. *message*
Cause	None of the configured DNS resolver IP addresses are reachable. None of the configured DNS resolver IP addresses are resolving the file server domain name.
Impact	File server services may be unavailable and new clients may not be authenticated.
Resolution	Ensure that the servers specified by DNS resolver IP addresses are functional. Refer to KB article 8406 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8406
Severity	Warning

Table 878: File Server PD action to incompatible Remote Site AOS [160057] [A160057]

Name	File Server PD action warning for non compatible Remote Site AOS
Description	Remote Site AOS version is not compatible with File Server version.

Alert message	File server *file_server_name* protection domain *protection_domain_name* enabled on non-compatible Remote site *remote_site_name*
Cause	Remote Site AOS version is not compatible with File Server version.
Impact	Migrate-to or Activate-on Remote Site would fail.
Resolution	Upgrade Remote Site AOS version similar or above of primary site. Please refer Software Product Interoperability page on portal for reference. Contact Nutanix support if issue still persists or assistance needed.
Severity	Critical
Severity	Warning

Table 879: File Server NTP servers connectivity failure [160058] [A160058]

Name	File Server NTP servers connectivity failure
Description	File server cannot connect to NTP server.
Alert message	File Server *file_server_name* cannot reach NTP servers specified. *reason*
Cause	None of the configured NTP servers are reachable.
Impact	File server services may become inaccessible.
Resolution	Ensure that the NTP servers are functional and are reachable from the file server VMs. Refer to KB article 8135 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8135
Severity	Critical
Severity	Warning
Severity	Info

Table 880: File server services got interrupted [160059] [A160059]

Name	File Server services got interrupted
Description	One of the fileserver services has restarted.
Alert message	File server *file_server_name* services got interrupted. Error Message: *message*
Cause	One of the fileserver services has restarted.
Impact	Some fileserver service would be unavailable to NAS clients momentarily.
Resolution	If the service continues to be down, consider restarting minerva_ha service.
Severity	Warning

Table 881: Share usage reaching to configured limit [160061] [A160061]

Name	File Server Share Space Usage Reaching Configured Limit
Description	Share will be no longer be writable.

Alert message	Share *share_name* on file server *file_server_name* is approaching capacity
Cause	Share usage reaching its configured limit
Impact	Writes to the share will fail if share gets full.
Resolution	Increase storage allocation to the share or direct users to free up space by removing unneeded files. Please refer KB article 8690 for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
KB Article	8690
Severity	Warning

Table 882: Pulse cannot connect to REST server endpoint on File Server [160062] [A160062]

Name	File Server Failed to Reach Pulse REST Endpoint
Description	Pulse cannot connect to REST server endpoint on File Server.
Alert message	Pulse cannot connect to REST server endpoint on File Server *file_server_name*. Connection Status: *connection_status*, Pulse Enabled: *enabled*, Error Message: *message*
Cause	REST server endpoint is not reachable from Pulse.
Impact	Data driven serviceability and customer support cannot be performed for affected File Server.
Resolution	Ensure that the REST server endpoint is reachable from Pulse on File Server.
Severity	Info

Table 883: Storage pool usage reaching its limit [160063] [A160063]

Name	File Server Storage Pool Usage Reaching Configured Limit
Description	Shares on storage pool will be no longer writable.
Alert message	*share_names* shares on file server *file_server_name* are about to run out of space. *command*
Cause	Storage pool usage is reaching its limit
Impact	Writes to the shares may fail if their storage pool gets full.
Resolution	Increase storage allocation to the share or direct users to free up space by removing unneeded files. Please refer KB article 8690 for more details. If you are unable to resolve the issues, contact Nutanix Support for assistance.
Severity	Warning

Table 884: Duplicate IP address detected for a File Server VM. [160068] [A160068]

Name	File Server Duplicate IP Address Detected
Description	File server VMs have IP address conflict.
Alert message	Duplicate IP address detected for a file server VMs for *file_server_name*. Error Message: *message*
Cause	Another host has started using a file server VM IP address.

Impact	Some file server operations may fail due to invalid network configuration.
Resolution	Update IP address on host which has an IP address conflict with file server VM. Refer to KB article 8951 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8951
Severity	Warning

Table 885: File Server Unique Fsid failure [160069] [A160069]

Name	File Server Unique Fsid Failure
Description	Assigning Unique Fsid failed
Alert message	Found unique fsid failure *file_server_name*, error information: *failure_msg*
Cause	Duplicate Unique Fsid present
Impact	NFS handles may not be unique.
Resolution	Contact Nutanix support if issue still persists or assistance needed.
Severity	Warning

Table 886: File Server VM Time Drift from NTP Servers [160072] [A160072]

Name	File Server VM Time Drift from NTP Servers
Description	Difference between FSVM clock and NTP servers is too large
Alert message	File Server VM *fsvm_name* clock is different from NTP servers for *time_diff_secs* seconds. Acceptable time difference *time_diff_limit_secs* seconds.
Cause	NTP servers are not publishing time correctly. File server VM time is not syncing with NTP servers.
Impact	File server cluster may become inaccessible
Resolution	Verify NTP servers are functional and reachable. Refer to KB article 8135 for more details. Contact Nutanix support if the issue still persists or assistance is needed.
KB Article	8135
Severity	Warning
Severity	Info

Table 887: File Server Service in Crash Loop [160074] [A160074]

Name	File server service in crash loop
Description	File server service in crash loop
Alert message	File server service *crashed_service_name* has crashed *crash_count* times in *crash_interval_secs* seconds.

Cause	File server service has crashed multiple times in short interval
Impact	Some file server features may not be functioning
Resolution	Refer to KB article 8418 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	8418
Severity	Critical

Table 888: File server service got interrupted [160101] [A160101]

Name	File server service got interrupted
Description	One of the fileserver services has restarted.
Alert message	File server *file_server_name* service got interrupted. Error Message: *message*
Cause	One of the fileserver services has restarted.
Impact	Some fileserver service would be unavailable to NAS clients momentarily.
Resolution	If the service continues to be down, consider restarting minerva_ha service.
Severity	Warning

Table 889: File Server Active Directory operation failed during Disaster Recovery [160103] [A160103]

Name	File Server Active Directory operation failed during Disaster Recovery
Description	Active Directory operations during failover or failback operation as part of Disaster Recovery failed.
Alert message	File server *file_server_name* Disaster Recovery operation failed to update Active Directory. *message*
Cause	Various. See message for details.
Impact	Clients connected to previous file server may not work with the current file server after Disaster Recovery.
Resolution	Various. See message for details.
KB Article	9805
Severity	Critical

Table 890: File Server Disaster Recovery - target site not reachable [160105] [A160105]

Name	File Server Disaster Recovery - target site not reachable
Description	File server target site is not reachable for Disaster Recovery
Alert message	File server *file_server_name* Disaster Recovery - Target site is not reachable. *message*
Cause	Network connectivity issue between the source and the target.

Impact	Disaster Recovery operation may be impacted. Share data may not be replicated on time.
Resolution	Please check the network connectivity between the source and the target.
KB Article	9857
Severity	Critical

Table 891: File Server Disaster Recovery - not able to meet RPO [160106] [A160106]

Name	File Server Disaster Recovery - not able to meet RPO
Description	File server Disaster Recovery feature is not able meet current RPO standard.
Alert message	File server *file_server_name* Disaster Recovery feature is not able meet current RPO standard. *message*
Cause	An error on source or target file server can cause replication to proceed slowly.
Impact	Disaster Recovery related share replication may not be up to date as per the RPO standard.
Resolution	Please check the alert message for details.
KB Article	9872
Severity	Warning

Table 892: File Server Disaster Recovery - both source and target are active file servers [160107] [A160107]

Name	File Server Disaster Recovery - both source and target are active file servers
Description	File server Disaster Recovery feature has both source and target file servers set as active with read-write shares.
Alert message	File server *file_server_name* Disaster Recovery feature has both source and target file servers set as active with read-write shares. *message*
Cause	Configuration error during Disaster Recovery setup on source and/or target file servers.
Impact	Data inside shares may be modified on both source and target and may not remain in sync.
Resolution	Please make the target file server as standby with read-only shares.
KB Article	9873
Severity	Warning

Table 893: File Server connected share path was not found [160108] [A160108]

Name	File Server Connected Share Path not found
Description	Path for connected share is not available.

Alert message	File server *file_server_name* connected share *child_share_name* path *submount_path* was not found.
Cause	The submount path doesn't exist.
Impact	Share won't be accessible from submount path
Resolution	Remove submount path and reconfigure the share with new or recreated submount path
Severity	Warning

Table 894: File Server Disaster Recovery - source file server has more nodes than target file server [160110] [A160110]

Name	File Server Disaster Recovery - source file server has more nodes than target file server
Description	File server Disaster Recovery - number of nodes on source and target file servers must be the same.
Alert message	File server *file_server_name* Disaster Recovery - Target site expansion required. *message*
Cause	Number of file server nodes on source and target has changed due to scale-out operation.
Impact	Data inside shares may not remain in sync.
Resolution	Expand or reduce the node count on source and target file servers to make the node count identical on both file servers.
KB Article	9858
Severity	Critical

Table 895: File Server Disaster Recovery - target file server has more nodes than source file server [160111] [A160111]

Name	File Server Disaster Recovery - target file server has more nodes than source file server
Description	File server Disaster Recovery - number of nodes on source and target file servers must be the same.
Alert message	File server *file_server_name* Disaster Recovery - Source site expansion required. *message*
Cause	Number of file server nodes on source and target has changed due to scale-out operation.
Impact	Data inside shares may not remain in sync.
Resolution	Expand or reduce the node count on source and target file servers to make the node count identical on both file servers.
KB Article	9858
Severity	Warning

Table 896: File Server Disaster Recovery - failover operation failed [160112] [A160112]

Name	File Server Disaster Recovery - failover operation failed
Description	File server - failed to take over primary file server during Disaster Recovery Failover operation.
Alert message	Failed to take over primary file server during Disaster Recovery Failover operation on *file_server_name*. *message*
Cause	Various. Please see alert message for more details.
Impact	Primary file server clients may not access the shares/exports.
Resolution	Correct the error and retry the operation.
Severity	Warning

Table 897: File Server Disaster Recovery - failback operation failed [160113] [A160113]

Name	File Server Disaster Recovery - failback operation failed
Description	File server - failed to take over current file server resource during Disaster Recovery Failback operation.
Alert message	Failed to take over current file server resources during Disaster Recovery Failback operation on *file_server_name*. *message*
Cause	Various. Please see alert message for more details.
Impact	Current file server clients may not access the shares/exports.
Resolution	Correct the error and retry the operation.
Severity	Warning

Table 898: File Server Disaster Recovery - Source and Target Have Mismatched Protocols Configured [160114] [A160114]

Name	File Server Disaster Recovery - Source and Target Have Mismatched Protocols Configured
Description	File server - failed to establish source and target relationship for shares due to mismatched protocol configuration.
Alert message	One or more shares are not protected for *file_server_name*. *message*
Cause	Source and target file servers are not configured to support the same protocols.
Impact	One or more shares will not be protected.
Resolution	Configure source and target file servers with the same protocols.
Severity	Warning

Table 899: File Server Disaster Recovery - SSR snapshot schedule migration operation failed [160116] [A160116]

Name	File Server Disaster Recovery - SSR snapshot schedule migration operation failed
Description	File server - failed to migrate SSR snapshot schedule.
Alert message	Failed to migrate SSR snapshot schedules on *file_server_name*. *message*
Cause	The Alert message includes the cause and the details of the SSR schedule migration failure.
Impact	SSR snapshots will not be replicated to target.
Resolution	Run the SSR snapshot migration via afs cli. Please refer to KB article 11593 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	11593
Severity	Warning

Table 900: File Server Domain is Offline [160120] [A160120]

Name	File Server Domain is Offline
Description	File server domain is offline.
Alert message	File Server *file_server_name* domain *domain_name* is offline on node *node_name*.
Cause	There is a network connectivity issue for the AD server.
Impact	Domain operations cannot be performed for the file server. Client connectivity with file-server is disabled.
Resolution	Check that the AD server is reachable. Refer to KB article 10074 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	10074
Severity	Critical

Table 901: File Server Domain Controller is detected slow. [160121] [A160121]

Name	File Server Domain Controller is detected slow
Description	File Server domain controller is detected slow.
Alert message	File Server *file_server_name* domain controller *dc_name* is detected slow on node *node_name*.
Cause	There could be network connectivity issue the AD server or the kerberos user ticket is getting resolved after multiple redirections to child domain having the information.
Impact	Domain operations cannot be performed for the file server. Client connectivity with file-server is disabled.
Resolution	Check that the AD server is reachable. Refer to KB article 10076 for more details. Contact Nutanix support if issue still persists or assistance needed.
KB Article	10076

Severity	Warning

Table 902: AIDE Baselines configuration changes on File Server. [160130] [A160130]

Name	File Server AIDE Baseline Configuration Changes Check
Description	Checks the log files of aide for changes to baseline configs if aide is enabled
Alert message	Detected changes to baseline configs on a File Server. *message*
Cause	Baseline configurations changed on a File Server.
Impact	Aide execution may fail depending on the files that were changed
Resolution	Check changed files and revert their changes if needed
KB Article	10461
Severity	Warning

Table 903: Remote syslog server forwarding failed on a File Server. [160131] [A160131]

Name	File Server Remote Syslog Server Forwarding Check
Description	Check remote syslog server log forwarding on a File Server.
Alert message	Remote syslog server forwarding failures detected on a File Server. *message*
Cause	Remote syslog server forwarding failed on a File Server.
Impact	Remote syslog server will not be receiving log output from the File Server.
Resolution	Check for forwarding failure causes with rsyslog server.
KB Article	10735
Severity	Warning

Table 904: File Server Tiering Service is not healthy [160135] [A160135]

Name	File Server Tiering Service is not healthy
Description	Tiering service is not in a healthy state and can lead to temporary unavailability of tiered files.
Alert message	*message*
Cause	Tiering service might be overloaded or is unavailable.
Impact	Delayed processing of tiering operations and can lead to temporary unavailability of tiered files.
Resolution	Check file server cluster health and refer to KB11073 for more details.
KB Article	11073
Severity	Warning

Table 905: File server Tiering Service failed to mark the bucket for cleanup [160136] [A160136]

Name	File server Tiering Service failed to mark the bucket for cleanup
Description	Tiering service failed to mark the bucket for garbage collection.
Alert message	*message*
Cause	The following causes are possible:File Server not in healthy stateProblems with Object Store Connectivity
Impact	Lifecycle policies are not accurately configured on the bucket and can retain objects longer than intended.
Resolution	Please cleanup the Object Store. Refer to KB11074 for more details. Contact Nutanix support for assistance.
KB Article	11074
Severity	Warning

Table 906: File Server/Share(s) are deleted; Object Store cleanup is required [160137] [A160137]

Name	File Server/Share(s) are deleted; Object Store cleanup is required
Description	File Server/Share(s) are deleted, please cleanup the Object Store.
Alert message	*message*
Cause	File Server/Share(s) are deleted on the cluster.
Impact	Objects can remain longer than intended. After the file server is deleted, we cannot accurately track the objects for garbage collection.
Resolution	Please cleanup the Object Store. Refer to KB11075 for more details. Contact Nutanix support for assistance.
KB Article	11075
Severity	Info

Table 907: Disk space usage for audit volume high on a file server. [160138] [A160138]

Name	File Server System Audit Volume Space Usage High
Description	Checks if FSVM system audit volume usage is within threshold to ensure uninterrupted operations.
Alert message	System Audit Volume Space Usage High on a file server. *message*
Cause	Increased FSVM system audit volume usage due to excessive logging, or miscellaneous files being placed inside.
Impact	If FSVM system audit volume is 100% utilized, then auditd will stop and further system auditing will not be recorded.
Resolution	Remove files other than audit.log* from /home/log/audit. If unsure, please engage Nutanix Support for assistance.
KB Article	1523

Severity	Critical
Severity	Warning

Table 908: File Server Share Conversion Alert [160139] [A160139]

Name	File Server Share Conversion Alert
Description	File server share conversion - see message for details
Alert message	Share conversion for *file_server_name*: *message*
Cause	See details in the alert message
Impact	See details in the alert message
Resolution	See details in the alert message
Severity	Info

Table 909: Versioning is suspended for the bucket in the file server tiering profile. [160142] [A160142]

Name	File Server Tiering - Bucket versioning is suspended
Description	Versioning is suspended for the bucket in the file server tiering profile.
Alert message	*message*
Cause	Versioning is suspended for the bucket in the file server tiering profile.
Impact	Tiering is not allowed if versioning is suspended for the bucket in the file server tiering profile.
Resolution	Enable versioning for the bucket in the file server tiering profile. Please refer to the vendor-specific documentation for bucket versioning configuration.
Severity	Warning

Other

Table 910: CVM renamed incorrectly. [3035] [A3035]

Name	CVM name check
Description	Check that CVM has been named properly.
Alert message	CVM renamed incorrectly.
Cause	CVM was renamed after installation
Impact	Cluster upgrade will get stuck.
Resolution	Rename CVM back to its original name. CVM name is nonconfigurable by the end user.
KB Article	5226
Severity	Warning

Table 911: Storage container(s) {ctr_list} is/are built on top of removed storage pools [20024] [A20024]

Name	Storage container on removed storage pool
Description	Check if any storage container is provisioned on storage pools that is marked for removal.
Alert message	The following storage containers are built on removed storage pools: *description_place_holder*
Cause	There are some storage containers that are built on top of removed storage pools. This can potentially block disk removal.
Impact	Can potentially cause a disk removal to get stuck.
Resolution	Refer to the KB-9491 for the cleanup procedure.
KB Article	9491
Severity	Warning

Table 912: Network adapter setting check [103068]

Name	Network adapter setting check
Description	Check Network adapter setting
Cause	Network adapter is not set to not connect on power on
Impact	CVM may not boot up after the host reboot.
Resolution	Check if "ethernet0.startConnected = "true"" is present in the CVM's .vmx file.
KB Article	2050

Table 913: PYNFS dependency check [106431]

Name	PYNFS dependency check
Description	Check PYNFS dependency
Cause	PYNFS is in use and is not present.
Impact	CVM may not boot up after the host reboot.
Resolution	Validate PYNFS configuration.
KB Article	2050

Table 914: vim command check [106433]

Name	vim command check
Description	Check that vim-cmd vmsvc/power.on command is present in local script
Cause	vim-cmd vmsvc/power.on\ command not present in local script.
Impact	CVM may not boot up after the host reboot.
Resolution	Check if 'vim-cmd vmsvc/power.on' entry is present in 'local.sh'.

KB Article	2050

Table 915: Autobackup check [106436]

Name	Autobackup check
Description	Check that /sbin/auto-backup.sh has run successfully
Cause	/sbin/auto-backup.sh has not run successfully.
Impact	CVM may not boot up after the host reboot.
Resolution	Make sure '/bootbank/state.tgz' has a newer timestamp.
KB Article	2050

Table 916: EOF check [106437]

Name	EOF check
Description	Check that there is no line "EOF" at end of rc local script
Cause	"EOF" statement at end of rc local.
Impact	CVM may not boot up after the host reboot.
Resolution	Check that 'local.sh' does not have 'EOF'
KB Article	2050

Table 917: RC local script exit statement present [106438]

Name	RC local script exit statement present
Description	Check that top level exit statement is not present in RC local script
Cause	Top level exit statement present in script rc local preventing script lines from being run.
Impact	CVM may not boot up after the host reboot.
Resolution	Check if 'local.sh' has an 'exit' statement. Generate INFO if the exit statement is NOT within the if..fi statement
KB Article	2050

Table 918: .dvsData directory in local datastore [106439]

Name	.dvsData directory in local datastore
Description	Check that .dvsData directory is present on pynfs mounted local datastore
Cause	.dvsData directory is not persistent yet.
Impact	CVM may not boot up after the host reboot.
Resolution	Check if .dvsData directory exists in the local datastore.

KB Article	2050

Table 919: Incorrectly configured local datastore. [106440] [A106440]

Name	Svmboot.iso mount check
Description	Checks if CVM ISO is on mounted local datastore
Alert message	*alert_msg*
Cause	ServiceVM_Centos.iso is not on mounted local datastore.
Cause	The local datastore name may be renamed to non-default.
Impact	CVM may not start after the host reboot.
Impact	If local datastore name is non-default, ESXi upgrade may stall.
Resolution	Check if ServiceVM_Centos.iso exists in the local datastore.
Resolution	Please review the local datastore name on the host by logging into vCenter. The datastore should be named as NTNX-local-ds-xxxxxx.
KB Article	2050
Severity	Warning

Table 920: Incorrect vSphere cluster configuration [106448] [A106448]

Name	ESXi vSphere Cluster config check
Description	Check if the nodes in a Nutanix cluster are configured correctly on a vSphere Cluster
Alert message	*warn_msg*
Cause	One or more ESXi hosts which are part of the Nutanix cluster may be disconnected from the registered vCenter Server, not part of inventory, or registered across multiple vSphere Clusters
Impact	Prism UVM management operations, 1-click hypervisor upgrade, DRS, HA and other vSphere automation may be hampered.
Resolution	Review KB 5369
KB Article	5369
Severity	Warning

Table 921: Protection domain contains volume group(s) using client IP address to whitelist a client. [110262] [A110262]

Name	Protected Volume Group with IP address in Whitelist Check
Description	Check if any volume group with IP address in whitelist is protected in protection domain
Alert message	Volume group(s) *vg_names* in protection domain *pd_name* use client IP address to whitelist a client.

Cause	Protection domain contains volume group(s) using client IP address to whitelist a client.
Impact	The attachment between VM and VG is not recoverable from snapshot
Resolution	Update the client properties to include the client IQN instead of the IP address to configure the client list for the volume group(s).
KB Article	6036
Severity	Warning

Table 922: Protected VM Name Too Long or Contains Restricted Characters. [110267] [A110267]

Name	VM name does not meet requirements.
Description	Check if the protected VM name is either too long, or contains restricted characters.
Alert message	Restore of the protected VM *vm_name* will fail due to VM name containing restricted characters or VM name exceeding the character limit.
Cause	VM name is either too long (longer than 80 characters) or contains restricted characters. Restricted characters include [,], ?, *, ", ', ,.
Impact	The VM will not be restored during failovers.
Resolution	Rename the VM.
KB Article	6097
Severity	Warning

Table 923: Nearsync staging area stale files check [110272]

Name	Nearsync staging area stale files check
Description	Check for stale files in Nearsync staging area.
Cause	Nearsync temporary files in staging area were not properly cleaned up.
Impact	Cluster storage capacity is unnecessarily reduced. Cluster may run out of space.
Resolution	Stale files in the temporary staging area need to be removed manually.
KB Article	9728

Table 924: Too Many Recovery Points in the System [110404] [A110404]

Name	Recovery Point limit is approaching system set threshold.
Description	Checks if the number of recovery points is approaching or has exceeded the supported limit.
Alert message	There are *num_snapshot* recovery points in the system which is approaching or has exceeded system supported limit of *snapshot_limit*. Crossing the supported limit may result in the DR service on the cluster becoming unstable.

Cause	Number of recovery points in the system is approaching or has exceeded the supported limit.
Impact	If number of recovery points exceeds supported limit, DR service may become unstable.
Resolution	Delete all the unwanted recovery points in the system including any out of band recovery points or reduce retention of recovery points for existing policies or remove recovery points existing from older VMs after unplanned failover.
KB Article	8935
Severity	Warning
Severity	Info

Table 925: Stale Recovery Points Check [110405]

Name	Stale Recovery Points Check
Description	Checks if there are any recovery point entries in the Insights database whose underlying snapshots have been deleted.
Cause	Snapshot was deleted without removing the associated stale entries in Insights database.
Impact	Upcoming replications might be impacted because of the remnant stale entries in Insights database.
Resolution	Check the running status of the Polaris service.
KB Article	10420

Table 926: Detected snapshots on cluster with high density nodes [110450] [A110450]

Name	High density nodes snapshots check
Description	Checks if snapshots are present on a cluster with high density nodes
Alert message	Protection domain snapshots are detected on cluster with high density nodes
Cause	Detected snapshots on the cluster with high density nodes
Impact	Snapshots are not currently supported on clusters with high density nodes. They may impact cluster operations
Resolution	Delete snapshots on the cluster with high density nodes
KB Article	7003
Severity	Info

Table 927: Data protection is configured on cluster with high density nodes [110451] [A110451]

Name	High density nodes DR configuration check
Description	Checks if disaster recovery is configured on a cluster with high density nodes

Alert message	**Disaster recovery entities like protection domains, remote sites are detected on cluster with high density nodes**
Cause	Remote sites or protection domains are detected on cluster with high density nodes
Impact	Data protection is currently not supported on clusters with high density nodes, this may impact cluster operations
Resolution	Delete disaster recovery entities like remote sites, protection domains on cluster with high density nodes
KB Article	7003
Severity	Info

Table 928: NGT CA Setup Check [111046]

Name	NGT CA Setup Check
Description	Checks that the NGT CA setup is same on all nodes.
Cause	/home/ngt/ca.tar does not have the same checksum on all CVMs.
Impact	NGT operations may not function properly.
Resolution	Contact Nutanix support.
KB Article	4254

Table 929: CVM active upgrade check [111062]

Name	CVM active upgrade check
Description	Check if there are any active upgrades
Cause	An upgrade was requested for this entity
Impact	Maintenance activities should wait until upgrade is finished.
Resolution	Manually upgrade AOS. Contact Nutanix support for assistance.
KB Article	5277

Table 930: AOS upgrades are disabled on {cvm_ip} [111063] [A111063]

Name	AOS upgrades are disabled check
Description	Check if AOS upgrades are disabled on CVM
Alert message	**AOS upgrades are disabled on *cvm_ip***
Cause	AOS upgrade has been paused manually.
Impact	Upgrades might fail because of stale entries
Resolution	Manually upgrade AOS. Contact Nutanix support for assistance.
KB Article	5277

Severity	Info

Table 931: Metro Availability Operation Failed [130124] [A130124]

Name	Metro Availability Operation Failed
Description	Metro availability operation failed
Alert message	For the protection domain '*protection_domain_name*', *operation* failed to the remote site '*remote_name*'. Reason: '*reason*'.
Cause	Check the alert message for the reason of failure
Impact	Metro availability operation could not be started.
Resolution	Resolve the issue as stated in the alert message and retry the Metro operation. If the issue persists contact Nutanix Support
KB Article	10111
Severity	Critical

Table 932: Recovery Point Objective Cannot Be Met [130138] [A130138]

Name	Recovery Point Objective Cannot Be Met
Description	Recovery Point Objective Cannot Be Met
Alert message	Recovery point objective cannot be met because '*reason*'
Cause	Various
Impact	Recovery plan could be affected.
Resolution	Contact Nutanix support if this condition persists
Severity	Warning

Table 933: Network Segmentation Config Update Failed [130363] [A130363]

Name	NS Config Update Failed
Description	Network Segmentation Config update during AOS upgrade failed
Alert message	*message*
Cause	Network Segmentation Configuration update during AOS upgrade failed
Impact	All Network Segmentation operations will be blocked
Resolution	Try CLI command network_segmentation fix_config_after_upgrade on any CVM. Contact Nutanix Support in case of command failure.
Severity	Warning

Table 934: Free Xi Account Expired [200701] [A200701]

Name	Free Xi Account Expired
Description	Free period using Xi expired. If this issue persists, you may experience an interruption in service.
Alert message	Free period using Xi expired. If this issue persists, you may experience an interruption in service.
Cause	The free period is limited by time.
Resolution	Select a payment plan and method in the billing section under the top right gear menu.
Severity	Critical

Table 935: Xi Subscription Expired [200702] [A200702]

Name	Xi Subscription Expired
Description	Xi subscription period expired. If this issue persists, you may experience an interruption in service.
Alert message	Xi subscription period expired. If this issue persists, you may experience an interruption in service.
Cause	The subscription period is limited by time.
Resolution	Select a new payment plan in the billing section under the top right gear menu.
KB Article	8255
Severity	Critical

Table 936: Xi Payment Missed [200703] [A200703]

Name	Xi Payment Missed
Description	Unable to process payment. If this issue persists, you may experience an interruption in service.
Alert message	Unable to process payment. If this issue persists, you may experience an interruption in service.
Cause	This is likely due to an expired or canceled credit card.
Resolution	Update payment method in the billing section under the top right gear menu.
KB Article	8254
Severity	Critical

Table 937: Xi Payment Missed [200704] [A200704]

Name	Xi Payment Missed

Description	Unable to process payment. If this issue persists, you may experience an interruption in service.
Alert message	Xi tenant *tenant_uuid* invoice payment for *plan_name* plan failed at *timestamp_usecs*
Cause	This is likely due to a change in address or a forgotten or misplaced invoice.
Resolution	Ensure invoices are sent to the correct contact information provided in the billing section under the top right gear menu. Confirm with your team that the invoice wasn't misplaced or forgotten. Contact support if needed.
KB Article	8253
Severity	Critical

Table 938: Nucalm Internal Service has Stopped Working [400101] [A400101]

Name	Nucalm Service Down
Description	Nucalm internal service is down
Alert message	Discovered that the nucalm internal service is not working: '*service_names*'
Cause	Nucalm internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-7080
KB Article	7080
KB Article	8621
Severity	Critical

Table 939: Epsilon Internal Service has Stopped Working [400102] [A400102]

Name	Epsilon Service Down
Description	Epsilon internal service is down
Alert message	Discovered that the epsilon internal service is not working: '*service_names*'
Cause	Epsilon internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-7080
KB Article	7080
KB Article	8622
Severity	Critical

Table 940: Applications Archive ready for download. [400104] [A400104]

Name	Applications Archive Ready
Description	Archive ready for download.

Alert message	Application logs from '*start_time*' to '*end_time*' are now available for download and archival. Please go to Calm > Settings to download them.
Cause	Archival process archived the run logs.
Impact	You will not be able to download it after a week.
Resolution	Make sure you download the archive.
Severity	Critical

Table 941: Calm showback is unable to reach beam service [400107] [A400107]

Name	Beam is unreachable
Description	Beam service is unreachable. Calm showback calculations will be affected
Alert message	Beam service is unreachable. Calm showback calculations will be affected
Cause	Prism central may not have internet connectivity or beam service is unavailable.
Impact	Impact on Calm Showback data.
Resolution	Review KB-6664 for further troubleshooting.
KB Article	6664
Severity	Critical

Table 942: AHV Prism Element attached [400108] [A400108]

Name	AHV Prism Element attached
Description	A new AHV Prism Element is attached.
Alert message	AHV Prism Element attached name: *pe_name*
Cause	A new AHV Prism Element is attached.
Impact	New Calm provider setting is added without showback price list.
Impact	Prism Element unusable in Calm if not associated with Calm project.
Resolution	(A) Associate the Prism Element in Calm project. (B) Define the price list for added Calm provider setting.
Severity	Info

Table 943: AHV Prism Element detached [400109] [A400109]

Name	AHV Prism Element detached
Description	One of the existing AHV Prism Element is detached.
Alert message	AHV Prism Element detached name: *pe_name*
Cause	One of the existing AHV Prism Element is detached.

Impact	Actions other than soft delete cannot be run on Calm applications that are provisioned on this Prism Element.
Impact	Calm provider settings corresponding to this Prism Element is deleted.
Resolution	In case the Prism Element is permanently detached run soft delete on the Calm applications that are provisioned on this Prism Element.
Severity	Info

Table 944: Calm Version Mismatch [400110] [A400110]

Name	Calm Version Mismatch
Description	Calm Version Mismatch
Alert message	All nodes are not running on same Calm version
Cause	Failed LCM upgrade
Impact	Calm will not work as expected
Resolution	Refer to KB 7528 for resolution
KB Article	7528
Severity	Critical

Table 945: Epsilon Version Mismatch [400111] [A400111]

Name	Epsilon Version Mismatch
Description	Epsilon Version Mismatch
Alert message	All nodes are not running on same Epsilon version
Cause	Failed LCM upgrade
Impact	Epsilon will not work as expected
Resolution	Refer to KB 7528 for resolution
KB Article	8694
Severity	Critical

Table 946: Calm Policy Engine Internal Service has Stopped Working [400114] [A400114]

Name	Calm Policy Engine Service Down
Description	Calm Policy Engine internal service is down
Alert message	Discovered that the Calm policy engine internal service running on *policy_engine_ip* is not working: '*service_names*'
Cause	Calm Policy Engine internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-9456

KB Article	9456
Severity	Critical

Table 947: Compute-only cluster sizing: hyperconverged node count check [405000] [A405000]

Name	Suggest min 4 HC nodes when using CO node(s)
Description	Check that cluster with compute-only nodes has at least 4 hyperconverged nodes.
Alert message	Cluster with compute-only nodes should have min of 4 hyperconverged nodes
Cause	Compute-only node was added to a cluster with fewer than 4 hyperconverged nodes.
Impact	Can lead to cluster not delivering sufficient storage IO.
Resolution	Remove all compute-only nodes from cluster or add additional hyperconverged nodes.
KB Article	7012
Severity	Warning

Table 948: Compute-only cluster sizing: compute-only to hyperconverged node ratio check [405001] [A405001]

Name	Ratio of CO to HC nodes to not exceed 1:2
Description	Ratio of compute-only to hyperconverged nodes to not exceed 1:2
Alert message	Cluster with compute-only nodes should not have the ratio of compute-only to hyperconverged nodes exceed 1:2
Cause	High ratio of compute-only to hyperconverged nodes.
Impact	Can lead to cluster not delivering sufficient storage IO.
Resolution	Remove all compute-only nodes from cluster or add additional hyperconverged nodes.
KB Article	7012
Severity	Warning

Table 949: Compute-only cluster sizing: CVM size check [405002] [A405002]

Name	Not large enough CVM (vcpus) for the CO node(s)
Description	Check that cluster with compute-only: Total CVM vCPUs must be greater than total cores on all compute-only nodes.
Alert message	Cluster with compute-only nodes should have total CVM vCPUs greater than total cores on all compute-only nodes
Cause	Compute-only node was added to a cluster with small CVMs.

Impact	Can lead to cluster not delivering sufficient storage IO.
Resolution	Remove all compute-only nodes from cluster or add additional hyperconverged nodes or resize CVMs
KB Article	7012
Severity	Warning

Table 950: Insufficient host bandwidth [405003] [A405003]

Name	HC network bandwidth is not enough to handle CO node(s)
Description	Checks whether the total hyperconverged network bandwidth is twice the compute-only bandwidth.
Alert message	For a cluster with compute-only nodes the total amount of hyperconverged host bandwidth must be double the amount of compute-only host bandwidth.
Cause	Cluster is configured with too many compute-only nodes.
Cause	Cluster is configured with less number of hyperconverged nodes.
Cause	Hyperconverged/compute-only host NIC is misconfigured.
Impact	Cluster will not be able to deliver sufficient storage I/O.
Resolution	Remove compute-only nodes from the cluster.
Resolution	Add hyperconverged nodes to the cluster.
Resolution	Reconfigure hyperconverged host NICs to increase the amount of bandwidth.
KB Article	7012
Severity	Warning

Storage

Table 951: Data Disk Space Usage High [1003] [A1005]

Name	Data Disk Usage
Description	Checks if any individual disk (SSD or HDD) storage space utilization is sustained for long above a high threshold.
Alert message	Disk space usage for *mount_path* on *entity* *ip_address* has exceeded the threshold of *disk_usage_threshold%*. Disks list: *disk_usage_pct_list*. *action_str*
Cause	There is more data than the underlying hardware can handle.
Cause	Compression has been disabled.
Cause	Aggressive DR policy settings.
Cause	Sustained High intensity IO.
Cause	Cassandra metadata ring imbalance.
Cause	Vdisk_block_map metadata threshold reached and vdisk severing not happening.

Cause	Individual disk storage is being highly utilized, possibly due to overall storage pool usage pressure.
Cause	Issues with disk/node balancing.
Impact	In isolation a spike in individual disk usage is not an immediate cause for concern but if sustained or seen across multiple disks it can indicate an underlying issue with data balancing and may lead to significant cluster performance degradation or issues with storage maintenance.
Resolution	Delete unnecessary data in the VM, or move VM out.
Resolution	Check if the Compression disabled alert was fired.
Resolution	Decrease data retention window for DR.
Resolution	Ensure Curator service is running and periodic scans are completing.
Resolution	Balance the Cassandra metadata ring.
Resolution	If height of snapshot_tree_printer chain is > 25, find the size of the vdisk block map across the cluster by running: allssh 'du -h ~/data/cassandra/data/medusa_vdiskblockmap/' (check for value of gflag --curator_max_vdisk_block_size_for_copy_blockmap_mb, it is default 40GB).
Resolution	Ensure storage pool usage is below critical threshold.
Resolution	In isolation a spike in individual disk usage is not an immediate cause for concern but if sustained or seen across multiple disks it can indicate an underlying issue with data balancing and may lead to significant cluster performance degradation or issues with storage maintenance.
Disk Usage Threshold	Warning Threshold: 90 % Critical Threshold: 95 %
KB Article	1523
Severity	Critical
Severity	Warning

Table 952: Disk Inode Usage High [1004] [A1027]

Name	Inode Usage
Description	Checks if current inode usage is high.
Alert message	Inode usage for one or more disks on Controller VM *svm_ip_address* has exceeded *inode_usage_threshold*%.
Cause	Inode usage is high.
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	Reduce disk usage or replace disk.
Inode Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	1532

Severity	Critical
Severity	Warning

Table 953: Disk Unused [1009] [A1063]

Name	Disk Configuration
Description	Checks if all disks in the cluster are assigned to a storage pool.
Alert message	Drive with disk id *disk_id* on Controller VM *service_vm_external_ip* is not part of any storage pool.
Cause	The cluster was not configured correctly during the installation or a drive replacement was not completed correctly.
Impact	The cluster was not configured correctly during the installation or a drive replacement was not completed correctly.
Resolution	Add the unused disk to the storage pool.
KB Article	1496
Severity	Info

Table 954: Hypervisor Disk usage is above the recommended threshold. [1015] [A1161]

Name	Host Disk Usage
Description	Checks that host disk usage is not above the recommended threshold%.
Alert message	Disk '*dev_name*' space usage on host *ip_address* has exceeded its threshold.
Cause	Too much data is stored on the host disk.
Impact	Cluster performance will be significantly degraded. In the case of multiple nodes with the same condition,the cluster may be unable to service I/O requests.
Resolution	Delete unneeded data from the host disk.
KB Article	2490
Severity	Critical
Severity	Warning

Table 955: Storage Pool Space Usage Exceeded the Configured Threshold [1016] [A1128]

Name	StoragePool Space Usage
Description	Checks the storage pool space usage to determine if it's above the threshold
Alert message	Storage Pool space usage for *storage_pool_name* is at *usage_pct*%. *alert_msg*
Cause	Space usage in the storage pool has exceeded the configured threshold.
Impact	Storage pool will run out of space, and the cluster will be unable to service I/O requests.
Resolution	Add more storage to the cluster or delete data to free up storage space.

Storage Pool Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	2475
Severity	Critical
Severity	Warning

Table 956: Storage Container Space Usage Exceeded [1017] [A1130]

Name	Storage Container Space Usage
Description	Check high space usage on Storage Containers
Alert message	Storage Container space usage for *container_name* is at *usage_pct*%.
Cause	Compression got disabled.
Cause	There is more data than the underlying hardware can handle.
Cause	Aggressive DR policy settings.
Cause	User deleted data but curator scan is slow.
Cause	Cassandra metadata ring imbalance.
Cause	Vdisk_block_map metadata threshold reached and vdisk severing not happening.
Impact	Storage Container will run out of space, and the cluster may become unable to service I/O requests.
Resolution	Check the compression disabled alert.
Resolution	Delete unnecessary data in the VM, or move VM out.
Resolution	Decrease data retention window for DR.
Resolution	Make scan faster using gflags, increase active req count from 10, or increase number of tasks generated by curator.
Resolution	Balance the Cassandra metadata ring
Resolution	If height of snapshot_tree_printer chain is > 25, find the size of the vdisk block map across the cluster by running: allssh 'du -h ~/data/cassandra/data/medusa_vdiskblockmap/' (check for value of gflag --curator_max_vdisk_block_size_for_copy_blockmap_mb, it is default 40GB).
Storage Container Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	2476
Severity	Critical
Severity	Warning

Table 957: System Root Partition Space Usage High [1021] [A1036]

Name	System Root Partition Usage
Description	Checks if CVM system root partition usage is within threshold to ensure uninterrupted operations. Does not refer to cluster shared storage.
Alert message	Disk space usage for *mount_path* on *entity* *ip_address* has exceeded *threshold*%.
Cause	Increased CVM system root partition usage due to excessive logging or incomplete maintenance operation.
Cause	Non-standard use of / (root) partition and sub-directories on the affected CVM local filesystem.
Impact	If CVM system root partition is highly utilised then certain maintenance operations, such as upgrades, may be impacted. If CVM system root partition is 100% utilised then services may stop and impact cluster storage availability.
Resolution	Remove unneeded/temporary/non-standard introduced files from affected CVM root partition.
Resolution	Please refer to KB 1523 for more information.
CVM System Root Partition Usage Threshold	Warning Threshold: 80 % Critical Threshold: 90 %
KB Article	1523
Severity	Critical
Severity	Warning

Table 958: Flash Mode Usage Limit Exceeded [1023] [A1159]

Name	Flash Mode Usage
Description	Check the amount of usage by flash-mode-enabled vDisks with respect to the threshold limit.
Alert message	Flash Mode usage has exceeded the threshold limit.
Cause	Flash-mode-enabled vDisks are using more flash tier space than the allotted limit.
Impact	System may down-migrate flash mode data.
Resolution	Reduce the number of flash-mode-enabled vDisks or increase the capacity of the flash tier.
KB Article	3671
Severity	Warning

Table 959: Flash-mode-enabled VM Power Status [1024] [A1160]

Name	Flash Mode Enabled VM Power Status
Description	Check if any flash-mode-enabled VMs are Powered Off.

Alert message	One or more Powered Off flash-mode-enabled VMs detected.
Cause	One or more flash-mode-enabled VMs are Powered Off.
Impact	Flash tier space might be getting wasted.
Resolution	Disable Flash Mode on the Powered Off VMs, or power them on.
KB Article	3671
Severity	Warning

Table 960: High Garbage due to erasure coding [1026] [A1174]

Name	Erasure Code Garbage
Description	Checks if erasure coding garbage is below a safe threshold.
Alert message	Garbage due to erasure coding on *container_name* is currently *usage* which is above the safety threshold.
Cause	Suboptimal performance of erasure coding is generating garbage that is consuming additional storage space.
Impact	Erasure coding on Nutanix clusters may experience suboptimal storage savings.
Resolution	Review KB 3021 for corrective actions. Contact Nutanix support to discuss possible fixes before making changes to erasure coding config.
KB Article	3021
Severity	Critical
Severity	Warning

Table 961: Storage Pool Flash Mode Configuration [1029] [A1173]

Name	Storage Pool Flash Mode Configuration
Description	Flash Mode is not supported when multiple storage pools are in use.
Alert message	Flash Mode configuration with Multiple Storage Pools detected.
Cause	Multiple Storage Pools and Flash Mode are in use together.
Impact	This configuration is not supported.
Resolution	Disable Flash Mode for all VMs and Volume Groups, or consolidate all Storage Pools into one.
KB Article	3671
Severity	Warning

Table 962: VSS Snapshot of Container Share Failed [20020] [A20020]

Name	VSS snapshot of container share failed

Description	A VSS initiated snapshot of the container share failed due to too many files in the container.
Alert message	A backup application requires a container id *container_id* to be backed up, however the VSS snapshot operation on it failed as there are *files_present* files on the container. Only a container with a maximum of *files_max* files can be snapshotted through VSS.
Cause	Too many files in the container.
Impact	Third party backup applications using VSS to backup VMs from a container share will fail.
Resolution	Reduce the number of files in each container to less than or equal to 1000.
Severity	Warning

Table 963: Erasure coding pending check [101054]

Name	Erasure coding pending check
Description	Check if EC undo is pending on any of the Storage Containers
Cause	EC egroups are still pending to be converted to non-EC egroups after undo EC on the container.
Impact	Node/Disk remove will be delayed until the EC egroups are undone to non-EC egroups are completed.
Resolution	Changes to undo erasure coding can take many hours to several days to complete. No immediate action is required. Reach out to Nutanix Support if this check is not passing for two (2) or more days.
KB Article	3021

Table 964: Volume Group Space Usage Exceeded [101056] [A101056]

Name	Volume Group Space Usage
Description	Checks for high space usage on volume groups.
Alert message	Volume Group space usage for *volume_group_name* is at *usage_pct*%.
Cause	Excessive space usage in the volume group.
Impact	Volume Group will run out of space, and the cluster may become unable to service I/O requests to that volume group.
Resolution	Add more capacity to the volume group or delete some data to free storage space. The Guest OS and/or the application using the volume group should be configured to send the list of freed blocks to AOS using SCSI UNMAP command. This can be done as part of in-band or out of band TRIM operation(s). Refer to 'Thin Provisioning' section in the Nutanix Volumes - Nutanix Best Practices guide.
Volume Group Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	4600

Severity	Critical
Severity	Warning

Table 965: System Non-Root Partitions Space Usage High [101061] [A1031]

Name	System Non-Root Partition Usage
Description	Checks if CVM system non-root partition usage is within threshold to ensure uninterrupted operations. Does not refer to cluster shared storage.
Alert message	Disk space usage for *mount_path* on *entity* *ip_address* has exceeded *threshold%*. *ref_msg*
Cause	Increased CVM system non-root partition usage due to excessive logging, storage of large log bundles or unused 1-click binaries.
Cause	Non-standard use of /home directories on the affected CVM.
Impact	If CVM system non-root partitions are highly utilised then certain maintenance operations, such as upgrades, may be be impacted. If CVM system non-root partitions are 100% utilised then services may stop and impact cluster storage availability.
Resolution	Remove unneeded 1-click bundles via Prism, log collector bundles and temporary/non-standard introduced large files from affected CVM partition.
Resolution	Please refer to KB 1540 for more information. If unsure, please contact Nutanix Support.
CVM System Non-Root Partition Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	1523
Severity	Critical
Severity	Warning

Table 966: /scratch location Space Usage is High [101068] [A101068]

Name	Scratch Space Usage
Description	Checks that scratch space usage is not high.
Alert message	/scratch space usage (*usage%*) on host *ip_address* has exceeded *threshold%*.
Cause	Too much data is stored in scratch space.
Impact	If scratch space is not available, temporary data is stored on a RAM disk with limited space, causing issues during upgrade
Resolution	Delete unneeded data or add nodes to the cluster.
Scratch space Usage Threshold	Warning Threshold: 90 % Critical Threshold: 95 %
KB Article	2490

Severity	Critical
Severity	Warning

Table 967: Invalid hybrid nodes configuration for erasure coding [101070] [A101070]

Name	Erasure coding all-flash node intermixed check
Description	Checks for valid configuration of coexisting all-flash and hybrid nodes in cluster when EC is enabled.
Alert message	Number of hybrid nodes found are below the recommended threshold of *threshold*.
Cause	Current hybrid nodes configuration does not satisfy EC storage container requirements.
Impact	Unsupported configuration. The cluster may experience performance degradation.
Resolution	(A) Add more storage heavy or hybrid nodes to the cluster. (B) Consult Nutanix documentation on Erasure Coding requirements and follow KB3021 recommendations for any required configuration changes.
KB Article	3021
Severity	Warning

Table 968: System Audit Volume Space Usage High [101072] [A1035]

Name	System audit volume usage
Description	Checks if CVM system audit volume usage is within threshold to ensure uninterrupted operations.
Alert message	Disk space usage for *mount_path* on *entity* *ip_address* has exceeded *threshold*%.
Cause	Increased CVM system audit volume usage due to excessive logging, or miscellaneous files being placed inside.
Impact	If CVM system audit volume is 100% utilized, then auditd will stop and further system auditing will not be recorded.
Resolution	Remove files other than audit.log* from /home/log/audit. If unsure, please engage Nutanix Support for assistance.
CVM System Audit Volume Usage Threshold	Warning Threshold: 75 % Critical Threshold: 99 %
KB Article	1523
Severity	Critical
Severity	Warning

Table 969: SED Disk Configuration Misconfigured [101080] [A101080]

Name	Disk configuration v/s intent consistency check
Description	Check if SED disk configuration is consistent with the intent
Alert message	Either the cluster had SED protection enabled and the disks did not have a password set OR the cluster did not have SED protection enabled and disks had password set OR found tentative passwords.
Cause	Inconsistency in disk's zeus configuration and the intent
Impact	Data unavailability requiring SRE intervention.
Resolution	This issue is usually seen when the KMS is switched off when the disks are still in protected state. Please check whether the KMS is on or off. If the issue persists, contact Nutanix Support.
KB Article	11322
Severity	Critical
Severity	Warning

Table 970: SAS Connectivity Not Normal [106019] [A106019]

Name	SAS Connectivity
Description	Check SAS connectivity
Alert message	SAS cables missing or not properly connected on host.
Cause	SAS cables are not properly connected.
Impact	Cluster performance may be significantly degraded.
Resolution	Ensure that SAS cables are properly connected.
KB Article	2311
Severity	Warning

Table 971: Storage Container Replication Factor Low [110210] [A1074]

Name	Storage Container RF Status
Description	Check that the Storage Container region replication factor is not low
Alert message	Replication factor of *container_region* on Storage Container *container_name* is set to 1.
Cause	Replication factor on the Storage Container was changed to 1.
Impact	Only one copy of guest VM data exists on the cluster. If any node fails, the data stored on that node is unavailable. If any drive fails, the data stored on that drive is unrecoverable.
Resolution	Change the replication factor on the Storage Container with the nCLI command "container edit name='ctr_name' rf='2'" (or higher if appropriate for your environment).
KB Article	2487

Severity	Critical

Table 972: RF1 Container Check [111085]

Name	RF1 Container Check
Description	List all RF1 containers and active VMs in the cluster
Cause	Found RF1 containers and VMs in the cluster
Impact	RF1 is enabled in the cluster. VMs that have vDisks or VGs on RF1 containers have single copy of data in the cluster. Those VMs will experience data un-availability if the affinity host experiences failure.
Resolution	Not Applicable
KB Article	10663

Table 973: Fingerprinting Disabled [130021] [A1068]

Name	Fingerprinting On Write
Description	Fingerprinting Disabled
Alert message	Disabling fingerprinting (deduplication) for future writes. Calculated metadata usage of *metadata_usage* bytes exceeds safety limit of *metadata_limit* bytes on Controller VM *service_vm_id*. The metadata disk size is *metadata_disk_size*.
Cause	The metadata disk is used for metadata and data. Once the metadata portion on the disk exceeds the safety threshold AOS disables fingerprinting (deduplication) of future writes. Metadata sizes can increase due to a variety of reasons such as having too many snapshots, or simply because too much data is fingerprinted.
Impact	The advantages of deduplication are not available.
Resolution	Once the metadata size is brought under the safety limit (for example, by removing unneeded snapshots, or by expanding the cluster size), the fingerprinting will be auto-enabled.
Severity	Warning

Table 974: Metadata Usage High [130025] [A1101]

Name	Metadata Usage
Description	Metadata Usage High.
Alert message	Metadata usage on Controller VM *service_vm_external_ip* has exceeded *critical_pct*%.
Cause	The cluster either has too many snapshots or too much data is being fingerprinted.
Impact	Cluster performance may be significantly degraded.
Resolution	Reduce metadata size by removing unneeded snapshots, creating snapshots less frequently, creating snapshots of fewer VMs, or expanding the cluster size.

Severity	Warning

Table 975: Non Self Encryption Drive Disk Inserted [130031] [A1122]

Name	Non SED Disk Inserted Check
Description	Non-self-encrypting drive is inserted.
Alert message	Non encrypted disk with serial *disk_serial* was added in drive bay *disk_location* on Controller VM *service_vm_external_ip*. It is not usable because the rest of the cluster is protected using encrypted drives.
Cause	A cluster with self-encrypting drives has a non-self encrypting drive installed.
Impact	The new drive was not mounted because the cluster cannot be fully secured without all drives supporting encryption.
Resolution	Ensure all installed drives are of self-encrypting type. If more help is needed call Nutanix support.
KB Article	8417
Severity	Critical

Table 976: On-Disk Dedup Disabled [130032] [A1089]

Name	On-Disk Dedup Status
Description	On-disk deduplication is disabled.
Alert message	Disabling further on-disk deduplication (fingerprinting). Controller VM with *ram_size*GB RAM and *ssd_size*GB SSD does not meet the minimum requirements of *required_ram_size*GB RAM and *required_ssd_size*GB SSD to support the on-disk deduplication feature.
Cause	To support on-disk deduplication (fingerprinting) feature, the Controller VM must meet minimum requirements for RAM and SSD. If this threshold is not satisfied, on-disk deduplication is disabled.
Impact	The advantages of on-disk deduplication are not available.
Resolution	Once the Controller VMs are upgraded to meet the minimum requirements on RAM and SSD, on-disk deduplication can be enabled. Contact Nutanix support for assistance.
Severity	Warning

Table 977: Disk Failed/Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Drive is marked offline.
Alert message	Disk *disk_position* with id *disk_id* on node *node_position* of block *block_position* has been marked offline due to IO errors. Serial number of the disk is *disk_serial* in host *host_ip* of block *block_serial*.
Cause	A drive on the node has failed and was marked offline.

Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Replace the failed drive. Refer to the Nutanix documentation for instructions.
KB Article	8453
Severity	Critical

Table 978: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	Drive *disk_id* with serial *disk_serial* and model *disk_model* in drive bay *disk_location* on Controller VM *service_vm_external_ip* has failed. Node serial is *node_serial_number* and Node position *node_position*. #rf1_custom_message@
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 979: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical drive is removed from the slot.
Alert message	Drive with serial *disk_serial* and model *disk_model* was removed from drive bay *disk_location* on Controller VM *service_vm_external_ip*. Node serial is *node_serial_number* and Node position *node_position*. #rf1_custom_message@
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the drive will start.
Resolution	Remove the drive from the slot and reseat it. Contact the Hardware Vendor if issue persists.
KB Article	8031
Severity	Critical

Table 980: Space Reservation Violated [130053] [A1021]

Name	Space Reservation Status
Description	Space reservation violated.
Alert message	Space reservation configured on vdisk *vdisk_name* belonging to Storage Container id *container_id* could not be honored due to insufficient disk space resulting from a possible disk or node failure.
Cause	A drive or a node has failed.

Impact	The space reservations on the cluster can no longer be met. Writes by guest VMs may fail if expected storage space is not available.
Resolution	Change space reservations to total less than 90% of the available storage, and replace the drive or node as soon as possible. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 981: Volume Group Action Error [130071] [A1163]

Name	Volume Group Action Status
Description	Volume group action error.
Alert message	Failed to *action* volume group with name '*vg_name*' and internal ID '*vg_uuid*' because *reason* in protection domain '*protection_domain_name*'
Cause	A volume group could not be restored or could not be deleted because it is still in use.
Impact	The requested volume group action (restore or delete) could not be completed.
Resolution	Detach the volume group from VMs and external initiators before recovery. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10120
Severity	Critical

Table 982: Unable To Remount Datastore [130076] [A130076]

Name	Datastore Remount Status
Description	Unable to remount a datastore.
Alert message	Unable to *operation* Datastore '*datastore_name*'*host_info*.
Cause	Datastore information unavailable
Impact	VM HA and DRS functionality might be impaired.
Resolution	Check container exists on ESXi, rescan datastore and mount the container on the hosts from Prism Element. If issue persists contact Nutanix Support
Severity	Warning

Table 983: Remounted Datastore [130080] [A130080]

Name	Datastore Remount Success
Description	Remounted datastore.
Alert message	Remounted Datastore '*datastore_name*' as metro availability is being enabled.
Cause	Datastore has been remounted

Impact	VM HA and DRS functionality might be impaired.
Resolution	As the datastore has been mounted successfully rescan the datastore in vCenter in order to complete the changes necessary to support the enhanced planned failover functionality
Severity	Info

Table 984: Protected VM(s) Not Found [130083] [A130083]

Name	Protected VMs Not Found
Description	Protected VM(s) Not Found.
Alert message	Unable to locate VM(s) *vm_names* protected by protection domain '*protection_domain_name*'.
Cause	Protected VM(s) cannot be found and may have been deleted.
Impact	Any data associated with the VMs may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Verify that the VM(s) are registered on the host. If the VM(s) were unregistered or deleted, remove VM(s) from the protection domain. If the CVM on the host the VM(s) are registered on was shutdown or rebooting, retry the snapshot operation.
Severity	Warning

Table 985: Protected Volume Groups Not Found [130084] [A130084]

Name	Protected Volume Groups Not Found
Description	Protected volume groups not found.
Alert message	Unable to locate volume group(s) *vg_names* protected by protection domain '*protection_domain_name*'.
Cause	Protected volume group(s) cannot be found and may have been deleted.
Impact	Any data associated with the volume group(s) may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Remove volume group(s) from the protection domain.
Severity	Warning

Table 986: VStore Snapshot Status [130086] [A130086]

Name	VStore Snapshot Status
Description	vStore snapshot status.

Alert message	Snapshot status for vstore *vstore_name*: *reason*.
Cause	Various
Impact	The requested VStore snapshot action could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
Severity	Warning

Table 987: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	Drive *disk_id* with serial *disk_serial* in drive bay *disk_location* on Controller VM *service_vm_external_ip* has failed diagnostic test. This drive is failing or has failed. #rf1_custom_message@
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.
Severity	Warning

Table 988: Associated entities are not protected together. [130090] [A130090]

Name	Related Entity Protection Status
Description	Protection status of a related entity.
Alert message	*error_message*.
Cause	Related entity is not protected in the same protection domain.
Impact	Related VM/Volume Group will not be snapshotted and recovered.
Resolution	Protect the related entity in the same consistency group.
Severity	Warning

Table 989: iSCSI Configuration Failed [130100] [A130100]

Name	iSCSI Configuration Failed
Description	iSCSI Configuration Failed
Alert message	Failed to re-configure iSCSI settings on the recovered VM '*vm_name*'. *reason*.
Cause	Nutanix Guest Tools failed to execute some iSCSI commands on the guest VM.
Impact	iSCSI disks may become unavailable on the guest VM.

Resolution	(A) If IQN and iSCSI target IP addresses of the VM have been updated by Nutanix Guest Agent, Discover and connect to new targets after rebooting the VM. (B) Manually configure iSCSI settings on the guest VM. (C) Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	10118
Severity	Critical

Table 990: Recovered VM Disk Configuration Update Failed [130107] [A130107]

Name	Disk Configuration Update Failed
Description	Disk Configuration Update Failed
Alert message	Failed to make some disks online on the recovered VM '*vm_name*'. *reason*.
Cause	Nutanix Guest Tools failed to automatically bring the disks online.
Impact	Disks may become offline on the recovered VM.
Resolution	Manually bring the disks online on the recovered VM. Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.
KB Article	8208
Severity	Critical

Table 991: System-Defined Flash Mode Usage Limit Exceeded [130120] [A130120]

Name	System Defined Flash Mode Usage Limit
Description	Check that usage for flash-mode-enabled vDisks is within system limits.
Alert message	System has down-migrated the data of flash-mode-enabled vDisks.
Cause	Too many vDisks are assigned to the flash tier, or the vDisks assigned to the flash tier are too large.
Impact	The advantages of flash mode may not be available.
Resolution	Reduce the number of flash-mode-enabled vDisks or increase the capacity of the flash tier.
Severity	Warning

Table 992: NFS metadata usage high [130122] [A130122]

Name	NFS Metadata Size Overshoot
Description	NFS Metadata Usage High
Alert message	NFS metadata usage is too high. Calculated metadata usage of *metadata_usage* bytes exceeds safety limit of *metadata_limit* bytes on Controller VM *service_vm_id*.
Cause	Either a NFS datastore is used as a file store which is not supported, or too many files are present in the datastore.

Impact	Cluster performance and stability may be significantly degraded.
Resolution	If a NFS datastore is used as a file store, remove the corresponding files and wait for up to 24 hours to allow for compaction.
KB Article	6012
Severity	Warning

Table 993: Guest Power Operation through NGT Failed [130149] [A130149]

Name	VM Guest Power Op Failed
Description	Guest Power Operation Failed
Alert message	Failed to perform operation '*operation*'on VM '*vm_name*'. *reason*.
Cause	Power operation failed in guest operating system.
Impact	Desired power operation may not be completed.
Resolution	(A) Manually shutdown the VM by logging into the VM and running the appropriate shutdown command. (B) Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.
KB Article	10242
Severity	Warning

Table 994: Protected VM(s) May Not Be Recoverable [130183] [A130183]

Name	Protected VMs May Not Be Recoverable
Description	Protected VM(s) May Not Be Recoverable
Alert message	VM(s) *vm_names* protected by protection domain '*protection_domain_name*' are running from VMWare snapshots. Recovery of the VMs from the PD snapshot '*snapshot_id*' may not succeed. Manual intervention is required to boot up the VMs.
Cause	Protected VM(s) are running from VMWare snapshots and may not be recoverable.
Impact	Recovery of the VMs from the PD snapshot may not succeed
Resolution	Ensure that the VMs are not running from VMWare snapshot at the time of scheduled Protection Domain snapshot. Refer Nutanix KB: 7023 for further details.
Severity	Warning

Table 995: External Repository Access Failure [130208] [A130208]

Name	External Repository Access Failure
Description	Failed to access an NFS v3 external repository.
Alert message	Failed to access external repository *external_repository_name*. Failure reason: *message*.

Cause	Network connectivity issues between the AOS CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	The NFS server may not have white-listed access to the CVM subnet.
Cause	The NFS server may not support the v3 protocol version.
Cause	The repository is no longer exported by the NFS server.
Cause	The NFS server may not allow clients accessing it using non-privileged TCP ports (port number higher than 1023).
Cause	The NFSv3 external repository does not have the correct access permission.
Cause	I/O error on the NFS server while mounting the exported directory.
Impact	Mounting an NFS v3 external repository fails
Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure the CVM subnet is (all AOS CVMs have) white-listed for the export on the NFS server.
Resolution	Make sure that the NFS server support NFS v3.
Resolution	Make sure the repository is exported by the NFS server.
Resolution	Make sure the NFS server allows clients accessing it using non-privileged TCP ports.
Resolution	Make sure that the exported directory of the NFS external repository is mountable.
Resolution	Make sure the NFS server is in a good state.
Severity	Warning

Table 996: IO failures to a data source in an external repository [130209] [A130209]

Name	IO failures to data sources in external repositories
Description	IO to an external data source within an external repository failed.
Alert message	IO to the external data source with path *relative_uri* in the external repository *external_repository_name* failed. Failure reason: *message*.
Cause	I/O to the external datasource failed with I/O error from the NFS server.
Cause	Network connectivity issues between the AOS CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	Write failed because of no disk space available or hitting quota limitation in the NFS external repository.
Cause	Write request is attempted while the external repository is configured in READ_ONLY mode.
Cause	The datasource within the NFSv3 external repository does not have the correct access permission.
Cause	The datasource is not available in the NFS server.
Impact	IO to an data source within an external repository fails
Resolution	Make sure the NFS server is in a good state.

Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure that NFS server has enough space/quota for the write to this datasource.
Resolution	Change external repository access mode from READ_ONLY to READ_WRITE.
Resolution	Make sure that the datasource is available for the CVM subnet and allows requests such as read and/or write.
Resolution	Make sure the datasource is available or recreate the VM disk with the correct datasource configuration.
KB Article	8362
Severity	Warning

Table 997: External datasource storage migration suspended due to unrecoverable errors [130347] [A130347]

Name	External datasource storage migration suspended due to unrecoverable errors
Description	Storage migration from an external datasource within an external repository suspended due to unrecoverable errors.
Alert message	Storage migration from the external datasource *datasource_uuid* with path *relative_uri* in the external repository *external_repository_name* suspended. Suspension reason: *message*.
Cause	Insufficient free storage space.
Cause	Reads to the external datasource failed with I/O error from the NFS server.
Cause	Network connectivity issues between the CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	The datasource within the NFSv3 external repository does not have the correct access permission.
Cause	The datasource is not available in the NFS server.
Impact	Storage migration from a datasource within an external repository is suspended
Resolution	Make sure the destination container has enough free capacity.
Resolution	Make sure the NFS server is in a good state.
Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure that the datasource is available for the CVM subnet and allows requests such as read and/or write.
Resolution	Make sure the datasource is available or recreate the VM disk with the correct datasource configuration.
KB Article	10864
Severity	Warning

System Indicator

Table 998: Tomcat is restarting frequently. [700101] [A700101]

Name	Tomcat Frequent Restart
Description	Tomcat is restarting frequently.
Alert message	Tomcat is restarting frequently on *cvm_ip*. Reason: *message*.
Cause	Tomcat is restarting frequently, please check prism, prism_monitor and catalina logs for exact cause of failure.
Impact	Performance and availability impact for Prism
Resolution	Check the logs for error information or contact Nutanix support.
KB Article	8524
Severity	Critical

VM

Table 999: Non Compliance with Host Affinity policies. [200403] [A200403]

Name	VM not in compliance with the defined affinity policies.
Description	VM not in compliance with the defined affinity policies.
Alert message	VM *vm_uuid* in cluster *cluster* not in compliance with policies *policy_list_str*
Cause	A new policy was added, an existing policy was updated, or a category was attached to or detached from the VM or host.
Impact	VM might not run on hosts it is expected to run on.
Resolution	Make appropriate updates to affinity policies OR VM or Host categories.
Severity	Warning

13

VIEW TASK STATUS

The web console displays detailed information about all tasks that have been performed on the cluster.

Task Page Navigation

- To view the Task dashboard, select **Home** > **Tasks**.
- An icon also appears in the main menu when one or more tasks are active (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 active tasks; clicking the **View All Tasks** button at the bottom of that list displays a details screen with information about all tasks for this cluster.

  > Note: The drop-down list of active tasks may include a **Clean Up** button (top right). Clicking this button removes from the list any tasks that are no longer running. However, this applies to the current session only. The full active list (including the non-running tasks) appears when you open a new Prism web console session.

- When multiple tasks are active, you can filter the list by entering a name in the filter by field.
- You can also filter the list by clicking the **Filters** button and selecting the desired filter options

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.

View Task Status Dashboard

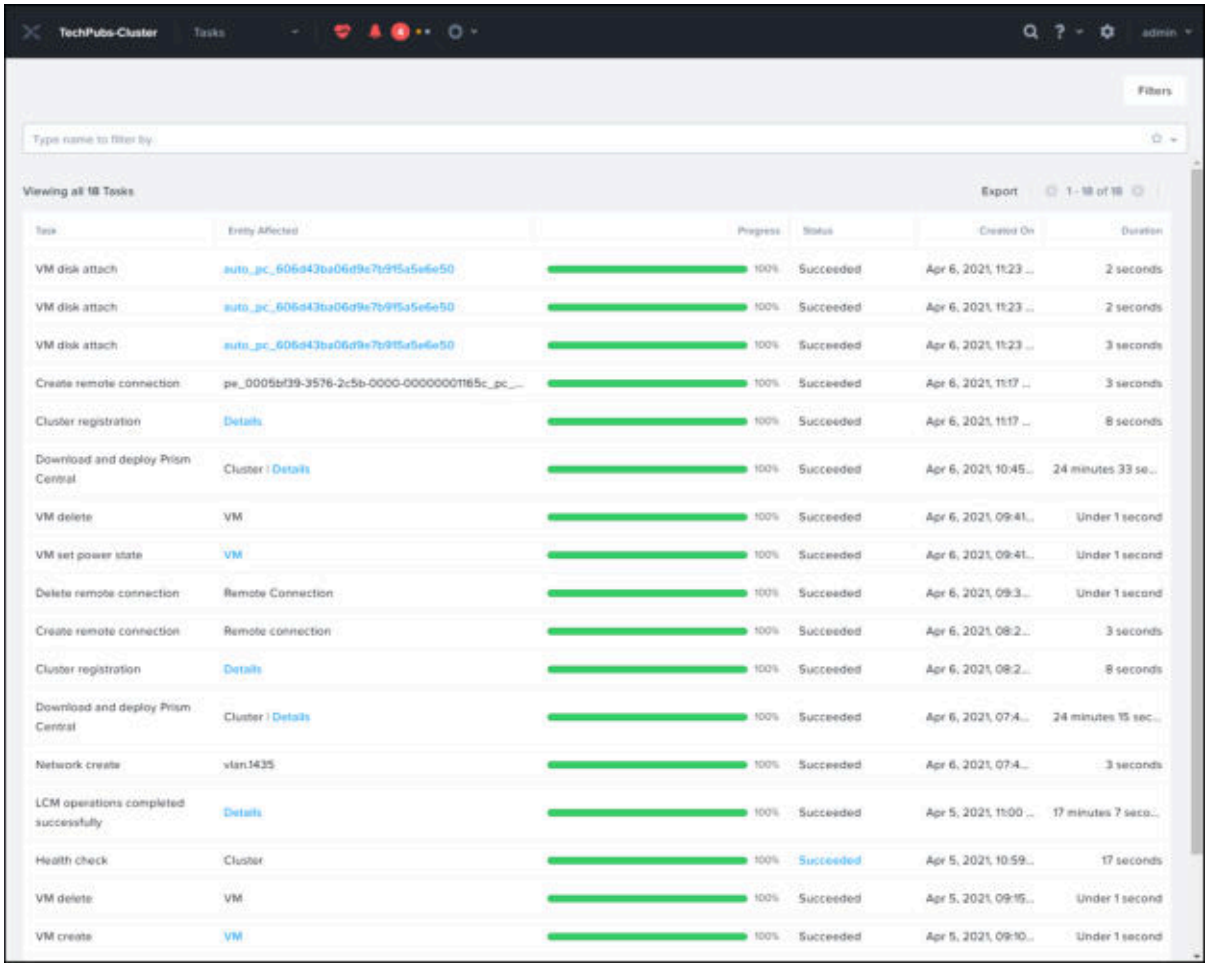

Figure 266: Task Dashboard

Table 1000: Tasks List Fields

Parameter	Description	Values
Task	Specifies which type of operation the task is performing.	Any cluster operation you can perform in the web console
Entity Affected	Display the entity on which task has been performed. If the link appears on the entity, click it to display the details.	Entity description
Percent	Indicates the current percentage complete for the task.	0%-100%
Status	Indicates the task status, which can be pending, running, completed, or failed.	pending, running,completed,failed
Created On	Displays when the task began.	seconds, minutes, hours

Parameter	Description	Values
Duration	Displays how long the task took to complete.	seconds, minutes, hours

14

SYSTEM MANAGEMENT

The web console allow you to configure various system settings.

- You can specify one or more name servers (see Configuring Name Servers on page 751)
- If Acropolis is enabled, you can configure one or more network connections (see Network Configuration for VM Interfaces on page 194).
- You can create a whitelist of IP addresses that are allowed access (see Configuring a Filesystem Whitelist on page 750).
- You can specify one or more NTP servers for setting the system clock (see Configuring NTP Servers on page 753).
- You can configure one or more network switches for statistics collection (see Configuring Network Switch Information on page 202).
- You can specify an SMTP mail server (see Configuring an SMTP Server on page 754).
- You can configure SNMP (see Configuring SNMP on page 755).
- You can configure a login banner page (see Configuring a Banner Page on page 769).

Configuring a Filesystem Whitelist

About this task

A whitelist is a set of addresses that are allowed access to the cluster. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

> CAUTION: Using a Nutanix storage container as a general-purpose NFS or SMB share is not supported. Because the Nutanix solution is VM-centric, the preferred mechanism is to deploy a VM that provides file share services.

To add (or delete) an address to (from) the filesystem whitelist, do the following:

Procedure

1. Click the gear icon in the main menu and then select Filesystem Whitelists in the **Settings** page.
   The Filesystem Whitelists dialog box appears.

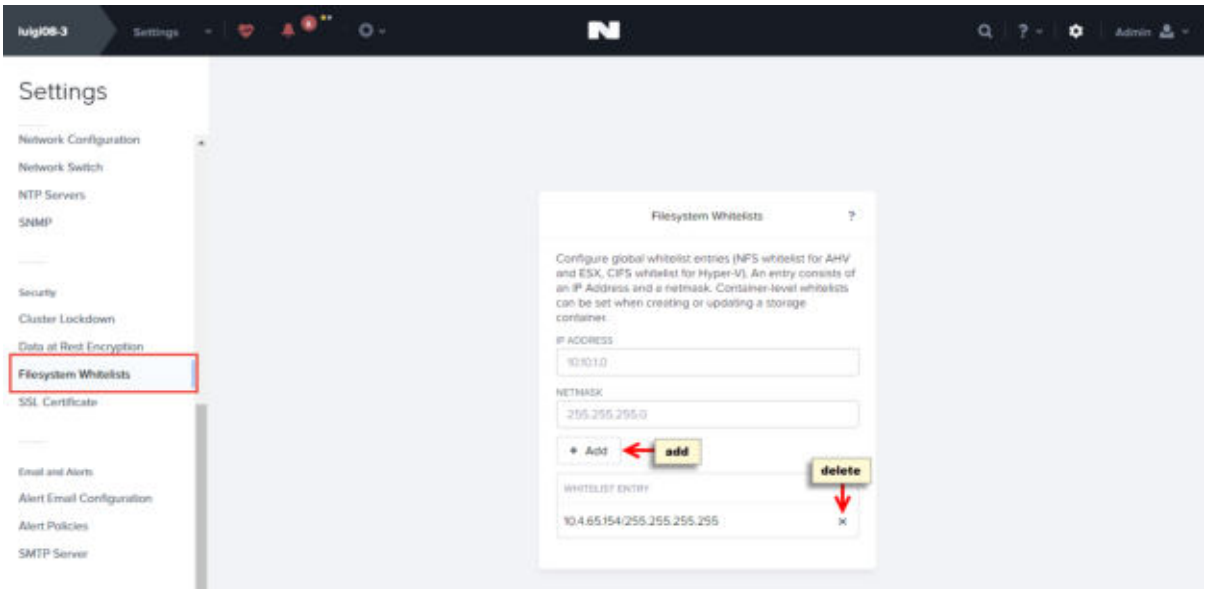

Figure 267: Filesystem Whitelists Window

2. To add an address to the whitelist, do the following in the indicated fields:

   a. **IP address**: Enter the IP address.

   b. **Netmask**: Enter the netmask value.

   c. Click the **+Add** button.

   The entry is added to the **Whitelist Entry** list. An NFS whitelist is created when the hypervisor is ESXi or AHV; a CIFS whitelist is created when the hypervisor is Hyper-V.

3. To delete an entry from the whitelist, click the X icon for that entry in the **Whitelist Entry** list.

   A window prompt appears to verify the action; click the **OK** button. The entry is removed from the 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.

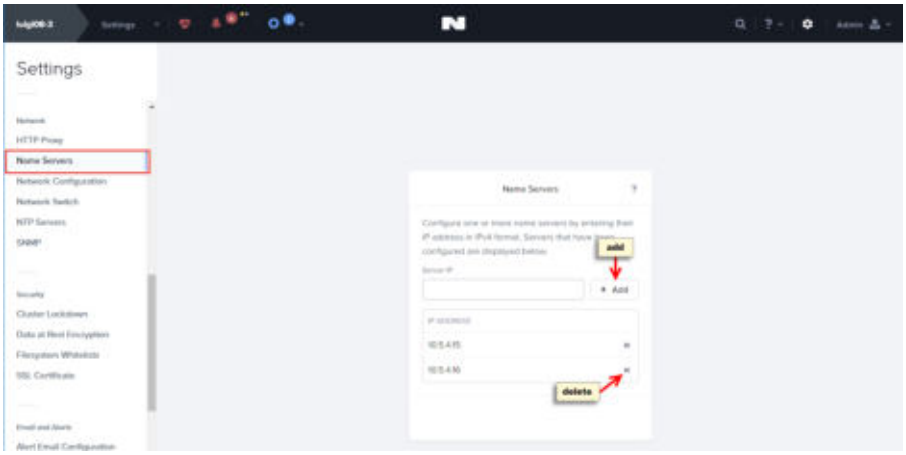

Figure 268: Name Servers Window

2. To add a name server, enter the server IP address in the **Server IP** field and then click the **Add** button to the right of that field.
   The server is added to the **IP Address** list (below the **Server** field).

   Note: Changes in name server configuration may take up to 5 minutes to take effect. Functions that rely on DNS may not work properly during this time. You can configure a maximum of three name servers.

3. To delete a name server entry, click the X icon for that server in the **IP Address** list.
   A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

Cluster Time Synchronization

Network Time Protocol (NTP) is a protocol for clock synchronization between computers. The hosts and CVMs in a Nutanix cluster must be configured to synchronize their system clocks with a list of stable NTP servers. Accurate timestamps are important for troubleshooting interactions with third-party software products such as Veeam or CommVault, which might require time synchronization between the hypervisor and the Controller VM to determine which files to back up. Accurate time synchronization between Nutanix clusters paired in Disaster Recovery (DR) configurations also ensures that snapshots do not expire too quickly or too late.

Graphs in the Prism interface rely on CVM time, and incorrect time skews graphs, especially in relation to other monitoring platforms such as vCenter, which rely on other clock sources.

Recommendations for Time Synchronization

Adhere to the following guidelines when configuring time synchronization on a Nutanix cluster:

- Where possible, synchronize Nutanix clusters with internal NTP sources to ensure stability from both a network and a security vulnerability perspective. When you cannot avoid using an external NTP source, prefer a time source maintained by the government.

- Make sure to specify at least five stable time sources that have a high degree of accuracy and that can be reached over a reliable network connection. Generally, the lower the stratum of an NTP source, the higher its accuracy.

  Note that three is the minimum to identify one time source as a false ticker but provides no redundancy, four is the minimum for redundancy, and five is the recommended minimum for a good configuration.

  If you want to use off-site NTP servers, see http://support.ntp.org/bin/view/Support/SelectingOffsiteNTPServers for various recommendations. Review all the information in this section before choosing pool.ntp.org servers for your NTP use.

- Do not use rate-limited NTP servers.

- Synchronizing a Nutanix cluster with a Windows time source is known to cause issues over a period of time, so Nutanix recommends that you not synchronize a cluster's time with Windows NTP sources. Use reliable non-Windows time sources instead. In an Active Directory domain, the best practice (a design that both works around and improves upon having to include domain controllers in the list of NTP sources) is to bypass the domain controllers and to synchronize the Nutanix hosts and CVMs directly with the NTP sources with which the domain controllers synchronize their time. Specify a common list of at least five reliable non-Windows NTP sources for both the domain controllers and the Nutanix cluster.

  Bypassing the domain controller as a time source is not an option for Hyper-V clusters owing to Kerberos requirements. When being joined to a domain, Nutanix clusters running Hyper-V detect local domain controllers and add them to all CVMs as NTP sources. For Hyper-V clusters, supplement the list of detected domain controllers with as many reliable non-Windows NTP sources as are required to meet the recommendation of a minimum of five NTP time sources. For example, if the Nutanix cluster adds two domain controllers as time sources, specify at least three reliable non-Windows NTP sources in the NTP server list. Specifying additional non-Windows NTP sources is necessary even if the domain controllers synchronize their time with a time source that is considered to be reliable.

- Specify public NTP servers by using their FQDN to help mitigate issues caused by IP address changes.

Configuring NTP Servers

About this task

To add (or delete) an NTP server entry, do the following:

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.

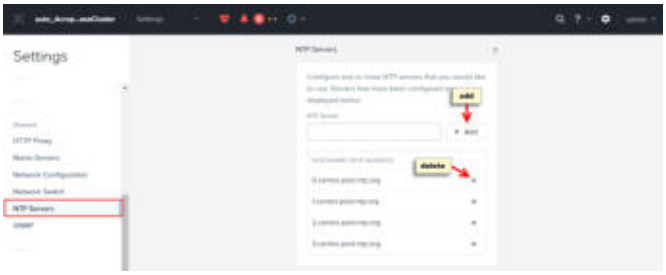

Figure 269: NTP Servers Dialog Box

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 **HOST NAME OR IP ADRESS** list (below the **NTP Server** field).

3. To delete an NTP server entry, click the cross 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.

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 Nutanix systems use SMTP to send alert emails.

If you have an external firewall deployed in your organization, see “” () to allow traffic from the cluster to Nutanix Support servers.

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.

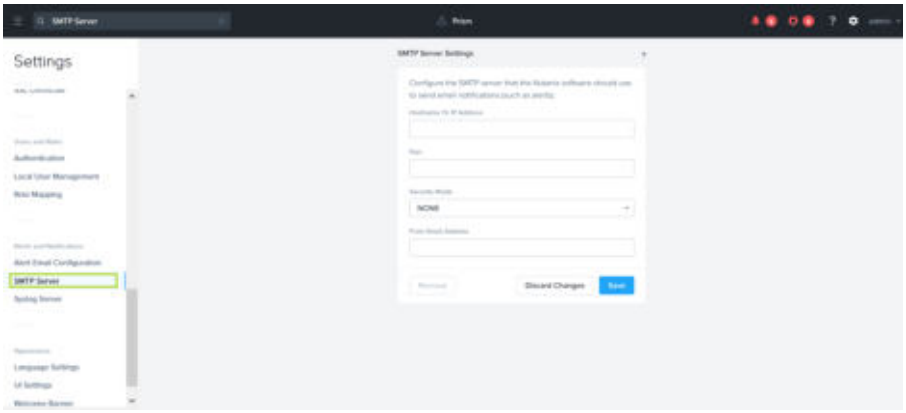

Figure 270: SMTP Server Settings Window

2. Do the following in the indicated fields:

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

   b. **Port**: Enter the port number to use.

      The standard SMTP ports are 25 (unencrypted), 587 (TLS), and 465 (SSL). For the complete list of required ports, see Port Reference.

   c. **Security Mode**: Enter the desired security mode from the pull-down list.

      The options are NONE (unencrypted), STARTTLS (use TLS encryption), and SSL (use SSL encryption).

   d. **User**: Enter a user name.

      The **User** and **Password** fields apply only when a secure option (STARTTLS or SSL) is selected. The user name might need to include the domain 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 (see Configuring Alert Emails on page 420) and cluster status information (see Configuring Pulse on page 798) 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. Nutanix systems include an SNMP agent that provides interoperability with industry standard SNMP manager systems. Nutanix also provides a custom Management Information Base (MIB) for Nutanix-specific information.

> Note: Nutanix recommends that you configure your SNMP manager to use the virtual IP address of the cluster to communicate with the Nutanix SNMP agent instead of the Controller VM IP address.

To configure SNMP, do the following:

Procedure

1. Click the gear icon in the main menu and then select **SNMP** in the **Settings** page.
   The SNMP Configuration dialog box appears.

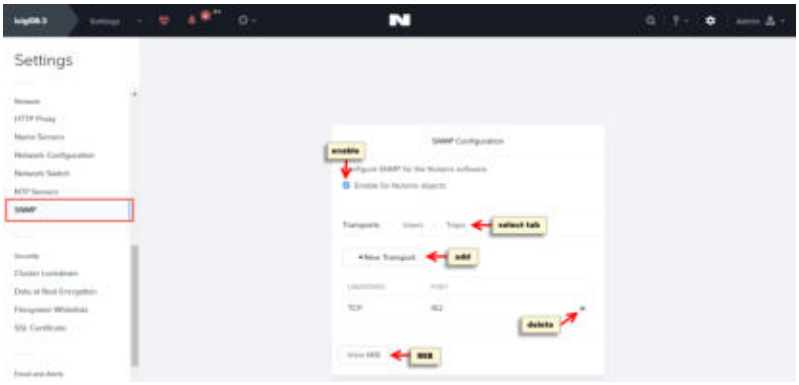

Figure 271: SNMP Configuration Window

2. To enable SNMP for this cluster, check the **Enable SNMP** box. To disable SNMP, uncheck the box.

   Note:

   SNMP traps are sent by the Controller VM that functions as the Alert Manager leader. If you need to open your firewall to receive the traps, keep in mind that the Alert Manager leader can rotate during tasks like AOS or host upgrades. Therefore, it might be necessary to open all the Controller VM IP addresses to ensure that the traps are received.

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 and then copy NUTANIX-MIB.txt to your SNMP manager systems.

   See your SNMP manager documentation for instructions on how to install the Nutanix MIB.

4. To add an SNMP transport, click the **Transports** tab and the **New Transport** button, and then do the following in the indicated fields. An SNMP transport is a combination of the transport protocol and port number on which you want the Nutanix SNMP agent to receive queries. SNMP transports enable you to combine transport protocols and port numbers

other than the default port number. The port numbers that are specified in SNMP transports are unblocked on the Controller VM, making them available to receive queries:

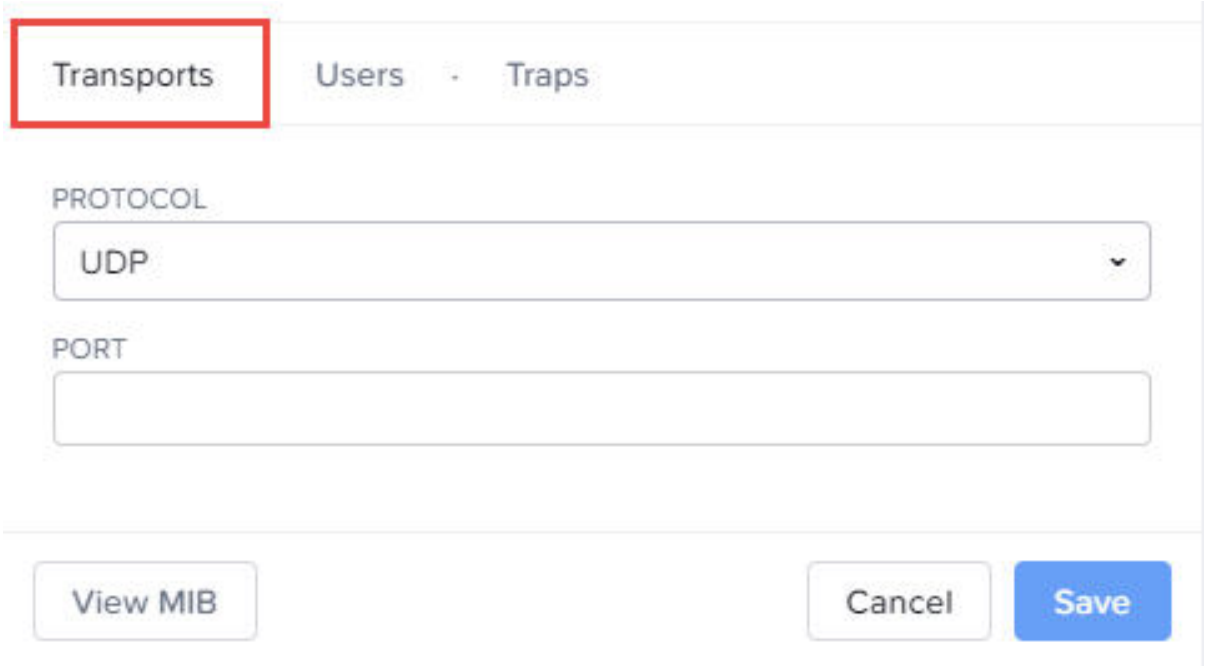

Figure 272: SNMP Transport Fields

a. **Protocol**: Select the protocol to use from the pull-down list.
   The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.

b. **Port**: Enter the port number to use.
   The standard SNMP port number is 161. For the complete list of required ports, see Port Reference.

c. When the fields are correct, click the **Save** button (lower right).
   This saves the configuration and redisplays the dialog box with the new transport appearing in the list.

> Note: To return to the SNMP Configuration window without saving, click the **Cancel** button.

5. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:

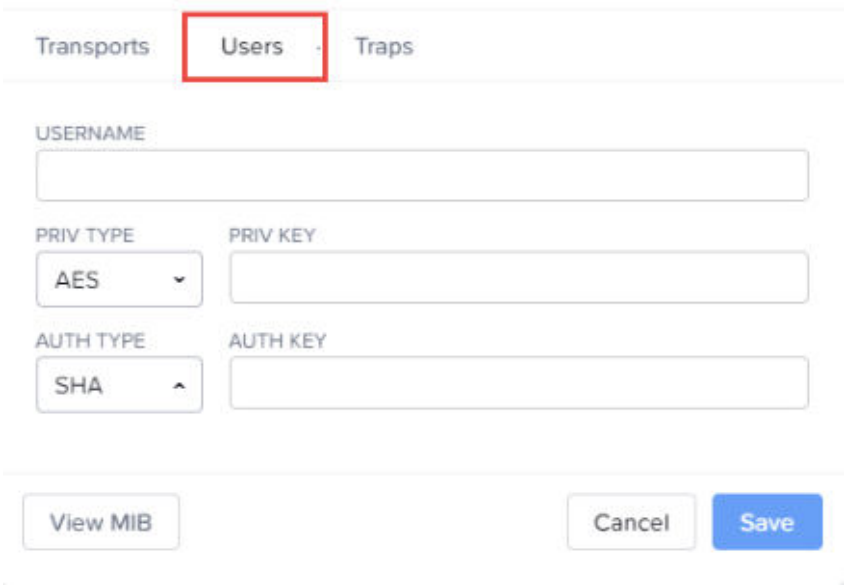

Figure 273: SNMP User Fields

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). In the nCLI, this setting is optional.

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.

6. To add an SNMP trap receiver, click the **Traps** tab and the **New Trap Receiver** button, and then do the following in the indicated fields:

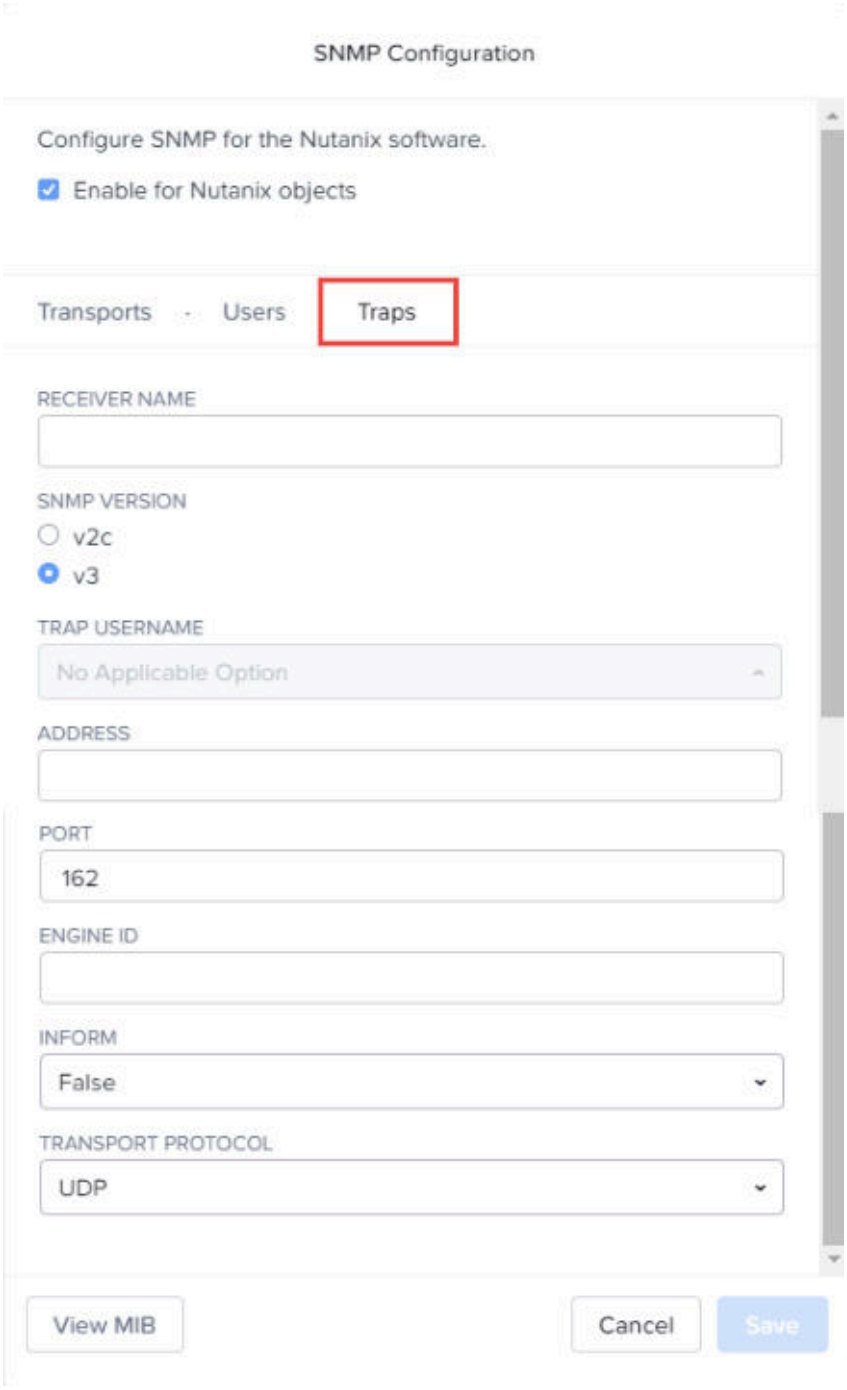

Figure 274: SNMP Trap Fields

a. **Receiver Name**: Enter the receiver name.

b. **SNMP Version**: Select (click the radio button for) the SNMP version, either v3 or v2c. (Version 1 is not supported.)

c. **Trap Username**: Select a user from the pull-down list.

All users added previously (see step 5) 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. For the complete list of required ports, see Port Reference.

f. **Engine ID**: Optionally, 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**.

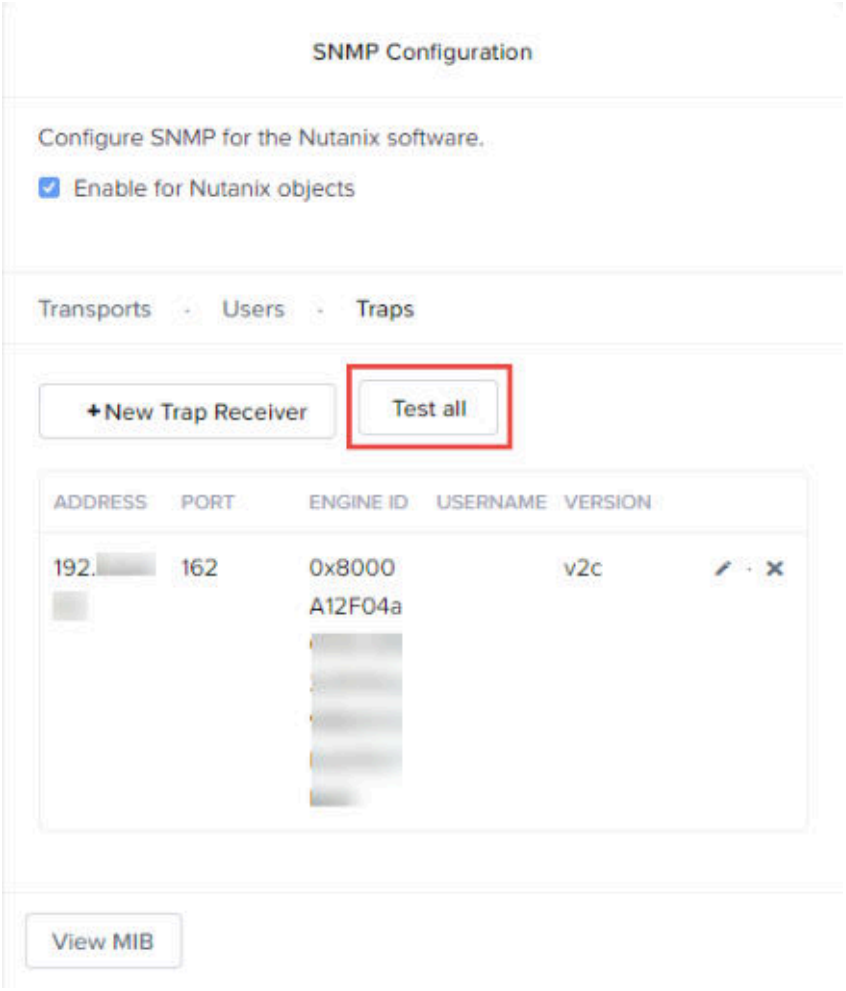

Figure 275: Test All Button

The Nutanix cluster sends test alerts to all the SNMP trap receivers configured on the cluster.

7. 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.

8. To delete an SNMP entry, click the appropriate tab (**Transports**, **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.

Nutanix MIB

Overview

The Simple Network Management Protocol (SNMP) enables administrators to monitor network-attached devices for conditions that warrant administrative attention. In the Nutanix SNMP implementation, information about entities in the cluster is collected and made available through the Nutanix MIB (NUTANIX-MIB.txt). The Nutanix enterprise tree is located at 1.3.6.1.4.1.41263.

The Nutanix MIB is divided into the following sections:

- Cluster information. Status information about the cluster as a whole.
- Software version information. Version information about the software packages that comprise the Controller VM.
- Service status information. Information about the status of essential services on each Controller VM.
- Hypervisor information. Information about each hypervisor instance.
- Virtual machine information. Information about hosted virtual machines.
- Disk information. Status information about the disks in the cluster.
- Controller VM resource information. Indicate how much CPU and memory capacity is available to a Controller VM.
- Storage pool information. Status information about the storage pools in the cluster.
- Storage Container information. Status information about the disks in the cluster.
- Alerts information. Information about generated alerts that can be captured through the SNMP trap (or inform) mechanism.

> Important: A statistic (counter) value resets to zero and starts increasing again after it reaches the maximum limit defined for its corresponding data type. The counter reinitialization is compliant with the *RFC 2578* standard.
>
> For example, the *vmRxBytes* statistic monotonically increases until it reinitializes on reaching a maximum value of its data type (Counter64).

Table 1001: Cluster Information Fields

Name	Description	Data Type
clusterName	Cluster name.	Display string
clusterVersion	Cluster version number. This is the Nutanix core package version expected on all the Controller VMs.	Display string

Name	Description	Data Type
clusterStatus	Current status of the cluster. Possible values are "started" and "stopped".	Display string
clusterTotalStorageCapacity	Total storage capacity of the cluster in bytes.	Unsigned 64-bit integer
clusterUsedStorageCapacity	Storage used on the cluster, in bytes.	Unsigned 64-bit integer
clusterIops	Average I/O operations per second (IOPS) in the cluster.	Unsigned 64-bit integer
clusterLatency	Average I/O latency in the cluster, in milliseconds.	Unsigned 64-bit integer
clusterIOBandwidth	Cluster-wide I/O bandwidth in kilobytes per second (KBps).	Unsigned 64-bit integer

Table 1002: Software Version Information Fields

Name	Description	Data Type
svtIndex	Unique index that is used to identify an entry in the software version information table.	Signed 32-bit integer
svtControllerVMId	Nutanix Controller VM identification number.	Display string
svtNutanixBootstrap	Nutanix bootstrap software package version.	Display string
svtNutanixInfrastructure	Nutanix infrastructure software package version.	Display string
svtNutanixCore	Nutanix core software package version.	Display string
svtNutanixToolchain	Nutanix toolchain software package version.	Display string
svtNutanixServiceability	Nutanix serviceability software package version.	Display string
svtLinuxKernel	Linux kernel version currently installed.	Display string

Table 1003: Service Status Information Fields

Name	Description	Data Type
cstIndex	Unique index that is used to identify an entry in the service status information table.	Signed 32-bit integer
cstControllerVMId	Nutanix Controller VM identification number.	Display string
cstControllerVMStatus	Status of the Nutanix node.	Display string
cstDataServiceStatus	Status of the core data services on the Controller VM.	Display string

Name	Description	Data Type
cstMetadataServiceStatus	Status of the metadata services on the Controller VM.	Display string

Table 1004: Hypervisor Information Fields

Name	Description	Data Type
hypervisorIndex	Number that is used to uniquely identify an entry in the hypervisor information table.	Signed 32-bit integer
hypervisorID	System-generated string that Nutanix uses to uniquely identify a hypervisor instance.	Display string
hypervisorName	Name of the hypervisor instance.	Display string
hypervisorVmCount	Number of VMs configured on the hypervisor instance.	Unsigned 32-bit integer
hypervisorCpuCount	Number of CPU cores available to the hypervisor instance.	Unsigned 32-bit integer
hypervisorCpuUsagePercent	Percentage of CPU resources in use by the hypervisor instance.	Unsigned 32-bit integer
hypervisorMemory	Total memory available to the hypervisor instance, in bytes.	Counter64
hypervisorMemoryUsagePercent	Memory in use by the hypervisor instance, as a percentage of the total available memory.	Unsigned 64-bit integer
hypervisorReadIOPerSecond	Total number of read I/O operations per second (IOPS) being performed by the hypervisor.	Unsigned 32-bit integer
hypervisorWriteIOPerSecond	Total number of write I/O operations per second (IOPS) being performed by the hypervisor.	Unsigned 32-bit integer
hypervisorAverageLatency	Average I/O latency of the hypervisor in microseconds (µs).	Counter64
hypervisorIOBandwidth	I/O bandwidth of the hypervisor in kilobytes per second (KBps).	Counter64
hypervisorRxBytes	Total number of bytes received by the hypervisor.	Counter64
hypervisorTxBytes	Total number of bytes transmitted by the hypervisor.	Counter64
hypervisorRxDropCount	Total number of packets dropped by the hypervisor when receiving data.	Counter64
hypervisorTxDropCount	Total number of packets dropped by the hypervisor when transmitting data.	Counter64

Table 1005: Virtual Machine Information Fields

Name	Description	Data Type
vmIndex	Number that is used to uniquely identify an entry in the VM information table.	Signed 32-bit integer
vmId	System-generated string that Nutanix uses to uniquely identify a virtual machine.	Display string
vmName	Name of the VM.	Display string
vmHypervisorId	System-generated ID of the hypervisor on which the VM is provisioned.	Display string
vmPowerState	Power state of the VM.	Display string
vmCpuCount	Number of CPU cores available to the VM.	Unsigned 32-bit integer
vmCpuUsagePercent	Percentage of CPU resources in use by the VM.	Unsigned 32-bit integer
vmMemory	Total memory allocated to the VM, in bytes.	Counter64
vmMemoryUsagePercent	Memory in use by the VM, as a percentage of the total allocated memory.	Unsigned 64-bit integer
vmReadIOPerSecond	Total number of read I/O operations per second (IOPS) being performed by the VM.	Unsigned 32-bit integer
vmWriteIOPerSecond	Total number of write I/O operations per second (IOPS) being performed by the VM.	Unsigned 32-bit integer
vmAverageLatency	Average I/O latency of the VM, in microseconds (µs).	Counter64
vmIOBandwidth	I/O bandwidth of the VM in kilobytes per second (KBps).	Counter64
vmRxBytes	Total number of bytes received by the VM.	Counter64
vmTxBytes	Total number of bytes transmitted by the VM.	Counter64
vmRxDropCount	Total number of packets dropped by the VM when receiving data.	Counter64
vmTxDropCount	Total number of packets dropped by the VM when transmitting data.	Counter64

Table 1006: Disk Information Fields

Name	Description	Data Type
dstIndex	Number that is used to uniquely identify an entry in the disk information table.	Signed 32-bit integer
dstDiskId	Disk identification number. The number is unique for each disk.	Display string

Name	Description	Data Type
dstControllerVMId	Nutanix Controller VM identification number.	Display string
dstSerial	Disk serial number.	Display string
dstNumRawBytes	Physical storage capacity on the device, in terms of number of raw bytes.	Unsigned 64-bit integer
dstNumTotalBytes	Usable storage on the device through its file system, in terms of number of usable bytes.	Unsigned 64-bit integer
dstNumFreeBytes	Available storage on the device through its file system for non-root users, in terms of number of free bytes.	Unsigned 64-bit integer
dstNumTotalInodes	Total number of usable inodes on the device through its file system.	Unsigned 64-bit integer
dstNumFreeInodes	Total number of available (free) inodes on the device through its file system for non-root users.	Unsigned 64-bit integer
dstAverageLatency	Average I/O latency of the disk in microseconds (µs).	Unsigned 64-bit integer
dstIOBandwidth	I/O bandwidth of the disk in kilobytes per second (KBps).	Unsigned 64-bit integer
dstNumberIops	Current number of I/O operations per second (IOPS) for the disk.	Unsigned 64-bit integer
dstState	State of the disk.	Signed 32-bit integer

Table 1007: Controller VM Resource Information Fields

Name	Description	Data Type
crtIndex	Number that is used to uniquely identify an entry in the Controller VM resource information table.	Signed 32-bit integer
crtControllerVMId	Nutanix Controller VM identification number.	Display string
crtMemory	Total memory available to the Controller VM in bytes.	Unsigned 64-bit integer
crtNumCpus	Total number of CPUs allocated to the Controller VM.	Signed 32-bit integer
crtName	Name of the Controller VM.	Display string

Table 1008: Storage Pool Information Fields

Name	Description	Data Type
spitIndex	Number that is used to uniquely identify an entry in the storage pool information table.	Signed 32-bit integer
spitStoragePoolId	Storage pool identification number.	Display string
spitStoragePoolName	Storage pool name.	Display string
spitTotalCapacity	Total storage pool capacity in bytes.	Unsigned 64-bit integer
spitUsedCapacity	Used storage pool capacity in bytes.	Unsigned 64-bit integer
spitIOPerSecond	Current number of I/O operations per second (IOPS) for this storage pool.	Signed 32-bit integer
spitAvgLatencyUsecs	Average I/O latency for the storage pool in microseconds.	Unsigned 64-bit integer
spitIOBandwidth	I/O bandwidth of the storage pool in kilobytes per second (KBps).	Unsigned 64-bit integer

Table 1009: Storage Container Information Fields

Name	Description	Data Type
citIndex	Number that is used to uniquely identify an entry in the storage container information table.	Signed 32-bit integer
citContainerId	Storage Container identification number.	Display string
citContainerName	Storage Container name.	Display string
citTotalCapacity	Total storage container capacity in bytes.	Unsigned 64-bit integer
citUsedCapacity	Used storage container storage in bytes.	Unsigned 64-bit integer
citIOPerSecond	Current number of I/O operations per second (IOPS) for this storage container.	Signed 32-bit integer
citAvgLatencyUsecs	Average I/O latency for the storage container in microseconds.	Unsigned 64-bit integer
citIOBandwidth	I/O bandwidth of the storage container in kilobytes per second (KBps).	Unsigned 64-bit integer

Trap Resolution

In addition to generating an SNMP trap when an alert condition is detected, Nutanix clusters generate a trap when an alert condition is resolved. The resolved trap, named ntxTrapResolved, is generated regardless of whether the alert condition is resolved manually or automatically. If a Prism Central alert is resolved in Prism Central, Prism Central sends a resolved trap. If a Nutanix

cluster alert is resolved in either the Prism web console or Prism Central, both Prism Central and the cluster send a resolved trap.

In this section, the terms original alert and original trap are used to refer, respectively, to the alert and trap that are generated when the alert condition is detected, and the terms resolved alert and resolved trap are used to refer, respectively, to the alert and trap that are generated when the alert condition is resolved.

To enable you to associate a resolved trap with the original alert and original trap, the UUID of the original alert is included in the original trap, as part of the alert message (ntxAlertDisplayMsg), and in the resolved trap, as a MIB object (ntxAlertUuid).

A resolved trap has the following MIB objects:

- ntxTrapName (name of the original trap)
- ntxAlertUuid (UUID of the original alert)
- ntxAlertResolvedTime
- ntxAlertDisplayMsg (message of the original alert)
- ntxAlertTitle (title of the original alert)
- ntxAlertSeverity (severity of the original alert)

The fields are described in the rows that follow.

Table 1010: Alerts Information Fields

Name	Description	Data Type
ntxAlertCreationTime	Time of alert creation. The value is the number of seconds since the UNIX epoch (01/01/1970).	Unsigned 64-bit integer
ntxAlertDisplayMsg	Alert message text.	Display string
ntxAlertTitle	Alert title.	Display string
ntxAlertSeverity	Alert severity, which has one of the following values: 1. Informational 2. Warning 3. Critical 4. Audit	Integer
ntxTrapName	Name of the trap that was generated when the alert condition was detected. This MIB object is included in resolved traps.	Display string
ntxAlertUuid	UUID of the alert that was generated when the alert condition was detected. This MIB object is included in resolved traps.	Display string
ntxAlertResolvedTime	Time at which the alert was resolved. The value is the number of seconds since the UNIX epoch (01/01/1970). This MIB object is included in resolved traps.	Unsigned 64-bit integer

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 the web console. 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.
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.
   You can disable (or enable) the banner at any time by unchecking (checking) the **Enable Banner** box.

Example

The following figure is a sample banner page configuration.

> Note: 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.

Figure 276: Welcome Banner Window

VM Management through Prism Element (ESXi)

From AOS 5.0 or later releases, you can perform your core VM management operations directly from Prism without using any hypervisor management interface (for example, vCenter Server). The VM Management through Prism for ESXi feature provides an unified management interface for all of the ESXi hypervisors. For this functionality to work, you need to register vCenter Server with the Prism Element or multiple vCenter Servers with the Prism Central. For more

information about registering vCenter Server to your cluster, see Registering a Cluster to vCenter Server on page 771.

Note:

By using this feature you can perform following operations directly through Prism.

- Create, clone, update, and delete VMs.
- Create and delete NICs.
- Attach and delete disks.
- Power operations: Power on or off, reset, suspend, resume, guest shutdown, and guest restart.
- Open and launch VM console.
- Manage VM guest tools (mounting VMware guest tools, mounting NGT).

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 is managed by a single vCenter Server.
- Ensure that DRS is enabled on the vCenter Server.
- Ensure that you are running ESXi and vCenter Server 5.5 or later releases.
- Ensure that you have homogeneous network configuration. 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 on page 773.

CAUTION: If multiple vCenter Servers are managing the hosts, you will not be able to perform the VM management operations. Move all the hosts into one vCenter Server.

Requirements and Limitations

- SCSI, IDE, and SATA disks are supported. PCI disks are not supported.
- The E1000, E1000e, PCnet32, VMXNET, VMXNET 2, VMXNET 3 network adapter types (NICs) are 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 gets deleted.
- Network configuration (creation of port groups or VLANs) is not supported.

Registering a Cluster to vCenter Server

To perform core VM management operations directly from Prism without switching to vCenter Server, you need to register your cluster with the vCenter Server.

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

Following are some of the important points about registering vCenter Server.

- Nutanix does not store vCenter Server credentials.
- Whenever a new node is added to Nutanix cluster, vCenter Sever registration for the new node is automatically performed.
- Nutanix does not support vCenter Enhanced Linked Mode.

Procedure

1. Log into the Prism web console.
2. Click the gear icon in the main menu and then select **vCenter Registration** in the **Settings** page.

   The vCenter Server that is managing the hosts in the cluster is auto-discovered and displayed.
3. Click the **Register** link.

   The IP address is auto-populated in the **Address** field. The port number field is also auto-populated with 443. Do not change the port number. For the complete list of required ports, see Port Reference.

4. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.

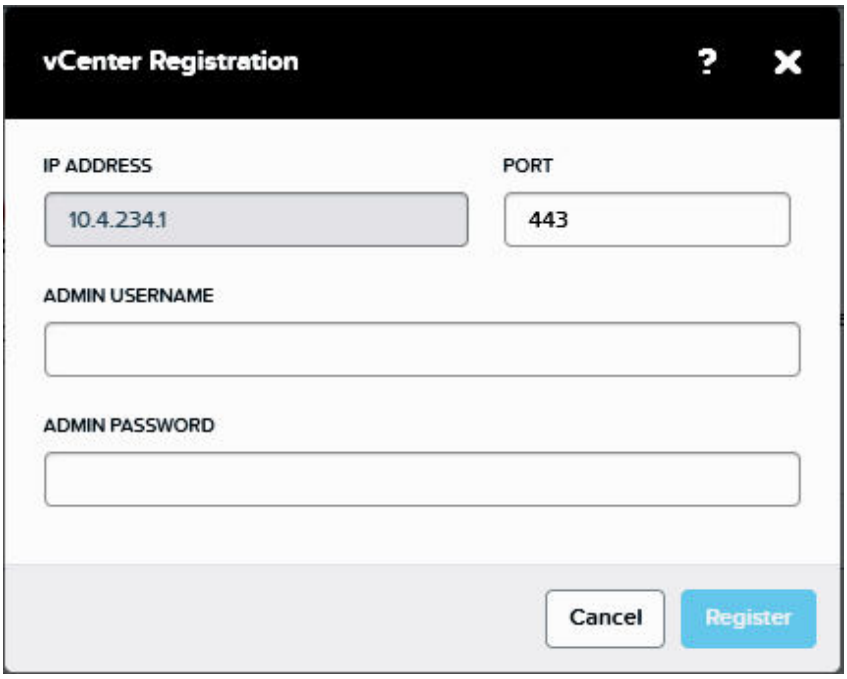

Figure 277: vCenter Registration Figure 1

5. Click **Register**.

   During the registration process a certificate is generated to communicate with the vCenter Server. If the registration is successful, 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.

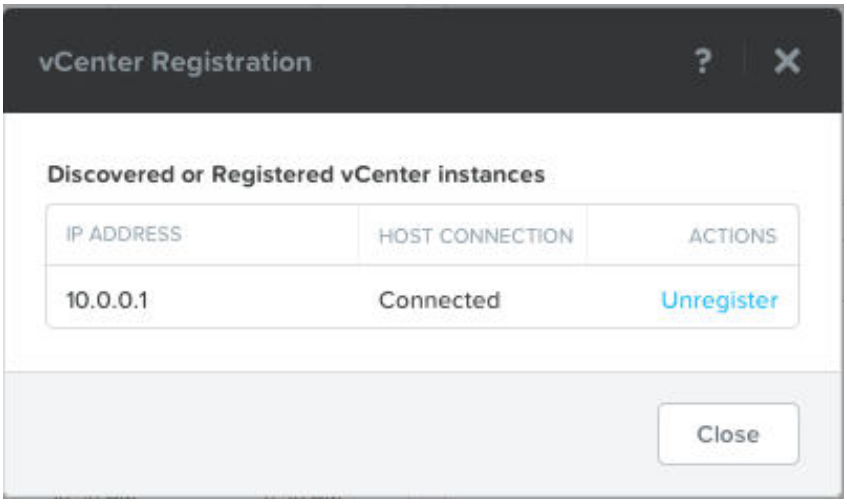

Figure 278: vCenter Registration Figure 2

Unregistering a Cluster from the vCenter Server

To unregister the vCenter Server from your cluster, 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.

Procedure

1. Log into the Prism web console.
2. Click the gear icon in the main menu and then select **vCenter Registration** in the **Settings** page.

   A message that cluster is already registered to the vCenter Server is displayed.
3. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.
4. 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

After you change the IP address of the vCenter Sever or if current vCenter Server is not managing the hosts then 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. Log into the Prism web console.

2. Click the gear icon in the main menu and then select **vCenter Registration** in the **Settings** page.
   The vCenter Server that is managing the hosts in the cluster is displayed along with the vCenter Server that you must unregister.

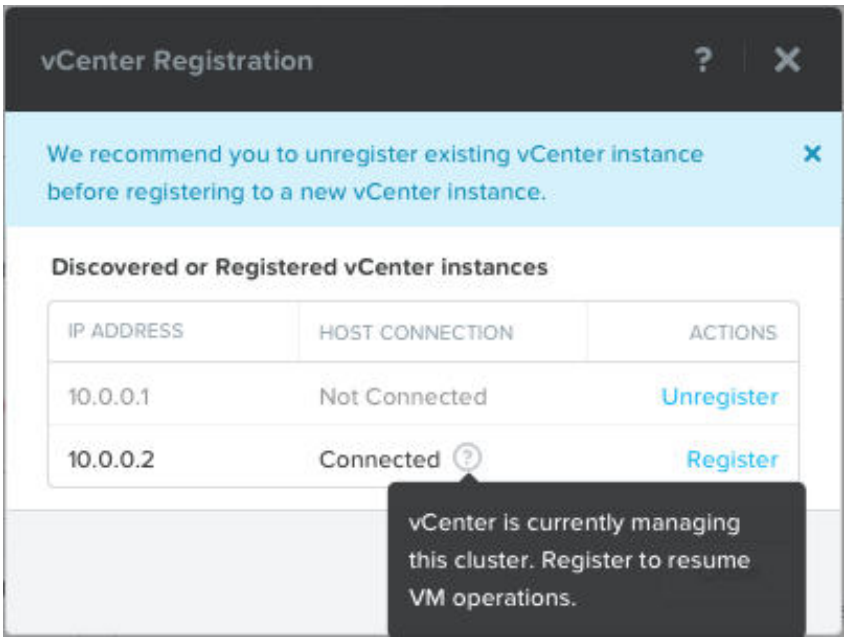

Figure 279: vCenter Registration Management

3. Unregister the existing vCenter Server instance by clicking **Unregister**.
   For more information about unregistering vCenter Server, click Unregistering a Cluster from the vCenter Server on page 772.

4. Register the new vCenter Server instance by clicking **Register**.
   For more information about registering vCenter Server, click Registering a Cluster to vCenter Server on page 771.

In-Place Hypervisor Conversion

The In-place hypervisor conversion feature provides you with an option to convert your existing ESXi cluster to an AHV cluster. All the VMs that are running in the ESXi cluster are converted so that they can run on the AHV cluster. The VM conversion is independent of the host that the VM is running on. The conversion process occurs in such a way that the VM down time is minimised. You can also convert the cluster back to ESXi. This conversion from AHV to ESXi is a file system restore of the host. Hence the cluster will be similar to what it was before the conversion.

Note:

- This feature converts your existing ESXi cluster to an AHV cluster. You cannot start the conversion process on the AHV cluster.
- Do not remove the hosts from the vCenter Server if you want to perform the reverse conversion process (AHV to ESXi).
- This feature is supported on the clusters with multiple hypervisors (combination of ESXi nodes with one AHV node or multiple AHV nodes).

- When you are converting the cluster from ESXi to AHV, remove the unused adapters displayed under **Teaming and Failover** in vCenter Server and keep only those adapters that are active and of the same speeds as the standby and active adapters.

Following are enhancements to In-Place Hypervisor Conversion feature.

- Decreased VM downtime: With the implementation of the new workflow of converting the nodes in the cluster in a rolling manner, the VM downtime is reduced to only shut down time and conversion time required for that particular VM. Approximately, the VM downtime has been reduced from 3 to 4 hours to less than 5 minutes.
- Prism state: The Prism console is responsive during the conversion process. However, the Prism goes into Read-Only state.
- State of the VM: The current state of the VM is preserved and the VM is brought back into the same state post conversion. For example, the VM is automatically powered on post conversion if that VM was powered on before you started the conversion.
- Preservation of the MAC addresses of the VM NICs after conversion: After the cluster is converted to AHV or ESXi, MAC addresses of the VM NICs are preserved. As a result of the MAC address preservation, both the static and dynamically configured IP addresses are preserved after conversion.

  Note: IP addresses are not preserved for Windows VMs.

- VMs that were created after conversion to AHV are retained post conversion back to ESXi.

Requirements and Limitations for In-Place Hypervisor Conversion

Following are the prerequisites with the requirements and limitations for an in-place hypervisor conversion.

Prerequisites

- Before performing In-place hypervisor conversion, ensure that you resolve all the NCC health check alerts (warnings, failures and errors) and upgrade the firmware to the latest version using LCM firmware upgrade.
- HA and DRS must be enabled.
- All the hosts must be managed by the same vCenter Server.
- The vCenter Server VM that is managing the cluster must not be running on the cluster which is being converted.
- Install NGT on all the VMs. See Enabling NGT and Mounting the NGT Installer in a VM on page 369 for more information. If NGT is not installed, the VM does not boot after conversion.
- All hardware platforms are supported.
- Ensure that you have all the ESXi ISOs (if you have more than one version of ESXi running in your cluster) at the **foundation/isos/hypervisor/esx/** location. This ISO image or images are used to bring the cluster back to its pre-conversion state if you abort the cluster conversion operation from ESXi to AHV.

General Limitations

In-place hypervisor conversion has the following limitations:

- Metro availability protection domains cannot be enabled in your environment.
- If you are using the free ESXi license for your cluster, you can convert the cluster to AHV. However, you cannot revert the cluster back to ESXi.
- You cannot have Nutanix Files deployed in your cluster.
- Affinity policies (if configured) are not honored after conversion to AHV.
- In-place hypervisor conversion is not supported for single-node clusters and two-node clusters.
- You cannot add multiple networks with the same VLAN ID to the same virtual switch like vs0. The conversion ignores such duplicate networks and adds only one network for every unique VLAN ID.

ESXi Supported Versions

Table 1011: Supported ESXi Versions

Version	Support
5.5	All the versions that are included in the ISO Whitelist (see ~/foundation/config/iso_whitelist.json file) that is bundled with AOS are supported. If you have upgraded foundation by using standalone foundation, the ISO Whitelist is invalidated. In this case, you need to see the original Whitelist that was present before you performed the upgrade. To verify the json file that was bundled with AOS, see the nutanix-packages.json file.
6.0	
6.5	
6.7	
7.0	

ESXi Network Requirements and Limitations

Table 1012: ESXi Network Requirements and Limitations

Component	Description
NIC	Minimum of two or more homogenous NICs of 1 GB, 10 GB, or 40 GB are supported.
vSwitch	• Each host should have only one external vSwitch. If you have more than one external vSwitch conversion validation fails. • Ensure you have one active and multiple passive failover configuration of NICs on vSwitch. • Active/active load balancing policies are not supported and is converted to active/passive on AHV. • For a standard switch configuration it is recommended that all the port groups are present on all the hosts because there might be a possibility that the conversion process might move the VMs to a different host.

Component	Description
Distributed vSwitch	Supported
Internal vSwitch	Internal vSwitch apart from Nutanix vSwitch is not supported.
LACP	Not supported.

- The vSwitch is converted to an open vSwitch on the AHV side. If you perform any configuration changes on AHV, these changes may not be maintained after converting the cluster back to ESXi.
- Networks not in use by a VM prior to the cluster conversion may not be converted.
- All the VMs may end up on a single host during conversion back to ESXi. The network configuration that you might have defined for the VMs may not be maintained.
- After migration to AHV, NIC may not have same type of virtualized hardware device (for example VMXNET3 may be E1000 after conversion).
- Serial ports and virtual graphics processing unit (vGPU) configurations may be lost after conversion.

Virtual Machine Requirements and Limitations

- Only VMs with flat disks are supported. The delta disks are not supported.
- Only IDE and SCSI storage controllers are supported for automatic conversion. SATA and PCI disks are not supported.
- On Windows VMs, set the SAN policy to **OnlineAll** for non-boot SCSI disks so that they can be automatically brought online. For more information about setting SAN policy, see “” () .
- VMs that have NFS datastore folders with vSphere tagging can not be converted..
- Ensure that you do not have VMs with UEFI boot firmware enabled.
- Virtual machines with attached volume groups or shared virtual disks are not supported.
- After reverting back to ESXi from AHV, the VMs are converted to the maximum hardware version that is supported by that specific ESXi version.
- Guest OS network interfaces on Linux VMs may change to generic type during ESXi to AHV conversion (for example, in RHEL 7 network interface enoXXXX changes to eth0). You may have to reconfigure the network settings according to changes post conversion.
- Guest OS type for the Linux VMs may change to a more generic type (for example RHEL 7 may change to Other Linux 64-bit) during the conversion back from AHV to ESXi.

Supported Operating Systems

Table 1013: Supported Operating Systems

Operating System	Versions	Limitations
Windows	• Windows 2008 R2 or later versions • Windows 7 or later versions	• Only 64-bit operating systems are supported.
Linux	• CentOS 6.5 and 7.0 • Red Had Enterprise Linux (RHEL) 6.5 or later and RHEL 7.0 or later. • Oracle Linux 6.5 and 7.0 • Ubuntu 14.04	SUSE Linux Enterprise Server (SLES) operating system is not supported.

In-Place Hypervisor Conversion Process

ESXi to AHV

After you start the conversion process, all the nodes in the cluster are converted in a rolling manner to AHV one node at a time. During conversion, the first node is placed in the maintenance mode and all the VMs that are running on the node are migrated to other ESXi nodes in the cluster using the HA and DRS feature. After the VMs are migrated, the node is converted to AHV. After node is successfully converted, all VMs that were migrated to ESXi are migrated one at a time to AHV. Similar steps are performed for the rest of the nodes in the ESXi cluster until the last ESXi node. The VMs that are running on the last ESXi node in the cluster are converted and migrated to the AHV hosts and then the ESXi host is converted to AHV. If any error occurs during VM conversion, appropriate alerts or error messages are displayed. When converting a VM the source vdisk is not modified. Therefore, if there are any fatal errors during imaging, storing, or restoring of the configuration, the conversion is stopped and you are prompted to abort the conversion.

> Note: After conversion, do not remove the host from the vCenter Server until you are sure that the conversion is successful, because it may impact the reverse conversion process.

Port Groups and VLAN Transformation

Before the conversion all of the ESXi port groups, VLAN IDs, and virtual machines that belong to a particular group are captured. On AHV, a corresponding Acropolis virtual network is created for every unique ESXi port group on the cluster including any corresponding VLAN ID, and then associated with the VMs. If the ESXi host has a VLAN set on the management port group, after conversion the Acropolis management interface and the Controller VM public interface are placed in that same VLAN.

AHV to ESXi

During the reverse conversion (AHV to ESXi), the process of conversion is similar. Additionally, if the cluster does not have the ESXi ISO stored on the cluster, you need to provide the ESXi ISO image during the conversion process.

Note:

- The image that you provide should be of the same major ESXi version that you have used during ESXi to AHV conversion.
- If new nodes were added to the cluster after conversion to AHV, then these nodes need to be removed before starting the reverse conversion process.

After conversion to ESXi, all the hosts are automatically registered to the vCenter Server.

Converting Cluster (ESXi to AHV)

Perform the following procedure to convert the cluster from ESXi to AHV.

Before you begin

Verify that you have met all the networking and virtual machine requirements as described in Requirements and Limitations for In-Place Hypervisor Conversion on page 775.

About this task

Procedure

1. Log into the web console.
2. Click the gear icon in the main menu and then select **Convert Cluster** in the **Settings** page.
3. From the **Available Hypervisor** drop-down menu, select **AHV**.

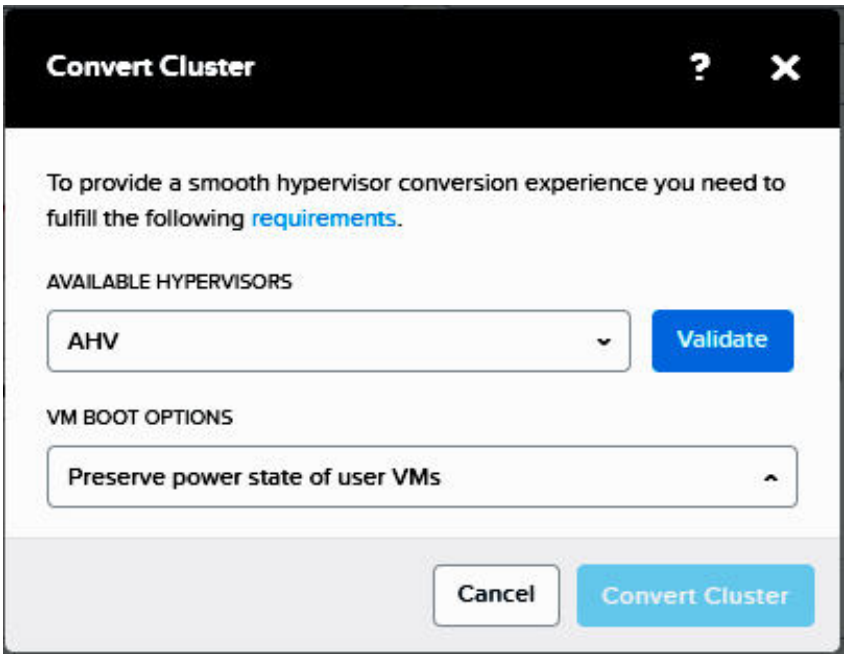

Figure 280: Convert Cluster

4. Select the state of the VMs that you want post conversion from the **VM Boot Options** drop-down menu.

   » **Preserve power state of the user VMs**: Select this option if you want to keep the original power state of the VMs. For example, if you want the VMs to be in a running state post conversion automatically, select this option.

   » **Power Off User VMs**: Select this option if you want power off all the VMs running on ESXi cluster before you start the conversion process. After conversion, these VMs will be in the powered off state.

5. Click **Validate** to enter vCenter Server credentials and to verify whether you have met all the requirements.

6. Enter the IP address of the vCenter Server in the **vCenter Sever IP Address** along with the administrator user name and password of the vCenter Server in the **Username** and **Password** fields.

7. Click **Yes**.

   A validation that you have met all the requirements is performed. Once the validation is successful the conversion process proceeds. If validation fails (for any reason), a relevant message to take appropriate action is displayed.

8. Click **Convert Cluster**.

   A confirmation message is displayed.

   > Note: The cluster changes to Read Only mode. A message Oops - server Error may be displayed when the Prism node is undergoing conversion. Wait for the conversion process to complete. You can access other Controller VM for Read-Only Prism operations.

   The entire conversion process may take 3 to 4 hours depending on the nodes that are present in your cluster. However, the VM downtime will be less than 5 minutes because all the nodes in the cluster are converted in a rolling manner. You can also track the progress of the conversion by logging again into the web console.

Converting Cluster (AHV to ESXi)

Perform the following procedure to convert the cluster back to ESXi from AHV.

About this task

Perform this procedure only if you have converted your cluster from ESXi to AHV. You cannot start the conversion process on the AHV cluster.

Before you begin

Verify that you have met all the networking and virtual machine requirements as described in Requirements and Limitations for In-Place Hypervisor Conversion on page 775 topic.

Procedure

1. Log into the web console.
2. Click the gear icon in the main menu and then select **Convert Cluster** in the **Settings** page.
3. From the **Available Hypervisor** drop-down menu, select **ESXi**.

4. Select the state of the VMs that you want post conversion from the **VM Boot Options** drop-down menu.

   » **Preserve power state of the user VMs**: Select this option if you want to keep the original power state of the VMs. For example, if you want the VMs to be in a running state post conversion automatically, select this option.

   » **Power Off User VMs**: Select this option if you want power off all the VMs running on ESXi cluster before you start the conversion process. After conversion, these VMs will be in the powered off state.

5. Click **Validate** to enter vCenter Server credentials and to verify whether you have met all the requirements.

6. Enter the IP address of the vCenter Server in the **vCenter Sever IP Address** along with the administrator user name and password of the vCenter Server in the **Username** and **Password** fields.

7. Click **Yes**.
   A validation that you have met all the requirements is performed. Once the validation is successful the conversion process proceeds. If validation fails (for any reason), a relevant message to take appropriate action is displayed.

8. Click **Convert Cluster** button to start the conversion process.

9. (Optional) If you have not saved the ESXi ISOs at the foundation/isos/hypervisor/esx/ location, click **Choose File** and select the ESXi ISO.

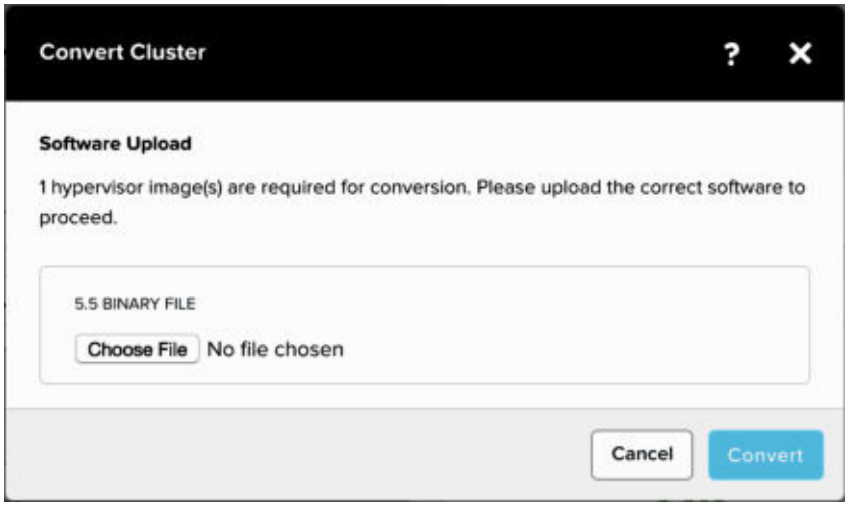

Figure 281: Software Upload

Note: If you have different versions of ESXi running in your cluster, you have to perform this step for every version of ESXi ISO.

10. Click **Convert**.
    All cluster operations stops and you will not be able to perform any operations on the cluster while the conversion is in progress. A message to this effect is displayed.

11. Click **Yes** to proceed with the conversion.

    The conversion begins. The time taken to finish the conversion is dependent on your environment. For example, it might take between 2 to 3 hours for the conversion to finish for a 3 or 4 nodes cluster. However, the VM downtime will be less than 5 minutes.

Stopping Cluster Conversion

If any issues occur with the cluster conversion process, a message to stop the cluster conversion or to continue with the conversion process is displayed. Depending on the type of message, you can appropriate action.

Before you begin

Ensure that you have saved the ESXi ISO at the foundation/isos/hypervisor/esx/ location. This ISO image is used to bring the cluster back to its pre-conversion state if you stop the cluster conversion operation from ESXi to AHV.

About this task

Issues in the cluster conversion process might occur because of the following reasons.

- Imaging Issue: If this issue occurs, you can only stop the cluster conversion process and a relevant message is displayed. Aborting the conversion reverts the cluster back to the its original state.
- VM Conversion Issue: If this issue occurs, you can either stop the conversion or can continue with the process. If you decide to continue with the process, cluster conversion is completed keeping the current state of the VM. After conversion is completed, you must perform appropriate actions on the VMs to bring back the VMs.

Procedure

Click **Abort** when a message to stop the cluster conversion is displayed.

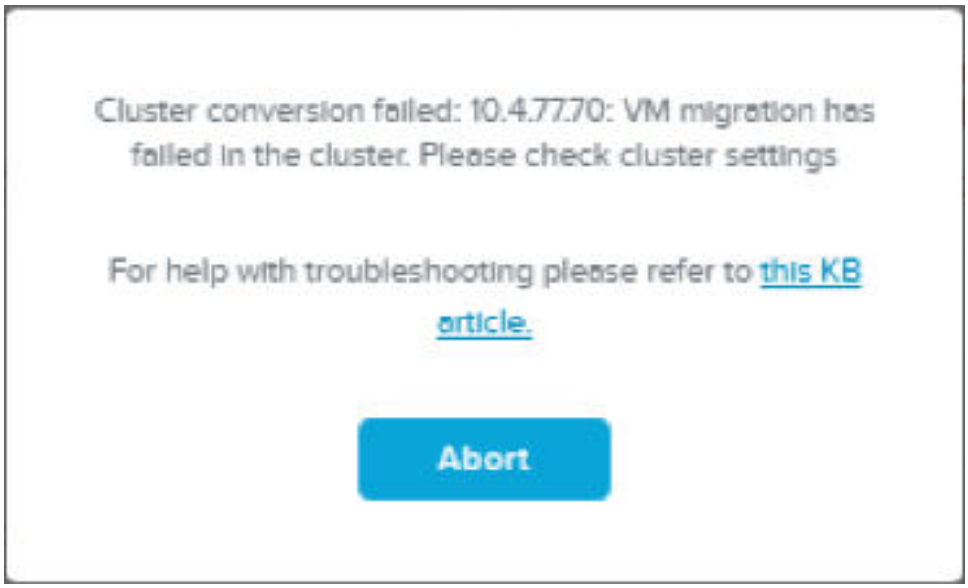

Figure 282: Abort Operation

The cluster goes back to its original state. You may need to manually power on the VMs depending on its state after you perform this operation.

Internationalization (i18n)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 1014: Internationalization Support

Supported Entities	Unsupported Entities
Cluster name	Acropolis file server
Storage Container name	Share path
Storage pool	Internationalized domain names
VM name	E-mail IDs
Snapshot name	Hostnames
Volume group name	Integers
Protection domain name	Password fields
Remote site name	Any Hardware related names (for example, vSwitch, iSCSCI initiator, vLAN name)
User management	
Chart name	

CAUTION: The creation of none of the above entities are supported on Hyper-V because of the DR limitations.

Entities Support (ASCII or non-ASCII) for the Active Directory Server

- In the New Directory Configuration, **Name** field is supported in non-ASCII.
- In the New Directory Configuration, **Domain** field is not supported in non-ASCII.
- In Role mapping, **Values** field is supported in non-ASCII.
- User names and group names are supported in non-ASCII.

Localization (L10n)

Nutanix localizes the user interface in Simplified Chinese and Japanese language. All the static screens are translated to the selected locale language. All the dashboards (including tool tips) and menus of the Prism Element are localized.

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 on page 784.

If the Prism Central instance is launched from the Prism Element, language settings of the Prism Central takes precedence over Prism Element.

You can also create new users with the specified language setting. For more information, see “” () .

Conditions and Limitations

- Logical entities that do not have a contextual translation available in the localized language are not localized.
- 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

Perform the following procedure to set the Prism language settings to a chosen language (English, simplified Chinese, or Japanese). You also have an option to set the calendar, date, and time format to a selected region.

Procedure

1. Click the gear icon in the main menu and then select **Language Settings** in the **Settings** page.
   The Language Settings dialog box appears.

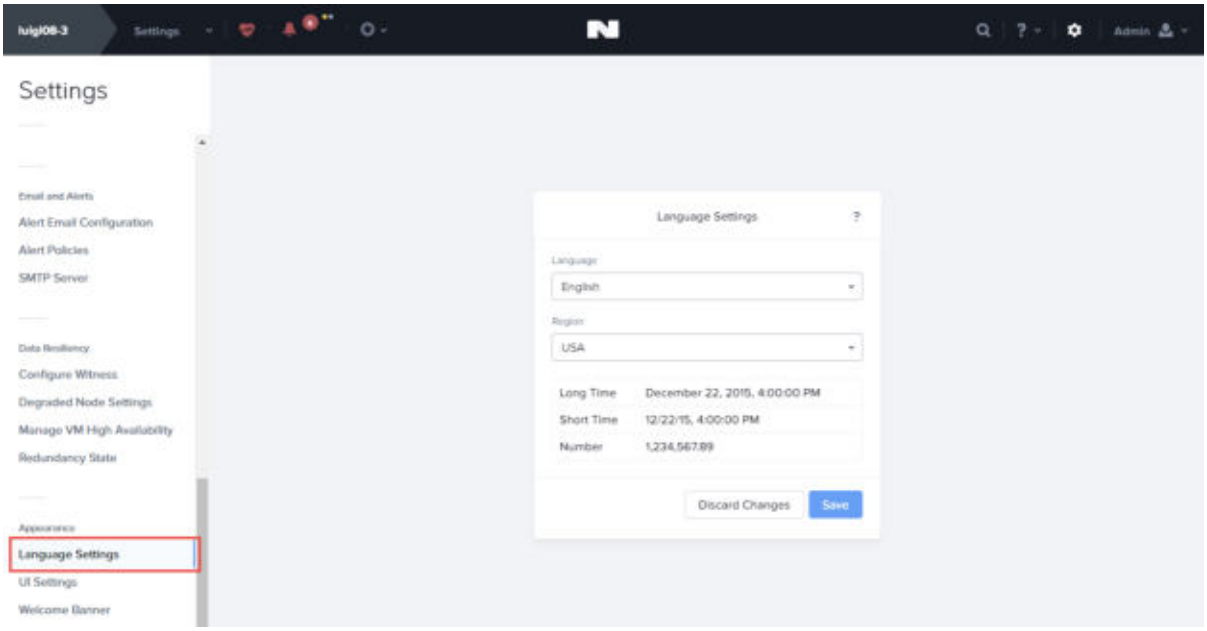

Figure 283: Language Settings Window

2. To change the language, select the desired language from the **Languages** field pull-down menu.

   The **English** language is selected by default, but you can change that to either **Simplified Chinese** or **Japanese**.

3. To change the locale settings (date, time, calendar), select the appropriate region from the **Region** field drop-down menu.

   A default locale is set based on the language setting. However, you can change the region to display the date, time, and calendar in some other format. This format for date, time, and calendar is applied to the entire cluster.

4. When the settings are correct, click the **Save** button.

   The language and locale settings are changed according to the selection. For example, in the following screenshot, the language setting for the cluster is changed to Chinese and the locale setting is changed to Russian.

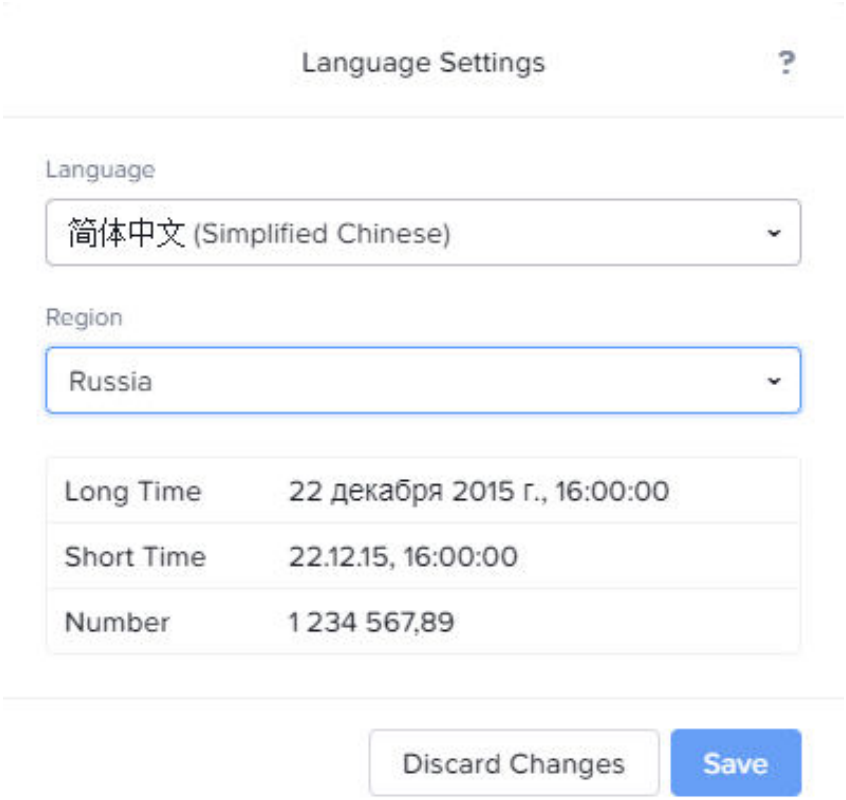

Figure 284: Locale Settings

For more information on the entities that are supported, see Internationalization (i18n) on page 783. For more information about localization, see Localization (L10n) on page 783.

Hyper-V Setup

Joining the Cluster and Hosts to a Domain

After completing foundation of the cluster, you need to join the cluster and its constituent hosts to the Active Directory domain. The joining of cluster and hosts to the domain facilitates centralized administration and security through the use of other Microsoft services such as Group Policy and enables administrators to manage the distribution of updates and hotfixes.

Before you begin

- If you have a VLAN segmented network, verify that you have assigned the VLAN tags to the Hyper-V hosts and Controller VMs. To configure VLANs for the Controller VM, see the *Advanced Setup Guide*.
- Ensure that you have valid credentials of the domain account that has the privileges to create a new or modify an existing computer accounts in the Active Directory domain. Active Directory domain created by using non-ASCII text may not be supported. For more

information about usage of ASCII or non-ASCII text in Active Directory configuration, see *Internationalization (i18n)* section.

Procedure

1. Log into the Web console by using one of the Controller VM IP address or by using cluster virtual IP address.
2. Click the gear icon in the main menu and select **Join Cluster and Hosts to the Domain** from the **Settings** page.

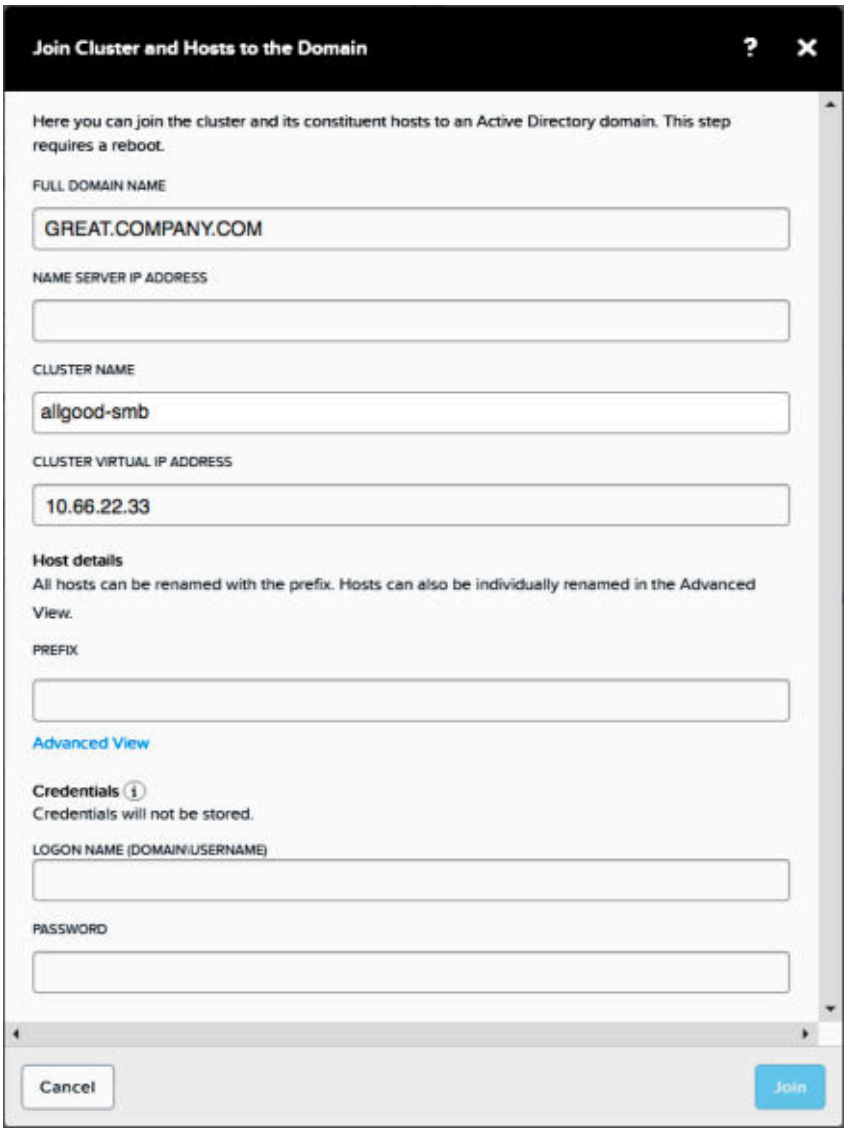

Figure 285: Join Cluster and Hosts to the Domain

3. Type the fully qualified name of the domain that you want to join the cluster and its constituent hosts to in the **Full Domain Name** text box.
4. Type the IP address of the name server in the **Name Server IP Address** text box that can resolve the domain name that you have entered in the **Full Domain Name** text box.

5. Type a name for the cluster in the **Cluster Name** text box.
   The maximum length of the cluster name should not be more than 15 characters and it should be a valid NetBIOS name.
6. Type the virtual IP address of the cluster in the **Cluster Virtual IP Address** text box.
   If you have not already configured the virtual IP address of the cluster, you can configure it by using this field.
7. Type the prefix that should be used to name the hosts (according to your convention) in the **Prefix** text box.
   - The prefix name should not end with a period.
   - The maximum length of the prefix name should not be more than 11 characters.
   - Should be a valid NetBIOS name.
     For example, if you enter prefix name as Tulip, the hosts are named as Tulip-1, Tulip-2, and so on, in the increasing order of the external IP address of the hosts.

   If you do not provide any prefix, the default name of NTNX-*block-number* is used. Click the **Advanced** link to see the expanded view of all the hosts in all the blocks of the cluster and to rename them individually.
8. In the **Credentials** field, type the logon name and password of the domain account that has the privileges to create a new or modify an existing computer accounts in the Active Directory domain.
   This logon name must be in the format *DOMAIN\USERNAME*. The credentials are required to join the cluster and its constituent hosts to the domain. Nutanix does not store the credentials.
9. When all the information is correct, click **Join**.
   The cluster is joined to the domain. All the hosts are renamed and joined to the domain and are restarted. Wait for about 5 to 10 minutes before proceeding to allow the hosts and Controller VMs time to start. After the cluster is ready, the logon page is displayed.

What to do next

Create a Microsoft failover cluster. For more information, see Creating a Failover Cluster for Hyper-V on page 787.

Creating a Failover Cluster for Hyper-V

Before you begin

Perform the following tasks before you create a failover cluster:

- Join the hosts to the domain as described in "Joining the Cluster and Hosts to a Domain" in the *Hyper-V Administration for Acropolis Guide* or the *Prism Web Console Guide*.
- Ensure that the Windows failover cluster IP address and hypervisor or host IP address are in the same subnet.

Perform the following procedure to create a failover cluster that includes all the hosts in the cluster.

Procedure

1. Log on to the Prism Element web console by using one of the Controller VM IP addresses or by using the cluster virtual IP address.

2. Click the gear icon in the main menu and select **Configure Failover Cluster** from the **Settings** page.

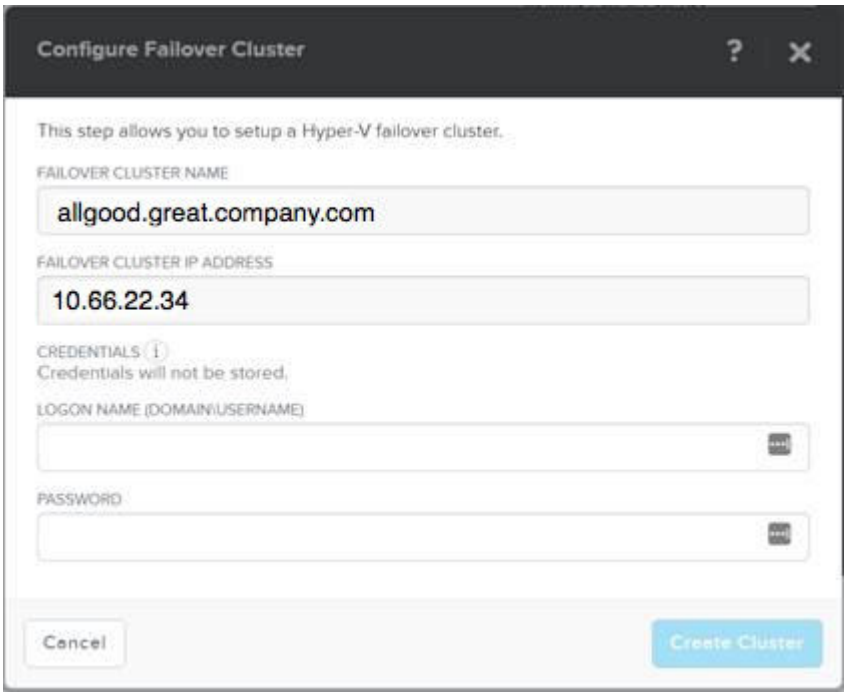

Figure 286: Configure Failover Cluster

3. Type the failover cluster name in the **Failover Cluster Name** text box.

   The maximum length of the failover cluster name must not be more than 15 characters and must be a valid NetBIOS name.

4. Type an IP address for the Hyper-V failover cluster in the **Failover Cluster IP Address** text box.

   This address is for the cluster of Hyper-V hosts that are currently being configured. It must be unique, different from the cluster virtual IP address and from all other IP addresses assigned to the hosts and Controller VMs. It must be in the same network range as the Hyper-V hosts.

5. In the **Credentials** field, type the logon name and password of the domain account that has the privileges to create a new account or modify existing accounts in the Active Directory domain.

   The logon name must be in the format *DOMAIN\USERNAME*. The credentials are required to create a failover cluster. Nutanix does not store the credentials.

6. Click **Create Cluster**.

   A failover cluster is created by the name that has been provided and it includes all the hosts in the cluster.

   See Manually Creating a Failover Cluster (SCVMM User Interface) on page 788 for information on manually creating a failover cluster.

Manually Creating a Failover Cluster (SCVMM User Interface)

Join the hosts to the domain as described in "Joining the Cluster and Hosts to a Domain" in the Hyper-V Administration for Acropolis guide or the Prism Web Console guide.

About this task

Perform the following procedure to manually create a failover cluster for Hyper-V by using System Center VM Manager (SCVMM).

If you are not using SCVMM or are using Hyper-V Manager, see Creating a Failover Cluster for Hyper-V on page 787.

Procedure

1. Start the Failover Cluster Manager utility.
2. Right-click and select **Create Cluster**, and click **Next**.
3. Enter all the hosts that you want to add to the Failover cluster, and click **Next**.
4. Select the **No. I do not require support from Microsoft for this cluster, and therefore do not want to run the validation tests. When I click Next continue creating the cluster** option, and click **Next**.

   > Note:
   >
   > If you select **Yes**, two tests fail when you run the cluster validation tests. The tests fail because the internal network adapter on each host is configured with the same IP address (192.168.5.1). The network validation tests fail with the following error message:
   >
   > Duplicate IP address
   >
   > The failures occur despite the internal network being reachable only within a host, so the internal adapter can have the same IP address on different hosts. The second test, Validate Network Communication, fails due to the presence of the internal network adapter. Both failures are benign and can be ignored.

5. Enter a name for the cluster, specify a static IP address, and click **Next**.
6. Clear the **All eligible storage to the cluster** check box, and click **Next**.
7. Wait until the cluster is created. After you receive the message that the cluster is successfully created, click **Finish** to exit the Cluster Creation wizard.

8. Go to **Networks** in the cluster tree and select **Cluster Network 1** and ensure it is in the internal network by verifying the IP address in the summary pane. The IP address must be 192.168.5.0/24 as shown in the following screen shot.

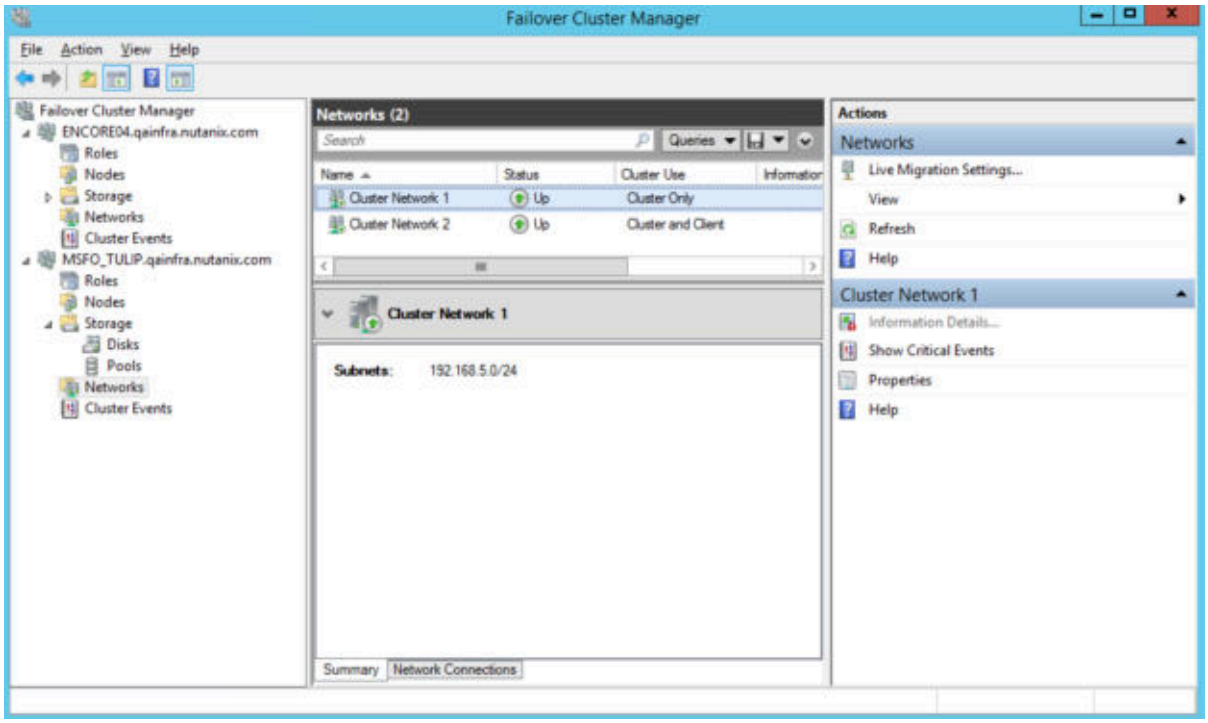

Figure 287: Failover Cluster Manager

9. Click the **Action** tab on the toolbar and select **Live Migration Settings**.

10. Remove **Cluster Network 1** from **Networks** for Live Migration and click **OK**.

> Note: If you do not perform this step, live migrations fail because the internal network is added to the live migration network lists. Log on to SCVMM, add the cluster to SCVMM, check the host migration setting, and ensure that the internal network is not listed.

Enabling Kerberos for Hyper-V

If you are running Windows Server 2012 R2, perform the following procedure to configure Kerberos to secure the storage. You do not have to perform this procedure for Windows Server 2016 because Kerberos is enabled automatically during failover cluster creation.

Before you begin

- Join the hosts to the domain as described in Joining the Cluster and Hosts to a Domain on page 785.
- Verify that you have configured a service account for delegation. For more information on enabling delegation, see the *Microsoft documentation*.

Procedure

1. Log on to the web console by using one of the Controller VM IP addresses or by using the cluster virtual IP address.

2. Click the gear icon in the main menu and select **Kerberos Management** from the **Settings** page.

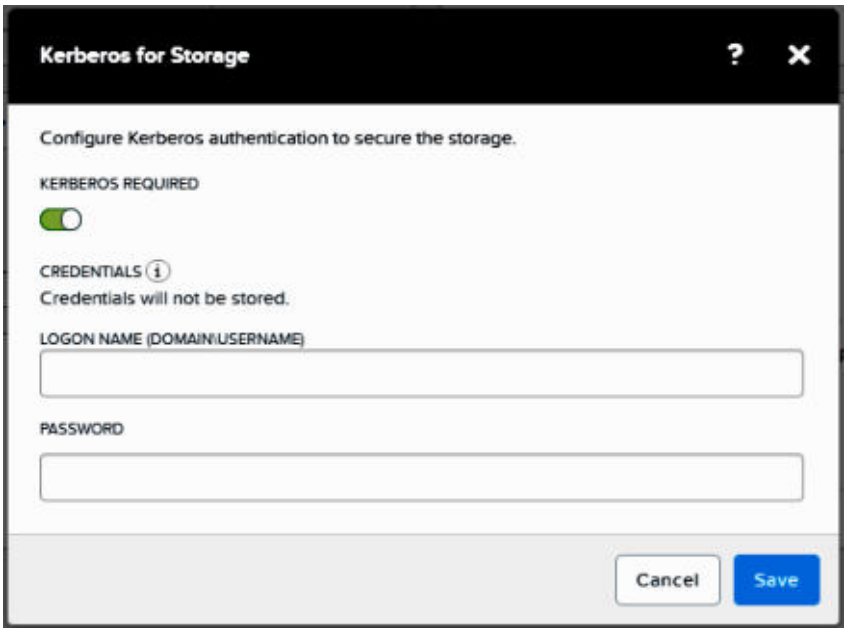

Figure 288: Configure Failover Cluster

3. Set the **Kerberos Required** option to enabled.
4. In the **Credentials** field, type the logon name and password of the domain account that has the privileges to create and modify the virtual computer object representing the cluster in Active Directory. The credentials are required for enabling Kerberos.

   The logon name must be in the format *DOMAIN\USERNAME*. Nutanix does not store the credentials.
5. Click **Save**.

Configuring the Hyper-V Computer Object by Using Kerberos

About this task

Perform the following procedure to complete the configuration of the Hyper-V Computer Object by using Kerberos and SMB signing (for enhanced security).

> Note: Nutanix recommends you to configure Kerberos during a maintenance window to ensure cluster stability and prevent loss of storage access for user VMs.

Procedure

1. Log on to Domain Controller and perform the following for each Hyper-V host computer object.

   a. Right-click the host object, and go to **Properties**. In the **Delegation** tab, select the **Trust this computer for delegation to specified services only** option, and select **Use any authentication protocol**.

   b. Click **Add** to add the **cifs** of the Nutanix storage cluster object.

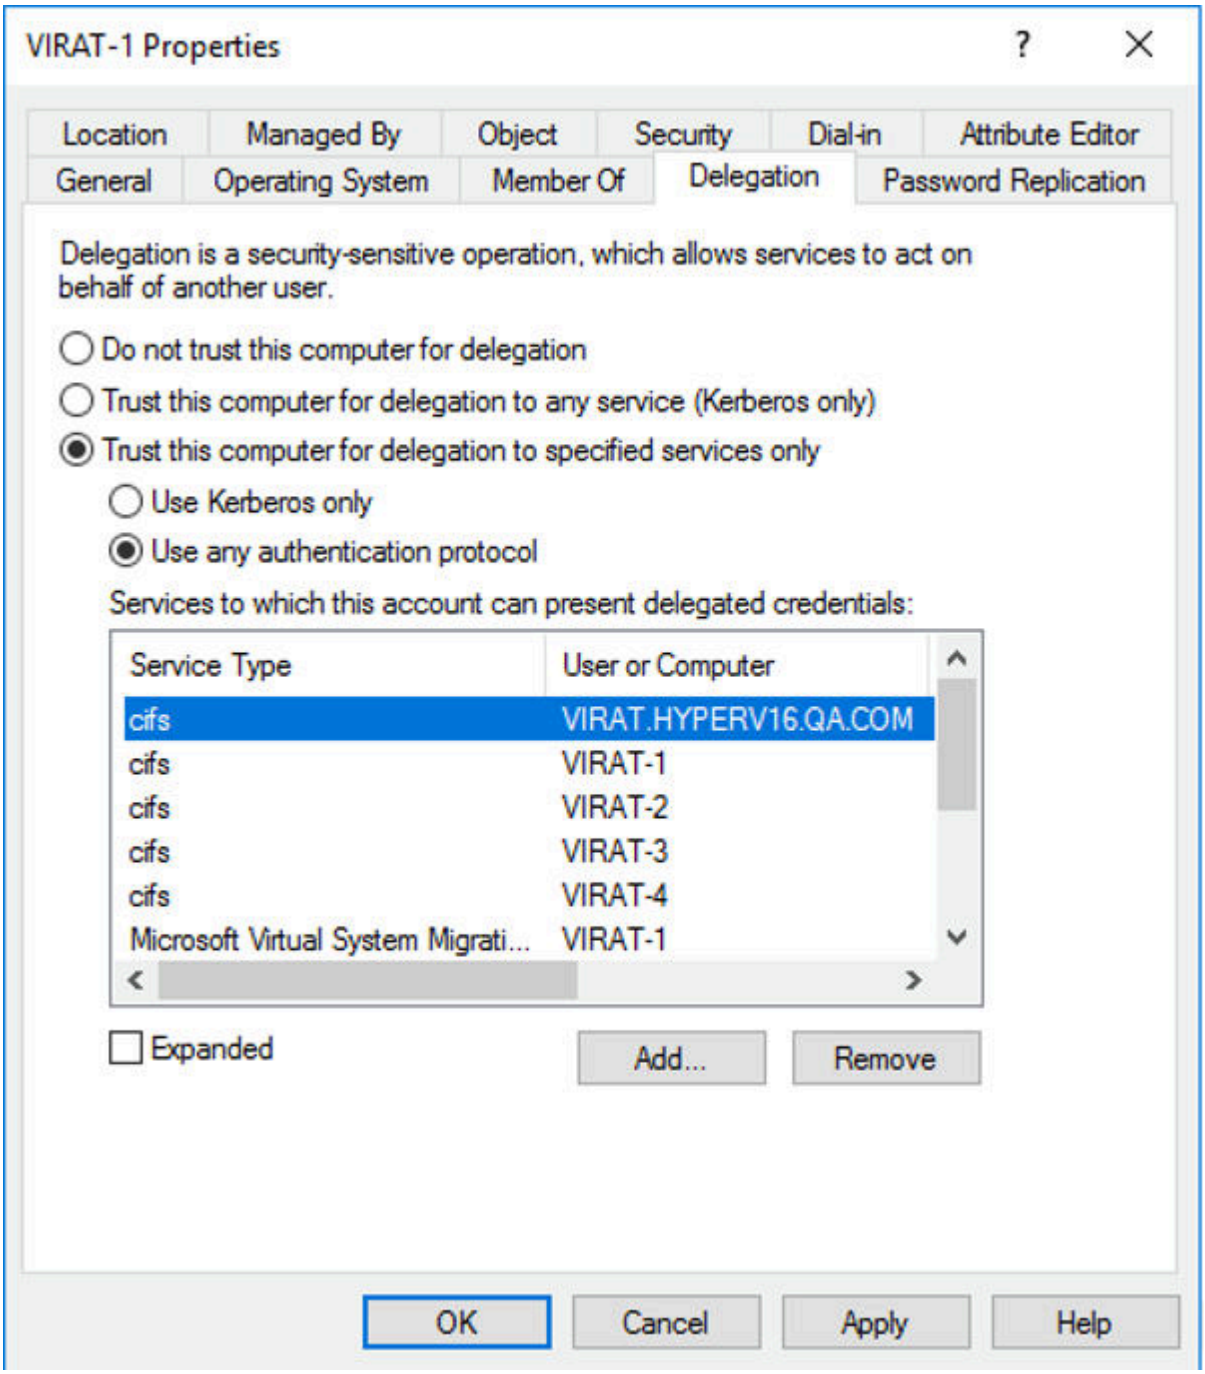

Figure 289: Adding the cifs of the Nutanix storage cluster object

2. Check the Service Principal Name (SPN) of the Nutanix storage cluster object.

```
> Setspn -l name_of_cluster_object
```

Replace *name_of_cluster_object* with the name of the Nutanix storage cluster object.

Output similar to the following is displayed.

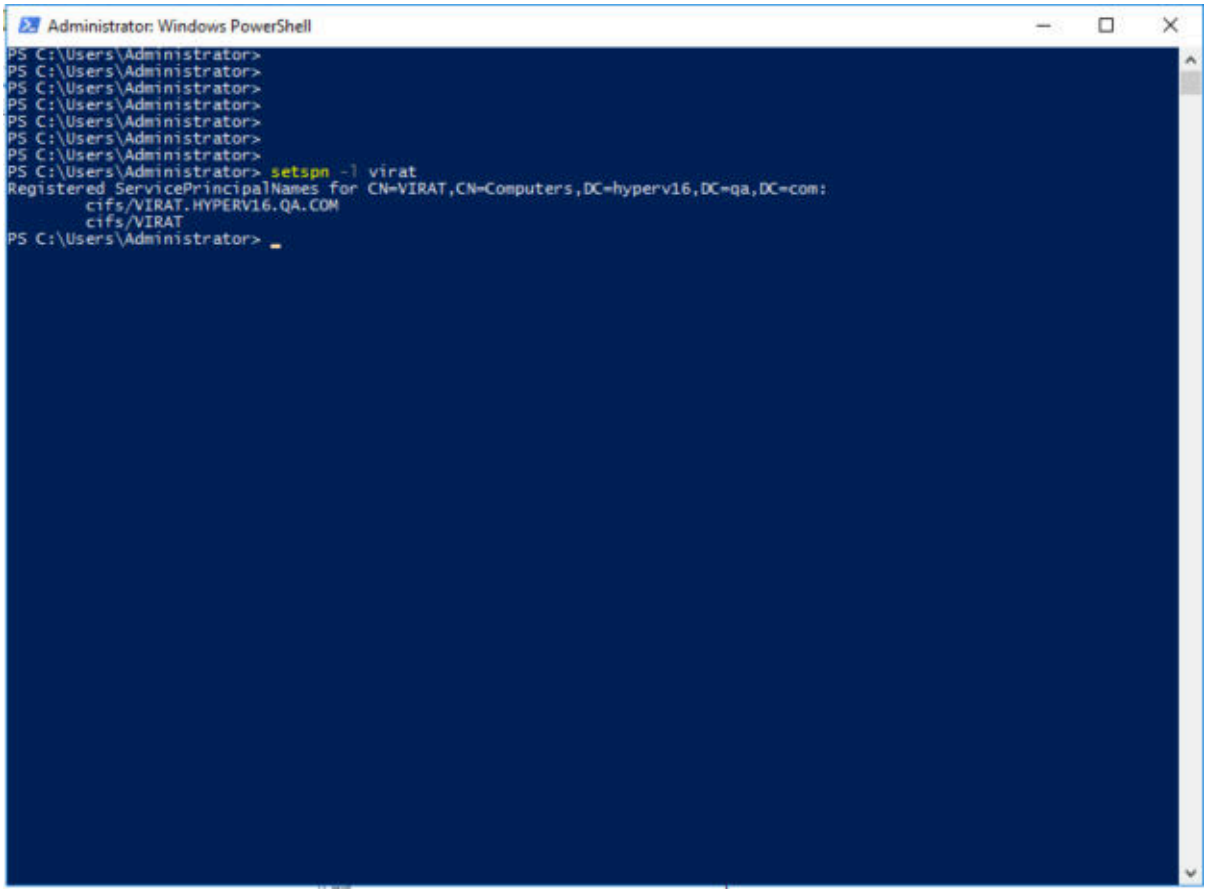

Figure 290: SPN Registration

If the SPN is not registered for the Nutanix storage cluster object, create the SPN by running the following commands.

```
> Setspn -S cifs/name_of_cluster_object name_of_cluster_object
> Setspn -S cifs/FQDN_of_the_cluster_object name_of_cluster_object
```

Replace *name_of_cluster_object* with the name of the Nutanix storage cluster object and *FQDN_of_the_cluster_object* with the domain name of the Nutanix storage cluster object.

Example

```
> Setspn -S cifs/virat virat
> Setspn -S cifs/virat.sre.local virat
```

3. [Optional] To enable SMB signing feature, log on to each Hyper-V host by using RDP and run the following PowerShell command to change the **Require Security Signature** setting to **True**.

```
> Set-SMBClientConfiguration -RequireSecuritySignature $True –Force
```

CAUTION: The SMB server will only communicate with an SMB client that can perform SMB packet signing, therefore if you decide to enable the SMB signing feature, it must be enabled for all the Hyper-V hosts in the cluster.

15

SECURITY AND USER MANAGEMENT

Prism provides several mechanisms to maintain security in a cluster and control user access.

> Note: The following links are to topics in the "Security Management using Prism Element (PE)" chapter in the *Nutanix Security Guide*.

- Set the user authentication method (see “” ())
- Add, edit, or delete local user accounts (see “” ())
- Install or replace an SSL certificate (see “” ())
- Control SSH access to the cluster (see “” ())
- Enable data-at-rest encryption (see “” ())
- Enable network segmentation (see “” ())
- Review authentication best practices (see “” ()) and firewall requirements (see “” ())

16

SUPPORT SERVICES

Nutanix provides support services in several ways.

- Nutanix technical support can monitor your clusters and provide assistance when problems occur (see Controlling Remote Connections on page 800, Configuring HTTP Proxy on page 801, and Configuring Pulse on page 798).
- Nutanix technical support maintains a portal that you can access to request assistance, download AOS updates, or view documentation (see Accessing the Nutanix Support Portal on page 804).
- Nutanix supports a REST API that allows you to request information or run administration scripts for a Nutanix cluster (see Accessing the REST API Explorer on page 807).

Pulse Health Monitoring

The feature known as Pulse provides diagnostic system data 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 a Nutanix cluster. Information includes:

- system alerts
- current Nutanix software version
- Nutanix processes and Controller VM information
- hypervisor details such as type and version

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 Pulse Access Requirements on page 800). For the complete list of required ports, see Port Reference.

> Note: When logging into Prism the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a screen appears to enable Pulse (see *Logging Into the Web Console* topic in the *Prism Web Console Guide*).

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" in the *Prism Central Guide*), 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 HTTP Proxy on page 801). 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 Pulse Access Requirements on page 800). 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 796 for Pulse configuration recommendations.
- Pulse sends messages through ports 80/8443/443. For the complete list of required ports, see Port Reference.
- When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse. To enable Pulse, click the **Continue** button in the message and follow the prompts; to continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

You can enable (or disable) Pulse at any time. To configure Pulse, do the following:

Procedure

1. Go to the Settings menu (see *Settings Menu* in the *Prism Web Console Guide*) and select **Pulse**.

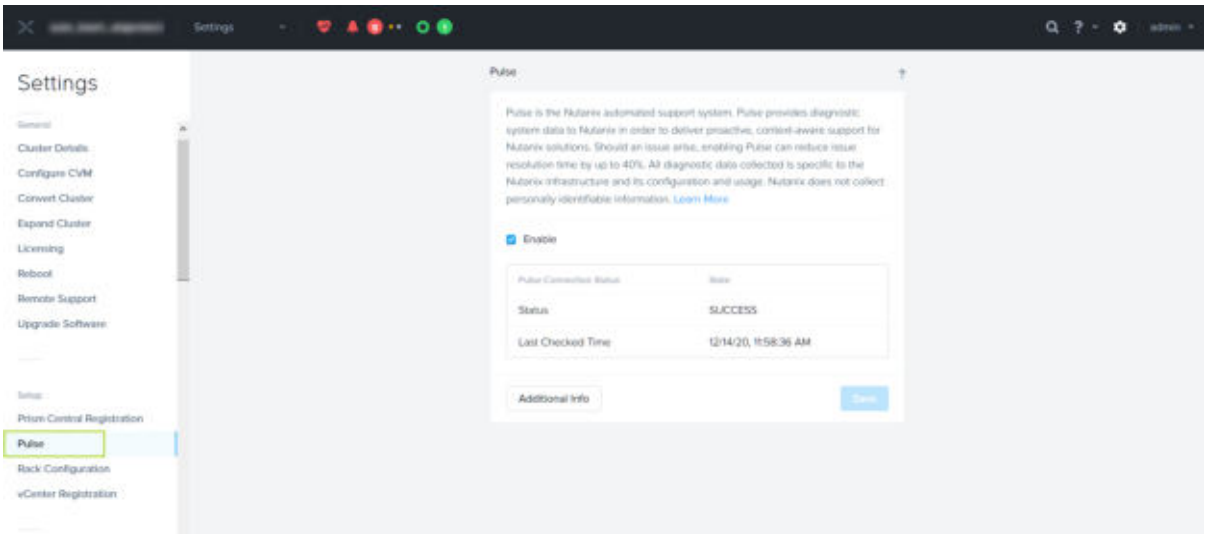

Figure 291: Pulse Window

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 to display additional connection and recipient information.

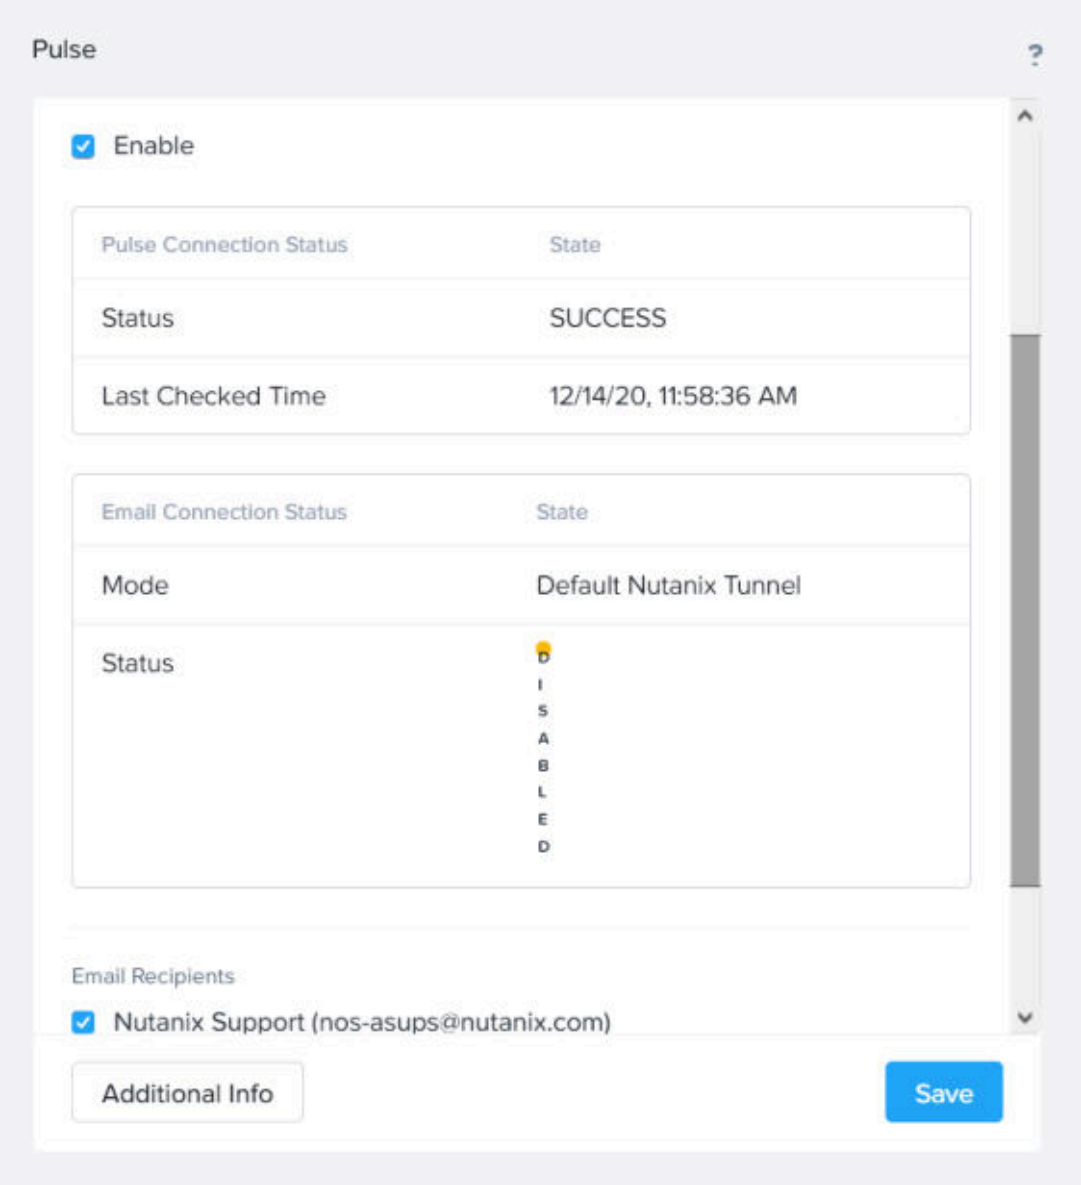

Figure 292: Pulse Additional Info

4. Check the **Pulse Connection Status** field for connection information.

   - **Status**: Displays the transport mechanism status, which is **Success** when the HTTP client can access the Pulse server successfully or **Failure** when it cannot (or **Unknown** when the status is not known).
   - **Last Checked Time**: The last time Pulse sent data and checked the connectivity with the Pulse server.

5. Check the **Email Connection Status** field for connection 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 **Disabled** when neither is available (or **Unknown** when the status is not known).

     If the status is **Success** the following two fields also appear.
   - **Service Center**: Displays the address of the Nutanix support service center to which the Pulse data is sent.
   - **Established Since**: Displays the date when the current Pulse session was established.

6. To send the cluster information to additional recipients, enter one or more e-mail addresses in the **Comma Separated Email List** field.

   You can e-mail the information to yourself or others by entering recipient e-mail addresses in this field. Use a comma to separate multiple addresses.

7. Click the **Save** button to save the new settings.

Pulse Access Requirements

In order to successfully send Pulse and alert messages from a cluster to the Nutanix support servers, Pulse requires the following access:

- Messages are sent from the Zeus (cluster configuration manager) leader, so the firewall must allow the Zeus leader IP address. Because the Zeus leader can change, it is recommended that the IP addresses for all Controller VMs in the cluster be open in the firewall.
- Required firewall rules configured to allow traffic from the cluster to Nutanix support servers. See Firewall Requirements in the *Security Guide*.

Pulse data is collected through the following methods:

- By default, emails are sent daily through port 8443 via nsc01.nutanix.net or nsc01.nutanix.net to nos-asups@nutanix.com.
  - If a SMTP server is configured on a cluster, emails will be sent via the SMTP server to nos-asups@nutanix.com.
- Throughout the day, each Controller VM in the cluster sends Pulse data to insights.nutanix.com over port 443 using the HTTPS REST endpoint.

> Note:
>
> Firewall port requirements are automatically met if you have a Prism Central deployment, see "Prism Central Proxy for Pulse Data" in the *Prism Central Guide*.

Controlling Remote Connections

About this task

Nutanix technical support can remotely assist with problems by logging into the cluster through an SSH connection. This remote connection is disabled by default, but it may be enabled during the site installation procedure or at a later date through the web console.

You must configure rules in your external firewall to allow traffic from the cluster to Nutanix remote support servers. See Firewall Requirements in the *Security Guide*.

To enable (or disable) Nutanix technical support remote access to your cluster through this connection, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Remote Support** in the **Settings** page. The Remote Support dialog box appears.
2. Select (click the radio button for) the desired access state.
   - To allow remote access (temporarily) for Nutanix support, select **Enable** and enter the number and duration (hours or minutes) in the two boxes. When remote support is enabled temporarily, an icon appears in the screen header that includes a dynamic countdown until remote support is disabled.

   - To prevent remote access by Nutanix support, select **Disable**.
3. Click the **Save** button to save the new setting and close the window.

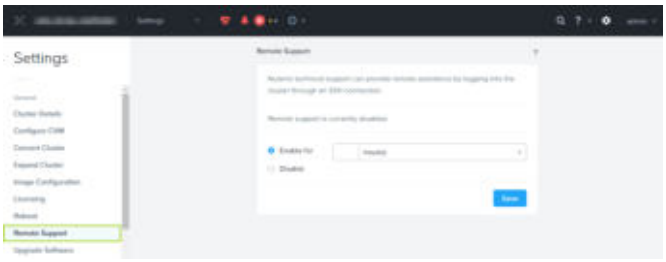

Figure 293: Remote Support Window

Configuring HTTP Proxy

About this task

If the customer site 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 Proxy dialog box appears.

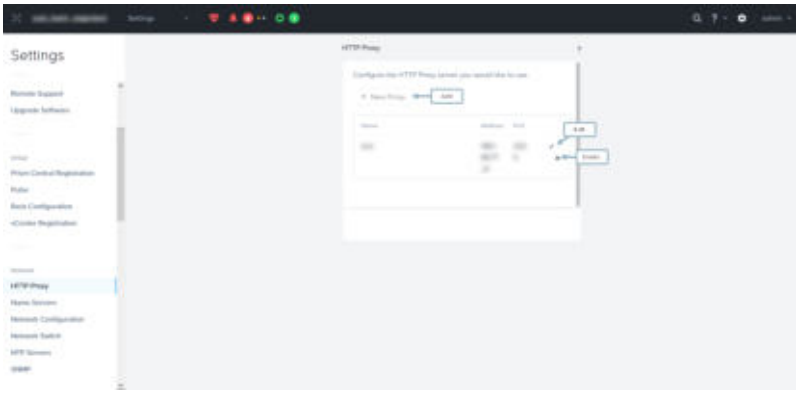

Figure 294: HTTP Proxy Window

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**: Check the box(es) for the protocol to proxy (HTTP, HTTPS, or both).

   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.

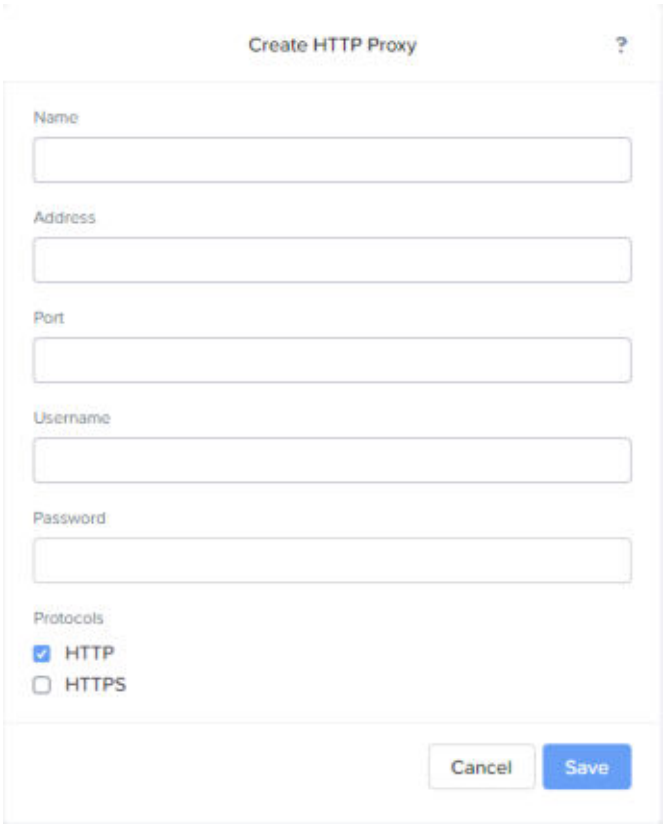

Figure 295: Create HTTP Proxy Window

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.

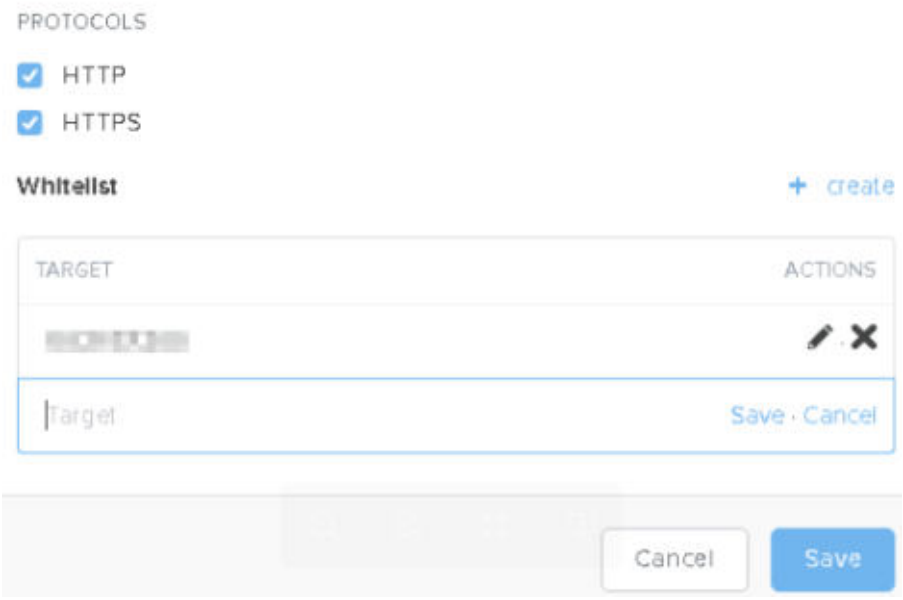

Figure 296: Whitelist Targets

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.

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, select **Support Portal** from the question mark icon

pull-down list.
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 1015: Main Menu Options

Category	Option	Description
Documentation	Product Documentation	Displays a page from which you can view the Nutanix product 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.

Category	Option	Description
	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.
	Acropolis Upgrade Paths	Displays a table of 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 & Forums	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.
	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) and Prism Central VM.
	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.

Nutanix REST API

The Nutanix REST APIs allow you to create scripts that run system administration commands against the Nutanix cluster. The API enables the use of HTTP requests to get information about the cluster as well as make changes to the configuration. Output from the commands are returned in JSON format.

There are two versions of the Nutanix REST API.

- v1: The original Nutanix REST API.
- v2: An update of the v1 API. Users of the v1 API are encouraged to migrate to v2.

A complete list of REST API functions and parameters is available in the REST API Explorer section of the Nutanix web console. In addition, the complete reference for the v2 Nutanix API, including code samples in multiple languages, and tutorials are available at http://developer.nutanix.com/.

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 v1 and v2 APIs can both be viewed in the REST API Explorer.

Procedure

1. Open the explorer for the desired version of the API.

   » v1: Connect to the Nutanix web console, click the user icon in the upper-right corner of the web console, and select **REST API Explorer**. In the explorer, select **Version 1** from the menu.

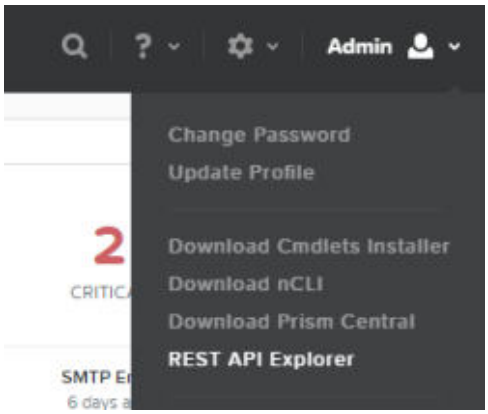

Figure 297: User menu

   » v2: Connect to the Nutanix web console, click the user icon in the upper-right corner of the web console, and select **REST API Explorer**. In the explorer, select **Version 2** from the menu.

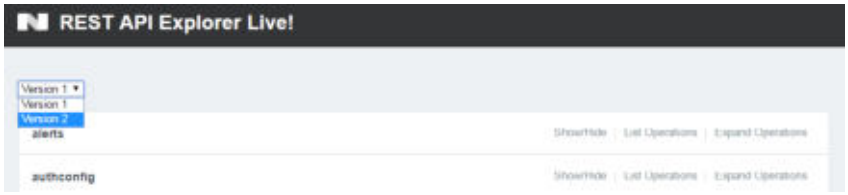

Figure 298: API version selection

The REST API Explorer displays a list of the cluster 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

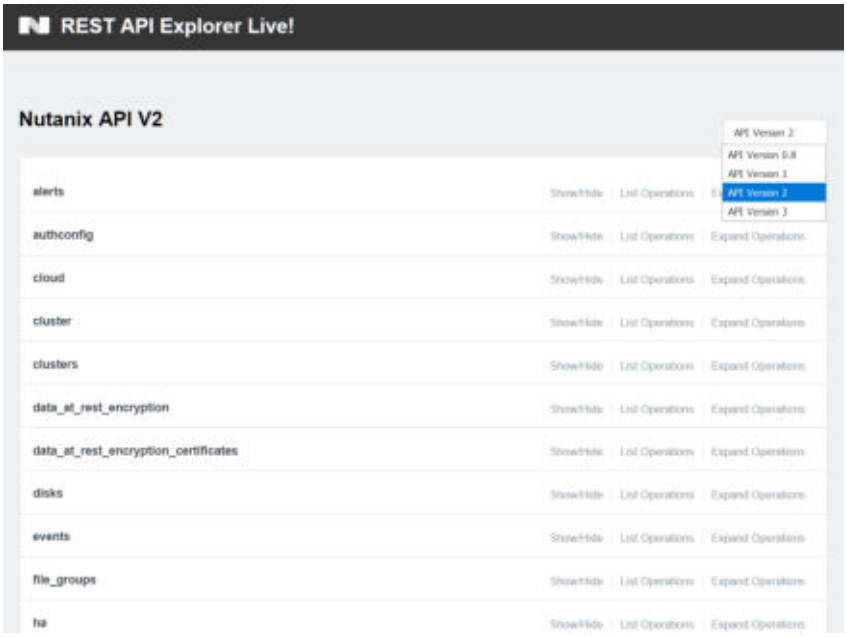

Figure 299: REST API Explorer

> Tip: The objects are listed by a relative path that is appended to the base URL https://*management_ip_addr*:9440/PrismGateway/services/rest/v[1,2,3]/api, where *management_ip_addr* is the IP address of any Nutanix Controller VM in the cluster.

2. Find the line for the object you want to explore and click **Expand Operations**. For this example, you will operate on a storage pool.

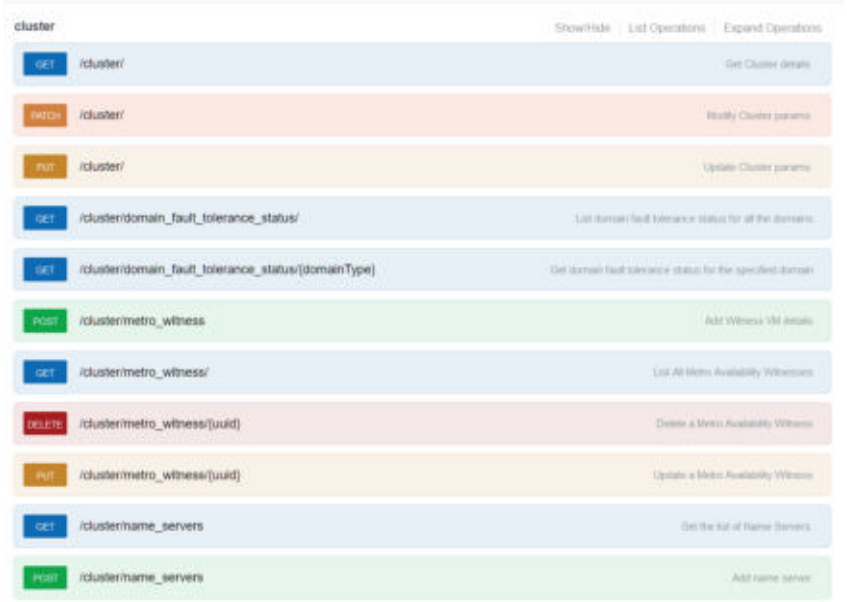

Figure 300: Cluster Operations

3. Click **GET** on the first line to show the details for this API call.

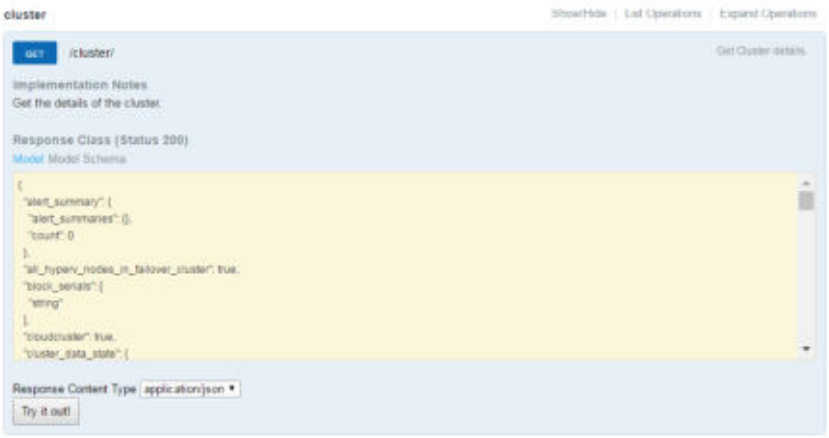

Figure 301: Get Cluster Details

The explorer displays the parameters that can be passed when this action is used.

4. Click **Try it out!** to test the API call when used with your cluster.

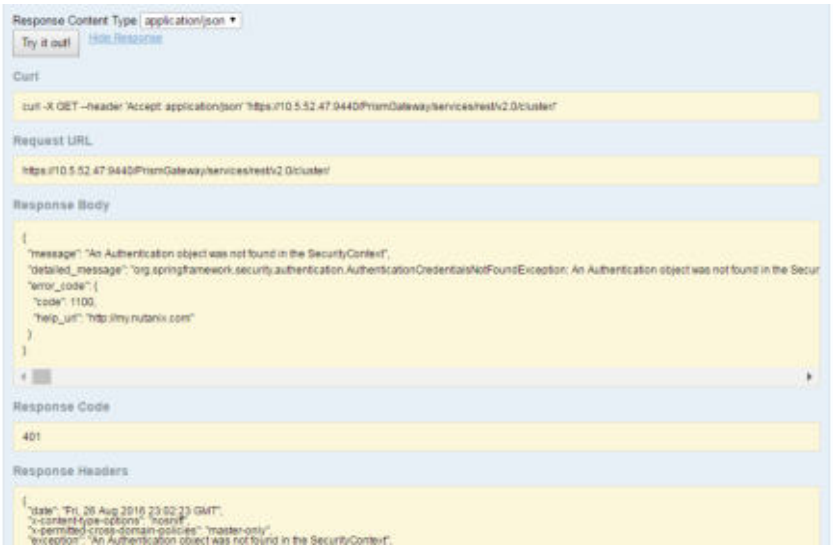

Figure 302: Get Cluster Details Response

The test displays the request URL required for scripting, as well as sample responses expected from the API call.

17

HELP RESOURCES

There are several information sources that you can access at any time when you need help:

- Context-sensitive help documentation (see Accessing Online Help on page 811).
- Health dashboard tutorial (see *Health Dashboard* in the *Prism Web Console Guide*).
- Customer support portal (see Accessing the Nutanix Support Portal on page 804).
- Nutanix community forum (see Accessing the Nutanix Next Community on page 812).
- REST API explorer (see Accessing the REST API Explorer on page 807).
- Glossary of terms (see Glossary on page 812).

Accessing Online Help

About this task

The web console 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 (see Main Menu Options in Cluster Management in the *Prism Web Console Guide*):

   » 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 **General Help** 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.

2. To select a topic from the table of contents, click the 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.

3. To display all the help contents as a single document (*Web Console Guide*), click the **epub** or **pdf** button in the upper right.

   You can view the *Web Console 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 **Other** 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.

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, select **Nutanix Next Community** from the question mark icon pull-down list of the main menu (see Main Menu Options in Cluster Management in the Prism Web Console Guide).
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.

Glossary

aCLI

Acropolis command-line interface.

Acropolis

The Nutanix converged software fabric for virtualization and storage management. It consists of the Acropolis base software, Acropolis Distributed Storage Fabric, AHV, App Mobility Fabric, Prism, and Acropolis APIs.

Acropolis App Mobility Fabric

Provides virtualization management, volume management, and other distributed management functions for Acropolis.

AHV

Nutanix-specific hypervisor host providing core server virtualization and optimized for Nutanix cluster and guest VM use.

block

A set of Nutanix nodes contained in a single enclosure.

block fault tolerance

When certain conditions are met, redundant copies of any data required to serve I/O are placed on nodes that are not in the same block. In the case where multiple nodes in a block fail, the cluster can continue to run because cluster configuration data exists on other blocks. Guest VMs can continue to run because redundant copies of guest VM data and metadata exist on other blocks.

clone

A writeable copy of a vDisk.

Cloud Connect

A feature that enables you to back up and restore copies of virtual machines and files to and from an on-premise cluster and a Nutanix Controller VM located on the Amazon Web Service (AWS) cloud. The cloud-based cluster is managed as a remote site from the Prism Web Console or nCLI.

cluster

A group of nodes contained in one or more Nutanix blocks.

cold data

Data that did not have many rewrites or overwrites for a long time. For example, snapshots, file servers, archives, or backups.

compression

An optional storage container setting that enables one of two types of compression.

compression, inline

Data is compressed as it is written. This type of compression is recommended for workloads that perform batch processing.

compression, post-process

Data is compressed after it is written. The delay time between write and compression is configurable. Nutanix has recommends a delay value of 60. However, every workload has a different I/O profile, so you might require a different delay value otherwise. This type of compression is recommended for most workloads.

Consistency Group

A subset of VMs in a Protection Domain. All VMs within a Consistency Group for that Protection Domain will be snapshotted in a crash-consistent manner. For all VMs in a Consistency Group, a snapshot creates one snapshot for all VMs in the group.

Controller VM

A Nutanix VM that manages storage and other cluster functions on a node.

data resiliency

A measure of the types of failures a cluster can withstand; affected by block awareness and redundancy factor.

datastore

A logical storage container for files necessary for VM operations.

deduplication

The sharing of identical guest VM data on premium tiers (RAM and Flash) for improved performance or on capacity tiers (HDD) for storage space savings. Enabled by properties of a storage container or vDisk.

Distributed Storage Fabric

All storage functionality, including snapshots and clones, data protection, disaster recovery, data path redundancy, replication factors, deduplication, compression, erasure coding, and so on.

Enterprise Cloud Platform

The Nutanix solution that natively converges compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence

erasure coding

Optional algorithm included in the Acropolis base software to help reduce the storage used for fault tolerance. It helps to increase the effective or usable capacity on a cluster, depending on cluster size and data coldness.

failback

Planned failover initiated from recovery site.

failover

Moving VMs from a primary site to a recovery site.

failover, disaster

Failover when the primary site is down.

failover, planned

Failover when both sites are up.

guest VM

A VM running on a Nutanix cluster that executes a workload, such as VDI or Exchange, as opposed to a VM that is involved in cluster operations, such as a Controller VM.

host

An instance of the hypervisor that runs on a Nutanix node.

image service

A workflow in the Prism web console that enables a user to upload ISO or disk images (in ESXi or Hyper-V format) to a Nutanix AHV cluster by specifying a remote repository URL or by uploading a file from a local machine.

local replication

Multiple copies of data within a storage container. These copies exist for fault tolerance: if a physical disk fails, the cluster can recover data from another copy. The cluster manages the replicated data, and the copies are not visible to the user.

local snapshot

Snapshots stored on the same cluster where they originated.

nCLI

Nutanix command-line interface.

node

A physical server contained in a Nutanix block; runs a hypervisor host.

oplog

A write cache on a faster, more expensive storage tier.

Prism

Web-based management interface for managing Nutanix clusters.

Prism Central

Centralized management tool that runs as a separate VM configured as a single-node cluster to monitor and manage multiple clusters through a single web console.

Prism Element

A single cluster being managed by and available through the Prism Central web console.

Protection Domain

A group of VMs to be backed up locally on a cluster or replicated on the same schedule to one or more clusters. Protection Domains may be associated with remote sites. Protection Domain names must be unique across sites. A VM can be in at most one Protection Domain.

Protection Domain, active

A Protection Domain that manages live VMs and makes, replicates, and expires snapshots.

Protection Domain, inactive

A Protection Domain that receives snapshots from a remote cluster.

redundancy factor

The number of nodes plus 1 that the cluster can tolerate being down at one time. By default, Nutanix clusters have a redundancy factor of 2, which means that they can tolerate 1 node being down. They are configurable to redundancy factor 3 to enable tolerating 2 nodes being down.

remote replication, one-to-one

Replicates a production cluster with one idle cluster as a DR target.

remote replication, reciprocal

Cross replication within running (non-idle) production clusters.

remote site

A pair of clusters that can replicate data to each other. A configured remote site can also be located in the cloud (based on Amazon AWS, for example).

remote snapshot

A snapshot copied asynchronously from one cluster to another.

replication factor

The number of copies of data kept by a storage container. The default is 2. Storage Containers on clusters with redundancy factor of 3 can have replication factor of 3.

reserved capacity

A property of a storage container or vDisk that guarantees that a certain amount of storage space is available.

schedule

A property of a Protection Domain that specifies the intervals to take snapshots and how long the snapshots should be retained. A schedule optionally specifies which remote site to replicate to.

Self-Service Restore

Allows a user to restore a file within a virtual machine from the Nutanix protected snapshot with minimal Nutanix administrator intervention.

shadow clone

A cache of a vDisk on all the nodes in the cluster. When a vDisk is read by multiple VMs (such as the base image for a VDI clone pool), the cluster creates shadow clones of the vDisk.

snapshot

A read-only copy of the state and data of a VM at a point in time. Snapshots for a VM are crash consistent, which means that the VMDK on-disk images are consistent with a single point in time. That is, the snapshot represents the on-disk data as if the VM crashed. The snapshots are not, however, application consistent, meaning that application data is not quiesced at the time of snapshot, unless the Protection Domain is configured to use application-consistent snapshots.

storage container

A subset of available storage within a storage pool.

storage pool

A group of physical disks from one or more tiers.

storage replication adapter

A Nutanix-provided module that allows VMware Site Replication Manager (SRM) to use native remote replication.

tier

A type of physical storage in a Nutanix node.

vDisk

Data associated with a VM represented as a set of files on a datastore.

VM high availability

In virtualization management VM high availability, when a node becomes unavailable, VMs that are running on that node are restarted on another node in the same cluster.

VM mobility

The ability to export your existing VMs from one non-AHV cluster to an AHV cluster. This option requires that you install the Nutanix VM Mobility installer on all the VMs. The Nutanix VM Mobility installer deploys the drivers that are required at the destination AHV cluster. After you prepare the source VMs, they can be exported to the AHV cluster.

vStore

A separate mount point within a storage container which has its own NFS namespace. This namespace maps to a Protection Domain. Each vStore is exported as a device through the Nutanix SRA.

COPYRIGHT

Copyright 2022 Nutanix, Inc.

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

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

License

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

Conventions

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

Version

Last modified: March 21, 2022 (2022-03-21T11:52:09+05:30)

BOOKLIST

Related Guides

Refer to the following guides for more information:

- *Acropolis Upgrade Guide*
- *Security Guide (AOS Security 5.18)*
- *Data Protection and Recovery with Prism Element*