Prism Central Guide

Prism pc.2022.6

January 11, 2024

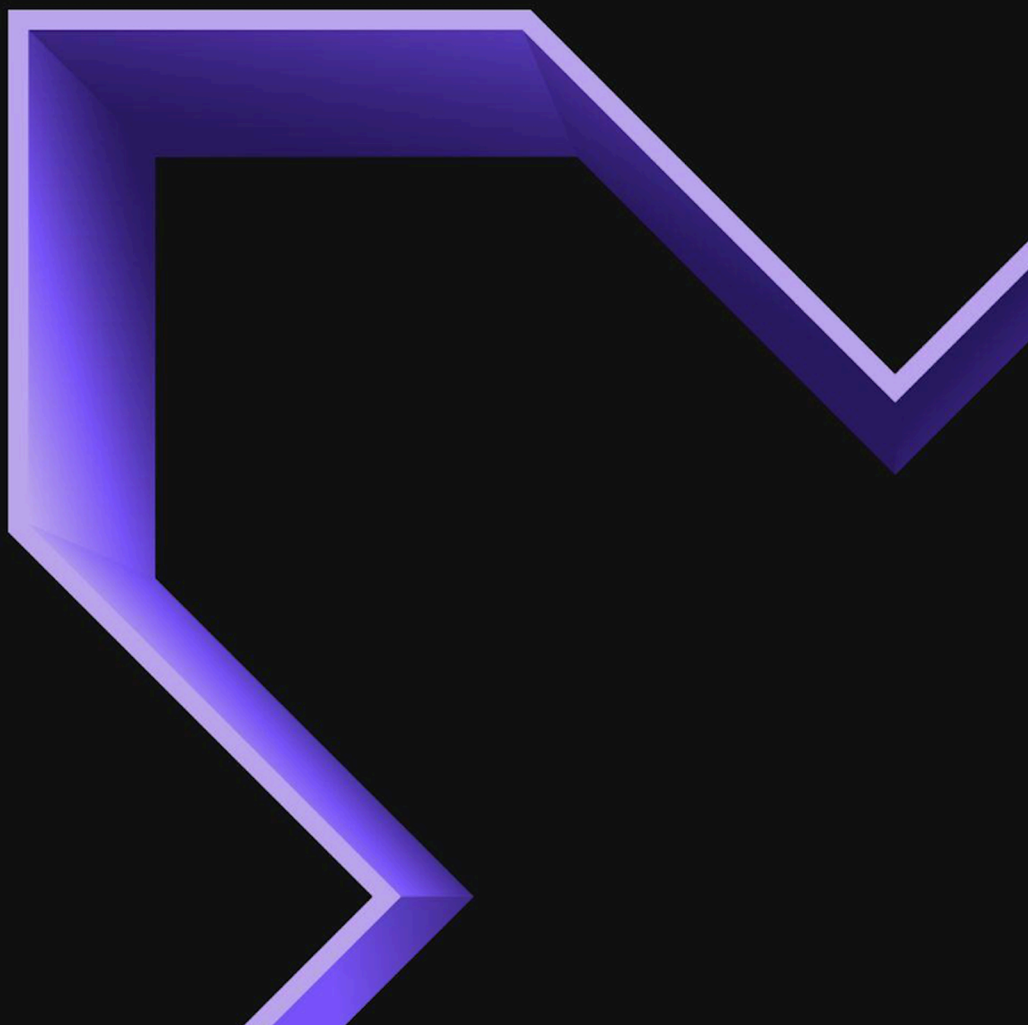

NUTANIX

Contents

Welcome 8

Help Organization 9

Prism Central Overview 10
- Main Menu (Prism Central) 10
- Entities Menu 13
- Settings Menu (Prism Central) 16
- Searching for Information 19
- Keyboard Shortcuts in Prism Central 28
- Understanding Displayed Statistics (Prism Central) 28

Prism Central Administration 30
- Installing or Upgrading 30
  - Prism Central Upgrade vDisk 30
- Licensing 31
- Prism Central Management 31
  - Logging Into Prism Central 31
  - Logging Out of Prism Central 34
  - Managing Prism Central 34
  - Finding the Prism Central Version 45
  - Finding the AHV Version on Prism Central 46
  - Finding the AOS Version Using Prism Central 47
  - Registering or Unregistering a Cluster with Prism Central 48
  - Expanding (Scale Out) Prism Central 56
  - Shutting Down (Starting Up) Prism Central VM 60
  - Prism Central Disaster Recovery 61
  - IP Address Reconfiguration 65
- Pairing AZs (Nutanix Disaster Recovery) 68

Main Dashboard 70
- Main Dashboard Display 72
- Widgets on a Dashboard 74
- Creating a New Dashboard 79
- Adding Dashboard Widgets 80
- Modifying a Dashboard 85
- Sharing a Dashboard 86

Entity Exploring 89
- Compute and Storage Entities 93
  - VMs Summary View 94
  - VM Template Summary View 118
  - OVAs View 122
  - Images Summary View 125

Catalog Items View..... 129
Storage Containers Summary View..... 130
Storage Policies Summary View..... 141
Volume Groups Summary View..... 147
vCenter Datastores Summary View..... 156
Catalog Items View..... 163
Categories Summary View..... 164
OVAs View..... 167
Network and Security View..... 170
Subnets..... 170
Virtual Private Clouds..... 174
Floating IPs..... 179
Connectivity..... 179
Security Policies Summary View..... 187
Policies..... 189
Affinity Policies Summary View..... 190
NGT Policies View..... 192
Image Placement Policies Summary View..... 193
Bandwidth Throttling Policies Summary View..... 194
Security Policies Summary View..... 195
Data Protection and Recovery Entities..... 198
Protection Policies View..... 198
Recovery Plans View..... 199
VM Recovery Points View..... 200
Hardware Entities..... 201
Clusters Summary View..... 202
Hosts Summary View..... 221
Disks Summary View..... 230
GPUs Summary View..... 237
Activity Entities..... 240
Audits Summary View..... 240
Tasks View..... 244
Operations Entities..... 246
Reports View..... 246
Administration Entities..... 248
Users Summary View..... 248
Roles Summary View..... 252
Projects Summary View..... 256
Availability Zones View..... 266
LCM View..... 267
Services Entities..... 267

Alert and Event Monitoring (Prism Central)..... 269
Alerts Summary View (Prism Central)..... 269
Alert Details..... 276
Events Summary View (Prism Central)..... 279
Event Details..... 280
Configuring Alert Emails (Prism Central)..... 281
Configuring Alert Policies (Prism Central)..... 285
Modifying System Alert Policies..... 286
Adding Custom Alert Policies..... 287
Viewing External Alert Policies..... 302
Prism Central Logs..... 302
Audit Log Events..... 305
Alerts/Health checks..... 307

Cluster....307
Controller VM....359
DR....371
License....429
Guest VM....430
Hardware....439
Node....442
Network....447
Nutanix Files....457
Other....482
Prism Central VM....494
Storage....494
System Indicator (Prism Central)....506
VM....506

System Management....507
Configuring Prism Central When a Cluster Uses Proxy Servers....507
Who Needs to Use the Whitelist Method....508
Configuring an HTTP Proxy....509
Whitelisting Prism Central and Its Managed Clusters (nCLI)....513
Configuring Name Servers (Prism Central)....514
Configuring NTP Servers (Prism Central)....515
Configuring an SMTP Server (Prism Central)....516
Configuring SNMP (Prism Central)....517
Configuring Syslog Monitoring....523
Syslog Modules....524
Modifying UI Settings (Prism Central)....527
Configuring a Banner Page (Prism Central)....531
Internationalization (i18n) (Prism Central)....531
Localization (L10n) (Prism Central)....532
Changing the Language Settings (Prism Central)....532

Security and User Management (Prism Central)....534

Compute and Storage (Cluster) Administration....535
VM Management....535
Creating a VM through Prism Central (AHV)....535
Creating a VM (Self Service)....545
Creating a VM from Catalog Items (Self Service)....546
Managing a VM through Prism Central (AHV)....552
Managing a VM (Self Service)....566
Adding Multiple vGPUs to the Same VM....573
Live Migration of vGPU-enabled VMs....575
Creating a VM through Prism Central (ESXi)....580
Managing a VM through Prism Central (ESXi)....585
Nutanix Guest Tools....590
Storage Quality of Service (QoS)....601
Memory Overcommit Management....604
VM Policy Management....608
VM Template Management....613
Limitations of VM Template Feature....613
Creating a VM Template....614
Deploying VM from a Template....615

Managing a VM Template....618
Storage Management (Prism Central)....620
Storage Components....621
Storage Container Management....628
Creating a Volume Group....644
Modifying a Volume Group....648
Managing Volume Group Connections....649
Managing Volume Group Virtual Disks....653
Cluster RBAC for Volume Group....654
Storage Policy Management....657
Storage Policy Based Encryption....659
Creating or Updating a Storage Policy....660
Managing Associations....662
Deleting a Storage Policy....664
Catalog Management....666
Adding a Catalog Item....667
Deleting a Catalog Item....669
Image Management....670
Requirements....670
Limitations....671
Creating an Image....671
Modifying an Image....688
Importing Images to Prism Central....689
Image Policy Management....690
OVA Management....698
Exporting a VM as an OVA....699
Uploading an OVA....700
Concatenating Upload using APIs....706
Deploying an OVA as VM....709
Downloading an OVA....715
Renaming an OVA....716
Deleting an OVAs....716
Performing Other Administrative Tasks....717
vCenter Server Integration....718
Registering vCenter Server (Prism Central)....719
Managing vCenter Server Registration Changes (Prism Central)....722

Prism Self Service Administration....723
Prism Self Service Overview....723
Configuring Prism Self Service....724
Projects Overview....726
Creating a Project....729
Adding Users to a Project....733
Infrastructure in Projects....736
Modifying a Project....744
Deleting a Project....745
Environments in Projects....745
Quota Policy Overview....773
Creating a Snapshot Policy....777

Network and Security Management....782
Network Connections....782
Configuring Network Connections....784
Creating or Updating Virtual Switch....787

Enabling LACP and LAG (AHV Only)....791
Modifying Network Connections....793
Security Policies....794
Flow Networking....794

Category Management....795
Creating a Category....795
Modifying a Category....796
Assigning a Category....797

Operations Management....799
Performance Monitoring....799
Analysis Dashboard (Prism Central)....799
Sessions Dashboard....824
Closing the Session....831
Switching Between Sessions....831
Application Discovery....831
Application Discovery View....832
Enabling Application Discovery....841
Connecting to the Cloud....847
Adding (Removing) Clusters to Monitor....848
Configuring vCenter Authentication....850
Discovering Applications....851
Publishing Applications....852
Creating Discovery Policies....852
Cost Management (Xi Beam)....855
Application Monitoring....856
Application Monitoring Summary View....856
Configuring Application Monitoring....868
Modifying Application Monitoring....872
Operations Policy Management....873
Operations Policies View....873
Modifying an Operations Policy....874
Resource Planning....876
Capacity Runway View....877
Scenarios View....878
Updating Capacity Configurations....898
Behavioral Learning Tools....899
Task Automation....902
Overview....902
Action Gallery....903
Plays....903
Playbook Actions....903
X-Play Integrations....912
Creating Playbooks using Triggers....922
Running a Playbook (Manual Trigger)....932
Configuring Manual Parameters....932
Exporting or Importing Playbooks....938
PagerDuty Integration with Prism....938
Reports Management....941
Creating a New Report....943
Managing a Report....944
Viewing Report Instances....947
Views in the Reports....957

Scheduling a Report....986
Downloading Reports....987
Configuring Report Settings....988
Role Based Access Control....990

Services Enablement....992
Enabling Calm....992
Enabling Files....992
Enabling Foundation Central....993
Enabling Karbon....993
Enabling Objects....993

Customer Support Services....995
Pulse Health Monitoring....995
Configuring Pulse....997
Pulse Health Monitoring Data Collection....999
Network Configuration for Pulse Health Monitoring....1005
Prism Central Proxy for Pulse Data....1005
Creating a Support Case....1006
Viewing Case Status....1010
Accessing the Nutanix Support Portal (Prism Central)....1011
Accessing the REST API Explorer....1014

Help Resources....1016
Accessing Online Help (Prism Central)....1016
Accessing the Nutanix Next Community....1018
Glossary....1018

Copyright....1024

WELCOME

You can monitor and manage entities across Nutanix clusters through a web console called *Prism Central*. This documentation describes how to use Prism Central.

- For a list and description of major topics, see Help Organization on page 9.
- For help with viewing this documentation, see Accessing Online Help (Prism Central) on page 1016.
- To access other Nutanix documents, see Nutanix support portal.

HELP ORGANIZATION

This documentation is organized as follows:

- Prism Central Overview on page 10 provides an overview of Prism Central including menu options and navigation tools.
- Prism Central Administration on page 30 describes how to install and manage Prism Central; how to upgrade Prism Central, NCC, and individual clusters through Prism Central; and how to manage licenses.
- Main Dashboard on page 70 describes the main (home) dashboard and how to customize it.
- Entity Exploring on page 89 describes how to use the entity browser to view detailed information about VMs, containers, disks, and other entities across the registered clusters.
- Resource Planning on page 876 describes how to review and analyze current and potential resource needs in a Nutanix cluster.
- Performance Monitoring on page 799 describes how to monitor and analyze performance in a cluster.
- Services Enablement on page 992 describes how to enable certain applications (Calm, Karbon, and Objects) through Prism Central.
- Alert and Event Monitoring (Prism Central) on page 269 describes how to monitor alerts and events across the registered clusters and how to configure alert policies and notification.
- System Management on page 507 describes how to configure various system settings (NTP, SMTP, and so on) for Prism Central.
- Security and User Management (Prism Central) on page 534 describes how to configure user authentication, manage local user accounts, control user access using roles, install SSL certificates, and control remote (SSH) access.
- Compute and Storage (Cluster) Administration on page 535 describes how to administer individual clusters through Prism Central such as how to create and manage VMs.
- Prism Self Service Administration on page 723 describes how to implement self-service features through Prism Central.
- Policies on page 189 describes how to create and apply policies to manage and control access and operations in a cluster.
- Reports Management on page 941 describes how to create and run reports about infrastructure resources.
- Task Automation on page 902 describes how to automate routine administrative tasks through Prism Central by using the X-play feature.
- Customer Support Services on page 995 describes how to configure Pulse (system diagnostics reporting feature), access the Nutanix customer support portal, and create a support ticket through Prism Central when you need help.
- Help Resources on page 1016 describes how to use the online help to find what you need and how to access the Nutanix Next Community portal.

PRISM CENTRAL OVERVIEW

Nutanix provides an option to monitor and manage multiple clusters through a single web console. This multi-cluster view, known as Prism Central, is a centralized management tool that runs as a separate instance comprised of either a single VM or a set of VMs. Prism Central provides the following features:

- Single sign on for all registered clusters.
- Customizable main dashboard that displays summary information across the registered clusters (see Main Dashboard on page 70).
- Virtual infrastructure dashboards with drill-down options to manage VMs, storage containers, catalog items, images, categories, recoverable entities, and subnets across the registered clusters (Compute and Storage Entities on page 93).
- Dashboards to manage security policies, protection policies, recovery plans, Nutanix Guest Tools (NGT) policies, and image placement (see Policies on page 189).
- Hardware component dashboards with drill-down options to view detailed information about individual clusters, hosts, disks, and GPUs across the registered clusters (see Hardware Entities on page 201).
- Activity monitors for alerts, events, audits, and tasks (see Activity Entities on page 240).
- Tools to analyze system activity, plan for resource needs, create usage reports, and automate routine administrative tasks (see Operations Entities on page 246).
- Dashboards to manage inventory (LCM), projects, roles, users, and availability zones (see Administration Entities on page 248).
- Service connection pages to enable Calm, Karbon, and Objects (see Services Entities on page 267).
- Settings menu from which you can configure Prism Central functions (see Settings Menu (Prism Central) on page 16).

Main Menu (Prism Central)

The main menu at the top of every screen provides access to the features of Prism Central. This section describes each of the main menu options.

Viewing Choices

Clicking the collapse menu button (also know as a "hamburger" button) in the far left of the main menu displays the entities menu (see Entities Menu on page 13). The Main dashboard is the first screen that appears after logging into Prism Central (see Main Dashboard on page 70).

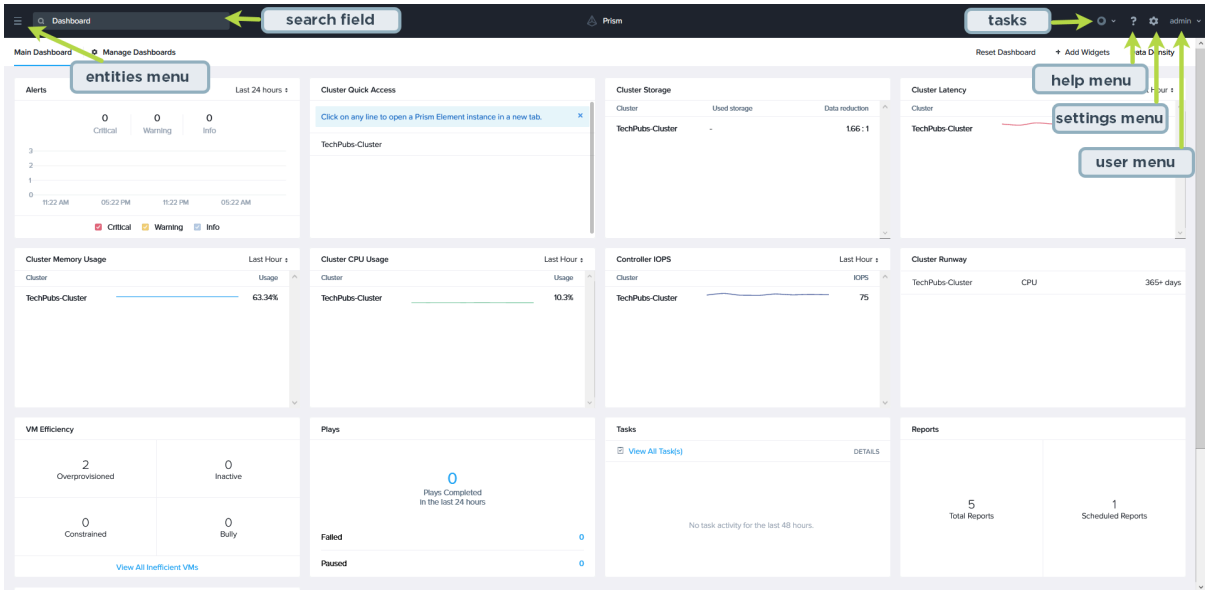

Figure 1: Prism Central Main Menu

Search, Alert, and Task Options

The main menu provides status information about active alerts and tasks plus access to advanced search capability:

- A field with a magnifying glass icon appears on the left of the main menu. Enter a string in this field to search for relevant content in Prism Central (see Searching for Information on page 19).
- An alerts icon appears on the right of the main menu when there are current critical (red), warning (yellow), or informational (gray) alert messages. The number of active alerts is displayed in the icon. Click the alert icon to display information about those alerts in the Alerts dashboard (see Alerts Summary View (Prism Central) on page 269).
- A tasks icon appears to the right of the alerts when there are current tasks (running or completed within the last 48 hours). The icon is blue when all the tasks are running properly or have completed successfully. If one or more of those tasks do not complete successfully, the icon turns yellow (warning) or red (failure) . Click the tasks icon to see a drop-down list of the current tasks; click **View All Tasks** at the bottom of the list to display the tasks dashboard (see Tasks View on page 244).

Help Menu ("?" Icon)

A question mark icon appears on the right side of the main menu. Clicking the question mark displays a list of help resource options that you can select. The following table describes each option in the pull-down menu.

Table 1: Help Menu Options

Name	Description
Learn about search	Displays search guidelines (see Searching for Information on page 19).
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help (Prism Central) on page 1016).
Online Documentation	Opens the online help at the introduction page (see Accessing Online Help (Prism Central) on page 1016).

Name	Description
Support Portal	Opens a new browser tab (or window) at the Nutanix Support portal logon page (see Accessing the Nutanix Support Portal (Prism Central) on page 1011).
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 1018). The portal is an online community site for customers and partners to exchange ideas, tips, and information about Nutanix technologies and related data center topics.
Create Support Case	Opens the **Create a new support case** page to view or create support cases with Nutanix customer support (see Creating a Support Case on page 1006).

Settings Menu (Gear Icon)

A gear icon appears on the right side of the main menu. Clicking the gear icon launches the **Settings** page and displays a menu of tasks you can perform (see Settings Menu (Prism Central) on page 16).

User Menu (<user_name> Icon)

A user icon appears on the far right side of the main menu with the current user logon name. Clicking the user icon displays a list of options to update your user account, log off from Prism Central, and other miscellaneous tasks. The following table describes each option in the pull-down menu.

Table 2: User Menu Options

Name	Description
Change Password	Opens the **Change Password** window to update your password (see Updating My Account in *Security Guide*).
Update Profile	Opens the **Update Profile** window to update your user name and email address (see Updating My Account in *Security Guide*).
Download Cmdlets Installer	Downloads the PowerShell installer for the Nutanix cmdlets. For information about installing the cmdlets locally and for cmdlet descriptions, see Powershell Cmdlets Reference.
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). For more information about installing the nCLI locally and for nCLI command descriptions, see Command Reference.
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 1014).
About Nutanix	Opens the **About Nutanix** window that displays Nutanix operating system (AOS) and other version information (see Finding the Prism Central Version on page 45).
Nothing To Do?	Opens a game that is strictly for entertainment. To quit the game, click the "X" at the upper right of the screen.
Sign Out	Logs you out of Prism Central (see Logging Out of Prism Central on page 34).

Name	Description
Adjust Contrast (Chrome only)	Displays a contrast setting box at the bottom of the screen where you can set the Prism Central display to **Normal** (default) or **High** contrast. This option is available only when using the Chrome browser.

Entities Menu

Clicking the collapse menu (hamburger) button in the main menu (see Main Menu (Prism Central) on page 10) displays a menu of entities on the left. The following table describes each entity in the menu.

- The primary menu options are **Dashboard**, category names (**Compute & Storage**, **Network & Security**, **Data Protection**, **Hardware**, **Activity**, **Operations**, **Administration**, and **Services**), and **Global Settings**. Clicking a category name displays a submenu of entities in that category.
- Clicking the entity name displays a dashboard or other appropriate page for that entity.
- For quicker access to one or more submenu entities, you can promote a submenu option to the primary menu by bookmarking that entity (clicking the star icon to the right of the name). This displays that entity name below the **Dashboard** entry. You can remove the bookmarked entry by clicking the star again.
- The entity menu pane closes automatically when you click in the display area to the right. To keep the entity menu pane open all the time, click the **Lock Navigation Bar** button at the bottom. The button is a toggle; to unlock the navigation bar, click the button again.
- When the entities menu is displayed, the collapse menu button changes to an X button. Click the X button to hide the menu.

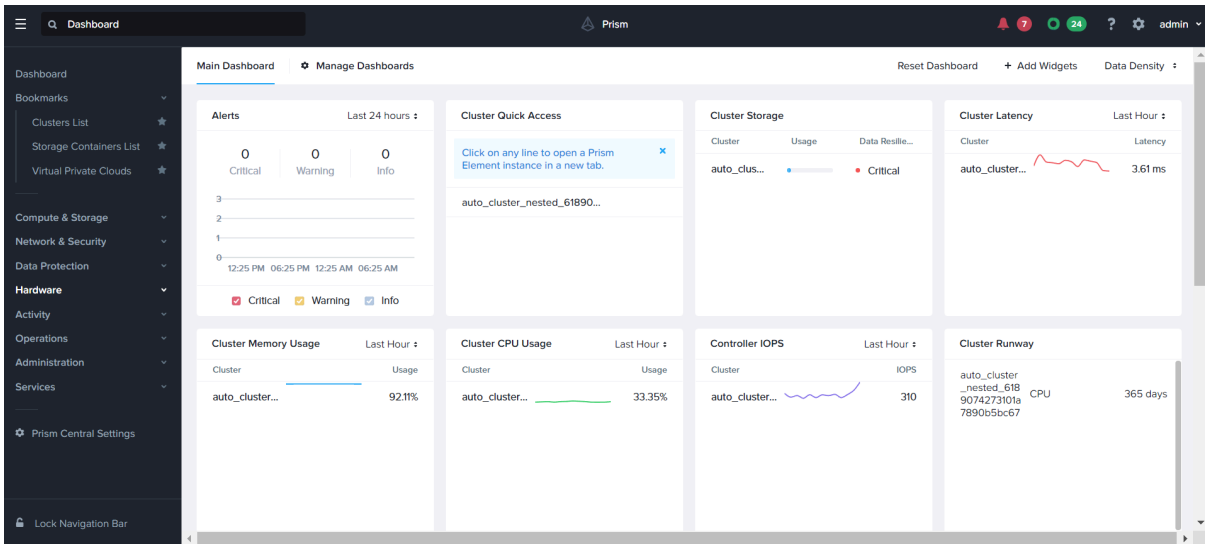

Figure 2: Entities Menu

Table 3: Entities Menu Options

Category	Name	Description
Dashboard		Displays the main dashboard (see Main Dashboard on page 70).

Category	Name	Description
Compute & Storage	VMs	Displays the VMs dashboard (see VMs Summary View on page 94).
	Templates	Displays the Templates dashboard (see VM Template Summary View on page 118).
	OVAs	Displays the OVAs dashboard (see OVAs View on page 122).
	Images	Displays the Images dashboard (see Images Summary View on page 125).
	Catalog Items	Displays the Catalog dashboard (see Catalog Items View on page 129).
	Storage Containers	Displays the Storage Containers dashboard (see Storage Containers Summary View on page 130).
	Volume Groups	Displays the Volume Groups dashboard (see Volume Groups Summary View on page 147).
	Storage Policies	Displays the Storage Policies dashboard(see Storage Policies Summary View on page 141).
	vCenter Datastores	Displays the vCenter Datastores dashboard (see vCenter Datastores Summary View on page 156).
Network & Security	Subnets	Displays the subnets dashboard (see Subnets on page 170).
	Virtual Private Clouds	Displays the virtual private clouds dashboard (see Network and Security View on page 170).
	Floating IPs	Displays the floating IPs dashboard (see Network and Security View on page 170).
	Connectivity	Displays the VPNs dashboard (see Network and Security View on page 170).
	Security Policies	Displays the Security Policies dashboard (see Security Policies Summary View on page 195).
Data Protection and Recovery	Protection Summary	Displays the protection summary dashboard.
	Protection Policies	Displays the Protection Policies dashboard (see Protection Policies View on page 198).
	Recovery Plans	Displays the Recovery Plans dashboard (see Recovery Plans View on page 199).
	VM Recovery Points	Displays the VM Recovery Points dashboard (see VM Recovery Points View on page 200).
Hardware	Clusters	Displays the Clusters dashboard (see Clusters Summary View on page 202).
	Hosts	Displays the Hosts dashboard (see Hosts Summary View on page 221).

Category	Name	Description
	Disks	Displays the Disks dashboard (see Disks Summary View on page 230).
	GPUs	Displays the GPUs dashboard (see GPUs Summary View on page 237).
Activity	Alerts	Displays the Alerts dashboard (see Alerts Summary View (Prism Central) on page 269).
	Events	Displays the Events dashboard (see Events Summary View (Prism Central) on page 279).
	Audits	Displays the Audits dashboard (see Audits Summary View on page 240).
	Tasks	Displays the Tasks dashboard (see Tasks View on page 244).
Operations	Analysis	Displays the Analysis dashboard (see Analysis Dashboard (Prism Central) on page 799).
	Cost Management	Provides access to the Xi Beam site (see Cost Management (Xi Beam) on page 855).
	Discovery	Displays the App Discovery dashboard or the enable wizard if application discovery is not enabled (see Application Discovery on page 831).
	Monitoring Configurations	Displays the Monitoring Integrations dashboard or the enable wizard if monitoring integrations is not enabled (see Application Monitoring on page 856).
	Ops Policies	Displays the Operations Policies dashboard (see Operations Policy Management on page 873)>.
	Planning	Displays the capacity Planning dashboard (see Resource Planning on page 876).
	Playbooks	Displays the list of playbooks (see Task Automation on page 902).
	Reports	Displays the Reports dashboard (see Reports View on page 246).
Administration	Categories	Displays the Categories dashboard (see Categories Summary View on page 164).
	LCM	Displays the life cycle manage (LCM) dashboard (see LCM View on page 267).
	Projects	Displays the Projects dashboard (see Projects Summary View on page 256).
	Roles	Displays the Roles dashboard (see Roles Summary View on page 252).
	Users	Displays the Users dashboard (see Users Summary View on page 248).

Category	Name	Description
	Availability Zones	Displays the Availability Zones dashboard (see Availability Zones View on page 266).
Services	Calm	Displays the Calm feature page or the enable wizard if Calm is not enabled (see Enabling Calm on page 992).
	Files	Displays the Files dashboard or the enable wizard if Files is not enabled (see Enabling Files on page 992).
	Foundation Central	Displays the Foundation Central dashboard or the enable wizard if Foundation Central is not enabled (see Enabling Foundation Central on page 993).
	Karbon	Displays the Karbon feature page or the enable wizard if Karbon is not enabled (see Enabling Karbon on page 993).
	Objects	Displays the Objects feature page or the enable wizard if Objects is not enabled (see Enabling Objects on page 993).
Prism Central Settings		Displays the settings menu (see Settings Menu (Prism Central) on page 16).

Settings Menu (Prism Central)

Prism Central includes a Settings page from which you can configure various system services. You can access the Settings page by doing either of the following:

- Click the gear icon in the main menu (see Main Menu (Prism Central) on page 10).
- Click the collapse menu (hamburger) button and then select **Prism Central Settings** from the entities menu (see Entities Menu on page 13).

The Settings page displays a menu of tasks (on the left) you can perform. Click the task to open the window or page for that option in the pane to the right. The following table describes each menu option.

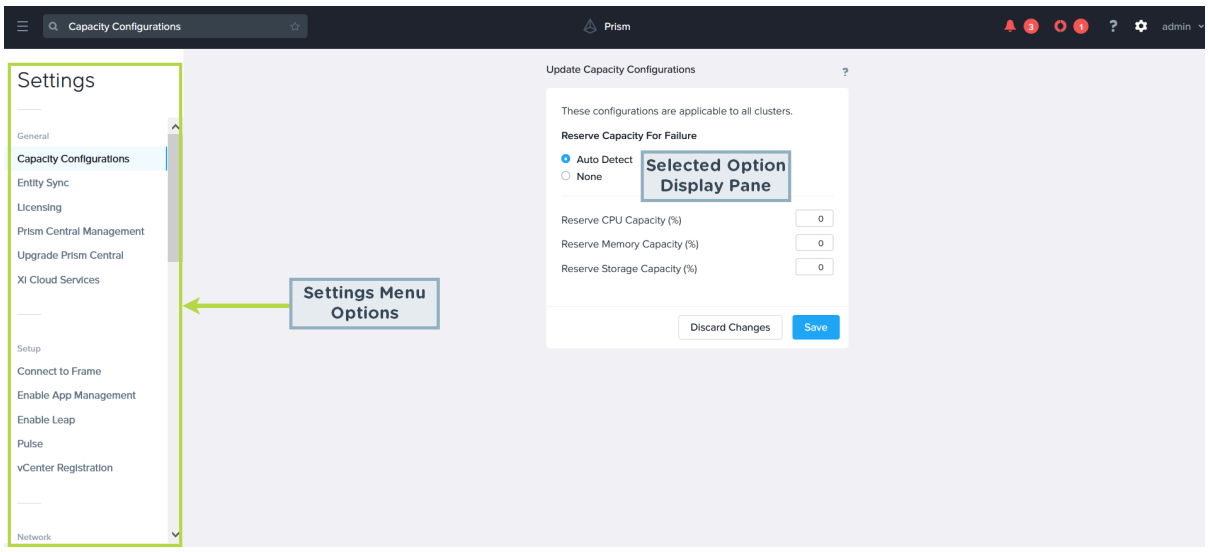

Figure 3: Settings Page

Table 4: Settings Menu

Category	Name	Description
General	Capacity Configurations	Opens the **Update Capacity Configurations** window, which allows you to configure sizing rules for cluster runway analysis and alerting (see Updating Capacity Configurations on page 898). This menu option does not appear if Prism Pro is disabled.
	Entity Sync	Opens the **Force Entity Sync** window, which synchronizes entities with the specified availability zones. For more information, see Nutanix Disaster Recovery Guide.
	Licensing	Opens the **Licensing** page to install or update the Prism Central license that enables entitled Nutanix features (see Licensing on page 31).
	Prism Central Management	Opens the **Manage Prism Central** page for viewing information about Prism Central and optionally expanding (scale out) Prism Central across multiple VMs (see Managing Prism Central on page 34).
	Upgrade Prism Central	Opens the **Upgrade Prism Central** window to upgrade the Prism Central VM to a newer version (see Installing or Upgrading on page 30).
	Xi Cloud Services	Opens the **Xi Cloud Services** windows, which allows you to log in to the Xi portal. See the Xi Cloud Services Administration Guide for more information.
Setup	Connect to Frame	Opens the **Connect to Frame** page to connect to Xi Frame. For more information, see Xi Frame documentation.
	Enable App Management	Opens the **Enable App Management** page to enable the Nutanix Calm features (see Enabling Calm on page 992).
	Enable Nutanix Disaster Recovery	Opens the **Enable Nutanix Disaster Recovery** page to enable the Disaster Recovery as a Service (DRaaS) capability. For more information, see Nutanix Disaster Recovery Guide.
	Pulse	Opens the **Pulse** page to enable (disable) the Nutanix automated support system (see Configuring Pulse on page 997).
	Self-Service Admin Management	Opens the Self-Service Admin Management page to modify the self-service configuration (see Configuring Prism Self Service on page 724). This menu option does not appear until a directory service is configured.
	vCenter Registration	Opens the **vCenter Registration** window to register (or unregister) clusters with vCenter (see Registering vCenter Server (Prism Central) on page 719).

Category	Name	Description
Network	HTTP Proxy	Opens the **HTTP Proxies** window to configure an HTTP proxy to which the Nutanix software can connect (see Configuring an HTTP Proxy on page 509).
	Name Servers	Opens the **Name Servers** window to configure name servers (see Configuring Name Servers (Prism Central) on page 514).
	NTP Servers	Opens the **NTP Servers** window to specify which NTP servers to access (see Configuring NTP Servers (Prism Central) on page 515).
	SNMP	Opens the **SNMP Configuration** window to configure SNMP (see Configuring SNMP (Prism Central) on page 517).
Flow	ID Based Security	Opens the **ID Based Security** page to add Active Directory domain services configurations and to import user groups for identity-based security policies. For more information, see Nutanix Flow Guide.
	Microsegmentation	Opens the **Microsegmentation** page to enable this feature, which is disabled by default. Before you can configure and use application security policies, isolation environment policies, and quarantine policies, you must enable the feature. For more information, see Nutanix Flow Guide.
Security	Cluster Lockdown	Opens the **Cluster Lockdown** window, which allows you to delete (or add) public authentication keys used for SSH access into Prism Central. Removing all public keys locks down Prism Central from external access. For more information, see Controlling Remote (SSH) Access in *Security Guide*.
	SSL Certificate	Opens the **SSL Certificates** window to create a self-signed certificate. For more information, see Installing an SSL Certificate in *Security Guide*.
User and Roles	Authentication	Opens the **Authentication Configuration** window to configure Prism Central authentication. For more information, see Configuring Authentication in *Security Guide*.
	Local User Management	Opens the **Local User Management** window. This window lists current users and allows you to add, update, and delete local user accounts. For more information, see Managing Local User Accounts in *Security Guide*.
	Role Mapping	Opens the **Role Mapping** window to configure role mappings that apply in the user authentication process. For more information, see Configuring Role Mapping in *Security Guide*.

Category	Name	Description
Alerts and Notifications	Alert Email Configuration	Opens the **Alert Email Configuration** window, which allows you to configure rules and templates that apply to the alert messages from selected clusters (see Configuring Alert Emails (Prism Central) on page 281).
	SMTP Server	Opens the **SMTP Server Settings** window to configure an SMTP server (see Configuring an SMTP Server (Prism Central) on page 516).
	Syslog Server	Opens the **Syslog Servers** window to configure a syslog server (see Configuring Syslog Monitoring on page 523).
Appearance	Language Settings	Opens the **Language Settings** window to localize the Prism Central screen text (see Changing the Language Settings (Prism Central) on page 532).
	UI Settings	Opens the **UI Settings** window to configure Prism UI background themes, disable (or re-enable) the login screen background video and configure the automatic logout timeout (see Modifying UI Settings (Prism Central) on page 527).
	Welcome Banner	Opens the **Edit Welcome Banner** window to create a welcome banner message that appears before users login to Prism Central (see Configuring a Banner Page (Prism Central) on page 531).

Searching for Information

The main menu includes a search field on the left (see Main Menu (Prism Central) on page 10) that allows you to find information about selected entities in various ways. If you need help at any time to navigate Prism Central or apply a search filter, you can click the question mark icon in the main menu and select **Learn about search**. This displays the **Search Guidelines** page that explains the search rules and options.

Search Basics

An *entity* is an object type such as a VM, cluster, node, security policy, project, report, event, alert, or audit. The search field is context-sensitive, which means it is populated automatically based on where you are in Prism Central. The first screen after logging on is the main dashboard, so "Dashboard" appears in the search field initially.

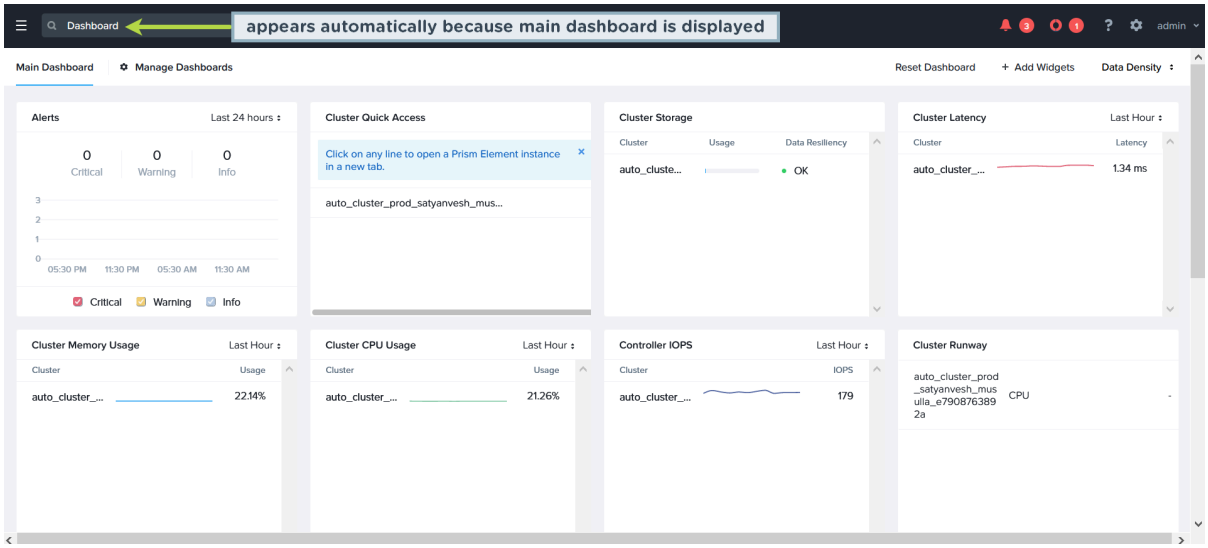

Figure 4: Search Field (initial screen)

As you go to different Prism Central screens, the search string changes automatically to match the current location, as illustrated in the following example from the clusters page.

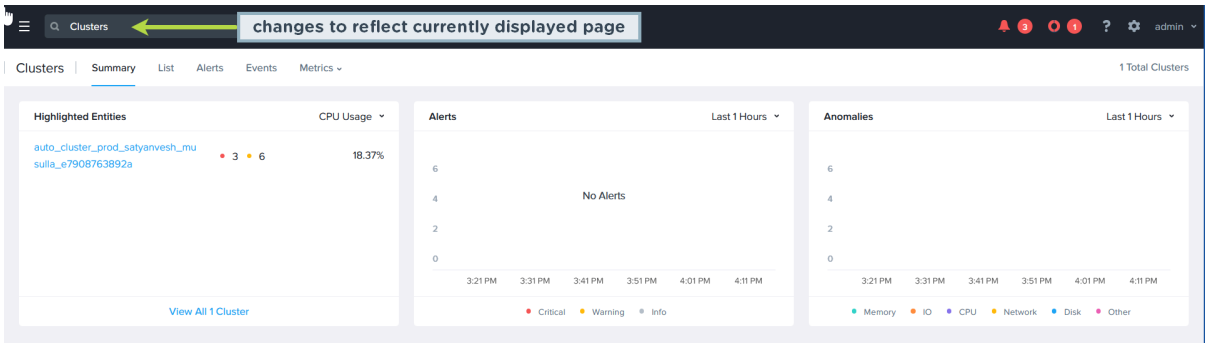

Figure 5: Search Field (from a different screen)

When you enter a string in the search field, a drop-down list appears with relevant matches across Prism Central. (Search strings are case insensitive.) For example, entering the string "VM" displays the following results:

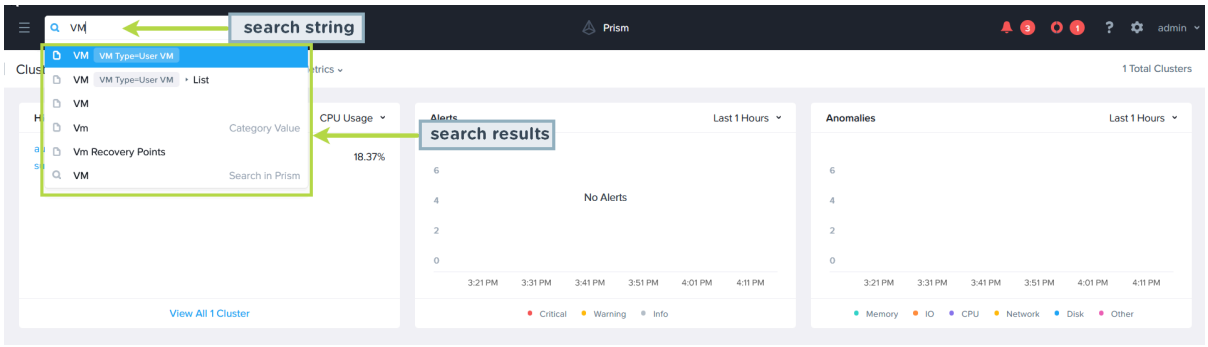

Figure 6: Search Results (for an entity type)

- "Category Value". Click this option to display the Category page with any VM-related entries.
- "VM Type=User VM". Click this option to display the VMs summary page filtered for user VMs.
- [blank]. Click this option to display the VMs summary page (unfiltered).
- "VM Type=User VM > List". Click this option to display the List tab of the user VMs page.
- "VM Type=User VM > Alerts". Click this option to display the Alerts tab of the user VMs page.
- "Search in Prism". Click this option to search across Prism Central for any information about the target entity. In the following example, two tables appear in the results, a list of the top VMs and a list of the top VM-related alerts. The top VMs list includes a link to the full list of VMs, and the top alerts list includes a link to the full list of alerts (which in this case is the same). Clicking the VM or alert name takes you to the details page for that VM or alert.

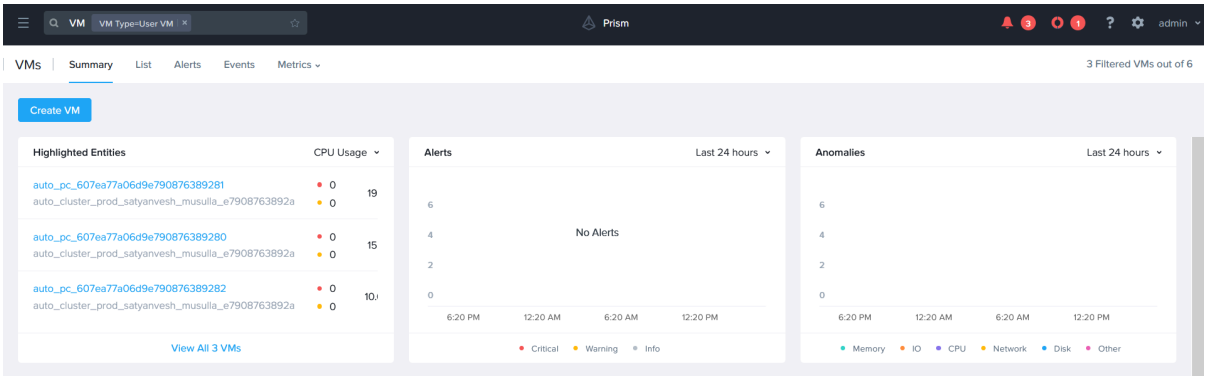

Figure 7: "Search in Prism" Results

Because search is context-sensitive, you can do a quick search from your current location without entering an explicit string. For example, if you are on the user VMs summary page, clicking in the search field (or clicking the forward slash [/] character) displays the search results for user VMs. In the following example, clicking the List, Alerts, or Events entry is the same as clicking those tabs on the page. Clicking "Memory Swap" displays the Memory Swap metrics page, and "Search in Prism" displays a results page for user VMs.

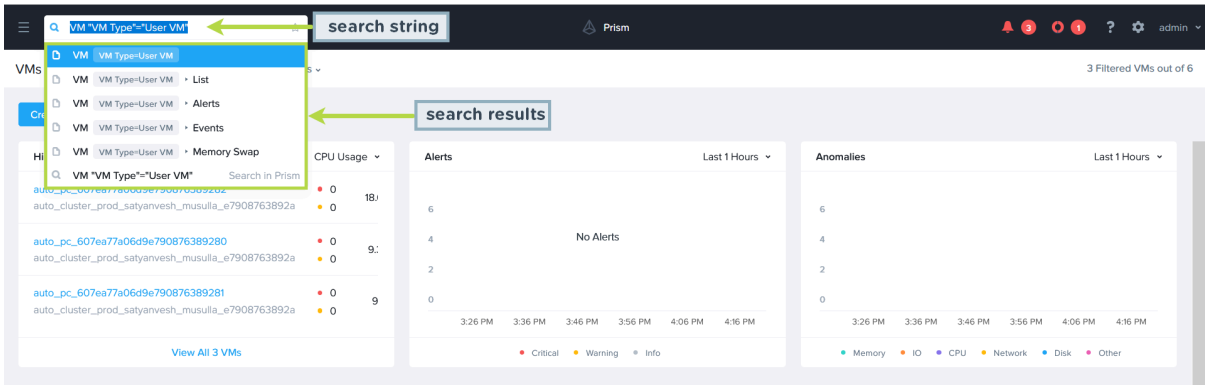

Figure 8: Search Results (for the current location)

Search provides intelligent responses based on whatever you enter in the field. For example, entering "scale out" returns a link to the Manage Prism Central page from which you can scale out Prism Central. If you enter a more generic string such as "version", the search engine returns results with links to pages across entities with version information.

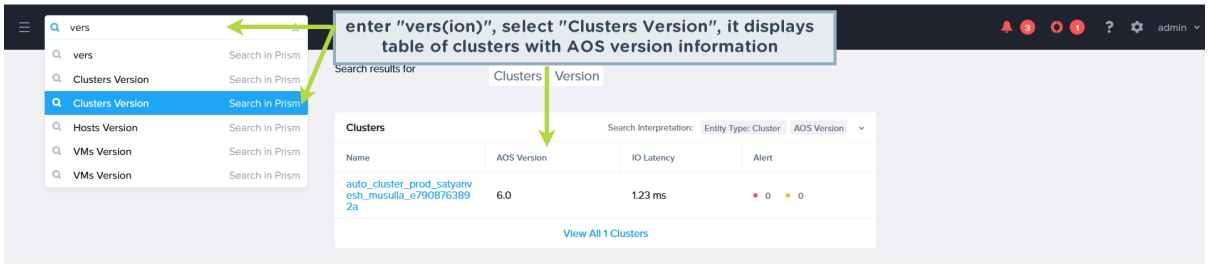

Figure 9: Search Results (for generic string)

You can bookmark a filtered page (see Entity Exploring on page 89) by clicking the star icon in the search field. This adds a bookmark for that page (including any applied filters) to the entities menu (see Entities Menu on page 13).

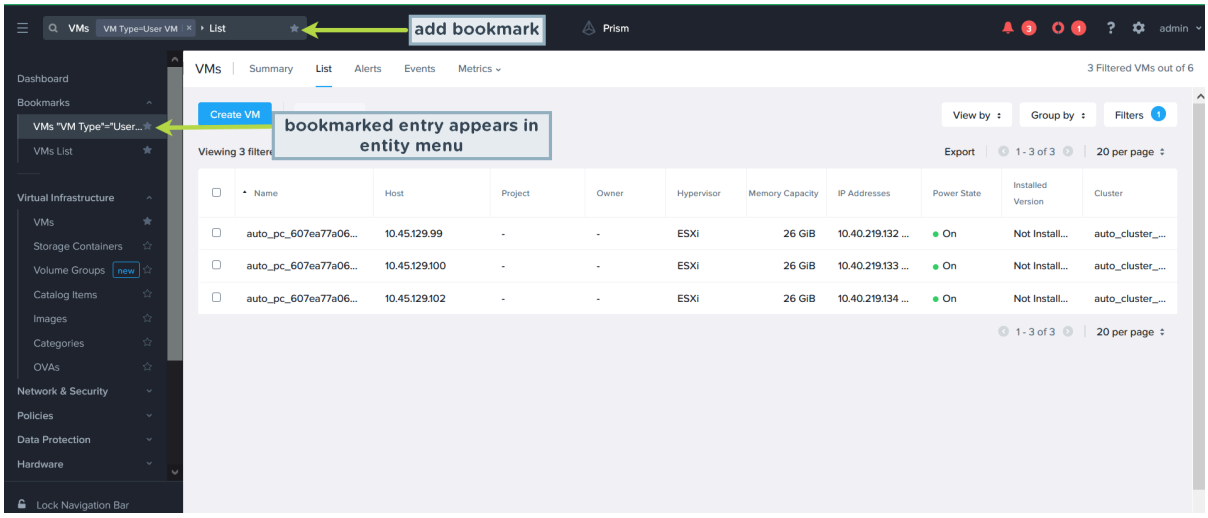

Figure 10: Bookmark Entry

On some pages, applied filters are not reflected in the main menu search field. For those pages, a local filters field appears. When you select one or more filters on such a page, the applied filters appear in the local filters field. You can do the following from this field:

- Remove a filter by clicking the X for that filter.
- Remove all filters by clicking **Clear** (on the right).
- Save the filter list by clicking the star icon. You can save a maximum of 20 filter lists per entity type.
- Use a saved filter list by selecting from the pull-down list.

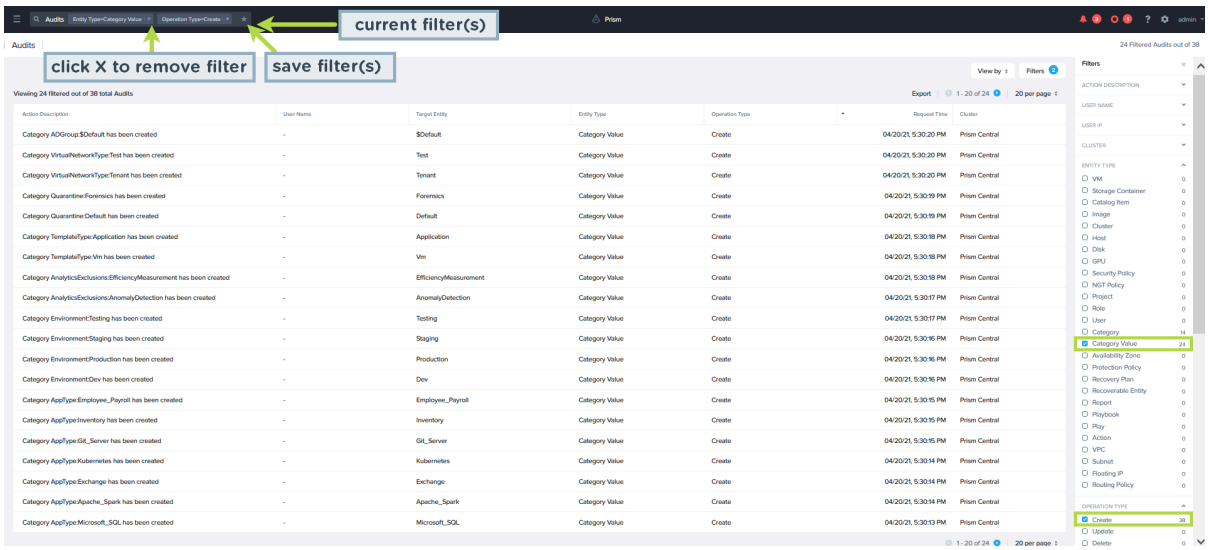

Figure 11: Local Filters Field

Query Syntax

The search field supports a range of simple and more complex query options in the following syntax forms:

- [entity|metric]: Enter an entity or metric type such as "`cluster`", "`vm`", or "`alerts`" to return information about the instances of that entity or metric across the registered clusters.
- <entity> <metric>: To refine an entity search, add a metric to the query. For example, "`vm iops`" returns IOPS information for the VMs.
- <entity> <attribute> <metric>: To further refine the search, add an attribute for the metric. For example, "`node failure alerts`" returns any failure alerts about the nodes.
- [metric|attribute] operator <value>: Enter an expression for a metric or attribute to return instances that match the expression. For example, "`"block model"=1050`" returns information about the NX-1050 nodes. See the following "Filter Expressions" section for the supported expression operators.
- <complex expression>: Combine the syntax forms to create complex queries. For example, "`clusters hypervisor = AHV "cpu usage" < 30`" returns information about clusters running the AHV hypervisor with CPU usage below 30%. Complex expressions have an implied AND so only instances that satisfy all conditions are returned.
- <action expression>: In addition to search queries, you can initiate certain actions from the query bar. For example, "`<vm_name> launch console`" launches a console window for the specified VM (if you are allowed to do so).

The following table describes the syntax rules for search queries.

Table 5: Syntax Rules

Rule	Example
RULES THAT GO TO WELL DEFINED PAGES	
entity type	vms
entity type + metric perspective (io, cpu, memory)	vms io

Rule	Example
entity type + alerts	vm alerts
entity type + alerts + alert filters	vm alerts severity=critical
entity type + events	vm events
entity type + events + event filters	vm events classification=anomaly
entity type + filters (both metric and attribute)	vm “power state”=on
entity type + filters + metric perspective (io, cpu, memory)	vm “power state”=on io
entity type + filters + alerts	vm “power state”=on alerts
entity type + filters + alerts + (alert filters)	vm “power state”=on alerts severity=critical
entity type + filters + events	vm “power state”=on events
entity type + filters + events + event filters	vm “power state”=on events classification=anomaly
entity instance (name, ip address, disk serial, ...)	vm1, 10.1.3.4, BHTXSPWRM
entity instance + metric perspective (io, cpu, memory)	vm1 io
entity instance + alerts	vm1 alerts
entity instance + alerts + alert filters	vm1 alerts severity=critical
entity instance + events	vm1 events
entity instance + events + event filters	vm1 events classification=anomaly
entity instance + pages	vm1 nics, c1 capacity
parent instance + entity type	c1 vms
alert title search	disk bad alerts
page name search	analysis, tasks
RULES THAT ONLY GO TO SEARCH RESULT PAGES (exclusively)	
entity type + metric	vm iops
entity type + attribute	vm power state
entity instance + metric	vm1 iops
entity instance + attribute	vm1 power state
entity type + filters + metric	vm “power state”=on iops
entity type + filters + attribute	vm “power state”=on hypervisor
help text search	upgrade cluster

Keywords

The following table lists the keywords or phrases you can use when formulating a query.

- The *Object* column identifies the type of object.
  - *Entities*: Lists the entities you can specify. The list is limited. For example, you cannot search for information about remote sites or network switches.
  - *Fields*: Lists the parameters (fields) you can specify.
  - *Alerts*: Lists the alert conditions you can specify.
  - *Events*: Lists the event conditions you can specify.
  - *Actions*: Lists the actions you can specify.
- The *Entity* column specifies the entities for which you can use that keyword in a query. For example, Alert queries can apply to any entity, but Fields has multiple rows broken down by entity type (such as cluster, VM, and disk) because there are select keywords that apply to each entity type.
- The *Keywords* column is a comma-separated list of the keywords or phrases you can use in a query for the specified object/entity type.

Table 6: Search Keywords

Object	Entity	Keywords
Entities	(n/a)	vm, cluster, node, container, disk
Fields	(not specified)	cpu usage, memory usage, disk usage, free physical storage, storage logical usage, saving ratio, savings, iops, read iops, write iops, io bandwidth, read io bandwidth, write io bandwidth, io latency, read io latency, write io latency, memory capacity, hypervisor
	Cluster	cluster name, ip address, version, number of hosts, cpu count, memory capacity, runway, storage runway, cpu runway, memory runway
	VM	vm name, ip address, host ip, virtual cpus count, power state, reserved memory, os, virtual hardware version
	Disk	serial, tier, mode, iops
	Container	container name, rf, on disk deduplication, perf-tier deduplication, compression, erasure coding
	Node	host name, ip address, service vm, hypervisor name, cpu capacity, cpu model, cpu sockets count, cpu cores count, cpu thread count, serial number, block model

Object	Entity	Keywords
Alerts	(any)	alert, alert title • Severity levels: critical, warning, info • Categories (impact types): capacity, performance, configuration, availability, system indicator • Dispositions types: resolved, unresolved, acknowledged, unacknowledged
Events	(any)	event type, title
Actions	VM	clone, migrate, delete, power on, power off, suspend, create vm, launch console, create network config, resume, snapshot, update, configure vm host affinity
	Cluster	launch prism element, unregister

Filter Expressions

You can use any of the following operators in an expression.

Table 7: Query Operators

Description	Operator
contains	~
does not contain	!~
starts with (“starts with” x)	=x*
ends with (“ends with” x)	=*x
equal to	=
not equal to	!=
greater than	>
greater than or equal to	>=
less than	<
less than or equal to	<=
a metric range, for example CPU usage between 10 and 30	=[10 to 30]

The following implicit operator rules apply to expressions:

- In most cases, an AND is applied to all filters in an expression. For example, "hypervisor=AHV iops>100" means the hypervisor is AHV and IOPS are over 100.
- If the filters are on the same attribute, an OR is applied. For example, “hypervisor=AHV hypervisor=ESXi” means the hypervisor is either AHV or ESXi.

- If the filters are on the same metric, an AND is applied to create a range. For example, "iops>100 iops<500" means IOPS in the range 100-500.
- If multiple range filters are defined in a single query, an OR is applied across the ranges. For example, "iops>100 iops<500 cpu>20 cpu<40" means IOPS in the 100-500 range or CPU in the 20-40 range.

Search Units

The data for search queries comes from the following parameters.

Table 8: Data Parameters

Data Type	Unit	Parameter
bytes usage related properties	GiBiBytes	memory_capacity_bytes, storage.usage_bytes, storage.free_bytes, storage.logical_usage_bytes, data_reduction.saved_bytes, storage.usage_bytes, storage.user_container_own_usage_bytes, data_reduction.saved_bytes, memory_size_bytes
runway properties	days	capacity.runway, capacity.storage_runway, capacity.cpu_runway, capacity.memory_runway
disk capacity related properties	TebiBytes	storage.capacity_bytes, storage.user_capacity_bytes
CPU/memory usage properties	percent	hypervisor_cpu_usage_ppm, hypervisor_memory_usage

Example Queries

Here are examples of various query types.

- Entity queries:

```
<cluster_name>
<cluster_ip_address>
<disk_serial#>
VMs "Power State"=On List
VMs Hypervisor=AHV List
powered on vms "memory capacity" > 32
```

- Performance queries:

```
clusters running out of cpu
clusters hypervisor = AHV "cpu usage" < 30
vm iops
```

- Alert queries:

```
node failure alerts
<cluster_name> alerts
<cluster_name> critical availability alerts
<alert title>
Alerts "Create Time"="08-Nov-2018 9:46 AM to 08-Nov-2018 10:46 AM" Severity=Critical
```

- Action queries:

```
<vm_name> launch console
<cluster_name> launch prism element
```

```
create vm
```

- Exploration queries:

```
clusters hypervisor=AHV
vm os=Linux
<cluster_name> vms
"block model"=1050
"cpu model"=Intel
containers Rf > 2
clusters version=4.6.2
hosts iops < 1000
<cluster_name> powered off vms
disks tier=ssd
vms "cpu usage"
vms “power state”
```

Keyboard Shortcuts in Prism Central

You can use the following keyboard shortcuts to invoke important menu options or views in Prism Central:

Table 9: Keyboard Shortcuts for Prism Central Menu Options

Shortcut Key	Menu Option/View
s	Settings Menu
f	Spotlight (search bar)
u	User Menu
h	Help menu (? menu)
p	Recent tasks

You can use the arrow keys to select a particular menu option.

Understanding Displayed Statistics (Prism Central)

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. For 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 10: Source for Displayed Statistics

Hypervisor Type	Overview, VM, and Storage	Hardware	Analysis
ESXi	hypervisor (controller for some storage stats)	disk	metric dependent
Hyper-V	controller	disk	metric dependent
AHV	controller	disk	metric dependent
Citrix Hypervisor	controller	disk	metric dependent
Mixed (ESXi + AHV)	hypervisor	disk	metric dependent
Prism Central	cluster dependent (hypervisor or controller)	disk	metric dependent

PRISM CENTRAL ADMINISTRATION

Administering Prism Central includes the following:

- Installing Prism Central and then upgrading Prism Central and registered clusters (see Installing or Upgrading on page 30)
- Managing licenses (see Licensing on page 31)
- Accessing Prism Central, managing the Prism Central instance (VM or set of VMs), and registering clusters with Prism Central (see Prism Central Management on page 31)

Installing or Upgrading

A Prism Central instance consists of either a single VM or a set of VMs that functions as a multi-cluster monitoring and control interface. For information about the maximum tested and supported values for entities in Prism Central, see Nutanix Configuration Maximums.

You can install a Prism Central instance in either a Nutanix cluster or a non-Nutanix environment, although there are some limitations when you choose a non-Nutanix environment.

> **Note:** Nutanix does not support Prism Central and its managed clusters in NAT-deployed environments.

- To install or upgrade a Prism Central instance, see Prism Central Upgrade and Installation topic in *Acropolis Upgrade Guide*.

  > **Note:** Do not add any additional vDisks or NICs to a Prism Central VM other than what is specified in the installation or upgrade instructions.
  >
  > See Prism Central Upgrade vDisk on page 30 for information about the dedicate Prism Central upgrade vdisk that the Prism Central upgrade process deploys and uses for upgrade operations.

- To scale out an existing single VM Prism Central to multiple VMs, see Expanding (Scale Out) Prism Central on page 56.
- To upgrade Nutanix Cluster Check (NCC) through Prism Central, see Upgrading NCC on Prism Central in *Acropolis Upgrade Guide*.
- To upgrade one or more managed clusters through Prism Central, see Prism Central-Managed Clusters Upgrade in *Acropolis Upgrade Guide*.
- If you use a proxy server, see Configuring Prism Central When a Cluster Uses Proxy Servers.

Prism Central Upgrade vDisk

Each Prism Central VM includes a `/home` disk partition that Prism Central (one partition per Prism Central VM) uses for storage when upgrading. The Prism Central upgrade process can fail if the available disk space in this partition is near capacity.

To resolve this, with a minimum version of Prism Central version of pc.2022.6, the upgrade process uses a separate, dedicated vDisk of 30 GB capacity. Instead of using the `/home` partition as part of the upgrade process, Prism Central uses this dedicated vDisk to store the downloaded Prism Central software bundle, decompressing the bundle and hosting the installer.

When you upgrade the Prism Central version to a minimum version of pc.2022.6, the Prism Central upgrade process uses the vDisk image in the bundle to create the Prism Central upgrade vDisk. The upgrade process uses this vDisk for the upgrade operations. The same vDisk is used for subsequent upgrades.

Note: For non-Nutanix deployments, the Prism Central upgrade vDisk is not supported during Prism Central deployment. For more details, see the Prism Central Release Notes.

Requirements

- Ensure that the Prism Central instance is registered to a host in the cluster or microservices infrastructure is enabled on the Prism Central instance.
- Ensure that the Prism Element cluster hosting Prism Central is registered to this Prism Central instance
- Ensure that the Prism Element cluster hosting Prism Central has sufficient disk resources to provide the 30 GB vDisk.

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the Nutanix License Manager Guide for information about how to license your cluster, apply add-on licenses, and manage your licenses in a variety of situations.

Prism Central Management

After installing Prism Central (see Installing or Upgrading on page 30), you can

- Log in and out through any supported browser (see Logging Into Prism Central on page 31 and Logging Out of Prism Central on page 34).
- Monitor the status of this Prism Central instance (see Managing Prism Central on page 34).
- Check version information (see Finding the Prism Central Version on page 45).
- Modify login page and automatic log off settings (see Modifying UI Settings (Prism Central) on page 527).
- Register clusters with this Prism Central instance (see Registering or Unregistering a Cluster with Prism Central on page 48).
- Expand Prism Central from a single VM to a multi-VM instance (see Expanding (Scale Out) Prism Central on page 56).
- Shutdown or Start Up a Prism Central VM (see Shutting Down (Starting Up) Prism Central VM on page 60).
- Protect Prism Central to up to three AHV or ESXi clusters registered to it (see Protecting Prism Central on page 62).
- Recover Prism Central from Prism Element running on the registered AHV or ESXi cluster (see Recovering Prism Central (1-Click Disaster Recovery) on page 63.

Logging Into Prism Central

About this task

To log in to Prism Central, do the following. Knowledge base article KB 1661 lists default cluster credentials.

Procedure

1. Open a web browser, enter http://*management_ip_addr* in the address field, and press **Enter**. Replace *management_ip_addr* with the Prism Central VM IP address.

   > **Note:** Prism Central supports the latest version, and the two preceding major versions of Firefox, Chrome, Safari, and Microsoft Edge browsers. The browsers must support TLS 1.2.

   The browser redirects to the encrypted port (9440) and may display an SSL certificate warning. Acknowledge the warning and proceed to the site. If user authentication is enabled and the browser does not have the correct certificate, a denied access message may appear. For the complete list of required ports, see Port Reference.

2. If a welcome screen appears, read the message and then click the **Accept terms and conditions** bar at the bottom.

3. In the login screen, do one of the following: enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

   » Enter your Nutanix login credentials and press **Enter** or click the right arrow icon.

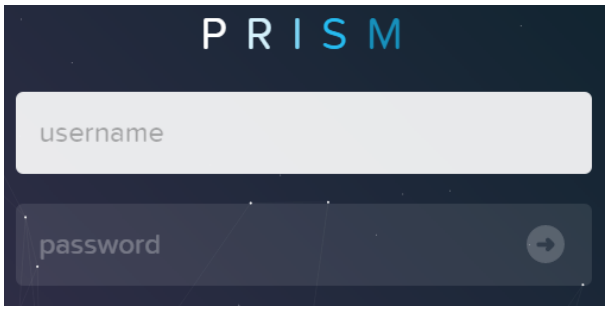

Figure 12: Login Screen

> **Note:** If LDAP authentication is used, enter the user name in *username@domain* format; the *domain\username* format is not supported. The user principle name [UPN] attribute is used to find the user account in Active Directory, so your user name must be in that format.

   » If SAML authentication is enabled (see Configuring Authentication in *Security Guide*), enter your user name according to the configured identify provider in the indicated field. (This field appears on the login page only when SAML authentication is enabled.)

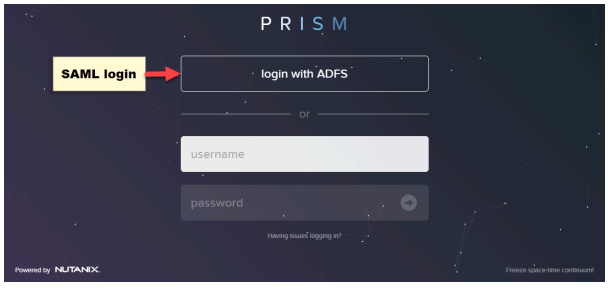

Figure 13: Login Screen (SAML)

> **Note:** The login page includes background animation that is enabled by default. Click the **Freeze space time continuum!** link at the bottom right of the login screen to disable the animation (or the **Engage the warp drive!**

link to enable the animation). To permanently disable (or enable) the animation, see Modifying UI Settings (Prism Central) on page 527).

4. If you are logging in as an administrator for the first time (user name admin and default password Nutanix/4u), which requires that the default password be changed, enter a new password in the **password** and **re-type password** fields and then press **Enter** or click the right arrow icon.

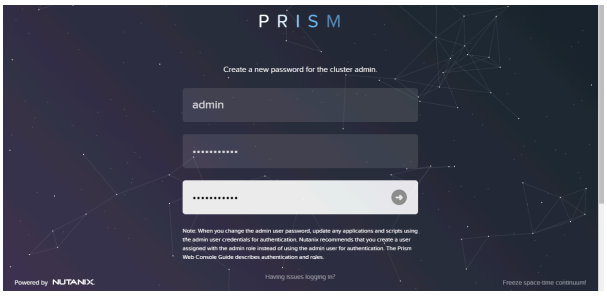

Figure 14: Login Screen (first admin login)

The password must meet the following complexity requirements:

- At least 8 characters long
- At least 1 lowercase letter
- At least 1 uppercase letter
- At least 1 number
- At least 1 special character (allowed special characters are: "#$%&'()*+,-./:;<=>@[]^_`{|}~!)
- At least 4 characters different from the old password
- Should not be among the last 10 passwords

After you have successfully changed the password, the new password is synchronized across all Controller VMs and interfaces (Prism web console, nCLI, and SSH).

Note:

- The password expiration age for the `admin` user is 60 days. You can configure the minimum and maximum password expiration days based on your security requirement.
  - `nutanix@cvm$ sudo chage -M <MAX-DAYS> admin`
  - `nutanix@cvm$ sudo chage -m <MIN-DAYS> admin`
- After you upgrade from an AOS earlier version and then attempt to log in to the Prism web console as the admin user, you are prompted to create a new admin user password.
- When you change the admin user password, update any applications and scripts using the admin user credentials for authentication. Nutanix recommends that you create a user assigned with the admin role instead of using the admin user for authentication. For more information about authentication and roles, see Prism Web Console Guide.

5. If a license agreement screen appears (typically on the first login or if the EULA changed since the last login), which indicates the current EULA has not been acknowledged yet, do the following:

   a. Read the license agreement (on the left).

   b. Enter appropriate information in the **Name**, **Company**, and **Job Title** fields (on the right).

   c. Check the **I have read and agree to the terms ...** box.

   d. Click the **Accept** button.

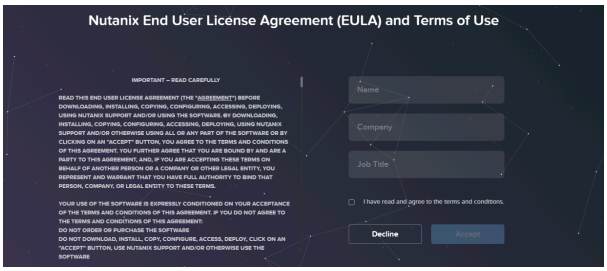

Figure 15: EULA Screen

Logging Out of Prism Central

Procedure

To log off from Prism Central, click the user icon in the main menu and then select the **Sign Out** option from the pull-down list. You are logged out immediately after selecting the option (no prompt or message).

Managing Prism Central

Prism Central consists of one or more VMs treated as a single instance that monitors and manages registered clusters. You can view information about the Prism Central VMs by clicking the gear icon in the main menu and then selecting **Prism Central Management** from the **Settings** menu (see Settings Menu (Prism Central) on page 16). This displays the **Manage Prism Central** page.

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

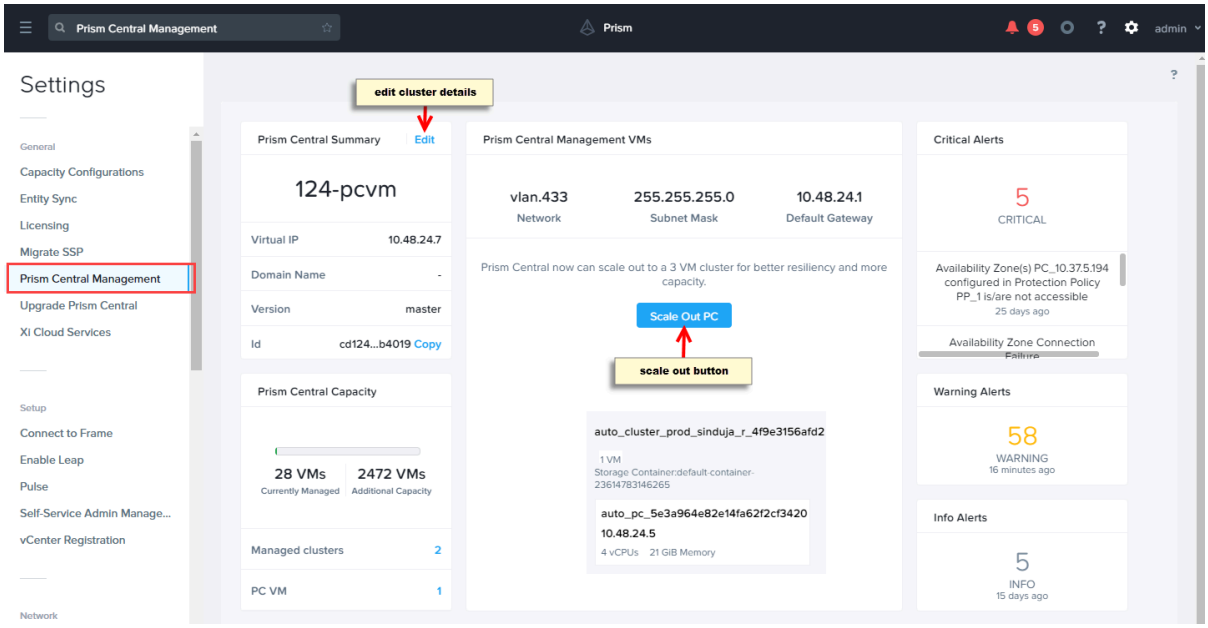

Figure 16: Manage Prism Central Page (1-VM instance)

The **Prism Central Management** page contains the following:

- A **Prism Central Summary** pane on the upper left that displays fields for
  - Cluster name, which is **Unnamed** by default.
  - Domain name. A dash appears if a domain name is not specified.
  - Version. This is the Prism Central version number.
  - ID number. Click the copy link to copy the ID number so you can paste it elsewhere.
- A **Prism Central Capacity** pane on the lower left that displays fields for
  - Number of VMs monitored currently by this Prism Central instance.
  - Number of additional VMs (remaining capacity) this Prism Central instance can monitor.
  - Number of clusters registered to this Prism Central instance; click the number to display the clusters page (see Clusters Summary View on page 202).
  - Number of VMs that comprise this Prism Central instance; click the number to display the VMs page (see VMs Summary View on page 94).
- A **Prism Central VMs** pane in the middle that displays network address information (network name, subnet mask, and default gateway address) at the top and information below about each VM in the Prism Central instance (VM name and name of storage container in which it is located, IP address, number of vCPUs, and memory size). In addition, single VM instances include a **Scale Out PC** button to scale out this Prism Central instance (see Expanding (Scale Out) Prism Central on page 56).
- An **Alerts** pane on the right that displays a list of Prism Central-related alerts broken into Critical, Warning, and Info sections.

When the Prism Central instance consists of three VMs, the display changes slightly. The **Scale Out PC** button disappears, and a new **Add PCVM** button appears (top right). Prism Central instances are limited to a maximum of three VMs, but if you lose one of those VMs for any reason, you can add a replacement by clicking the **Add PCVM**

button. This displays the **Add PC VMs** page (similar to the **Scale Out PC** page) from which you can add the new VM.

> **Note:** Verify that the prerequisites in Expanding (Scale Out) Prism Central on page 56 are satisfied before adding a Prism Central VM.

To add or change the Prism Central domain name or virtual IP, click the **Edit** link in the Prism Central Summary pane. The **Cluster Details** window appears. Do the following:

1. In the **Cluster Name** field, enter a name for the cluster (default "Unnamed"). After naming the cluster, the field no longer appears when subsequently opening the window. Naming the cluster is optional.
2. In the **FQDN** field, enter the fully qualified domain name (FQDN) for the Prism Central cluster. This requires an administrator to configure the domain name in the DNS server so that the DNS server can resolve to all the external IPs of the Prism Central VMs.
3. In the **Virtual IP** field, enter an IP address that will be used as a virtual IP for the cluster. This is relevant if you have a multi-VM Prism Central.

   > **Note:** FQDN and virtual IP settings in Prism Central are mutually exclusive. Do not enter a virtual IP in this field if you specified a FQDN in the previous field. A virtual IP provides resiliency but not load balancing. A FQDN does both.

4. When the fields are correct, click the **Update** button to save the changes and close the window.

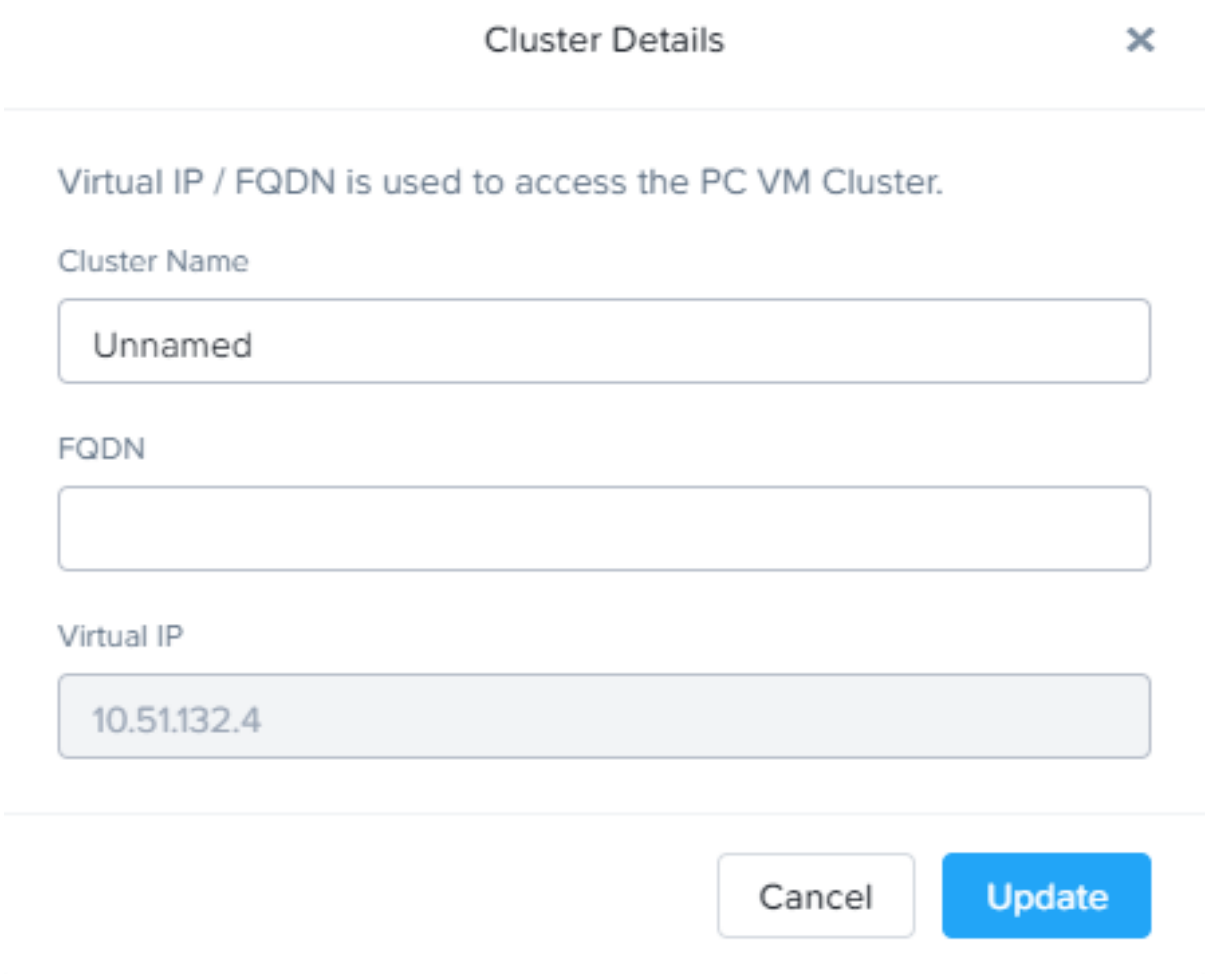

Figure 17: Cluster Details

Microservices Infrastructure

Microservices Infrastructure (sometimes referred to as MSP) provides a common framework and services to deploy the container-based services associated with Prism Central based components like Flow Virtual Networking and Objects. It deploys services like Identity and Access Management (IAM), Load Balancing (LB), and Virtual Private Networking (VPN). Such services are packaged in containers as microservices and the control plane for the microservices platform enables microservices infrastructure.

Microservices infrastructure is disabled by default, but you can enable it at any time in the Prism Central **Settings** page.

> **Caution:** You cannot disable microservices infrastructure after it is enabled.

Microservices Infrastructure Prerequisites and Considerations

Make sure you meet the requirements listed before you enable the microservices infrastructure.

Microservices Infrastructure Prerequisites

- **Hypervisor**

  Clusters running AHV or ESXi only support microservices infrastructure. For ESXi clusters, enter your vCenter credentials (user name and password) and a network for deployment.

- **Prism Element cluster and Prism Central IP addresses**

  Ensure that you have configured the Prism Element and Prism Central cluster virtual IP address and ISCSI data services IP address (see Modifying Cluster Details information in the *Prism Element Web Console Guide*).

  > **Note:**
  >
  > - You can only change these IP addresses of Prism Central or Prism Element cluster after microservices infrastructure is enabled (whether enabled by default or manually enabled). You cannot remove any of these IP addresses. If you remove the current IP address, you must add another IP address before you save the configuration.
  > - If a Prism Central instance (on which Microservices Infrastructure is enabled or will be enabled during upgrade) manages multiple Prism Element clusters, configure the iSCSI Data Service IP address on the Prism Element cluster hosting the Prism Central instance. Ensure that the iSCSI Data Service IP address is configured on the hosting Prism Element cluster before you install the or upgrade to the Prism Central that enables Microservices Infrastructure by default.

- **Name and NTP servers**

  Ensure that you have configured the Name and NTP servers in both the host Prism Element cluster and Prism Central.

- **Prism Central memory**

  KB 8932 lists the Prism Central VM memory requirements. If your Prism Central VM resources do not meet the recommended amount, manually increase the memory.

- **Prism Central registration**

  Ensure that you have registered the Prism Central instance to the Prism Element cluster. See Register (Unregister) Cluster with Prism Central.

  > **Note:** If you unregister the Prism Element cluster from the Prism Central instance after you enable microservices infrastructure, then trust is automatically configured between the Prism Element cluster and the Prism Central

instance. There is no impact to Prism Central. This behavior is supported with a minimum Prism Central version of pc.2022.1 with Prism Element cluster running minimum AOS versions 6.1 or 6.5.

If you do not want to register the Prism Central instance with the Prism Element, then you must complete the trust configuration procedure to enable microservices infrastructure.

The workflow to enable microservices infrastructure includes steps to set up trust between the unregistered Prism Central and the Prism Element. For more information, see Configuring Trust Between Prism Central and Host Cluster.

Trust setup is supported on Prism Element cluster running minimum AOS versions 6.1 or 6.5 and Prism Central running minimum version pc.2022.1.

Note: Ensure that Prism Element cluster is running a minimum version of AOS 6.5 or AOS 6.6 before upgrading the AOS or Prism Central, for which trust is already set up.

If the Prism Element cluster is running a version of AOS older than *minimum supporting AOS version* and the Prism Central is not registered to the host Prism Element, then, before initiating the Prism Central upgrade, register Prism Central with the Prism Element cluster to enable microservices infrastructure by default.

If you upgrade the Prism Central version even though it is not registered to the Prism Element cluster, you need to enable microservices infrastructure manually after the upgrade is complete.

- **Default Private VXLAN**

  If the Prism Central that you want to upgrade is registered to the Prism Element cluster, then the upgrade process enables microservices infrastructure by default using private VXLAN network. If the private VXLAN is used, the ports 9440, 3205, and 3260 are required to be opened from Prism Central to all the Controller VM IP addresses and Controller VM virtual IP addresses in the clusters managed by the Prism Central, and the iSCSI Data Services IP (DSIP) addresses configured for the clusters managed by the Prism Central.

- **Firewall Access to URLs**

  Microservices infrastructure requires access to the following URLs:

  - *.docker.io
  - *.production.cloudflare.docker.com
  - *.nutanix.github.io
  - .quay.io
  - .cloudfront.net
  - *.us-west-2.amazonaws.com

  Microservices Infrastructure requires the following IP address ranges or subnets open (in the firewall) to access the Prism Element Controller VM IP addresses, Controller VM virtual IP addresses, and the Prism Central VIP addresses of port 0, for ICMP traffic. Microservices Infrastructure also requires these IP address ranges or subnets to access the Prism Element Controller VM IP addresses and Controller VM virtual IP addresses, the Prism Element iSCSI Data Services IP (DSIP), and the Prism Central virtual IP addresses on ports 9440, 3205, and 3260. Therefore, ports 0, 3205, 3260 and 9440 must be open in the firewall for these IP address ranges or subnets

to reach the all the Prism Element Controller VM IP addresses and Controller VM virtual IP addresses, Prism Element iSCSI Data Services IP (DSIP) address, and the Prism Central virtual IP addresses as required.

- Default private networks for Microservices Infrastructure: 10.100.0.0/16, 10.200.0.0/16.
- Other IP address ranges or subnets such as 10.200.32.0/24, 10.100.0.0/24, 10.100.1.0/24, and 10.100.2.0/24.
- Any specific IP address range or subnet (if configured) as the private network for Microservices infrastructure.

> **Note:** Do not use the IP addresses in the subnet 10.200.32.0/24, 10.100.0.0/24, 10.100.1.0/24 and 10.100.2.0/24 for operational purposes such as DNS or CVM networks. This prevents any negative impact and maintain seamless operations in the networking.

For more information, see Ports and Protocols to see the access requirements for all the Nutanix software.

- **Prism Central Service Domain Name**.

  During the Microservices Infrastructure deployment process, a DNS service is deployed within Prism Central. This DNS service is utilized for discovery of available Prism Central Microservices. Nutanix recommends that the DNS infrastructure of the customer's domain be updated to forward or redirect all queries for records within the Prism Central Service Domain to the DNS server deployed during Microservices deployment. This ensures future scalability by avoiding any future need for duplication of these service location records on your primary DNS infrastructure.

  Ensure that the following caveats are considered when selecting a Prism Central Domain Name:

  - This value refers to the Prism Central Service Domain, and should not exactly match any FQDN currently in use, as this will complicate configuration of forwarding or redirection of queries.
  - The top-level domain cannot exceed 6 characters.
  - The top-level domain cannot end with "test".
  - The subdomain cannot exceed 16 characters.
  - The selected FQDN must contain at least three labels.
  - All labels within the FQDN must be at least two characters long.
  - If Prism Central has a proxy configuration, then ensure that the domain name is added to the allow-list with a * preceding the domain name (*.defined_domian_name).

    For example, if you use the domain name test.nutanix.com, then allow-list *.test.nutanix.com.

  The following are examples of valid domain configurations:

  - my.cluster.domain
  - my.test.cluster.test.domain
  - test.nutanix.com

  The following are examples of domains not supported:

  - my.cluster.test
  - my.cluster.domain.test

Limitations

The following limitations apply after enabling the microservices infrastructure:

- **Hypervisor Support**

  Nutanix does supports Microservices Infrastructure only on an on-premises Prism Central (PC) deployment hosted on an AOS cluster running AHV or ESXi. Clusters running other hypervisors are not supported.

- Scaleout three-node Prism Central VMs deployment: Small, large, or x-large Prism Central VMs are supported. Reverting this deployment to a single Prism Central VM deployment is not supported.
- Single Prism Central VM deployment: Small, large, or x-large Prism Central VM is supported. You can expand a single PCVM deployment to a scaleout three-node PC deployment.
- Ensure that the IP addresses in the subnet that you configure for microservices infrastructure do not conflict with the IP addresses in the management subnet.

  Note: Do not use the IP addresses in the subnet 10.100.0.0/24 for operational purposes like DNS.

Trust Configuration and Manually Enable Microservices Infrastructure

The following two workflows are closely associated with each other:

- Trust setup between the Prism Central that Microservices Infrastructure is enabled in, and the Prism Element cluster that hosts the Prism Central.
- Manually enabling Microservices Infrastructure in Prism Central.

The trust configuration workflow occurs as the starting point in the workflow to manually enable microservices infrastructure (see Enabling Microservices Infrastructure Manually on page 42).

The steps to manually enable Microservices Infrastructure follows the trust configuration steps.

The starting point for the workflows is the **Enable Now** in the **Prism Central on Microservices Infrastructure** widget on the **Prism Central Management** page in **Prism Central Settings**page.

Configuring Trust Between Prism Central and Host Cluster

If Prism Central is not registered to the host cluster Prism Element, it must have a trust relationship setup with the Prism Element cluster.

Before you begin

Ensure that you have satisfied these requirements.

- The Prism Central instance is deployed with a minimum Prism Central version of pc.2022.1. The Prism Element cluster is deployed with a minimum AOS version of 6.1 (STS versions) or 6.5 (LTS versions). If required, upgrade the AOS version of the Prism Element cluster AOS version of the host cluster before you manually enable microservices infrastructure.
- Add the domain used (like *.prism-central.cluster.local) by the Prism Central to the allow-list.
- Ensure that requirements provided in Microservices Infrastructure Prerequisites and Considerations on page 37 are met.

About this task

If the Prism Central instance is not registered to the Prism Element cluster that hosts it, then after you click **Continue** on the **Important Information** window, Prism Central displays the **Enable Microservices Infrastructure** window with the trust configuration parameters.

Procedure

1. Go to the **Prism Central Management** page (see Managing Prism Central on page 34).
2. In the **Prism Central on Microservices Infrastructure** widget, click **Enable Now**.

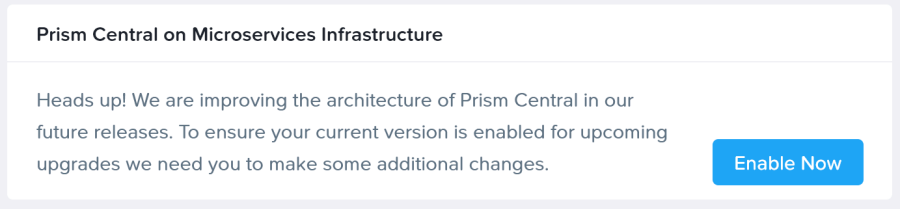

Figure 18: Prism Central on Microservices Infrastructure widget

3. Read the information and then click **Continue**.
4. On the **Enable Microservices Infrastructure** window, provide the following information:
   - Virtual IP Address or FQDN of the host Prism Element.
   - Admin user name used to log on to the host Prism Element.
   - Password for the admin user name provided to log on to the host Prism Element.

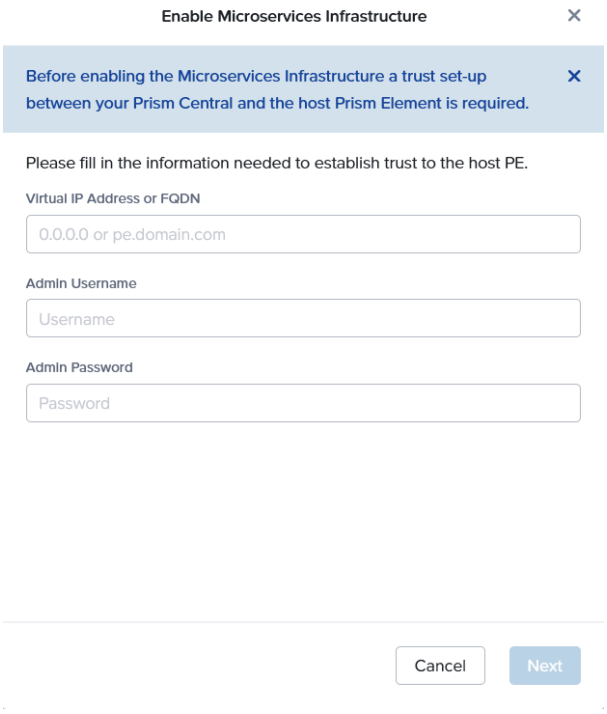

Figure 19: Enable Microservices Infrastructure - Trust setup

5. Click **Next**.
   The **Enable Microservices Infrastructure** main dialog box provides the parameters to enable microservices infrastructure.
6. Complete manually enabling microservices infrastructure. See Enabling Microservices Infrastructure Manually on page 42.

Enabling Microservices Infrastructure Manually

The microservices infrastructure enabled by default. This topic provides the procedure to enable microservices infrastructure manually.

Before you begin

Ensure that all the requirements provided in Microservices Infrastructure Prerequisites and Considerations on page 37 are fulfilled before enabling the microservices infrastructure.

If the Prism Central instance is not registered to the Prism Element cluster hosting the Prism Central, then set up trust between the Prism Central and the Prism Central cluster by completing the steps provided in Configuring Trust Between Prism Central and Host Cluster on page 40.

About this task

To enable the microservices infrastructure, do the following:

Procedure

1. Read the information, and then click **Continue**.

   > **Important:** If the Prism Central deployment is not registered with the Prism Element, then the **Enable Microservices Infrastructure** window provides the parameters for establishing trust between Prism Central and Prism Element that hosts the Prism Central instance.
   >
   > Complete Configuring Trust Between Prism Central and Host Cluster on page 40. Then, the main **Enable Microservices Infrastructure** dialog box is displayed.

   The main **Enable Microservices Infrastructure** dialog box is also displayed if the Prism Central instance is already registered with the host Prism Element.

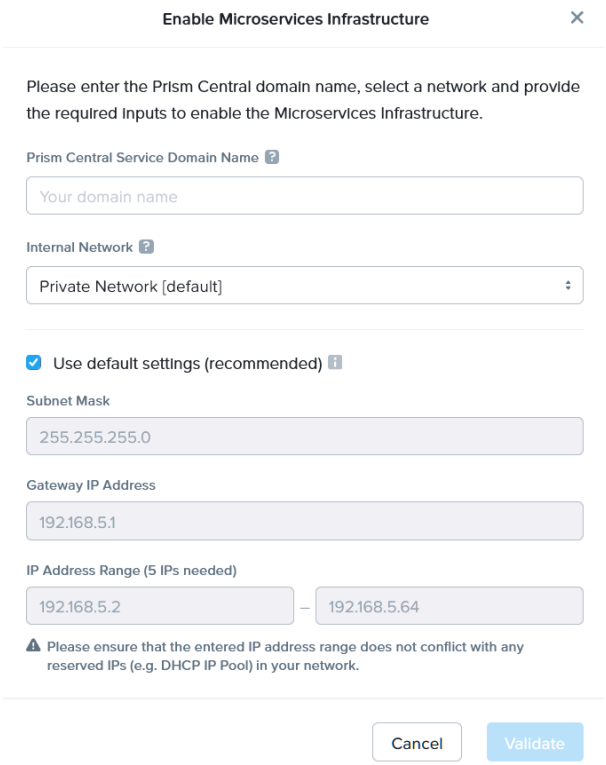

Figure 20: Enable Microservices Infrastructure

2. Provide the following information in the main **Enable Microservices Infrastructure** box.

   - **Prism Central Service Domain Name**: Enter a domain name for the service. See **Prism Central Service Domain Name** restrictions in Microservices Infrastructure Prerequisites and Considerations on page 37.
   - **Internal Network**: Select the network to use for Prism Central micro services communication from the dropdown list.

     The default selection **Private Network [default]** is a pre-configured private VxLAN network. Instead, if you want microservices infrastructure to use a different network, you can select the network (managed or

unmanaged) from the drop-down list. If the network you want microservices infrastructure to use does not appear in the list, you must first configure it (see Network Connections on page 782).

- The **Use default settings (recommended)** checkbox is available only when you have retained the default selection (**Private Network [default]**) for **Internal Network** in the preceding step. This checkbox is checked by default. Do one of the following:

  - Retain the check mark for the **Use default settings (recommended)** checkbox and click **Validate**. (Go to step 3 on page 44.)

    Retaining the check mark for the **Use default settings (recommended)** checkbox allows Prism Central to use the **Private Network [default]** with the default values for **Subnet Mask**, **Gateway IP Address** and **IP Address Range**.

  - Clear the **Use default settings (recommended)** checkbox, if you want Prism Central to use the **Private Network [default]** setting with specific (non-default or custom) values for **Subnet Mask**, **Gateway IP Address** and **IP Address Range**.

Configure the **Internal Network** for microservices infrastructure if you did one of the following:

- Selected a managed or unmanaged network other than **Private Network [default]** for **Internal Network**.

  If you selected a managed network, the values in the **Subnet Mask**, **Gateway IP Address** and **IP Address Range** fields are already configured. If you selected an unmanaged network, you must enter the necessary values in the respective fields.

- Cleared the **Use default settings (recommended)** checkbox with the **Private Network [default]** selection for **Internal Network**.

Enter the values for the following parameters to configure **Internal Network**:

Parameter	Description
Subnet Mask	Enter the subnet mask.
Gateway IP Address	Enter the IP address for the gateway.
IP Address Range	Enter a range of IP addresses that the network can use. Enter a range of at least 10 available (unreserved) addresses for a 3-VM (scale out) Prism Central instance or at least 5 addresses for a 1-VM Prism Central instance. The addresses must be consecutive or sequential. For a managed network, the range of addresses for microservices infrastructure must be outside the range of reserved IP addresses (for example, DHCP IP Pool) in the selected network.

3. After you check that the values entered for all the fields are correct, click **Validate**.
   The system checks the settings. When the settings are confirmed, a Validation was successful message appears and the **Validate** changes to **Enable**.

4. Click **Enable**.
   The process begins, which takes about 45 minutes. You can monitor progress from the tasks dashboard (see Tasks View on page 244). When the microservices infrastructure is enabled, the contents of the enable box in the **Manage Prism Central** page changes to display configuration information such as the domain name, the used IP addresses, and the service domain's DNS IP address.

What to do next

After the tasks are completed, including the **IAM Migration & Bootstrap** task, log out of Prism Central. Wait at least 15 minutes before logging on to Prism Central.

Enabling Microservices Infrastructure (dark site)

Enabling the Microservices Infrastructure when there is no Internet access (dark site) requires you to perform these additional steps.

About this task

To enable the Microservices Infrastructure at a dark site, do the following:

Procedure

1. From a device that has public Internet access, log on to the Nutanix Support portal, select **Downloads** > **LCM** from the entity menu (hamburger icon), and then download the LCM Dark Site Bundle tar file.
2. Configure a local web server, upload LCM Dark Site Bundle to the server, and extract the files in a directory in the base of the web server.

   See the *Setting Up a Local Web Server* section in the *Life Cycle Manager Dark Site Guide* for more information. Set the following file types in the LCM web server when you are using Windows Web Server: .json, .xz, .qcow2, .sign, .gz

   For the mime types associated with the above file types, see Setting Up a Local Web Server
3. In the Nutanix support portal **Downloads** page, select **Microservices Platform (MSP)**, download the **Upgrade (Darksite)** bundle to the web server base directory you used in step 2, and extract the files.
4. Follow the procedure for enabling the Microservices Infrastructure through Prism Central described in Enabling Microservices Infrastructure Manually on page 42.

Finding the Prism Central Version

About this task

You can check the version of your Prism Central instance using the Prism Central user interface or by running the following command.

Figure 21: Finding the Prism Central Version

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275778971001 on Brightcove.

```
nutanix@cvm$ ncli cluster info
```

Do the following to check the Prism Central version in the UI.

Procedure

1. Click the user icon in the main menu and then select the **About Nutanix** option from the pull-down list.
   An **About Nutanix** window appears that includes the AOS version number. It also includes a link to Nutanix patent information.

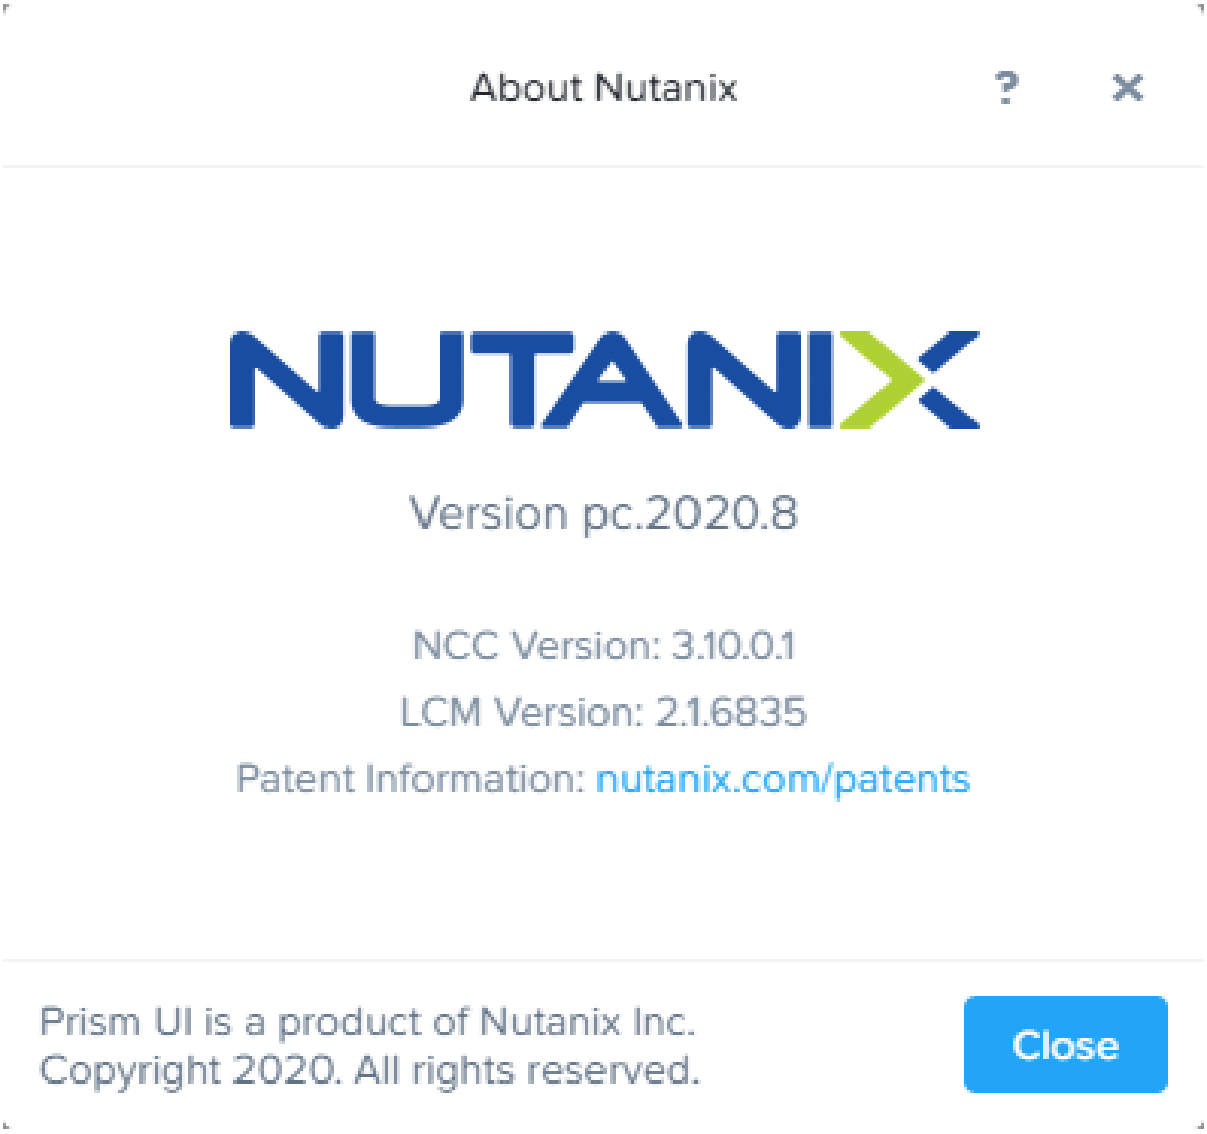

Figure 22: Example About Nutanix Window

2. Click **Close** to close the window.

Finding the AHV Version on Prism Central

You can see the installed AHV version in the Prism Central console.

About this task

To view the AHV version installed on any host in the clusters managed by the Prism Central, do the following.

Figure 23: Finding the AHV Version on Prism Central

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275777446001 on Brightcove.

Procedure

1. Log on to Prism Central.
2. In side bar, select **Hardware** > **Hosts** > **Summary** tab.
3. Click the host you want to see the hypervisor version for.
4. The **Host** detail view page displays the **Properties** widget that lists the Hypervisor Version.

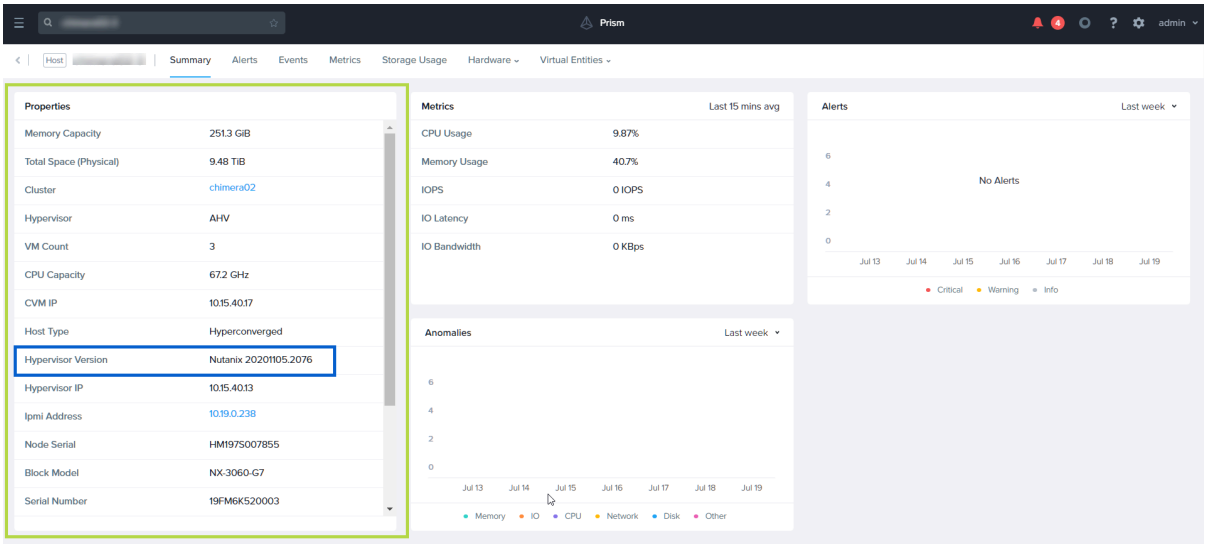

Figure 24: Hypervisor Version in Host Detail View

Finding the AOS Version Using Prism Central

To view the Nutanix AOS version running in the Prism Central, do the following:

Procedure

1. Log on to Prism Central.
2. Click the entities menu (hamburger menu) near the search field and then select **Hardware > Clusters** to view the cluster dashboard.
3. In the **List** tab (default), the AOS version of the cluster is listed against the cluster name.

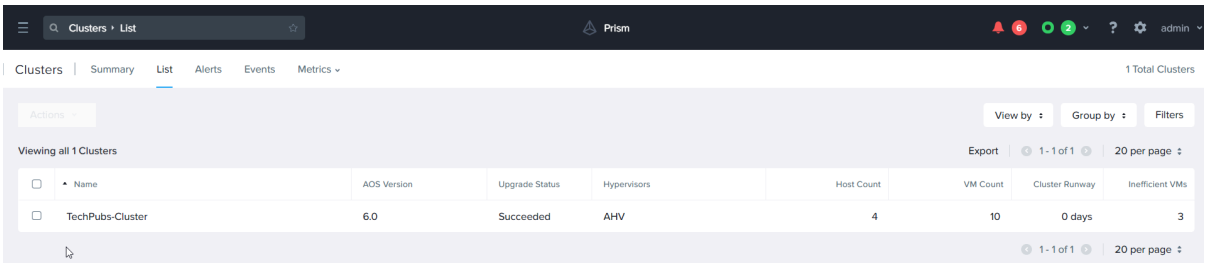

Figure 25: AOS Version in a Cluster

Registering or Unregistering a Cluster with Prism Central

Before you begin

- If you have never logged into Prism Central as the user admin, you need to log in and change the password before attempting to register a cluster with Prism Central.
- Do not enable client authentication in combination with ECDSA certificates on a registered cluster since it causes interference when communicating with Prism Central.
- Ports 9440 and 80 need to be open in both directions between the Prism Central VM and all the Controller VMs (and the cluster virtual IP address if configured) in each registered cluster. For the complete list of required ports, see Port Reference.
- A cluster can register with just one Prism Central instance at a time. To register with a different Prism Central instance, first unregister the cluster.

About this task

To register a cluster with Prism Central, do the following:

Procedure

1. Log in to the Prism web console on the target cluster as the user admin.
2. In the Health dashboard, from the **Actions** drop-down menu, select **Run Checks**. to run Nutanix Cluster Checks.

3. Do one of the following:

   » On the **Home** dashboard, click **Register or create new** from the **Prism Central** widget.

   » Click the gear icon and select **Prism Central Registration** from the **Settings** menu.

Figure 26: Home Dashboard

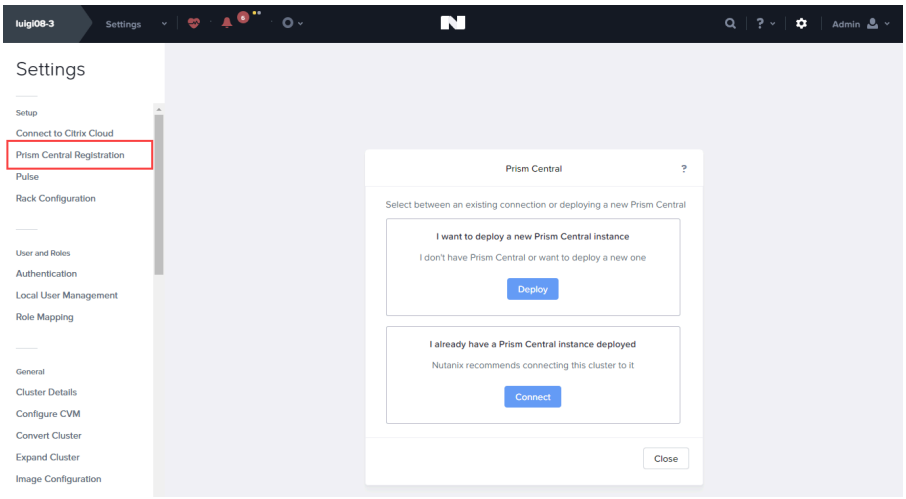

Figure 27: Settings Menu

4. In the first screen of the **Prism Central** dialog box, click the **Connect** button.

   This screen includes two options, **Deploy** and **Connect**. This procedure describes how to connect to an existing Prism Central instance. See *Installing Prism Central (1-Click Method)* section in Acropolis Upgrade Guide for

instructions on how to deploy a new Prism Central instance. For more information on how to deploy a new Prism Central, see Installing Prism Central (1-Click Method) in *Acropolis Upgrade Guide*.

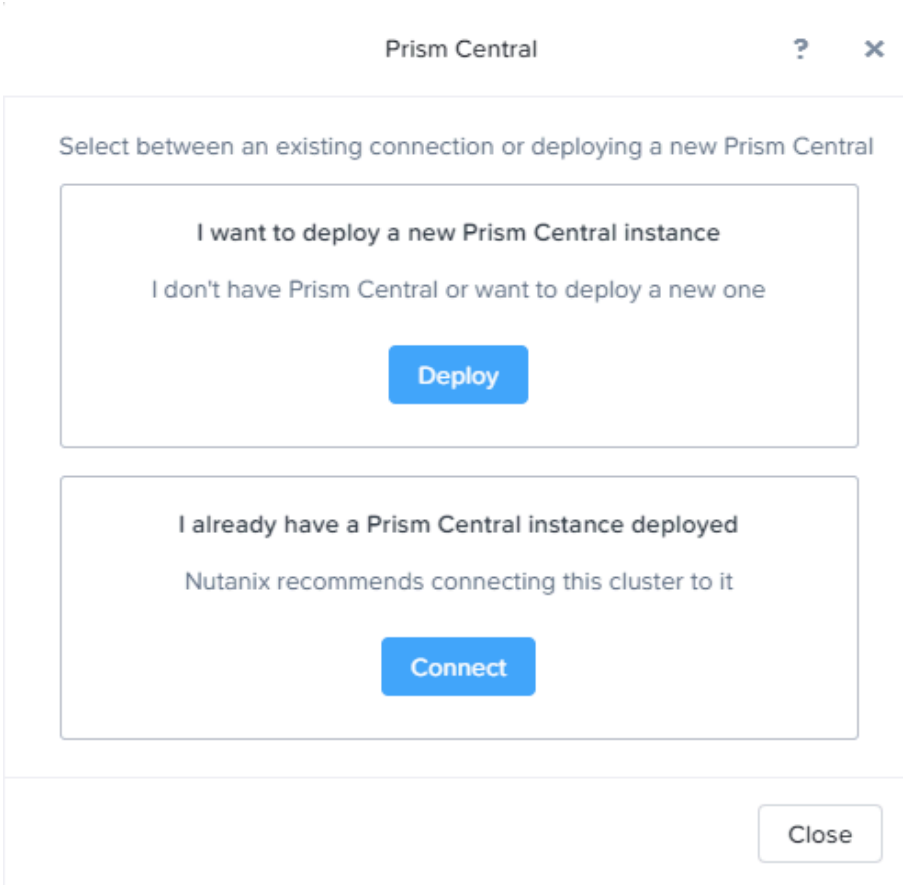

Figure 28: Prism Central Dialog Box (deploy/connect screen)

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

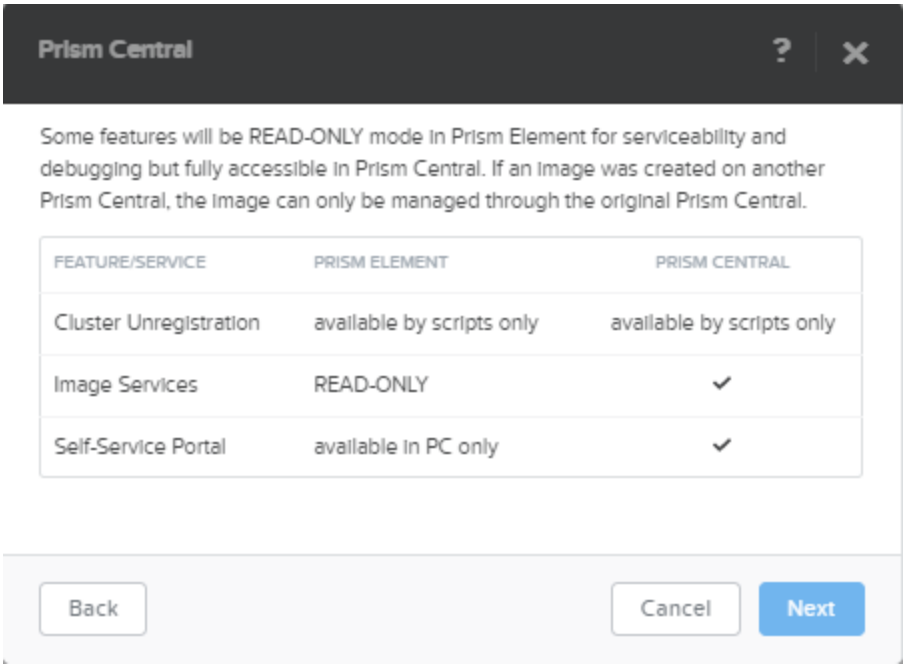

Figure 29: Prism Central Dialog Box (services screen)

6. In the registration screen, do the following in the indicated fields:

   a. **Prism Central IP**: Enter the IP address of the Prism Central VM.

   b. **Port**: The default port number is 9440. This is an optional field. For the complete list of required ports, see Port Reference.

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

   d. **Password**: Enter the Prism Central admin user password.

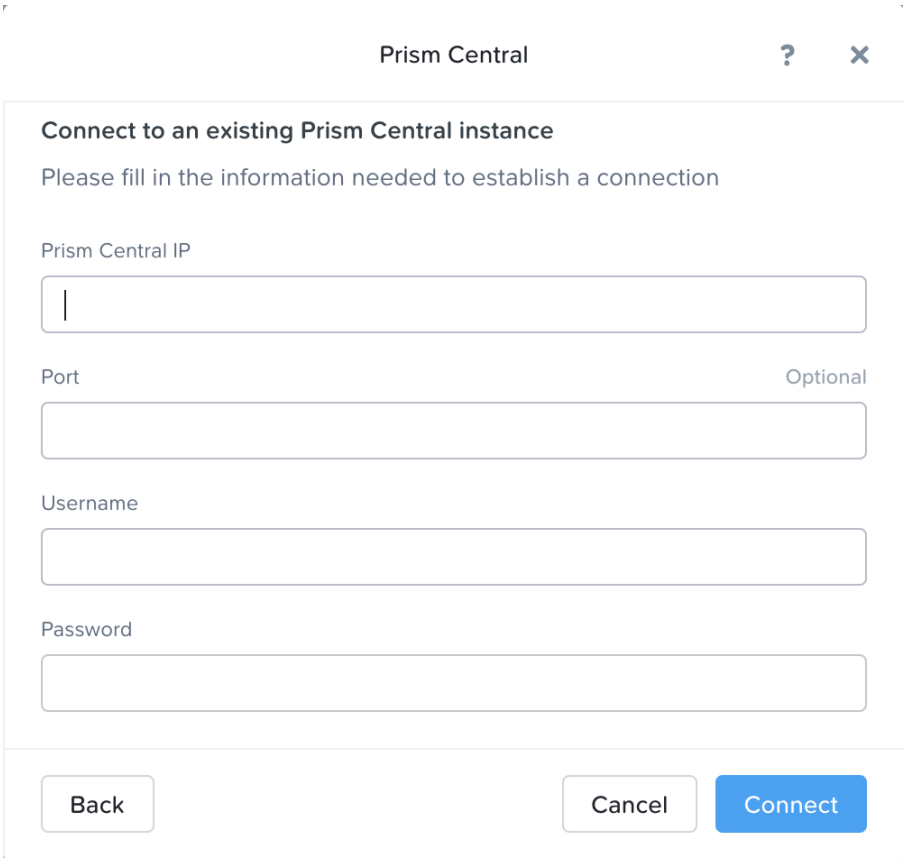

Figure 30: Prism Central Dialog Box (registration screen)

7. When all the fields are correct, click the **Connect** button to save the values and close the window.
   This registers the cluster on the specified Prism Central VM and allows the passing of information between the cluster and Prism Central.

   Note:

   - The user credentials provided when registering a cluster (Prism Element) with Prism Central are only used once. After registration, modifying the admin password would not impact any communication between Prism Central and the cluster.
   - On small, large, and x-large Prism Central deployments, when you register a new cluster to Prism Central, Prism Central synchronises the past 90 days of data (including multiple metrics) from the cluster. To view the list of metrics that are synced during registration, see */home/nutanix/config/arithmos/data_sender/arithmos_history.json*. To view the list of metrics that are synced during a regular synchronisation between Prism Central and the cluster, see */home/nutanix/config/arithmos/data_sender/arithmos.json*.

Unregistering from Prism Central

Before you begin

- Unregistering a cluster through the Prism web console is no longer available. This option is removed to reduce the risk of accidentally unregistering a cluster. Several features such as 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, it leads to features unavailability and configuration erasure. You can only use the following procedure from Controller VM (CVM) to unregister a cluster.
- Perform the entire registration process, followed by the cleanup process.
- Do not remove the IPs of the cluster and Prism Central from the whitelists on both sides until the unregistration process completes successfully.

> **Important:** If you have enabled additional applications or features in Prism Central, see the following table for recommendations before you unregister a cluster. For more information, see KB 4944.

Table 11: Unregistering a cluster with additional applications enabled

Application/Feature	Recommendation
Nutanix Disaster Recovery (Leap)	Before unregistering the cluster from Prism Central, you must remove any Nutanix Disaster Recovery (Leap) configuration involving virtual machines or volume groups for the cluster being unregistered. Otherwise you will not be able to manage the snapshot creation or replication policies configured on the cluster. Refer to KB 12749 for more details. **Note:** Do not proceed with unregistration, if the stretch config (AHV Synchronous replication) is present.
Flow Networking	You must disable Flow Networking on the cluster before unregistering it from Prism Central, using the steps mentioned in the Unregistering a PE from the PC in *Flow Virtual Networking Guide*. If Flow Networking is not disabled on the cluster prior to unregistering from Prism Central, attempts to enable Flow Networking in the same cluster does not work as expected. Refer to KB 12449 for more details.
NuCalm/App Management	You must clean up Calm entities after unregistration. Contact Nutanix Support for assistance in cleaning up the CALM entities.
Prism Self Service configuration	Changes that have been made to the Prism Self Service configuration in Prism Central are lost after unregistration. Ensure that you follow the extra cleanup steps mentioned in KB 4944.
NKE Kubernetes	Do not unregister the cluster hosting an NKE Kubernetes cluster from Prism Central. Unregistration of cluster from Prism Central will prevent the management of the NKE clusters.

About this task

To unregister a cluster from an existing Prism Central instance, perform the following steps:

Procedure

1. Log on to any Controller VM of the registered cluster through an SSH session.
2. Run the cluster status command and verify that all services are in a healthy state.

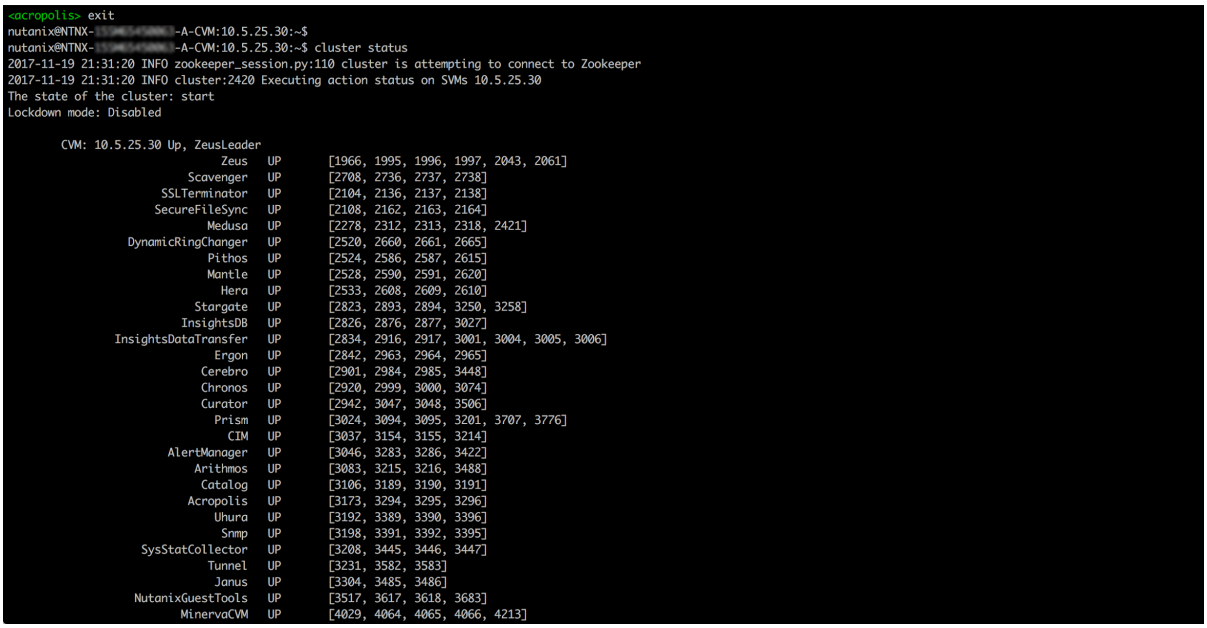

Figure 31: Cluster Status Command

3. Unregister the cluster from Prism Element.

```
nutanix@cvm$ ncli multicluster remove-from-multicluster external-ip-address-or-svm-ips=pc-name-or-ip username=pc-username password=pc-password force=true
```

Replace *pc-name-or-ip* with the Prism Central name or IP address and *pc-username* and *pc-password* with the login credentials for your Prism Central administrator account. This step can take some time (though typically just a few seconds). To check if the unregistration has completed, run the following command.

```
nutanix@cvm$ ncli multicluster get-cluster-state
```

4. Get the UUID for the cluster.

   ```
   nutanix@cvm$ ncli cluster info
   ```

   Find the "Cluster UUID" value in the displayed information.

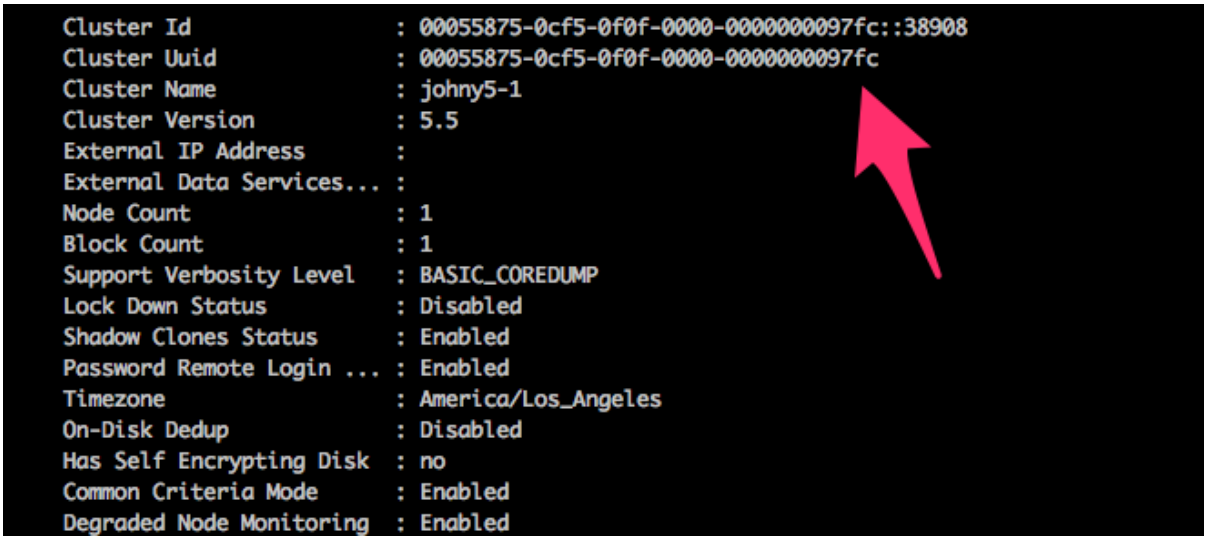

Figure 32: Cluster Info Output

5. Log in to the Prism Central VM through an SSH session (as the "nutanix" user) and perform the following steps:

   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.

      > **Note:** If you do not run the clean-up script, some artifacts continue to retain references to entities that are no longer managed by the cluster. Some artifacts that might have lost references due to the unregistration process might not be able to recover their references.

   b. Get the UUID for Prism Central.

      ```
      [pcvm]$ ncli cluster info
      ```

      Find the "Cluster UUID" value in the displayed information (see step 4), which in this case is the UUID for Prism Central.

6. Go back to the Controller VM and run the unregistration_cleanup.py script to complete the unregistration process on the cluster.

   ```
   nutanix@cvm$ python /home/nutanix/bin/unregistration_cleanup.py uuid
   ```

   In this case the *uuid* is the Prism Central UUID obtained in step 5b.

What to do next

After performing these steps you can re-register the cluster with a new or re-created Prism Central instance.

If the clean up process does not complete successfully, try the following:

- Check the logs to indicate if there are any input errors when calling the script. The logs for the unregistration cleanup script can be found under ~/data/logs/unregistration_cleanup.log.

- If errors occur during script execution, run the cluster status command and check that the cluster services are up and running. Rerun the script and check if it succeeds.

> **Note:** If Prism Self Service is enabled and the unregistered cluster is set as the default cluster, you must update the default project to select a new (registered) cluster. See the topic *Modifying A Project* in the Prism Central Guide.

Expanding (Scale Out) Prism Central

Before you begin

The following requirements must be met before you can expand Prism Central or add a Prism Central VM:

- The specified gateway must be reachable.
- No duplicate IP addresses can be used.
- The container used for deployment is mounted on the hypervisor hosts.
- When installing on an ESXi cluster:
  - vCenter and ESXi cluster must be configured properly. For more information about vCenter and ESXi configuration, see the vCenter Configuration topic in *vSphere Administration Guide for Acropolis*.
  - vCenter must be registered in Prism.
  - DRS must be enabled in vCenter.
  - vCenter is up and reachable during the deployment.

About this task

If Prism Central is just a single VM currently, you can expand it to three VMs. This increases both the capacity and resiliency of Prism Central (at the cost of maintaining two additional VMs). To expand this Prism Central instance across multiple VMs, do the following:

> **Note:**
>
> - Scale out Prism Central is supported on AHV and ESXi clusters only. If you are expanding a pre-AOS 5.17 Prism Central "large" VM that contains just a single vDisk, the upgrade process adds three new vDisks to each VM. This allows Prism Central to shard Cassandra metadata to improve performance. While the listed disk capacity of a Prism Central VM includes all the vDisks, the additional vDisks do not impact overall capacity. During the scale out you can monitor progress ("VM disk attach" tasks) from the Tasks dashboard.
> - All scale out Prism Central VMs must run on the same cluster. For example, running two VMs in cluster_1 and one VM in cluster_2 is not supported.

Procedure

1. Click the gear icon and then select **Prism Central Management** from the **Settings** menu (see Settings Menu (Prism Central) on page 16).
   The **Manage Prism Central** page appears. This page provides information about this Prism Central instance.

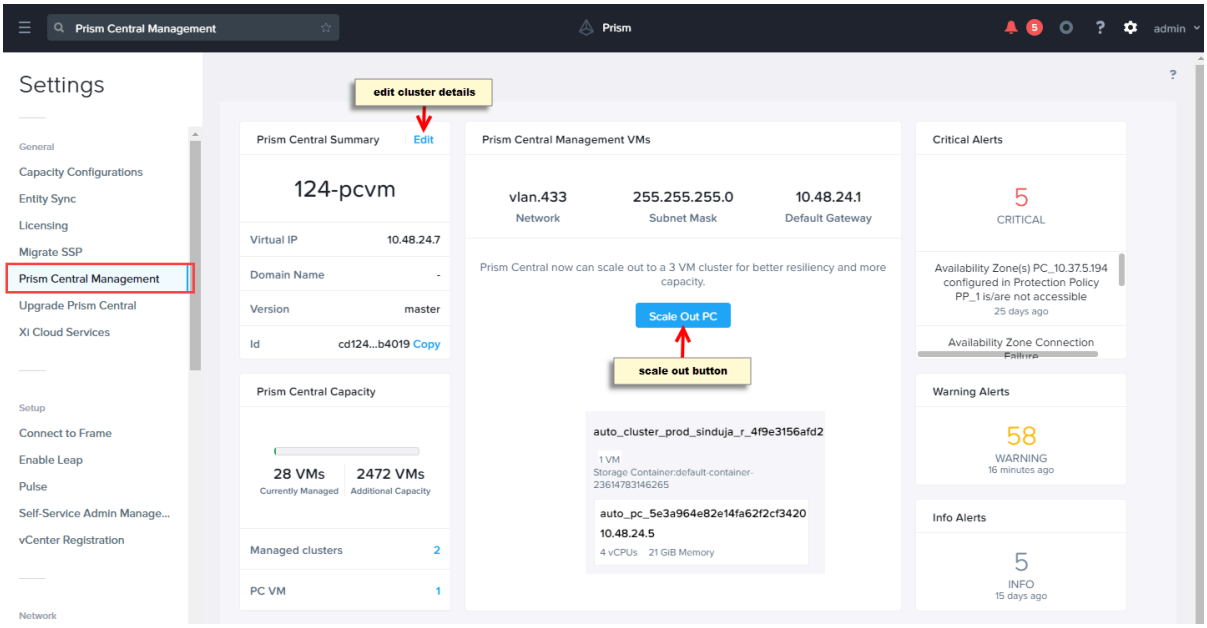

Figure 33: Manage Prism Central Page

2. To expand this Prism Central instance from one to three VMs, click the **Scale Out PC** button to display the **Scale Out PC** page and do the following:

> **Note:** A pop-up window appears explaining that scale out is a one-way process. Click the **Continue** button to display the **Scale Out PC** page. Once you scale out a Prism Central instance from a single VM to multiple VMs, you cannot revert back. Deleting any of the Prism Central VMs may result in data loss.

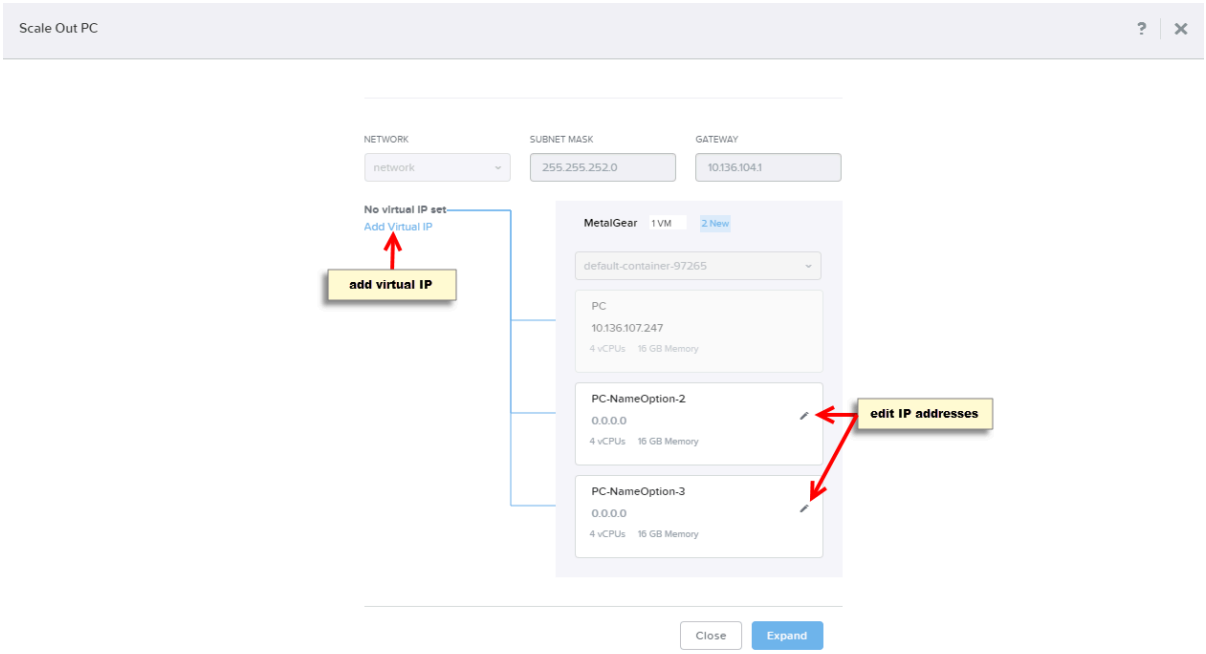

Figure 34: Scale Out PC Page

> **Note:** The Prism Central VM image is deployed from the target cluster. If the required image cannot be accessed, typically because there is no Internet access (such as at a dark site), a pop-up error message appears. In this case you have the option to manually download and deploy the image as follows:
>
> 1. Log on to the Nutanix customer support portal, click **Downloads** > **Prism Central**, and click the download link for the target version to save the Prism Central binary .TAR and metadata .JSON

files on your local media. The binary .TAR file is from the Prism Element bundle, used for 1-click deployment.

2. Log on (using SSH) to any Controller VM in the cluster specified in the error message and copy the Prism Central binary .TAR and metadata .JSON files to the Controller VM.
3. Run the following command to deploy the Prism Central image:

   ```
   nutanix@cvm$ ncli software upload software-type=PRISM_CENTRAL_DEPLOY file-path=file_path meta-file-path=metadata_file_path
   ```

   The *file_path* is the full (absolute) path to the .TAR file, and the *metadata_file_path* is the full path to the .JSON file. After this step completes, you can continue the scale out procedure.

a. Review the **Network**, **Subnet Mask**, and **Gateway** fields, which display the network name, subnet mask value, and gateway IP address applied to this Prism Central instance.

   The values in these three fields are read-only and cannot be changed.

   Note: The scale-out process uses the IP addresses from the network configured. If you used the default network configuration being the Private Network [default] while enabling microservices infrastructure, Prism Central displays and uses Private Network [default] details in this section.

b. [optional] To specify a virtual IP address for Prism Central, click the **Add Virtual IP** link.

   This displays the **Add PC Virtual IP** window. A virtual IP can be used as a single point of access for Prism Central. Enter the IP address in the Virtual IP field and then click the **Update** button.

c. Specify IP addresses for the two new Prism Central VMs.

   1. Click the pencil icon for one of the new VMs. (The VM names are set automatically.) The **IP field** opens for editing. Enter the IP address and then click the green check mark icon to save that address.

Figure 35: Edit IP Field

   2. Repeat this step for the second new VM.

   The vCPU count and memory size in the new VMs are fixed and match the current Prism Central VM values; you cannot change these values.

d. When all the parameters are correct, click the **Expand** button.

   This starts the process of creating the new VMs and deploying this Prism Central instance. You can monitor progress from the **Tasks** page (see Tasks View on page 244).

What to do next

In some cases, especially when using Calm, a version mismatch with a linked product could occur after expanding Prism Central. To fix this situation, perform a life cycle manager (LCM) inventory after completing the Prism Central expansion procedure. See the LCM documentation for instructions on how to perform an LCM inventory.

Shutting Down (Starting Up) Prism Central VM

About this task

A Prism Central VM (PC VM) is managed like any other VM in a cluster through the Prism Element web console for the cluster in which the Prism Central VM is running. However, shutting down a PC VM requires extra caution.

> **Note:**
>
> Prism Central supports features that could be damaged by shutting down the PC VM abruptly. You must follow the steps carefully to avoid any issues

To shut down a single PC VM in a scale-out PC configuration, log on to the web console of the host and go to the VM dashboard (see VM Dashboard in Prism Element Web Console Guide). In the VM Table view, select the PC VM you want to shut down, and then select **Power Off Actions >> Power Off** or **Power Off Actions >> Guest Shutdown** from the action links. Alternatively, you can SSH to the IP address of the PC VM and run the command to shut down the PC VM.

To shut down the PC VM in single VM configuration, or to shut down all the three PC VMs in a scale out configuration, perform the following steps:

Procedure

1. SSH to the IP address of the PC VM that you want to shut down.
2. Stop the cluster.

   ```
   nutanix@pcvm$ cluster stop
   ```

3. Check the status of the services running on the cluster.

   ```
   nutanix@pcvm$ cluster status
   ```

   The command output should show the following three services still running: Zeus, Scavenger, and VipMonitor.

4. To shut down a PC VM, log on to the web console of the host and go to the VM dashboard (see VM Dashboard in Prism Element Web Console Guide). In the VM Table view, select the PC VM you want to shut down (one at a time), and then select **Power Off Actions >> Power Off** or **Power Off Actions >> Guest Shutdown** from the action links.

   Alternatively, you can SSH to the IP address of the PC VM and run the command to shut down the PC VM.

   ```
   nutanix@pcvm$ sudo shutdown -h now
   ```

   You must run this command for each PC VM that you want to shut down.

5. To power on a PC VM, log on to the web console of the host and go to the VM Dashboard (see VM Dashboard). In the VM Table view, select the PC VM you want to power on (one at a time), and then select **Power On**.
6. Once all the PC VMs are powered on, check the status of the services from any one of the PC VMs.

   ```
   nutanix@pcvm$ cluster status
   ```

   In the command output, verify that all the three services (Zeus, Scavenger, and VipMonitort) are running.

7. Once all the three services start running, start the cluster from any one of the PC VMs.

   ```
   nutanix@pcvm$ cluster start
   ```

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.

- Nutanix Disaster Recovery
- Flow microsegmentation
- Prism Pro (AIOps)
- VM management
- Cluster management
- Identity and Access Management (IAM)

  > **Note:** IAM is recovered only after you enable microservices infrastructure.

- Categories
- Networking
- Licensing
- Reporting templates

Although the following services continue to run, they are not protected and therefore data associated with them not recovered.

- Flow Virtual Networking (Virtual Private Clouds (VPCs) and Virtualized Networking based Advanced Networking Controller (ANC))
- Calm
- LCM
- Karbon
- Objects
- Files
- Catalog
- Images
- VM templates

For Nutanix Disaster Recovery Recovery Jobs in progress, see this Note.

> **Note:** If you have enabled PCDR, Prism Central pauses the backup of Prism Central configuration data when Prism Central upgrade is initiated. The backup resumes after the upgrade process is complete.

You can protect a Prism Central deployment with the GUI-based Prism Central disaster recovery solution described in this document.

> **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 or back up a Prism Central deployment to its registered Nutanix clusters that fulfill the requirements. The system enables you to back up a Prism Central deployment to at most three AHV or ESXi clusters. Clusters running Hyper-V are not supported.

Before you begin

Ensure that the following requirements are met before protecting Prism Central.

- Prism Central must run on a minimum version of pc.2021.7 or later.
- Prism Central must have an NTP configuration to synchronize time between Prism Central and the registered clusters. For more information on NTP configuration, see Configuring NTP Servers (Prism Central) in the *Prism Central Guide*.
- The Nutanix clusters used to back up the Prism Central deployment must be running a minimum version of AOS 6.0 or later.
- > **Note:** At least one cluster used to protect (back up) Prism Central must be running AOS 6.0 or later because Prism Central can be recovered only on clusters running AOS 6.0 or later. For more information, see Recovering Prism Central (1-Click Disaster Recovery) on page 63.
- The Prism Central instance must be registered to the cluster hosting the Prism Central VM. If PC is not registered to the cluster then, enabling Prism Central Disaster Recovery fails.

About this task

To protect a Prism Central deployment, perform the following steps.

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 the Prism Central deployment, and then click **Proceed**. The system synchronizes Prism Central configuration data to the selected clusters. The first backup creation on the selected clusters takes at least 30 minutes. After the first backup, the system synchronizes Prism Central configuration data with the selected clusters every 30 minutes.
6. (optional) Click **+ Add Backup** after the configuration, to backup Prism Central to more clusters.

   However, ensure that you do not backup Prism Central to more than three clusters.
7. (optional) Click **Remove** after the configuration, to remove a cluster from backing up Prism Central.

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 the cluster you use to recover Prism Central meets the following requirements:

- Registered to the protected Prism Central.
- Running AOS 6.0 or newer.
- Configured with ISCSI data service IP address for efficient recovery of Nutanix Disaster Recovery or Calm service configurations.

> **Note:** The cluster takes about 10 minutes to stabilize Prism Central after the recovery task on Prism Element is displayed as complete. This is because Microservices Infrastructure rotates the certificates after completion of PCDR and restarts services like IAM and Flow Virtual Networking.

About this task

To restore a Prism Central deployment, perform the following steps.

> **Important:** To do a planned restore of a Calm-enabled Prism Central to a recovery cluster, perform the following steps:
>
> - **1.** Manually backup the Calm-related entities in Prism Central at the primary cluster. For more information, see Backing up Calm Data in *Calm Administration and Operations Guide*.
>   **2.** Restore Prism Central in the recovery cluster.
>
>   > **Caution:** Before you restore Prism Central in a recovery cluster, ensure the Prism Central backup is synched at least once with the recovery cluster you select to restore Prism Central.
>
>   **3.** Restore the Calm-related entities in the new Prism Central at the recovery cluster. For more information, see Restoring Calm Data in *Calm Administration and Operations Guide*.

Procedure

1. Log on to any Prism Element registered to Prism Central that you want to recover.
   The Prism Element dashboard shows the Prism Central widget that contains the Prism Central information (IP address and connection status).
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 those that will not.

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**.

   The Prism Central deployment restores in at least 1 to 2 hours, depending on the configuration data of the hosts. For example, CMSP enabled Prism Central restores in at least 2 hours. The restored Prism Central takes an additional 30 to 40 minutes to show all the guest VMs, disks, and metrics. Wait to perform any actions on the restored Prism Central until all the recovery tasks are completed on the cluster. You can see the restoration status and the related processes in the **Tasks** window.

What to do next

Consider the following after Prism Central restoration.

- Use the newly restored Prism Central deployment 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.

  > **Note:** If the Prism Central restoration fails, contact Nutanix Support. Refrain from bringing up the old Prism Central deployment.

- Reset the credentials.

  Prism Central restores with the default credentials. Nutanix recommends changing the default credentials. For information about changing the default credentials, see Logging Into Prism Central in the *Prism Central Infrastructure Guide*

- Reconfigure the proxy server. For information on how to configure the HTTP proxy through the Prism Central web console, see Configuring an HTTP Proxy in the *Prism Central Admin Center Guide*

  If the old Prism Central deployment had a proxy server, reconfigure the proxy server so that the recovered Prism Central deployment 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.

- Recovery Plan Jobs (RPJ) in progress: Perform the steps mentioned in KB-10962.

  If the old Prism Central had a failover task running (Nutanix Disaster Recovery) 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 deployment had Microservices infrastructure enabled, see Enabling Microservices Infrastructure to enable microservices infrastructure on the recovered Prism Central.

IP Address Reconfiguration

You can use this procedure to reconfigure the IP address and gateway of single Prism Central instances as well as Prism Central VMs in a Scale out Prism Central (clustered PC VMs or PC cluster).

> **Note:** Reconfiguring the IP address and gateway of Prism Central VMs does not require additional steps when using Nutanix Disaster Recovery (formerly Leap).

Preparing to Reconfigure the IP Address and Gateway of PC VMs

Perform the tasks described in this topic before you start the IP address reconfiguration procedure.

About this task

You must perform the following tasks while the Prism Central VMs are still on the existing IP addresses.

Procedure

1. Coordinate Prism Central downtime, because the features and functionality of Prism Central will be unavailable for the entire duration of the IP address reconfiguration procedure.
2. Create a table to map the existing IP addresses with the new IP addresses of the Prism Central VMs for your reference.
3. Use SSH to log on to any running Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.
4. Verify if the PC VM or PC cluster is in a stable state.

   ```
   nutanix@pcvm$ cluster status
   ```

   This command is valid even if you only have a single PC VM and do not have a PC cluster.
5. Run the NCC health checks to make sure that the PC cluster is in a healthy state.
6. If you want to reconfigure the IP addresses of the entities in the AOS clusters (such as IP addresses of CVMs and hypervisor hosts), on which the PC is hosted, reconfigure those IP addresses first before you reconfigure the IP addresses of the PC VMs.
7. Stop the PC VM or PC cluster.

   ```
   nutanix@pcvm$ cluster stop
   ```

   This command is valid even if you only have a single PC VM and do not have a PC cluster.

   Wait to proceed until an output showing all the services as DOWN is displayed, except the Zeus and Scavenger services.

Reconfiguring the IP Address and Gateway of Prism Central VMs

Before you begin

Ensure that you have completed the tasks described in Preparing to Reconfigure the IP Address and Gateway of PC VMs on page 65.

About this task

The procedure to reconfigure the IP addresses of PC VMs differs depending on whether or not IP Address Management (IPAM) is enabled in an AHV cluster. The procedure to reconfigure the IP addresses of PC VMs in an ESXi cluster is the same as the procedure of an AHV cluster in which IPAM is disabled.

IP Address Management (IPAM) is a feature of AHV that allows it to assign IP addresses automatically to VMs by using DHCP. For more information, see IP Address Management section in the *AHV Administration Guide*.

In this document, managed network refers to an AHV network that has IPAM enabled, and unmanaged network refers to an AHV network that does not use IPAM.

> **Note:**
>
> - Reconfiguration of PC VM IP addresses is not supported for Prism Central scale-out deployments used in Nutanix Cloud Clusters (NC2) on AWS or Azure environment.
> - Do not reconfigure the IP address of the PC VM if any Nutanix Kubernetes Engine (NKE) or Objects are deployed on the cluster since changing the PC VM IP address is not supported for NKE or Objects after deployment.

Perform the following procedure to reconfigure the IP address of a single PC VM or PC VMs in a PC cluster.

Procedure

1. Log on as nutanix user and run the *external_ip_reconfig* script.

   ```
   nutanix@pcvm$ external_ip_reconfig
   ```

2. Follow the prompts to type the new netmask, gateway, and external IP addresses.

   A message similar to the following is displayed when the procedure begins:

   ```
   External IP reconfig started
   ```

   A message similar to the following is displayed if the procedure is completed successfully:

   ```
   External IP reconfig finished successfully. Restart all the CVMs and start the
    cluster.
   ```

   The message indicates that you must restart the CVMs. In this case, you are reconfiguring the IP addresses of the PC VMs, so you must restart the PC VMs and not the CVMs.

   > **Note:**
   >
   > If your PC VMs are running in a managed network of an AHV cluster, do not start the PC VMs yet but proceed to step 3.
   >
   > If your PC VMs are in an unmanaged network of an AHV cluster or are in an ESXi cluster, skip steps 3, 4, and 5, and proceed to step 6 directly.
   >
   > Perform steps 3, 4, and 5 only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   A message similar to the following is displayed if the procedure fails:

   ```
   External IP reconfig Failed
   ```

   If the procedure fails, run the following command to check the log for troubleshooting purposes:

   ```
   nutanix@pcvm$ cat ~/data/logs/ip_reconfig.log
   ```

3. (Optional) Use SSH to log onto any CVM of the AOS cluster that is hosting the PC cluster or PC VM.

   > **Note:** Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

4. (Optional) Update the NIC attached to the PC VM with the new IP address of the PC VM.

   > **Note:** Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   ```
   nutanix@cvm$ acli vm.nic_update pc-vm-name pc-vm-mac request_ip=true ip=pc-vm-ip
   ```

   Replace the variables with their appropriate values as follows:

   - *pc-vm-name*: Name of the Prism Central VM whose IP address you modified.
   - *pc-vm-mac*: MAC address of the Prism Central VM whose IP address you modified (press **Tab** after you type the name of the PC VM to automatically populate the MAC address).
   - *pc-vm-ip*: New IP address of the Prism Central VM.

5. (Optional) Log on to the Prism Element web console and modify the IP address range in the IP pool for the PC network to work with the modified IP address of Prism Central VMs. In the Prism Element web console, perform the following:

   > **Note:** Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your PC VM is in a managed network of an AHV cluster.

   a. In the **VM** dashboard, click the **Table** view.

   b. Select a **PC VM** and click **Update**.

   c. In the **Update VM** dialog box, scroll down to the **Network Adaptors (NIC)** section.

   d. In the **VLAN Name** column, note the names of the networks attached to your PC VM.

   e. Click **Close**.

   f. Click the gear icon in the top-right corner, and under **Settings**, click **Network Configuration**.

   g. In the **Network Configuration** dialog box, identify the network names you noted in step d.

   h. Select each network and click the pencil icon next to the network.

   i. Under **IP Address Pools**, modify the IP address range if the new IP address of the PC VM is in a different IP address range than the previous IP address range.

6. Restart each Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.

7. Once the PC VMs are powered on, use SSH to log on to each PC VM or the single PC VM if you do not have a PC cluster and verify if the output of the following verification script displays the new IP addresses:

   ```
   nutanix@pcvm$ python cluster/bin/external_ip_reconfig_verify.py
   ```

8. Start the PC cluster.

   ```
   nutanix@cvm$ cluster start
   ```

9. If you have enabled Microservices Infrastructure on the Prism Central, then perform the steps mentioned in KB-15469.

10. Perform the following verification steps:

    a. Log on to the Prism Element web console of every cluster attached to the PC VM whose IP address you changed.

    b. On the **Home** page, verify if the **PC Registration** status is displayed as **OK**. This might take a few minutes after the PC VM or PC cluster is powered on.

Pairing AZs (Nutanix Disaster Recovery)

To replicate entities (protection policies, recovery plans, and recovery points) to different on-prem AZs bidirectionally, pair the AZs with each other. To replicate entities to different Nutanix clusters at the same AZ bidirectionally, you need not pair the AZs because the primary and the recovery Nutanix clusters are registered to the same AZ (Prism Central). Without pairing the AZs, you cannot perform DR to a different AZ.

About this task

To pair an on-prem AZ with another on-prem AZ, perform the following procedure at either of the on-prem AZs.

Procedure

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

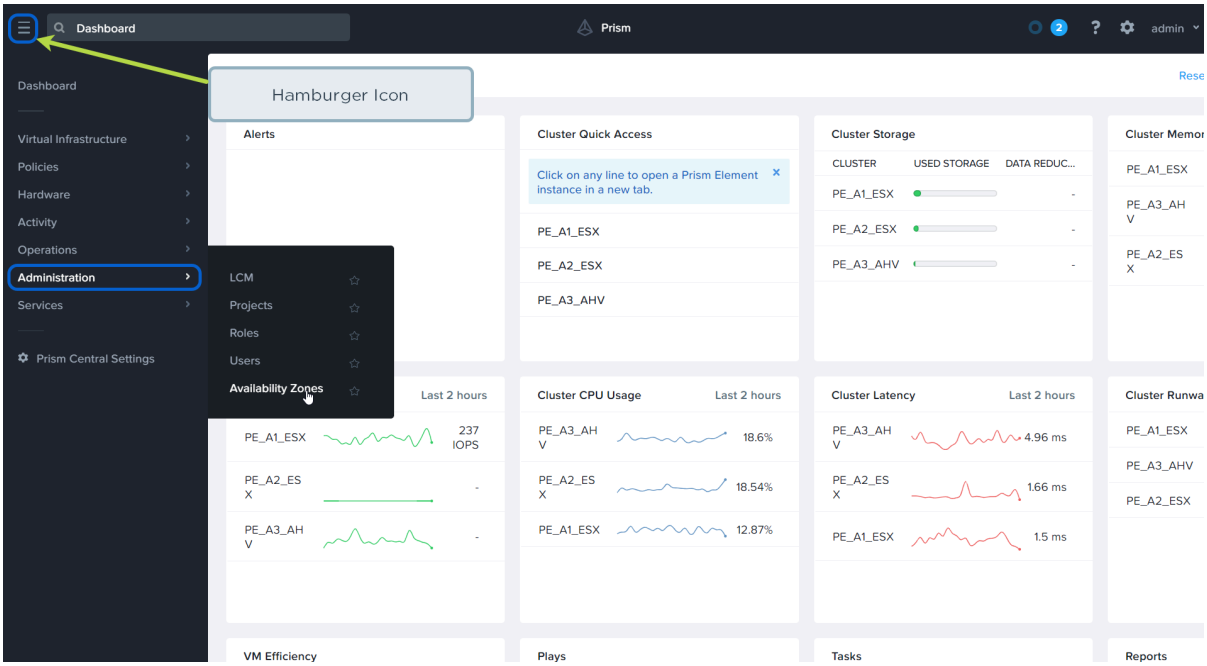

Figure 36: Pairing AZ

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

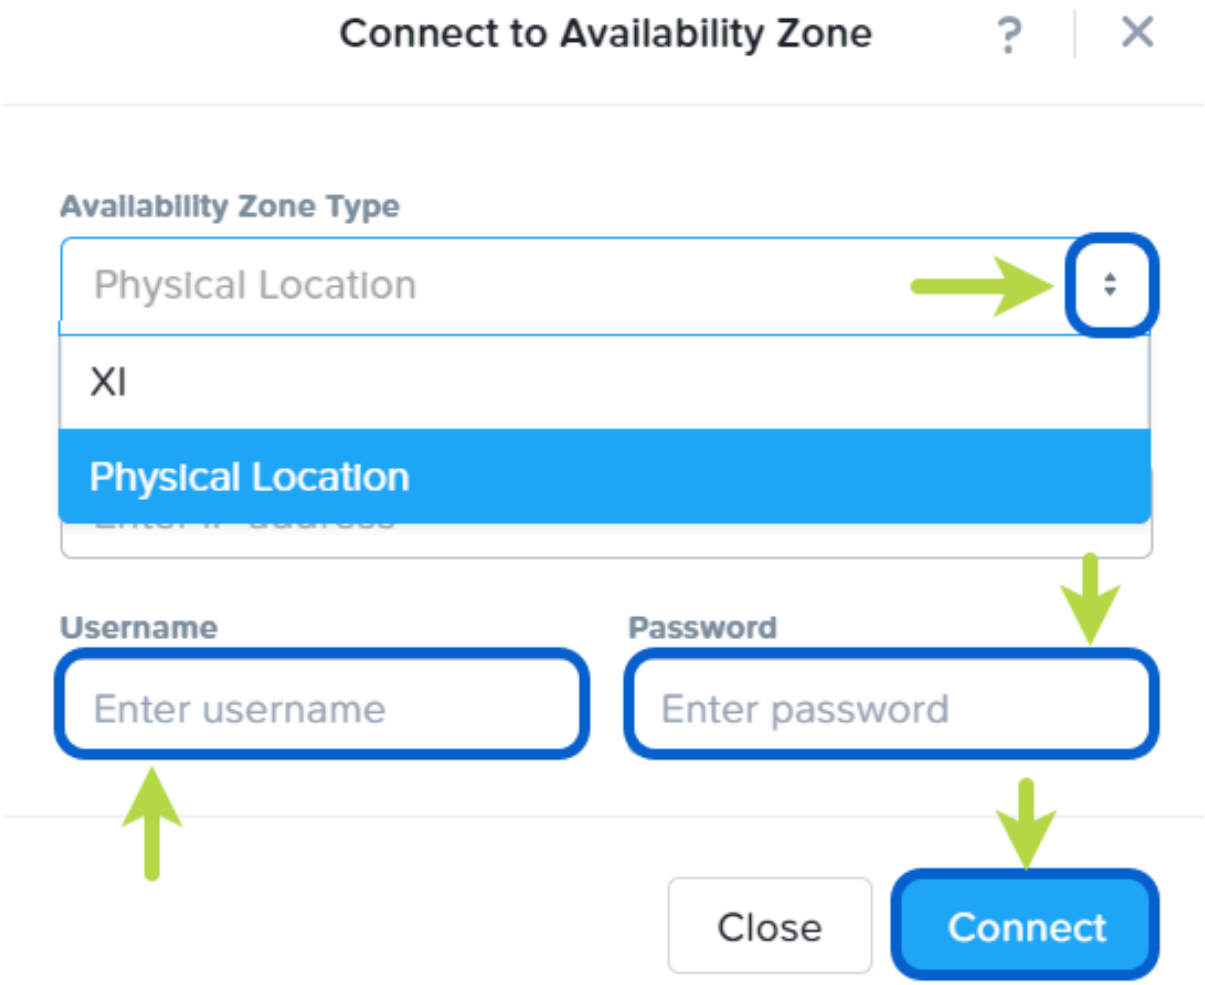

Figure 37: Connect to AZ

   a. **AZ Type**: Select **Physical Location** from the drop-down list.
      A physical location is an on-prem AZ. To pair the on-prem AZ with Xi Cloud Services, select **XI** from the drop-down list, and enter the credentials of your Xi Cloud Services account in step c and set d.
   b. **IP Address for Remote PC**: Enter the IP address of the recovery AZ Prism Central.
   c. **Username**: Enter the username of your recovery AZ Prism Central.
   d. **Password**: Enter the password of your recovery AZ Prism Central.

4. Click **Connect**.
   Both the on-prem AZs are paired to each other.

MAIN DASHBOARD

The main dashboard is the opening screen that appears after logging into Prism Central. It provides a dynamic summary view across the registered clusters. To view the main dashboard at any time, select **Dashboard** from the Entities menu (see Entities Menu on page 13 section).

> **Important:** Nutanix recommends using Google Chrome web browser to access Prism Central. The display of dashboard and widgets is optimized for Google Chrome browser.

Main Screen Layout

In addition to the main menu options (see Main Menu (Prism Central) on page 10 section), the main dashboard includes a menu bar and a display section for a collection of widgets that appear as tiles with targeted information about the registered clusters in each tile. The menu bar includes the following options:

- **Main Dashboard** tab. Click this tab to display the main dashboard. Additional tabs appear for any custom dashboards you create.
- **Manage Dashboard** button. Click this button to create a custom dashboard, edit the name of a custom dashboard, or delete a custom dashboard (see Creating a New Dashboard on page 79).
- **Reset Dashboard** button. Click this button to reset the main dashboard to the default set of widgets.
- **Add Widget** button. Click this button to add a widget to the displayed dashboard (see Adding Dashboard Widgets on page 80).

- **Data Density** button. Click this button to select the density of data representation in the widgets. Select one of the following three options:
  - Light—The information is lightly packed in the widget with more spacing between the elements.
  - Default—The information is moderately packed with sufficient spacing between the elements. This scheme of data density is based on user research by Nutanix.
  - Dense—The information is densely packed with minimal spacing between the elements for essential legibility.

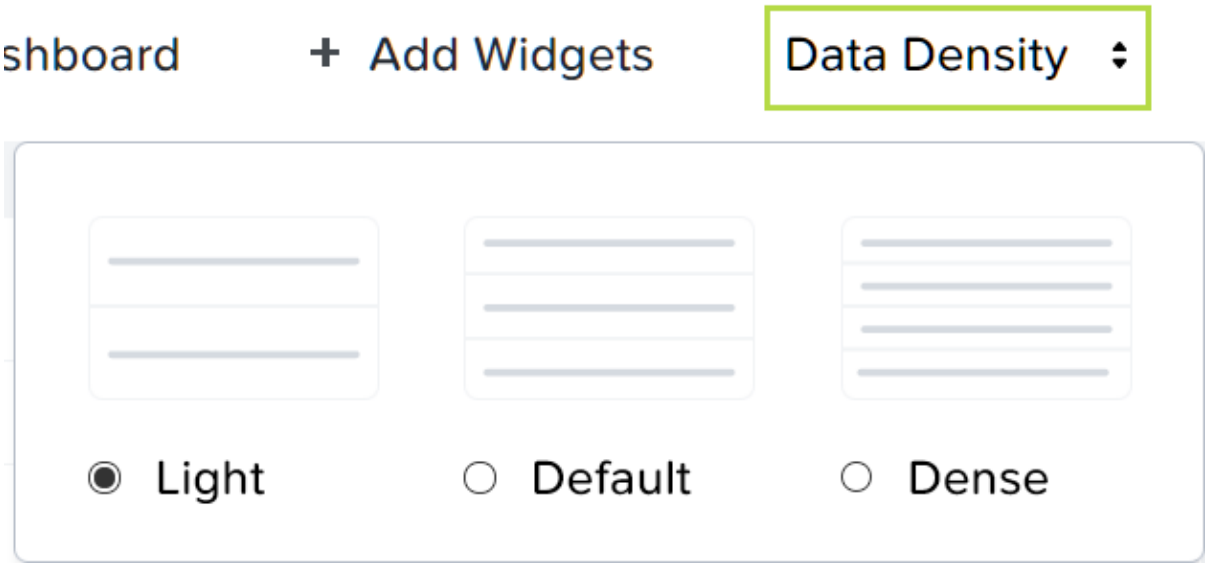

Figure 38: Data Density

Note: **Manage Dashboards** and **Add Widget** are Prism Pro license features. If Prism Pro is disabled, these buttons disappear, and you cannot customize or add dashboards. Any custom dashboards created before Prism Pro was disabled remain, but clicking the **Reset All** button (which replaces the **Reset Dashboard** button) deletes all custom dashboards and returns the main dashboard to the default state.

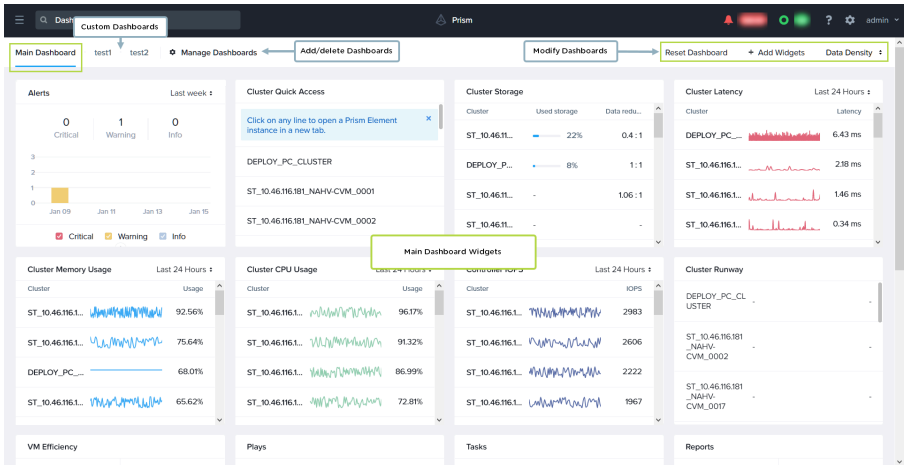

Figure 39: Main Dashboard

Main Dashboard Display

The main dashboard displays thirteen information tiles (widgets) by default. You can customize this view by adding or deleting widgets as desired (see Modifying a Dashboard on page 85).

The widgets are lined up in rows of four widgets per row. The default widgets are described in this section as they appear from left to right, top to bottom on the **Main Dashboard** page. A sample view of the **Main Dashboard** with the default widgets is displayed in the figure.

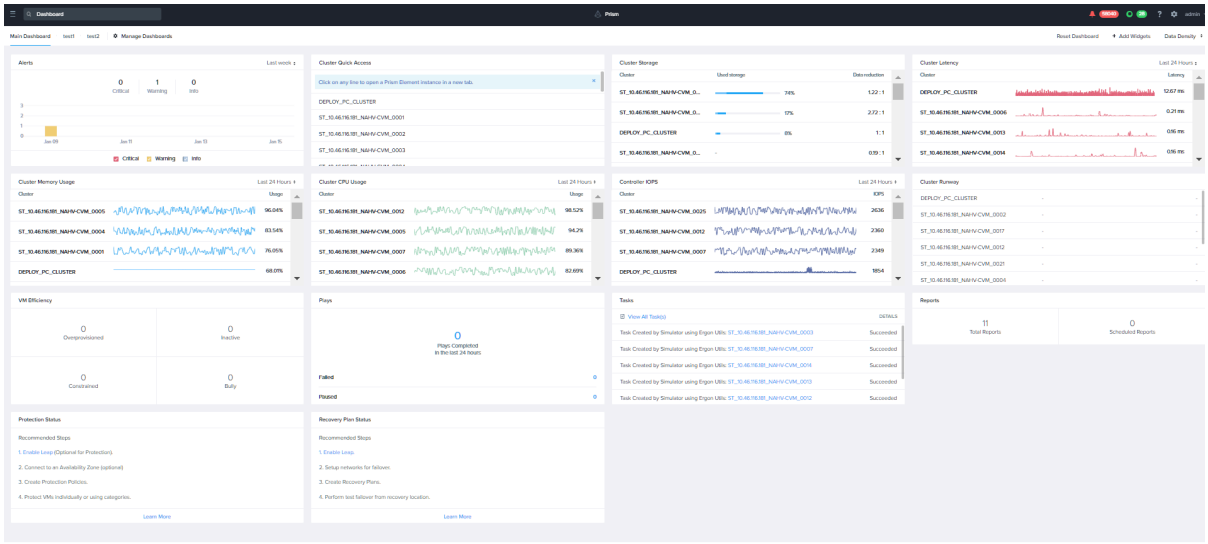

Figure 40: Main Dashboard with Default Widgets

The following figure displays the sample view of a widget with a period (time) selection:

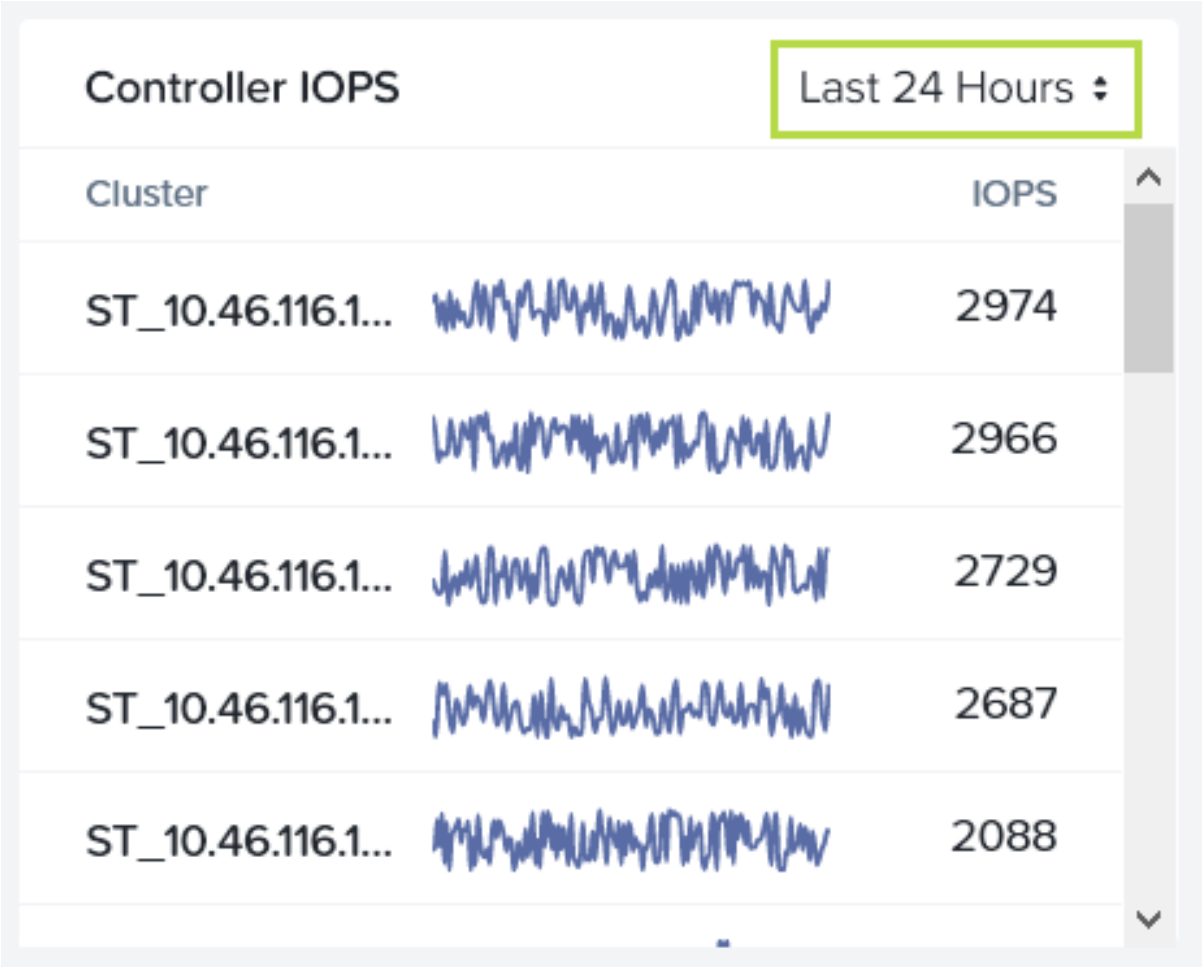

Figure 41: Sample Widget

- **Alerts**. This widget displays colored bar graph representing the alerts raised in the period selected. You can select the period on the drop down list at the top right corner of the widget. Each category of alerts is represented as a colored bar. Click any bar on the graph to see a list of the alerts represented in the bar. Click the **View All Alerts** link below the list to see all the alerts on the **Alerts** page.

  The Alerts categorized as **Critical**, **Warning** and **Info** are represented as red, yellow and blue bars respectively in the graphs.

- **Cluster Quick Access**. This widget displays a list of the clusters managed by the Prism Central. It allows you to open the Prism Element for a registered cluster in a new tab or window by simply clicking the line for that cluster. The cluster name appears on the left of each line while current health and alert icons (as appropriate) appear on the right.

- **Cluster Storage**. This widget displays storage and resiliency information for the highest usage clusters. Each line includes the cluster name, a usage column with a bar that visually indicates how much capacity is used currently, and a data resiliency column that displays the current status (critical, warning, OK, or unknown). The cluster list is ordered by data resiliency status with critical clusters at the top of the list. Click the cluster name to see the cluster details page. Hovering the cursor over the usage bar displays a box with usage details including data reduction ratio, usage percentage, used space, total space, fault domain (disk, node, or rack), and fault tolerance level (0, 1, or 2). The data reduction ratio indicates the data savings due to data reduction techniques such as deduplication, compression, and erasure coding. (A 1:1 ratio indicates none of these data reduction techniques are in use currently.)

- **Cluster Latency**. This widget displays the total (read and write) IO latency average for the highest latency clusters. Clicking on a cluster name displays the summary page for that cluster.
- **Cluster Memory Usage**. This widget displays the percentage of total memory in use currently for the highest usage clusters. Clicking on a cluster name displays the summary page for that cluster.
- **Cluster CPU Usage**. This widget displays the percentage of total CPU in use currently for the highest usage clusters (or all clusters if there are fewer than five). Clicking on a cluster name displays the summary page for that cluster.
- **Controller IOPS**. This widget displays the total (read and write) controller IOPS for the highest volume clusters. Clicking on a cluster name displays the summary page for that cluster. The IOPS number comes from the controller when the hypervisor is AHV or Hyper-V and from the hypervisor when the hypervisor is ESXi.
- **Cluster Runway**. This widget alerts you to potential storage, CPU, or memory resource constraints across the clusters and provides an estimated *runway* (time remaining) before the resources are maxed out based on current usage. Clicking on a cluster name displays the capacity page for that cluster.
- **VM Efficiency**. This widget displays the number of VMs that are considered inefficient broken down by category: overprovisioned, inactive, constrained, and bully. It includes a link to the VMs dashboard for details about these VMs (see VMs Summary View on page 94). For more information about VM efficiency, see Behavioral Learning Tools on page 899.
- **Plays**. This widget displays a list of the Plays running on the cluster. The list is categorised to completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the **Plays** dashboard page.
- **Tasks**. This widget displays a list of recent tasks with the current status of each task. Click on the View All Tasks(s) link to view the **Tasks** page.
- **Reports**. This widget displays a table that lists the number of total and scheduled reports with a link to the reports dashboard (see Reports Management on page 941).
- **Protection Status**. This widget displays the protection plan status if one has been created. If not, it displays a set of recommended steps to create a protection plan.
- **Recovery Plan Status**. This widget displays the recovery plan status if one has been created. If not, it displays a set of recommended steps to create a recovery plan.

Widgets on a Dashboard

Resize the Widgets

You can vertically resize of a widget by clicking and dragging down the resize handle at the bottom of the widget. You cannot resize a widget horizontally.

The following figure provides an example of the resize handle.

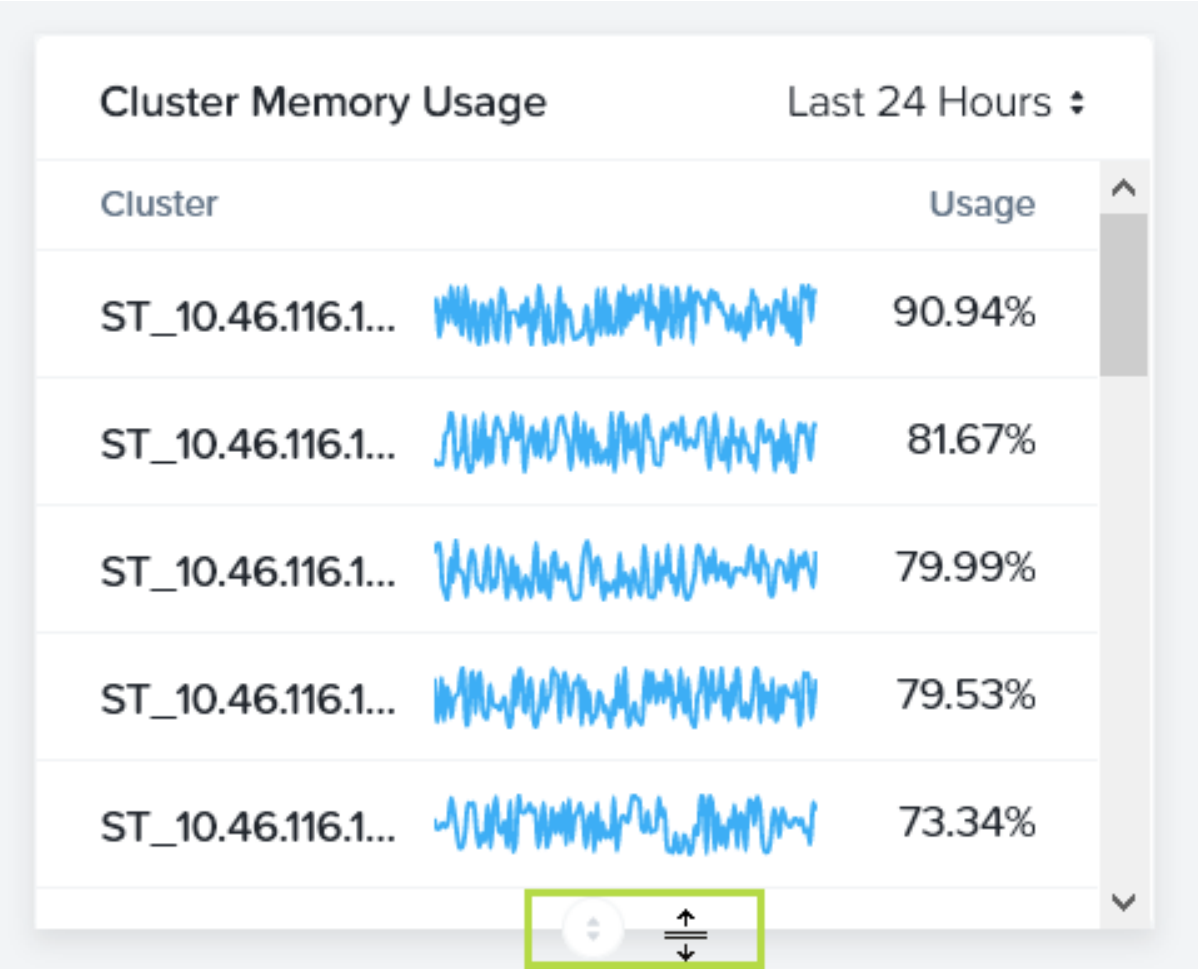

Figure 42: Vertically Resizing the Widget

When you click and drag up or down the resize handle to resize a widget, the widget placed below the resized widget shifts upward or downward.

You can resize a widget in multiples of half the default length of the widget, such as half, one-and-half, twice, two-and-half times the default length. Therefore, you can reduce the length to half or expand it to one-and-half times or twice the default length. Even if you attempt to resize it in any other variation of the default length, the widget auto-sizes itself to the nearest half length.

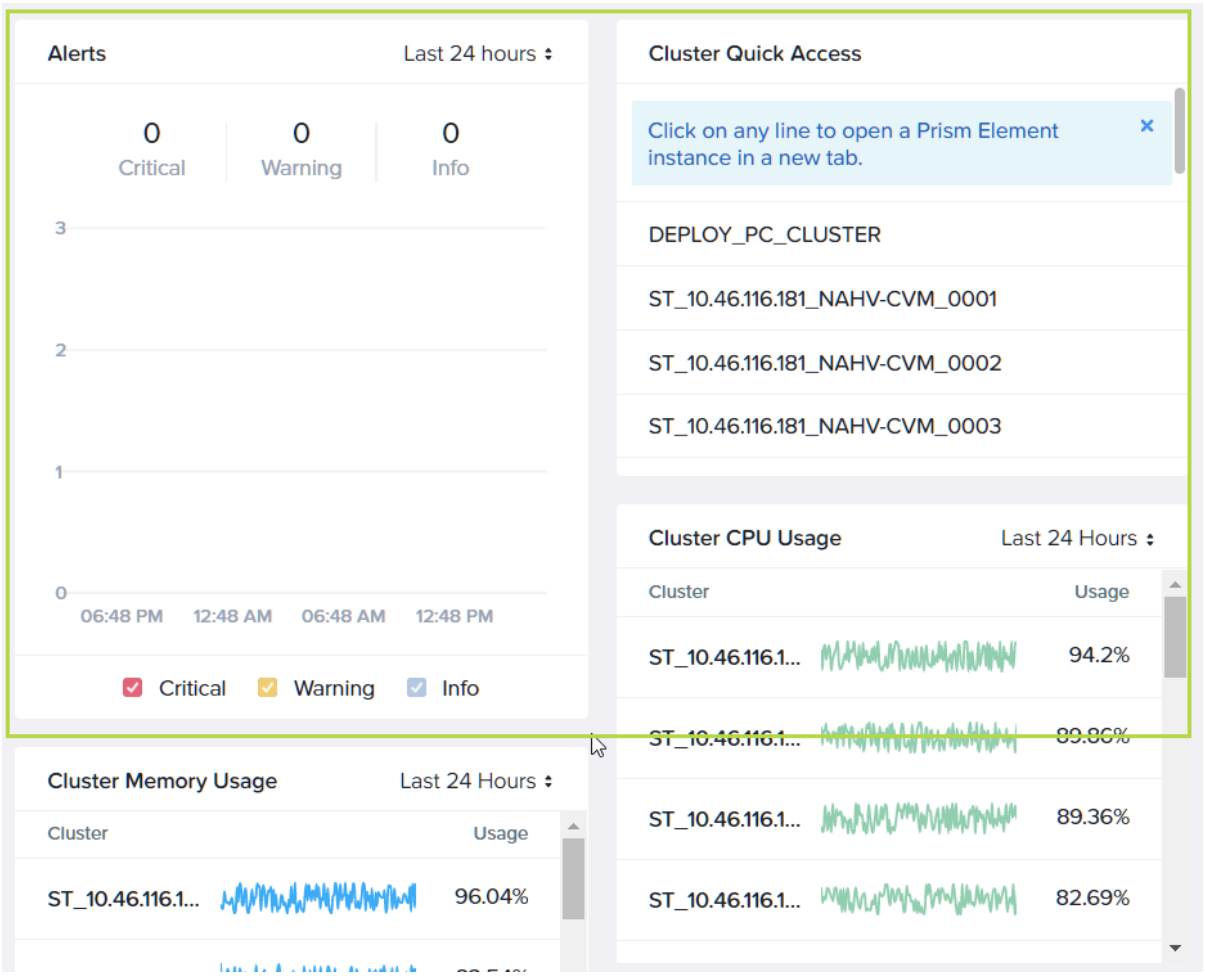

Figure 43: Widget Resize to One-and-half Times the Default Length

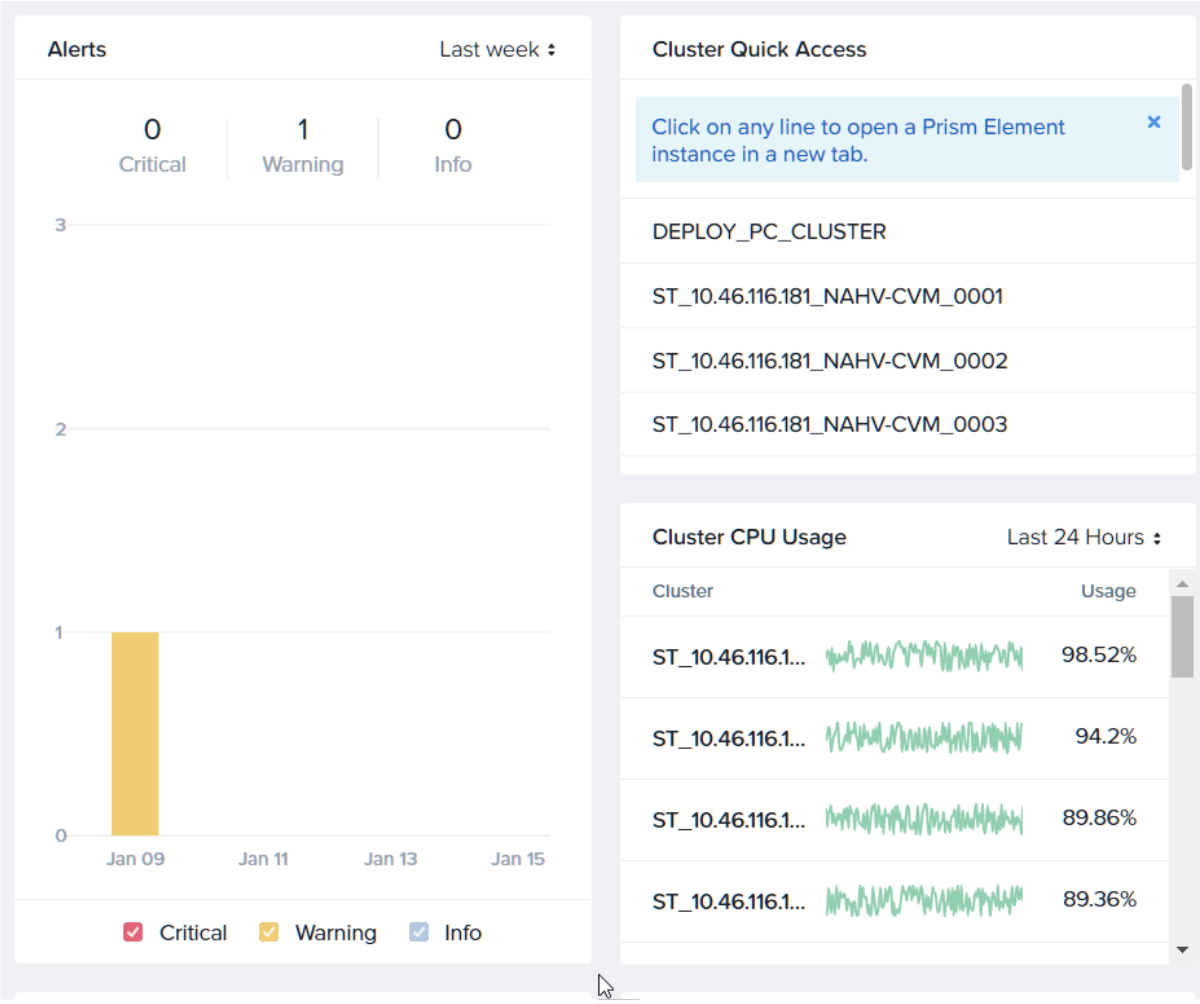

Figure 44: Widget Resize to Twice the Default Length

Drill-down or Expanding Data Elements in a Widget

You can expand or drill-down the data presented in the widgets by clicking on the data element. The following conditions are applicable:

- Expanding a data element in a widget is not possible in all widgets.
- In widgets that provide a list, clicking the list item or row opens an entity page in Prism Central or another application such as Prism Web Console in a new browser tab. For example, in the **Cluster Quick Access** widget, clicking the cluster name in the list opens the Prism Element Web Console for that cluster in a new browser tab. In the **Cluster Storage** widget, clicking a cluster name opens the cluster details view for that cluster.
- In widgets that provide numbers, clicking a number opens the respective entity page in Prism Central. For example, in the **Reports** widget, clicking the **Total Reports** or the **Scheduled Reports** number opens the **Reports** page. Similarly, clicking

- If the widget presents charts, click on the chart to display an expanded or drill-down view of the chart. Do not click on the entity (such as cluster name). If you do so, instead of displaying the drill-down view of the chart, the entity details page is displayed.

  The figure displays an example of expanding a chart in a widget.

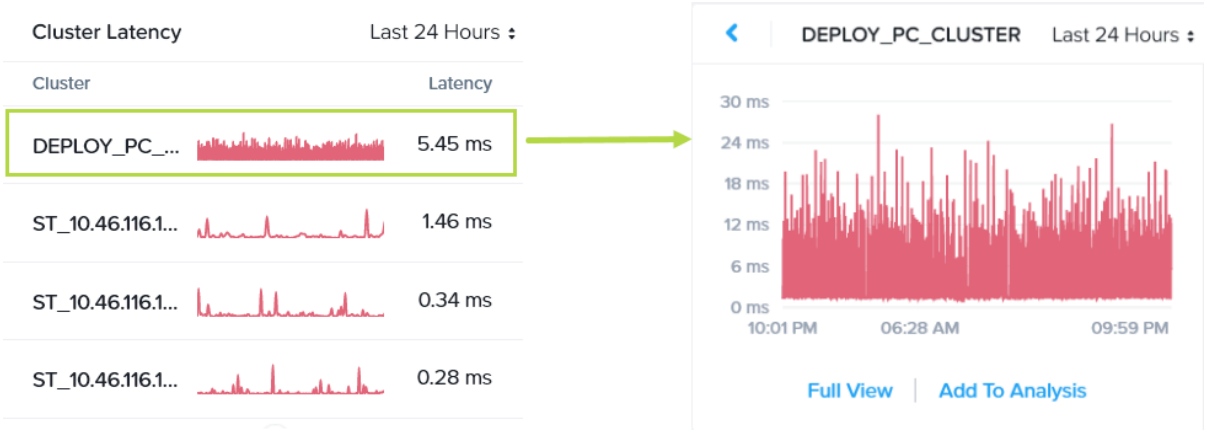

Figure 45: Widget Data Drill-down

At the bottom of the drill-down view, you can click **Full View** to open the details page of the entity(cluster or VM). Click **Add To Analysis** to add the chart to the **Analysis** dashboard (See Analysis Dashboard (Prism Central) on page 799.

Hover Details in Drill-down Views

After you click on a chart in a widget, the drill-down view of the chart is displayed. You can hover on any point on the time chart to display the parameters at that point.

In the same example as presented in the Drill-down or Expanding Data Elements in a Widget on page 77, if you hover over the chart, the following sample details are displayed:

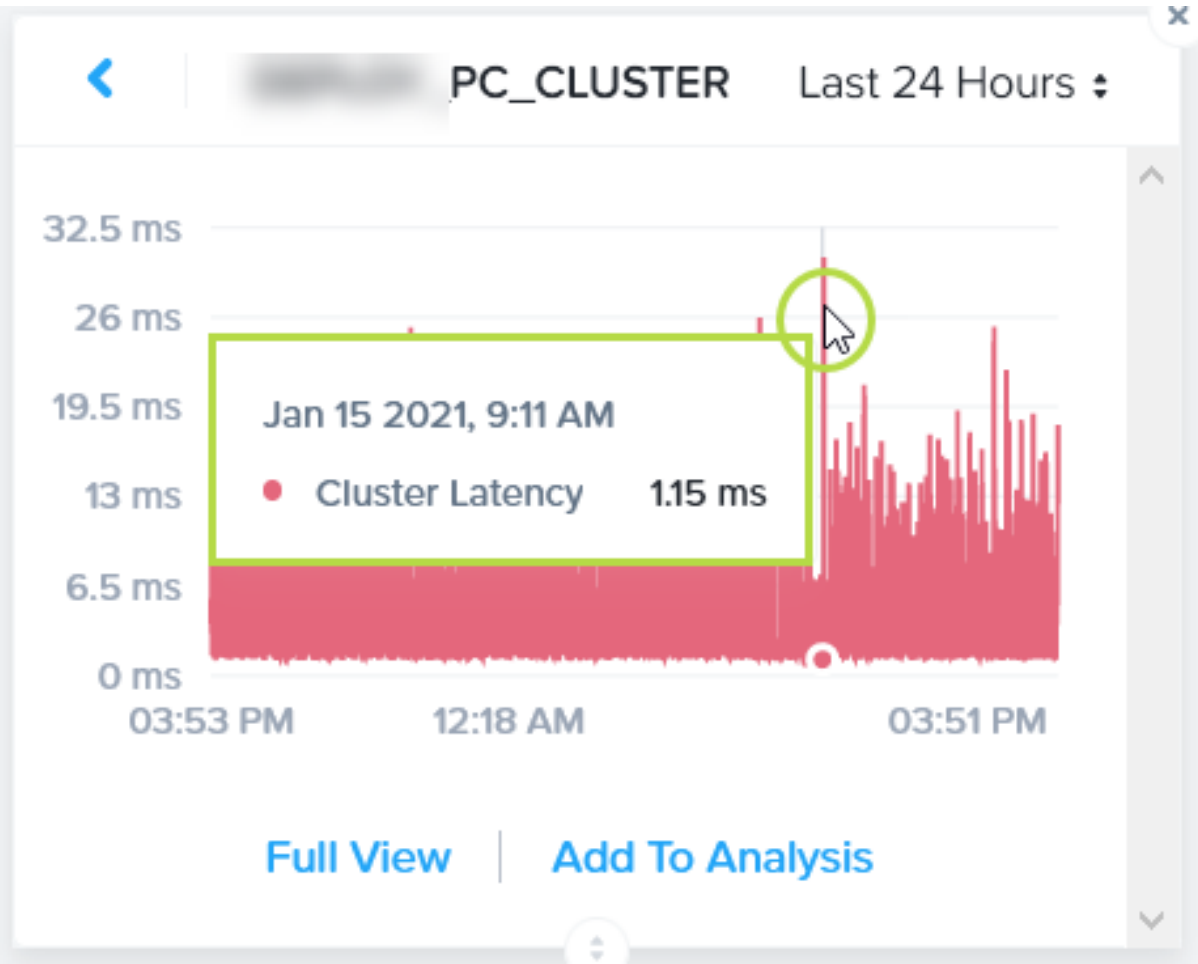

Figure 46: Hover Details

Creating a New Dashboard

About this task

The Main dashboard provides a default view into the clusters, but you can add custom views by creating custom dashboards. To create a custom dashboard, do the following:

> **Note:** Creating a dashboard requires a Prism Pro license. The **Manage Dashboards** button does not appear if Prism Pro is disabled.

Procedure

1. In the Main dashboard (see Main Dashboard on page 70), click the **Manage Dashboards** button.

2. In the **Manage Dashboards** window (see Modifying a Dashboard on page 85), do the following:

   a. Click the **New Dashboard** button.

   b. Enter a name for the dashboard in the displayed field.

   c. Click the **Save** button.

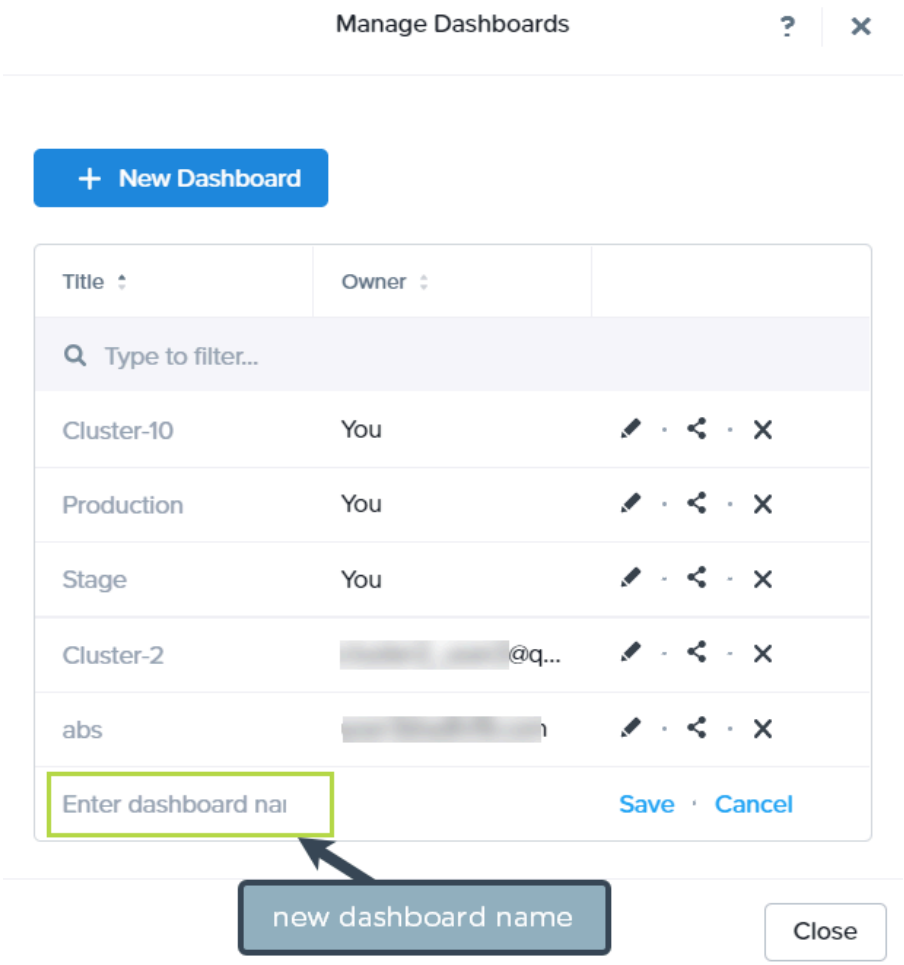

Figure 47: Manage Dashboards Windows (new dashboard)

3. Click the **Close** button to close the **Manage Dashboards** window.
   A tab for the new dashboard appears next to the **Main Dashboard** (or other custom dashboard) tab.

4. Click the tab for the new dashboard (to display that dashboard) and then click the **Add Widgets** button.
   The new dashboard is empty (no widgets) initially.

5. In the **Add Widgets** screen, select a widget and add it to the dashboard (see Adding Dashboard Widgets on page 80).
   Repeat this step until all the desired widgets have been added to the dashboard.

Adding Dashboard Widgets

Prism Central provides a selection of widgets that can be added to the home or a custom dashboard. Each widget provides information about a specific resource or usage topic that is displayed in a tile on the dashboard.

To add a widget to a dashboard, do the following:

> **Note:** Adding a widget to a dashboard requires a Prism Pro license. The **Add Widgets** button does not appear if Prism Pro is disabled.

1. Go to the target (main or custom) dashboard and click the **Add Widgets** button on the far right.
2. In the **Add Widgets** page, select the desired widget in the left column. (To find a particular widget, enter the name in the search field.) A preview of the selected widget appears in the middle column.
3. Enter values for the configurable parameters in the right column. The parameters vary by widget; some widgets do not have any configurable parameters.
4. Click the **Add to Dashboard** (or **Add & Return to Dashboard**) button to add the widget to the dashboard.

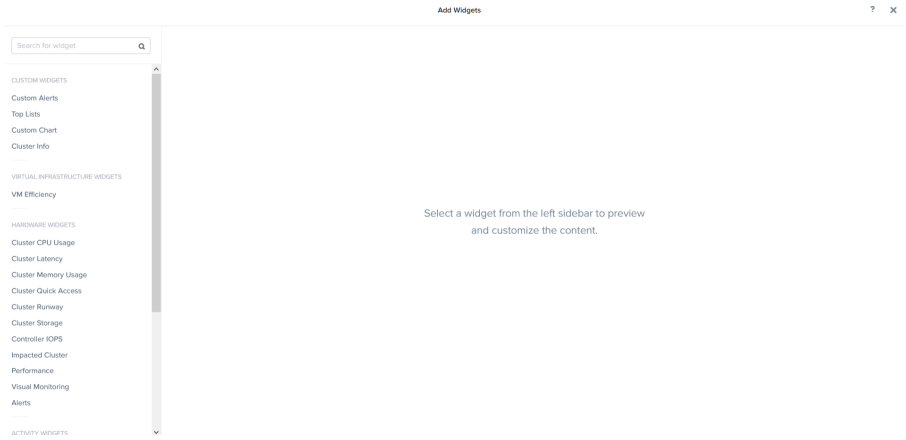

Figure 48: Add Widget Screen

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

Table 12: Widget Parameters

Widget	Description	Parameter	Values
Custom Widgets			
Custom Alerts Widget	Displays a custom list of alerts.	Widget Name - Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster - Select the cluster(s) to monitor from the pull-down list. The default is all clusters.	All Clusters, <cluster name>
		Choose Size - Select the size (width/height dimensions) of the widget. The default is 1 x 1.	1 x 1, 2 x 1, 2 x 2, 2 x 3, 2 x 4, 4 x 4
		Severity. The default is critical.	Critical, Warning, Info

Widget	Description	Parameter	Values
		Entity. Select the entity of interest. The default is cluster.	Cluster, Storage, VM, Hardware, DR
Top Lists Widget	Displays a list of the top performers for a selected metric.	Widget Name - Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster - Select the cluster(s) to monitor from the pull-down list. The default is all clusters.	All Clusters, <cluster name>
		Choose Size - Select the size (width/height dimensions) of the widget. The default is 2 x 1.	2 x 1, 2 x 2, 2 x 3
		Entity. Select the entity of interest. The default is host.	host, VM
		Metric. Select the metric of interest. The default is IOPS.	IOPS, Memory Usage, CPU Usage, Bandwidth, Latency
Custom Chart Widget	Displays a performance graph for a selected entity and metric pair.	Widget Name - Enter a name for the widget. The default is "New Chart Widget".	(name)
		Entity Type. The default is host.	Host, Disk, Storage Pool, Storage Container, Virtual Machine, Virtual Disk
		Entity. (There is no default value.)	(entity name)
		Metric. (There is no default value.)	(metric name)
Cluster Info Widget	Displays cluster summary information about alerts, anomalies, runway, and inefficient VMs.	Widget Name - Enter a name for the widget. The default is "New Widget".	(n/a)
		Select a Cluster. Select the target cluster from the pull-down list.	(registered cluster names)
Compute & Storage Widgets - This widget appears on the Main dashboard by default (see Main Dashboard on page 70).			

Widget	Description	Parameter	Values
VM Efficiency	Displays a four-cell table that lists the number of overprovisioned, inactive, constrained, and bully VMs in the cluster with links to the details.	(no customizable parameters)	(n/a)
Hardware Widgets - All these widgets (except Performance) appear on the Main dashboard by default (see Main Dashboard on page 70).			
Cluster CPU Usage	Displays the percentage of total CPU in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Latency	Displays the total (read and write) IO latency average for the highest latency clusters.	(no customizable parameters)	(n/a)
Cluster Memory Usage	Displays the percentage of total memory in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Quick Access	Displays a list of registered clusters. Health and alert icons appear for each cluster. Clicking the cluster line opens Prism (element) for that cluster in a separate tab or window.	(no customizable parameters)	(n/a)
Cluster Runway	Displays storage, CPU, and memory runway estimates (time remaining before the resource reaches capacity). See the *Capacity Runway Tab* section in *Resource Planning* for more runway information.	(no customizable parameters)	(n/a)
Cluster Storage	Displays storage statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Controller IOPS	Displays the total (read and write) controller IOPS for the highest volume clusters.	(no customizable parameters)	(n/a)
Impacted Cluster	Displays information about any clusters that are impacted (performance, capacity, or other potential issues) and may need attention.	(no customizable parameters)	(n/a)
Performance	Displays latency, bandwidth, and IOPS statistics for the highest usage clusters.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Alerts	Displays colored bar graph representing the alerts raised in the period selected. Clicking the **View All Alerts** link below the list to see all the alerts on the **Alerts** page.	(no customizable parameters)	(n/a)
Activity Widgets - The Tasks widget appears on the Main dashboard by default (see Main Dashboard on page 70).			
Tasks	Displays a list of recent tasks with the current status of each task and a link to the **Tasks** dashboard.	(no customizable parameters)	(n/a)
Discovered Apps	Displays the number of applications discovered in a specified set of monitored clusters. Click the **App instances** number to display the application discovery dashboard. Click the **Identified** or **Unidentified** number to display the application discovery dashboard filtered for that condition. This widget displays the discovered apps information only if you enable App Discovery.	(no customizable parameters)	(n/a)
Operations Widgets - These widgets appear on the Main dashboard by default (see Main Dashboard on page 70).			
Plays	Displays a list of the Plays running on the cluster. The list is categorised to completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the **Plays** dashboard page.	(no customizable parameters)	(n/a)
Reports	Displays a table that lists the number of total and scheduled reports with a link to the **Reports** dashboard.	(no customizable parameters)	(n/a)
Disaster Recovery Widgets - These widgets appear on the Main dashboard by default (see Main Dashboard on page 70).			

Widget	Description	Parameter	Values
Protection Status	Displays the protection plan status if one has been created. If not, it displays a set of recommended steps to create a protection plan.	(no customizable parameters)	(n/a)
Recovery Plan Status	Displays the recovery plan status if one has been created. If not, it displays a set of recommended steps to create a recovery plan.	(no customizable parameters)	(n/a)

Modifying a Dashboard

About this task

The Main dashboard provides a default view into the registered clusters, but you can customize that view at any time. To modify the Main dashboard or any other dashboard you create, do the following:

> **Note:** Customizing a dashboard requires a Prism Pro license. The **Add Widgets** and **Manage Dashboards** buttons do not appear if Prism Pro is disabled.

Procedure

1. Go to the Main dashboard (see Main Dashboard on page 70) and click the tab of the dashboard to modify (main or previously added custom dashboard).
2. To add a widget to the displayed (main or custom) dashboard, click the **Add Widgets** button, select a widget from the **Add Widgets** screen, and add it to the dashboard (see Adding Dashboard Widgets on page 80).
3. To delete a widget, click the X icon in the upper right of the displayed tile.

   A prompt appears to verify the delete; click the **OK** button. The widget then disappears from the dashboard.
4. To reset the Main dashboard to the default set of widgets (after you have previously added or deleted widgets), click the **Reset Dashboard** button.

   A prompt appears to verify the reset; click the **OK** button. The main dashboard returns to its default view.

5. To rename a custom dashboard, do the following:
   a. Click the **Manage Dashboards** button.
   b. Click the pencil icon for that dashboard.
   c. Enter a new name in the displayed field.
   d. Click the **Save** button.

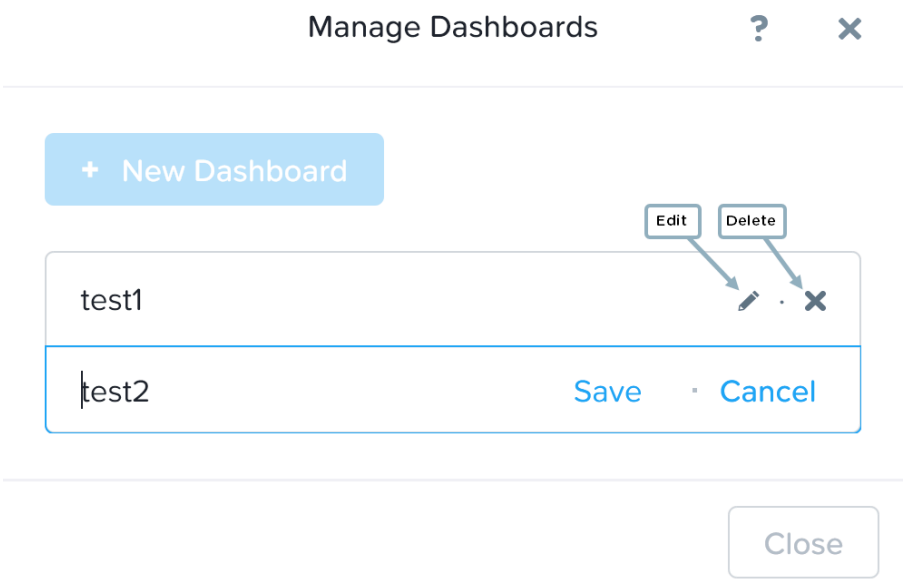

Figure 49: Manage Dashboards Window

6. To delete a custom dashboard, click the **Manage Dashboards** button and then click the X icon for that dashboard.
   A prompt appears to verify the delete; click the **OK** button. The tab for that dashboard disappears from the screen.

Sharing a Dashboard

About this task

Any user with edit permission can share a custom dashboard with other users and provide view or edit access to them.

> **Note:** Enable CMSP to use this feature.

Procedure

1. In the main menu, click **Manage Dashboard**.
   The **Manage Dashboard** window appears. This window lists all existing custom dashboards.

2. Select the custom dashboard that you want to share and click the share icon.

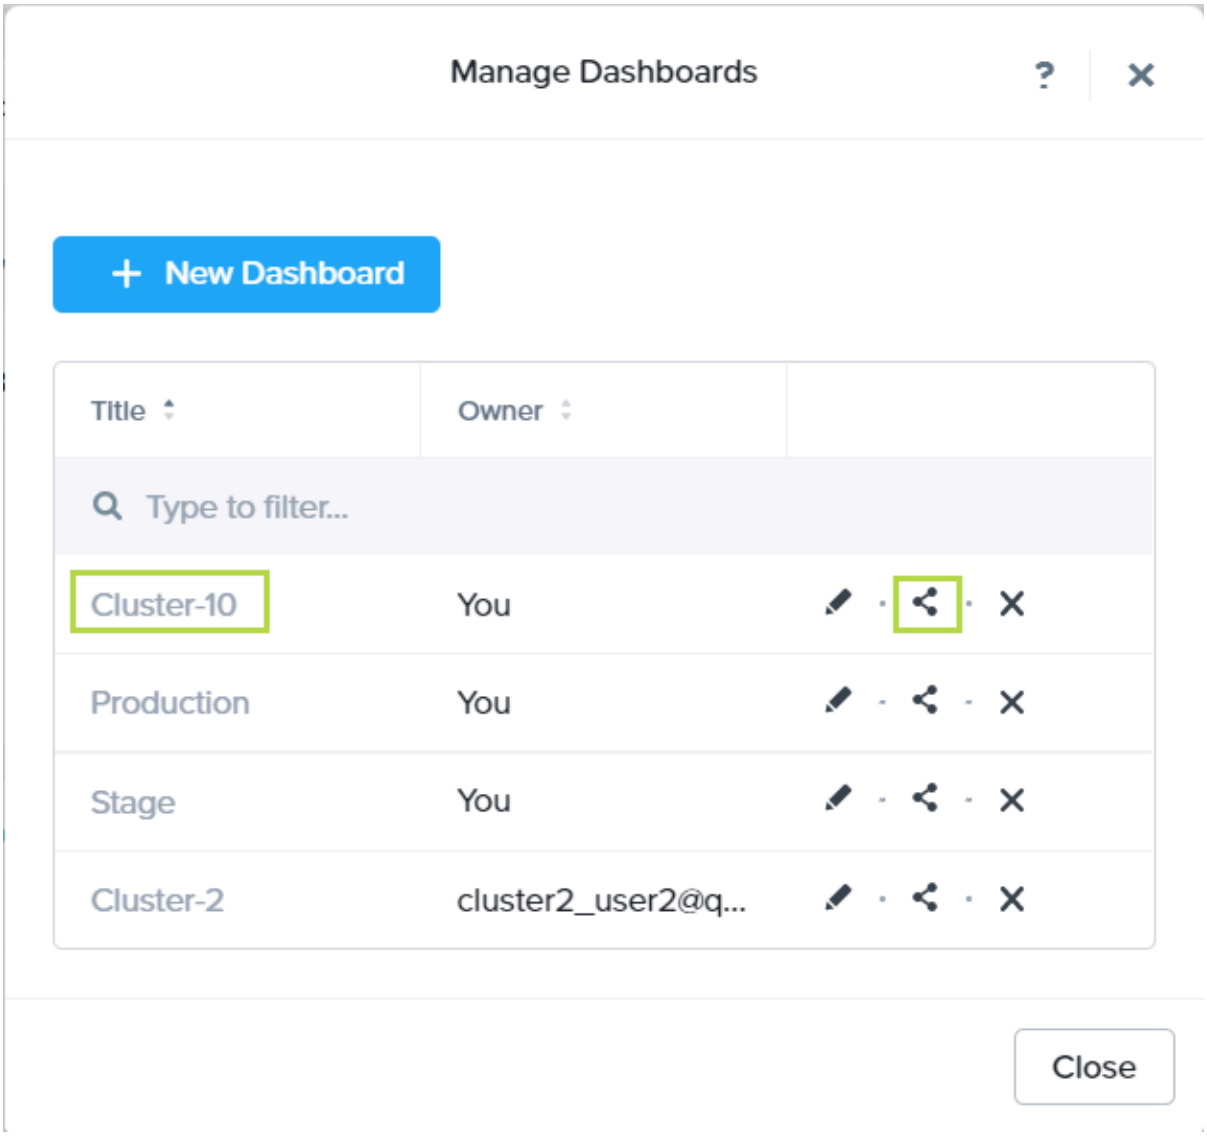

Figure 50: Manage Dashboards

The **Sharing Dashboard** window appears.

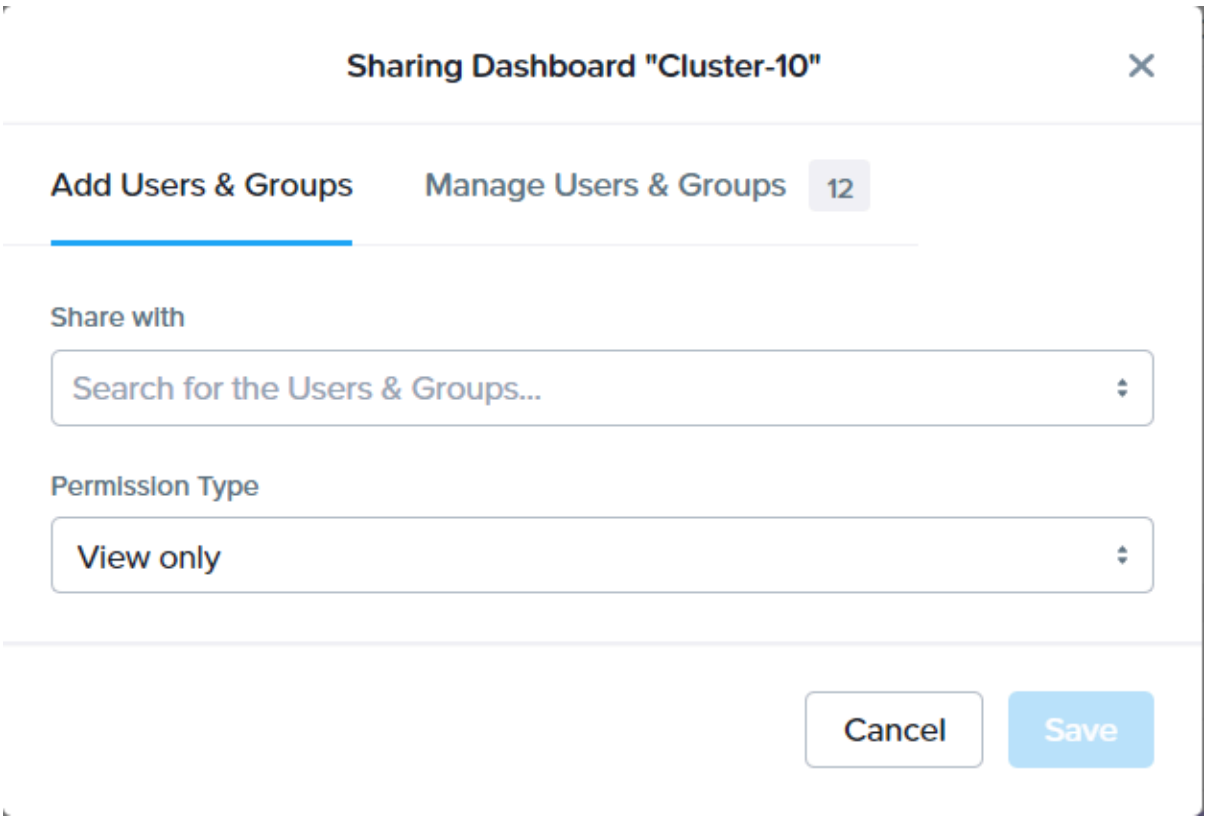

Figure 51: Sharing Dashboard

3. In the **Add Users & Groups** tab, do the following:

   a. **Share with**. Add the users or groups to share the custom dashboard. You can type a first few characters of a user or group, as you type a pull-down list appears. Click the username or group name to add. You can add more than one user or a group.

   b. **Permission Type**. You can provide View only or Edit permissions to the user or group. User with Edit permission can do these tasks: View, Edit or Share the custom dashboard.

   The **Manage Users & Groups** tab allows you to change the permission type for users or groups with whom you have shared the dashboard earlier. You can also remove the users or groups.

4. Click **Save**.

   A message that the dashboard is shared successfully appears.

   After sharing the dashboard, wait for few minutes for changes to take effect.

   Users will not be able to view any information about entities in the dashboard chart without entity permissions.

ENTITY EXPLORING

The entities menu (see Entities Menu on page 13) provides access to dashboards that allow you to view statistics about a range of entities across clusters and to organize that information in a variety of ways. An *entity* is an object type such as a VM, cluster, security policy, project, or report. Many of the entity dashboards include common elements in a summary view with drill-down options to view detailed information about individual entities of that type.

Note: This section describes common dashboard elements and how you can customize the view. See the appropriate **<entity> Summary View** section for a description of the specific fields and options applicable to each entity view. The Entities menu is organized into group categories, and the **<entity> Summary View** sections are located under one of the following group category headings:

- Compute and Storage Entities on page 93 (VMs, OVAs, images, catalog items, storage containers, storage policies, volume groups, vCenter datastores)
- Network and Security View on page 170 (subnets, virtual private clouds, floating IPs, network connectivity with VPNS and VTEPS, and security policies)
- Policies on page 189 (NGT policies, image placement policies)
- Data Protection and Recovery Entities on page 198 (protection summary, protection policies, recovery plans, VM recovery points)
- Hardware Entities on page 201 (clusters, hosts, disks, GPUs)
- Activity Entities on page 240 (alerts, events, tasks, audits)
- Operations Entities on page 246 (analysis, cost management, app discovery, monitoring integrations, operations policies, planning, playbooks, reports)
- Administration Entities on page 248 (categories, LCM, projects, roles, users, availability zones)
- Services Entities on page 267 (Calm, Files, Foundation Central, Karbon, Objects)

These views reflect that Prism Central retains alerts and events for 90 days, raw metric values for 14 days, and metric data that is downsampled every hour for a period of 365 days.

Screen Layout

Common display elements in many dashboards include the following:

- A selection pane on the left that lists the entity type and number at the top.
  - For hardware and virtual infrastructure entity types like clusters and VMs, the pane includes five tabs: Summary, List, Alerts, Events, Metrics. Click the desired tab to display that content in the main section of the screen.
  - For most other entity types, the pane is blank (no tabs), and the relevant content appears automatically in the main section of the screen.
  - For some entity types, the pane does not appear.

- When a list appears in the main section, the display may include:
  - A menu bar with (left to right) selection options, action options (if any), grouping options, and display options for the list, and a **Filters** button on the far right that when clicked displays the **Filters** pane (hidden by default).
  - A query field that identifies the filters (if any) being applied to the list. This field displays all filter options that are currently in use. It also allows for basic filtering on the entity name.
  - A table (list) of entities. What is included in the table is determined by the grouping, display, and filter options. Entities are displayed 20 per page by default (use the scroll bar to scroll a page), but you can specify a different number per page (10 to 60) from the pull-down list above the table (click "`X - XX of XXX`" to display list).
  - A download icon (just above the table on the right); click the icon to download the table contents in CSV format. You can download a maximum of 1000 rows.

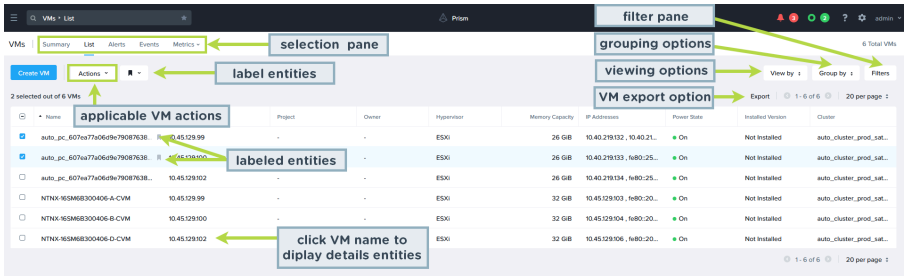

Figure 52: Dashboard Layout (example)

Filter Options

When the main section displays a list of entities, you can filter the list by clicking the far right icon in the menu bar to display the **Filters** pane. This pane includes a set of fields that vary according to the type of entity. Select the desired field values to filter the list on those values. An entry appears in the search field (see Searching for Information on page 19) for each value you select. You can save a filter by clicking the start icon in the search field, or you can remove a filter by clicking the X for that value.

Numeric filters have **To/From** fields to specify a range. These fields can take numeric values along with units. For example, the filter adjusts the scale accordingly when you type in "10 K" or "100 M".

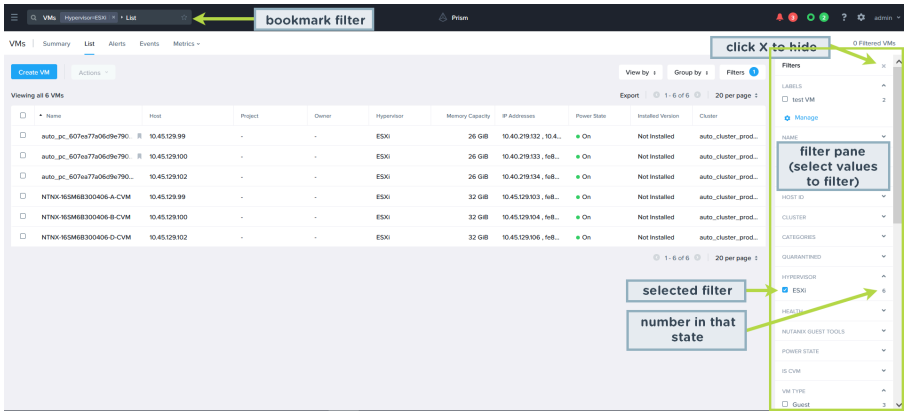

Figure 53: Filter Pane

Grouping Options

There are three (or four) drop-down menus on the right of the menu bar that allow you to organize the entity information in several ways.

- The **View by** menu allows you to specify the type of displayed information.

  > **Note:** A general focus is available for all entities, but the other options are available only when appropriate for that entity.

  - Select **General** to display a set of general information parameters, as illustrated in the Explore Dashboard figure.
  - Select **Performance** to display a set of performance-specific parameters.
  - Select **Efficient** to display a set of efficiency-related parameters. "Efficiency" is determined through the VM behavioral learning engine (see Behavioral Learning Tools on page 899).
  - Select **GPU** to display a set of GPU-specific parameters.
  - Select **Data Protection** to display a set of data protection-related parameters.

  In addition to the preset **View by** displays, you can create one or more custom displays as follows:

  1. Select the **Add Custom** button at the bottom of the **Focus** menu.
  2. In the **<entity> Column** window, do the following:
     1. Enter a name for the custom display in the first (top) field.
     2. Select (click the blue + circle for) each entity property in the left column you want in your display. The selected entity properties appear in the right column. You can filter the entity-property list by entering a

string in the search field above the left column. You can select a maximum of 10 columns in one custom display including the Name column.

3. When the entity-property list in the right column is correct (you can remove an entity property from this list by clicking the red - circle for that entity), click the **Save** button.

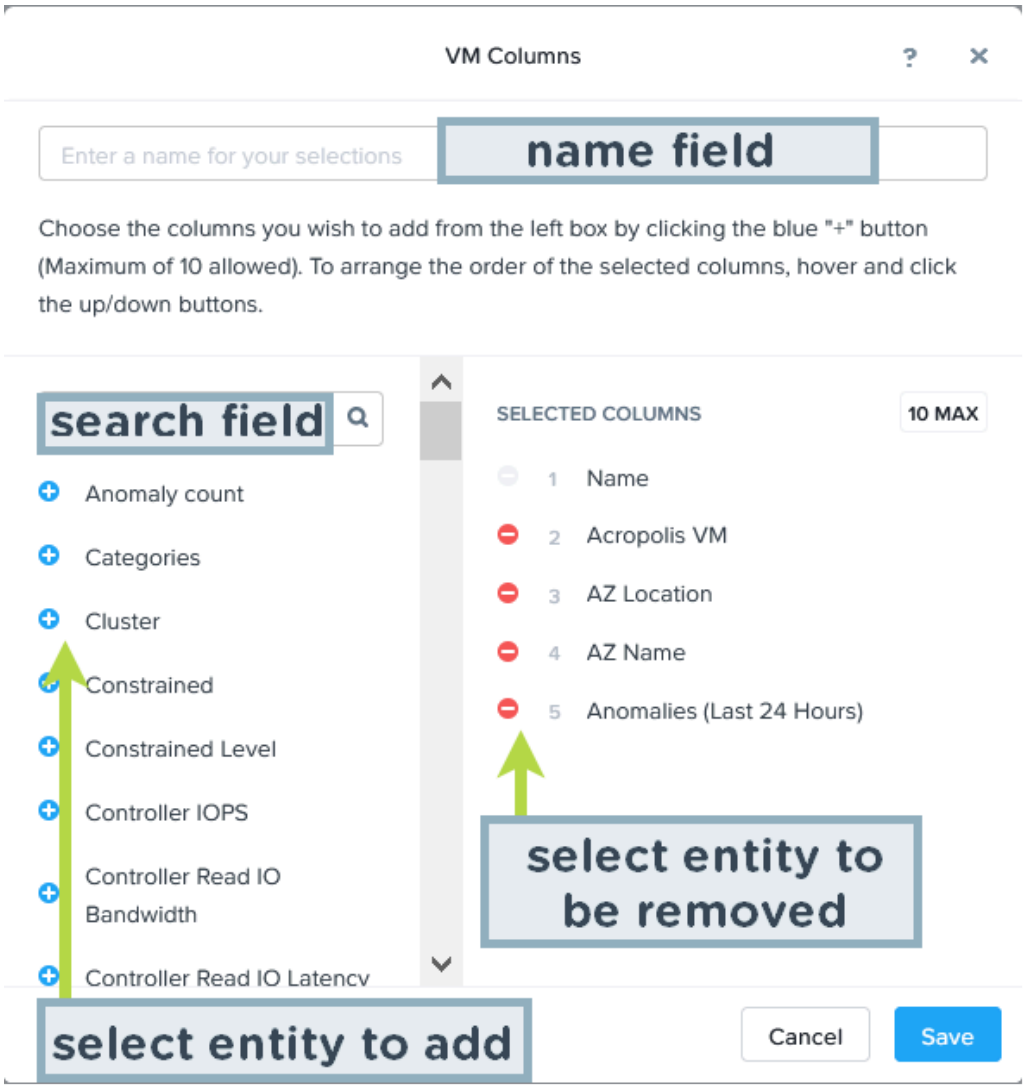

Figure 54: VM Column Window (example)

The **<entity> Column** window disappears and the new custom display appears (by name) in the **Focus** menu. To view that display at any time, select it from the menu.

> **Note:** Only the user that created a custom display can view that custom display. Other users and the administrative user cannot view that custom display.

- The **Color** menu allows you to color code the entries based on a selected parameter, as illustrated in the Circles View figure.
- The **Group** menu allows you to group the entries based on a selected parameter, as illustrated in the Circles View figure.
- The **Sort** menu allows you to sort the entries based on a selected parameter. This menu appears for the Tiles and Circles views only; ordering in the Grid (tabular) view is accomplished by clicking on a column header.

Action (and Selection) Options

The dashboard provides options to administer and categorize the entities. Actions are applied to selected entities that are currently visible. You can select a set of entities (rows) by selecting the first one and then using <Shift> select to select the last one.

- While some administrative tasks for a cluster must be done through Prism Element (see the Prism Element Web Console Guide), you can perform other administrative tasks directly from Prism Central. Action buttons for such tasks appear in the menu bar when an entity type is selected. For example, buttons for creating a VM and configuring the network appear in the VMs view. (No buttons appear when there are no relevant actions available for the selected entity type.)
- When one or more entities are selected from the list, an **Actions** menu appears in the menu bar. The menu includes actions that you can do to the selected entities. For example, the VMs menu includes actions such as power on or clone which will be executed on all the selected VMs. (Only currently valid actions are available; other menu options are inactive and appear gray.)
- When one or more entities are selected from the list, a label icon appears (next to the **Actions** menu). Labels allow you to create custom groupings for the entities. To add a label, click the icon and enter the label name in the field. The label is applied to all the selected entities in the list.

  Note:

  - Only super admin and Prism admin roles have permissions to perform this operation.
  - Labels are available for VMs and clusters only.

- Use the selection menu to
  - Select all entries in the list by checking the box or choosing **Select all** from the pull-down list.
  - Unselect all by unchecking the box or choosing **Clear selections** from the pull-down list.
  - Show just the selected entries by choosing **Show selected entities** from the pull-down list.

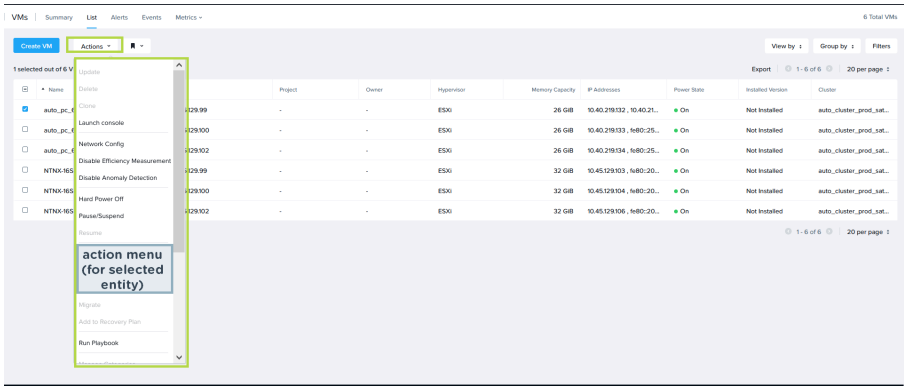

Figure 55: Action Options

Compute and Storage Entities

You can access dashboards for the following virtual infrastructure entities from the **Compute and Storage** category of the entities menu (see Entities Menu on page 13):

- VMs (see VMs Summary View on page 94)

- OVAs (see OVAs View on page 122)
- Images (see Images Summary View on page 125)
- Catalog Items (see Catalog Items View on page 129)
- Storage Containers (see Storage Containers Summary View on page 130)
- Volume Groups (see Volume Groups Summary View on page 147)
- vCenter Datastores (see vCenter Datastores Summary View on page 156)

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

VMs Summary View

To access the VMs dashboard, select **Compute & Storage** > **VMs** from the entities menu (see Entities Menu on page 13). The VMs dashboard summary view displays information about VMs across the registered clusters and allows you to access detailed information about each VM. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> **Note:** This section describes the information and options that appear in the VMs dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.
> - See VM Management on page 535 for information about creating and managing VMs.

Summary Tab

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

- **Highlighted Entities**: Displays a list of the VMs with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **Memory Usage**, and **IOPS**. Click the **View All XX VMs** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of VM-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click the graph, which then displays a list of those alerts. Clicking an alert displays the details page for that alert.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.
- **VM Efficiency**: Displays the number of VMs that are considered inefficient broken down by category (overprovisioned, inactive, constrained, and bully). This widget is the same that appears in the main dashboard (see Main Dashboard on page 70). See Behavioral Learning Tools on page 899 for more information about VM efficiency.

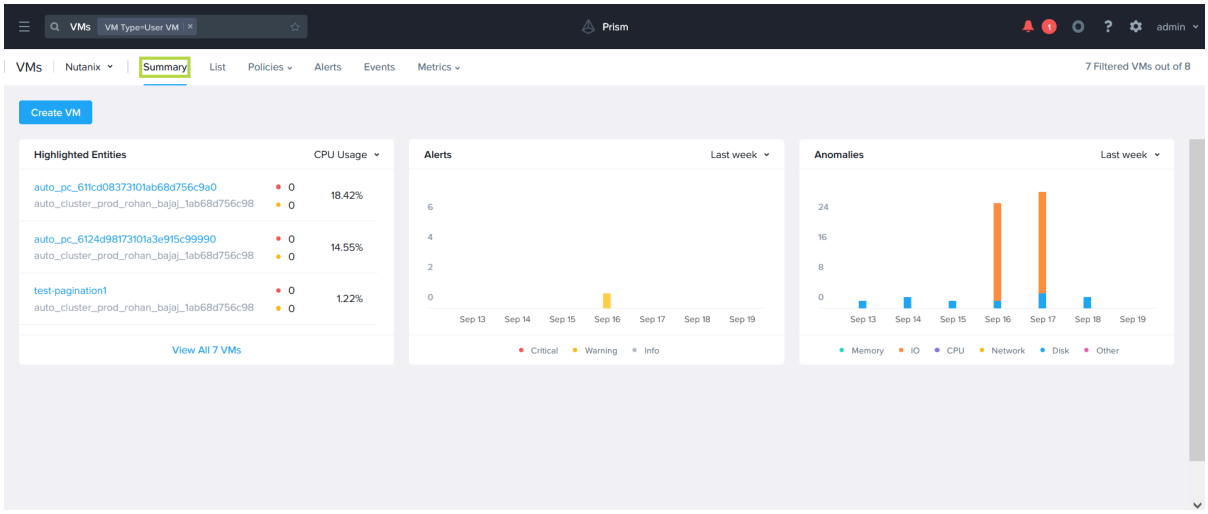

Figure 56: VMs Summary Tab (Nutanix)

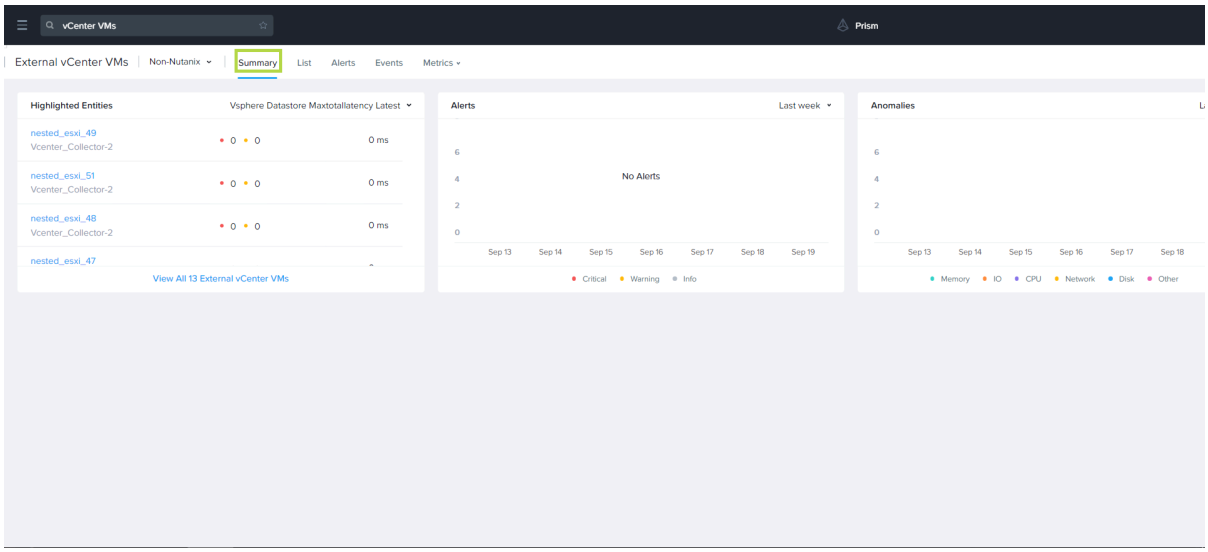

Figure 57: VMs Summary Tab (Non-Nutanix)

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the VMs across the registered clusters. The following table describes the fields that appear in the VMs list. The fields vary based on the **View by** menu selection, which is **General**, **Performance**, **Efficiency**, or **GPU**. A dash (-) is displayed in a field when a value is not available or applicable.

> **Note:** Create your own customized view and add the necessary columns to that view.

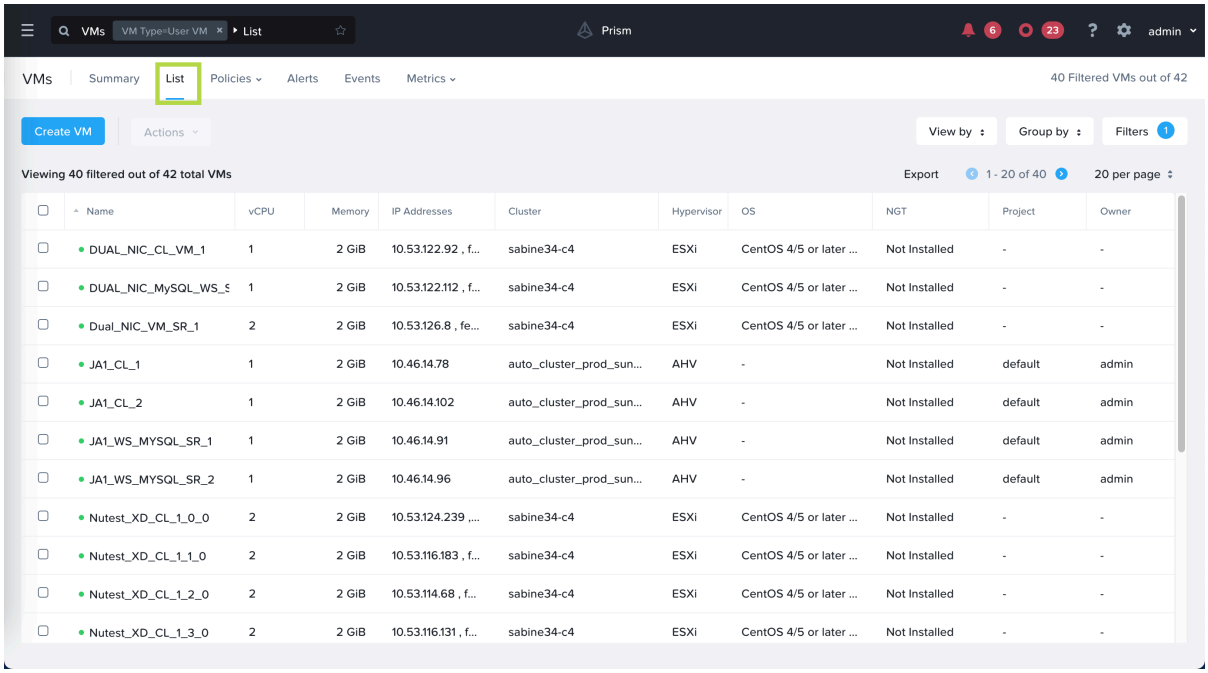

Figure 58: VMs List Tab

Table 13: VMs List Fields

Parameter	Description	Values
"General" Focus fields		
Name	Displays the VM name. Clicking the name displays the details page for that VM (see VM Details View on page 105).	(VM name)
vCPU	Displays the virtual CPU count of the VM.	(vCPU count)
Memory	Displays the total amount of memory available to this VM.	xxx [MB|GiB]
IP address	Displays one or more VM IP addresses.	(IP address)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Hypervisor	Displays the hypervisor type on which the VM is running.	AHV, ESX, or Hyper-V
OS	Displays the guest operating system of the VM.	(OS)
NGT	Displays whether Nutanix Guest Tools is installed on the VM.	Installed, Not Installed
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
"Performance" Focus Fields		

Parameter	Description	Values
Name	Displays the VM name.	(VM name)
Memory Overcommit	Displays the Memory Overcommit state of the VM.	(enabled, disabled)
Memory Usage	Displays the percentage of allocated memory capacity currently being used by this VM.	0 -100%
CPU Usage	Displays the percentage of allocated CPU currently being used by this VM	0 -100%
Read IOPS	Displays read I/O operations per second (IOPS) for this VM.	(number)
Write IOPS	Displays write I/O operations per second for this VM.	(number)
I/O Bandwidth	Displays I/O bandwidth used per second for this VM.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for this VM.	xxx [ms]
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"Efficiency" Focus Fields		
Name	Displays the VM name.	(VM name)
Efficiency	Displays the assessed efficiency of the VM. If a VM is performing in an expected range, the efficiency is listed as "Good". If not, the type of inefficiency is displayed as determined by the VM behavioral learning engine (see Behavioral Learning Tools on page 899).	Good, Constrained, Overprovisioned, Inactive, Bully; ("NA" if insufficient baseline data for categorization, which is typically 21 days)
Reason	Displays why a VM is considered inefficient. A dash (-) appears for a "Good" VM.	(text message)
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
"GPU" Focus Fields		
Name	Displays the VM name.	(VM name)
GPU Configuration	Displays the GPU board (in case of Passthrough mode), vGPU software, type and number of vGPU instances in brackets (in case you configured the vGPU mode).	(GPU board/software-type name) Example: Passthrough - Nvidia Tesla M60 Example: vGPU - Nvidia GRID M60-8Q (2)

Parameter	Description	Values
GPU Type	Displays the GPU configuration made as Passthrough or vGPU or if you did not configure GPU then "-"	[Passthrough|vGPU|-]
GPU Usage	Displays the percentage of GPU capacity being used by the VM.	(percentage)
GPU Framebuffer Usage	Displays the percentage of GPU framebuffer (RAM) capacity being used by the VM.	(percentage)
"Storage Configuration" Focus Fields		
Name	Displays the VM name.	(VM name)
Replication Factor	Displays the Replication Factor setting for the VM in the storage policy applied to the VM.	Inherit from Container or 2 or 3
Encryption	Displays the Encryption setting for the VM in the storage policy applied to the VM.	Enabled or Inherit from Cluster
Compression	Displays the Compression setting for the VM in the storage policy applied to the VM.	Inline or Post Process (if enabled) , Disabled or Inherit from Cluster
Throttled Throughput (IOPS)	Displays the throttled throughput value in terms of IOPS set for the VM in the storage policy applied to the VM.	(Integer number)
Throttled Throughput (MB/s)	Displays the throttled throughput value in terms of MB/s set for the VM in the storage policy applied to the VM.	(Integer number)
Associated Policy	Displays the name of the storage policy applied to the VM.	(String) Storage Policy Name
Compliance State	Displays the compliance state of the VM based on realization of the storage policy applied to the VM. **Note:** The compliance state is displayed as **In Progress** until the extent groups are synchronized by the Curator service at full scan intervals. Full scan intervals are intervals of 6 hours. Depending up on the time elapsed in the full scan interval at the time when the Replication Factor was set or updated, it may take up to 6 hours for the **In Progress** state to change to **Compliant** or **Non Compliant**.	**In Progress**, **Compliant** or **Non Compliant**
Cluster	Displays the name of the cluster that the VM resides in.	(String) Cluster Name
"Anomalous Behavior" Focus Fields		

Parameter	Description	Values
Name	Displays the VM name.	(VM name)
Anomaly Count	Displays the Anomaly count for Memory, I/O, CPU, Network, and Disks	Integer Number
Anomaly Disabled	Displays whether Anomaly reporting is Disabled.	Yes/No
"Data Protection" Focus Fields		
Name	Displays the VM name.	(VM Name)
Categories	Displays the category mapped to the VM. For information about categories in Prism Central, see Categories Summary View on page 164.	String (Category name)
Consistency Group	Displays the consistency group of the VM. For information about Consistency Groups in Prism Central, see Consistency Groups information in *Nutanix Disaster Recovery Guide*.	Yes/No
Protection Status	Displays the Protection Status of the VM.	Protected/Unprotected
Sync Status	Displays the sync status of the VM.	Yes/No
Protection Type	Displays the type of protection defined for the VM.	Protection Policy
Protection Policy	Displays the protection policy that is applicable for the VM. For more information about protection policies, see Nutanix Disaster Recovery Guide.	(String) Protection Policy Name
Recovery Plans	Displays the recovery plans that is applicable for the VM. For more information about VM recovery plans, see Nutanix Disaster Recovery Guide.	(String) Recovery Plans Name

You can filter the VMs list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range. Click the plus (+) sign to the right of the parameter name to expand the fields for that parameter; click the minus (-) sign to contract those fields.

Table 14: Filter Pane Fields

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the pull-down list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)

Parameter	Description	Values
Name	Filters on the VM name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the VM name condition/string. **Note:** In this and the following two fields, the condition menu options are **Contains**, **Doesn't contain**, **Starts with**, **Ends with**, and **Equal to**.	(VM name string)
Host	Filters on the host name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the host name condition/string.	(host name string)
Cluster	Filters on the cluster name. Select a condition from the pull-down list and enter a string in the field. It will return a list of VMs that satisfy the cluster name condition/string.	(cluster name string)
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a pull-down list appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of VMs tagged to each selected category is displayed on the right of the line.	(category name)
Hypervisor	Filters on the hypervisor type. Check one or more of the boxes to filter on those hypervisors. The number of VMs currently on each hypervisor type is displayed on the right of the line.	AHV, ESX, HyperV
Health	Filters on the VM health state (good, warning, or critical). Select one or more states to return a list of VMs in that state(s). The number of VMs currently in each state is displayed on the right of the line.	Critical, Warning, Good
Power State	Filters on the VM power state. Select one or more states to return a list of VMs in that state(s). The number of VMs currently in each state is displayed on the right of the line.	On, Off, Suspended, Paused, Unknown
VM Type	Filters on the type of VM. Select either user VM or Controller VM (or both). The number of VMs of each type is displayed on the right of the line.	User VM, AHV Controller VM
Memory Usage	Filters on the amount of memory capacity being used. Check the box for the desired percentage range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of VMs utilizing memory in that range.	([xx] to [yy]% range)
Read IOPS	Filters on the read IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of VMs with read IOPS in that range.	([xx] to [yy] range)

Parameter	Description	Values
Write IOPS	Filters on the write IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of VMs with write IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of VMs with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of VMs with average I/O latency in that range.	([xx] to [yy] range)
Over Provisioned	Filters for over-provisioned VMs. Check the box(es) for the desired type (high and moderate). **Note:** This and the following two fields filter based on a VM efficiency algorithm, which is part of the VM behavioral learning capabilities (see Behavioral Learning Tools on page 899).	High, Moderate
Constrained	Filters for constrained VMs. Check the box(es) for the desired type (high and moderate).	High, Moderate
Efficiency	Filters for certain VM profiles. Check the boxes for the desired profile types. There is one for efficient VMs (good) and four for inefficient VMs (bully, over-provisioned, constrained, inactive).	Bully, Over Provisioned, Constrained, Inactive VM, Good
GPU Configuration	Filters for GPU configuration information such as model name. Enter the GPU configuration information in the field and then check the box. As you type a pull-down list appears to help you select the correct configuration information.	(configuration info)
GPU Type	Filters for GPU operational mode. Check the box for one or more of the GPU types.	vGPU, Passthrough, Passthrough(Compute)
GPU Usage	Filters on the amount of GPU capacity being used. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of GPUs in that range.	([xx] to [yy]% range)
GPU Framebuffer Usage	Filters on the amount of GPU framebuffer (RAM) capacity being used. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of GPUs in that range.	([xx] to [yy]% range)
Storage Configuration	Filters on the storage policy configuration parameters set for VM. This view displays the storage configurations parameters set for the VMs by storage policies applied to the VMs. For more information about storage policies, see Storage Policies Summary View on page 141.	(Storage policy name)

Parameter	Description	Values
vGPU Guest driver Version	Filters on the guest driver version. Enter the guest driver version number in the field.	(guest driver version number)

Export VM information. You can export the table containing the list of VMs and their information to a file in a CSV format by clicking the export icon on the right.

You can group the VM list in the following ways:

- The **Group** pull-down menu allows you to group the VM entries by cluster, hypervisor, power state, virtual CPU count, or health state. (You can only choose one.)
- When you select one or more VMs, the **Label** icon appears (to the left of the **Actions** menu). To assign the selected VMs a group label, click the **Label** icon and do one of the following:

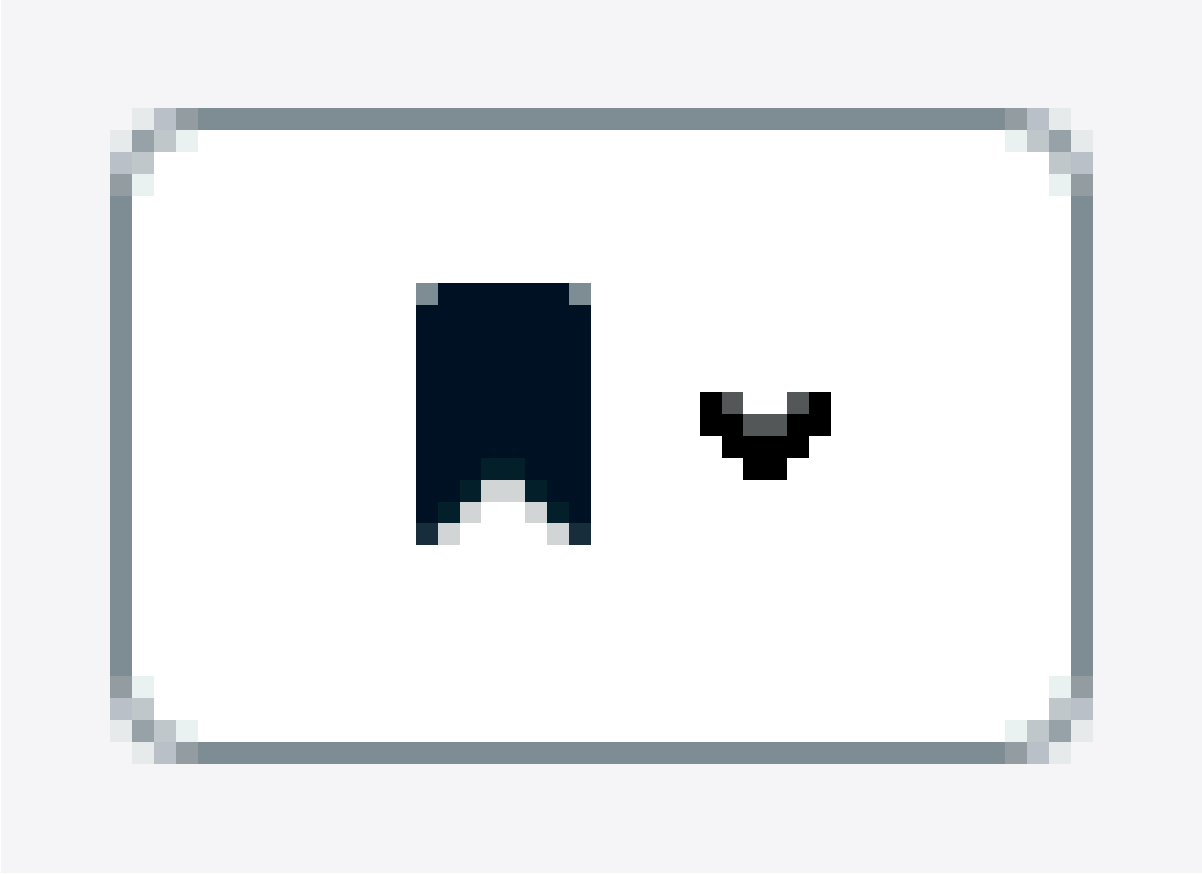

Figure 59: Label Icon

  - To assign an existing label (if there are existing labels), select the desired label and then click **Apply changes**.
  - To assign a new label, enter a meaningful label in the blank text field and then click **Create new label**.

  You can now perform any actions available on this labeled group from the **Actions** menu (see following section). You can edit or delete the label by clicking the **Label** icon and selecting **Manage Labels**.

The VMs view includes two action buttons:

- To create a VM, click the **Create VM** button. You can create a VM on a cluster running AHV or ESXi (see Creating a VM through Prism Central (AHV) on page 535 or Creating a VM through Prism Central (ESXi) on page 580).
- To configure the network, click the **Network Config** button (see Configuring Network Connections on page 784).

The **Actions** menu appears when one or more VMs are selected. It includes the following actions:

- These actions can be applied to multiple VMs: **Delete**, **Power on**, **Power off**, **Power Cycle**, **Pause/ Suspend**, **Resume**, **Disable Efficiency Measurement**, **Enable Efficiency Measurement**, **Disable Anomaly Detection**, **Enable Anomaly Detection**, **Protect**, **Unprotect**, **Add to Recovery Plan**, **Run Playbook**, **Manage Categories**, **Install NGT**, **Enable NGT**, **Disable NGT**, **Manage NGT Applications**, **Upgrade NGT**, **Set QoS Attributes**
- These actions can be applied to only one VM at a time: **Update**, **Clone**, **Create VM Template**, **Launch console**, **Reset**, **Guest Shutdown**, **Guest Reboot**, **Create Recovery Point**, **Migrate**, **Configure VM Host Affinity**, **Add to Catalog**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the current state of the selected VM(s). See Managing a VM through Prism Central (AHV) on page 552 or Managing a VM through Prism Central (ESXi) on page 585 for instructions on how to perform these actions.

Policies Tab

The **Policies** tab displays **Affinity Policies** (see Affinity Policies Defined in Prism Central on page 608) and **NGT Policies** (see NGT Policies on page 611).

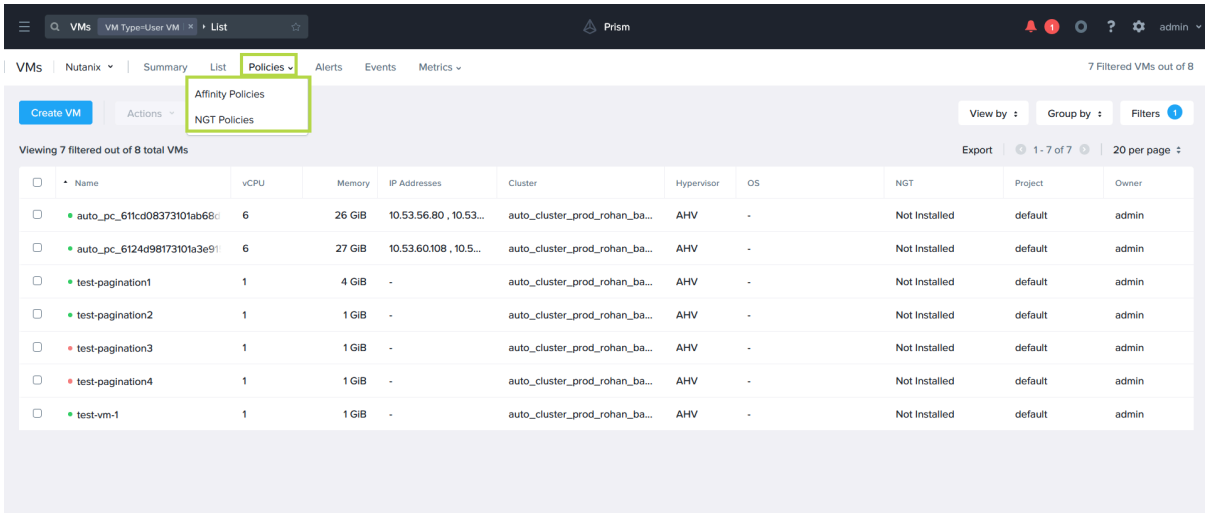

Figure 60: VMs Policies Tab

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just VM-related alerts across the registered clusters (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just VM-related events across the registered clusters (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the VMs. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant performance information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

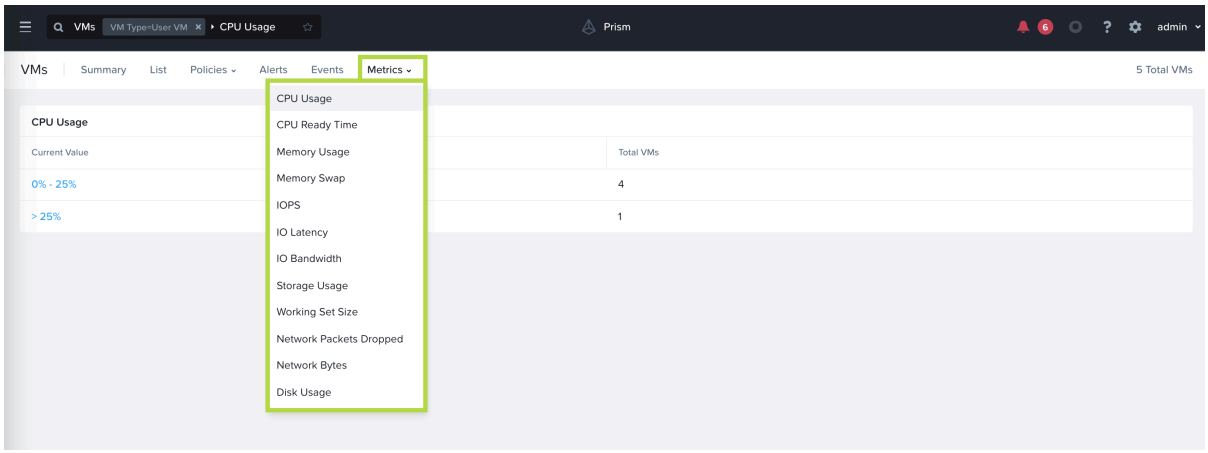

Figure 61: VMs Metrics Tab

Table 15: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total VMs (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking a percentile interval displays the Summary tab filtered to just those VMs. **Note:** The same format also applies to the other metrics in this table with either percentile or quantity intervals.
CPU Ready Time	Displays a CPU ready time percentage usage table.
Memory Usage	Displays a memory percentage usage table.
Memory Swap	Displays memory swap-out and swap-in rate tables.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate tables.
Usage	Displays total, snapshot, and shared storage size tables.

Metric	Description
Working Set Size	Displays total, snapshot, and shared working set size tables. Working set is the amount of memory that a VM requires in a given time interval.
Network packets dropped	Displays tables for the number of transmitted and received packets dropped.
Network bytes	Displays tables for the amount of transmitted and received bytes (in GiB).
Disk Usage	Displays a disk usage table listing current values and total VMs (number). The current values are split into percentile intervals (for example, l3.68% - 4.1%, more than 4.93%). Clicking a percentile interval displays the Summary tab filtered to just those VMs. **Note:** The same format also applies to the other metrics in this table with either percentile or quantity intervals.

VM Details View

To access the details page for a VM, go to the VMs **List** tab (see VMs Summary View on page 94) and click the VM name. You can also access the details page by clicking the VM name wherever that name appears, such as in a dashboard widget or search result.

The VM name and the following set of tabs appear on the left: Summary, Console, Data Protection, Alerts, Events, Metrics, NICs, Disks, Snapshots, Categories, and Apps & Relationships. Click a tab to display that information on the right. (Click the **Back to VMs** link to return to the VMs summary view.)

> **Note:** VirtIO must be installed in a VM for AHV to display correct VM memory statistics. For more information about VirtIO drivers, see Nutanix VirtIO for Windows in *AHV Administration Guide*.

Summary Tab

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

- A **Properties** widget that displays summary information about the VM (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

- A **Storage Properties - Associations** widget that displays the storage properties of the VM as defined in the storage policy associated with the VM and the other entities like Virtual Disks, Storage Containers, Categories and Storage Policies associated with the VM. See the Storage Properties Fields and the Associations Fields tables.

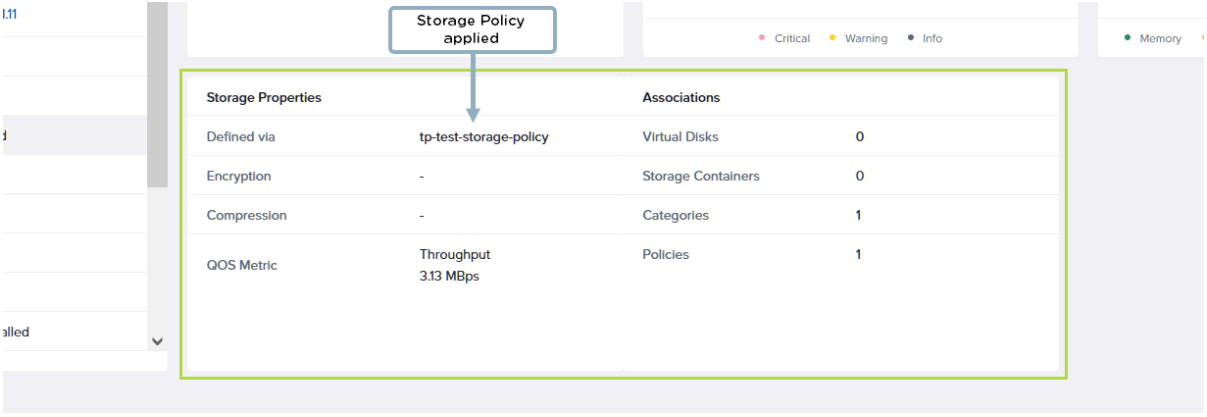

Figure 62: VM Storage Properties - Associations Widget

The widget displays the actual values for Encryption and Compression that is applied to the VM. For the values that are displayed in the **Storage Policy** dashboard as **Inherited from Cluster** also the **Storage Properties Association** widget displays the actual inherited values.

- Action button (above the widgets). Click the appropriate button to run that administrative action on the VM. The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM. See Managing a VM through Prism Central (AHV) on page 552 for instructions on how to perform each action.

> **Note:** You can perform administrative actions on VMs in Acropolis-managed clusters only. Unavailable actions are grayed out.

- 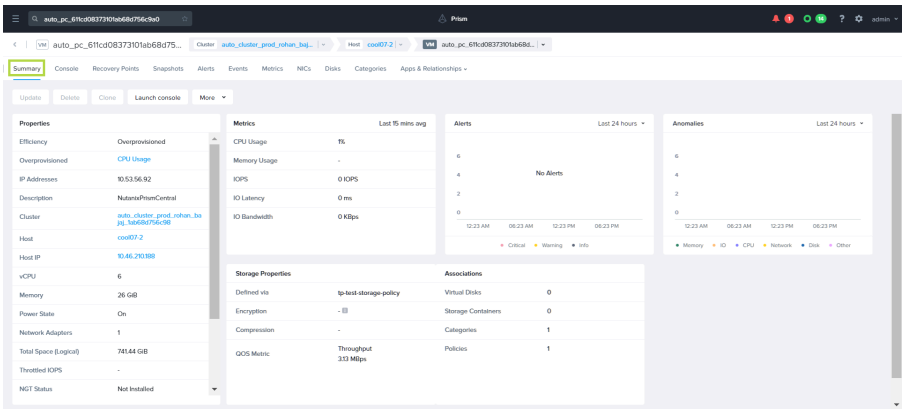

Figure 63: VM Summary Tab

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

Table 16: VM Properties Fields

Parameter	Description	Values
Efficiency	Displays the efficiency state for this VM. If the efficiency is not good, an additional field may appear that specifies the problem. For example, if the VM is constrained, a Constrained field appears that identifies the constrained resource such as the CPU or memory. **Note:** The Anomalies, Efficiency, Over provisioned, and Constrained parameters relate to the VM behavioral learning feature. See Behavioral Learning Tools on page 899 for a description of each parameter.	Bully, Over Provisioned, Constrained, Inactive VM, Good
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Host	Displays the host name. This field may be blank if the VM is powered off and a host is not assigned.	(host name)
Host IP	Displays the host IP address.	(IP address)
Virtual CPU Count	Displays the number of virtual CPUs assigned to this VM.	(number)
Memory Capacity	Displays the amount of memory available to this VM.	xxx [MB|GB]
IP Addresses	Displays one or more IP addresses assigned to the VM.	(IP address)
Power State	Displays whether the VM is powered on or powered off	On, Off
Network Adapters	Displays the number of network adapters available to this VM.	(# of adapter ports)
Disk Capacity	Displays the total disk capacity available to this VM.	xxx [GB|TB]
NGT Status	Displays if NGT is installed **Note:** The system displays the **Upgrade available** state if the **Ngt Cluster Version** is latest than the NGT installed version.	Installed, Not Installed, Latest, Upgrade available
Services Enabled	Displays the services enabled for this VM.	Calm, Karbon
NGT	Displays the NGT version installed on the VM.	<NGT version> such as 2.1.5
Ngt Cluster Version	Displays the NGT version available on the cluster.	<NGT cluster version> such as 2.3.2
(the following fields appear when the VM is allocated to a GPU)		
GPU Type	Displays the GPU operational mode. If it is vGPU, the following fields also appear.	vGPU, Passthrough, None

Parameter	Description	Values
GPU Configuration	Displays the vGPU profile used.	(vGPU profile name)
Framebuffer	Displays the size of the GPU framebuffer (RAM).	xxx GiB
Virtual Slice	Displays the virtual slice applied. The "virtual slice" reflects the approximate amount of physical GPU resources that the vGPU can receive. **Note:** The Virtual Slice and vGPU Guest Driver Version fields do not appear for passthrough GPUs.	(slice amount)
vGPU Guest Driver Version	Displays the version number of the vGPU guest driver.	(version number)

The following table describes the fields in the **Storage Properties** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. For more information about the parameters, see Storage Policy Details View on page 143.

Table 17: Storage Properties Fields

Parameter	Description	Values
Defined via	Displays the name of the storage policy that the VM is associated with.	Storage policy name
Encryption	Displays the status of encryption that the storage policy applies to the VM.	Enabled or Inherit from Cluster
Compression	Displays the status and type of compression that the storage policy applies to the VM.	Inline, Post Process, Off or Inherit from Cluster
QOS Metric	Displays the type of QOS metric that the storage policy applies to the VM with the throttled throughput value.	Throttled Throughput value in IOPS or MBps.

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

Table 18: Associations Fields

Parameter	Description	Values
Virtual Disks	Displays the number of virtual disks associated with the VM.	(number)
Storage Containers	Displays the number of storage containers associated with the VM.	(number)

Parameter	Description	Values
Categories	Displays the number of categories associated with the VM.	(number)
Policies	Displays the number of policies like storage policies and Image Placement policies that are associated with the VM.	(number)

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget allows quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective pull-down menus.

Example: Displaying Relationship between Entities

The following image indicates that VM *auto_pc_611cd08373101ab68d756c9a0* resides on host *cool-07-2* and cluster *auto_cluster_prod_rohan_bajaj_1ab68d756c98*.

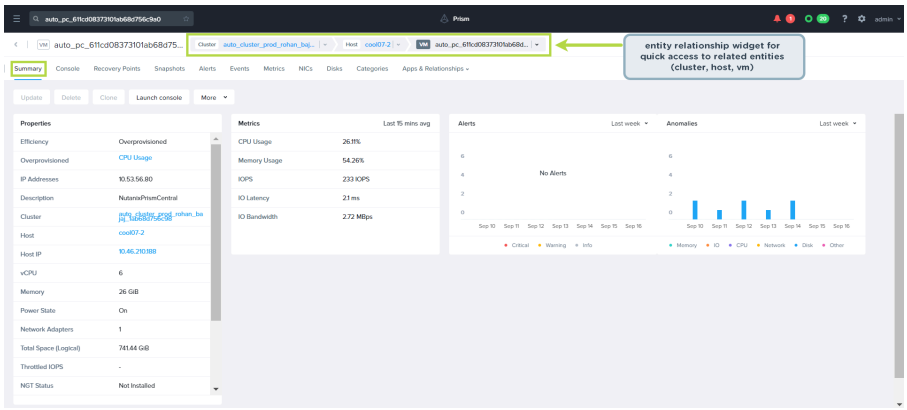

Figure 64: VM Summary: Entity Relationship Widget

Example: Viewing VMs on a Cluster

Click the **VM** pull-down menu to view the list of VM instances on the selected cluster. Alternatively, you search the VM instance name residing on the target cluster.

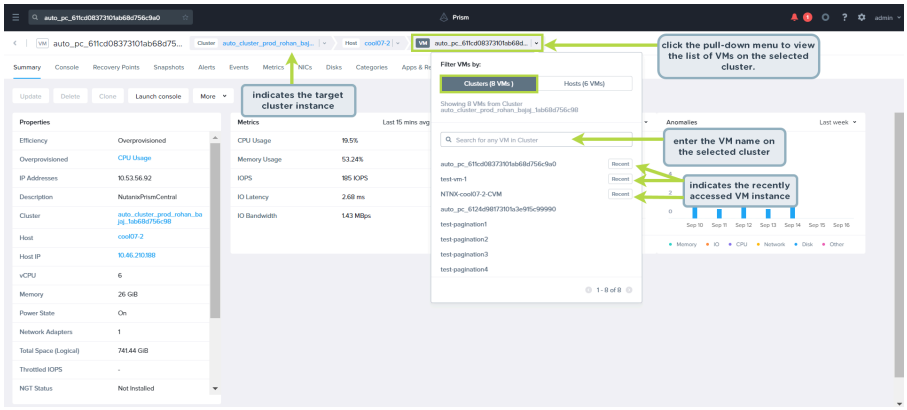

Figure 65: VMs on a Cluster

Example: Viewing VM Instances on a Host

Click the **VM** pull-down menu to view the list of VM instances on the selected host. Alternatively, you can search the VM instance name belonging to the target host.

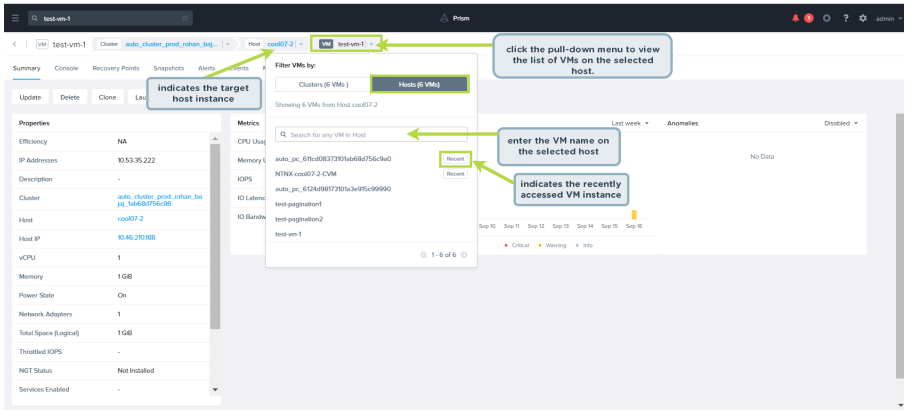

Figure 66: VMs on a Host

> **Note:**
>
> - The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
> - The filtered list of VMs display only the powered-on VM instances.
> - If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Console Tab

The **Console** tab displays the VM console screen. There are three icons above the console display (on the right).

- Click the left (three small boxes) icon to send a **Control-Alt-Delete** command to the console.
- Click the middle (camera) icon to take a screen shot of the console display.
- Click the right (box with an arrow in it) icon to open the console in a new window.

Recovery Points

The **Recovery Points** tab displays a list of recovery points (backup snapshots) when backups have been enabled. (The list is blank if there are no snapshots available.) The total number of recovery points and the latest and oldest recovery points are listed on the left. A list of all recovery points appears in a table on the right with the create time, location, expiry time, and recovery point type provided for each recovery point.

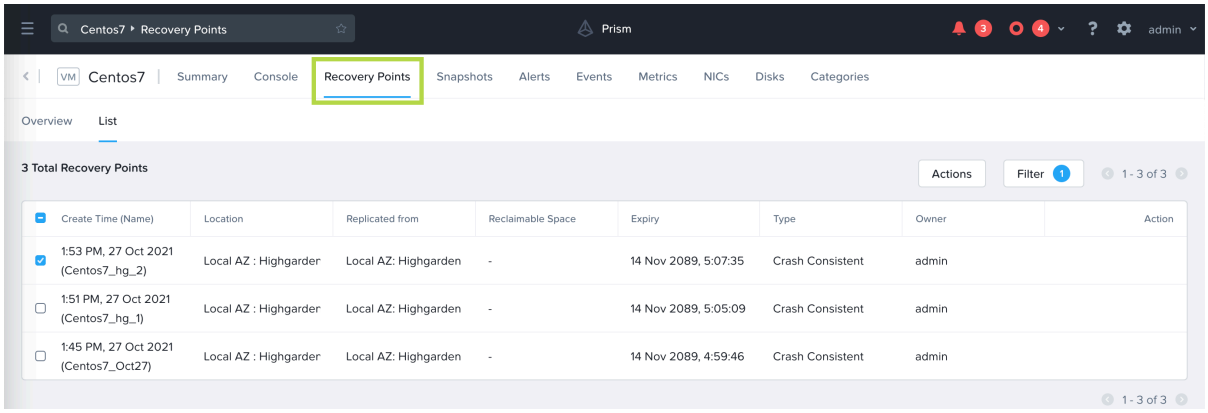

Figure 67: Recovery Points Tab

You can restore or replicate a VM from a recovery point. Select the recovery point and click the required action from the **Actions** menu:

- Click **Clone (Previously Restore)** to create a new cloned VM from the selected recovery point.
- Click **Replicate** to replicate a VM from the selected recovery point either locally or remotely in a state of a chosen recovery point.
- Click **Revert** to revert the VM to the selected recovery point and delete all the updated data after the recovery point creation.
- Click **Delete** to delete the recovery point.

Snapshots Tab

The **Snapshots** tab displays information in tabular form about backup snapshots of the VM taking from Prism Element. Each line represent a snapshot, and the following information is displayed for each snapshot:

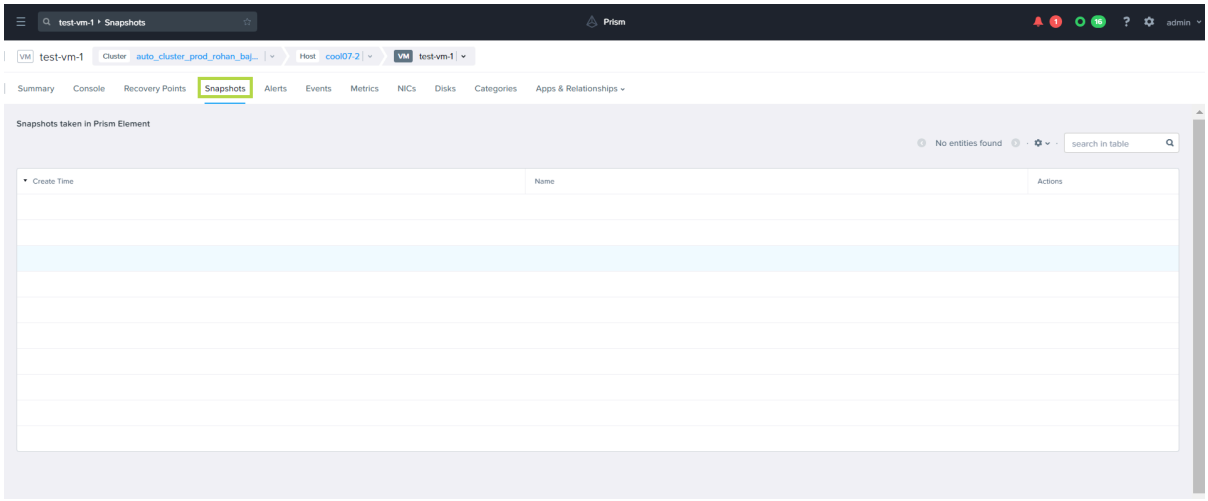

Figure 68: Snapshot Tab

- **Create Time**. Displays the time the backup snapshot was created (completed).
- **Name**. Displays a name for the backup if one was created.

- **Action**. Displays four action links:
  - Click the **Details** link to open a window that displays the snapshot details.

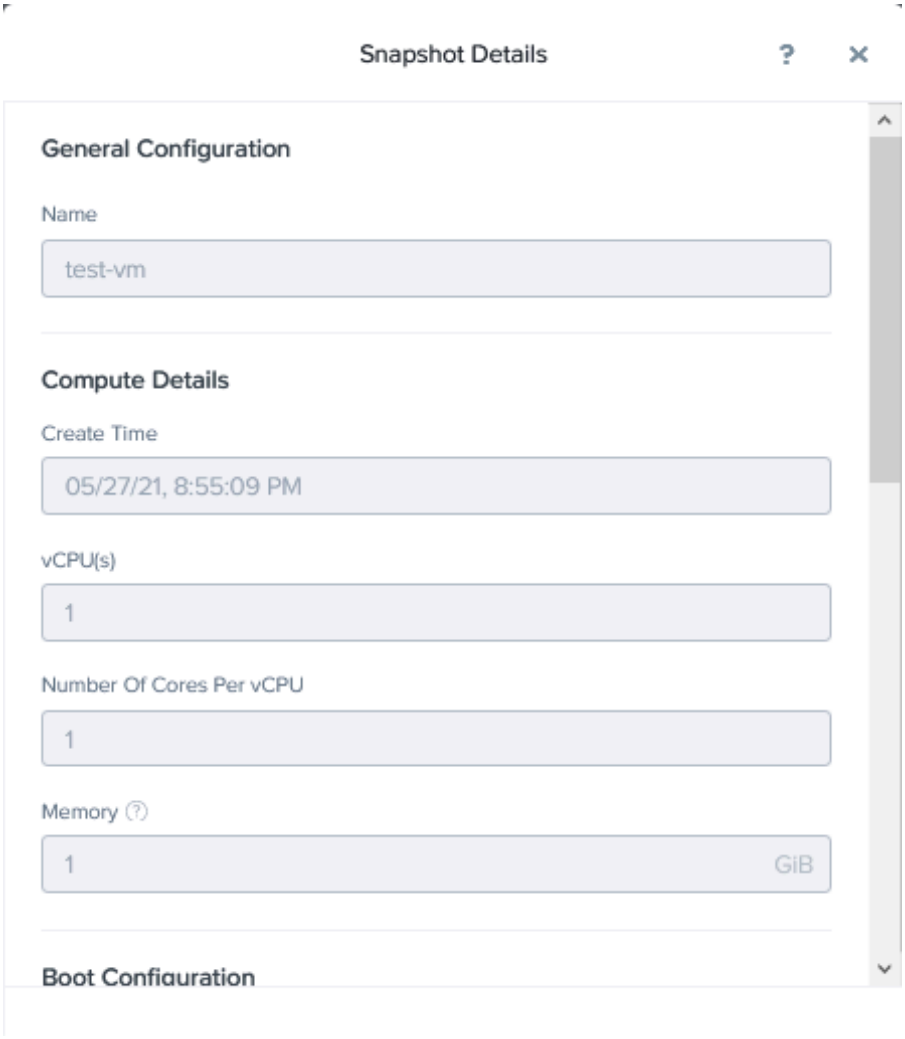

Figure 69: Snapshot Details Window

  - Click the **Clone** link to clone a VM from the snapshot.
  - Click the **Restore** link to restore the VM from the snapshot. This restores the VM back to the state of the selected snapshot.
  - Click the **Delete** link to delete the snapshot.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this VM (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this VM. (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the VM. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

> **Note:** The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 899). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check one or more appropriate boxes to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 288).

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

Table 19: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the VM (0–100%).
CPU Ready Time	Displays the current, high, and low percentage of CPU wait time (0–100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the VM (0–100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
IO Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
Usage	Displays separate graphs for current, snapshot, and shared storage usage (in GiBs) by the VM.
Working Set Size	Displays separate graphs for total, write, and read storage usage (in GiBs) for the VM working set size.
Network Packets Dropped	Displays separate graphs for the number of transmitted and received packets dropped.
Network Bytes	Displays separate graphs for the amount of transmitted and received bytes (in GiBs).

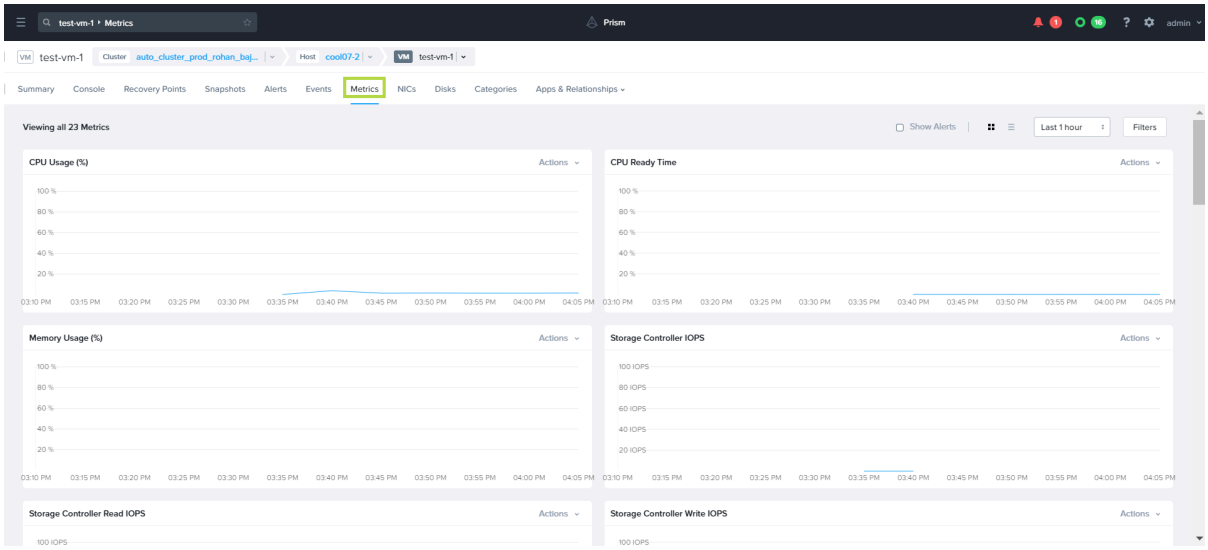

Figure 70: Metrics Tab: CPU Usage

NICs Tab

The **NICs** tab displays information in tabular form about the virtual NICs in the VM. Each line represent a virtual NIC, and the following table describes the fields.

Table 20: NIC Fields

Parameter	Description	Values
VLAN ID	Displays the VLAN name for this NIC.	(VLAN ID
MAC Address	Displays the virtual NIC MAC address.	(MAC address)
Network Connection State	Displays whether the NIC is connected to the network currently.	Connected, Disconnected
Requested IP Address	Displays the virtual NIC IP address	(IP address)
Action	Displays the available actions you can execute on this NIC.	Delete

Disks Tab

The **Disks** tab displays information in tabular form about the virtual disks in the VM. Each line represents a virtual disk, and includes the following fields.

- **Disk Address**: Displays the disk address (such as ide.0 or scsi.1).
- **Capacity**: Displays the disk capacity (in MiB or GiB).

Categories Tab

The **Categories** tab displays the categories and image placement policies associated with the VM. Each VM can have a one-to-many relationship with categories and the categories can have a many-to-one relationship with image placement policies.

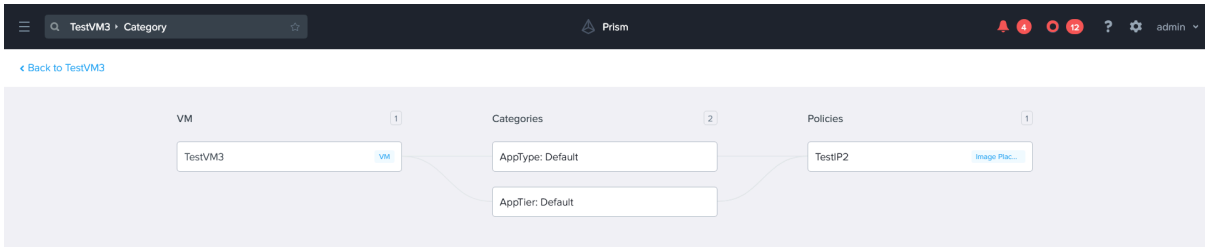

Figure 71: Categories Tab

For more information about categories management and image placement policies, see Category Management on page 795 and Image Placement Policies on page 690.

Apps & Relationships Tab

The **Apps & Relationships** tab displays a pull-down menu with two menu options: **Discovered Apps** and **App Relationships**.

The **Discovered Apps** menu option displays (up to) the top 20 applications that are communicating with other entities (usually client VMs). The discovered applications are sorted based on the number of clients communicating with the application. This menu option provides similar information as the application discovery dashboard, except it is filtered to display the details only for this VM (see Application Discovery View on page 832).

The **App Relationships** option displays three tabs: Visualization, Incoming List, and Outgoing List. All the three tabs display (up to) the top 20 results with respect to the incoming or outgoing client communication.

Visualization: Displays a visual representation of the VM and the applications running on it, along with its incoming and outgoing communication with other entities. Data for last 24 hours is fetched to display the information on the visualization tab. You can hover on an entity to see the applications it is communicating with. Hover on an application to see the application's incoming communication with external entities and outgoing communications of the VM on which the application is running.

> **Note:** In case of dual NIC VMs, click and expand the App listing to view the specific IP that is being used by the application.

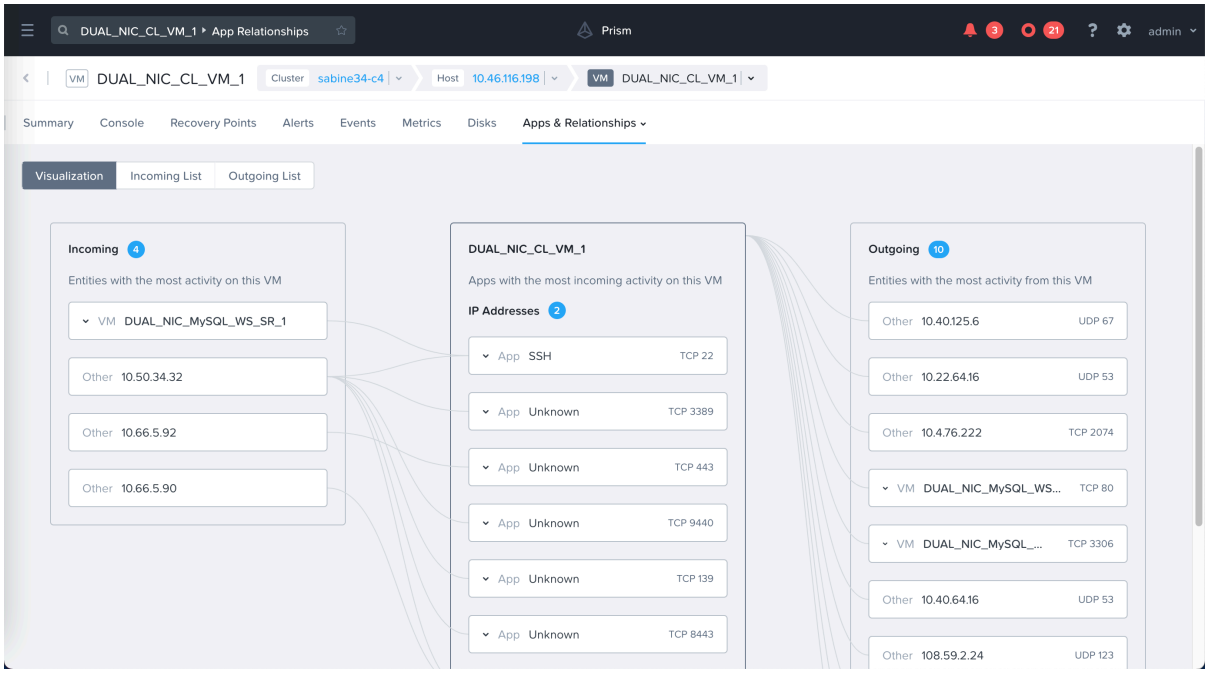

Figure 72: Visualization Tab

Incoming List: Displays a list of entities having incoming communication with this VM. Each line represents an incoming communication and includes the following fields.

Table 21: Incoming List Fields

Parameter	Description	Values
Source Entity Name	Displays the name of the communicating entity. If the entity is not identified by Prism Central (through registered PE or through vCenter using monitoring configurations feature), the name is displayed as unknown.	(VM name), Unknown
Source IP Address	Displays the IP Address of the communicating entity.	(IP Address)
Destination App	Displays the name of the application that is being communicated to. If the application is not an identified one, the Destination App is displayed as Unknown.	(application name), Unknown
Destination IP Address	Displays the IP Address of the VM.	(IP address)
Destination Port	Displays the port used by the application for communication.	(port number)
Destination App Identification Status	Indicates whether the application is an identified application type. An identified application is one that is a known type.	Identified, Unidentified

You can filter this list by opening the Filter pane to select a filter option. The following table describes the filter options for **Incoming List**. You can apply filters across multiple parameters. The filters fetch the latest discovery and communication data and accordingly display the result set, to the maximum of 20 results.

Table 22: Filter Pane Fields

Parameter	Description	Values
Source Entity Name	Filters on the source entity name. Select a condition from the pull-down list (**Contains**, **Equal to**, **Not equal to**, **Doesn't contain**, **Starts with**, **Ends with**) and enter a string in the field.	(source entity name string)
Entity Type	Filters on the entity type. A VM belongs to a cluster managed by Prism Central or vCenter (only if using monitoring configurations feature). An Unknown entity is one that is not identified by Prism Central.	VM, Unknown
Source IP Address	Filters on the source IP address.(same options as Source Entity Name).	(source IP address string)
Destination App	Filters on the destination app (same options as Source Entity Name).	(Destination App string)
Destination IP Address	Filters on the destination IP address (same options as Source Entity Name).	(destination IP address string)
TCP Port	Filters on the TCP port number.	(TCP port number)
UDP Port	Filters on the UDP port number.	(UDP port number)
Time Range	Filters on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

Outgoing List: Displays a list of entities having outgoing communication from this VM. Each line represents an outgoing communication and includes the following fields.

Table 23: Outgoing List Fields

Parameter	Description	Values
Destination Entity Name	Displays the name of the entity this VM is communicating to. If the entity is not identified by Prism Central (through registered PE or through vCenter using monitoring configurations feature), the name is displayed as unknown.	(entity name), Unknown
Destination IP Address	Displays the IP Address of the entity this VM is communicating to.	(IP address)
Destination Port	Displays the port on which this VM is communicating with the destination entity.	(port number)
Source IP Address	Displays the IP Address of the VM.	(IP address)

Parameter	Description	Values
Source App Identification Status	Indicates whether the communicating application is an identified application type. An identified application is one that is a known type.	Identified, Unidentified

You can filter this list by opening the Filter pane to select a filter option. The following table describes the filter options for **Outgoing List**. You can apply filters across multiple parameters. The filters fetch the latest discovery and communication data and accordingly display the result set, to the maximum of 20 results.

Table 24: Filter Pane Fields

Parameter	Description	Values
Destination Entity Name	Filters on the destination entity name. Select a condition from the pull-down list (**Contains**, **Equal to**, **Not equal to**, **Doesn't contain**, **Starts with**, **Ends with**) and enter a string in the field.	(destination entity name string)
Entity Type	Filters on the entity type. A VM belongs to a cluster managed by Prism Central or vCenter (only if using monitoring configurations feature). An Unknown entity is one that is not identified by Prism Central.	VM, Unknown
Destination IP Address	Filters on the destination IP address. (same options as Destination Entity Name).	(destination IP address string)
TCP Port	Filters on the TCP port number.	(TCP port number)
UDP Port	Filters on the UDP port number.	(UDP port number)
Source IP Address	Filters on the source IP address (same options as Destination Entity Name).	(source IP address string)
Time Range	Filters on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

VM Template Summary View

To access the templates dashboard, select **Compute & Storage** > **Templates** from the entities menu (see Entities Menu on page 13). The templates dashboard summary view displays information about templates across the registered clusters and allows you to access detailed information about each template.

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

- See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.
- See Creating a VM Template on page 614 for information about creating the templates.
- See Managing a VM Template on page 618 for information about managing the templates.

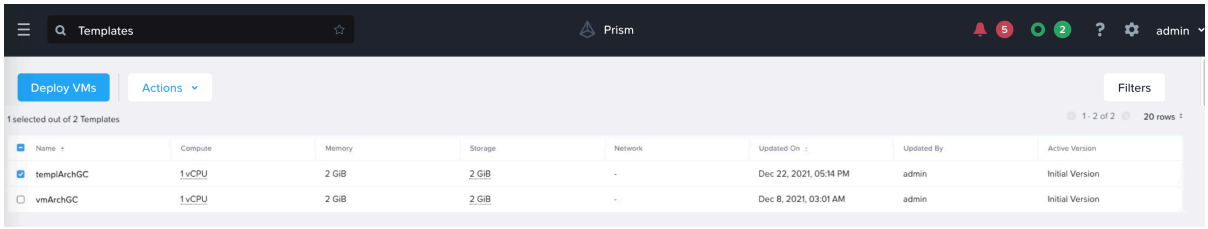

Figure 73: Templates View

Table 25: Template fields

Parameter	Description	Values
Name	Displays the template name. Clicking the name displays the details page for that template (see VM Template Details View on page 120).	(Template name)
Compute	Displays the virtual CPU count of the VM that is deployed using this template.	(vCPU count)
Memory	Displays the total amount of memory available to the VM that is deployed using this template.	xxx [MB|GiB]
Storage	Displays total amount of storage available to the VM that is deployed using this template.	xxx [MB|GiB]
Network	Displays the NIC count of the VM that is deployed using this template.	(NIC count)
Updated On	Displays the last date and time when the template has been updated.	(date and time)
Updated By	Displays the user who most recently updated the template.	(user)
Active Version	Displays the name of the active version of the template. An active version is the version of the template that by default gets deployed when you click the **Deploy VMs** button.	(version)

You can filter the template list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range. Click the plus (+) sign to the right of the parameter name to expand the fields for that parameter; click the minus (-) sign to contract those fields.

Table 26: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the template name. Select a condition from the pull-down list and enter a string in the field. It will return a list of templates that satisfy the template name condition/string. In this field, the condition menu options are **Contains**, **Doesn't contain**, **Starts with**, **Ends with**, and **Equal to**.	(template name string)

You can sort the list of templates by clicking on **Name** and **Updated On** columns.

The templates view includes an action button **Deploy VMs** along with the **Actions** menu. To deploy VMs using the VM template, click the **Deploy VMs** button. By default the active version of the template gets deployed.

The **Actions** menu appears when one or more templates are selected. It includes the following actions:

- Actions that can be applied to multiple templates: **Delete Template**.
- Actions that can be applied to only one template at a time: **Update Guest OS**, **Complete Guest OS Update**, **Cancel Guest OS Update**, **Update Configuration**, and **Delete Template**.

The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the selected template. See Managing a VM Template on page 618 for instructions on how to perform these actions.

VM Template Details View

To access the details page of a template, go to the templates view (see VMs Summary View on page 94) and click the template name. You can also access the details page by clicking the template name wherever that name appears, such as in a dashboard widget or search result.

The templates details page includes two views: **Summary** view and **Versions** view.

Summary View

The **Summary** view includes an action button **Deploy VMs** along with the **Actions** menu. There are only two options available in the **Actions** menu under this view: **Set as Active** and **Delete**. You can select a version and set it as active, or delete a selected version. You cannot delete an active version of the template.

The following set of widgets appears on the details page: **Template**, **Active Version**, **VM Properties**, and **VM Resources**. Click the **Back to Templates** link to return to the template summary view.

- **Template**: The widget displays the following:
  - **Available Versions**: Number of versions available
  - **Description**: User-provided description of the template
  - **Last Updated by**: User who has last updated the template
  - **Last Updated on**: Date and time of last update to the template
- **Active Version Configuration**: The widget displays the following:
  - **Version Name**: Name of the currently active version of the template
  - **Change Notes**: User-provided change notes for the active version

- **VM Properties**: The widget displays the following:
  - **Compute**: Compute resources of the VM such as CPU, core per CPU, and memory
  - **Boot Type**: Boot type of the VM such as legacy or UEFI
  - **Guest OS**: Guest OS of the VM. If the source VM of the template does not have NGT installed on it, then this field is displayed as unknown.
  - **Guest Customization**: Status of the guest customization application (applied/not applied)
  - **NGT Status**: Status of the NGT installation
- **VM resources**: The widget displays the following:
  - **GPU**: Number of GPU assigned to the VM
  - **Disks**: Number of disks and total disk capacity assigned to the VM
  - **Network**: Number of NICs assigned to the VM

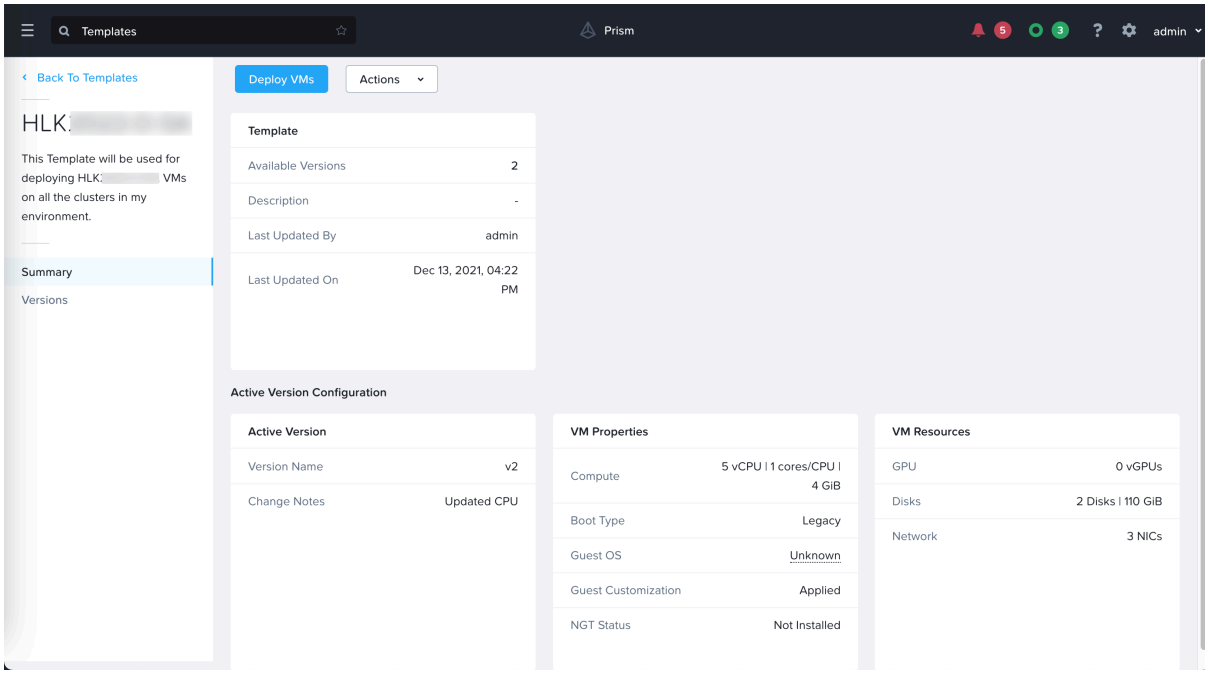

Figure 74: Summary View of a Template

Versions View

The **Versions** view includes an action button **Deploy VMs** along with the **Actions** menu. There are only two options available in the **Actions** menu under this view: **Set as Active** and **Delete**. You can select a version and set it as active, or delete a selected version. You cannot delete an active version of the template.

Table 27: Versions Fields

Parameter	Description	Values
Name	Displays the version name. This column also indicates if the version is an active version.	(template name)
Compute	Displays the virtual CPU count of the VM that gets deployed using this version.	(vCPU count)
Memory	Displays the total amount of memory available to the VM that gets deployed using this version.	xxx [MB|GiB]
Storage	Displays total amount of storage available to the VM that gets deployed using this version.	xxx [MB|GiB]
Network	Displays the NIC count of the VM that gets deployed using this version.	(NIC count)
Added On	Displays the last date and time when the version was added.	(date and time)
Added By	Displays the user who added the version.	(user)
Notes	Displays the user-specified change note for the version.	(version)

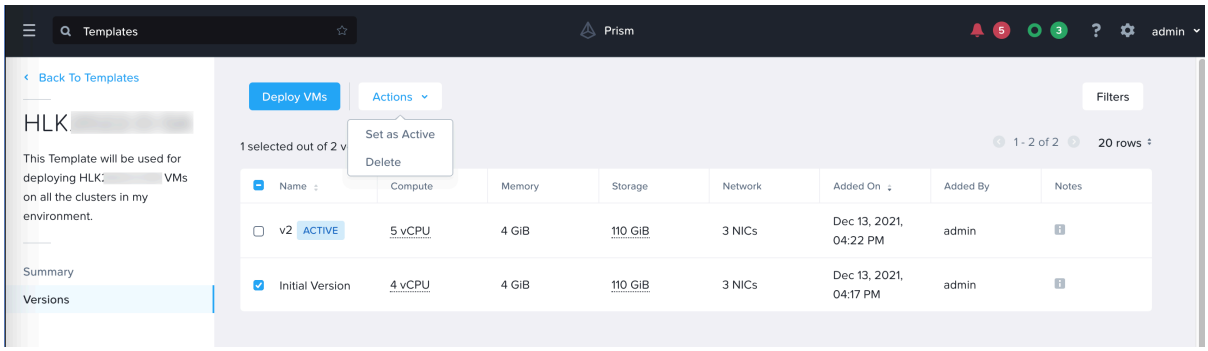

Figure 75: Versions View of a Template

You can sort the list of versions by clicking on **Name** and **Added On** columns.

OVAs View

To access the OVAs dashboard, select **Compute & Storage** > **OVAs** from the entities menu (see Entities Menu on page 13). The OVAs dashboard allows you to view information about the OVAs available in the cluster.

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

- See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
- See OVA Management on page 698 for information about the OVA operations.

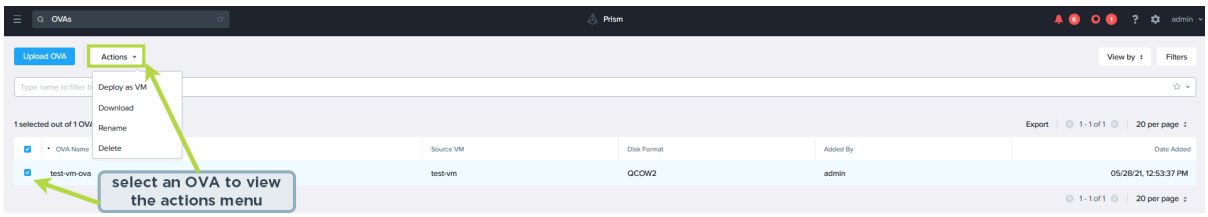

Figure 76: OVAs Dashboard

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

Table 28: OVA List Fields

Parameter	Description	Values
OVA Name	Displays the name of the OVA file.	(file name)
Source VM	Displays the name of source VM that was exported as an OVA. The alert indicators for the task status are displayed beside the source VM for which the task is run. See the Alert indicator table for more information.	(file name)
Disk Format	Displays the format of the disks in the OVA.	[QCOW2|VMDK]
Added by	Indicates which user added this OVA.	(text string)
Date added	Displays the date and time when the OVA was added.	(date in mm/dd/yyyy format, time in hr:min:sec [AM|PM] format) Example: 07/30/20, 3:00:20 PM

Table 29: OVA Alerts

OVA Alerts	Status Message sample	Indicates...
Figure 77:	`Incorrect file uploaded. Multiple OVA files present in OVA.`	That the upload process failed during validation. You must fix the issue in the OVA file or upload the correct, compatible OVA file.
Figure 78:	`Upload for OVA file has not started. Please start OVA file upload using Resume Upload option.` `Uploaded OVA file has not been verified. Please start verification using Resume Upload option.`	That there is an interruption in one of the three sub-tasks in the upload process. You can resume the upload. See the *Resume Upload* action procedure.
Figure 79:	`OVA file upload is in progress.`	That the OVA upload is in progress without any error or interrupts. Check the Tasks page after some time to verify the successful upload.

A sample of the error hover message is as follows:

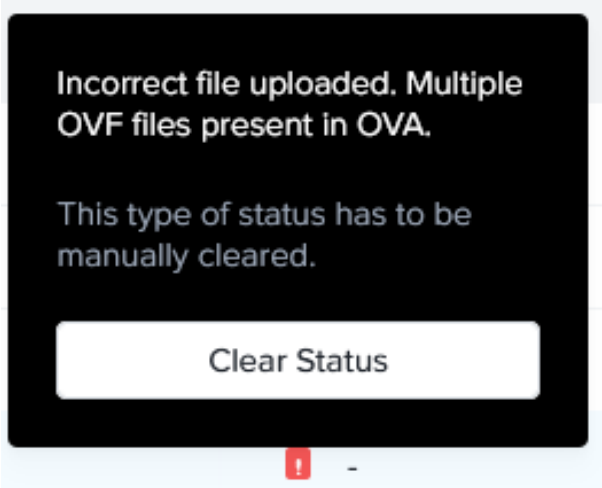

Figure 80: Sample error hover message

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

For each filter, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.

Table 30: Filter Pane Fields

Parameter	Description	Values
OVA Name	Filters on the OVA file name. It returns a list of OVAs that satisfy the name condition/string.	(OVA file name string)
Source VM	Filters on the Source VM name. It returns a list of OVAs for which the Source VM field satisfies the condition/string.	(Source VM name string)
Disk Format	Filters on the disk format. Check the box(es) for the OVAs of the required disk format.	VMDK, QCOW2

The **Actions** menu appears when a OVA is selected. It allows you to:

- Deploy an OVA file as a VM.
- Download an OVA file to your local machine.
- Rename an OVA file.
- Delete an OVA file.
- Resume upload when an upload is interrupted and shows an error.

> **Note:** The resume upload action is not available in case the concatenate process of the upload is interrupted. In such a case, you need to run the Concatenate API from the **REST API Explorer** to resume the upload.

Images Summary View

To access the images dashboard, select **Compute & Storage** > **Images** from the **Entities** menu (see Entities Menu on page 13). The images dashboard allows you to view information about images available from this Prism Central instance.

> **Note:** This section describes the information and options that appear in the images dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See Image Management on page 670 for information about adding and managing images through Prism Central.
> - See Image Placement Policies on page 690 for information about image placement policies.
> - See Bandwidth Throttling Policies on page 695 for information about bandwidth throttling policies.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the Images.

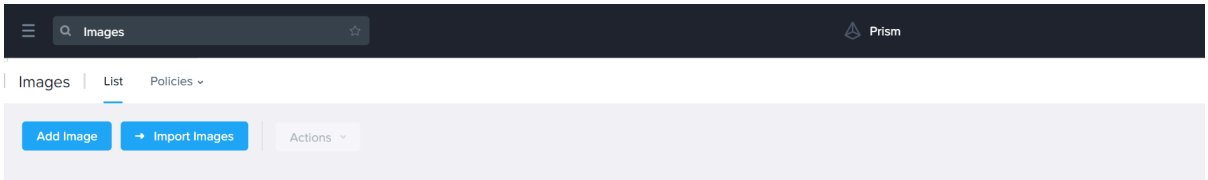

Figure 81: Images List View

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

Table 31: Image List Fields

Parameter	Description	Values
Name	Displays the image name. Clicking the name displays the details page for that image (see Image Details View on page 127).	(name)
Description	Indicates which user uploaded this image.	(text string)
Type	Displays the image type.	ISO, Disk
Size	Displays the image size.	xxx [MB|GB]
Creator	Displays who created the image.	(name)

You can filter the images list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 32: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the image name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of images that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It will return a list of images whose description field satisfy the condition/string.	(description string)
Type	Filters on the image type. Check the box(es) for the desired image types.	Disk, ISO

You can group the images list in the following ways:

- The **Group** pull-down menu allows you to group by type (disk or ISO).
- The **View by** pull-down menu allows you to view the entries by categories, creator, description, size, and type.

The images dashboard includes the following action buttons:

- **Add Image** Click to add an image (see Creating an Image on page 671).
- **Import Images** Click to import images from registered clusters to Prism Central (see Importing Images to Prism Central on page 689.

The **Actions** menu appears when one or more images are selected. It includes delete, update, and add image to catalog options (see Modifying an Image on page 688). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Policies Tab

The **Policies** tab displays the **Placement Policies** (see Image Placement Policies on page 690) and **Bandwidth Throttling Policies** (see Bandwidth Throttling Policies on page 695).

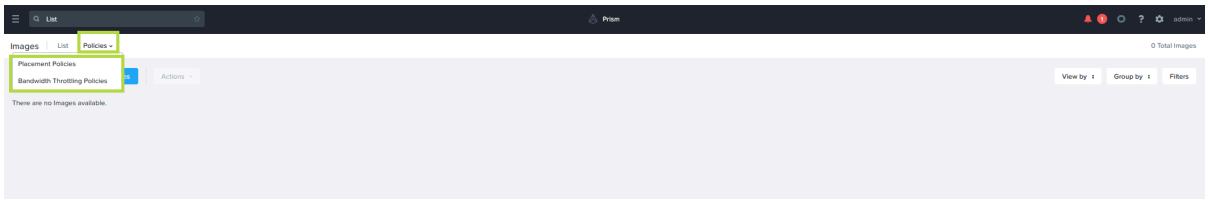

Figure 82: Images Policies Tab

Image Details View

You can view the detailed information about an image by going to the details page for that image.

To access the details page for an image, go to the images dashboard (see Images Summary View on page 125) and click the image name. An image details page includes the following:

- Image name (upper left). You can switch from one image to another by selecting a different image name from the pull-down list.
- Image information (middle left). The information includes the image description, type, size, and creator name (see Creating an Image on page 671).
- Tab options (lower left). Select the **Location** (default) or **Policies** tab to display that tab content to the right.
- Action buttons (upper right). When the window size is not large enough, some actions appear under a **More** menu button.
  - To delete the image, click the **Delete** button
  - To modify the name, description, or type, click the **Update Image** button (see Modifying an Image on page 688).
  - To add the image to the catalog, click the **Add Image to Catalog** button (see Adding a Catalog Item on page 667).
  - To assign the image to a category, click the **Manage Categories** button (see Assigning a Category on page 797).
  - To open a help page in a separate tab or window, click the question mark icon.
  - To close the details page, click the **X** icon.

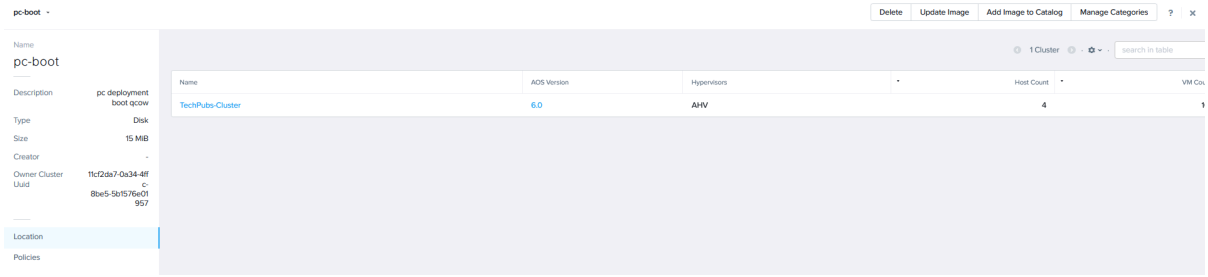

Figure 83: Image Details View (Location Tab)

Location Tab

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

Table 33: Location Fields

Parameter	Description	Values
Name	Displays the name of the cluster in which the image is stored. Each row represents another cluster that contains the image.	(cluster name)
AOS Version	Displays the AOS version running in the cluster.	(AOS version number)
Hypervisor	Displays the name of the hypervisor running in the cluster.	AHV, ESXi, Hyper-V
Host Count	Displays the number of hosts in the cluster.	(integer)
VM Count	Displays the number of VMs in the cluster	(integer)

Policies Tab

The **Policies** tab lists the policies that apply to the image (see Image Placement Policies on page 690). The following table describes the information displayed in the **Policies** tab.

Table 34: Policies Fields

Parameter	Description	Values
Policy Name	Displays the policy name.	(cluster name)
Status	Indicates whether the image is compliant according to the policy.	compliant, noncompliant
Details	Lists the clusters in which the policy is enforced.	(cluster names)

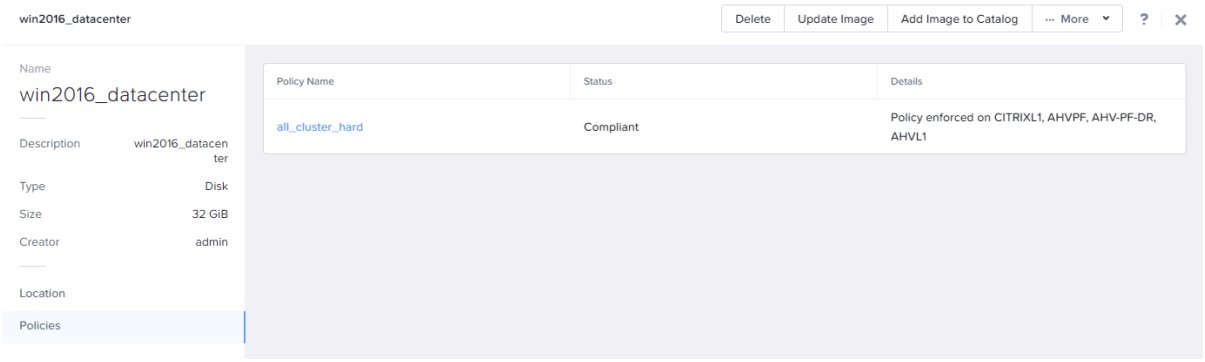

Figure 84: Image Details View (Policies Tab)

Catalog Items View

To access the catalog items dashboard, select **Compute & Storage** > **Catalog Items** from the **Entities** menu (see Entities Menu on page 13). The catalog items dashboard allows you to view information about the images and VMs in the catalog.

> **Note:** This section describes the information and options that appear in the catalog items dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See Catalog Management on page 666 for information about the catalog service and how to add (or delete) items.

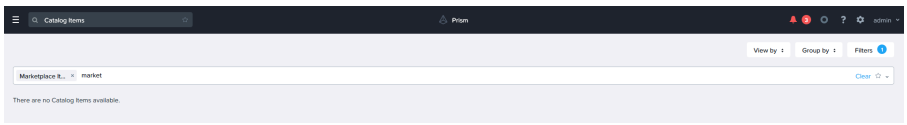

Figure 85: Catalog Items Dashboard

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

Table 35: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the item name.	(item name)
Type	Displays the type of item (image or VM).	Image, VM
Description	Indicates the description for the catalog item	(description string)

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 36: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the catalog item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of catalog items that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of catalog items whose description field satisfy the condition/string.	(description string)
Type	Filters on the item type. Check the box(es) for the desired item types.	VM, Image

You can group the catalog items list in the following ways:

- The **Group** pull-down menu allows you to group by type (VM and image).
- The **View by** pull-down menu allows you to view by name, description, message, percent, reason, state, and type.

Storage Containers Summary View

To access the storage containers dashboard, select **Compute & Storage** > **Storage Containers** from the entities menu (see Entities Menu on page 13). The storage containers dashboard summary view displays information about storage containers across the registered clusters and allows you to access detailed information about each storage container. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> **Note:** This section describes the information and options that appear in the storage containers dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.

Summary Tab

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

- **Highlighted Entities**: Displays a list of the storage containers with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that storage container. Click the **View all XX Storage Containers** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of storage container-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

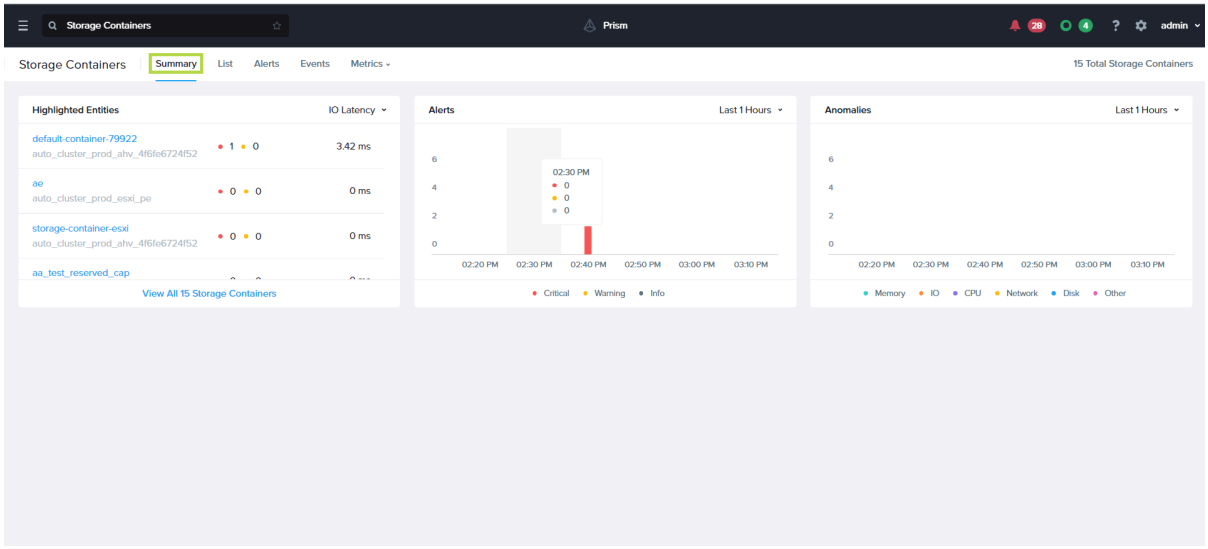

Figure 86: Storage Containers Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the storage containers across the registered clusters. The following table describes the fields that appear in the storage containers list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

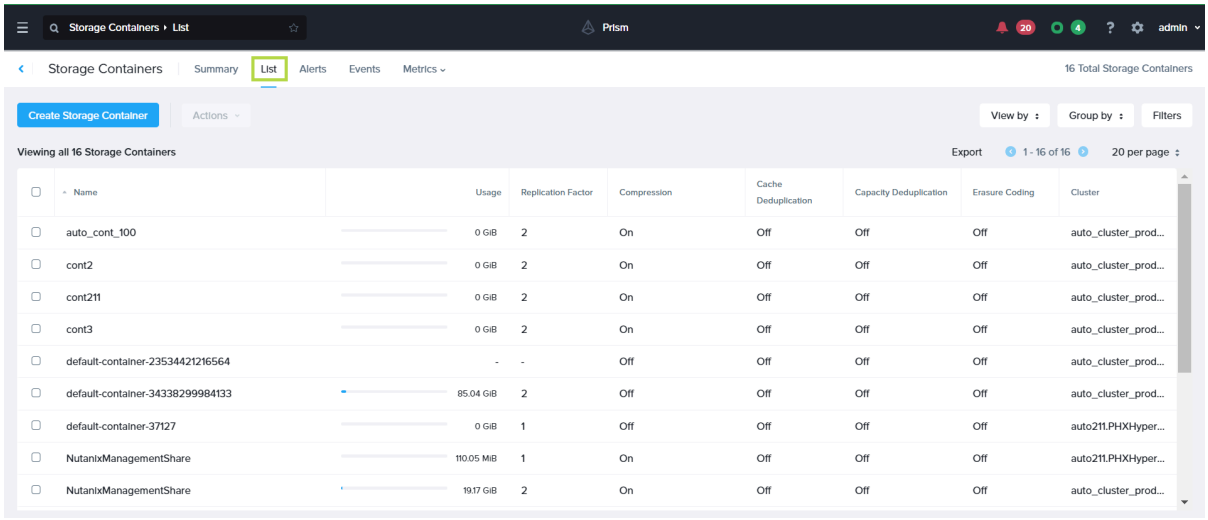

Figure 87: Storage Containers List Tab

Table 37: Storage Containers List Fields

Parameter	Description	Values
"General" Focus Fields		

Parameter	Description	Values
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 136).	(name)
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the storage container. Hovering the cursor over the bar displays a window listing the used space, free space, and total space in the storage container.	xxx [GB|TB]
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether capacity deduplication is enabled, that is deduplication compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Cluster	Displays the name of the cluster in which the storage container resides. Click the name to display the cluster details (see Cluster Details View on page 209).	(cluster name)
"Performance" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 136).	(name)
Free Space	Displays the amount of free space available in the storage container.	xxx [GB|TB]
Used Space	Displays the amount of used space in the storage container.	xxx [GB|TB]
Total Space	Displays the total amount of storage space in the storage container.	xxx [TB]

Parameter	Description	Values
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
"Optimization" Focus Fields		
Name	Displays the name of the storage container. Click the name to display the details page for that storage container (see Storage Container Details View on page 136).	(name)
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	*xx*:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GB|TB]
Compression Delay	Displays the time delay to perform compression on the data.	xx [m]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GB|TB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	*xx*:1

You can filter the storage containers list based on various parameter values. The following table describes the filter options available when you open the storage containers view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 38: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of storage containers that satisfy the name condition/string.	(storage container name string)

Parameter	Description	Values
Cluster	Filters on the cluster name. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of storage containers in the clusters that satisfy the name condition/string.	(cluster name string)
Cache Deduplication	Filters on the cache deduplication setting. Select one or more settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
Replication Factor	Filters on the replication factor. Check the box for the desired replication factor (1, 2, 3). Only existing RF values appear. For example, if all containers are at RF 2, just a single box for RF 2 would appear. The number of storage containers currently in each setting is displayed on the right of the line.	1, 2, 3
Capacity Deduplication	Filters on the capacity deduplication setting. Select one or more settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off, None, Inline, Post Process
Erasure Coding	Filters on the erasure coding setting. Select one or both settings to return a list of storage containers with those settings. The number of storage containers currently in each setting is displayed on the right of the line.	On, Off
Free Space	Filters on the available storage space for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with available capacity in that range.	([xx] to [yy] GiB range)
Health	Filters on the storage container health state. Select one or more states to return a list of storage containers in those states. The number of storage containers currently in each state is displayed on the right of the line.	Critical, Warning, Good
Used Space	Filters on the used storage space for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with used storage in that range.	([xx] to [yy] GiB range)
Total Space	Filters on the total available capacity for a storage container. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of storage containers with maximum capacity in that range.	([xx] to [yy] GiB range)

Parameter	Description	Values
IOPS	Filters on the current IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of storage containers with IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> MBps**" field. It will return a list of storage containers with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of storage containers with average I/O latency in that range.	([xx] to [yy] range)

You can group the storage containers list in the following ways:

- The **Group** pull-down menu allows you to group the host entries by replication factor, compression, erasure coding, cluster, or health setting. (You can only choose one.)
- The **View by** pull-down menu allows you to group the host entries by the information parameters (fields), which vary depending on whether you selected the **General**, **Performance**, or **Optimization** focus. (You can only choose one parameter.) You can also use the **+ Add custom** entry to create a custom view.

There are no action options available from the storage containers view (no action buttons and no **Actions** menu options when a storage container is selected).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just storage container-related alerts across the registered clusters (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just storage container-related events across the registered clusters (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the storage containers. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

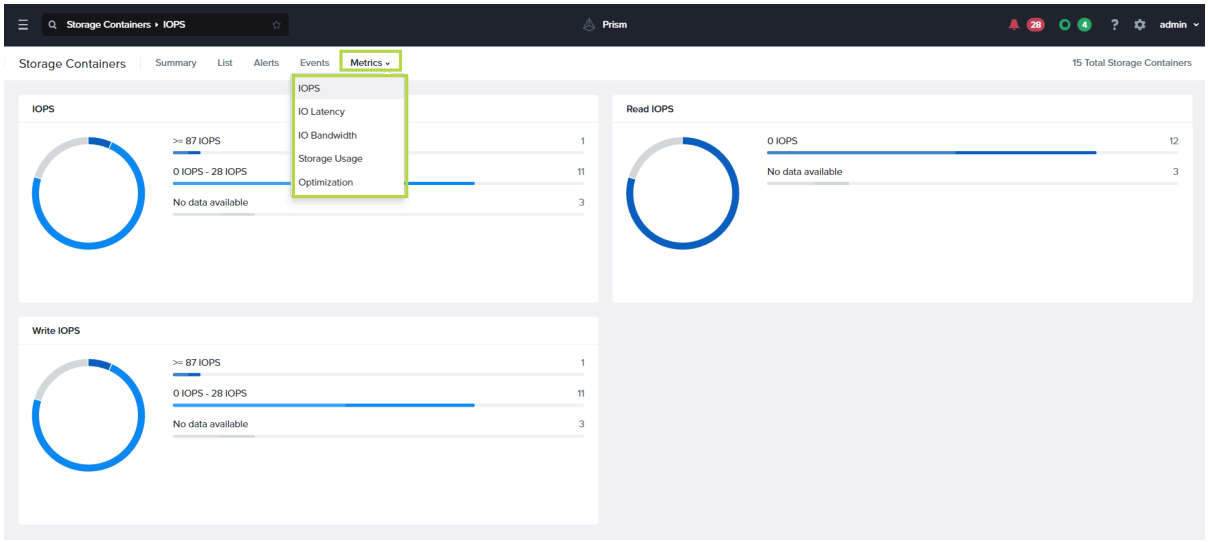

Figure 88: Storage Containers Metrics Tab

Table 39: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000). **Note:** The same format also applies to the other metrics in this table.
I/O Latency	Displays total, read, and write I/O latency rate graphs.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.
Storage Usage	Displays storage usage graph.
Optimization	Displays replication factor, data reduction ratio, and overall efficiency graphs.

Storage Container Details View

To access the details page for a storage container, go to the storage container dashboard **List** tab (see Storage Containers Summary View on page 130) and click the storage container name. You can also access the details page by clicking the storage container name wherever that name appears, such as in a dashboard widget or search result.

The storage container name and the following set of tabs appear at the top: **Summary**, **Alerts**, **Events**, **Metrics**, and **Usage**. Click a tab to display that information. (Click the arrow at the upper left to return to the summary view.)

Summary Tab

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

- A **Properties** widget that displays summary information about the storage container (see following table).
- A **Usage** widget that displays usage data for the storage container (see following table).

- An **Optimization & Performance** widget that displays optimization and performance data for the storage container (see following table).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last week** (default), **Last 1 Hours**, or **Last 24 hours** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

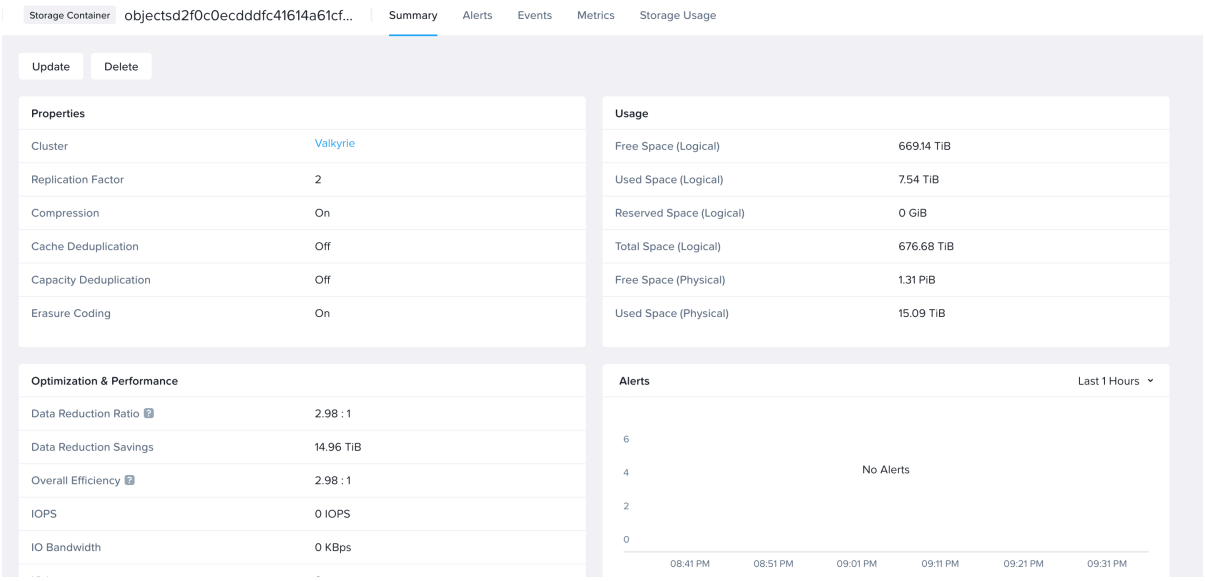

Figure 89: Storage Container Summary Tab

The following table describes the fields in the **Properties**, **Usage**, and **Optimization & Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 40: Storage Container Widget Fields

Parameter	Description	Values
"Properties" Widget		
Cluster	Displays the name of the cluster in which the storage container resides. Click the name to display the cluster details (see Cluster Details View on page 209).	(cluster name)
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off|On]

Parameter	Description	Values
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to **On** causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether capacity deduplication is enabled, that is deduplication compression applied to data on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]
Filesystem Whitelists	Displays the IP addresses for file systems that are whitelisted.	[IP addresses]
Compression Delay	Displays the delay (number of minutes) before data changes are compressed. A zero value indicates that compression is immediate (not delayed).	xx min
Thick Provision	Displays the reserved storage capacity for thick provisioned VMs.	xxx [GB|TB]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GB|TB]
"Usage" Widget		
Free Space (Logical)	Displays the amount of logical (effective) free space available in the storage container. **Note:** Logical space accounts for the storage container replication factor. RF1 means the logical and physical spaces are the same, RF2 means the logical space is half the physical space, and RF3 means the logical space is a third of the physical space.	xxx [GB|TB]
Used Space (Logical)	Displays the amount of logical used space in the storage container.	xxx [GB|TB]
Reserved Space (Logical)	Displays the amount of logical reserved space in the storage container.	xxx [GB|TB]
Total Space (Logical)	Displays the amount of logical total space in the storage container.	xxx [GB|TB]
Free Space (Physical)	Displays the amount of physical free space available in the storage container.	xxx [GB|TB]
Used Space (Physical)	Displays the amount of physical used space in the storage container.	xxx [GB|TB]
Total Space (Physical)	Displays the amount of physical total space in the storage container.	xxx [GB|TB]

Parameter	Description	Values
"Optimization & Performance" Widget		
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	*xx*:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GB|TB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	*xx*:1
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
Erasure Coding	Displays whether erasure coding is enabled for the storage container.	[On|Off]

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this storage container (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this storage container (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the storage container. Click the **Metrics** tab to display graphs for nine metrics: **Storage Controller IOPS**, **Storage Controller Read IOPS**, **Storage Controller Write IOPS Storage Controller Latency**, **Storage Controller Read Latency**, **Storage Controller Write Latency**, **Storage Controller IO Bandwidth**, **Storage Controller Read IO Bandwidth**, **Storage Controller Write IO Bandwidth**. Each graph is a rolling time interval performance or usage monitor. The baseline range (when established) appears as a blue band in the graph.

> **Note:** The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 899). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week).
- Click the **Actions** pull-down menu (upper right of any graph) to monitor the metric performance in greater detail (see Analysis Dashboard (Prism Central) on page 799):
  - Click **Add to Analysis** to add an active monitor for the metric on the Analysis dashboard.
  - Click **Select Analysis Session** to assign the metric to a target session.
- Click the filter button to display the filter panel; check (uncheck) the desired boxes to display (hide) those metrics.
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 288).

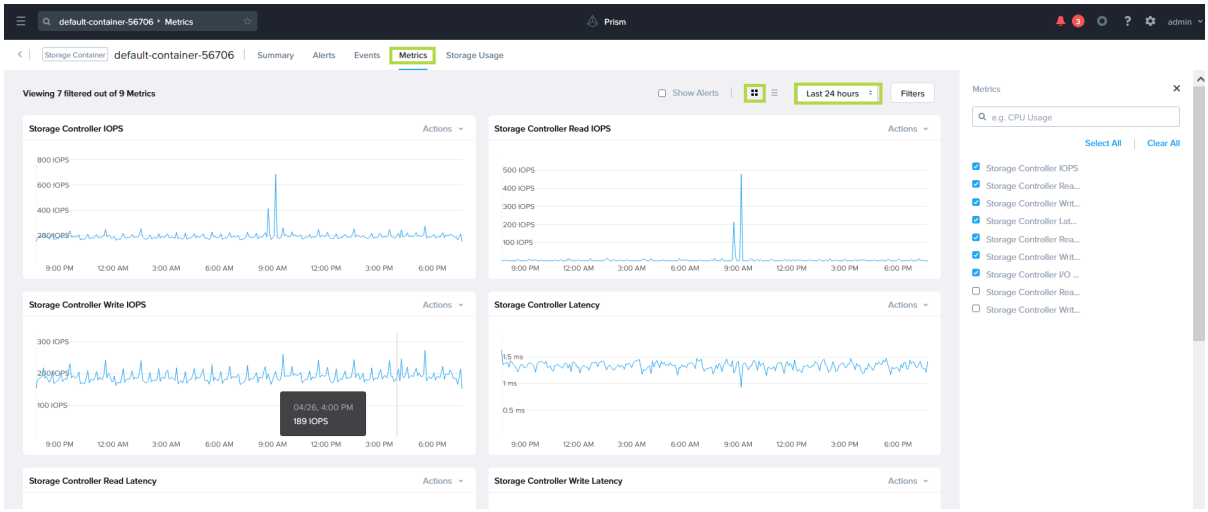

Figure 90: Storage Container Metrics Tab

Storage Usage Tab

The **Storage Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of storage container storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of container storage space used by each disk tier (SSD and DAS-SATA).

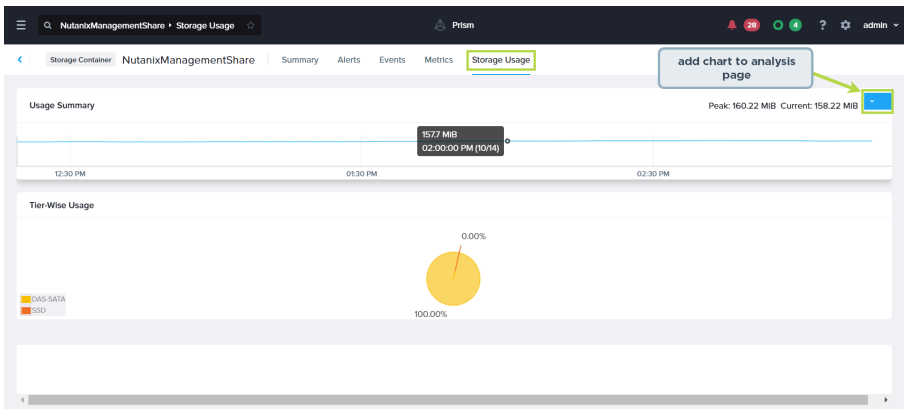

Figure 91: Storage Container Usage Tab

Storage Policies Summary View

To access the **Storage Policies** dashboard, select **Compute & Storage** > **Storage Policies** from the entities menu (see Entities Menu on page 13).

> **Note:** **Storage Policies** dashboard is not available if you log on as a view-only user.

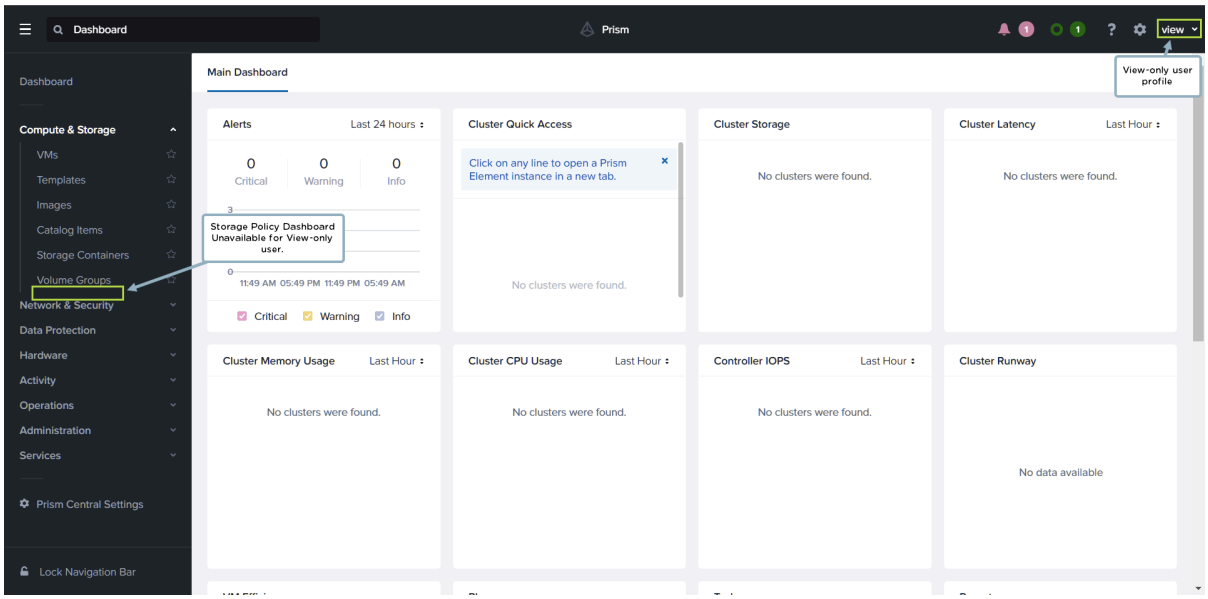

Figure 92: Storage Policy Dashboard - Unavailable for view-only user

The **Storage Policies** dashboard summary view displays information about the storage policies configured in Prism Central for the registered clusters and allows you to access detailed information about each storage policy. The dashboard has **List** tab with a display area below the selected tab that displays a list of the storage policies.

> **Note:** This section describes the information and options that appear in the **Storage Policies** dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.

- See Storage Policy Management on page 657 for information about storage policies, creating and managing storage policies.

List Tab

The **List** tab appears by default when you first open the **Storage Policies** dashboard, displays a list of the security policies across the registered clusters.

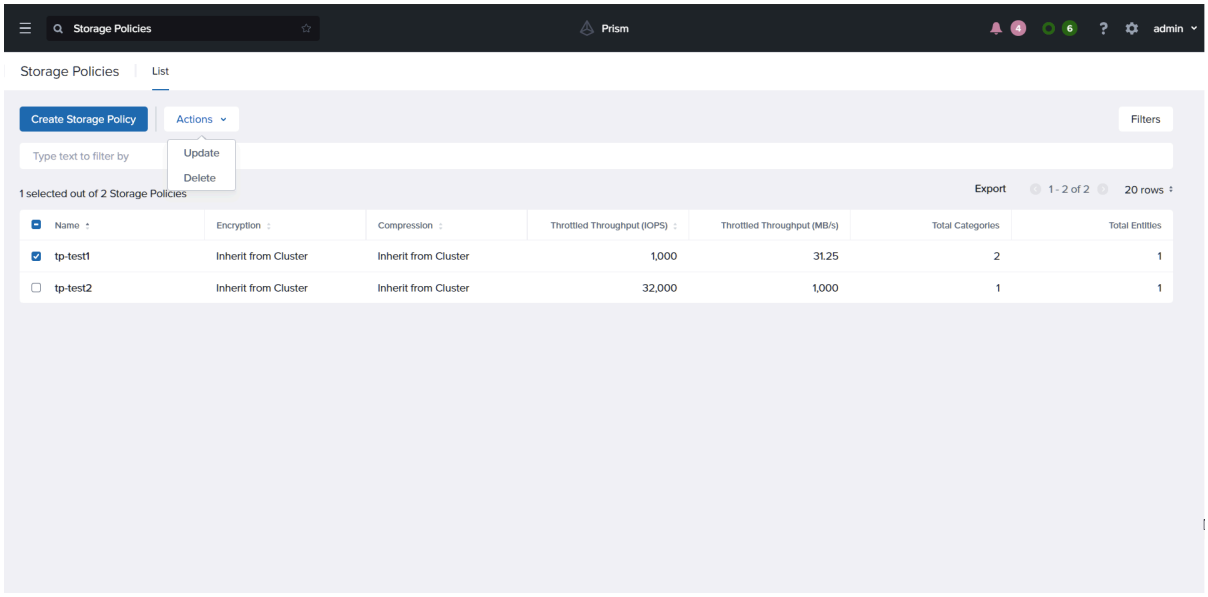

Figure 93: Storage Policy List Tab

The **List** tab displays the following:

- **Create Storage Policy**button: Click this button to open the **Create Storage Policy** page. See Creating or Updating a Storage Policy on page 660 for more information about creating a storage policy.
- **Actions** drop down menu: This menu includes two actions - **Update** and **Delete**.
- You can filter the list of storage policies based on the following properties:
  - Replication Factor: **2**, **3**, or **Inherit from Cluster**
  - Encryption: **Enabled**, or **Inherit from Cluster**
  - Compression: **On**, **Off**, or **Inherit from Cluster**
  - Throttled Throughput: Enter a range of throughput values to filter with, using the **From** and **To** input boxes.

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

Table 41: Storage Policies List Fields

Parameter	Description	Values
Name	Displays the name of the Storage Policy. Clicking the name displays the details page for that storage policy.	(Name of storage policy)
Encryption	Displays the encryption mode selected.	Enabled or Inherit from Cluster
Compression	Displays the compression mode selected.	Inline or Post Process (if enabled), Disabled or Inherit from Cluster
Throttled Throughput (IOPS)	Displays the throttled throughput value in terms of IOPS.	(Integer number)
Throttled Throughput (MBps)	Displays the throttled throughput value in MBps.	(Integer number)
Total Categories	Displays the total number of Categories that are associated with the storage policy.	(Integer number)
Total Entities	Displays the total number of entities that the storage policy is applied to. **Note:** A single storage policy can be applied to multiple entities. Multiple storage policies cannot be applied to a single entity.	(Integer number)

For more information about Storage Policies, see Storage Policy Management on page 657.

- For information about creating or updating storage policies, see Creating or Updating a Storage Policy on page 660.
- For information about deleting storage policies, see Deleting a Storage Policy on page 664.

Storage Policy Details View

To access the details page for a storage policy, go to the storage policy **List** tab (see Storage Policy Management on page 657) and click the name of the storage policy you want to view the details of. You can also access the details page by clicking the name of the storage policy wherever that name appears, such as in a dashboard widget or search result.

The storage policy details page includes the following tabs:

- **Summary**—displayed by default when you open the storage policy details page. This tab provides information about the storage policy summarized into widgets.
- **Entities**—displays the list of entities including VMs associated with the storage policy.

Summary Tab

The **Summary** tab displays the following:

- **Update** button: Click this button to open the **Update Storage Policy** page and modify the storage policy. This page has the same parameters as the **Create Storage Policy** page. See Creating or Updating a Storage Policy on page 660 for more information about creating a storage policy.
- **Delete** button: Click this button to delete the storage policy. When you delete a storage policy, the categories and entities associated with the policy are governed by any other policy that is applied to the category or entity.

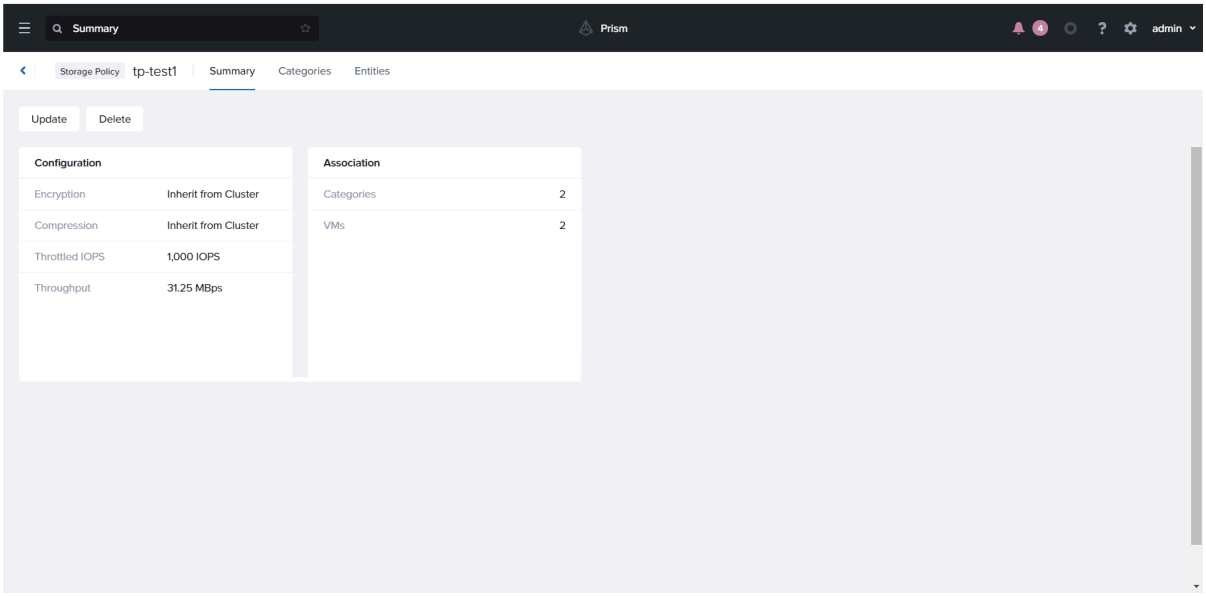

Figure 94: Storage Policy Details Summary Tab

The **Summary** tab displays information about the storage policy in following widgets:

- **Configuration**: This widget provides the following information.

Table 42: Configuration Properties

Properties	Description
Encryption	Indicates whether the storage policy involves encryption of the data. The possible values are: • **Enabled** where the encryption is enabled. • **Inherit from Cluster** where the entity inherits the encryption configuration from the cluster.

Properties	Description
Compression	Indicates whether the storage policy involves compression of the data. The possible values are: • **Inline** • **Post Process** • **Off** to disable compression. • **Inherit from Cluster** where the entity inherits the compression configuration from the container.
Throttled IOPS	Provides the throttled throughput value in IOPS.
Throughput	Provides the throttled throughput value in MBps.

- **Association**: This widget provides the following information.

Table 43: Association Properties

Properties	Description
Categories	Indicates the number of categories associated with the storage policy.
VMs	Indicates the number of entities or VMs associated with the storage policy.

Categories Tab

The Categories tab displays the following:

- **Manage Associations** button: Click this button to add or delete category associations.
- The list of categories associated with the storage policy. The following table provides the details of the categories associated.

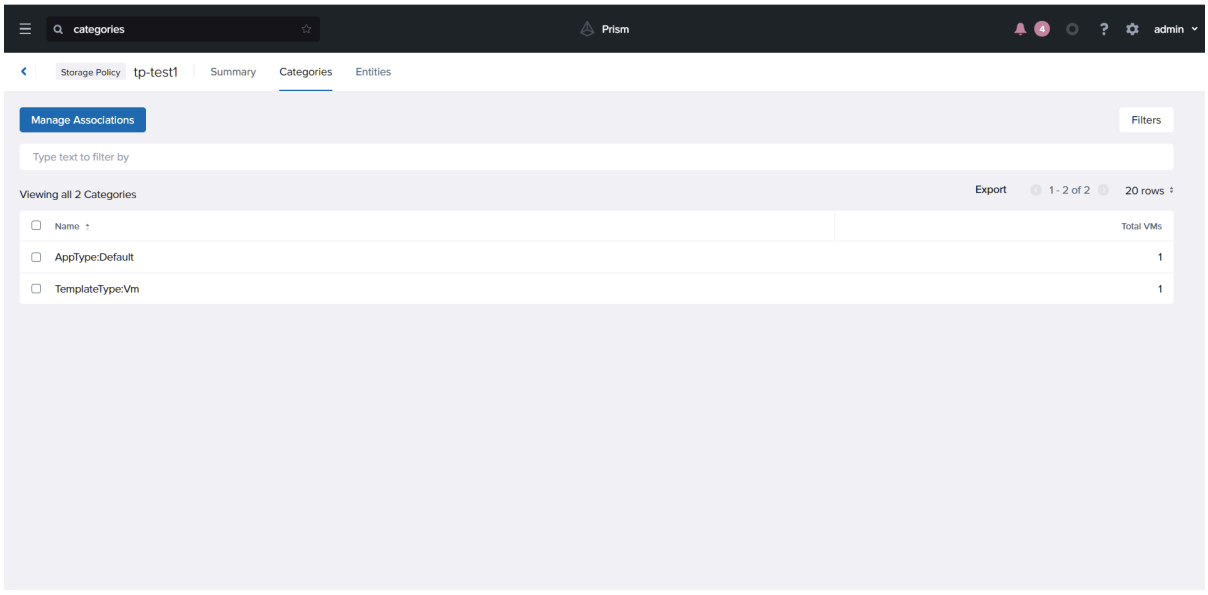

Figure 95: Storage Policy Details Categories Tab

Table 44: Categories List Fields

Properties	Description
Name	Displays the name of the category.
Total VMs	Displays the total number of VMs associated with the category.

Entities Tab

The Entities tab displays the following:

- **Manage Categories** button: Click this button to add or delete category associations. You can only associate entities to a storage policy if the entities are associated with one or more categories.

  > **Note:** Storage policies use categories to associate with entities like VMs. For more information about the requirements and limitations for storage policies, see Storage Policy Management on page 657.

- The list of entities (VMs) associated with the storage policy. The following table provides the details of the entities associated.

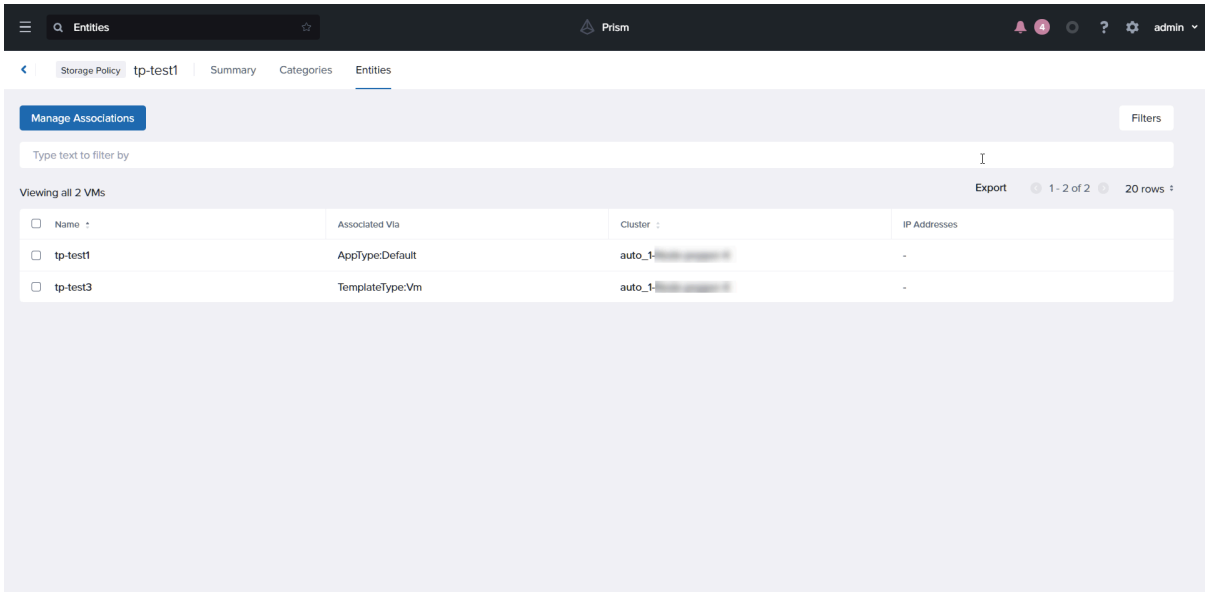

Figure 96: Storage Policies Details Entities Tab

Table 45: Entity List Fields

Properties	Description
Name	Displays the name of the category.
Associated via	Displays the categories that the entity and the storage policy are concurrently associated with. The storage policy is associated with the entity because the storage policy is associated with this category that is associated with the entity. For example, StoragePolicy1 is associated with CategoryA which is in turn associated with VM1. The category displayed in **Associated via** is CategoryA. The other categories associated with VM1 but not associated with StoragePolicy1 are not displayed here.
Cluster	Displays the name of the cluster that the entity is configured on.
IP Addresses	Displays the IP addresses configured on the entity.

Volume Groups Summary View

To access the volume groups dashboard, select **Compute & Storage** > **Volume Groups** from the entities menu (see Entities Menu on page 13). The volume groups dashboard summary view displays information about volume groups across the registered clusters and allows you to access detailed information about each volume group. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, and **Metrics**) with a display area below the selected tab.

Note:

- This section describes the information and options that appear in the volume groups dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.
- See Storage Management (Prism Central) on page 620 for information about creating and managing volume groups.

- Some volume group details and actions are unavailable for clusters running a pre-6.0 AOS version.

Summary Tab

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

- **Highlighted Entities**: Displays a list of the volume groups with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that volume group. Click the **View all XX Volume Groups** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of storage container-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.

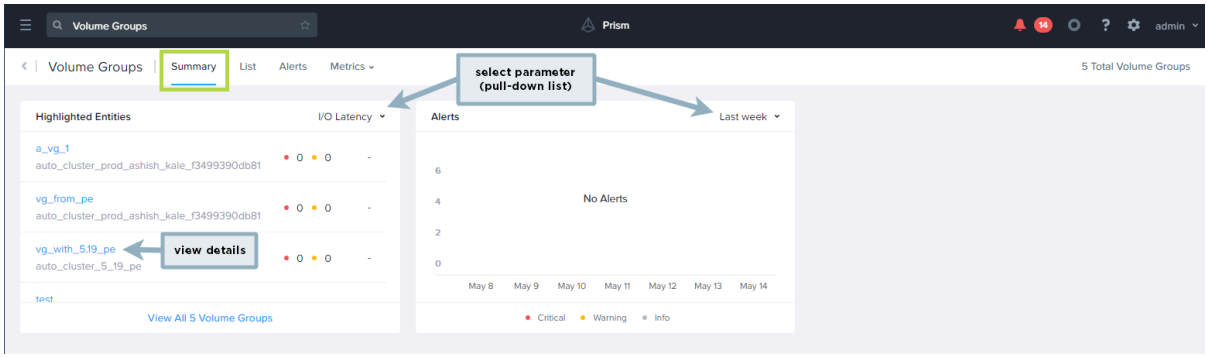

Figure 97: Volume Groups Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of volume groups across the registered clusters. The following table describes the fields that appear in the volume groups list. A dash (-) is displayed in a field when a value is not available or applicable.

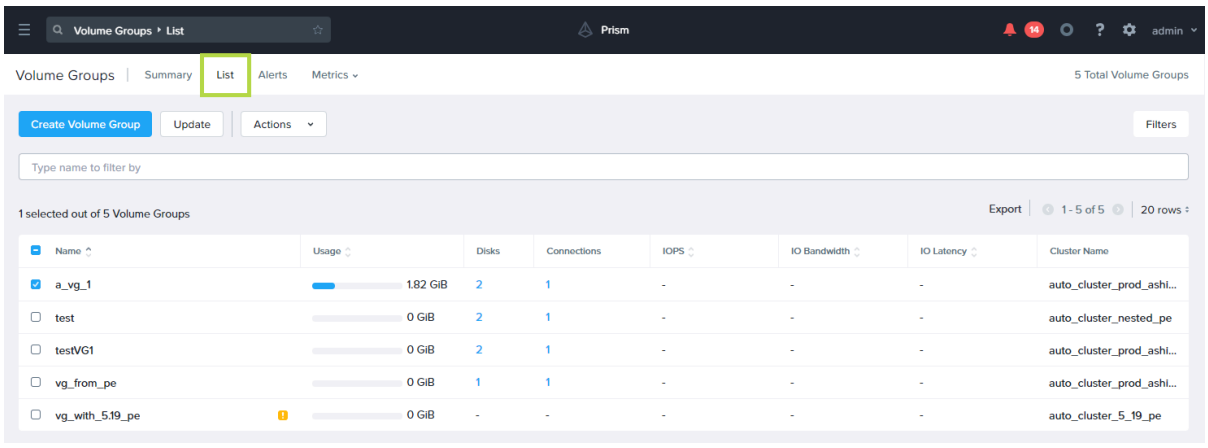

Figure 98: Volume Groups List Tab

Table 46: Volume Groups List Fields

Parameter	Description	Values
Name	Displays the name of the volume group. Click the name to display the details page for that volume group (see Volume Group Details View on page 151). An exclamation point icon appears if the cluster is running a pre-6.0 AOS version, which indicates the available information and actions are limited.	(name)
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the volume group. Hovering the cursor over the bar displays a window listing the used space, free space, and total space in the volume group.	xxx [GB|TB]
Disks	Displays the number of virtual disks in the volume group. Clicking the number displays the **Virtual Disks** tab of the details view (see Volume Group Details View on page 151).	(integer)
Connections	Displays the number of external client connections to the volume group. Clicking the number displays the **Connections** tab of the details view (see Volume Group Details View on page 151).	(integer)
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Cluster Name	Displays the name of the cluster in which the volume group resides. Click the name to display the cluster details (see Cluster Details View on page 209).	(cluster name)

You can filter the volume groups list based on various parameter values. The following table describes the filter options available when you open the volume groups view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 47: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the volume group name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of volume groups that satisfy the name condition/string.	(volume group name string)
Usage	Filters on the used storage space for a volume group. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of volume groups with used storage in that range.	([xx] to [yy] GiB range)
Cluster Name	Filters on the cluster name. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of volume groups in the clusters that satisfy the name condition/string.	(cluster name string)
IOPS	Filters on the current IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It returns a list of volume groups with IOPS in that range.	([xx] to [yy] range)
I/O Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It returns a list of volume groups with I/O bandwidth usage in that range.	([xx] to [yy] range)
I/O Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of volume groups with average I/O latency in that range.	([xx] to [yy] range)

The volume groups view includes two action buttons and an actions menu:

- To create a volume group, click the **Create Volume Group** button (see Creating a Volume Group on page 644).
- To update or delete a volume group, click the **Update** or **Delete** button (see Modifying a Volume Group on page 648).
- To manage external client connections for a volume group, select the volume group and then select **Manage Connections** from the **Actions** menu (see Managing Volume Group Connections on page 649).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just volume group-related alerts across the registered clusters (see Alerts Summary View (Prism Central) on page 269).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the volume groups. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

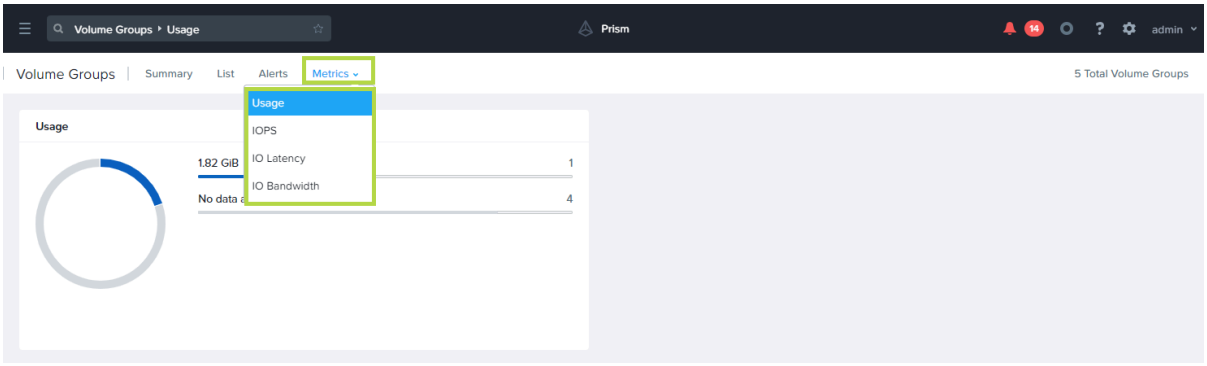

Figure 99: Volume Groups Metrics Tab

Table 48: Metrics Tab Fields

Metric	Description
Usage	Displays volume group usage graph.
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000). **Note:** The same format also applies to the other metrics in this table.
I/O Latency	Displays total, read, and write I/O latency rate graphs.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.

Volume Group Details View

To access the details page for a volume group, go to the volume groups dashboard **List** tab (see Volume Groups Summary View on page 147) and click the volume group name. You can also access the details page by clicking the volume group name wherever that name appears, such as in a dashboard widget or search result.

The volume group name and the following set of tabs appear at the top: **Summary**, **Virtual Disks**, **Connections**, **Alerts**, and **Metrics**. Click a tab to display that information. (Click the arrow at the upper left to return to the summary view.)

Summary Tab

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

- A **Properties** widget that displays summary information about the storage container (see following table).
- A **Usage and Performance** widget that displays usage data for the storage container (see following table).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last week** (default), **Last 1 Hours**, or **Last 24 hours** from the pull-down menu.

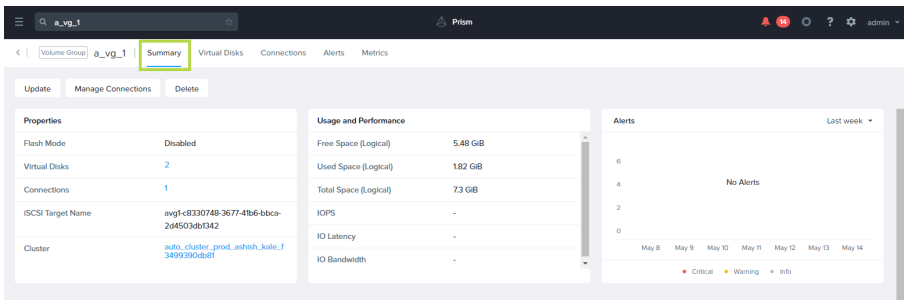

Figure 100: Volume Group Summary Tab

The following table describes the fields in the **Properties** and **Usage and Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 49: Volume Group Widget Fields

Parameter	Description	Values
"Properties" Widget		
Flash Mode	Displays whether Flash Mode is enabled for the volume group.	[Enabled, Disabled]
Virtual Disks	Displays the number of virtual disks in the volume group. Clicking the number displays the **Virtual Disks** tab.	(integer)
Connections	Displays the number of external client connections to the volume group. Clicking the number displays the **Connections** tab.	(integer)
iSCSI Target Name	Displays the name of the iSCSI target.	(ISCSI name)
Cluster	Displays the name of the cluster in which the volume group resides. Click the name to display the cluster details (see Cluster Details View on page 209).	(cluster name)
"Usage and Performance" Widget		
Free Space (Logical)	Displays the amount of logical (effective) free space available in the volume group. **Note:** Logical space accounts for the replication factor. RF1 means that the logical and physical spaces are the same, RF2 means that the logical space is half the physical space, and RF3 means that the logical space is a third of the physical space.	xxx [GB|TB]
Used Space (Logical)	Displays the amount of logical used space in the volume group.	xxx [GB|TB]
Total Space (Logical)	Displays the amount of logical total space in the volume group.	xxx [GB|TB]

Parameter	Description	Values
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
I/O Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
I/O Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps|KBps]

The **Summary** tab includes three action buttons:

- To update or delete the volume group, click the **Update** or **Delete** button (see Modifying a Volume Group on page 648).
- To manage external client connections for the volume group, click the **Manage Connections** button (see Managing Volume Group Connections on page 649).

Virtual Disks Tab

Clicking the **Virtual Disks** tab displays a list of virtual disks for the volume group. The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable.

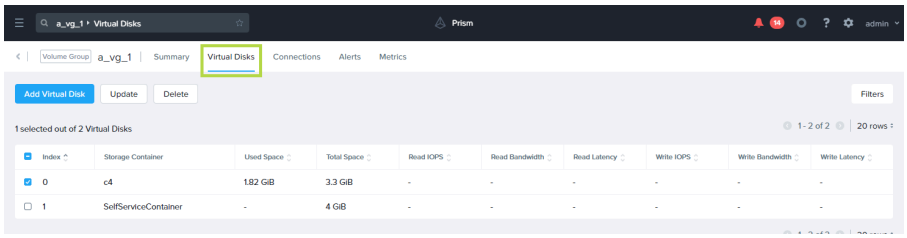

Figure 101: Volume Group Virtual Disks Tab

Table 50: Virtual Disks Fields

Parameter	Description	Values
Index	Displays the index number assigned to the virtual disk.	(integer)
Storage Container	Displays the name of the storage container in which the virtual disk is located.	(container name)
Used Space	Displays the amount of used space in the virtual disk.	xxx [GB|TB]
Total Space	Displays the total amount of storage space in the virtual disk.	xxx [GB|TB]
Read IOPS	Displays the current read I/O operations per second (IOPS) for the virtual disk.	[0 - unlimited]

Parameter	Description	Values
Read Latency	Displays the average read I/O latency for the virtual disk.	xxx [ms]
Read Bandwidth	Displays read I/O bandwidth used per second for the virtual disk.	xxx [MBps|KBps]
Write IOPS	Displays the current write IOPS for the virtual disk.	[0 - unlimited]
Write Latency	Displays the average write I/O latency for the virtual disk.	xxx [ms]
Write Bandwidth	Displays write I/O bandwidth used per second for the virtual disk.	xxx [MBps|KBps]

You can filter the virtual disks list based on various parameter values. The following table describes the filter options available when you open the virtual disks view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require more context such as a constraint string or a range.

Table 51: Filter Pane Fields

Parameter	Description	Values
storage container	Filters on the storage container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of volume groups that satisfy the name condition/string.	(storage container name string)
[Total|Used] Space	Filters on the total or used storage space for a virtual disk. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It returns a list of virtual disks with total or used storage in that range.	([xx] to [yy] GiB range)
[Read|Write] IOPS	Filters on the current read or write IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It returns a list of virtual disks with IOPS in that range.	([xx] to [yy] range)
[Read|Write] Bandwidth	Filters on the read or write I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It returns a list of virtual disks with I/O bandwidth usage in that range.	([xx] to [yy] range)
[Read|Write] Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It returns a list of virtual disks with average I/O latency in that range.	([xx] to [yy] range)

The **Virtual Disks** tab includes three action buttons:

- To add a virtual disk to the volume group, click the **Add Virtual Disk** button (see Managing Volume Group Virtual Disks on page 653).

- To update or delete a virtual disk, select the target disk and then click the **Update** or **Delete** button (see Managing Volume Group Virtual Disks on page 653).

Connections Tab

Clicking the **Connections** tab displays a list of external client connections to the volume group. The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable.

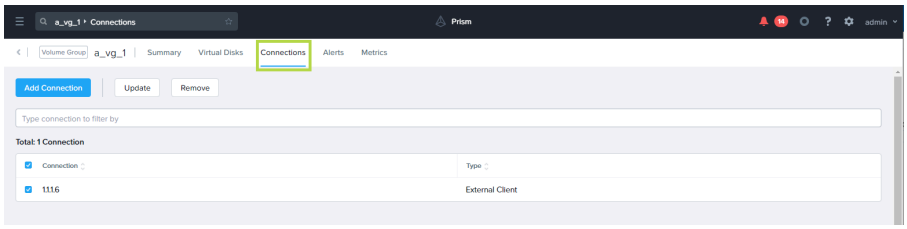

Figure 102: Volume Group Connections Tab

Table 52: Connections Fields

Parameter	Description	Values
Connection	Displays the IP address or IQN of the external client.	(IP or IQN address)
Type	Displays the type of connection.	external client

The **Connections** tab includes three action buttons:

- To add an external client connection, click the **Add Connection** button (see Managing Volume Group Connections on page 649).
- To update or remove a connection, click the **Update** or **Remove** button (see Managing Volume Group Connections on page 649).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this volume group (see Alerts Summary View (Prism Central) on page 269).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the volume group. Click the **Metrics** tab to display graphs for 10 metrics: **Usage**, **Controller IOPS**, **Controller Read IOPS**, **Controller Write IOPS Controller Average I/O Latency**, **Controller Average Read I/O Latency**, **Controller Average Write I/O Latency**, **Controller I/O Bandwidth**, **Controller Read I/O Bandwidth**, **Controller Write I/O Bandwidth**. Each graph is a rolling time interval performance or usage monitor. The baseline range (when established) appears as a blue band in the graph.

> **Note:** The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 899). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.

- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- Click the **Actions** pull-down menu (upper right of any graph) to monitor the metric performance in greater detail (see Analysis Dashboard (Prism Central) on page 799):
  - Click **Add to Analysis** to add an active monitor for the metric on the Analysis dashboard.
  - Click **Select Analysis Session** to assign the metric to a target session.
- Click the filter button to display the filter panel; check (uncheck) the desired boxes to display (hide) those metrics.
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 288).

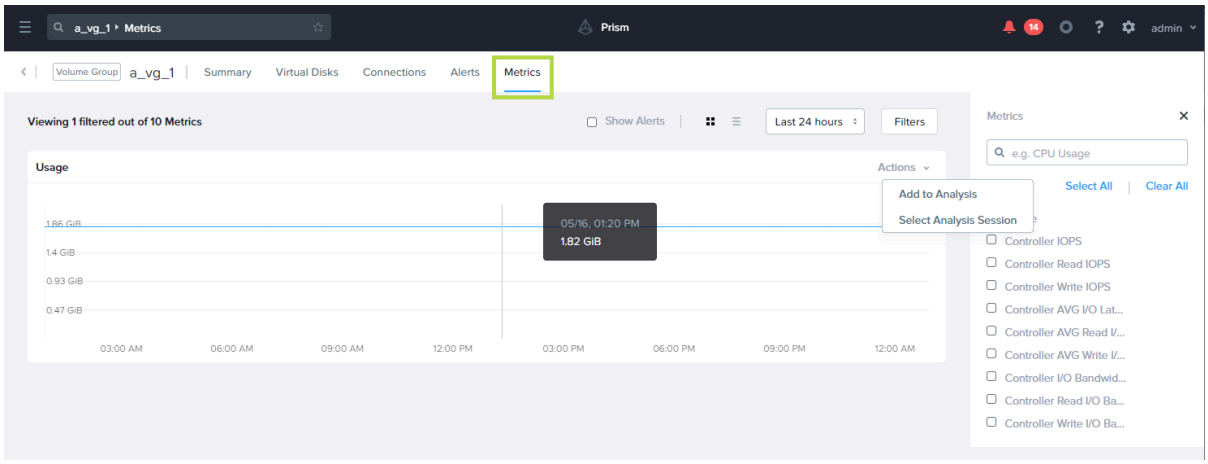

Figure 103: Volume Group Metrics Tab

vCenter Datastores Summary View

To access the vCenter datastores dashboard, select **Compute & Storage** > **vCenter Datastores** from the entities menu (see Entities Menu on page 13). The vCenter datastores dashboard summary view displays information about datastores across the registered vCenter instances and allows you to access detailed information about each datastore. The dashboard includes five tabs at the top (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below for the selected tab.

Note:

- This section describes the information and options that appear in the vCenter datastores dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
- The vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 856 for more information.

Summary Tab

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

- **Highlighted Entities**: Displays a list of the datastores with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The only option is **vSphere Datastore Numberreadaveraged**

Average. Click the **View All XX vCenter Datastores** link at the bottom to display the **List** tab (following section).

- **Alerts**: Displays a list of datastore-related alerts that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

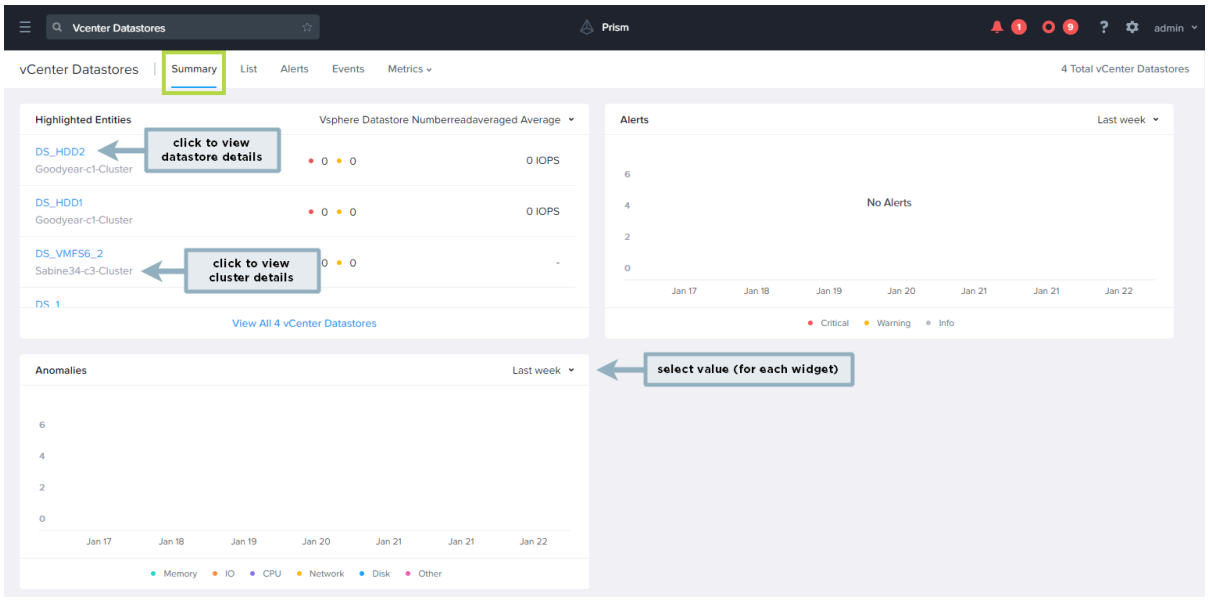

Figure 104: Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the registered vCenter datastores. The following table describes the fields that appear in the datastores list. A dash (-) is displayed in a field when a value is not available or applicable.

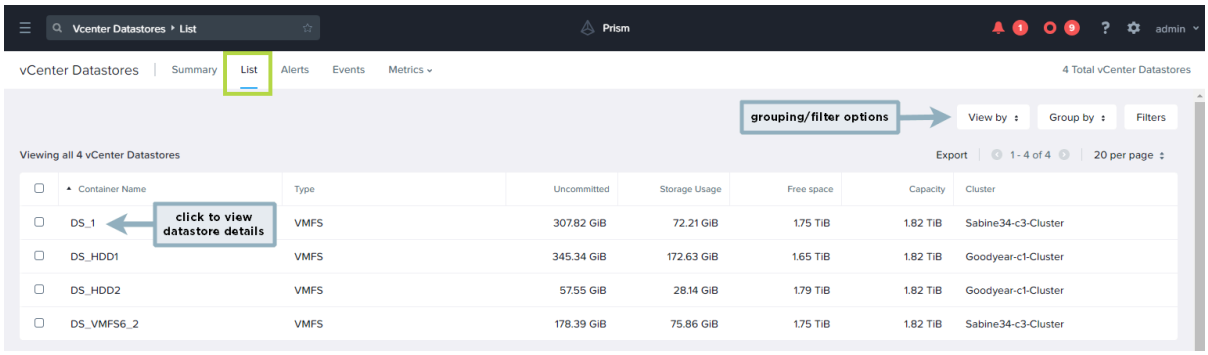

Figure 105: List Tab

Table 53: Datastores Fields

Parameter	Description	Values
Container Name	Displays the name of the datastore container. Click the name to display the details for that container (see vCenter Datastore Details View on page 160).	(container name)
Type	Displays the file system type.	VMFS, NFS
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB|TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB|TiB]
Free Space	Displays the amount of free storage space.	xx [GiB|TiB]
Capacity	Displays the total amount of storage capacity.	xx [GiB|TiB]
Cluster	Displays the name of the cluster in which the datastore resides.	(cluster name)

You can filter the datastores list based on a variety of parameter values. The following table describes the filter options available when you open the datastores view **Filters** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 54: Filter Pane Fields

Parameter	Description	Values
Container Name	Filters on the container name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of containers that satisfy the name condition/string.	(name string)
Cluster	Filters on the cluster name. Enter a string in the field to display a list of clusters that satisfy the name string.	(cluster name)
Storage Usage	Filters on the amount of used storage space. Check the box(es) for the desired range or enter a range in the "**from <low> to <high> TiB**" field. The number of containers in each range is displayed on the right of the line.	([xx] to [yy] range)
Capacity	Filters on the amount of total storage capacity.	([xx] to [yy] range)
Free Space	Filters on the amount of free storage space.	([xx] to [yy] range)
Uncommitted	Filters on the amount of uncommitted storage space.	([xx] to [yy] range)
Type	Filters on file system type.	VMFS, NFS

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

You can group the datastore containers list in the following ways:

- The **View by** pull-down menu allows you to select a view for the list, either **General** (default) or a custom view.
- The **Group by** pull-down menu allows you to group the datastore entries by cluster or type. (You can only choose one.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just datastore-related alerts across the registered clusters (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just datastore-related events across the registered clusters (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the datastores. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant performance information. The following table describes the available metrics.

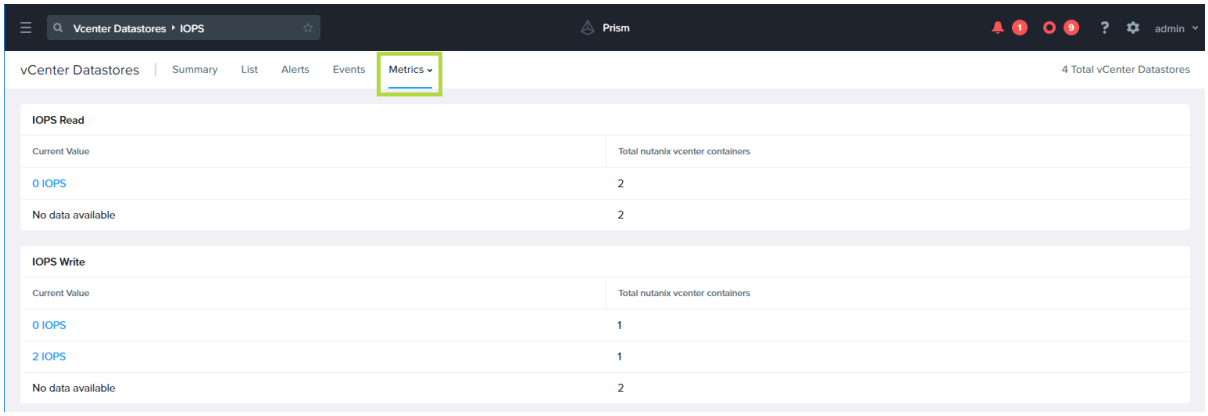

Figure 106: Metrics Tab (IOPS)

Table 55: Metrics Tab Fields

Metric	Description
IOPS	Displays the number of read and write IOPS for the listed number of containers. (Clicking an IOPS value redisplays the Summary tab filtered for those containers.)
Disk Usage	Displays the disk usage amount for the listed number of containers. (Clicking the usage value redisplays the Summary tab filtered for those containers.)
Disk Capacity	Displays the disk capacity for the listed number of containers. (Clicking the capacity value redisplays the Summary tab filtered for those containers.)

Metric	Description
Disk Provisioned	Displays the provisioned disk space for the listed number of containers. (Clicking the provisioned value redisplays the Summary tab filtered for those containers.)

vCenter Datastore Details View

To access the details page for a vCenter datastore, go to the vCenter datastores **Summary** tab (see vCenter Datastores Summary View on page 156) and click the datastore name. You can also access the details page by clicking the datastore name wherever that name appears, such as in a dashboard widget or search result.

The datastore name and the following set of tabs appear at the top: Summary, Alerts, Events, and Metrics. Click a tab to display that information below the tabs. (Click the **Back to vCenter Datastores** link to return to the datastores summary view.)

Summary Tab

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

- A **Properties** widget that displays summary information about the VM (see following table).
- A **Metrics** widget that displays average metrics for the past 15 minutes (IOPS reads and writes, and disk usage, disk capacity, and disk provisioned amounts).
- An **Alerts** widget that displays a list of related alerts that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select **Last 1 hours**, **Last 24 hours** or **Last week** (default) from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
- 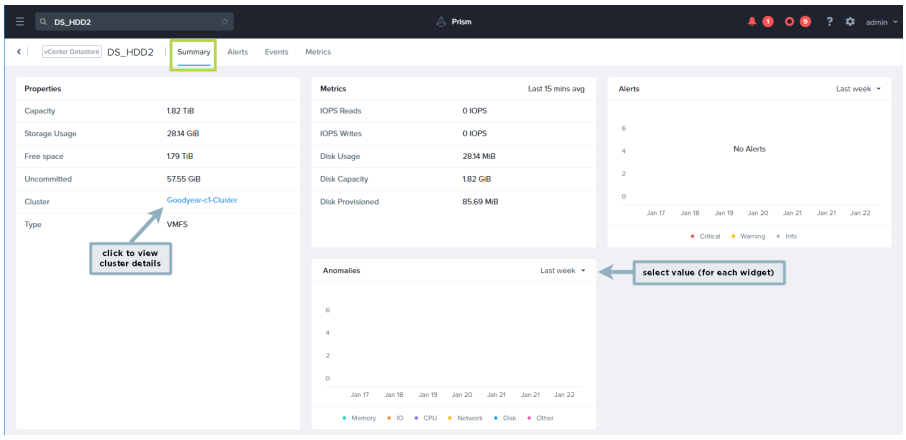

Figure 107: Summary Tab

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 56: Datastore Properties Fields

Parameter	Description	Values
Capacity	Displays the total amount of storage capacity.	xx [GiB|TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB|TiB]
Free Space	Displays the amount of free storage space.	xx [GiB|TiB]
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB|TiB]
Cluster	Displays the name of the cluster in which the datastore resides. Click on the cluster name to displays details about that cluster.	(cluster name)
Type	Displays the file system type.	VMFS, NFS

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this vCenter data store (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this vCenter data store (see Events Summary View (Prism Central) on page 279).

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.

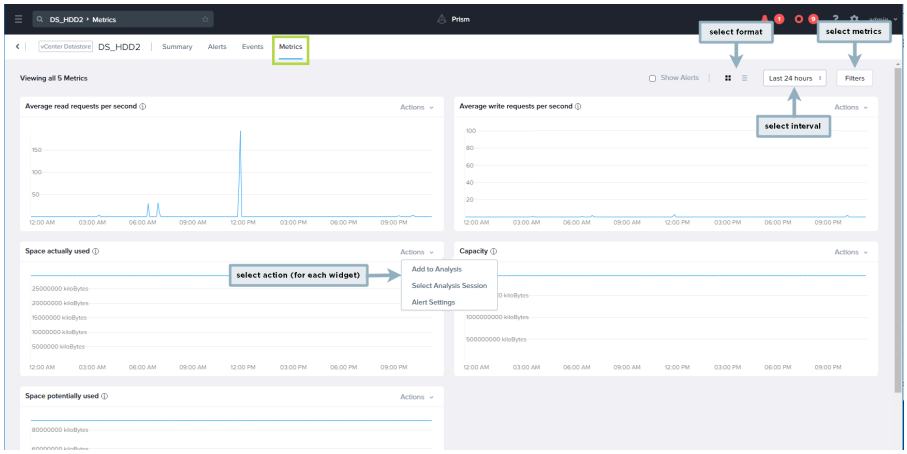

Figure 108: Metrics Tab

- Each graph includes a pull-down actions menu with some or all of the following options:
  - Select **Add to Analysis** to add the metric to the analysis page (see Analysis Dashboard (Prism Central) on page 799).
  - 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.

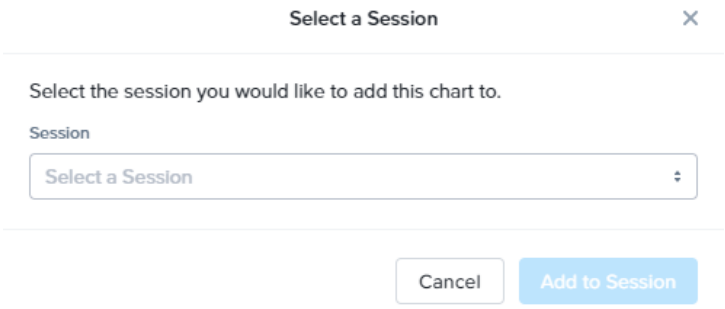

Figure 109: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric (see Creating Custom Alert Policies on page 288).
- To view a subset of the graphs, click the **Filters** button and check (uncheck) the desired metrics.
- To specify a view format for the graphs, click the tile (default) or stack icon.
- To specify a time interval for the graphs, select **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the pull-down list.
- Click the **Back to vCenter Datastores** link to return to the vCenter datastores dashboard.

The following table describes the metric graphs.

Table 57: Metric Performance Graphs

Metric	Description
Average read requests per second	Average number of read commands issued per second to the datastore during the collection interval.
Average write requests per second	Average number of write commands issued per second to the datastore during the collection interval.
Space actually used	Amount of space actually used by the virtual machine or the datastore. May be less than the amount provisioned at any given time, depending on whether the virtual machine is powered-off, whether snapshots have been created or not, and other such factors. Available from datastore and virtual machine target entities.
Capacity	Configured size of the datastore. Available from a datastore entity only.

Metric	Description
Space potentially used.	Amount of storage set aside for use by a datastore or a virtual machine. Files on the datastore and the virtual machine can expand to this size, but not beyond it. Available from datastore and virtual machine target entities.

Catalog Items View

To access the catalog items dashboard, select **Compute & Storage** > **Catalog Items** from the **Entities** menu (see Entities Menu on page 13). The catalog items dashboard allows you to view information about the images and VMs in the catalog.

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

- See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
- See Catalog Management on page 666 for information about the catalog service and how to add (or delete) items.

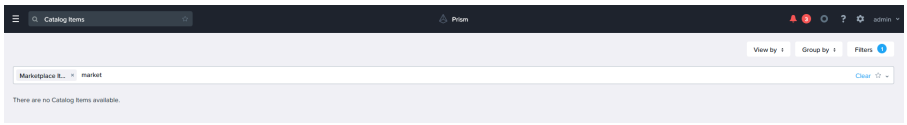

Figure 110: Catalog Items Dashboard

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

Table 58: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the item name.	(item name)
Type	Displays the type of item (image or VM).	Image, VM
Description	Indicates the description for the catalog item	(description string)

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 59: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the catalog item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It returns a list of catalog items that satisfy the name condition/string.	(image name string)
Description	Filters on the description. Select a condition from the pull-down list (same options as for name) and enter a string in the field. It returns a list of catalog items whose description field satisfy the condition/string.	(description string)
Type	Filters on the item type. Check the box(es) for the desired item types.	VM, Image

You can group the catalog items list in the following ways:

- The **Group** pull-down menu allows you to group by type (VM and image).
- The **View by** pull-down menu allows you to view by name, description, message, percent, reason, state, and type.

Categories Summary View

To access the categories dashboard, select **Compute & Storage** > **Categories** from the entities menu (see Entities Menu on page 13). The categories dashboard allows you to view summary information about existing categories and access detailed information about each category.

> **Note:** This section describes the information and options that appear in the categories dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See Category Management on page 795 for information on how to create, modify, and apply categories.

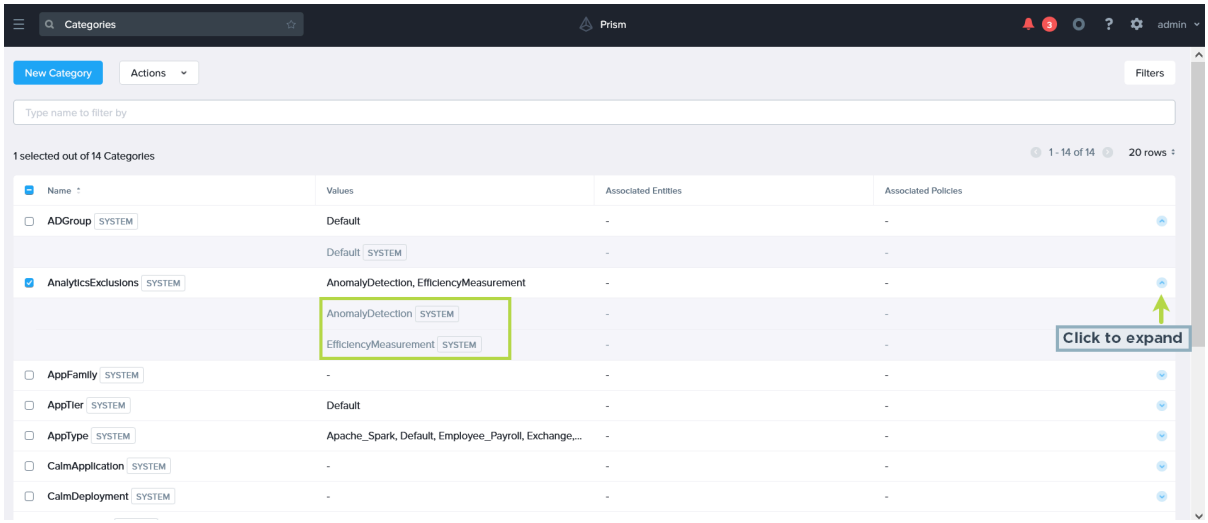

Figure 111: Categories Dashboard

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

Table 60: Categories List Fields

Parameter	Description	Values
Name	Displays the name of the category.	(category name)
Value	Displays the values defined for the category. Click **Show more** (right of line) to see a line for each value. (Click**Show fewer** to collapse the list.) You may see an icon with a "showing X of Y" message to indicate there are additional values; click the icon to display the full list.	(value names)
Assigned Entities	Lists the number of entities assigned to this category.	(number)
Assigned Policies	Lists the number of policies assigned to this category.	(number)

You can filter the category list based on several parameter values. The following table describes the options available when you click the **Filter** button, which displays the Categories view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 61: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the category name. Select a condition from the pull-down list and enter a string in the field. It will return a list of categories that satisfy the name condition/string. **Note:** In this and the Value field, the condition menu options are **Contains**, **Does not contain**, **Starts with**, **Ends with**, and **Equal to**.	(name string)
Entities	Filters on the entity type. Check the box for one or more entity types.	VM, Host, App, Blueprint, Cluster, Image, Marketplace Item, Reports, Subnet
Policies	Filters on the policy type. Check the box for one or more entity types.	Security Policy, Affinity Policy, Image Placement Policy, NGT Policy, Protection Policy, QoS Policy

There is a **New Category** action button to create a new category (see Creating a Category on page 795). The **Actions** menu appears when one or more categories are selected and includes the following options:

- This action can be applied to multiple categories: **Delete**
- This action can be applied to only one category at a time: **Update**

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the selected category. See Modifying a Category on page 796 for instructions on how to perform these actions.

Category Details View

You can view detailed policy, value, and entity associations for a category by going to the details page for that category. To access the details page for a category, go to the categories dashboard (see Categories Summary View on page 164) and click the category name.

A category details page includes the following:

- Category name (upper left). You can switch from one category to another by selecting a different category name from the pull-down list in the upper left of the screen.
- Action buttons (upper right).
  - Click the **Update** button to update the category definition (see Modifying a Category on page 796); click the **Delete** button to delete the category. The button is grayed out if that action is not allowed. For example, you cannot delete system categories.
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.
- **Policies** (left), **Values** (middle), and **Entities** (right) columns that list the values defined for the category plus the policies and entities associated with that category. Placing the cursor over a policy, value, or entity displays

lines that graphically indicate the association among the three parameters. For example, in the following figure placing the cursor over the **Forensics** value shows that it is associated with the **Quarantine-Forensics** security policy, but it is not associated with any entities currently.

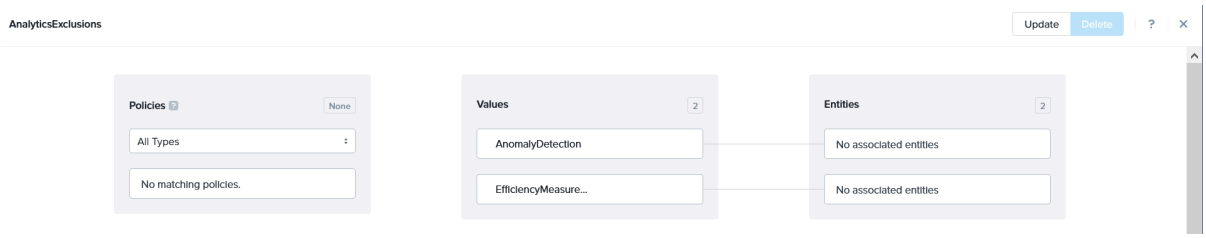

Figure 112: Category Details View

OVAs View

To access the OVAs dashboard, select **Compute & Storage** > **OVAs** from the entities menu (see Entities Menu on page 13). The OVAs dashboard allows you to view information about the OVAs available in the cluster.

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

- See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
- See OVA Management on page 698 for information about the OVA operations.

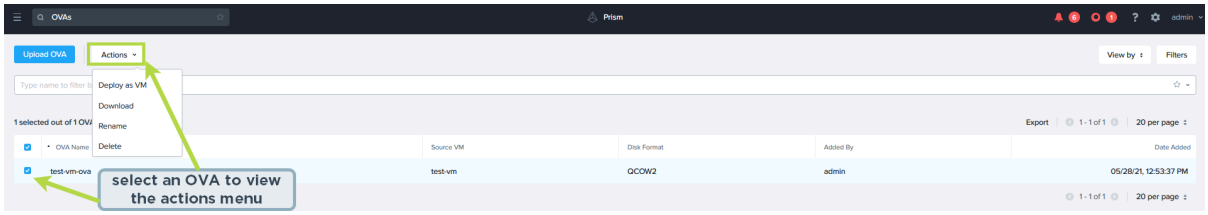

Figure 113: OVAs Dashboard

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

Table 62: OVA List Fields

Parameter	Description	Values
OVA Name	Displays the name of the OVA file.	(file name)
Source VM	Displays the name of source VM that was exported as an OVA. The alert indicators for the task status are displayed beside the source VM for which the task is run. See the Alert indicator table for more information.	(file name)
Disk Format	Displays the format of the disks in the OVA.	[QCOW2|VMDK]

Parameter	Description	Values
Added by	Indicates which user added this OVA.	(text string)
Date added	Displays the date and time when the OVA was added.	(date in mm/dd/yyyy format, time in hr:min:sec [AM|PM] format) Example: 07/30/20, 3:00:20 PM

Table 63: OVA Alerts

OVA Alerts	Status Message sample	Indicates...
 Figure 114:	`Incorrect file uploaded.` `Multiple OVA files present in OVA.`	That the upload process failed during validation. You must fix the issue in the OVA file or upload the correct, compatible OVA file.

Figure 115:
	`Upload for OVA file has not started.`
	`Please start OVA file upload using Resume Upload option.`
	`Uploaded OVA file has not been verified.`
	`Please start verification using Resume Upload option.`
		That there is an interruption in one of the three sub-tasks in the upload process.
		You can resume the upload. See the *Resume Upload* action procedure.

 Figure 116:	`OVA file upload is in progress.`	That the OVA upload is in progress without any error or interrupts. Check the Tasks page after some time to verify the successful upload.

A sample of the error hover message is as follows:

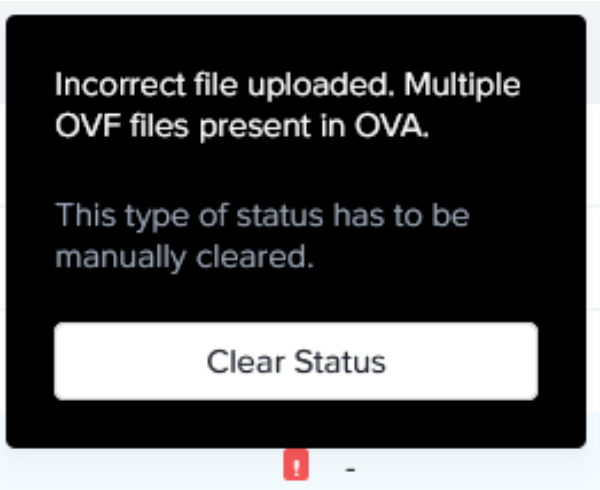

Figure 117: Sample error hover message

You can filter the catalog items list based on several parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

For each filter, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.

Table 64: Filter Pane Fields

Parameter	Description	Values
OVA Name	Filters on the OVA file name. It returns a list of OVAs that satisfy the name condition/string.	(OVA file name string)
Source VM	Filters on the Source VM name. It returns a list of OVAs for which the Source VM field satisfies the condition/string.	(Source VM name string)
Disk Format	Filters on the disk format. Check the box(es) for the OVAs of the required disk format.	VMDK, QCOW2

The **Actions** menu appears when a OVA is selected. It allows you to:

- Deploy an OVA file as a VM.
- Download an OVA file to your local machine.
- Rename an OVA file.
- Delete an OVA file.
- Resume upload when an upload is interrupted and shows an error.

> **Note:** The resume upload action is not available in case the concatenate process of the upload is interrupted. In such a case, you need to run the Concatenate API from the **REST API Explorer** to resume the upload.

Network and Security View

The **Network and Security** category in the **Entities Menu** expands on-click to display the following networking and security entities that are configured for the registered clusters:

1. **Subnets**: This dashboard displays the subnets and the operations you can perform on subnets.
2. **Virtual Private Clouds**: This dashboard displays the VPCs and the operations you can perform on VPCs.
3. **Floating IPs**: This dashboard displays a list of floating IP addresses that you are using in the network. It allows you to request for floating IP addresses from the free pool of I addresses available to the clusters managed by the Prism Central instance.
4. **Connectivity**: This dashboard allows you to manage the following networking capabilities:

    - **Gateways**: This tab provides a list of network Gateways you have created and configured, and the operations you can perform on the network Gateways. You can check and upgrade the Gateway bundle in **Administration** > **LCM** > **Inventory**.
    - **VPN Connections**: This tab provides a list of VPN connections you have created and configured, and the operations you can perform on VPN connections.
    - **Subnet Extensions**: This tab provides a list of subnets that you have extended at the Layer 2 level using VPN (point-to-point over Nutanix VPN) or VTEP (point-to-multi-point including third party).
5. **Security Policies**: This dashboard provides a list of security policies you configured using Flow Segmentation. For more information about Security Policies, see the Flow Microsegmentation Guide.

See Network Connections section for information on how to configure network connections.

Subnets (Overlay IP subnets), Virtual private clouds, floating IPs, and Connectivity are Flow virtual networking features. These features support flexible app-driven networking that focuses on VMs and applications instead of virtual LANs and network addresses. Flow virtual networking powers network virtualization to offer a seamless network experience with enhanced security. It is disabled by default. It is a software-defined network virtualization solution providing overlay capabilities for the on-premises AHV clusters.

Security policies drives the Flow Segmentation features for secure communications. See Flow Microsegmentation Guide.

Subnets

Manage subnets in the **List** view of **`Subnets`** dashboard in the **`Network and Security`** section.

To access the **Subnets** dashboard, select **Subnets** from the entities menu in Prism Central. The **Subnets** dashboard allows you to view information about the subnets configured for the registered clusters.

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

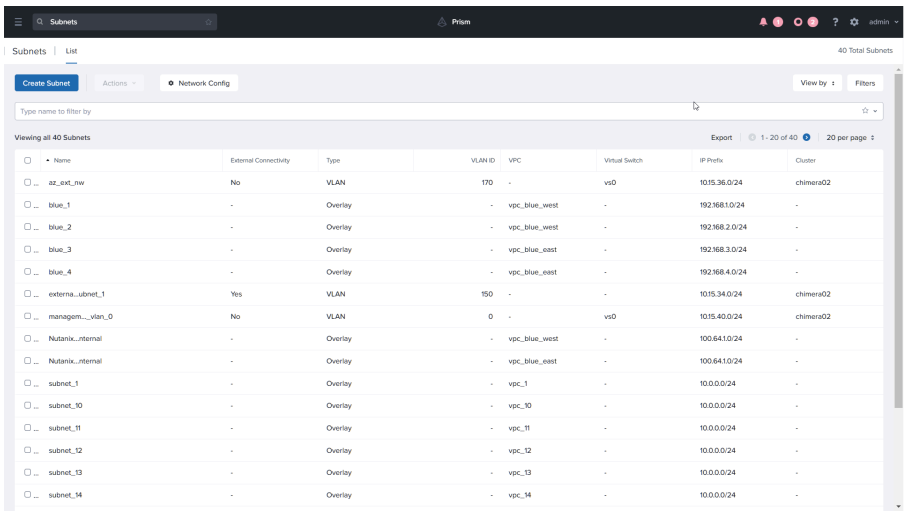

Figure 118: Subnets Dashboard

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

Table 65: Subnets Dashboard Fields

Parameter	Description	Values
Name	Displays the subnet name.	(subnet name)
External Connectivity	Displays whether or not the subnet has external connectivity configured.	(Yes/No)
Type	Displays the subnet type.	VLAN
VLAN ID	Displays the VLAN identification number.	(ID number)
VPC	Displays the name of the VPC that the Subnet is used in.	(Name of VPC)
Virtual Switch	Displays the virtual switch that is configured for the VLAN you selected. The default value is the default virtual switch `vs0`. **Note:** The virtual switch name is displayed only if you add a VLAN ID in the VLAN ID field.	(virtual switch name)
IP Prefix	Displays the IPv4 Address of the network with the prefix.	(IPv4 Address/Prefix)
Cluster	Displays the name of the cluster for which this subnet is configured.	(cluster name)
Hypervisor	Displays the hypervisor that the subnet is hosted on.	(Hypervisor)

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

To view focused fields in the List, select the focus parameter from the **Focus** drop down list. You can create your own customised focus parameters by selecting **Add custom** from the drop down list and selecting the necessary fields after providing a **Name**, in the **Subnet Columns**.

There is a **Network Config** action button to configure a new network (see Configuring Network Connections)

The **Actions** menu appears when one or more networks are selected and includes a **Manage Categories** option (see Assigning a Category).

Go to the **Subnets** list view by clicking **Network and Security** > **Subnets** on the left side-bar.

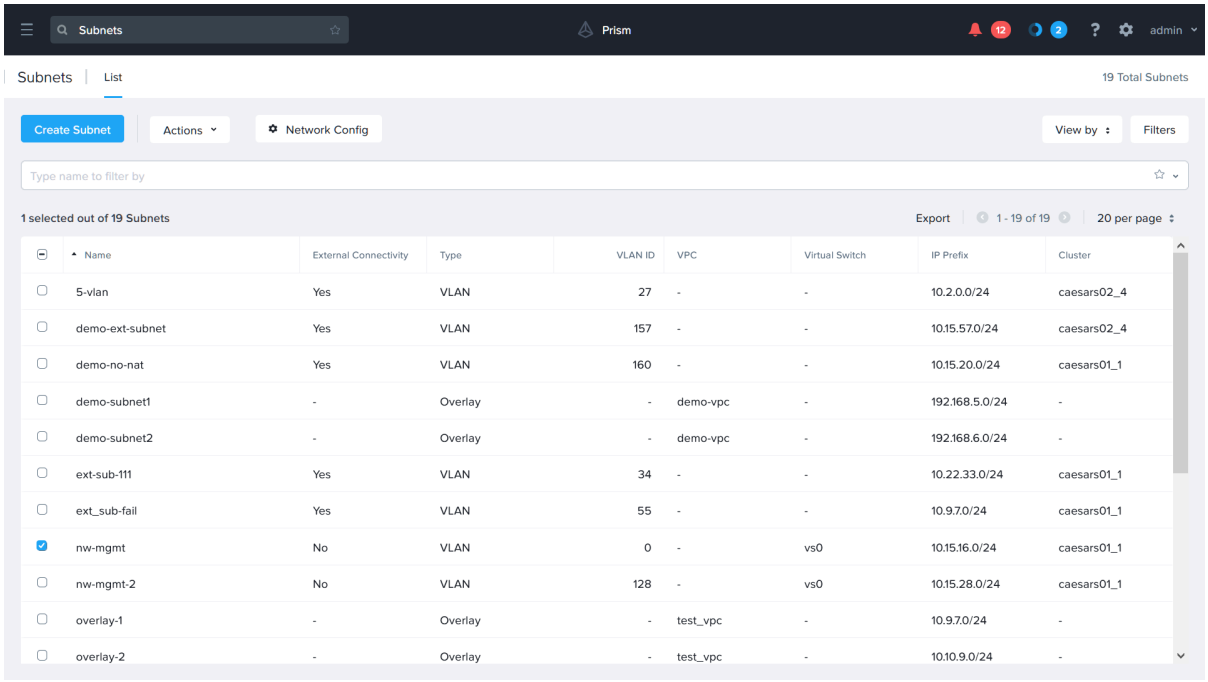

Figure 119: Subnets Page

To view or select actions you can perform on a subnet, select the subnet and click the **Actions** dropdown.

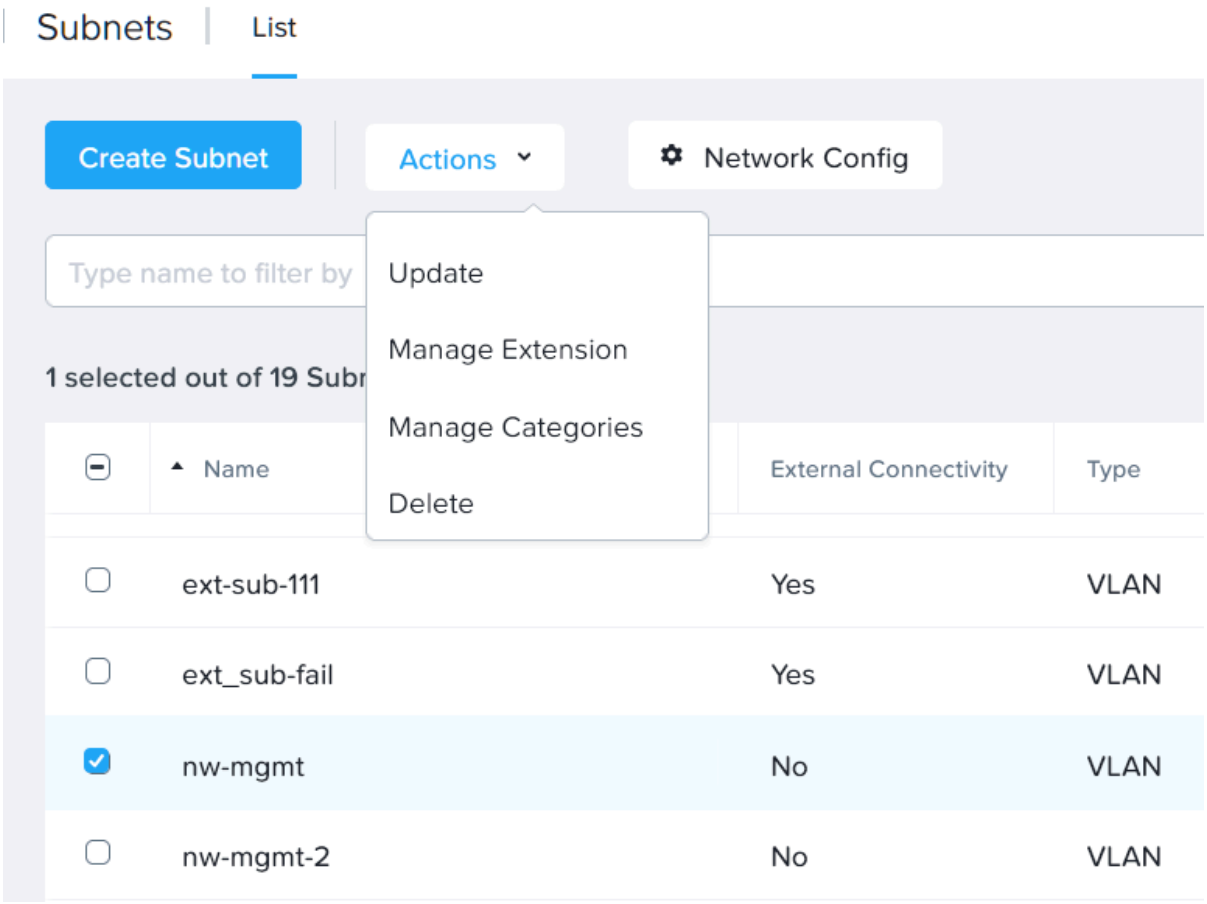

Figure 120: Subnet Actions

Table 66: Subnet Actions

Action	Description
Update	Click this action to update the selected subnet. see Updating a Subnet in the Flow Virtual Networking Guide.
Manage Extension	Click this action to create a subnet extension. A subnet extension allows VMs to communicate over the same broadcast domain to a remote Xi availability zone (in case of Xi-Leap based disaster recovery) via the extension.
Manage Categories	Click this action to associate the subnet with a category or change the categories that the subnet is associated with.
Delete	Click this action to delete the selected subnet. See Deleting Subnets, Policies, or Routes in the Flow Virtual Networking Guide.

You can also filter the list of subnets by clicking the **Filters** option and selecting the filtering parameters.

Subnet Summary View

View the details of a subnet listed on the Subnets page.

To view the details of a subnet, click the name of the subnet on the subnet list view.

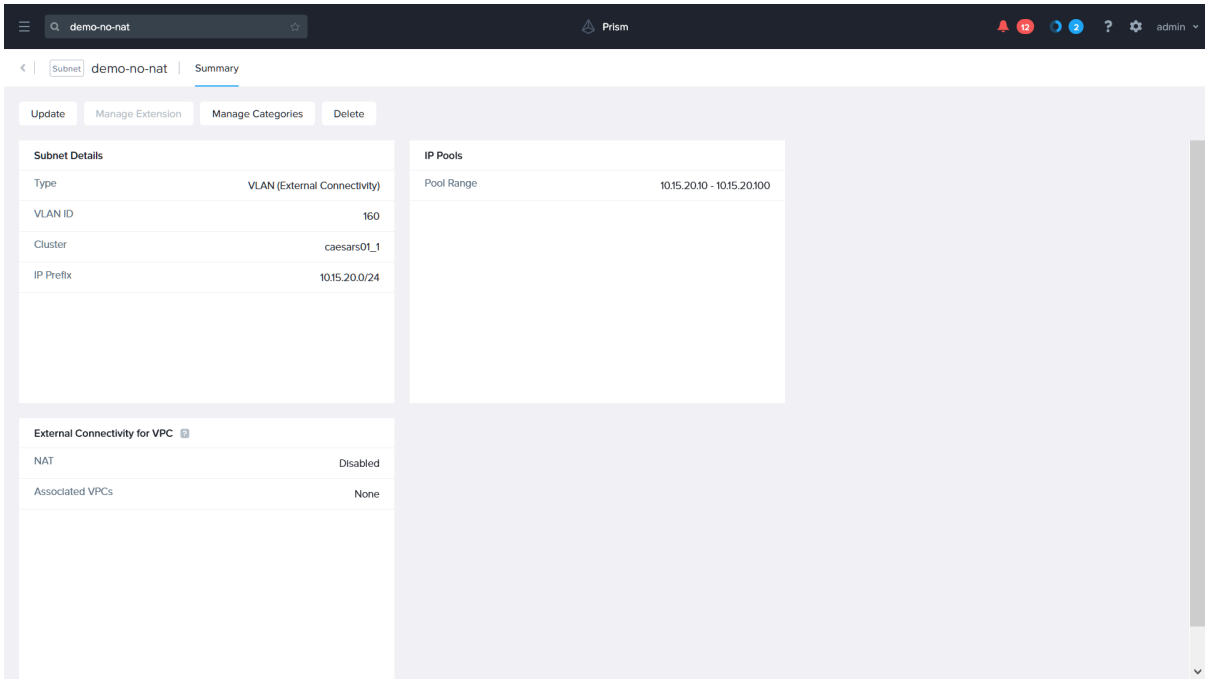

Figure 121: Subnet Summary Page

The **Summary** page provides buttons for the actions you can perform on the subnet, at the top of the page. Buttons for the following actions are available: **Update**, **Extend**, **Manage Categories**, and **Delete**.

The subnet **Summary** page has the following widgets:

Widget Name	Information provided
Subnet Details	Provides the following: • Type — Displays the type of network like VLAN or Overlay. • VLAN ID — Displays the VLAN ID. This parameter is displayed only for VLAN networks. • VPC — Displays the VPC name. This parameter is displayed only for Overlay networks. • Cluster — Displays the cluster that the VLAN network is configured on. This parameter is displayed only for VLAN networks. • IP Prefix — Displays the IP address prefix configured for the network. This parameter is displayed for both VLAN and Overlay networks.
IP Pool	Provides the IP address Pool Range assigned to the network.
External Connectivity	Provides the following: • NAT — Displays whether NAT is enabled or disabled for VPCs connecting to the network. When you hover on the Enabled/Disabled status, the hover message displays details of VPCs connected to the external subnet. • Associated VPCs — Displays the VPCs associated with this external subnet.

Virtual Private Clouds

You can manage Virtual Private Clouds (VPCs) on the **`Virtual Private Clouds`** dashboard.

Go to the **Virtual Private Clouds** dashboard by clicking **Network and Security** > **Virtual Private Clouds** on the left side-bar.

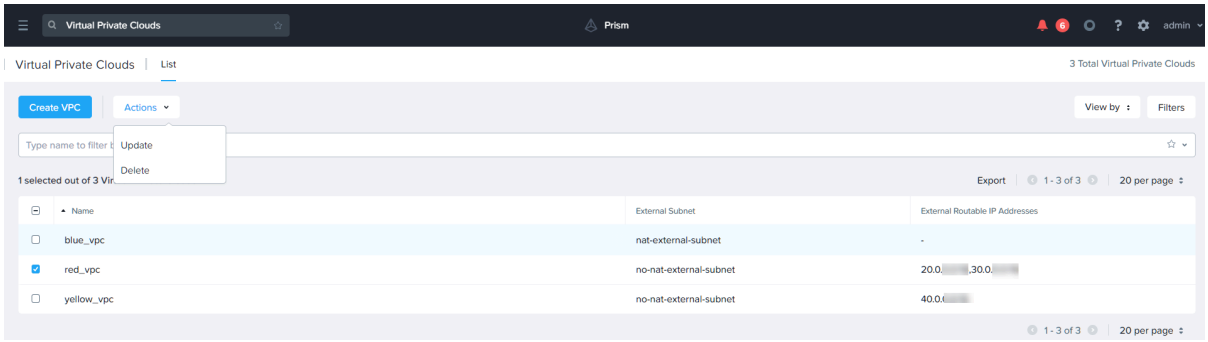

Figure 122: Virtual Private Clouds dashboard

You can configure the table columns for the VPC list table. The available column list includes Externally Routable IP Addresses that provides address space within the VPC that is reachable externally without NAT.. For the list of columns that you can add to the list table, see Customizing the VPC List View on page 175.

> **Note:**
>
> Ensure that the externally routable IP addresses (subnets with external connectivity without NAT) for different VPCs do not overlap.
>
> Configure the routes for the external connectivity subnets with next hop as the Router or SNAT IP address. Also configure the routes on the router for the return traffic to reach the VPC. See **External Connectivity** panel in VPC Details View on page 177.

To view or select actions you can perform on a VPC, select the VPC and click the **Actions** drop down.

You can also filter the list of VPC by clicking the **Filters** option and selecting the filtering parameters.

Customizing the VPC List View

About this task

You can customize the columns in the table. Click the **View by** drop down and select **+ Add custom**.

In the **Virtual Network Columns** dialog box, do the following.

Procedure

1. Enter a name for the view.

2. Select the columns you want displayed in the table.

   During the column selection, the columns you select are moved under the Selected Columns list. The ***Name*** (of the VPC) column is the default column already selected. You can add a maximum of 10 columns (including the ***Name*** column) to the Selected Column list.

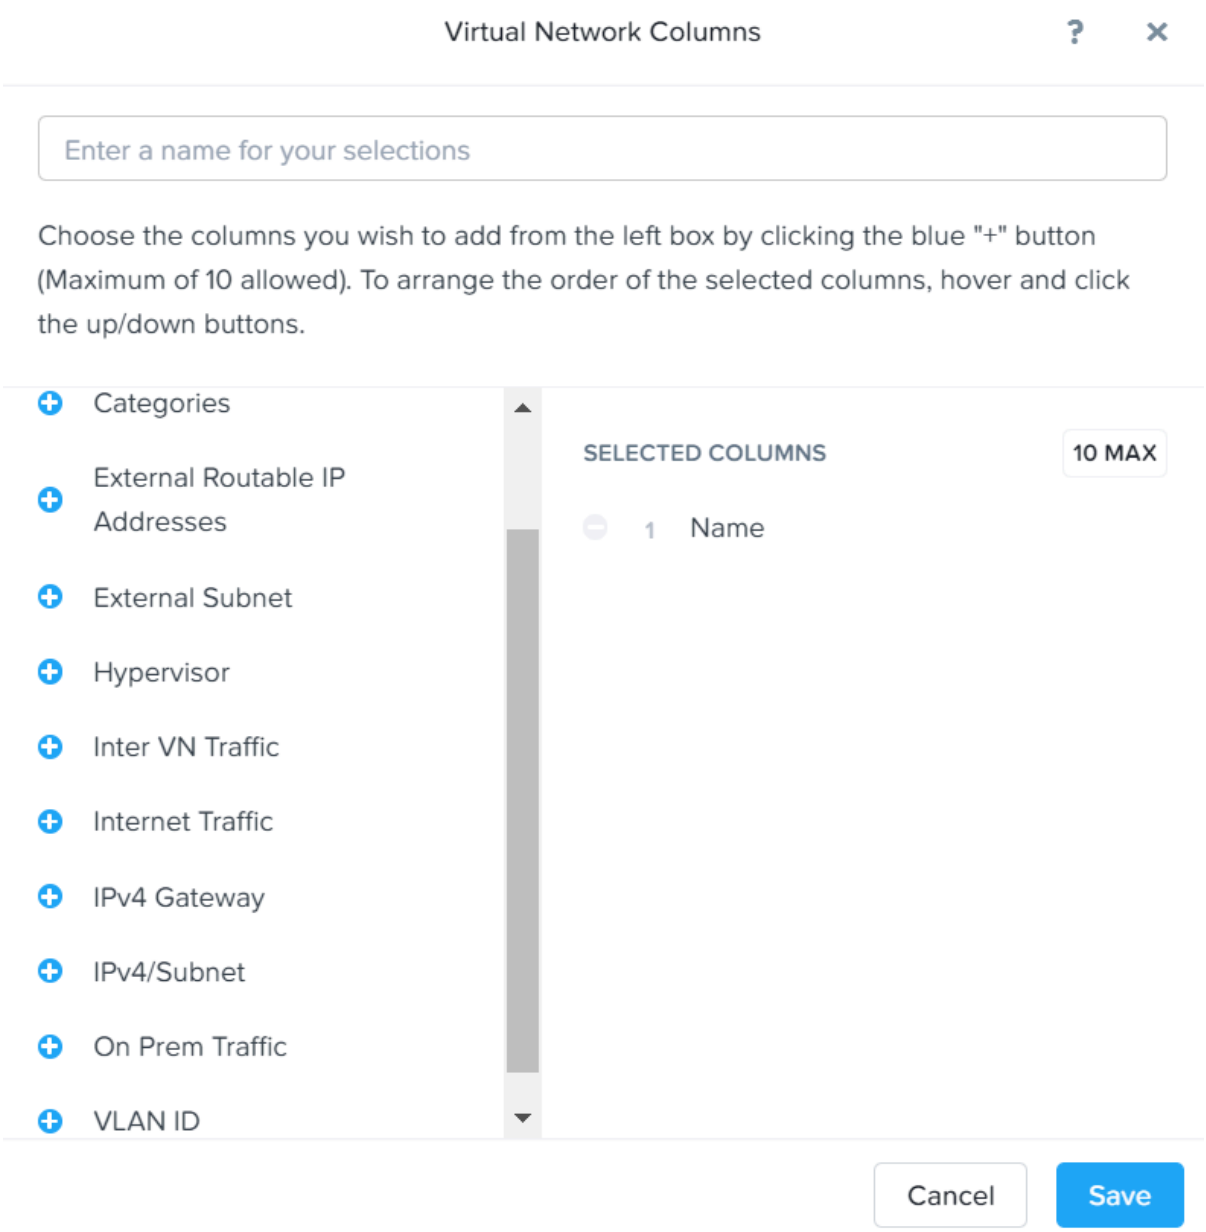

Figure 123: Customizing Columns in VPC View

   To arrange the order of the selected columns, hover on the column name and click the up or down arrow button as appropriate.

3. Click **Save**.

VPC Details View

To view the details of a VPC, click the name of the VPC on the VPC list view.

The VPC details view has the following tabs:

- **Summary**

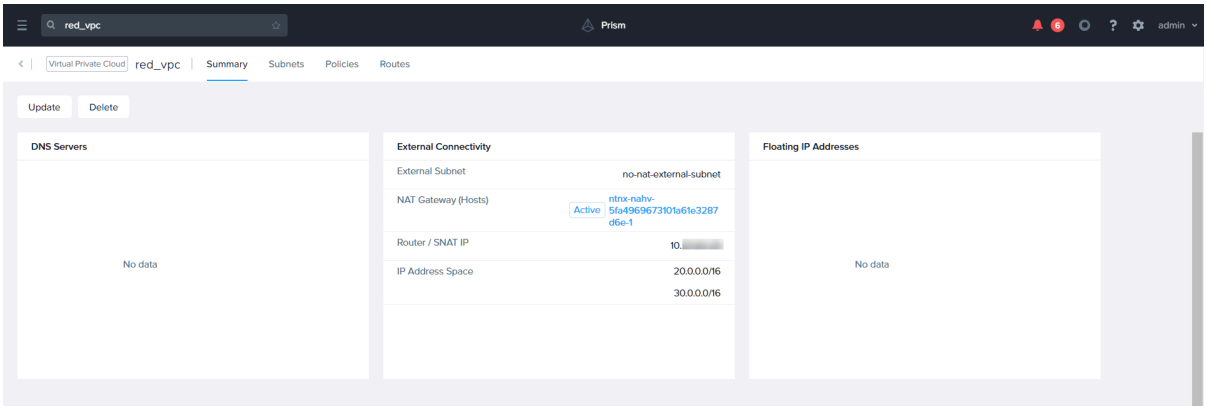

Figure 124: Summary Tab

The **Summary** tab provides the following panes:

- DNS Servers—Provides more information about the DNS Servers used by the VPC.
- External Connectivity—Provides the name of the external subnet, NAT Gateway host details, router/SNAT IP address and the IP address spaces or ranges configured for the VPC.
- Floating IP Addresses—Provides details of the floating IP addresses that the VPC uses.

- **Subnets**

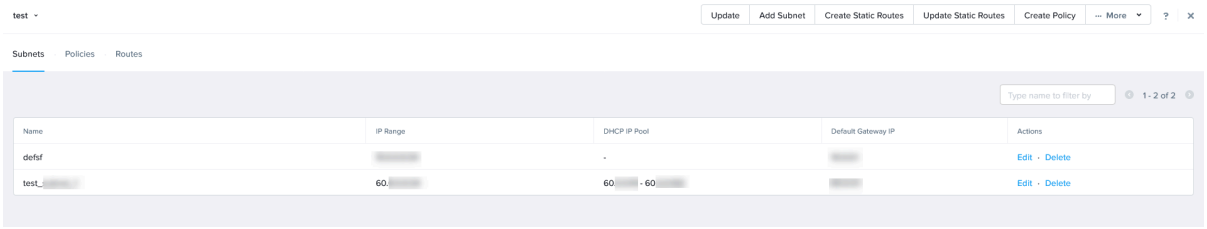

Figure 125: Subnet Tab

The **Subnet** tab provides the following information for the subnets:

- Name—Displays the name of the subnet.
- IP Range—Displays the IP address range configured for the subnet.
- DHCP IP Pool—Displays the DHCP IP address pool configured for the subnet.
- Default Gateway IP—Displays the IP address used as the default gateway by the entities in the subnet.
- Actions—Displays the actionable links to **Edit** or **Delete** the subnet.

- **Policies**

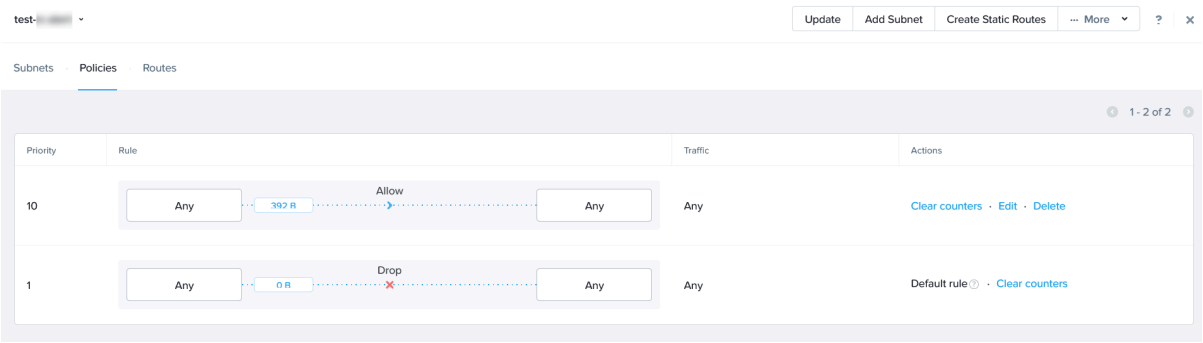

Figure 126: Policies Tab

The **Policies** tab maps the following information about the security-based traffic shaping policies you configure:

- Priority—The traffic priority.
- Rule—The **Allow** or **Deny** rule set for the priority.
- Traffic—The traffic type that the priority and rule should be applied to.
- Actions—Actions you can take on the policy. You can perform three actions: **Clear counters**, **Edit** the policy or **Delete** the policy.

- **Routes**

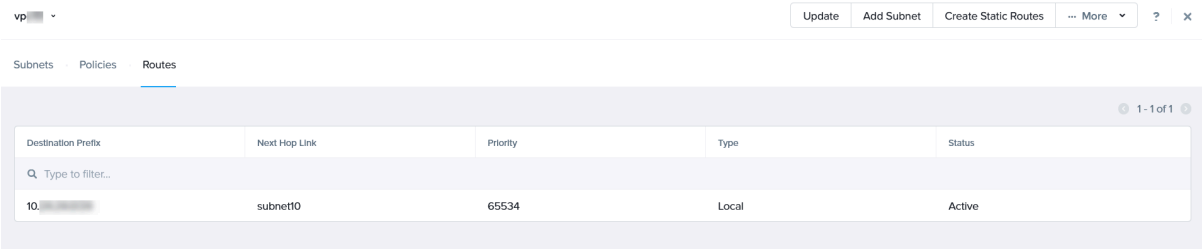

Figure 127: Routes Tab

The **Routes** tab provides the following information about the routes:

The VPC details view has the following configuration options for the VPC:

- **Update**: Use this option to update the VPC. For more information, see Updating Virtual Private Cloud.
- **Add Subnet**: Use this option to add a subnet to the VPC. For more information, see Creating a Subnet.
- **Create Static Routes**: Use this option to create a static route. For more information, Creating Static Routes.
- **Update Static Routes**: Use this option to update static route configurations that you already created. For more information, see Updating Static Routes.
- **Create Policy**: Use this option to create traffic policies in addition to the pre-configured default policy. When you create a VPC, there is one default policy that Advanced Networking creates for the VPC. This policy is pre-configured and cannot be edited. For more information, see Creating a Policy.
- **Clear All Counters**: Allows you to clear all the counters for the VPC.

- **Delete**: Allows you to delete the VPC. For more information, see Deleting a Virtual Private Cloud.

Floating IPs

You can access floating IPs on the **`Floating IPs`** dashboard or list view in the **`Network and Security`** section.

For information about floating IP addresses and their role in Flow virtual networking, see SNAT and Floating IP Address.

Go to the **Floating IPs** dashboard by clicking **Network and Security** > **Floating IPs** on the left side-bar.

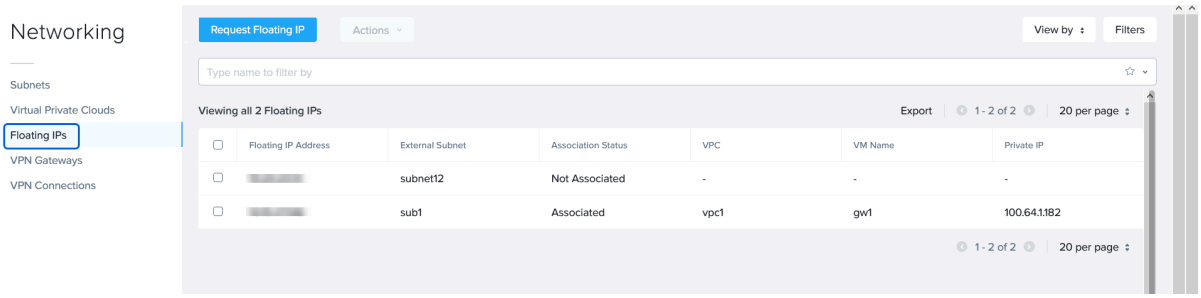

Figure 128: Floating IPs dashboard

To view or select actions you can perform on a floating IP address assigned, select the floating IP address and click the **Actions** drop down. The following actions are available for a selected floating IP address:

- Update—Assign or change the assignment of the floating IP address. You can assign the floating IP address to a IP address such as a private IP address in a VPC or the primary IP address of a VM or a secondary IP address created on a VM.
- Delete—Delete the floating IP address. The deleted IP address returns to the IP address pool as unused. Before you delete, ensure that it is not assigned to a private IP address or a VM. Change the assignment to None if it is already assigned, using the **Update** action.

> **Note:** Floating IP addresses are not reachable (Pings fail) unless you associate them to primary or secondary IP addresses of VMs. For more information about assigning floating IP addresses to secondary IP addresses of VMs, see Assigning Secondary IP Addresses to Floating IPs.

To filter the list of floating IP address assignments, click the **Filters** option and select the appropriate filtering parameters.

To request floating IP addresses, see Requesting Floating IPs.

Connectivity

You can access network Gateways, VPN connections and subnet extensions on the **`Connectivity`** dashboard.

Click **Network & Security** > **Connectivity** to see the **Connectivity** dashboard.

The **Connectivity** dashboard opens on the **Gateways** tab. To see the VPN connections, click the **VPN Connections** tab. To see the subnets extended across AZs, click the **Subnet Extensions** tab.

Gateways Summary View

The **Connectivity** dashboard opens on the **Gateways** dashboard or summary view.

The **Gateway** dashboard provides a list of gateways created for the clusters managed by the Prism Central.

The **Gateways** dashboard provides a **Create Gateway** dropdown menu that lets you create a **Local** or a **Remote** gateway. You can create a local or remote gateway with **VPN** or **VTEP** service. For more information, see Creating a Network Gateway.

You can select a gateway from the list (select the checkbox provided for the gateway) and then perform an action provided in the **Actions** dropdown list. The **Actions** dropdown list allows you to **Update** or **Delete** the selected gateway.

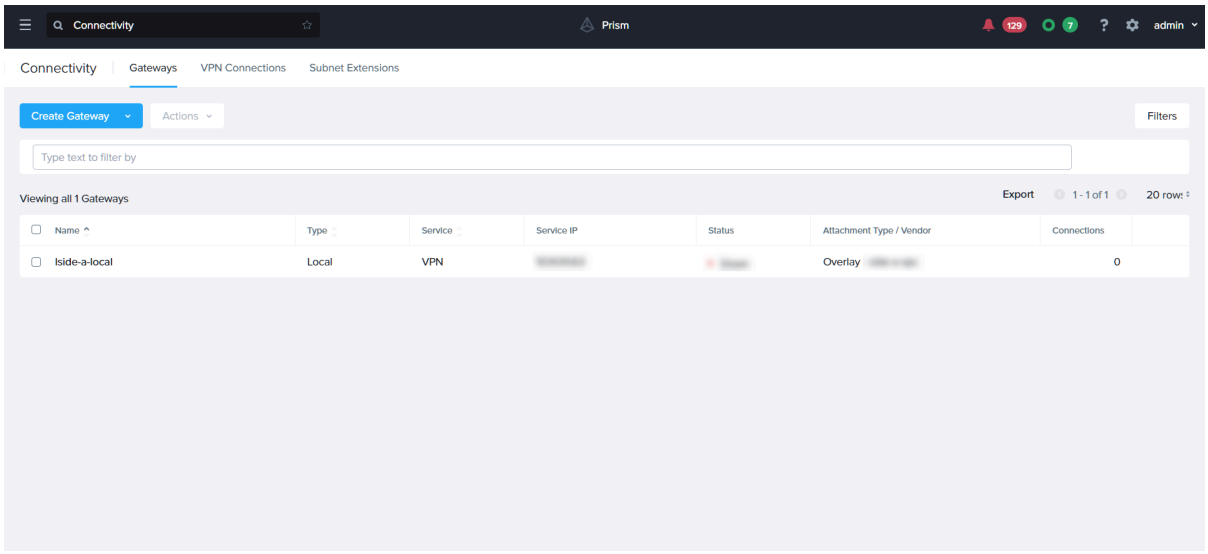

Figure 129: Gateways dashboard

The **Gateway** summary list view provides the following details about the gateway.

Table 67: Gateway List Fields

Parameter	Description	Values
Name	Displays the name of the gateway.	(Name of gateway)
Type	Displays the gateway type.	(Local or Remote)
Service	Displays the service that the gateway uses.	(VPN or VTEP)
Service IP	Displays the IP address used by the service.	(IP address)
Status	Displays the operational status of the gateway.	(Up or Down)
Attachment Type/Vendor	Displays the type of subnet associated with the gateway.	(VLAN or Overlay-VPC name)
Connections	Displays the number of service connections (such as VPN connections) configured and operational on the gateway.	(number)

You can click the name of a gateway to open the gateway details page that presents the information about the gateway in widgets.

Gateway Details View

You can click the name of a gateway in the **Gateway** dashboard list to open the gateway details page that presents the information about the gateway in widgets.

The gateway details page displays the name of the gateway on the top left corner.

- On the top right corner, the close button (X) allows you to close the details page.
- The **Update** button opens the **Update Gateway** page. For more information, see Updating a Network Gateway in *Flow Networking Guide*.
- The **Delete** button allows you to delete the gateway. For more information, see Deleting a Network Gateway in *Flow Networking Guide*.

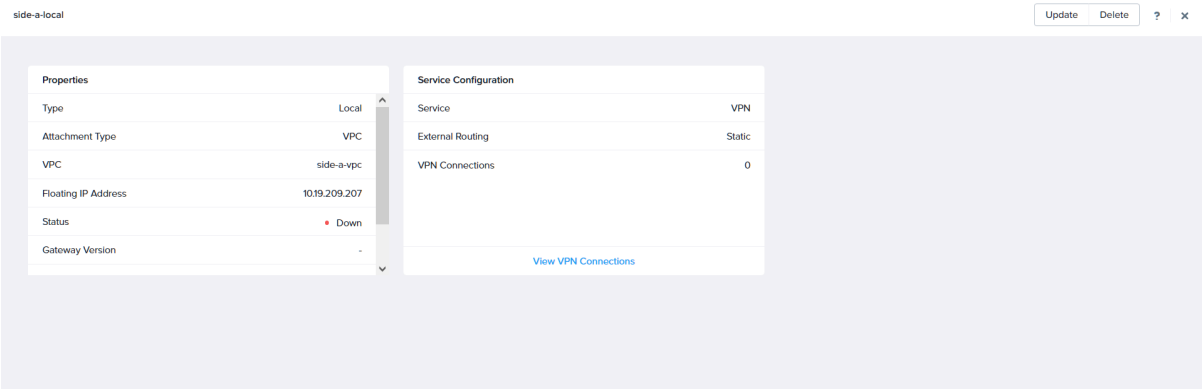

Figure 130: Gateway Details View

The details about the gateway are organized in widgets as follows:

Table 68: Gateway Details

Parameter	Description	Values
Properties widget		
Type	Displays the gateway type.	(Local or Remote)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address)
Status	Displays the operational status of the gateway.	(Up or Down)
Gateway Version	Displays the version of the Nutanix gateway appliance deployed.	(Version)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)

Parameter	Description	Values
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Service Configuration		
Service	Displays the service used by the gateway.	(VPN or VTEP)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
VPN Connections	Displays the total number of VPN connections associated with the gateway.	(Number - actionable link. Click the link to open the VPN connection details page for the associated VPN connection.)
View VPN Connections	Click this link to open the VPN Connections tab.	-

VPN Connections Summary View

The **Connectivity** dashboard allows you to open the **VPN Connections** dashboard or summary view.

VPN Connection: Represents the VPN IPSec tunnel established between local gateway and remote gateway. When you create a VPN connection, you need to select two gateways between which you want to create the VPN connection.

The **VPN Connections** dashboard provides a list of VPN connections created for the clusters managed by the Prism Central.

The **VPN Connections** dashboard provides a **Create VPN Connection** button that opens the **Create VPN Connection**. For more information, see Creating a VPN Connection in *Flow Networking Guide*.

You can select a VPN connection from the list (select the checkbox provided for the VPN connection) and then perform an action provided in the **Actions** dropdown list. The **Actions** dropdown list allows you to **Update** or **Delete** the selected VPN connection.

The **VPN Connections** summary list view provides the following details about the VPN connection.

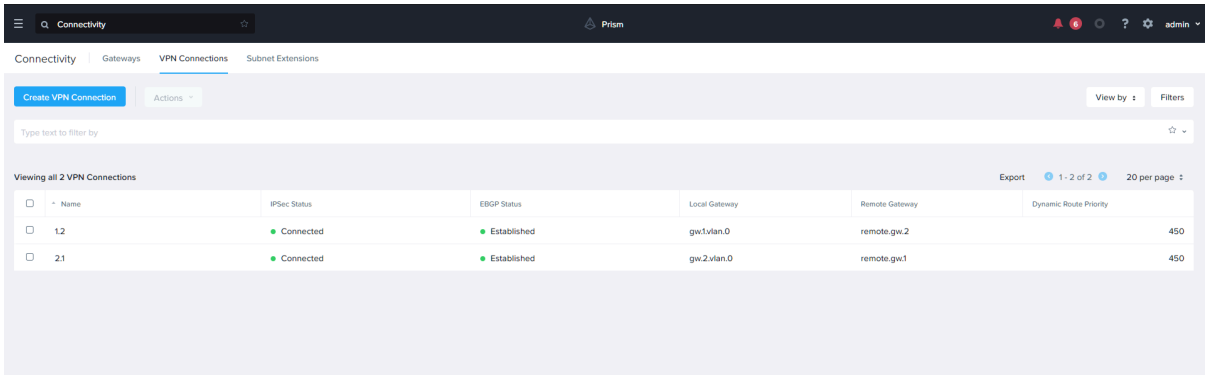

Figure 131: VPN Connections dashboard

Table 69: VPN Connections List Fields

Parameter	Description	Values
Name	Displays the name of the connection.	(gateway name)
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Local Gateway	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Remote Gateway	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Flow Virtual Networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)

VPN Connections Details View

You can click the name of a VPN connection in the **VPN Connections** dashboard list to open the VPN connection details page that presents the information about the VPN connection in widgets.

The VPN connection details page displays the name of the VPN connection on the top left corner.

- On the top right corner, the close button (X) allows you to close the details page.
- The **Update** button opens the **Update VPN Connection** page. For more information, see Updating a VPN Connection in *Flow Networking Guide*.
- The **Delete** button allows you to delete the VPN connection. For more information, see Deleting a VPN Connection in *Flow Networking Guide*.

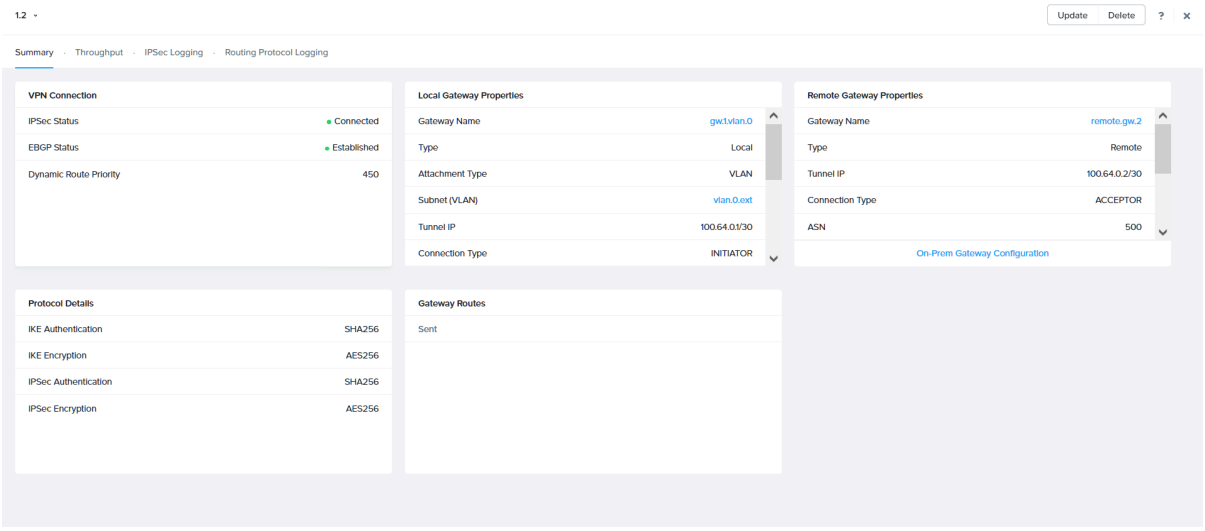

Figure 132: VPN Connection Details

The details about the VPN connection are organized in widgets as follows:

- Summary tab—See the *VPN Connection Summary Tab Details* table below.
- Throughput tab—See the *VPN Connection Throughput Tab Details* table below.
- IPSec Logging tab—Provides logs for the IPSec tunnel.
- Routing Protocol Logging tab—Provides logs for the routing protocol used in the VPN connection.

Table 70: VPN Connection Summary Tab Details

Parameter	Description	Values
VPN Connection widget		
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Flow Virtual Networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)
Local Gateway Properties		
Gateway Name	Displays the name of the local gateway used for the connection.	(Name of local gateway)

Parameter	Description	Values
Type	Displays the type of gateway.	(Local)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Tunnel IP	Displays the Tunnel IP address of the local gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address that you assigned to the local gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	(Up or Down)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Remote Gateway Properties		
Gateway Name	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Type	Displays the type of gateway.	(Remote)
Tunnel IP	Displays the Tunnel IP address of the remote gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)

Parameter	Description	Values
ASN	Displays the ASN of the EBGP route. This information is only displayed if you configured EBGP as the External Routing protocol.	(Number)
Vendor	Displays the name of the vendor of the gateway appliance at the remote site.	(Name of vendor of gateway appliance)
External IP	Displays the IP address assigned to remote the gateway.	(IP Address that you assigned to the remote gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	-
Protocol Details		
Service	Displays the service used by the gateway.	(VPN or VTEP)
Gateway Routes	Displays the status of the routes used by the gateways.	(Sent)

Viewing an Extended Subnet

The Connectivity dashboard opens on the Subnet Extensions dashboard or summary view.

About this task

Perform the following procedure to view an extended subnet at the on-prem site.

Procedure

1. Click the hamburger icon in the top-left corner of the **Dashboard**.
   The main feature list appears.
2. Click **Network & Security** > **Connectivity** > > **Subnet Extensions**.
   The **Subnet Extensions** tab displays a list of the extended subnets.

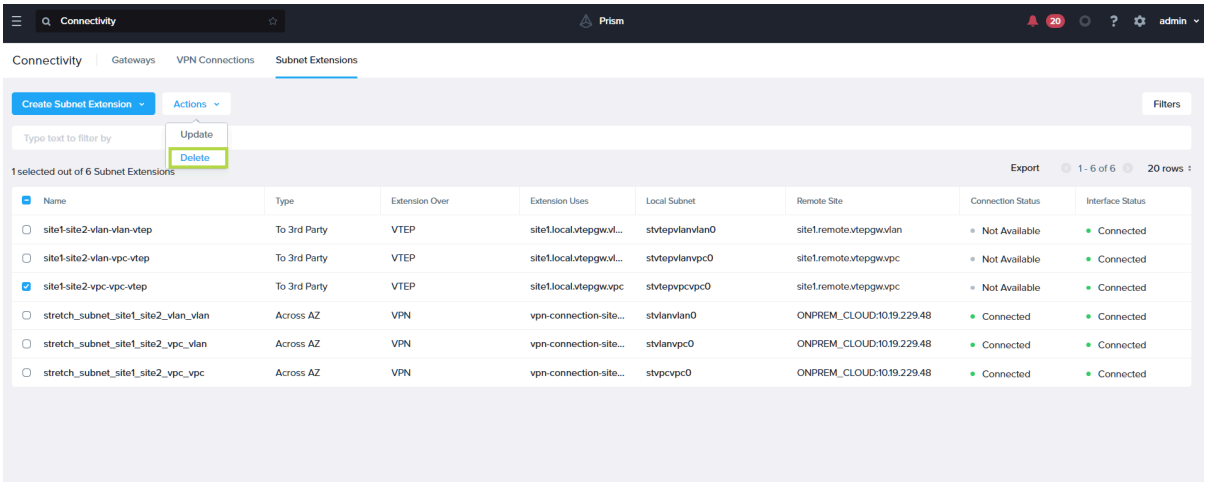

Figure 133: Extended Subnet

What to do next

You can also see which subnets are extended, in the **Subnets** dashboard. The **List** tab displays the list of configured subnets. The subnets that are extended have the **Extended Subnet** tag alongside their name.

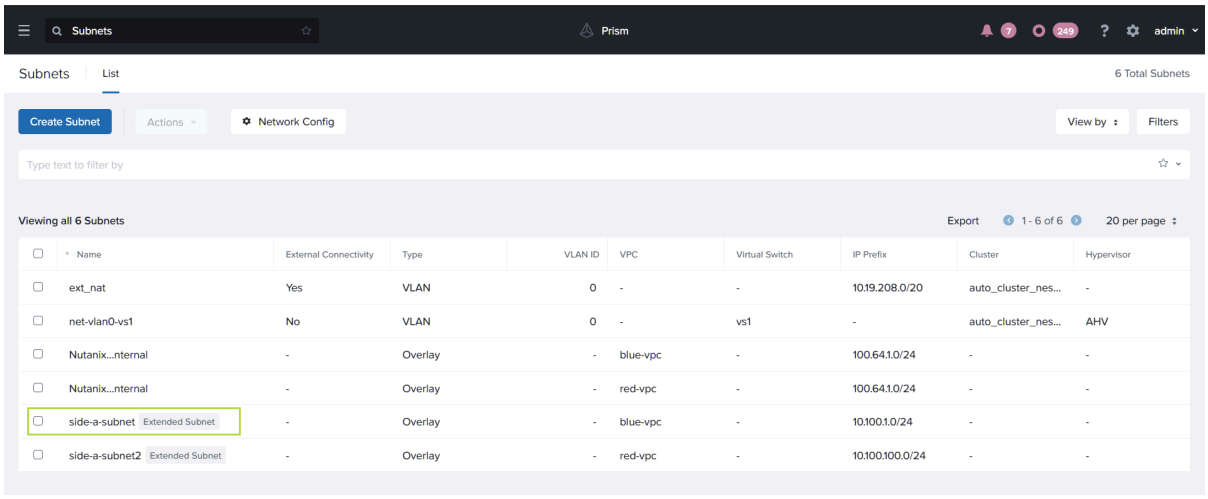

Figure 134: Subnet List View - Extended Subnets

Security Policies Summary View

To access the security policies dashboard, select **Policies** > **Security Policies** from the entities menu (see Entities Menu). The security policies dashboard allows you to view summary information about defined security policies.

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

- See Entity Exploring for instructions on how to view and organize that information in a variety of ways.
- See Flow Microsegmentation Guide for information about how to create and apply security policies.

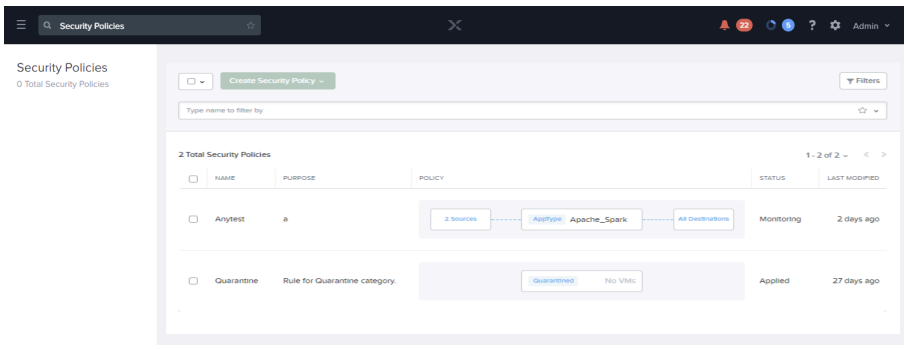

Figure 135: Security Policies Dashboard

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

Table 71: Security Policies List Fields

Parameter	Description	Values
Name	Displays the policy name. The policy is one of three types: application, quarantine, or isolation.	(name), Application, Quarantine, Isolation
Purpose	Describes (briefly) the policy's purpose.	(text string)
Policy	Displays (high level) what the policy does.	(boxed text)
Status	Displays the current status of the policy (either applied currently or in monitoring mode).	Applied, Monitoring
Last Modified	Displays the date the policy was last modified (or the creation date if the policy has never been modified).	(date)

You can filter the security polices list based on several parameter values. The following table describes the filter options available when you open the Security Policies view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 72: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of security policies that satisfy the name condition/string.	(policy name string)
Type	Filters on the policy type. Check the box for one or more of the policy types (application, quarantine, isolation). It will limit the list to just those policy types.	Application, Quarantine, Isolation
Status	Filters on the policy status. Check the box for applied or monitoring.	Applied, Monitoring

The security policies dashboard includes a **Create Security Policy** action button with a drop-down list to **Secure an Application** or **Isolation Environments**.

The **Actions** menu appears when one or more policies are selected. It includes options to update, apply, monitor, and delete. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Security Policy Details View

This Guide does not provide detailed information about *Security Policies*. It only describes the Summary and Details views for the same. For detailed information about *Security Policies*, see Flow Microsegmentation Guide

To access the details page for a security policy, click on the desired security policy name in the list (see Security Policies Summary View on page 187). The Security Policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.

- The rule status appears below the name and indicates whether the policy is being applied currently or is in monitoring mode.
- Three columns appear that specify the Inbound policy (on the left), the affected entities (in the middle), and the Outbound policy (on the right).
- There are three action buttons (upper right).
  - Click the appropriate button to update, apply, monitor, or delete the policy (see *Nutanix Security Guide* for details). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

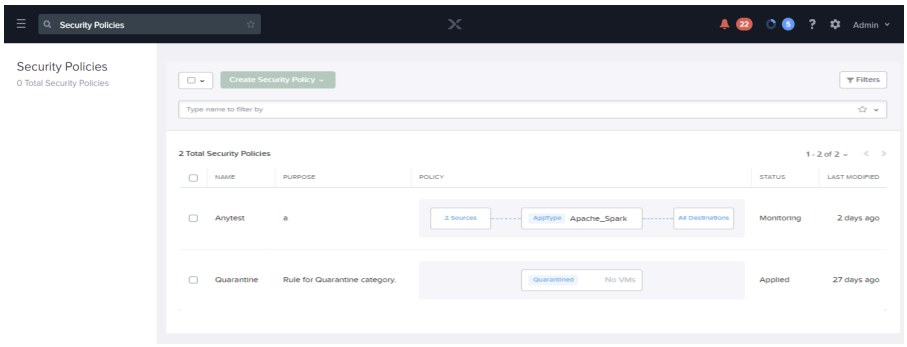

Figure 136: Security Policy Details View: Monitoring Rule Example

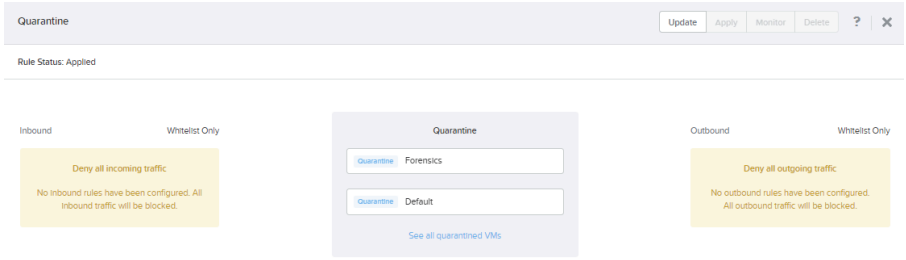

Figure 137: Security Policy Details View: Applied Rule Example

For more information about Security Policies, see Flow Microsegmentation Guide.

Policies

You can access dashboards for the following policy types from the **VM**, **Images**, **Storage**, and **Network & Security** entities.

> **Note:** For information about creating and applying protection policies and recovery plans, see Nutanix Disaster Recovery Guide.

VM Policies

- Affinity Policies (see Affinity Policies Summary View on page 190

- NGT Policies (see NGT Policies View on page 192)

Image Policies

- Image Placement Policies (see Image Placement Policies Summary View on page 193)
- Bandwidth Throttling Policies (see Bandwidth Throttling Policies Summary View on page 194)

Storage Policies

Storage Policies (see Storage Policies Summary View on page 141)

Security Policies

Storage Policies (see Security Policies Summary View on page 195)

Affinity Policies Summary View

The affinity policies dashboard displays a list of current policies that includes the name and type for each policy. To access the affinity policies dashboard:

1. Select **Compute & Storage** > **VMs** from the entities menu (see Entities Menu on page 13). The VMs dashboard page opens.
2. Select **Policies** > **Affinity Policies** from the drop-down menu.

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

- See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.
- See Affinity Policies Defined in Prism Central on page 608 for information on how to create or modify the affinity policies.

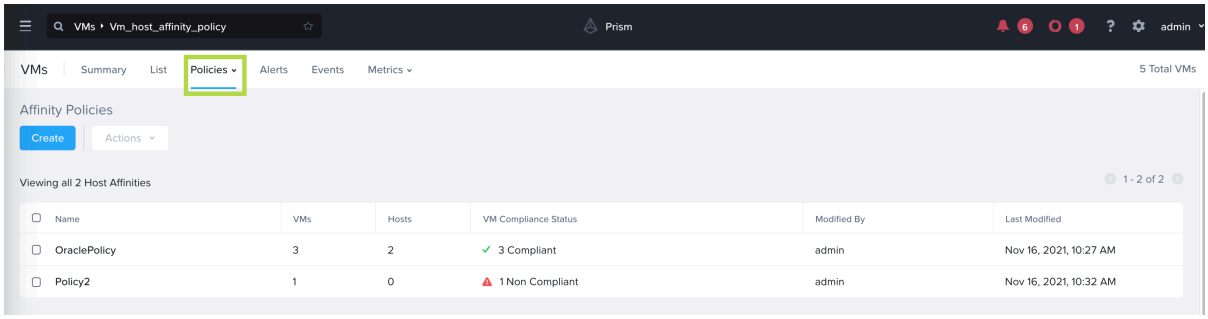

Figure 138: Affinity Policies Dashboard

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

Table 73: Affinity Policies List Fields

Parameter	Description	Values
Name	Displays the policy name.	(name)
VMs	Displays the count of VMs associated with this policy.	(number of VMs)
Hosts	Displays the count of hosts associated with this policy.	(number of hosts)
VM Compliance Status	Displays the compliance status of the VMs associated with this policy. If the policy is being applied and the compliance status is not yet known, the status is displayed as Pending. If a VM is part of multiple VM-Host affinity policies, the oldest policy is applied on the VM. For rest of the policies, the VM is displayed as non-compliant.	(number of VMs Compliant/Non Compliant/Pending)
Modified By	Displays the name of the user who modified the policy last time.	(user)
Last Modified	Displays the date and time when the policy was modified last time.	(date & time)

Affinity Policies Details View

To access the details page for an Affinity policy, click the desired policy name in the list (see Affinity Policies Summary View on page 190). The affinity policy details page includes the following tabs:

- **Summary**: On the **Summary** tab, you can view the **Overview**, **Associations**, and **Compliance Status** sections. All the three sections display information related to the policy.

  The **Summary** tab view also includes options to **Update**, **Delete**, and **Re-Enforce** the policy. If any of the VMs become non-compliant, you can use this option to re-enforce the policy after fixing the issue for the non-compliance.

- **Entities**: On the **Entities** tab, you can view the details of the **VMs** and **Hosts** entities that are associated with the policy. The details displayed on the **Hosts** tab include the host name, cluster name, and the category used to associate, while the details displayed on the **VMs** tab include the VM name, host name, cluster name, category used to associate, and VM compliance status. If a VM is non-compliant, the cause of non-compliance is also displayed along with the status.

- **Categories**: The **Categories** tab displays the VMs and hosts categories information.

- **Audits**: The **Audits** tab displays the user action-related information for the affinity policies.

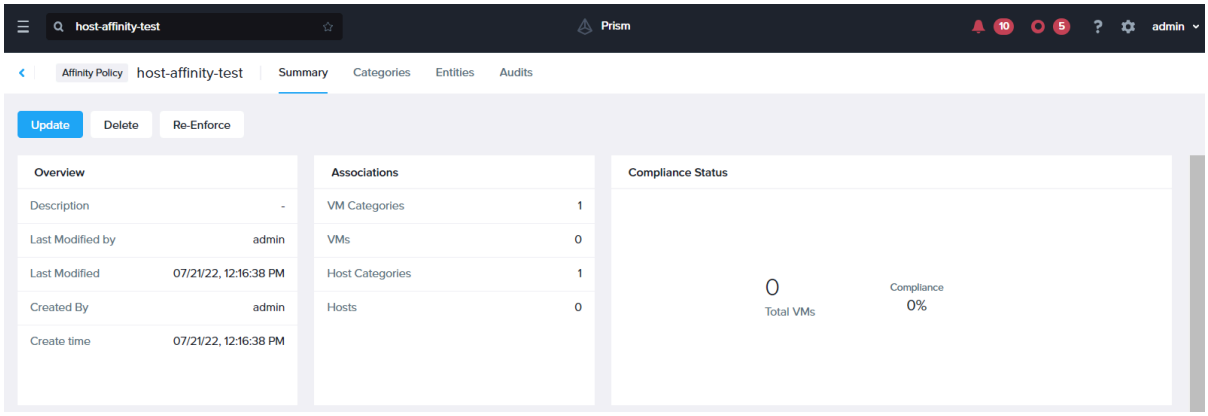

Figure 139: Affinity Policies Details view

NGT Policies View

The NGT policies dashboard allows you to view information about existing NGT policies and create new policies. To access the Nutanix guest tools (NGT) policies dashboard,

1. Select **Compute & Storage** > **VMs** from the entities menu (see Entities Menu on page 13). The VMs dashboard page opens.
2. Select **Policies** > **NGT Policies** from the drop-down menu.

> **Note:** This section describes the information and options that appear in the NGT policies dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See NGT Policies on page 611 for information on how to create or modify NGT policies.

The NGT dashboard displays a list of current policies that includes the name and type for each policy.

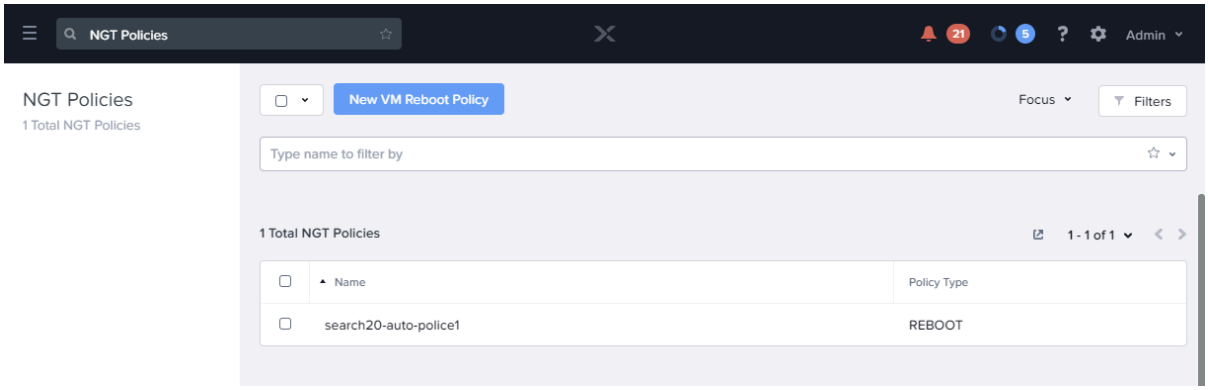

Figure 140: NGT Policies Dashboard

To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the NGT policies list based on the NGT policy name.

There is a **New VM Reboot Policy** action button to create a new reboot policy (see Creating a New Restart Policy on page 611). The **Actions** menu appears when one or more policies are selected and includes the following options:

- **Update** (see Modifying a Restart Policy on page 613)
- **Delete**

Image Placement Policies Summary View

The image placement policies dashboard allows you to view summary information about defined image placement policies.

To access the image placement policies dashboard, select **Images** > **Policies** > **Image Placement** from the **Entities** menu (see Entities Menu on page 13).

> **Note:** This section describes the information and options that appear in the image placement policies dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See Image Placement Policies on page 690 for information about how to create and apply image placement policies.

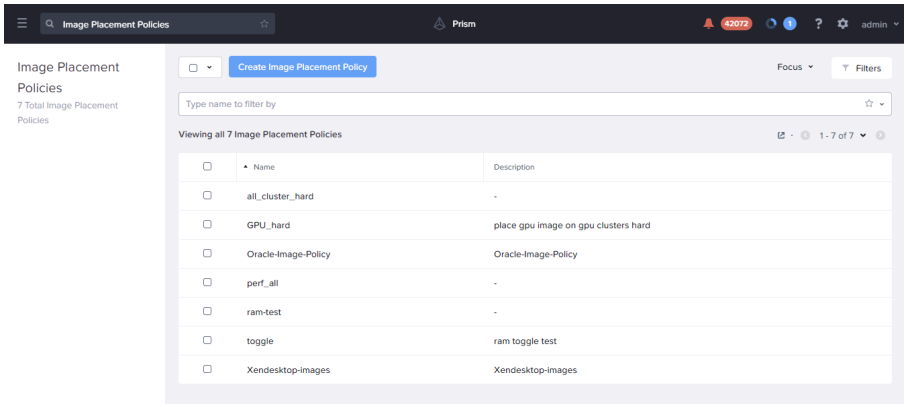

Figure 141: Image Placement Policies Dashboard

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

Table 74: Image Placement Policies List Fields

Parameter	Description	Values
Name	Displays the policy name.	(name)
Description	Describes (briefly) the purpose of the policy.	(text string)

You can filter the image placement polices list by opening the **Filter** pane. To apply a filter, check the box and enter a text string in the **Contains** field for the Name or Description parameter. You can apply filters across both parameters.

The image placement policies dashboard includes a **Create Image Placement Policy** action button (see Configuring an Image Placement Policy on page 692).

The **Actions** menu appears when one or more policies are selected. It includes options to **Update** and **Delete**. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Image Placement Policy Details View

To access the details page for an image placement policy, click on the desired policy name in the list (see Image Placement Policies Summary View on page 193). The image placement policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The policy enforcement status appears below the name and indicates whether the enforcement is hard or soft.
- Two columns appear that specify the images in question (on the left) and the target clusters (on the right). Click **Expand all** (or **Collapse all**) to display (or hide) the list of images or clusters.
- There are four action buttons (upper right).
  - Click the **Update** or **Delete** button to update or delete the policy (see Updating an Image Placement Policy on page 692). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

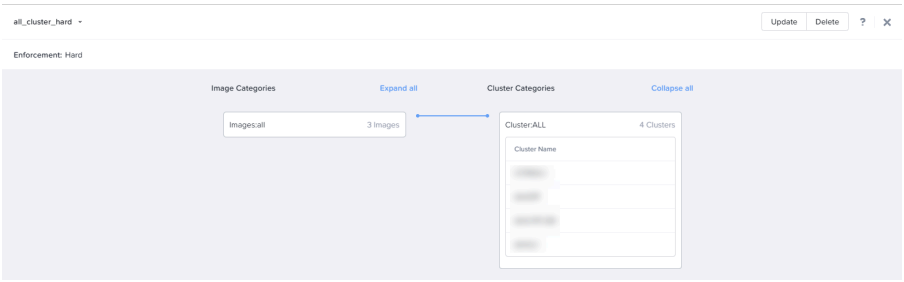

Figure 142: Image Placement Policy Details View

Bandwidth Throttling Policies Summary View

The bandwidth throttling policy allows you to limit the bandwidth consumed during image creation using the URL option in specific clusters. The bandwidth throttling policies dashboard allows you to view summary information about user-defined bandwidth throttling policies.

To access the bandwidth throttling policies dashboard, select **Images** > **Policies** > **Bandwidth Throttling Policies** from the **Entities** menu (see Entities Menu on page 13).

Note: This section describes the information and options that appear in the bandwidth throttling policies dashboard.

- See Entity Exploring on page 89 for instructions on how to view and organize that information in various ways.
- See Bandwidth Throttling Policies on page 695 for information about how to create and apply bandwidth throttling policies.

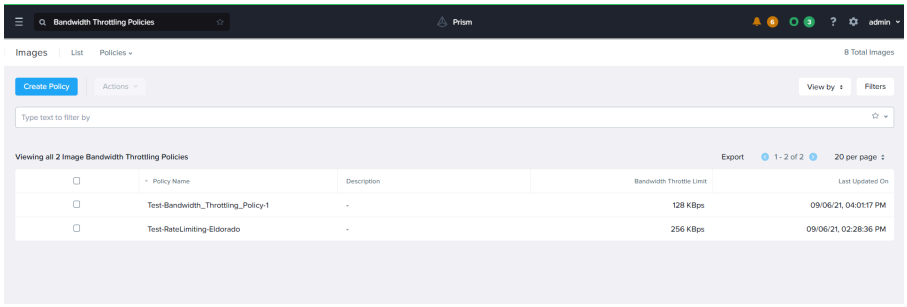

Figure 143: Bandwidth Throttling Policies Dashboard

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

Table 75: Bandwidth Throttling Policies List Fields

Parameter	Description
Policy Name	Displays the policy name.
Description	Describes (briefly) the purpose of the policy.
Bandwidth Throttle Limit	Describes the bandwidth throttle limit configured for the policy.
Last Updated On	Describes the date and time when the bandwidth throttling policy was created or updated.

Click the check box associated with the policy name to update or delete the bandwidth throttling policy.

Select specific policies and click **Actions** > **Update** to update the selected bandwidth throttling policies.

Select **Actions** > **Delete** to delete the selected bandwidth throttling policies.

You can filter the bandwidth throttling polices list by opening the **Filter** pane. To apply a filter, check the box and enter a text string in the **Contains** field for the name or description parameter. You can apply filters across both parameters.

The bandwidth throttling policies dashboard includes a **Create Policy** action button (see Creating Bandwidth Throttling Policy on page 696).

Security Policies Summary View

To access the security policies dashboard, select **Network & Security** > **Security Policies** from the entities menu (see Entities Menu on page 13). The security policies dashboard allows you to view summary information about defined security policies.

> **Note:** This section describes the information and options that appear in the security policies dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See Flow Microsegmentation Guide for information about how to create and apply security policies.

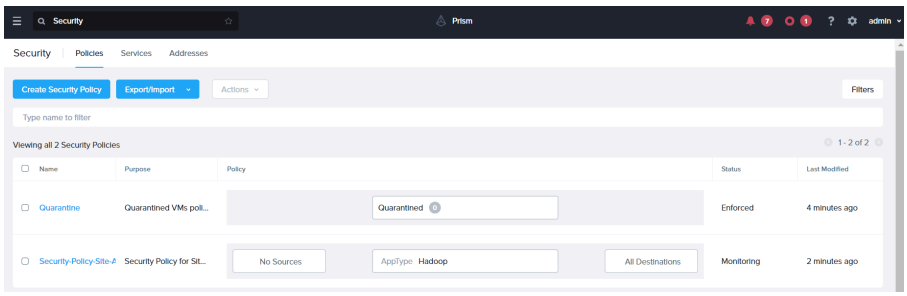

Figure 144: Security Policies Dashboard

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

Table 76: Security Policies List Fields

Parameter	Description	Values
Name	Displays the policy name. The policy is one of three types: application, quarantine, or isolation.	(name), Application, Quarantine, Isolation
Purpose	Describes (briefly) the policy's purpose.	(text string)
Policy	Displays (high level) what the policy does.	(boxed text)
Status	Displays the current status of the policy (either enforced currently or in monitoring mode).	Enforced, Monitoring
Last Modified	Displays the time-line for the security policy. Last modified time-line or the creation timeline if the policy has never been modified.	Time-line. For example, 4 days ago, 6 minutes ago

You can filter the security polices list based on several parameter values. The following table describes the filter options available when you open the Security Policies view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 77: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the item name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of security policies that satisfy the name condition/string.	(policy name string)
Type	Filters on the policy type. Check the box for one or more of the policy types (application, quarantine, isolation). It will limit the list to just those policy types.	Application, Quarantine, Isolation
Status	Filters on the policy status. Check the box for applied or monitoring.	Applied, Monitoring

The security policies dashboard includes a **Create Security Policy** action button with a drop-down list to **Secure an Application** or **Isolation Environments**.

The **Actions** menu appears when one or more policies are selected. It includes options to update, apply, monitor, and delete. The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)

Services Tab

Service is a group of protocol-port combination. You can use any of the default services or create a custom service. For more information about how to create a service group, and update, delete, or clone a service, see Flow Microsegmentation Guide.

Addresses Tab

Address Group is a way to group one or many IP address or range. It can be used when you define security policies. For more information about how to create, update or delete an address, see Flow Microsegmentation Guide.

Security Policy Details View

To access the details page for a security policy, click on the desired security policy name in the list (see Security Policies Summary View on page 195). The Security Policy details page includes the following:

- The policy name appears in the upper left. You can switch from one policy to another by selecting the policy name from the pull-down list.
- The rule status appears below the name and indicates whether the policy is being applied currently or is in monitoring mode.
- Three columns appear that specify the Inbound policy (on the left), the affected entities (in the middle), and the Outbound policy (on the right).
- There are three action buttons (upper right).
  - Click the appropriate button to update, apply, monitor, or delete the policy (see Security Policies configuration in *Flow Microsegmentation Guide* for details). The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.)
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

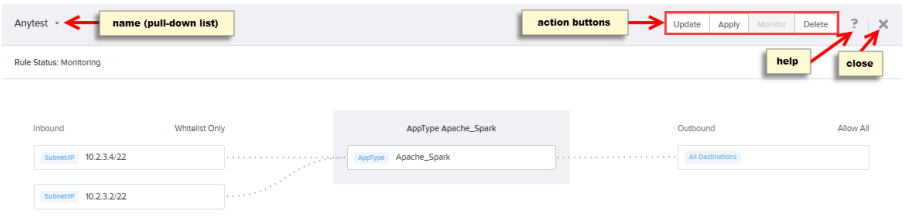

Figure 145: Security Policy Details View: Monitoring Rule Example

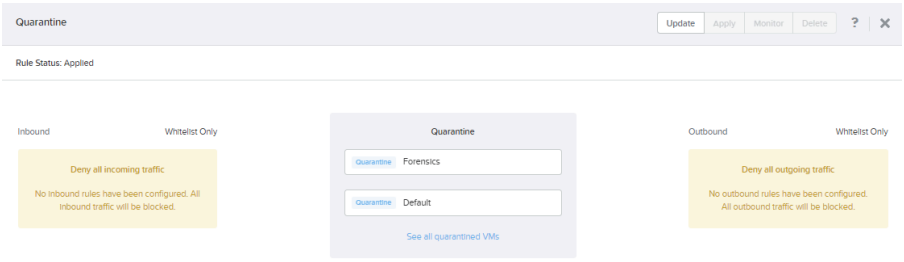

Figure 146: Security Policy Details View: Applied Rule Example

Data Protection and Recovery Entities

You can access dashboards for the following data protection types from the **Data Protection** category of the entities menu (see Entities Menu on page 13):

- Protection Summary
- Protection Policies (see Protection Policies View on page 198)
- Recovery Plans (see Recovery Plans View on page 199)
- VM Recovery Points (see VM Recovery Points View on page 200)

Protection Policies View

To access the protection policies dashboard, select **Policies** > **Protection Policies** from the entities menu (see Entities Menu on page 13). The protection policies dashboard allows you to view information about the current policies and create new policies.

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

- For instructions on how to view and organize that information in a variety of ways, see Entity Exploring on page 89.
- For information about how to create and use protection policies, see Nutanix Disaster Recovery Guide.

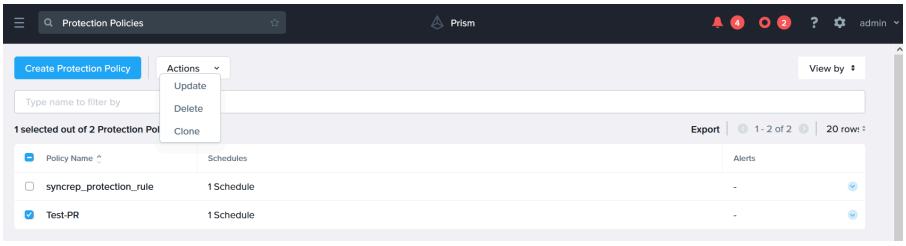

Figure 147: Protection Policies Dashboard

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

Table 78: Protection Policies Fields

Parameter	Description	Values
Name	Displays the protection policy name.	(item name)
Source	Displays the replication source name for the protection policy	(source name)
Destination	Displays the replication destination name for the protection policy.	(destination name)
RPO	Displays the recovery point objective (RPO) for the protection policy	(time interval)
Remote Retention	Displays the number of retention points at the remote availability zone.	*xx* Recovery Points
Local Retention	Displays the number of retention points at the local availability zone.	*xx* Recovery Points

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

Click the **Create Protection Policy** button to create a new protection policy. The **Actions** menu appears when one or more protection policies are selected and includes the following options:

- **Update**: Update the protection policy.
- **Clone**: Clone the protection policy.
- **Delete**: Delete the protection policy.

Recovery Plans View

To access the recovery plans dashboard, select **Policies** > **Recovery Plans** from the entities menu (see Entities Menu on page 13). The recovery plans dashboard allows you to view information about the current plans and create new plans.

> **Note:** This section describes the information and options that appear in the recovery plans dashboard.
>
> - For instructions on how to view and organize that information in a variety of ways, see Entity Exploring on page 89.
> - For information about how to create and use recovery plans, see Nutanix Disaster Recovery Guide.

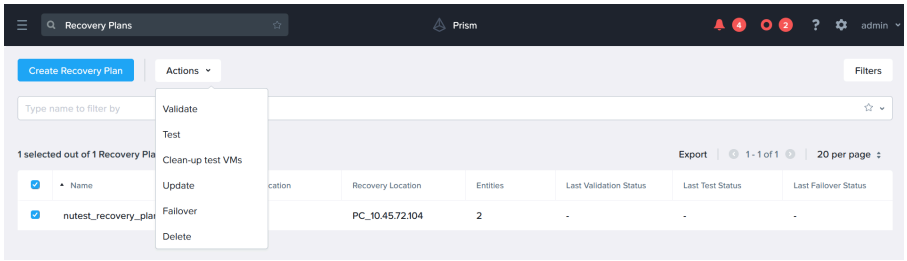

Figure 148: Recovery Plans Dashboard

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

Table 79: Recovery Plans Fields

Parameter	Description	Values
Name	Displays the recovery plan name.	(name)
Source	Displays the primary availability zone for the recovery plan.	(zone name)
Destination	Displays the recovery availability zone for the recovery plan.	(zone name)
Entities	Displays the number of VMs and snapshots associated with the recovery plan in the availability zone in which you are viewing the recovery plan.	(integer)
Last Validation Status	Displays the status of the most recent validation of the recovery plan.	
Last Test Status	Displays the status of the most recent test performed on the recovery plan.	

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

Click the **Create Recovery Plan** button to create a new recovery plan. The **Actions** menu appears when one or more recovery plans are selected and includes the following options:

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

VM Recovery Points View

To access the recoverable entities dashboard, select **Data Protection & Recovery** > **VM Recovery Points** from the entities menu (see Entities Menu on page 13). A *recoverable entity* is a VM that you can recover from a snapshot. The recoverable entities dashboard allows you to view information about the entities for which snapshots exist in the availability zone.

> **Note:** This section describes the information and options that appear in the recoverable entities dashboard.
>
> - For instructions on how to view and organize that information in a variety of ways, see Entity Exploring on page 89.
> - For information about recoverable entities and how to use them, see Nutanix Disaster Recovery Guide.

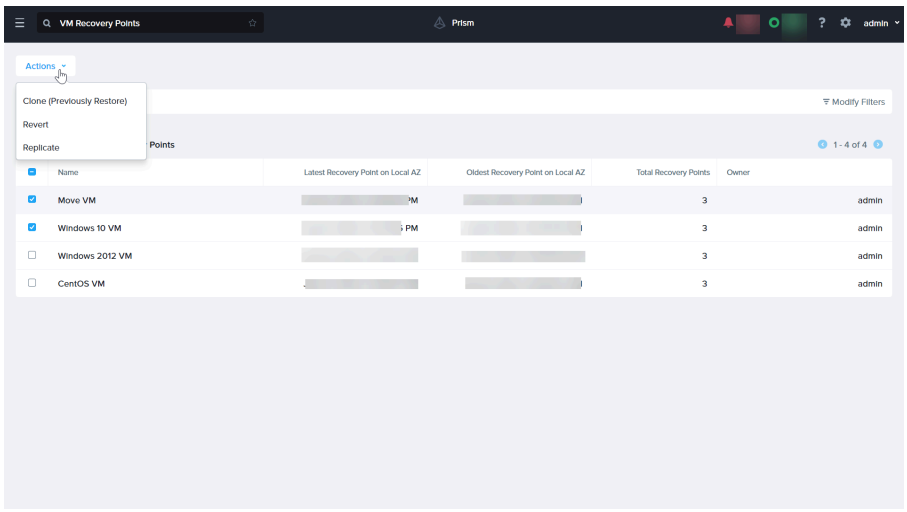

Figure 149: VM Recovery Points Dashboard

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

Table 80: VM Recovery Points List Fields

Parameter	Description	Values
Name	Displays the recoverable entity name.	(item name)
Latest Recovery Point on Local AZ	Displays the time at which the most recent snapshot was taken. This is specific to the local availability zone.	(time and date)
Oldest Recovery Point on Local AZ	Displays the time at which the oldest available snapshot was taken. This is specific to the local availability zone.	(time and date)
Total Recovery Points.	Displays the total number of recovery points available locally.	(integer)
Owner	Displays the owner who created the recovery points	(user name)

To filter the list by name, click the **Filters** button (upper right) and enter a name string in the field.

The **Actions** menu appears when one or more recoverable entities are selected and includes the following options:

- **Clone (Previously Restore)**: Clone a VM entity from the snapshot taken at the recovery point.
- **Revert**: Revert a VM entity to the snapshot taken at the recovery point. You can perform Revert operation only if you have admin role assigned to your local user or directory user via role mapping.
- **Replicate**: Replicate the snapshot to a paired availability zone.

Hardware Entities

You can access dashboards for the following hardware components from the **Hardware** category of the entities menu (see Entities Menu on page 13):

- Clusters (see Clusters Summary View on page 202)
- Hosts (see Hosts Summary View on page 221)
- Disks (see Disks Summary View on page 230)
- GPUs (see GPUs Summary View on page 237)

Clusters Summary View

To access the clusters dashboard, select **Hardware** > **Clusters** from the entities menu (see Entities Menu on page 13). The clusters dashboard allows you to view summary information about registered clusters and access detailed information about each cluster. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

> **Note:** This section describes the information and options that appear in the clusters dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

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

- **Highlighted Entities**: Displays a list of the clusters with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **IOPS**, and **Aggregate Hypervisor Memory Usage Pp**. Click the **View all XX Clusters** link at the bottom to display the **List** tab (following section).

- **Alerts**: Displays a list of cluster-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.

- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

- **Storage Usage**: Displays a list of clusters along with the percentage usage (with respect to total usable capacity of the cluster) and data reduction. Total usage is calculated as the sum of used capacity and free reserved capacity. Data reduction ratio displays the ratio of how much the data size is reduced by enabling compression, deduplication, and erasure coding. Placing the cursor anywhere on the horizontal axis displays a breakdown view of the storage capacity usage.

- **Data Resiliency**: Displays a list of clusters along with the fault domain and the current state of the fault tolerance. Fault tolerance of a cluster can have any one of the following four states:

  - OK: This state indicates that the fault tolerance domain is highly resilient to safely handle a node or a disk (in single or two node clusters) failure.
  - Warning: This state indicates that the fault tolerance level is almost reaching to 0. Warning state is displayed if the cluster is not fault tolerant at the configured domain, but is fault tolerant at a lower domain. For example, if you have configured rack as the configured domain and the cluster can no longer handle any rack failures due to some reason but can still handle node (lower domain) failures, then fault tolerance state is displayed as Warning.
  - Critical: This state indicates that the fault tolerance level is 0, and the fault tolerance domain cannot handle a node or a disk (in single or two node clusters) failure.

- Computing: This state indicates that the new fault tolerance level is being calculated. This state is displayed soon after a node or disk failure, before rebuild is initiated.

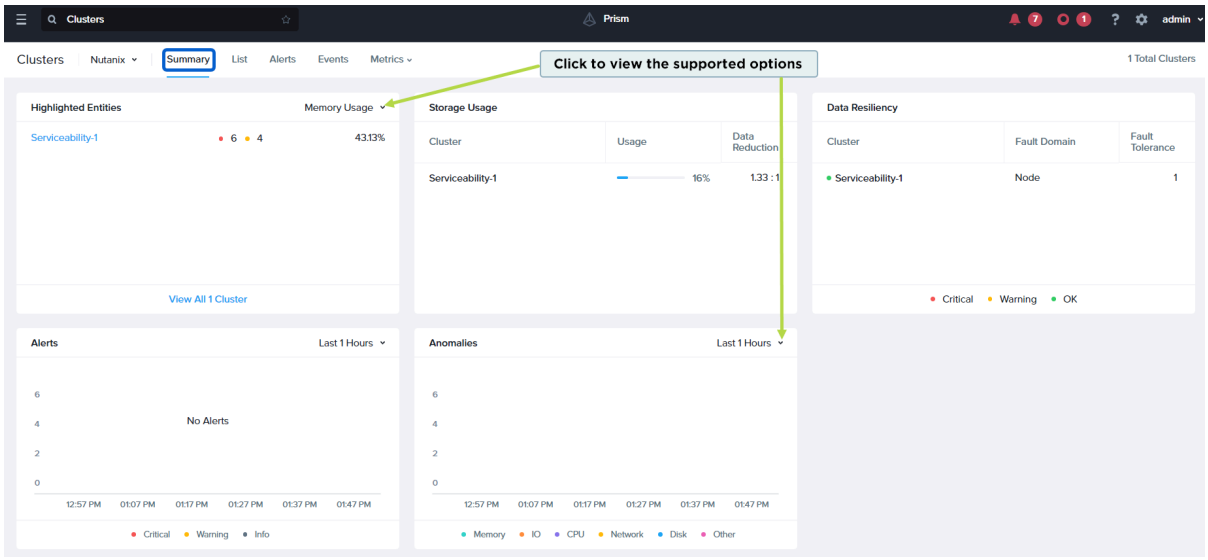

Figure 150: Clusters Summary Tab (Nutanix)

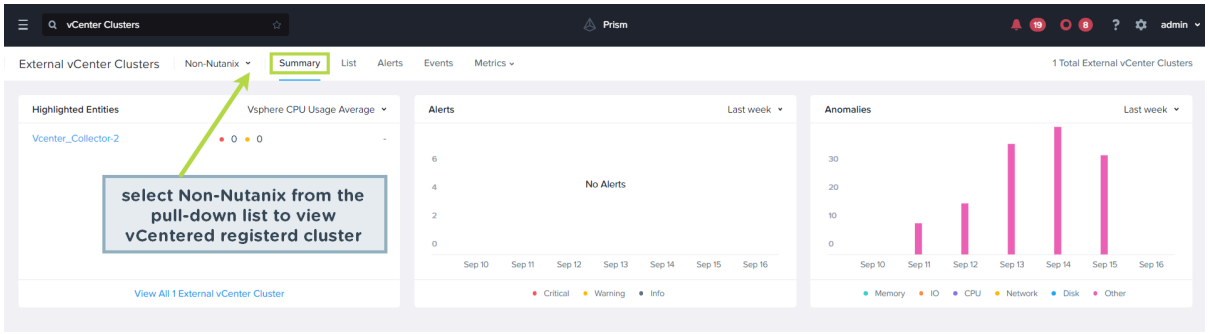

Figure 151: Clusters Summary Tab (Non-Nutanix)

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the registered clusters. The following table describes the fields that appear in the clusters list. The fields vary based on the **View by** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

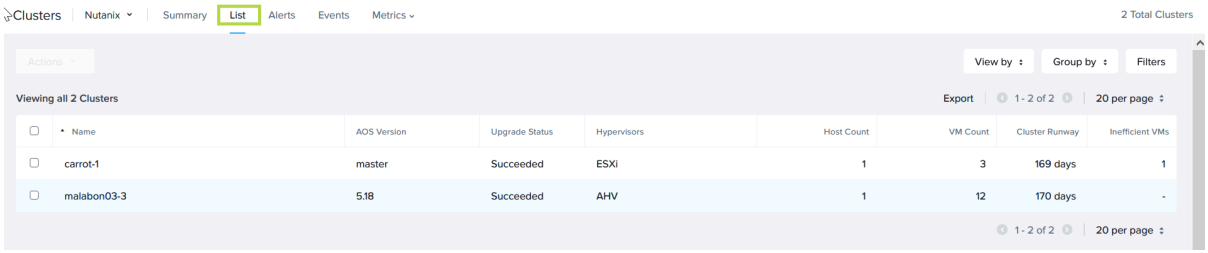

Figure 152: Clusters List Tab (Nutanix)

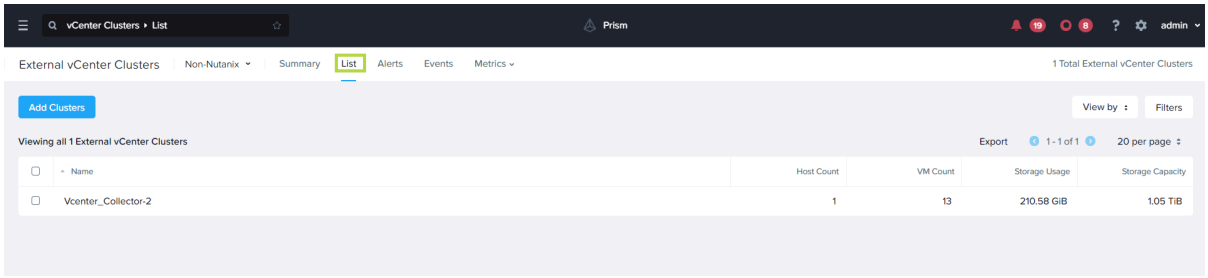

Figure 153: Clusters List Tab (Non-Nutanix)

Table 81: Clusters List Fields

Parameter	Description	Values
View by "General" Fields		
Name	Displays the cluster name. Clicking on the name displays the details page for that cluster (see Cluster Details View on page 209).	(cluster name)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Upgrade Status	Displays the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, Hyper-V
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)
VM Count	Displays the total number of VMs in the cluster (in any state).	(number of VMs)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 209 for more information).	(number of days)
Inefficient VMs	Displays the number of inefficient VMs in the cluster (see Behavioral Learning Tools on page 899 for more information).	(number)
View by "Performance" Fields		
Name	Displays the cluster name.	(cluster name)
CPU Usage	Displays the percentage of CPU capacity in the cluster currently being used.	0 -100%

Parameter	Description	Values
Memory Usage	Displays the percentage of memory capacity in the cluster currently being used.	0 -100%
IOPS	Displays total (both read and write) I/O operations per second (IOPS) for this cluster.	(number)
IO Bandwidth	Displays total I/O bandwidth used per second in this cluster.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency in this cluster.	xxx [ms]
View by "Capacity" Fields		
Name	Displays the cluster name.	(cluster name)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see "Capacity Tab" section in Cluster Details View on page 209 for more information).	(number of days)
CPU Runway	Displays the predicted CPU runway (time period) before the cluster requires additional resources.	(number of days)
Memory Runway	Displays the predicted memory runway (time period) before the cluster requires additional resources.	(number of days)
Storage Runway	Displays the predicted storage memory runway (time period) before the cluster requires additional resources.	(number of days)
View by "Encryption" Fields		
Name	Displays the cluster name.	(cluster name)
Encryption Type	Displays the cluster encryption type.	Not Encrypted, Software, Hardware, Software and Hardware
Encryption Scope	Displays the encryption scope.	Cluster, Entity
KMS Type	Displays the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)

You can filter the cluster list based on a variety of parameter values. The following table describes the filter options available when you open the Clusters view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 82: Filter Pane Fields

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the pull-down list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)
Name	Filters on the cluster name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of clusters that satisfy the name condition/string.	(cluster name string)
AOS Version	Filters on AOS version. Select one or more versions to return a list of clusters running those version(s). The number of clusters currently running each version is displayed on the right of the line.	(Acropolis version numbers across clusters currently)
Hypervisors	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, HyperV
Health	Filters on the cluster health state (good, warning, or critical). Select one or more states to return a list of clusters in that state(s). The number of clusters currently in each state is displayed on the right of the line.	Critical, Warning, Good
CPU Usage	Filters on the amount of total CPU being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
IOPS	Filters on the total (both read and write) IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of clusters with total IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the total I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of clusters with total I/O bandwidth usage in that range.	([xx] to [yy] range)

Parameter	Description	Values
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of clusters with average I/O latency in that range.	([xx] to [yy] range)
Upgrade Status	Filters on the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Categories	Filters on the category type. Search for the category name. For example, ADGroup:$Default	(category name string)
Encryption Type	Filters on the encryption type. Check the box for the desired encryption type. It will return a list of clusters with the corresponding encryption type applied.	Not Encrypted, Software, Hardware, Software and Hardware
Encryption Scope	Filters on the encryption scope. Check the box for the entities where encryption is applied on. It will return a list of clusters with the corresponding encryption scope applied.	Cluster, Entity
KMS Type	Filters on the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External

You can group the clusters list in the following ways:

- The **Group by** pull-down menu allows you to group the cluster entries by AOS version, host count, or health state. (You can only choose one.)

- When you select one or more clusters, the **Label** icon appears (to the left of the **Actions** menu). To assign the selected clusters a group label, click the **Label** icon and do one of the following:

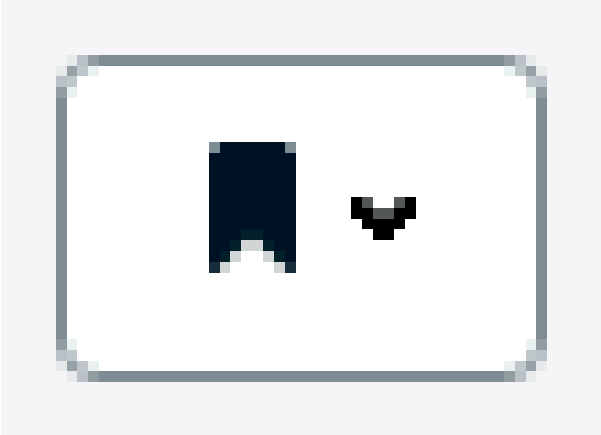

Figure 154: Label Icon

  - To assign an existing label (if there are existing labels), select the desired label and then click **Apply changes**.
  - To assign a new label, enter a meaningful label in the blank text field and then click **Create new label**.

  You can now perform any actions available on this labeled group from the **Actions** menu (see following section). You can edit or delete the label by clicking the **Label** icon and selecting **Manage Labels**.

The **Actions** menu appears when a cluster is selected. The pull-down list includes the following actions:

- Select **Launch Prism Element** to launch Prism element for that cluster in a separate tab or window (depending on your browser settings).

  > **Note:** When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user drop-down list are **REST API Explorer**, **About Nutanix**, **Support Portal**, **Help**, **Nutanix Next Community**, and **Sign Out**.

-
- Select **Upgrade Software** to upgrade the AOS version on that cluster (see Installing or Upgrading on page 30).
- Select **Rack Configuration** to configure the rack awareness feature. This option appears only for clusters that satisfy the conditions for rack awareness. See Rack Fault Tolerance in *Prism Web Console Guide* for instructions on configuring rack awareness.
- Select **Manage Categories** to configure cluster categories.
- Select **Disable Efficiency Measurement** to disable efficiency measurement for the cluster.
- Select **Disable Anomaly Detection** to disable anomaly detection for the cluster.
- Select **Run Playbook** to run or create a new playbook for the cluster entity type.
- Select **Manage & Backup Keys** to download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just cluster-related alerts (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just cluster-related events (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the clusters. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information below the tab. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

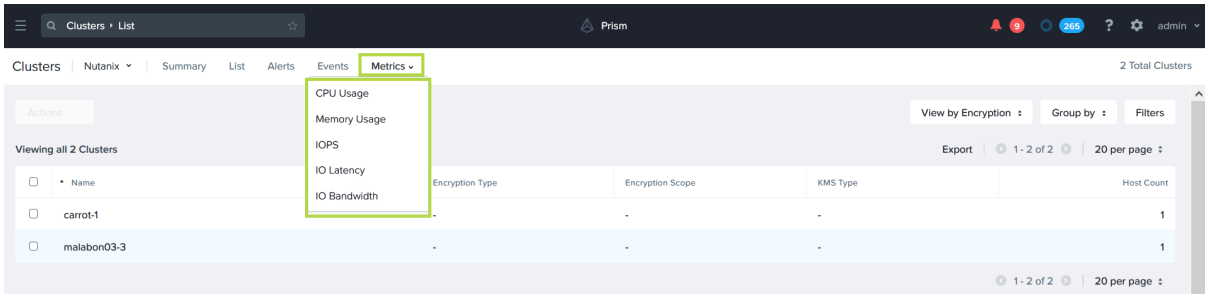

Figure 155: Clusters Metrics Tab

Table 83: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total clusters (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those clusters. **Note:** The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Cluster Details View

To access the details page for a cluster, go to the clusters dashboard **List** tab (see Clusters Summary View on page 202 for more information about VM efficiency) and click the cluster name. You can also access the details page by clicking the cluster name wherever that name appears, such as in a dashboard widget or search result.

The cluster name and the following set of tabs appear on the top: **Summary**, **Alerts**, **Events**, **Metrics**, **Storage Usage**, **Hardware**, **Virtual Entities**, and **Capacity**. Click a tab to display that information below the tab. (Click the **Back to Clusters** link to return to the summary view.)

> **Note:**
>
> For Non-Nutanix clusters, the cluster name and the following set of tabs appear on the top: **Summary**, **Alerts**, **Events**, **Metrics**, and **Capacity**. Click a tab to display that information below the tab. (Click the **Back to Clusters** link to return to the summary view.)

Summary Tab

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

- A **Properties** widget that displays summary information about the cluster (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the pull-down menu. When an anomaly appears, you can click the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.
- A **VM Efficiency** widget that displays the number of VMs, which are considered inefficient broken down by category (overprovisioned, inactive, constrained, and bully). See Behavioral Learning Tools on page 899 for more information about VM efficiency.
- Action buttons (only the applicable ones appear):
  - **Launch Prism Element**: Click this button to launch Prism Element for this cluster in a new tab or window.

    > **Note:** When you access a cluster from Prism Central, you are logging in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user drop-down list are **REST API Explorer**, **About Nutanix**, **Support Portal**, **Help**, **Nutanix Next Community**, and **Sign Out**.

  - **Upgrade Software**: Click this button to upgrade the AOS version on this cluster. See Installing or Upgrading on page 30 .
  - **Rack Configuration**: Click this button to configure the rack awareness feature (see Rack Fault Tolerance in *Prism Web Console Guide*).
  - Select **Manage Categories** to configure cluster categories (see Category Management on page 795).
  - Select **Disable Efficiency Measurement** to disable efficiency measurement for the cluster.
  - Select **Disable Anomaly Detection** to disable anomaly detection for the cluster.
  - Select **Run Playbook** to run or create a new playbook for the cluster entity type (see Creating Playbooks using Triggers on page 922).
  - Select **Manage & Backup Keys** to download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.

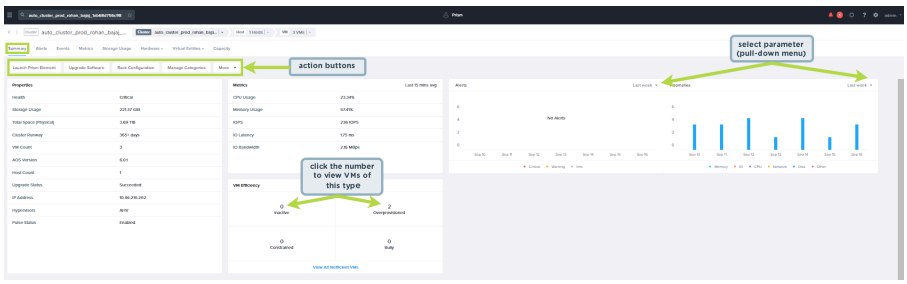

Figure 156: Cluster Summary Tab

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

Table 84: Cluster Properties Fields

Parameter	Description	Values
Health	Displays the cluster health state (good, warning, or critical).	Critical, Warning, Good
Storage Usage	Displays the amount of storage used in the cluster	xxx [GiB|TiB]
Storage Capacity	Displays the total amount of storage capacity in this cluster.	xxx [GiB|TiB]
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see the "Capacity Tab" section below).	(number of days)
VM Count	Displays the number of VMs in the cluster.	(number of VMs)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of hosts)
Upgrade Status	Displays the status of the last (or current) upgrade attempt.	(status condition)
IP Address	Displays the virtual IP address for the cluster (if defined).	(IP address)
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, or Hyper-V

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget allows quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective pull-down menus.

Example: Viewing Cluster Instances

Click the **Cluster** pull-down menu to view the list of cluster instances. Alternatively, you search a cluster name.

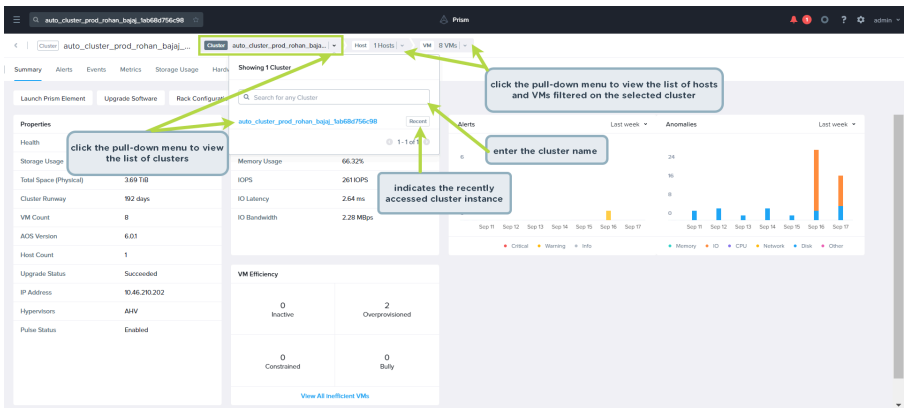

Figure 157: Clusters

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this cluster (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this cluster (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the cluster. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric below the tab. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities (see Behavioral Learning Tools on page 899). The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 288).

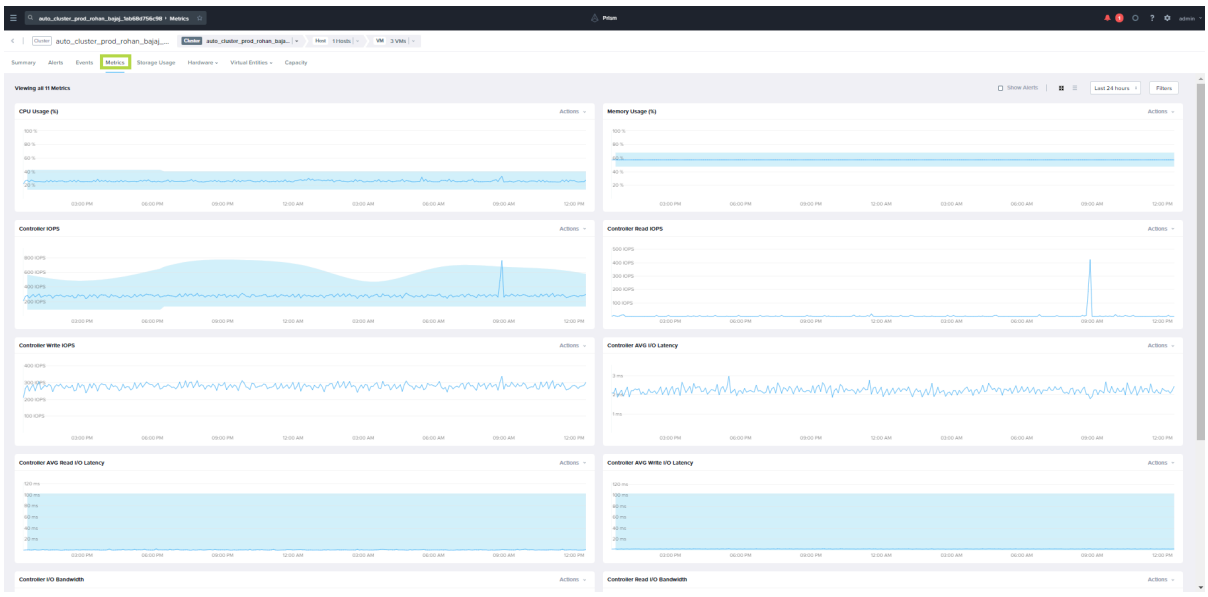

Figure 158: Cluster Metrics Tab

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

Table 85: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the cluster (0–100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the cluster (0–100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the cluster.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the cluster.
I/O Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the cluster.

Usage Tab

The **Usage** tab displays the following graphs:

- The **Cluster-wide Usage Summary** graph displays a rolling time interval monitor of total storage usage across the cluster that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of storage space used by each disk tier (SSD and DAS-SATA) across the cluster.

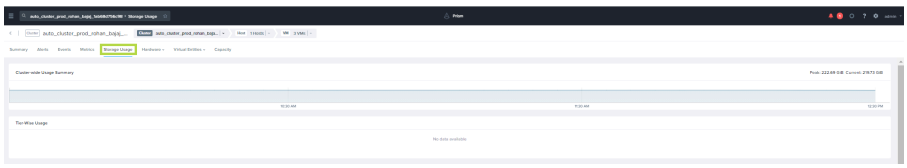

Figure 159: Cluster Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist in this cluster. Clicking a hardware entry such as **Hosts** or a virtual entity such as **Containers** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those in this cluster. See Hardware Entities on page 201 and Compute and Storage Entities on page 93 for more information.

Capacity Tab

The **Capacity** tab displays current and historical usage information and provides resource planning tools. It includes the following sections:

> **Note:**
>
> - The capacity planning feature requires a Prism Pro license. If Prism Pro is disabled, the **Capacity** tab is grayed out and not available.
> - Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
> - Capacity planning checks the past 90 days of data for monthly seasonality, it also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
> - For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance has been registered. See Application Monitoring on page 856 for more information.

- A **<cluster_name> Runway** area (upper left) that displays the current cluster runway overall and broken down by CPU, memory, and storage. *Runway* refers to how long the cluster can continue to run normally, based on the current consumption rate, before the existing resources are used to capacity. The overall runway value is the same as the lowest value for any single resource (storage, CPU, or memory). For example, if the storage runway is 14 days while the CPU and memory runways are 45 days, the overall runway is 14 days (the storage value).

  > **Note:** A plus sign in any runway value, for example the "89+" in the Cluster Capacity Tab figure, indicates that the predicted runway is more than that number of days, but the estimation stopped at that number.

  - Clicking **Storage Runway** displays a storage usage graph and table (to the right).
  - Clicking **CPU Runway** displays a CPU usage graph and table.
  - Clicking **Memory Runway** displays a memory graph and table.

- An **Optimize Resource** button (middle left). Clicking this button opens the **Optimize Resources** window, which identifies areas to improve cluster resource allocation or capacity. Recommendations might include
  - Table of "inefficient" (over-provisioned, inactive, constrained, and bully) VMs based on the VM behavioral learning engine (see Behavioral Learning Tools on page 899). Click the appropriate link to see a list of the VMs in that state.
  - An add node recommendation. The recommendation is based on the current (historical) CPU, memory, and storage usage demand across the cluster.

  Click the **Get Report** link to run the Cluster Efficiency report (see Reports Management on page 941).

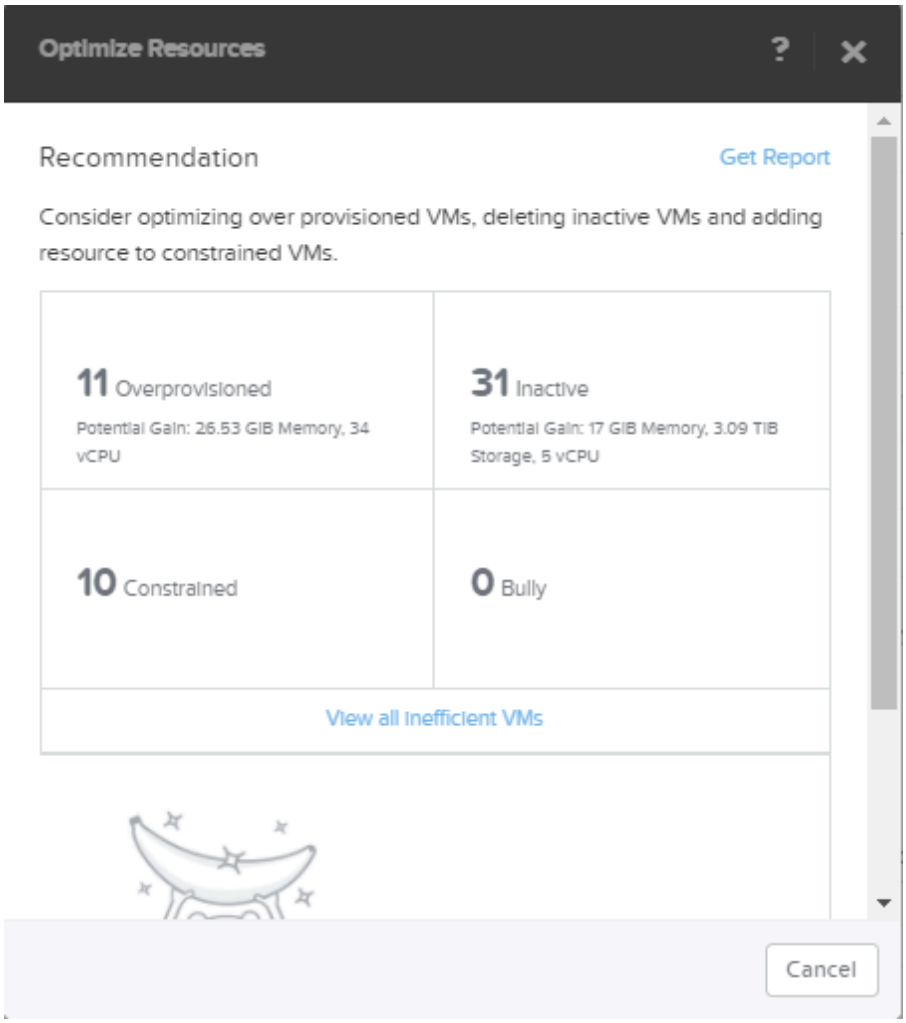

Figure 160: Optimize Resources Window

- A **Get Started** button (lower left). Clicking this button opens the resource planning scenario page. From this page you can create "what if" scenarios to test future workloads against current or hypothetical storage, CPU, and memory resource capacity (see Creating a Scenario on page 879).
- A usage history graph (upper right). The graph displays storage, CPU, or memory usage depending on which is selected in the left column. Placing the cursor anywhere on the horizontal axis displays the value(s) at that time. A (solid red) maximum capacity line appears near the top of the graph, and an additional (dotted red) effective capacity line appears on the storage graph. A (vertical) "Today" line is displayed to indicate where the current date is located on the graph. When there is an alert or event, a bell icon appears either on the Today line or at the

starting point of the alert. Clicking the bell icon (red for critical alert, orange for warning alert, or gray for event) displays more information about that alert or event.

- A table (lower right) that displays usage information. The information in the table varies depending on which runway (storage, CPU, or memory) is selected.

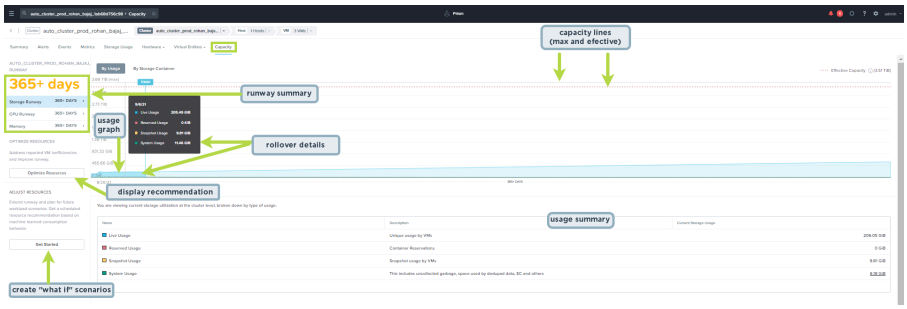

Figure 161: Cluster Capacity Tab (Nutanix Cluster)

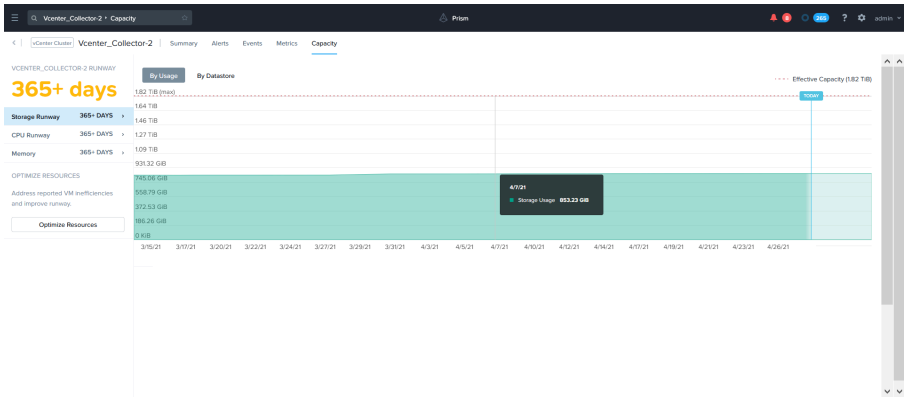

Figure 162: Cluster Capacity Tab (Non-Nutanix Cluster)

Capacity Tab: Storage View

When **Storage Runway** is selected, the graph and table display storage usage information. The maximum and effective storage capacity values (in TiBs) for the cluster appear above the graph. Effective capacity of a cluster is computed based on the configured reservations (from the capacity configurations settings) and the replication factor for all the containers in the cluster.

The storage view includes two tabs:

- Click the **By Usage** button to see cluster-wide storage information (see previous figure).

- Click the **By Storage Container** button to see storage information per storage container. You can select all storage containers from the pull-down list (upper left of graph) or a single storage container from the pull-down list or by clicking the storage container name in the table.

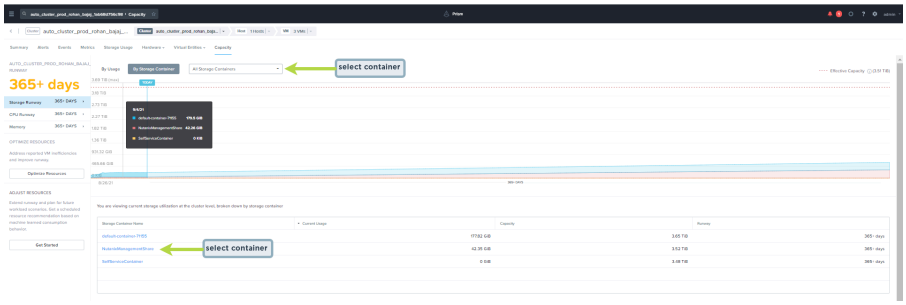

Figure 163: Storage Container View (Nutanix Cluster)

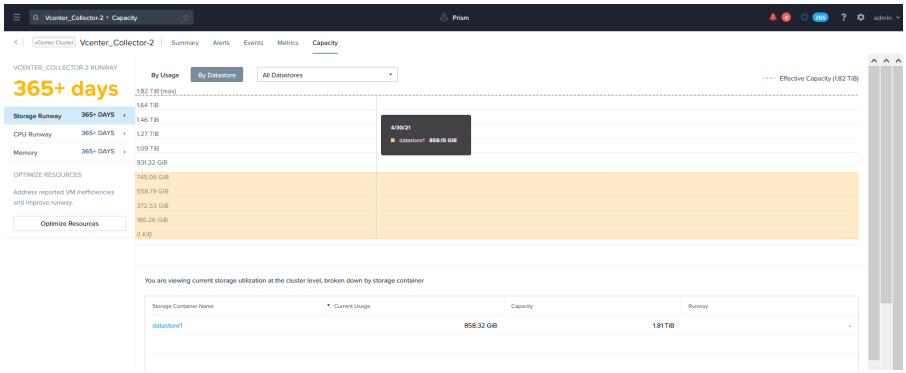

Figure 164: Storage Container View (Non-Nutanix Cluster)

The following table describes the fields in the storage table.

Table 86: Capacity Tab: Storage Fields

Parameter	Description	Values
"By Usage" tab (cluster-wide)		
Name	Displays a descriptive name for a type of storage. There are four types: • Live Usage: The amount of live storage. • Snapshot Usage: The amount of storage used for snapshots. • System Usage: The amount of storage for everything else (total usage - (reserved + snapshot + live)). System usage includes garbage, parity, and other miscellaneous uses. • Reserved Usage: The amount of storage reserved for thick provisioned VMs (ESXi).	Live Usage, Reserved Usage, Snapshot Usage, System Usage

Parameter	Description	Values
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB|TB]
"By Storage Container" tab: All Storage Containers		
Storage Container Name	Displays the name of the storage container.	(storage container name)
Current Usage	Displays the amount of storage space used by the storage container currently.	xxx [GB|TB]
Capacity	Displays the total amount of storage capacity available to the storage container.	xxx [TB]
"By Storage Container" tab: <storage container name>		
Storage Container Name	Displays one of the following: • If "All Storage Containers" is selected, the storage container names are displayed. Click a name to see the details for that storage container. • If a single storage container is selected, it displays the four storage types described for the Name field in the "By Usage" tab (Live Usage, Reserved Usage, Snapshot Usage, and System Usage).	(storage container name) OR Live Usage, Reserved Usage, Snapshot Usage, System Usage
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB|TB]

Capacity Tab: CPU View

When **CPU Runway** is selected, the graph and table display CPU usage information. The maximum and effective CPU capacity values (in GHz) for the cluster appear above the graph. The CPU view includes two tabs:

- Click the **Overall** button to see cluster-wide CPU information.
- Click the **By Host** button to see CPU information per host. You can select all hosts from the pull-down list (upper left of graph) or a single host from the pull-down list or by clicking the host name in the table.

The following table describes the fields in the CPU table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the CPU table.

Table 87: Capacity Tab: CPU Fields

Parameter	Description	Values
"Overall" tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM currently.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)
"By Host" tab: All Hosts		

Parameter	Description	Values
Host Name	Displays the name of the host.	(host name)
Average CPU Usage over last 1 hour	Displays the average host CPU usage during the last hour.	xxx [GHz]
Current CPU Capacity	Displays the CPU capacity of the host.	xxx [GHz]
Runway	Displays the current runway (estimated time interval that current CPU capacity is sufficient for the workload).	(number of days)
"By Host" tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)

Capacity Tab: Memory View

When **Memory Runway** is selected, the graph and table display memory usage information. The maximum and effective memory capacity values (in GiBs) for the cluster appear above the graph. The memory view includes two tabs:

- Click the **Overall** button to see cluster-wide memory information.
- Click the **By Host** button to see memory information per host. You can select all hosts from the pull-down list (upper left of graph) or a single host from the pull-down list or by clicking the host name in the table.

The following table describes the fields in the memory table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the memory table.

Table 88: Capacity Tab: Memory Fields

Parameter	Description	Values
"Overall" tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM.	(0–100%)
"By Host" tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average Memory Usage over last 1 hour	Displays the average host memory usage during the last hour.	xxx [GiB]

Parameter	Description	Values
Runway	Displays the current runway (estimated time interval that current memory capacity is sufficient for the workload).	(number of days)
Current Memory Capacity	Displays the memory capacity of the host.	xxx [GiB]
"By Host" tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM currently.	(0–100%)

Metrics Tab

The **Metrics** tab allows you to view performance metrics for the cluster, both for actual past performance and for predicted future performance. The **Metrics** tab displays a section on the left with a list of metrics.

> **Note:** The predicted metrics and identified anomalies are based on sophisticated machine-learning capabilities. See Behavioral Learning Tools on page 899 for information about these capabilities and how they are used.

- Clicking a metric displays a graph on the right. (Some metrics have multiple graphs.) The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph. Placing the cursor anywhere on the horizontal axis displays the value at that time. To set the time interval (last 24 hours, last week, last 21 days), select the duration from the pull-down list on the right.

  > **Note:** The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph may not appear if less than 21 days of data is available.

- To display predicted performance, check the **Prediction** box. This displays projected performance for the next week.
- To create an alert for this cluster based on either behavioral anomalies or status thresholds, click the **Set Alerts** link above the graph (see Creating Custom Alert Policies on page 288).

The following table describes the available metrics.

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

Table 89: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the VM (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the VM (0 - 100%).

Metric	Description
I/O Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.

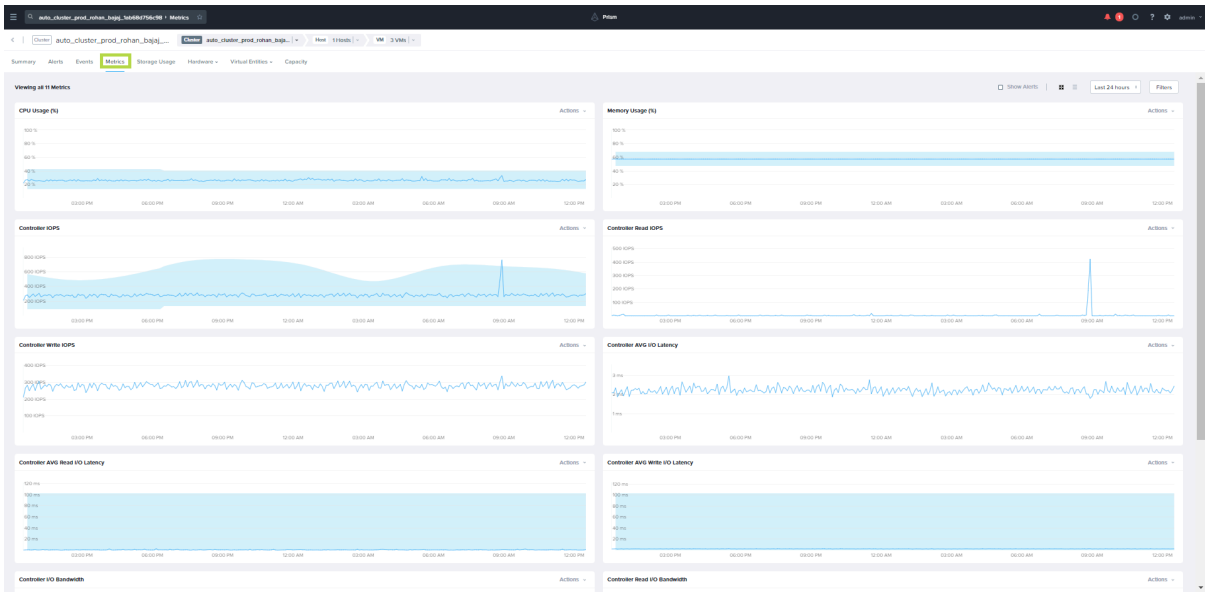

Figure 165: Metrics Tab: CPU Usage metric

Hosts Summary View

To access the hosts dashboard, select **Hardware** > **Hosts** from the entities menu (see Entities Menu on page 13). The hosts dashboard allows you to view summary information about hosts across the registered clusters and access detailed information about each host. The dashboard includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area to the right for the selected tab.

> **Note:** This section describes the information and options that appear in the Hosts dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

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

- **Suggested**: Displays a list of the hosts with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **CPU Usage**, **Memory Usage**, **IO Latency**, and **IOPS**. Click the **View all XX Hosts** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of host-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an

anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

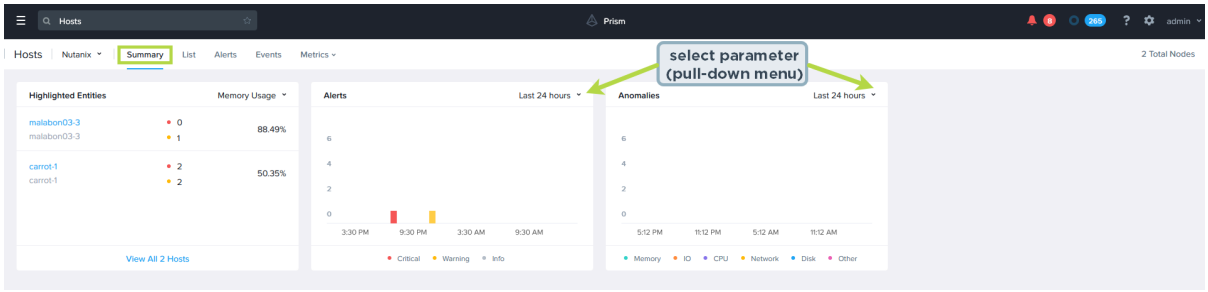

Figure 166: Hosts Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the hosts across the registered clusters. The following table describes the fields that appear in the hosts list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

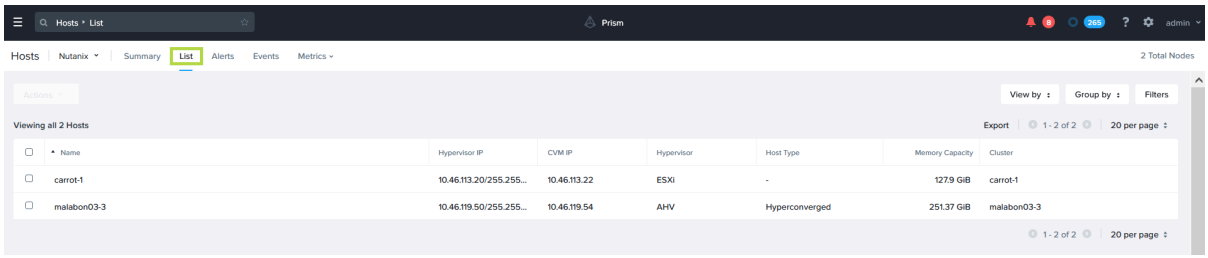

Figure 167: Hosts List Tab

Table 90: Hosts List Fields

Parameter	Description	Values
"General" Focus Fields		
Name	Displays the name of the host. Clicking on the name displays the details page for that host (see Host Details View on page 225).	(host name)
Host IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Hypervisor	Displays the hypervisor type running on the host.	AHV, ESX, or Hyper-V
Memory Capacity	Displays the memory capacity of the host.	xxx [MB|GB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)

Parameter	Description	Values
"Performance" Focus Fields		
Name	Displays the name of the host.	(host name)
CPU Usage	Displays the percentage of CPU capacity currently being used by this host.	0 - 100%
Memory Usage	Displays the percentage of memory capacity currently being used by this host.	0 -100%
IOPS	Displays I/O operations per second (IOPS) for this host.	[0 - unlimited]
Disk IO Bandwidth	Displays I/O bandwidth used per second for this host.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this host.	xxx [ms]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)

You can filter the hosts list based on a variety of parameter values. The following table describes the filter options available when you open the Hosts view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 91: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the host name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of hosts that satisfy the name condition/string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of hosts that reside in the clusters which satisfy the name string.	(cluster name string)
Health	Filters on the host health state (good, warning, or critical). Select one or more states to return a list of hosts in that state(s). The number of hosts currently in each state is displayed on the right of the line.	Critical, Warning, Good
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a pull-down list appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of hosts tagged to each selected category is displayed on the right of the line.	(category name)

Parameter	Description	Values
Hypervisor	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, HyperV
Memory Capacity	Filters on the host memory capacity. Check the box for the desired range or enter an amount range in the "**from <low> to <high> GiB**" field. It will return a list of hosts with memory capacity in that range.	([xx] to [yy] GiB range)
CPU Usage	Filters on the amount of CPU being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of hosts utilizing CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
GPUs	Filters for GPU configuration information such as model name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. As you type a pull-down list appears to help you select the correct configuration information. It will return a list of hosts that satisfy the GPU condition/string.	(configuration info)
IOPS	Filters on the IOPS. Check the box for the desired range or enter a range in the "**from <low> to <high> iops**" field. It will return a list of hosts with IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the "**from <low> to <high> bps**" field. It will return a list of hosts with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the "**from <low> to <high> ms**" field. It will return a list of hosts with average I/O latency in that range.	([xx] to [yy] range)

You can group the hosts list in the following ways:

- The **Group** pull-down menu allows you to group the host entries by cluster, hypervisor type, or health state. (You can only choose one.)
- The **View by** pull-down menu allows you to group the host entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

The **Actions** menu appears when one or more hosts are selected. It includes a **Manage Categories** option (see Assigning a Category on page 797).

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just host-related alerts across the registered clusters (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just host-related events across the registered clusters (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

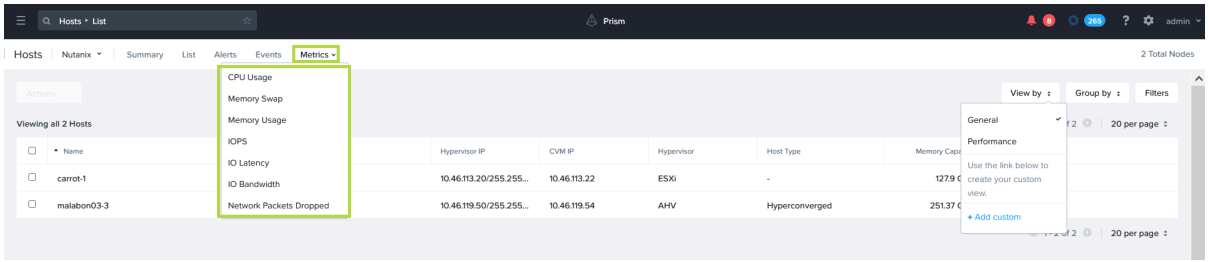

Figure 168: Hosts Metrics Tab

Table 92: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total hosts (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those hosts. **Note:** The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Swap	Displays memory swap-out and swap-in rate tables.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Host Details View

To access the details page for a host, go to the hosts dashboard **List** tab. See Hosts Summary View on page 221 and click the host name. You can also access the details page by clicking the host name wherever that name appear, such as in a dashboard widget or search result.

The host name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, **Usage**, **Hardware**, and **Virtual Entities**. Click a tab to display that information on the right. (Click the **Back to Hosts** link to return to the summary view.)

Summary Tab

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

- A **Properties** widget that displays summary information about the host (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.
- A **Manage Categories** button to manage (set) categories for the host. See Category Management on page 795.

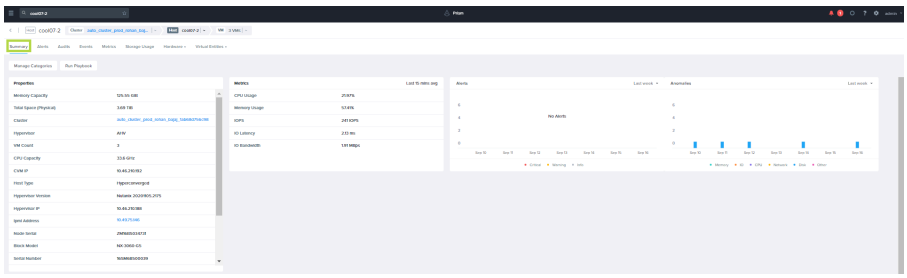

Figure 169: Host Summary Tab

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

Table 93: Host Properties Fields

Parameter	Description	Values
Memory Capacity	Displays the total memory capacity for this host.	xxx [MB|GB]
Disk Capacity	Displays the total amount of disk capacity on this host.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the host resides. Clicking the name displays the details page for that cluster (see Cluster Details View on page 209).	(cluster name)
Host IP	Displays the host IP address.	(IP address)
Hypervisor	Displays the hypervisor name.	(hypervisor name)
VM Count	Displays the number of VMs running on this host.	(number)
Block Model	Displays the block model number.	(model series number)

Parameter	Description	Values
Serial Number	Displays the block serial number.	(block serial number)
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Host Type		
IPMI Address	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console.	(IP address)
Node Serial	Displays the node serial number. The node serial is a unique number passed through from the manufacturer. (The form can vary because it is determined by each manufacturer.)	(manufacturer serial number)
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on the metadata disk.	[0 - 100%]
Oplog Disk Size	Displays the current size of the operations log. (The Oplog maintains a record of write requests in the cluster.) A portion of the metadata disk is reserved for the Oplog, and you can change the size through the nCLI.	xxx [GB]
Monitor Enabled	Displays whether the host is high availability (HA) protected. A **Yes** value means HA is active for this host. A **No** value means VMs on this host are not protected (will not be restarted on another host) if the host fails. Normally, this value should always be **Yes**. A **No** value is likely a sign of a problem situation that should be investigated.	[Yes|No]
Disks	Displays the number of disks in each storage tier in the host. Tier types vary depending on the Nutanix model type.	DAS-SATA: (number), SSD-SATA: (number), SSD-PCIe: (number)
GPUs	Displays the number and type of GPUs in the host. For example, if the host contains four Tesla M10 GPUs, this field displays "Tesla M10 (4)".	(GPU type and number)
Datastore(s)	Displays the names of any datastores.	(names)

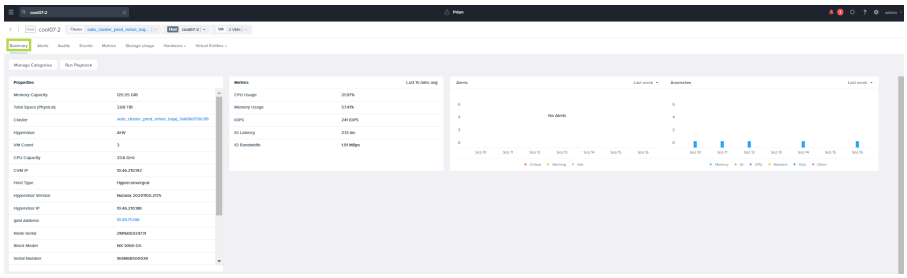

Figure 170: Host Summary Tab

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget provides quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective pull-down menus.

Example: Viewing Hosts on a Cluster

Click the **Host** pull-down menu to view the list of host instances on the selected cluster. Alternatively, you can search the host instance name belonging to the target cluster.

Figure 171: Hosts

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this host. See Alerts Summary View (Prism Central) on page 269 .

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this host. See Events Summary View (Prism Central) on page 279 .

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the host. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. See Behavioral Learning Tools on page 899. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.

- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Click the **Alert Settings** button to configure an alert for this metric. See Creating Custom Alert Policies on page 288.

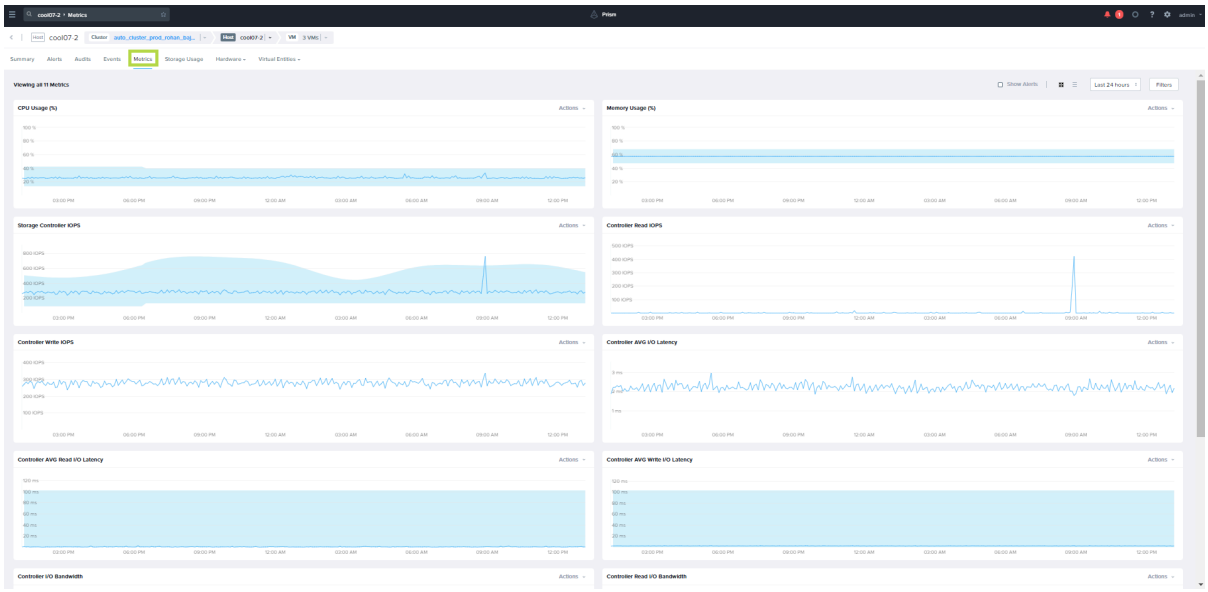

Figure 172: Host Metrics Tab

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

Table 94: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the host (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the host (0 - 100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the host.
IO Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the host.
IO Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the host.

Usage Tab

The **Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of host storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of host storage space used by each disk tier (SSD and DAS-SATA).

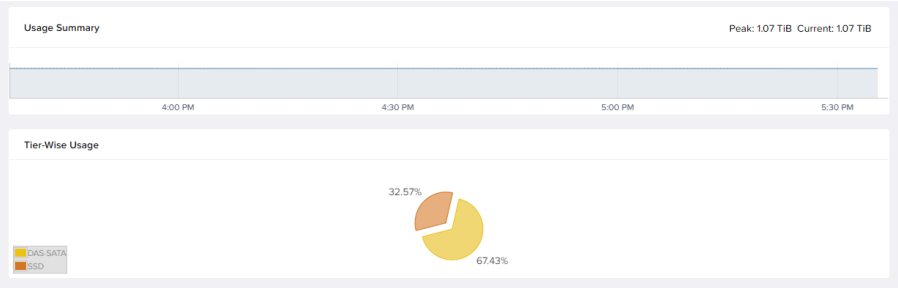

Figure 173: Host Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist on this host. Clicking a hardware entry such as **Disks** or a virtual entity such as **VMs** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those on this host. See Hardware Entities on page 201 and Compute and Storage Entities on page 93 for more information.

Disks Summary View

To access the disks dashboard, select **Hardware** > **Disks** from the entities menu (see Entities Menu on page 13). The disks dashboard allows you to view summary information about disks across the registered clusters and access detailed information about each disk.

> **Note:** This section describes the information and options that appear in the Disks dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.

Summary Tab

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

- **Suggested**: Displays a list of the disks with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Bandwidth**, **IOPS**, **IO Latency**, and **Disk Usage**. Click the **View all XX Disks** link at the bottom to display the **List** tab (following section).
- **Alert**: Displays a list of disk-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- **Anomalies**: Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

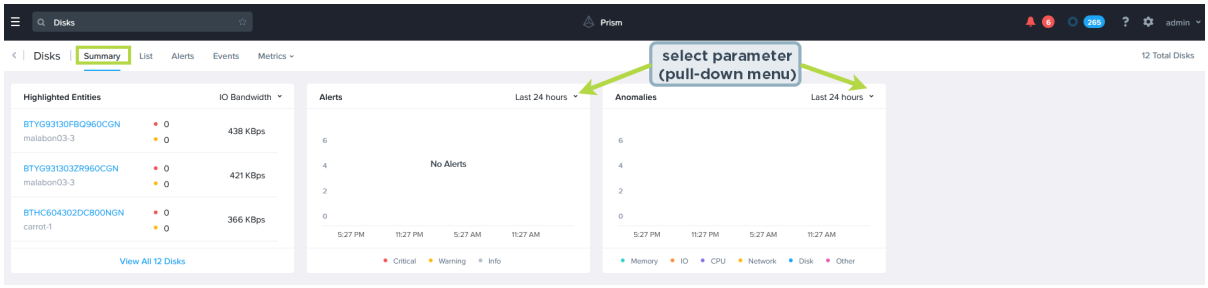

Figure 174: Disks Summary Tab

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the disks across the registered clusters. The following table describes the fields that appear in the disks list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

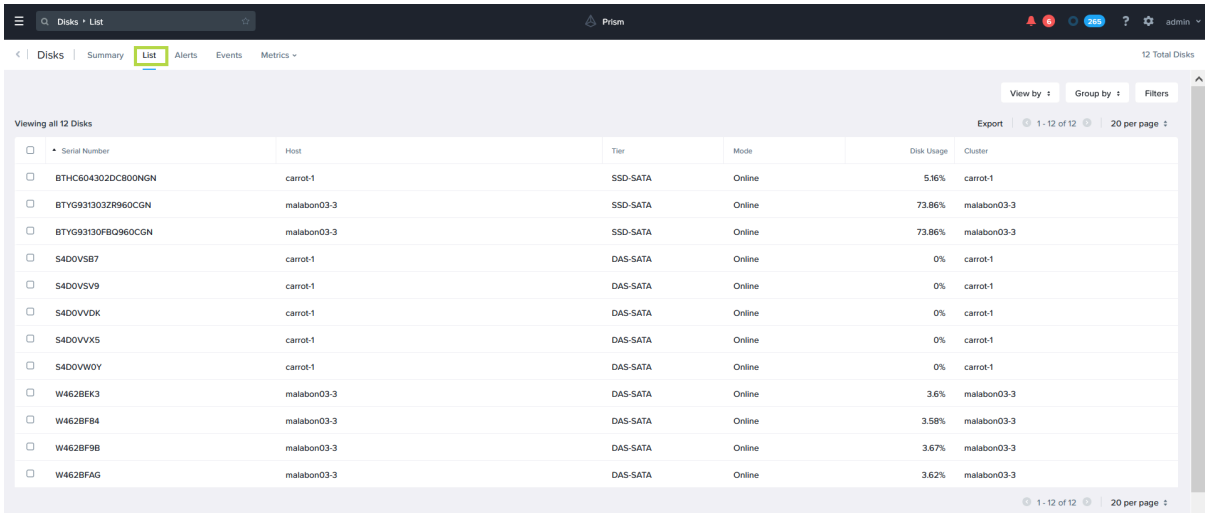

Figure 175: Disks List Tab

Table 95: Disks List Fields

Parameter	Description	Values
"General" Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Host	Displays the name of the host in which this disk resides.	(host name)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe | SSD-SATA | DAS-SATA]
Mode	Displays the operating state of the disk.	online, offline

Parameter	Description	Values
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB|TB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
"Performance" Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB|TB]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB|TB]
IOPS	Displays I/O operations per second (IOPS) for this disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for this disk.	xxx [MBps|KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this disk.	xxx [ms]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)

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

Table 96: Filter Pane Fields

Parameter	Description	Values
Serial Number	Filters on the disk serial number. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of disks that satisfy the serial number condition/string.	(serial number string)
Host	Filters on the host name. Enter a string in the field. It will return a list of disks in the hosts that satisfy the host name condition/string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of disks in the clusters that satisfy the cluster name condition/string.	(cluster name string)
Mode	Filters on whether the disk is online or offine. Check the box for one or both of these modes. The number of disks currently in each state is displayed on the right of the line.	Online, Offline

Parameter	Description	Values
Tier	Filters on whether the disk is in the solid state (SSD-SATA) or hard disk (DAS-SATA) tier. Check the box for one or both of these modes. The number of disks currently in each tier is displayed on the right of the line.	DAS-SATA, SSD-SATA
Health	Filters on the disk health state (good, warning, or critical). Select one or more states to return a list of disks in that state(s). The number of disks currently in each state is displayed on the right of the line.	Critical, Warning, Good
Disk Usage	Filters on the used capacity. Enter a percentage range in the "**from <low> to <high> %**" field. It will return a list of disks with used capacity in that range (0-100%).	([xx] to [yy]% range)
Disk Capacity	Filters on the total capacity. Enter an amount range in the "**from <low> to <high> GiB**" field. It will return a list of disks with total capacity in that range.	([xx] to [yy] GiB range)
IOPS	Filters on the IOPS. Enter a range in the "**from <low> to <high> iops**" field. It will return a list of disks with IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the I/O bandwidth used. Enter a range in the "**from <low> to <high> bps**" field. It will return a list of disks with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Enter a range in the "**from <low> to <high> ms**" field. It will return a list of disks with average I/O latency in that range.	([xx] to [yy] range)

You can group the disks list in the following ways:

- The **Color** pull-down menu allows you to color code the disk entries by tier type, mode, or health state. (You can only choose one.) A legend appears at the bottom to indicate what each color means in that grouping.
- The **Group** pull-down menu allows you to group the disk entries by host, tier type, mode, cluster, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** pull-down menu allows you to group the disk entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just disk-related alerts across the registered clusters (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just disk-related events across the registered clusters (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

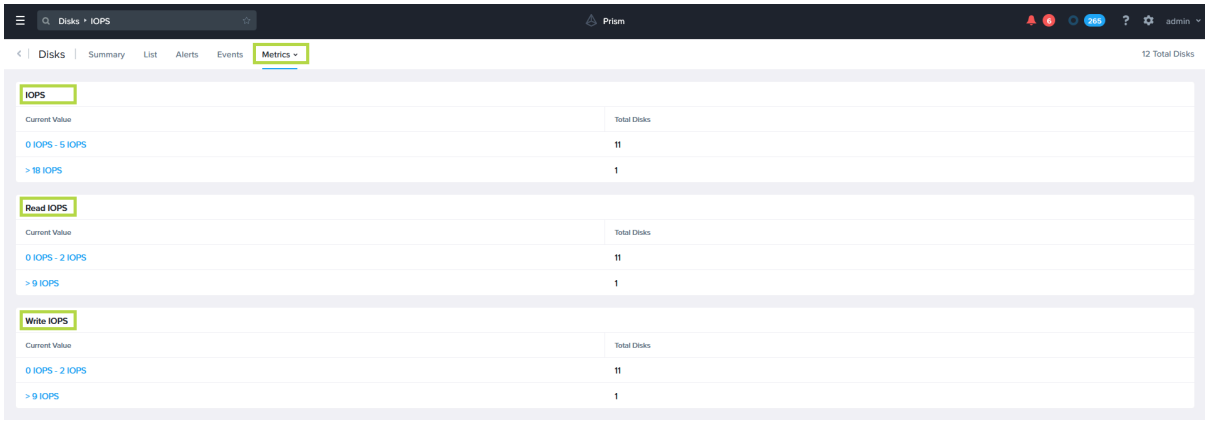

Figure 176: Disks Metrics Tab

Table 97: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS tables listing current values and total disks (number). The current values are split into intervals (for example, less than 20, 20-40, 40-60, more than 60). Clicking on an interval displays the Summary tab filtered to just those disk. **Note:** The same format also applies to the other metrics in this table.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Disk Details View

To access the details page for a disk, go to the disks dashboard **List** tab (see Disks Summary View on page 230) and click the disk name. You can also access the details page by clicking the disk name wherever that name appear, such as in a dashboard widget or search result.

The disk name and the following set of tabs appear on the left: **Summary**, **Alerts**, **Events**, **Metrics**, and **Usage**. Click a tab to display that information on the right. (Click the **Back to Disks** link to return to the summary view.)

Summary Tab

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

- A **Properties** widget that displays summary information about the disk (see following table).
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the pull-down menu.

When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

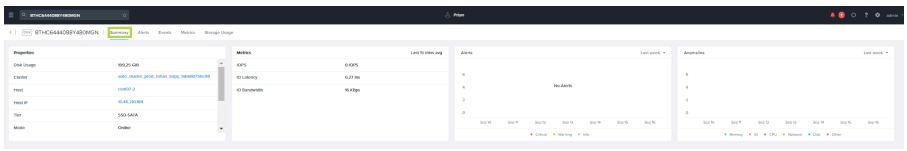

Figure 177: Disk Summary Tab

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

Table 98: Disk Properties Fields

Parameter	Description	Values
Disk Usage	Displays the amount of used space on the drive.	xxx [GB|TB]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)
Host	Displays the name of the host in which the disk resides.	(host name)
Host IP	Displays the IP address of the host.	(IP address)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe | SSD-SATA | DAS-SATA]
Mode	Displays whether the disk is currently online or offline.	[online|offline]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB|TB]
Status	Displays the operating status of the disk. Possible states include the following: • **Normal**. Disk is operating normally. • **Data migration initiated**. Data is being migrated to other disks. • **Marked for removal, data migration is in progress**. Data is being migrated in preparation to remove disk. • **Detachable**. Disk is not being used and can be removed.	Normal; Data migration initiated; Marked for removal, data migration is in progress; Detachable
Self Encryption Drive	Displays whether this is a self-encrypted drive.	Not Present, Present

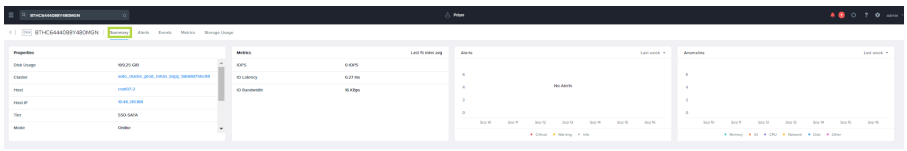

Figure 178: Disk Summary Tab

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this disk (see Alerts Summary View (Prism Central) on page 269).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this disk (see Events Summary View (Prism Central) on page 279).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the disk. Click the **Metrics** tab and then the desired metric name (**IOPS**, **IO latency**, and **IO Bandwidth**) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph.

> **Note:** The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the pull-down list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- Click the **Alert Settings** button to configure an alert for this metric (see Creating Custom Alert Policies on page 288).

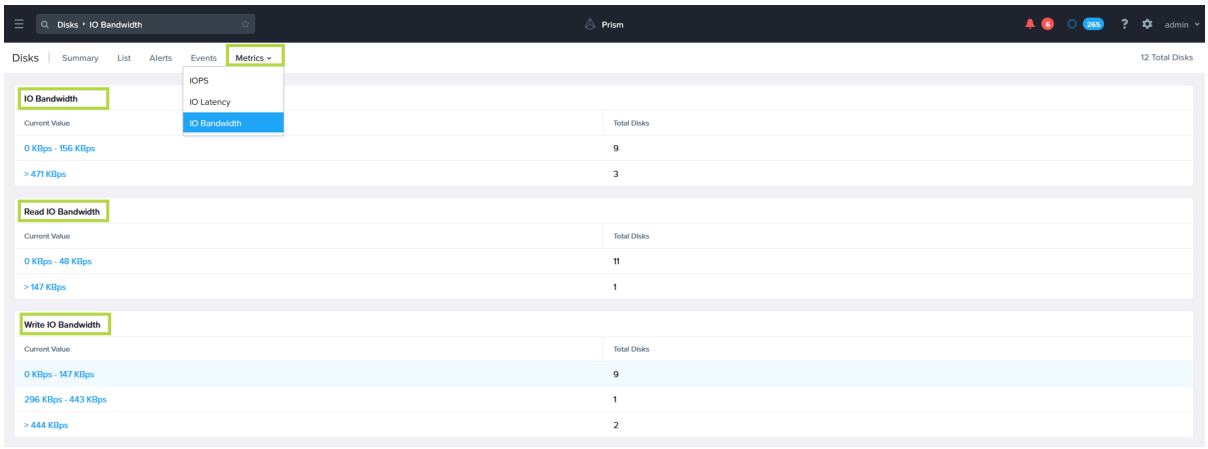

Figure 179: Disk Metrics Tab

Usage Tab

The **Usage** tab displays the following graph:

- The **Usage Summary** graph displays a rolling time interval monitor of disk storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

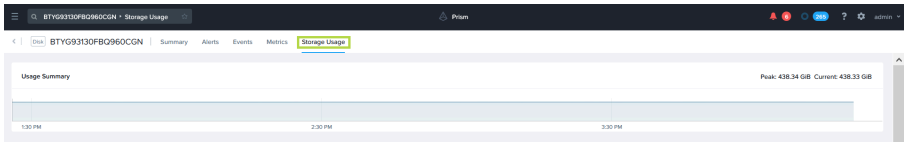

Figure 180: Disk Usage Tab

GPUs Summary View

To access the GPUs dashboard, select **Hardware** > **GPUs** from the entities menu (see Entities Menu on page 13). The GPUs dashboard allows you to view summary information about GPUs across the registered clusters and access detailed information about each GPU.

> **Note:** This section describes the information and options that appear in the GPUs dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.

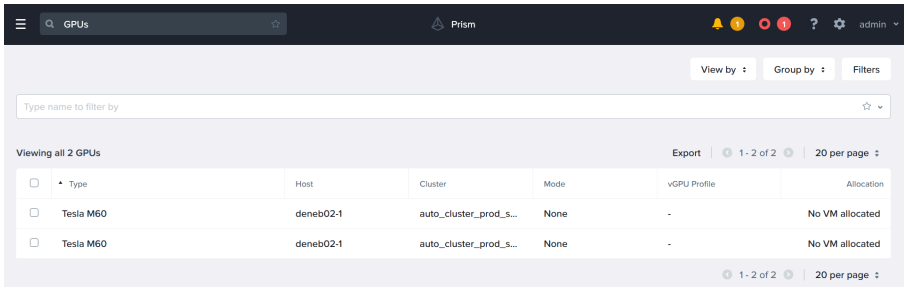

Figure 181: GPUs Dashboard

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

Table 99: GPU List Fields

Parameter	Description	Values
Type	Displays the GPU model type.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster	Displays the name of the cluster in which the GPU resides.	(cluster name)
Mode	Displays the mode in which the GPU is operating.	None, vGPU, passthrough

Parameter	Description	Values
Allocation	Indicates the number of VMs allocated to the GPU.	"No VM allocated", "*x* of *y* VMs allocated"

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

Table 100: Filter Pane Fields

Parameter	Description	Values
Type	Filters on the GPU model type. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of GPUs that satisfy the type condition/string.	Tesla M10, Tesla M60, Tesla M60 compute
Host	Filters on the host name. Enter a string in the field. It will return a list of GPUs in the selected hosts.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of GPUs in the selected clusters.	(cluster name string)
Mode	Filters on the GPU operation mode. Check the box for one or more of these modes. The number of GPUs currently in each mode is displayed on the right of the line.	None, vGPU, Passthrough

You can group the GPUs list in the following ways:

- The **Group** pull-down menu allows you to group the GPU entries by cluster, host, or mode. (You can only choose one.)
- The **Sort** pull-down menu allows you to group the GPU entries by type, cluster, mode, or allocation. (You can only choose one parameter.)

GPU Details View

To access the details page for a GPU, click on the desired GPU type entry in the list (see GPUs Summary View on page 237). You can also access the details page by clicking the GPU name wherever that name appear, such as in a dashboard widget or search result. The GPU details page includes **Summary** and **VMs** tabs.

Summary Tab

Clicking the **Summary** tab, which appears by default, displays the following:

- A section on the left that displays summary information about the GPU (see following table).

- A section of the right that displays GPU performance metrics. The graphs are rolling time interval performance monitors that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. This section includes the following graphs:
  - **GPU Usage**: Displays the percentage of GPU capacity being used.
  - **GPU Framebuffer Usage**: Displays the percentage of GPU framebuffer (RAM) capacity being used.
- Action button on the upper right (which appears regardless of which tab is selected):
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

The following table describes the GPU summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 101: GPU Summary Fields

Parameter	Description	Values
GPU Type	Displays the GPU type for this entry.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster Name	Displays the name of the cluster in which the GPU resides.	(cluster name)
Host	Displays the name of the host in which the GPU resides.	(host name)
Mode	Displays the GPU operational mode. If it is vGPU, an additional field appears that displays the vGPU profile used.	None, vGPU, Passthrough
Allocation	Displays the number of VMs allocated to this GPU. This field does not appear if no VMs are allocated.	"*x* of *y* VMs allocated"
Framebuffer (RAM) Per GPU	Displays the framebuffer (RAM) size per GPU.	xx GiB
ID	Displays the GPU ID number.	(ID number)

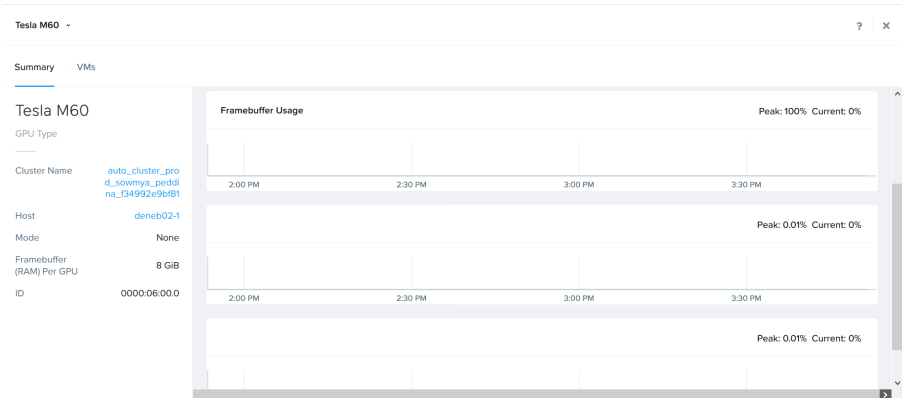

Figure 182: GPU Summary Tab

VMs Tab

Clicking the **VMs** tab displays a table of VMs allocated (attached) to the GPU. The table includes the following fields:

- **Name**: Displays the VM name. Click the name to display the details page for that VM (see VM Details View on page 105).
- **GPU Usage**: Displays the percentage of GPU capacity used by this VM.
- **GPU Framebuffer Usage**: Displays the percentage of GPU framebuffer (RAM) capacity used by this VM.

Activity Entities

You can access dashboards for the following activity monitors from the **Activity** category of the entities menu (see Entities Menu on page 13):

- Alerts (see Alerts Summary View (Prism Central) on page 269)
- Events (see Events Summary View (Prism Central) on page 279)
- Audits (see Audits Summary View on page 240)
- Tasks (see Tasks View on page 244)

Audits Summary View

To access the audits dashboard, select **Activity** > **Audits** from the entities menu (see Entities Menu on page 13). The audits dashboard allows you to view a list of actions performed across the registered clusters.

Note:

- This section describes the information and options that appear in the audits dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
- Audit information appears only for those registered clusters running AOS 5.10 or later.
- Audit logs with default values are generated when updates to VMs are initiated, either by Prism Central **Self Service** users or by using Nutanix v3 API calls for the first time.
- The retention period for audit entries is four weeks by default.
- See Prism Central Logs on page 302 for all Prism Central related events. Additionally, you can refer to the Audit Log Events on page 305 for the complete list of captured audit events.

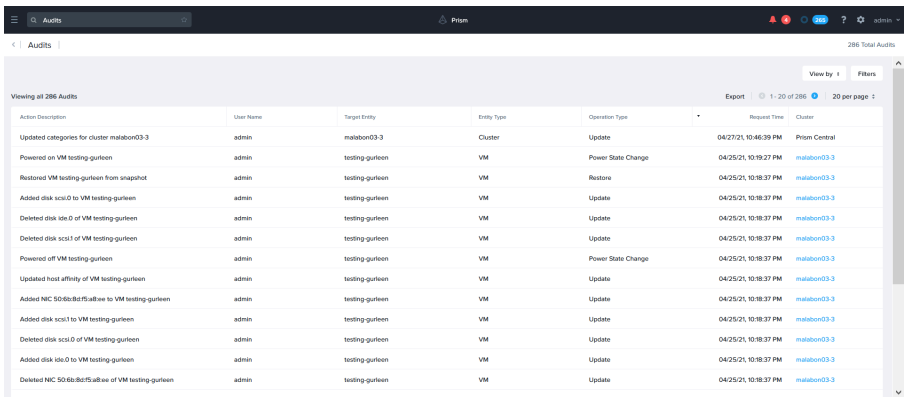

Figure 183: Audits Dashboard

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

Table 102: Recovery Plans Fields

Parameter	Description	Values
Action Description	Describes the action taken such as "deleted VM `vm-name`" or "added disk `disk-name`"	(action description)
User Name	Displays the name of the user who requested the action.	(user name)
Target Entity	Displays the entity name. Click the name to go to the details page for that entity.	(entity name)
Entity Type	Displays the entity type such as VM or host.	(entity type)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)

To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. The following table describes the filter options available. You can apply multiple filters together. For example, if you want to filter by a particular VM name, you can filter by *entity type = VM* and then use the *name contains* filter.

Table 103: Filter Pane Fields

Parameter	Description	Values
User Name	Enter a name string in the field to filter for users who requested an action.	(user name)

Parameter	Description	Values
Entity Type	Check the boxes of one or more entities to filter for actions on those entity types.	VM, Storage Container, Catalog Item, Image, Cluster, Host, Disk, GPU, Security Policy, NGT Policy, Project, Role, User, Category, Availability Zone, Protection Policy, Recovery Plan, Recoverable Entity, Report
Operation Type	Check the boxes of one or more operations to filter on those operations.	Create, Update, Delete, Power State Change
Request Time	Check an interval box to filter for actions that were requested during that time period. For the custom interval option (**from** *xxx* **to** *xxx*), click in each field and select a date from the pop-up calendar.	Last 1 hour, Last 24 hours, Last week, From *xxx* to *xxx*
Cluster	Enter a cluster name in the field to filter for actions in the cluster. **Note:** Entering *Prism Central* in the cluster filter criteria does not populate any result.	(cluster name)
User IP	Enter a user IP address and then click the Add button to filter for actions requested by that user. You can add multiple user IP addresses.	(IP address)

Audit Details View

To access the details page for an action, go to the audits dashboard (see Audits Summary View on page 240) and click the action description. The details page includes the following:

- The action description (upper left). You can switch from one action details page to another by selecting from the pull-down list.
- A section on the left that displays summary information about the action (see following table).
- A section on the right that displays a table of information specific to that action. The table shows the attributes that were changed during the action and the current value that is set for the attribute. The attributes vary depending on the specifics of the action (see following examples).

The following table describes the action summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 104: Action Summary Fields

Parameter	Description	Values
Action Description	Displays the action description.	(description)
User Name	Displays the name of the user who requested the action.	(user name)

Parameter	Description	Values
Target Entity	Displays the name of the entity that was the action target. Click the entity name to display the details for that entity.	(entity name)
Affected Entities	Displays the names of the entities that were affected by the action. Click an entity name to displays the details page for that entity.	(one or more entity names)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
User IP	Displays the IP address of the user.	(IP address)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)
Status	Displays the status of the action.	Succeeded, Failed

In the first example the action is adding more categories for the cluster, so the details simply include the names of the attached categories.

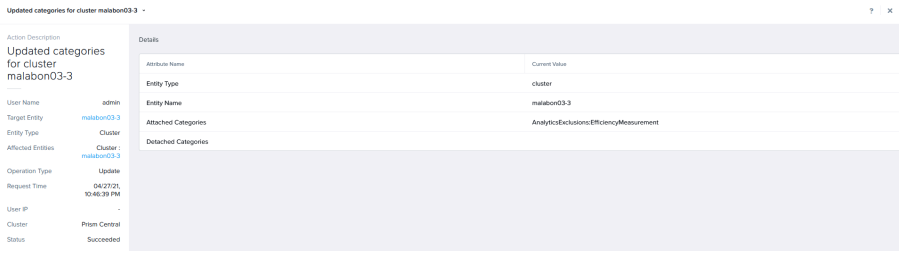

Figure 184: Audit Details Page (example 1)

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

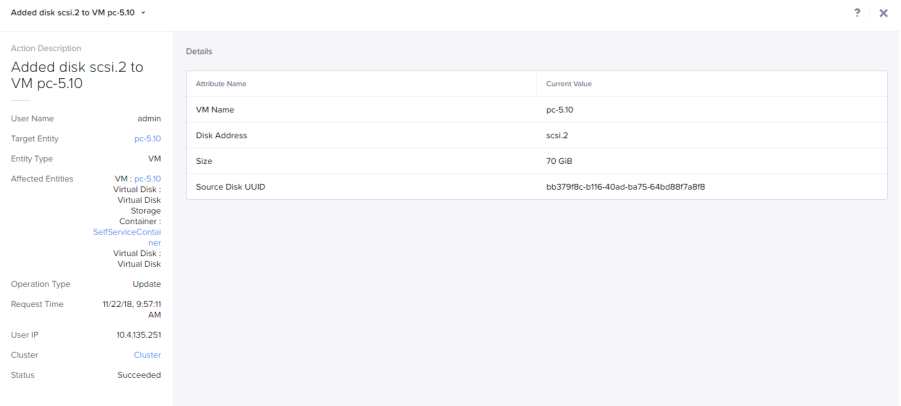

Figure 185: Audit Details Page (example 2)

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

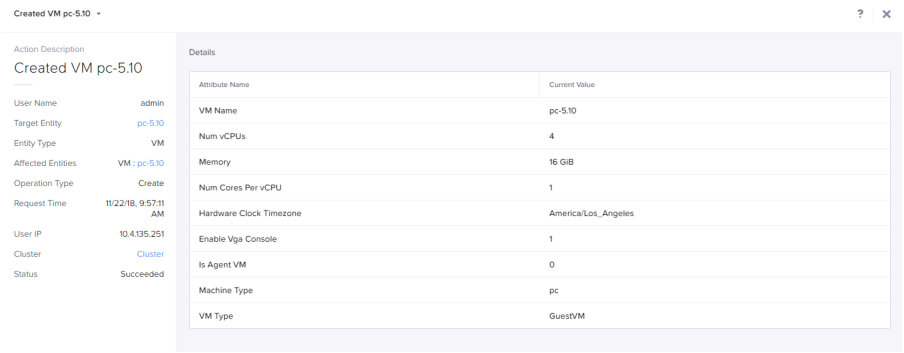

Figure 186: Audit Details Page (example 3)

Tasks View

An icon appears in the main menu when there are current tasks (running or completed within the last 48 hours). The icon appears blue when a task runs normally, yellow when it generates a warning, or red when it fails. Clicking the icon displays a drop-down list of current tasks. You can see the Tasks dashboard, which displays information about all tasks across the registered clusters, by doing one of the following:

- Clicking the **View All Tasks** link at the bottom of the current tasks drop-down list.
- Clicking the **View All Tasks** link in the Tasks widget in the main dashboard (see Main Dashboard on page 70).
- Selecting **Activity** > **Tasks** in the entities menu (see Entities Menu on page 13).

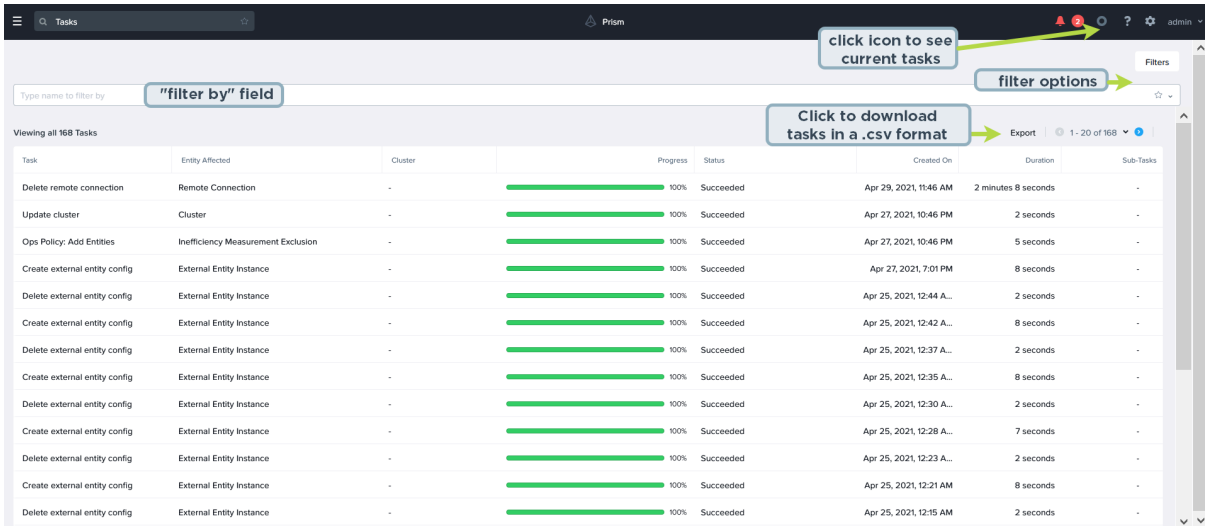

Figure 187: Tasks Screen

You can do the following in the Tasks dashboard:

- Filter the list by entering a name in the "filter by" field.
- Filter the list by clicking the **Filters** button. This displays a filter pane on the right of the screen. Click on one or more of the states to filter the list to just tasks in those states. The states are Aborted, Canceled, Failed, Queued,

Running, Paused, Skipped, Succeeded, and Scheduled. The number of tasks in each state currently appear to the right of the state name. Click the X icon (upper right) to close the filter pane.

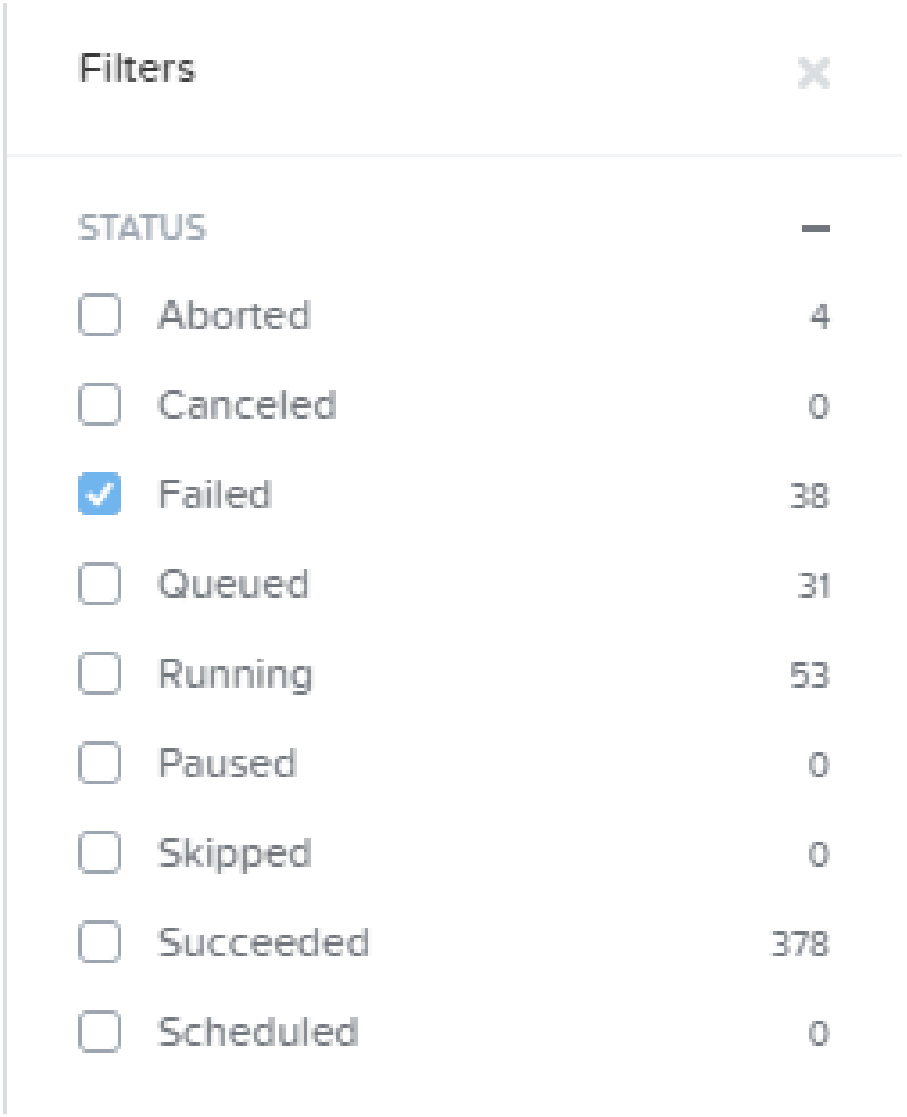

Figure 188: Filter Pane

- Download the task list in CSV format by clicking the download icon (on the right above the table).

The following table describes the fields in the tasks list. Each task appears in the list for a minimum of one hour after completion, but how long that task remains in the list depends on several factors. In general, the maximum duration is two weeks. However, tasks are rotated off the list as new tasks arrive, so a task might disappear from the list much sooner when activity is high. In some cases a task appears for longer than two weeks because the last task for each component is retained in the listing.

Table 105: Tasks List Fields

Parameter	Description	Values
Task	Specifies which type of operation the task is performing.	(name of operation such as "upgrade")

Parameter	Description	Values
Entity Affected	Lists the entity name or describes the operation. If this is a live link, click it to either see more details or go to the entity details page.	(name or description)
Cluster	Specifies the cluster name(s) on which the task is running. A dash (-) appears when the task applies to Prism Central rather than one of the registered clusters.	(cluster names)
Progress	Indicates the current percentage complete for the task.	0%-100%
Status	Indicates the task status, which can be pending, running, completed, or failed.	Succeeded, Running,Failed, Queued
Created On	Displays when the task began.	xx [seconds| minutes| hours| days] ago
Duration	Displays how look the task has been running or took to complete.	xx [seconds| minutes| hours| days]
Sub-Tasks		

Operations Entities

You can access dashboards for the following tools from the **Operations** category of the entities menu (see Entities Menu on page 13):

- Analysis: Analyze cluster activity (see Analysis Dashboard (Prism Central) on page 799)
- Cost Management: Analyze and manage costs (see Cost Management (Xi Beam) on page 855)
- App Discovery: Discover applications running in a cluster (see Application Discovery on page 831)
- Monitoring Configurations: Monitor select applications in a cluster (see Application Monitoring on page 856)
- Operations Policies: Manage operations policies (see Operations Policy Management on page 873)
- Planning: Plan for resource capacity and usage (see Resource Planning on page 876)
- Playbooks: Create playbooks for task automation (see Task Automation on page 902)
- Reports: Generate resource and activity reports for Nutanix and Non-Nutanix environments (see Reports View on page 246)

Reports View

The Reports dashboard allows you to create and run reports about the Nutanix and vCenter entities and metrics of your choice. To access the Reports dashboard, select **Operations** > **Reports** from the entities menu (see Entities Menu on page 13).

> **Note:** This section describes the information and options that appear in the Reports dashboard.
>
> - See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
> - See Reports Management on page 941 for detailed information on how to create, configure, and run reports.

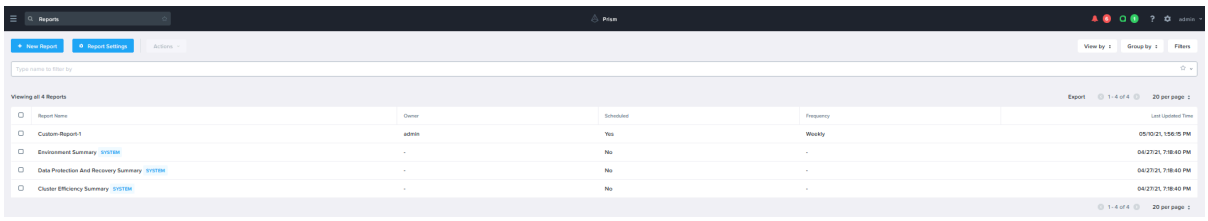

Figure 189: Reports Dashboard

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

Table 106: Reports List Fields

Parameter	Description	Values
Report Name	Displays the name of the report. **Note:** The **SYSTEM** label appears against the report name for pre-defined reports.	(report name)
Owner	Displays the name of the user who created the report.	(user name)
Scheduled	Indicates whether this report is scheduled to run at specified times. • "Yes" means the report is scheduled to run automatically at specified times (see Frequency). • "No" means the report is not scheduled to run. To run the report you must either manually run it or update the report to specify a schedule.	Yes, No
Frequency	Displays when (how frequently) the report is scheduled to run.	Daily, Weekly, Monthly, Yearly
Last Updated Time	Displays the last updated data and time of the report.	mm/dd/yy, hh:mm:ss AM/PM

You can filter the reports by report name. To apply a filter, click the **Filter** button to open the **Filter** pane. Enter a string in the report name field. It will return a list of reports that satisfy the name string.

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

The Reports view includes two action buttons: **+ New Report** and **Report Settings**.

- Click the **New Report** button to create a new custom report (see Creating a New Report on page 943).
- Click the **Report Settings** button to configure the report appearance, e-mail, and retention settings (see Configuring Report Settings on page 988).

The **Actions** menu appears when one or more reports are selected. It includes the following actions (see Managing a Report on page 944):

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

The available actions appear in bold; other actions are grayed out. (For grayed out options, a tool tip explaining the reason is provided.) The available actions depend on the selected report.

Administration Entities

You can access dashboards to manage the following objects from the **Administration** category of the entities menu (see Entities Menu on page 13):

- Users (see Users Summary View on page 248)
- Roles (see Roles Summary View on page 252)
- Projects (see Projects Summary View on page 256)
- Availability Zones (see Availability Zones View on page 266)
- LCM (see LCM View on page 267)

Users Summary View

To access the users dashboard, select **Administration** > **Users** from the entities menu (see Entities Menu on page 13). The users dashboard allows you to view a list of local and Active Directory-specified users.

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

- See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.
- See *Managing Local User Accounts* in Security Guide for information about creating and maintaining local user accounts.

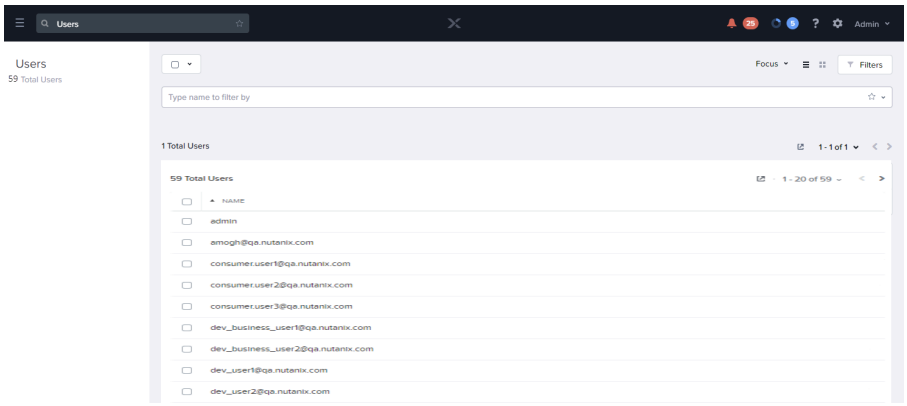

Figure 190: Users Dashboard

The Users view displays a list of local users and project-assigned Active Directory users. You can order the list alphabetically (A-Z or Z-A) by clicking (toggling) the **Name** column header. Clicking on a user name displays the details page for that user.

From the**Filter** pane you can filter the users list by name. Check the name box, select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**), and enter a string in the field. It will return a list of users that satisfy the name condition/string.

User Details View

To access the details page for a user, go to the users dashboard (see Users Summary View on page 248) and click the user name. The details page includes the following:

- The user name (upper left). You can switch from one user details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Trends**, **VM**, **Project**) to display that tab contents below.

Summary Tab

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

- A section on the left that displays summary information about the user (see following table).
- A section of the right that displays vCPU, storage, and memory usage graphs for the top five VMs for this user. Place the cursor in a metric graph to see which VM is using that portion of the resource (vCPU, storage, or memory).

The following table describes the project summary information fields. A dash (-) is displayed in a field when a value is not available.

Table 107: User Summary Fields

Parameter	Description	Values
Display Name	Displays the name that appears in the Prism Central displays. This might be a truncated version of the name (next parameter) when that name is long.	(user display name)
Name	Displays the user name.	(user name)
Projects Count	Displays the number of projects to which this user is a member.	(number)
Total VMs	Displays the number of VMs assigned to this user.	(user name)
vCPU Usage	Displays the number of vCPUs used by this user.	(number)
Memory Usage	Displays the amount of memory used by this user.	xxx [GiB]
Storage Usage	Displays the amount of storage used by this user.	xxx [GiB]

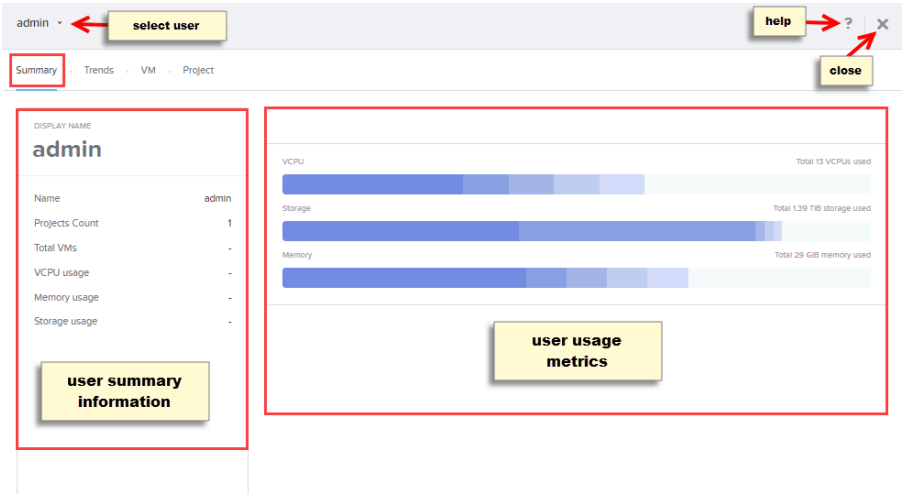

Figure 191: User Summary Tab

Trends Tab

The **Trends** tab displays the following graphs.

- The Total VMs graph displays a rolling time interval monitor of VMs owned by this user. Placing the cursor anywhere on the horizontal axis displays the number at that time.
- The **vCPU Usage** graph displays a rolling time interval monitor of user vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of user memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of user storage usage.

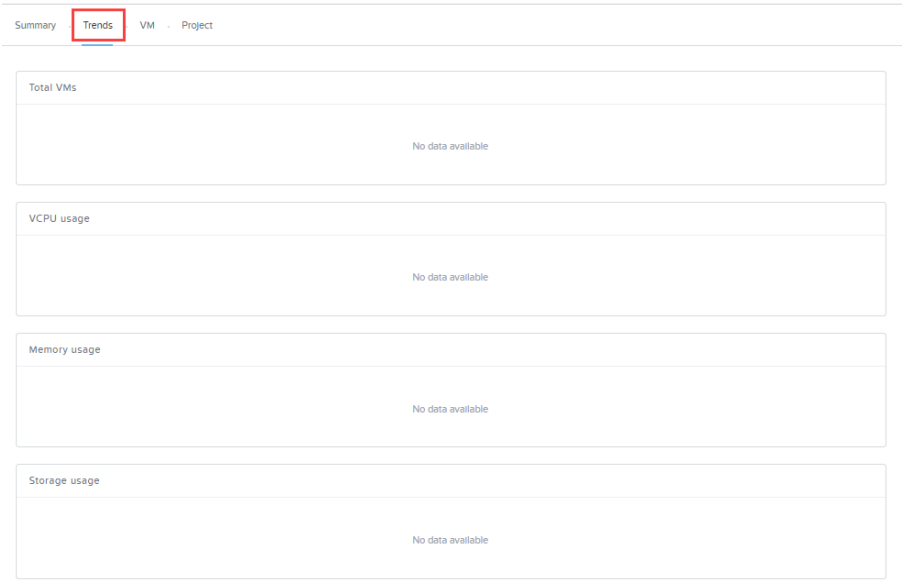

Figure 192: User Trends Tab

VM Tab

The **VM** tab displays a table of VMs owned by the user. You can filter the table list in the table by entering a string in the search field (upper right above the table). The following table describes the VM information fields.

Table 108: User VM Fields

Parameter	Description	Values
Name	Displays the name of the VM.	(VM name)
Project Name	Displays the name of the project associated with this VM.	(project name)
Virtual CPU Count	Displays the number of vCPUs allocated to the VM.	(number)
Memory Capacity	Displays the amount of memory capacity allocated to the VM.	xxx [GiB]
Disk Capacity	Displays the amount of disk capacity allocated to the VM.	xxx [GiB]

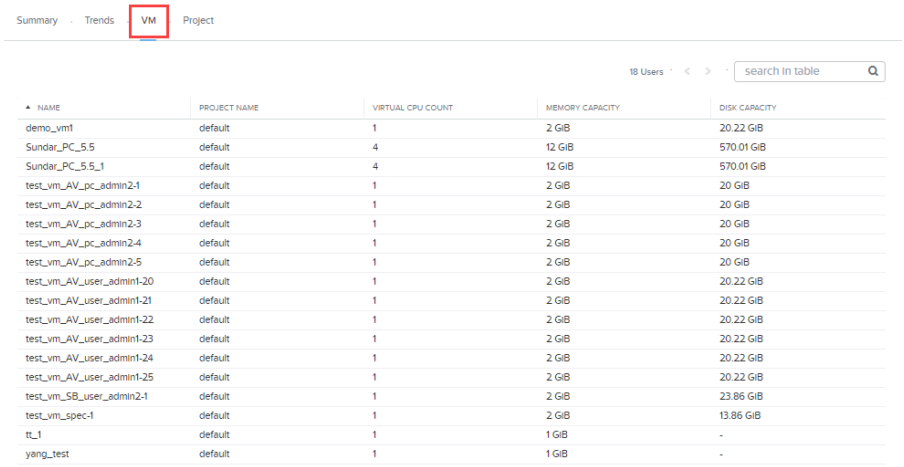

Figure 193: User VM Tab

Project Tab

The **Project** tab displays a table of projects to which the user is a member. You can filter the list by entering a string in the search field (upper right above the table). The following table describes the users information fields.

Table 109: User Project Fields

Parameter	Description	Values
Name	Displays the name of the project.	(project name)
vCPU Usage	Displays the number of vCPUs used by this project.	(number)
Memory Usage	Displays the amount of memory used by this project.	xxx [GiB]

Parameter	Description	Values
Storage Usage	Displays the amount of storage used by this project.	xxx [GiB]

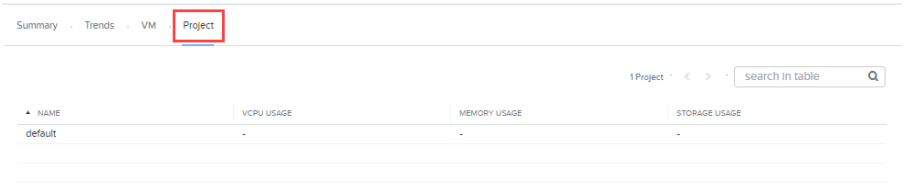

Figure 194: User Project Tab

Roles Summary View

To access the roles dashboard, select **Administration** > **Roles** from the entities menu (see Entities Menu on page 13). The roles dashboard allows you to view summary information about built-in (default) and custom (created) roles and to access detailed information about each role.

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

- See Entity Exploring on page 89 for instructions on how to view and organize this information in a variety of ways.
- See Controlling User Access (RBAC) in Security Guide for information on how to create, manage, and apply roles.

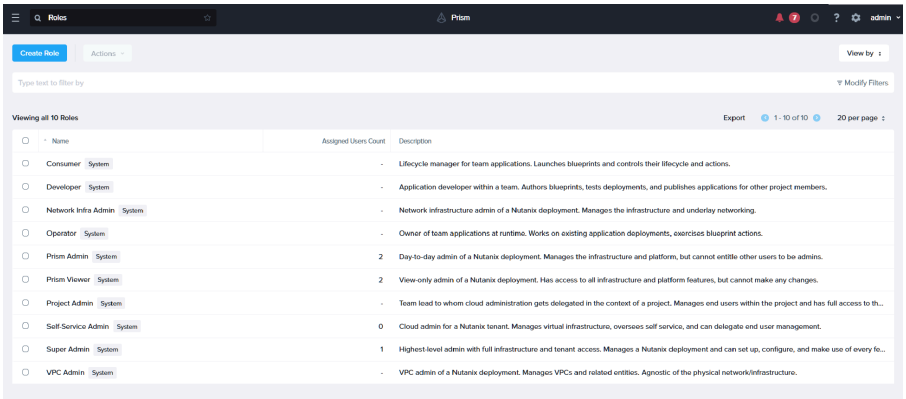

Figure 195: Roles Dashboard - View by General

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

Table 110: Roles List Fields

Parameter	Description	Values
View by General		
Name	Displays the role name.	(role name)

Parameter	Description	Values
Assigned Users Count	Displays the number of users assigned this role.	(number)
Description	Describes the role (if a description was provided when the role was created or updated).	(number)

You can create other custom views by adding other data columns, in addition to the **Name** column, to such custom views.

- Click **View by** > **+ Add Custom** to open the **Role Columns** window.
- Enter a name for your custom view, click the **+** icon for each column that you want to see in your custom view, and click **Save**.

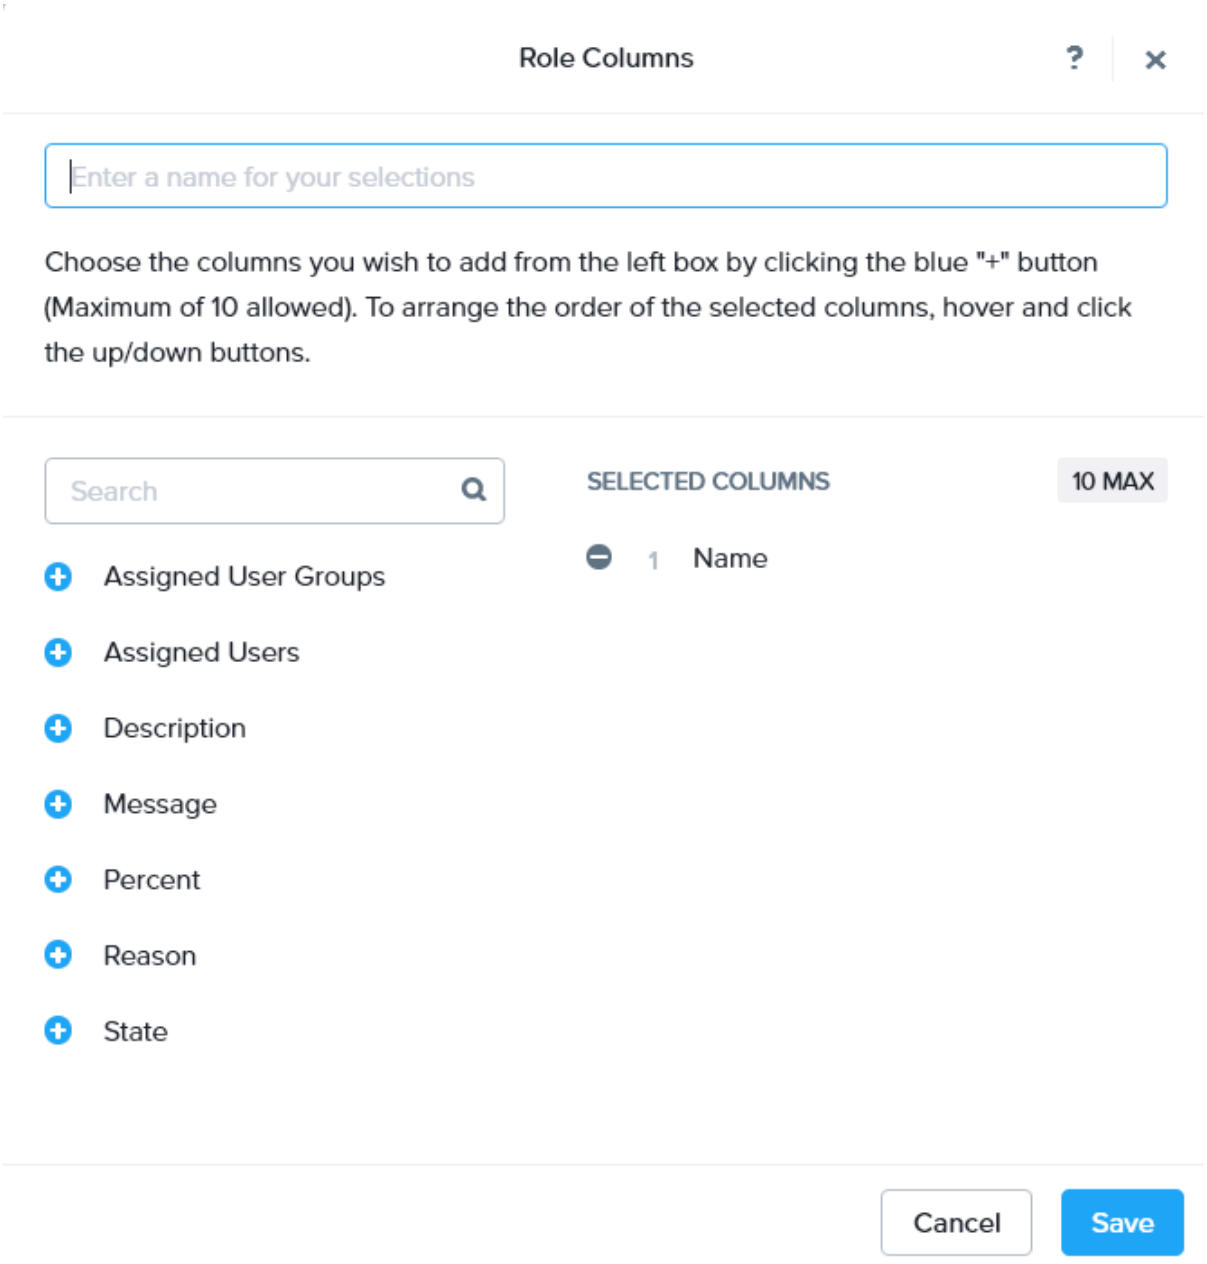

Figure 196: Roles List Fields - Custom Views

Table 111: Other Roles List Fields

Parameter	Description	Values
View by <custom>		
Assigned User Groups	Displays the number of user groups assigned this role.	(number)
Message	Displays specified message for this role.	(string)
Percent	Displays the percentage of users.	(number up to maximum of 100 or "-")
Reason	Displays the reason.	(string)
State	Displays the state of the role.	(string)

You can filter the roles list based on several parameter values. The following table describes the filter options available when you open the roles dashboard **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 112: Filter Pane Fields

Parameter	Description	Values
Name	Filters on the role name. Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field. It will return a list of role that satisfy the name condition/string.	(project name string)
Assigned Users Count	Filters on the number of assigned users. Enter a range in the "**from <low> to <high>**" field. It will return a list of roles that have an assigned user count within that range.	([xx] to [yy] range)

The dashboard includes a **Create Role** button (see Creating a Custom Role). The **Actions** menu appears when one or more roles are selected. It allows you to manage assignments for, duplicate, update, or delete a role (see Role Details View on page 254).

> **Note:** You cannot delete or modify the default roles.

Role Details View

To access the details page for a role, go to the roles dashboard (see Roles Summary View on page 252) and click the role name. The details page includes the following:

- The role name (upper left). You can switch from one role details page to another by selecting from the pull-down list.
- A set of tabs (upper left). Click the tab (**Summary**, **Users**, **User Groups**, **Role Assignment**) to display that tab contents below.

- Action buttons (upper right). The available actions appear in bold; other actions are grayed out. For more details, see the relevant topic in Security Guide.
  - **Manage Assignment**: Click this button to assign the role to users, groups, and entities (see Assigning a Role). This button does not appear for certain default roles such as Prism Admin that you are not allowed to assign.
  - **Duplicate**: Click this button to duplicate the role. It opens the role configuration screen preconfigured with the same permissions as this role (see Creating a Custom Role).
  - **Delete**: Click the button to delete the role. You cannot delete or modify the default roles.
  - **Update Role**: Click this button to modify the role permissions (see Modifying a Custom Role).
  - Click the question mark icon to open a help page in a separate tab or window.
  - Click the X icon to close the details page.

Summary Tab

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

- A section on the left that displays summary information about the role (see following table).
- A section in the middle which displays a list of actions that the role is permitted to execute. Click the **Actions** header to display additional actions for those roles with more than a single page of permitted actions. You can filter the actions list by entering a string in the search field (upper right above the list).

The following table describes the role summary information fields. A dash (-) is displayed in a field when a value is not available.

Table 113: Role Summary Fields

Parameter	Description	Values
Name	Displays the role name.	(role name)
Assigned Users	Displays the number of users assigned this role.	(number)
Assigned User Groups	Displays the number of user groups assigned this role.	(number)

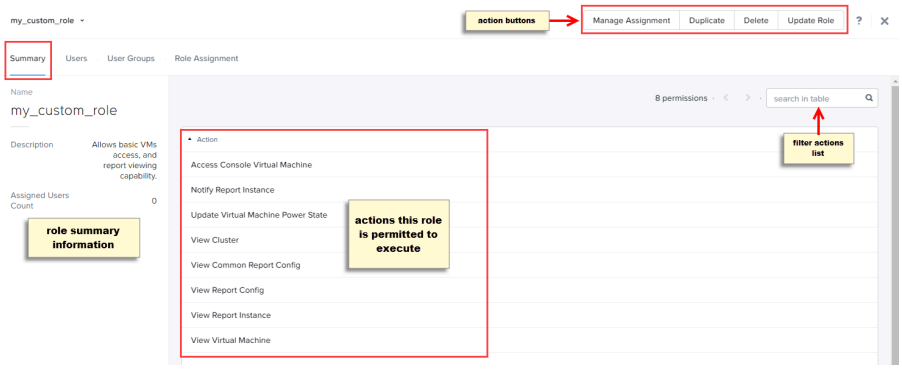

Figure 197: Role Summary Tab

Users Tab

The **Users** tab displays a list of users assigned to the role. You can filter the list by entering a string in the search field (upper right above the table).

User Groups Tab

The **User Groups** tab displays a list of user groups assigned to the role. You can filter the list by entering a string in the search field.

Role Assignment Tab

The **Role Assignment** tab displays the relationship between the users, groups, and entities assigned to this role. This tab does not appear for certain default roles such as Prism Admin that you are not allowed to assign.

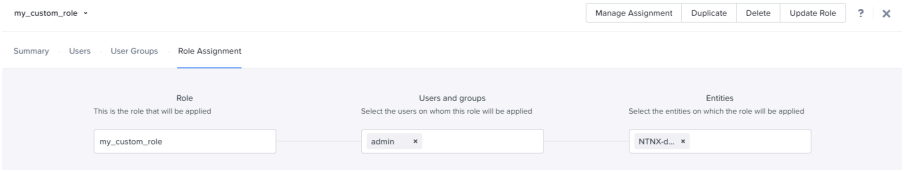

Figure 198: Role Assignment Tab

Projects Summary View

The Projects Summary page allows you to view the list of existing projects and their definitions in an at-a-glance view. You can click a project name to view and manage the detailed configuration of the project on the Project Details page. To access the Projects Summary page, click **Administration** > **Projects** in the Prism Central entities menu or the **Projects** icon in Calm.

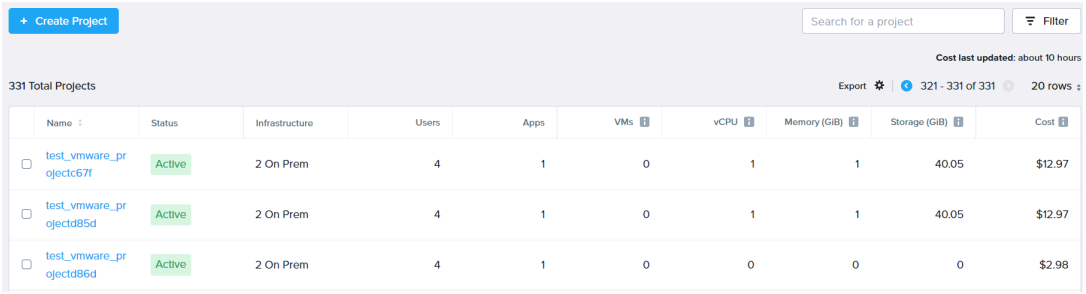

Figure 199: Projects Summary

The following table describes the definitions that appear for a project on the Projects Summary page.

Table 114: Projects List Fields

Definition	Description
Name	Displays the project name.
Status	Displays the status of the project.
Infrastructure	Displays the number of on-prem and cloud accounts added to the project.

Definition	Description
Users	Displays the number of users in the project.
Apps	Displays the number of applications deployed using the project.
VMs	Displays the number of on-prem AHV cluster VMs in the project.
vCPU	Displays the number of vCPUs used by the on-prem AHV clusters within the project.
Memory	Displays the amount of memory (in GiB) used by the on-prem AHV clusters within the project.
Storage	Displays the amount of storage (in GiB) used by the on-prem AHV clusters within the project.
Cost	Displays the resource cost of your applications for the last 30 days using Showback. Showback is supported only for Nutanix and VMware through vCenter platforms.

You can filter the projects list based on different parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 115: Filter Pane Fields

Parameter	Description
Name	Filters based on the project name. Select a condition from the pull-down list and enter a keyword in the field to get a list of projects that satisfy the condition and keyword combination. You can select **Contains**, **Does not contain**, **Starts with**, **Ends with**, or **Equals to** condition.
vCPU Usage	Filters based on the number of vCPUs used. Enter a range in the **from <low> to <high>** field to get a list of projects within the range of vCPU usage that you specify.
Memory Usage	Filters based on the amount of memory used. Enter a range in the **from <low> to <high> GiB** fields to get a list of projects within the range of memory usage that you specify.
Storage Usage	Filters based on the amount of storage space used. Enter a range in the **from <low> to <high> GiB** field to get a list of projects within the range of storage usage that you specify.
VM Count	Filters based on the number of VMs. Enter a range in the **from <low> to <high>** field to get a list of projects within that range of the total number of VMs.

You can click the Gear icon to select, deselect, or reorder the columns that you want to view on the Projects Summary page. You can click **Export** to export the data in a .csv file.

The Projects Summary page includes a **Create Project** button to create projects. For more information, see Creating a Project on page 729.

The **Actions** menu appears when you select one or more projects. The menu allows you to update or delete the project. You can update only one project at a time, but you can select multiple projects when deleting.

To view the details of a project, click the project name to open the Project Details page. For more information, see Project Details View on page 258.

Project Details View

The Project Details page allows you to view the entire setup of a project in different tabs and helps you manage their configuration. To access the Project Details page, click the project on the Project Summary page. For more information about the Project Summary page, see Projects Summary View on page 256.

A few tabs on the Project Details page appear only when you enable Calm in your Prism Central instance. The Project Details page has the following tabs:

- Dashboard
- Usage
- Workloads
- Users & Groups
- Infrastructure
- Environments (appears when you enable Calm in your Prism Central instance)
- Policies (appears when you enable Calm in your Prism Central instance)
- Tunnels (appears when you enable Calm in your Prism Central instance)

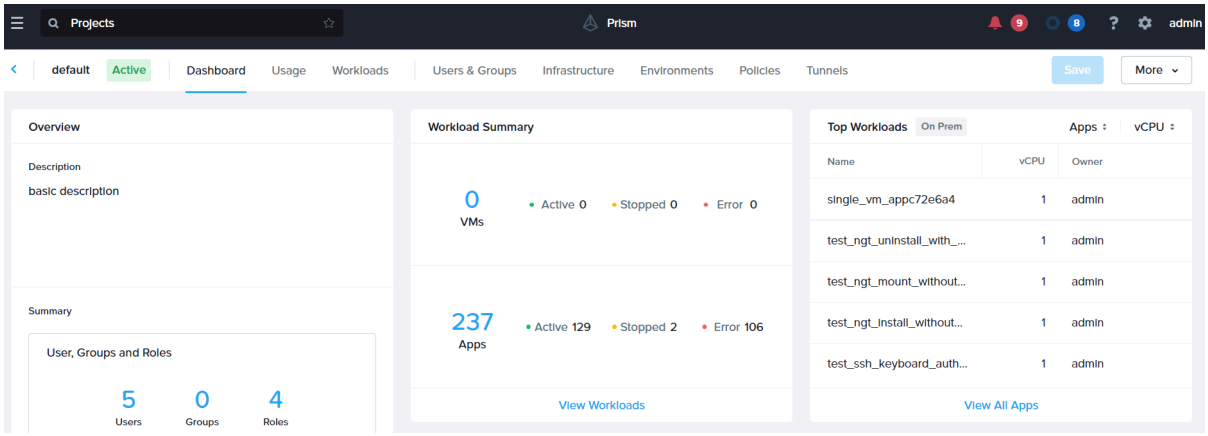

Figure 200: Project Details

Project Dashboard Tab

The **Dashboard** tab appears when you first open the Project Details page. The project dashboard contains tiles to display various configurations and usage summary of the project at one place. The availability of tiles and the appearance of the information on the tiles differ based on Calm enablement (see Enabling and Accessing Calm) and the overall project configuration and usage data.

When Calm is disabled, the project dashboard typically appears as follows:

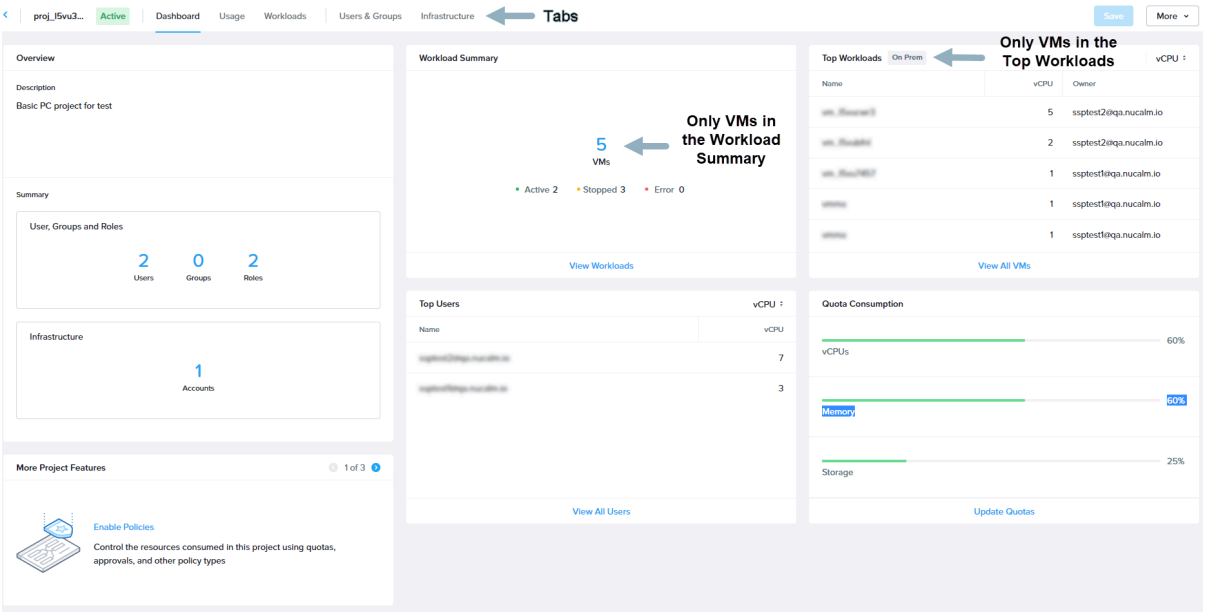

Figure 201: Project Dashboard when Calm is Disabled

When Calm is enabled, the project dashboard typically appears as follows:

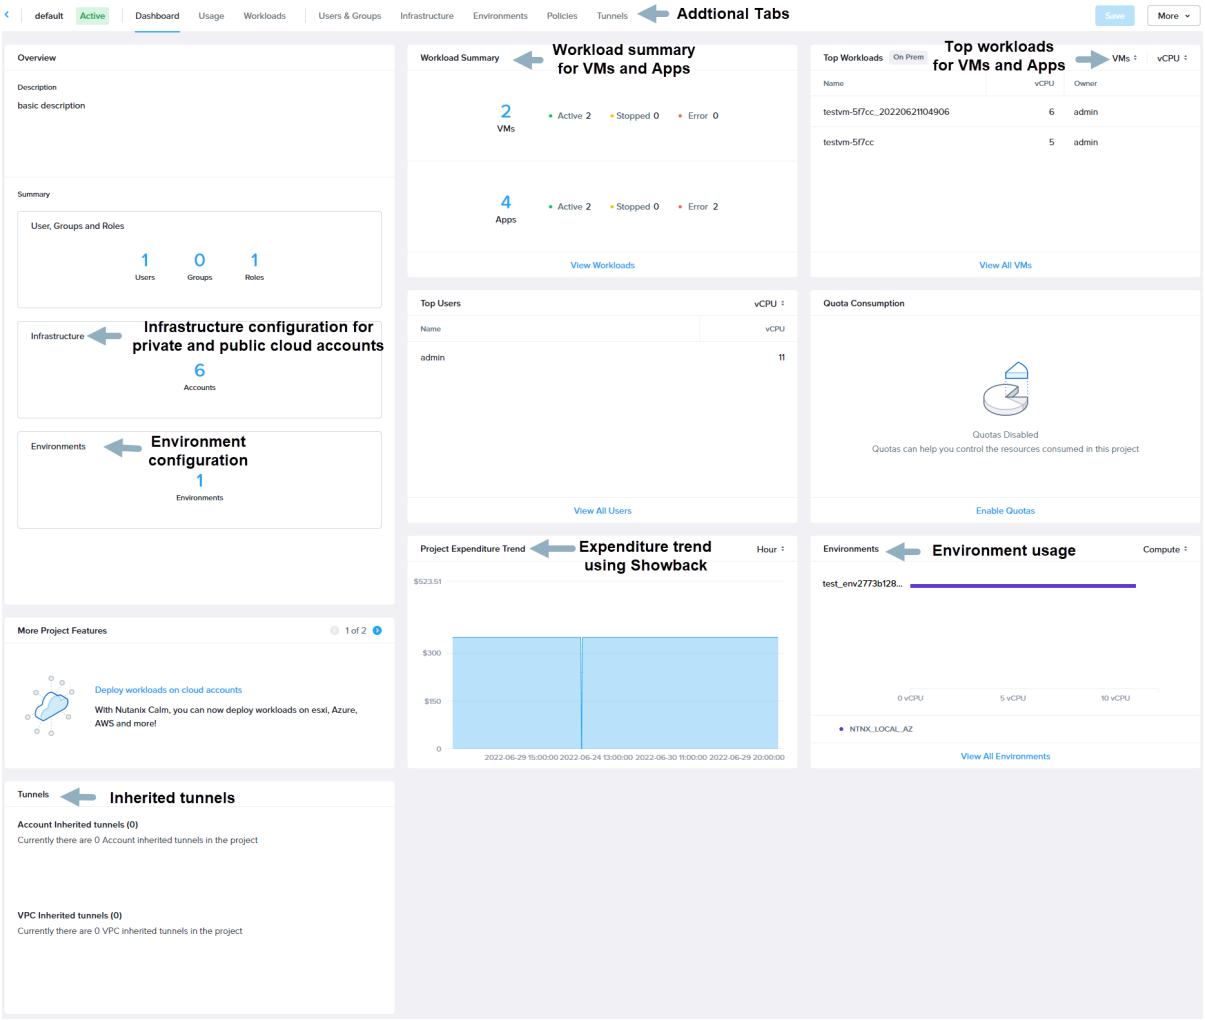

Figure 202: Project Dashboard when Calm is Enabled

The dashboard has the following tiles:

Table 116: Project Dashboard

Tiles	Description
Overview	The Overview tile has the following sections. • Description: Displays the description you provided for the project. You can hover your mouse on the tile and click **Edit** to add or modify the description of the project. • Summary: Displays the number of users, groups, roles, infrastructure (accounts), and environments configured for the project. You can hover your mouse and click the **Edit** icon to open the respective tab and edit the configuration. For more information about the infrastructure (accounts), see Infrastructure in Projects on page 736. For more information about environments, see Environments in Projects on page 745.
Workload Summary	Provides information related to the overall workloads associated with the project. When Calm is disabled, the tile displays the total number of VMs that are up and the number of VMs that are in the active, stopped, or error state within the project. When Calm is enabled, the tile also displays the total number of applications and the number of applications that are in the active, stopped, or error state within the project. You can click **View Workloads** to navigate to the **Workloads** tab and view the state, owner, and other details of the VMs or applications. For more information, see Project Workloads Tab on page 263.
Top Workloads	Provides information related to the workloads that consume the maximum resources. When Calm is disabled, the tile displays the VMs that consume the maximum resources within the project. When Calm is enabled, the tile also displays the applications (on-prem AHV and ESXi) that consume the maximum resources within the project. You can use the switcher at the top of the tile to view the list of VMs or applications. You can filter the list of VMs or applications based on the vCPU, memory, or storage consumption. You can click **View all VMs** or **View all Apps** to navigate to the **Workloads** tab and view the state, owner, and other details of the VMs or applications. For more information, see Project Workloads Tab on page 263.

Tiles	Description
More Project Features	Displays recommendations and options to enable additional features for your projects. You can click the options available on the tile to enable Calm and configure these additional features for your project. • Enable Policies: Policies in a project allows you to control the resources that are consumed in your project. With policies, you can configure resource quotas, approvals, and other policy types. For more information, see Policy Engine Overview. • Deploy workloads on cloud accounts: In addition to the on-prem private clouds with AHV or ESXi, you can also deploy your workloads on public clouds when you enable Calm. You can configure your public clouds such as AWS, Azure, or GCP in Calm. For more information, see Provider Account Settings in Calm. • Provision, scale, and manage multi-VM applications: You can design complex blueprints that span multiple environments and teams using Nutanix Calm. For more information, see Calm Blueprints Overview. • Further bucket your infrastructure: You can configure environments to segregate your infrastructure for cases such as development, testing, production,disaster recovery, and so on. For more information, see Environments in Projects on page 745.
Top Users	Display the users that consume the maximum resources within the project. You can filter the users based on the vCPU, memory, or storage consumption using the resource switcher on the tile. You can click **View all Users** to navigate to the **Users & Groups** tab and view the user details such as the group that the user belongs and the resources that the user is consuming. For more information, see Project Users & Groups Tab on page 265.
Quota Consumption	Displays the percentage of resource quotas (vCPU, memory, and storage) that you have defined for the Nutanix accounts associated with the project. You can click **Update Quotas** to navigate to the **Infrastructure** tab and edit the quota values for the account. When Calm and policy engine are enabled, you can also define quota values for any VMware accounts that you add to your project. You can click **View Quota Utilization Report** to view the resource utilization, quota utilization, and application quota utilization at the project level. To get better quota consumption data with the upgraded quota capabilities, it is recommended to enable Calm and then the policy engine. You do not need to license Calm to enable policy engine. For more information about policy engine, see Policy Engine Overview.
Project Expenditure Trends	Display a graph to report the expenditure trend within the project from an hour, day, or month viewpoint when Calm is enabled in your Prism Central instance. Expenditure trend is generated when you enable Showback in Calm. For more information, see Showback.

Tiles	Description
Environments	Displays the environment usage within the project when you have configured environments in your project and have launched applications using those environments. You can sort the environment usage on the basis of compute, memory, or usage using the resource switcher on the tile. You can click **View all Environments** to navigate to the **Environments** tab and view or update your environment configuration. For more information on environment configuration, see Environments in Projects on page 745.
Tunnels	Displays the number of tunnels inherited from the account and the VPCs associated with the project.

Project Usage Tab

The **Usage** tab provides the graphical representation of the resource usage within the project. You can select the time period for the graphs by selecting the 1 day or 1 week option in the **Show** list. The tab displays the following graphs:

- The **vCPU Usage** graph displays a rolling time interval monitor of project vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of project memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of project storage usage.

The graphs appears for on-prem AHV clusters only.

Project Workloads Tab

The **Workloads** tab displays a list of on-prem AHV cluster VMs associated with the project. When Calm is enabled, the tab also displays the applications launched within the project. The Workloads tab has the following tabs that you can access in the left pane.

Table 117: Tabs Within the Workloads Tab

Tabs	Description
VMs	Provides the details of all the VMs associated with the project. You can view the following details of a VM on this tab. • **Name**: Displays the name of the VM. You can click the VM name to navigate to the VM Details page where you can view and update the VM configuration. • State: Displays whether the VM is currently On (green circle) or off (red circle). • Owner: Displays the user who owns the VM. • vCPU: Displays the number of vCPUs allocated to the VM. • Memory: Displays the amount of memory (in GiB) allocated to the VM. • Storage: Displays the amount of disk space (in GiB) allocated to the VM.
Apps	This tab appears only when Calm is enabled in your Prism Central instance and you have launched applications ((on-prem AHV and ESXi) using the blueprints that are configured within the project. You can view the following details of an application on this tab. • **Name**: Displays the name of the application. You can click the application name to navigate to the Application Details page where you can view or manage the application configuration. • **Source Blueprint**: Displays the name of the blueprint that was used to launch the application. You can click the blueprint name to navigate to the Blueprint details page where you can view or manage the blueprint configuration. • State: Displays the state of the application. The state can be Running, Stopped, or Error. • Owner: Displays the user who owns the application. • Created on: Displays the date and timestamp of the application creation. • Last Updated at: Displays when the application was updated the last time. • Cost: Displays the service cost of the application incurred in the last 30 days. The information appears for Nutanix and VMware when the Showback is enabled in Calm.

Project Users & Groups Tab

The **Users & Groups** tab displays a list of users and groups assigned to the project.

Table 118: Users & Groups

Parameter	Description
Name	Displays the Active Directory name (typically in the form of `name@domain`) associated with the project.
User group	Displays the group because of which the user is added to the project.
Total VMs	Displays the number of VMs owned by the user.
vCPU	Displays the number of vCPUs consumed by the user.
Memory	Displays the amount of memory (in GiB) consumed by the user.
Storage	Displays the amount of storage (in GiB) consumed by the user.

You can click **Add/Edit Users & Groups** to add or edit users or group of users to your project. For more information, see Adding Users to a Project on page 733.

Project Infrastructure Tab

The **Infrastructure** tab allows you to add accounts to your project and displays the details of the added accounts. The tab also allows you to remove accounts from the project. The left pane on the tab lists all the accounts that you added. When you click an account, the right pane displays the account details.

You can click **Add Infrastructure** to add accounts to your project.

When Calm is disabled in your Prism Central instance, you can add your local Nutanix account and select clusters and subnets for the project. For more information, see Adding Infrastructure in a Calm-Disabled Prism Central on page 737.

When Calm is enabled, you can add your on-prem private clouds with AHV or ESXi or any public clouds (such as AWS, Azure, or GCP) that you configured in Calm. For more information, see Adding Infrastructure in a Calm-Enabled Prism Central on page 739.

For more information about account configuration in Calm, see Provider Account Settings in Calm.

Project Environments Tab

The **Environments** tab appears when you enable Calm in your Prism Central instance. The tab displays all the environments that you configured within the project.

An environment is a subset of a project that you configure to use during blueprint creation or application launch. For more information, see Environments in Projects on page 745.

You can use the icons at the top of the page to view the environment details in a list or tile mode. You can use the Search field to search environments on the page.

The tab displays the following details about an environment.

Table 119: Environments

Parameter	Description
Name	Displays the name of the environment. You can click the environment name to navigate to the Environment Configuration page where you can view or manage the configuration.
Created By	Displays the user (role) who created the environment.
Accounts	Displays the accounts associated with the environment.
Credentials	Displays the number of credentials associated with the environment.
Ready For	Displays the readiness of the environment for blueprint creation or application launch.

You can click the gear icon to select, deselect, or reorder the columns in the list mode.

You can click the **Create Environment** button to create and configure an environment. For more information, see Configuring Environments in a Project on page 746.

You can click the vertical ellipsis next to an environment to update or delete the environment.

Project Policies Tab

The **Policies** tab appears when you enable Calm in your Prism Central instance. You must enable policy engine to view or manage policy configurations on the **Policies** tab.

The **Quotas** tab in the left pane provides a unified view of the resource quota limits that you defined for the project and the accounts within the project. For information about Quota Policy, see Quota Policy Overview on page 773 or Managing Quota Limits for Projects on page 774.

The **Snapshot** tab in the left pane provides the details of all snapshot policies you created within the project. For information on snapshot policy creation, see Creating a Snapshot Policy on page 777.

Project Tunnels Tab

The **Tunnels** tab appears when you enable Calm in your Prism Central instance and displays the number of tunnels inherited from the account and the VPCs associated with the project. For more information about tunnels, see Tunnels for Orchestration within a VPC.

Availability Zones View

To access the availability zones dashboard, select **Administration** > **Availability Zones** from the entities menu (see Entities Menu on page 13). The availability zones dashboard allows you to view information about the availability zones connected to Prism Central currently and to connect to other zones.

> **Note:** This section describes the information and options that appear in the availability zones dashboard.
>
> - For instructions on how to view and organize that information in a variety of ways, see Entity Exploring on page 89.
> - For information about availability zones and how to connect to them, see Nutanix Disaster Recovery Guide.

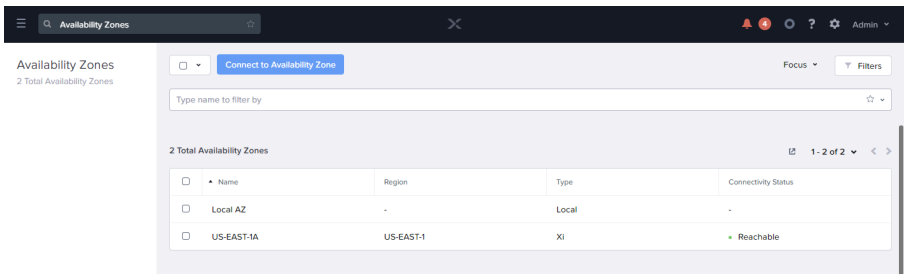

Figure 203: Availability Zones Dashboard

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

Table 120: Catalog Items List Fields

Parameter	Description	Values
Name	Displays the zone name.	(zone name)
Region	Displays the region in which the zone is located.	(region name)
Type	Displays the type or zone. The availability zone that you are logged in to is shown as a local availability zone. Availability zones that are backed by on-premises Prism Central instances are shown as type physical. Availability zones in Xi Cloud Services are shown as type Xi.	Local, Physical, Xi
Connectivity Status	Indicates whether Prism Central can access the zone currently.	Reachable, Unreachable

To filter the list based on the type of zone (local, physical, or Xi), click the **Filters** button and select the desired type(s) in the **Filters** pane.

Click the **Connect to Availability Zone** button to connect to a new availability zone. The **Actions** menu appears when one or more zones are selected. It allows you to disconnect from the selected zones.

LCM View

To access the LCM dashboard, select **Administration** > **LCM** from the entities menu (see Entities Menu on page 13). The life cycle manager (LCM) tracks software and firmware versions of the various components in a cluster. It allows you to view information about the current inventory and update the versions as needed. See the Life Cycle Manager Guide for more information.

Services Entities

You can enable the following services from the **Services** category of the entities menu (see Entities Menu on page 13):

- Calm (see Enabling Calm on page 992)
- Files (see Enabling Files on page 992)
- Foundation Central (see Enabling Foundation Central on page 993)
- Karbon (see Enabling Karbon on page 993)

- Objects (see Enabling Objects on page 993)

ALERT AND EVENT MONITORING (PRISM CENTRAL)

Prism Central provides several mechanisms to monitor and configure cluster alerts and events.

- The Alerts and Events dashboards, which you can access from the entities menu (see Entities Menu on page 13), allow you to view and monitor alert and event activity across the registered clusters.
  - The Alerts Summary View (Prism Central) on page 269 displays a list of alerts that you filter in various ways. You can also drill down for detailed information and corrective actions for an alert (see Alert Details on page 276)
  - The Events Summary View (Prism Central) on page 279 displays a list of events that you filter in various ways. You can also drill down for detailed information for an event (see Event Details on page 280).
- You can configure rules for who should receive email alerts and customize the alert messages (see Configuring Alert Emails (Prism Central) on page 281).
- You can customize the list of events that generate an alert (see Configuring Alert Policies (Prism Central) on page 285).
- You can review event log files when necessary (see Prism Central Logs on page 302).

> **Note:** These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default. (Some DR events are retained for just 7 days by default.)

Alerts Summary View (Prism Central)

The Alerts dashboard allows you to view summary information about alert messages across the registered clusters, access detailed information about each alert, and view alert policies from any source (user defined, system defined, or external defined).

To access the Alerts dashboard, select **Activity** > **Alerts** from the entities menu (see Entities Menu on page 13). The top pane includes **List** and **Alert Policies** tabs with **User Defined**, **System Defined**, and **External Defined** tabs under **Alert Policies**. Click the tab to display that content below the tab.

> **Note:** This section describes the information and options that appear in the Alerts dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways. See Alerts/Health checks on page 307 for a list of alert messages and corrective actions.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of active alerts across the registered clusters. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

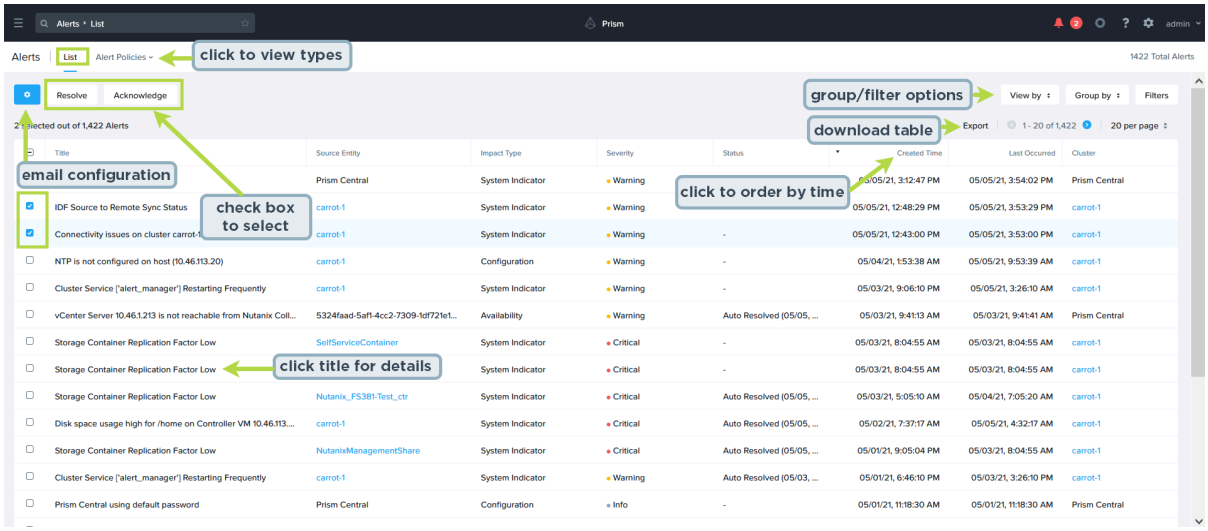

Figure 204: Alerts List Tab

- To order the alerts, click a column header, which toggles the alerts alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of alerts in CSV format, click **Export**. The maximum number of alerts you can export is 1000.
- To display the details page for an alert, click the alert title (see Alert Details on page 276).
- To enable alert emails and specify email addresses to which alerts should be sent, click the **Email Configuration** button (see Configuring Alert Emails (Prism Central) on page 281).
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 89).
- To group the alerts, select an option from the **Group** pull-down menu. You can group the alerts by cluster, severity, or impact type. (You can only choose one.)
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Cluster: Enter name in search field.
  - Provider: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Severity: Critical, Warning, Info
  - Resolved: Yes (Manual), Yes (Auto), No
  - Acknowledged: Yes, No
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Create Time: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*
  - Last Occurred: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*
  - Title: Enter a string in the field.

- To select (or unselect) alerts in the list, do one of the following:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all** (or **Clear selections**) to select (unselect) all the alerts in the list.
- When one or more alerts are selected, **Resolve** and **Acknowledge** appear in the menu; click one or both to resolve and/or acknowledge the selected alerts. The user name and time appears in the appropriate field when an alert is acknowledged or resolved.

Table 121: Alerts List Fields

Parameter	Description	Values
(selection box)	Click this box to select the message for acknowledgement or resolution.	n/a
Title	Displays the name of the alert. Click the name to displays the details page for that alert.	(alert name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the alert applies. A comma separate list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Impact Type	Displays the category in which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Severity	Displays the severity level of this condition. There are three levels: Critical An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. Warning An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. Info (Informational) An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational

Parameter	Description	Values
Status	Indicates whether the alert has been resolved. Resolving an error means you set that error as fixed. (The alert may return if the condition is scanned again at a future point.) If you do not want to be notified about the condition again, turn off the alert for this condition (see Configuring Alert Policies (Prism Central) on page 285). • A blank value means the alert is not resolved. • An "Acknowledged By" value means this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, check the box for that alert and then select the **Acknowledge** button. • An "Auto Resolved (*`date_time`*)" value means this alert was resolved automatically at the specified date and time. This requires that auto alert is enabled for this alert type. • A "Resolved By *`user`* (*`date_time`*)" value means this alert was resolved manually by the specified user at that date and time. To manually resolve an alert, check the box for that alert and then select the**Resolve** button.	(blank), Acknowledged By *`user`* (*`date_time`*), Auto Resolved (*`date_time`*), Resolved By *`user`* (*`date_time`*)
Create Time	Displays the date and time when the alert occurred.	(date and time)
Last Occurred	Displays the date and time when the alert last occurred (before this occurrence). If this is the first occurrence, both Create Time and Last Occurred display the same date and time.	(date and time)
Cluster	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 209).	(cluster name)

Alert Policies: User Defined Tab

Clicking the **User Defined** tab displays a list of the custom (user defined) alert policies. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

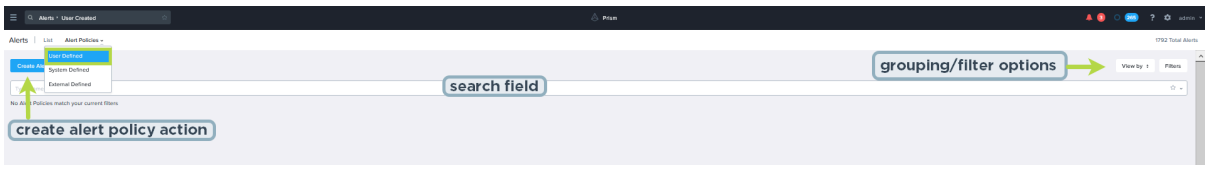

Figure 205: Alert Policies: User Defined Tab

- To create a custom policy, click the **Create Alert Policy** button (see Adding Custom Alert Policies on page 287).

- To view or modify a custom policy, click the policy name (see Modifying Custom Alert Policies on page 292).
- To delete a custom policy, select the policy and then click the **Delete** button.
- To order the polices, click a column header, which toggles the list alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 89).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Auto Resolve: Enabled, Disabled
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.

Table 122: Alert User Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to displays the policy definition page (see Modifying Custom Alert Policies on page 292).	(policy name)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Auto Resolve	Displays whether auto resolved is enabled.	Enabled, Disabled
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Policies: System Defined Tab

Clicking the **System Defined** tab displays a list of the default (system defined) alert policies. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

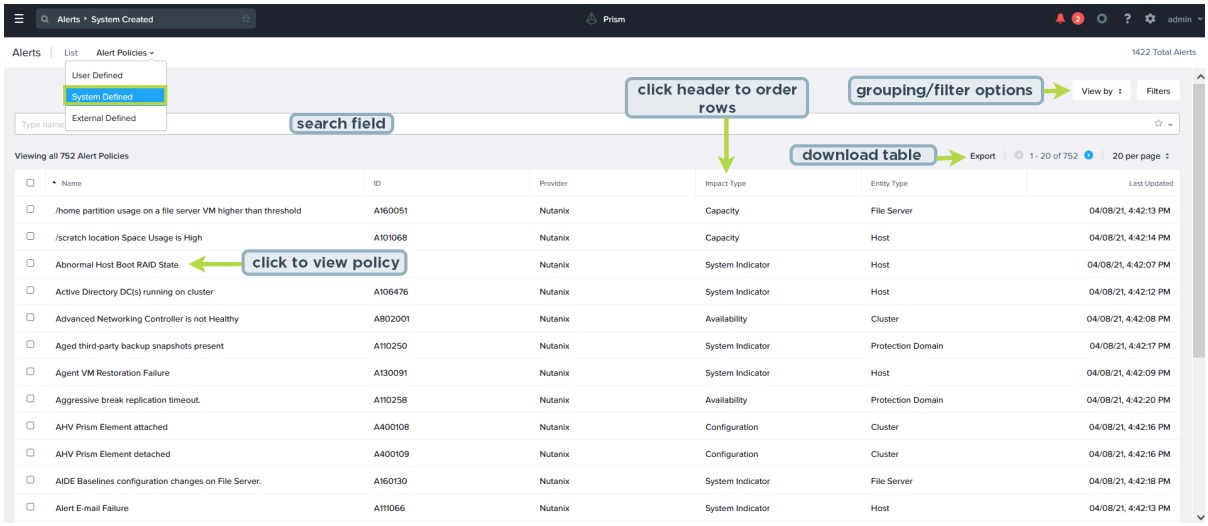

Figure 206: Alert Policies: System Defined Tab

- To view or modify a system policy, click the policy name (see Modifying System Alert Policies on page 286).
- To order the polices, click a column header, which toggles the list alphabetically, numerically, or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 89).
- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - ID: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 123: Alert System Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to display the **Update Policy** window (see Modifying System Alert Policies on page 286).	(policy name)
ID	Displays the identification number assigned to the policy.	(ID number)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Policies: External Defined Tab

Clicking the **External Defined** tab displays a list of the alert policies defined by external entities, typically an application through the API. The following figure is a sample view, and the table describes each field in this view. A dash (-) is displayed in a field when a value is not available or applicable.

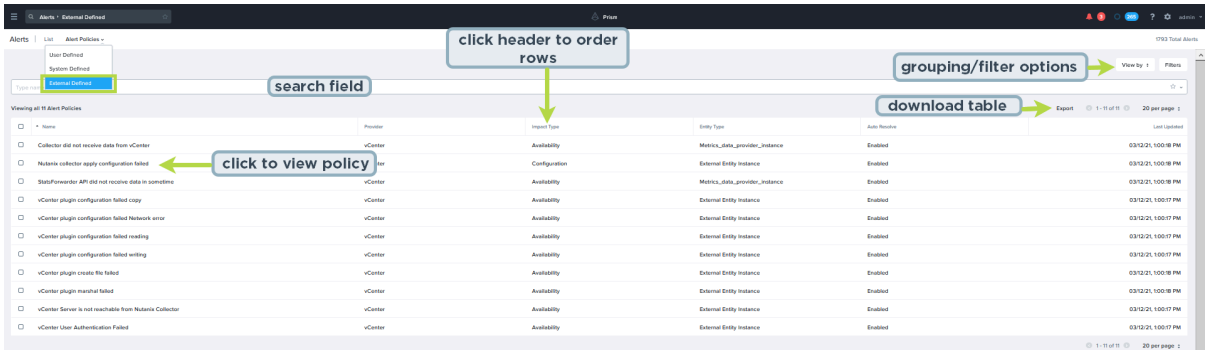

Figure 207: Alert Policies: External Defined Tab

- To view an external policy, click the policy name (see Viewing External Alert Policies on page 302).
- To order the polices, click a column header, which toggles the list alphabetically or chronologically (as appropriate) based on that parameter.
- To download the table of policies in CSV format, click the **Export** link.
- To create a custom view, select **Add custom** from the **Focus** pull-down menu (see Entity Exploring on page 89).

- To filter the list, either enter a string in the search field or click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
  - Name: Select a condition from the pull-down list (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
  - Provider: Select a condition from the pull-down list (same as Name) and enter a string in the field.
  - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
  - Auto Resolve: Enabled, Disabled
- To select (or unselect) policies in the list:
  - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
  - Click the check box in the header row to select (unselect) all the alerts in the list.
  - In the check box pull-down menu, select **Select all**, **Show Selections** (or **Show Selected Entities**), or **Clear selections**) to select (unselect) that set of policies.

Table 124: Alert System Defined Policies Fields

Parameter	Description	Values
(selection box)	Click this box to select the policy.	n/a
Name	Displays the name of the policy. Click the name to display the *policy_name* window (see Viewing External Alert Policies on page 302)..	(policy name)
Provider	Displays the name of the provider (organization, feature, or external source) that defined the policy.	(provider name)
Impact Type	Displays the category in which the policy applies.	Availability, Capacity, Configuration, Performance, System Indicator
Entity Type	Displays the entity type such as a cluster, host, or VM to which the policy applies.	(entity names)
Auto Resolve	Displays whether auto resolved is enabled.	Enabled, Disabled
Last Updated	Displays the date and time when the policy was last updated. (This is the creation date if the policy has never been updated.)	(date and time)

Alert Details

Clicking on an alert message in the dashboard (see Alerts Summary View (Prism Central) on page 269) or anywhere else the alert title appears, such as in a search list, displays detailed information about that alert. The alert details appear in the left column. Possible causes for the alert appear to the right. The most likely cause appears first with other possible causes (if any) appearing below in the order of likelihood. Each cause includes a Recommendation section that describes the recommended corrective action and in many cases a Details section that provides additional context and instructions.

- To switch among alert details pages, select the desired alert from the pull-down list in the upper left of the screen.
- To acknowledge or resolve the alert, click the **Resolve** or **Acknowledge** button in the upper right.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 125: Alert Details Fields

Parameter	Description	Values
Alert Title	Displays the title of the alert.	(alert title)
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the alert. Clicking an entity link opens the detail page for that entity.	(entity names)
Severity	Displays the severity level of this condition. There are three levels: Critical An actionable critical situation has been detected, and action is required immediately. The cluster may have the potential to stop running, or it could run into irreparable issues. Warning An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. Informational (Info) An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational
Cluster Name	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 209).	(cluster name)
Cluster Version	Displays the AOS version running on the cluster.	(version number)
Created Time	Displays the date and time when the alert first occurred.	(time and date)
Last Occurred	Displays the date and time for the most recent occurrence of the alert.	(time and date)
Impact Type	Displays the impact type to which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Policy	Displays the name of the alert policy.	(alert policy name)

Parameter	Description	Values
Status (for manually resolved alerts)	Displays the current alert status: acknowledged, resolved, or neither (blank).	(blank), Acknowledged, Resolved
Status (for auto resolved alerts)	Displays the current alert status: acknowledged, resolved, auto resolved, or neither (blank).	(blank), Acknowledged, Auto Resolved
Acknowledged By	Indicates whether the alert has been acknowledged. • A blank (dash) value means the alert is not acknowledged. • An "N/A *date_time*)" entry indicates this alert was acknowledged automatically at the specified date and time. This requires that auto resolve is enabled for this alert type (see Configuring Alert Policies (Prism Central) on page 285). • A "*user* *date_time*" entry indicates this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, click the **Acknowledge** button (upper right). The field value changes from blank to the user name and date/time the alert was acknowledged.	(blank), N/A *date_time*), *user* *date_time*
Resolved By	Indicates whether the alert has been resolved. The options (blank, N/A, or user) are the same as for Acknowledged. To manually resolve an alert, click the **Resolve** button (upper right). The field value changes from blank to the user name and date/time the alert was resolved.	(dash), N/A *date_time*), *user* *date_time*
Description	Describes the reason for the alert.	(event description)

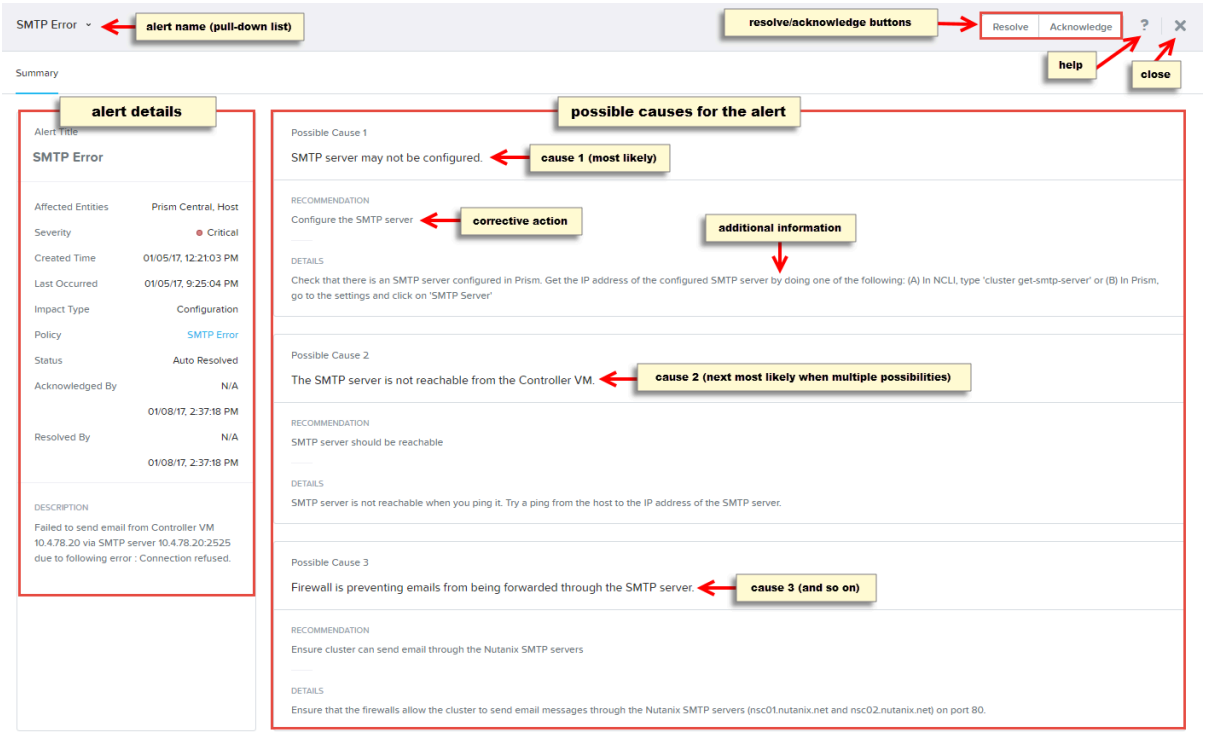

Figure 208: Alert Details Screen

Events Summary View (Prism Central)

The Events dashboard summary view displays a list of event messages across the registered clusters. Event messages describe cluster actions such as adding a storage pool or taking a snapshot. Unlike alerts, event messages are simply informational without the need to acknowledge or resolve. The following figure is a sample view, and the table describes each field in this view. To access the Events dashboard, select **Activity** > **Events** from the entities menu (see Entities Menu on page 13).

> **Note:** This section describes the information and options that appear in the Events dashboard. See Entity Exploring on page 89 for instructions on how to view and organize that information in a variety of ways.

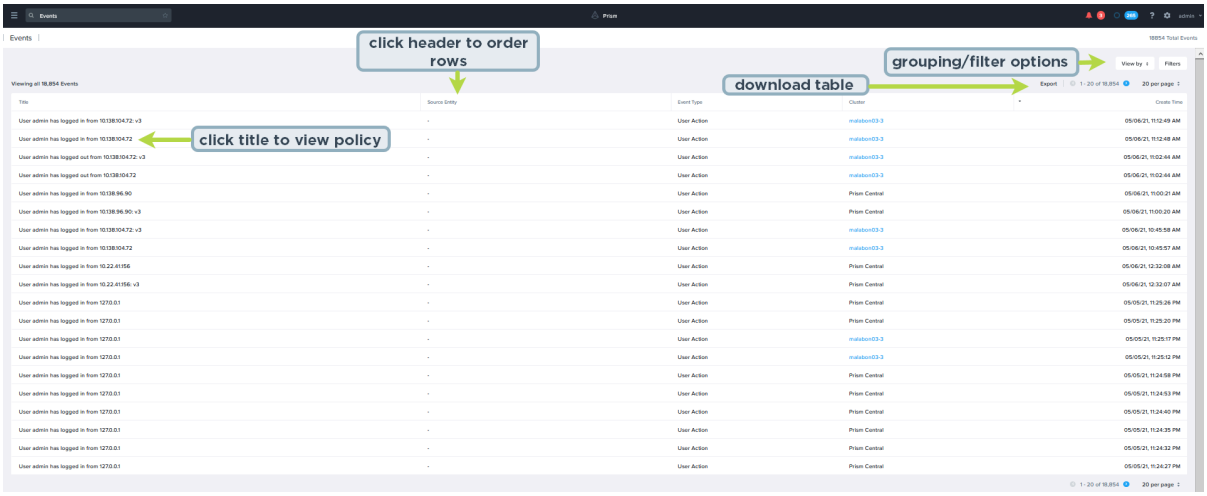

Figure 209: Events Dashboard

- To display the details page for an event, click the event title (see Event Details on page 280).
- To order the events chronologically, click the **Create Time** column header, which toggles the event list latest-to-earliest or earliest-to-latest.
- To download the table of alerts in CSV format, click the **Export** link.
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following event parameters and values.
  - Event Type: Behavioral Anomaly, System Action, User Action
  - Cluster: Enter name in search field.
  - Create Time: Last 1 hour, Last 24 hours, Last week, From *XXX* to *XXX*

Table 126: 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 280).	(event name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the event applies. A comma separated list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Event Type	Displays the category in which the event is classified.	System Action, User Action, Behavioral Anomaly (see Behavioral Learning Tools on page 899), DR
Cluster	Displays the name of the cluster in which the event was issued.	(cluster name)
Create Time	Displays the date and time when the event occurred.	(date and time)

Event Details

Clicking on an event message in the dashboard (see Events Summary View (Prism Central) on page 279) 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.

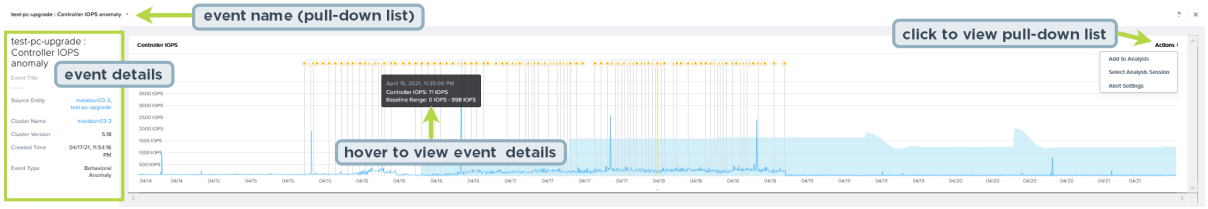

Figure 210: Event Details Screen

- 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 127: Event Details Fields

Parameter	Description	Values
Event Title	Displays the title of the event.	(alert title)
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the event. Clicking an entity link opens the detail page for that entity.	(entity names)
Cluster Name	Displays the name of the cluster in which the alert was issued. Click the name to display the details page for that cluster (see Cluster Details View on page 209).	(cluster name)
Cluster Version	Displays the AOS version running on the cluster.	(version number)
Created Time	Displays the date and time when the event occurred.	(time and date)
Event Type	Displays the type of event. Three types are defined, and each event is assigned to one of these types.	Behavioral Anomaly, System Action, User Action

Configuring Alert Emails (Prism Central)

About this task

Prism Central allows you to configure the alert messages sent by Prism Central. To configure alert settings, reporting rules, and message templates, do the following:

Note:

- Alert emails sent by Prism Central are in addition to any alert emails you might have configured on individual clusters through the Prism Element web console. You will receive email from both entities in this case. Prism Central alert emailing is not enabled by default; you must explicitly enable it and specify the recipients (Nutanix customer support and/or supplied email addresses). If you enable alerts through Prism Central and do not want to receive double email notifications for the same alert, disable customer email notification for those alerts on the individual clusters through Prism Element (but keep email notification for Nutanix customer support enabled).

- Prism Central requires an SMTP server to send alert email messages (see Configuring an SMTP Server (Prism Central) on page 516).

Procedure

1. Do one of the following:

   » Click the gear icon and select **Alert Email Configuration** from the **Settings** menu (see Settings Menu (Prism Central) on page 16).

   » In the **Alerts** dashboard, click the **Configure** button and select **Email Configuration** from the pull-down list.

   The **Alert Email Configuration** window appears.

2. Click the **Settings** tab and do the following:

   a. Enable Prism Central to sent alert emails.

      - Check the **Every Single Alert** box to send an email whenever the event occurs.
      - Check the **Daily Digest** box to send a cumulative (24 hour) list of alerts once a day. (If you click the **Skip empty digest email** box, an email is not sent for the days when no alerts occur.)

      You can check one or both boxes. If neither box is checked, no alert emails are sent by Prism Central.

   b. To send alert notifications to others, enter their email addresses in a comma separate list in the **Email Recipients** field.

   c. Click the **Save** button to apply the changes and close the window.

      If you want to configure rules or an email message template, continue to the following steps before clicking the **Save** button.

   The **Tunnel Connection** section displays mail transport status information.

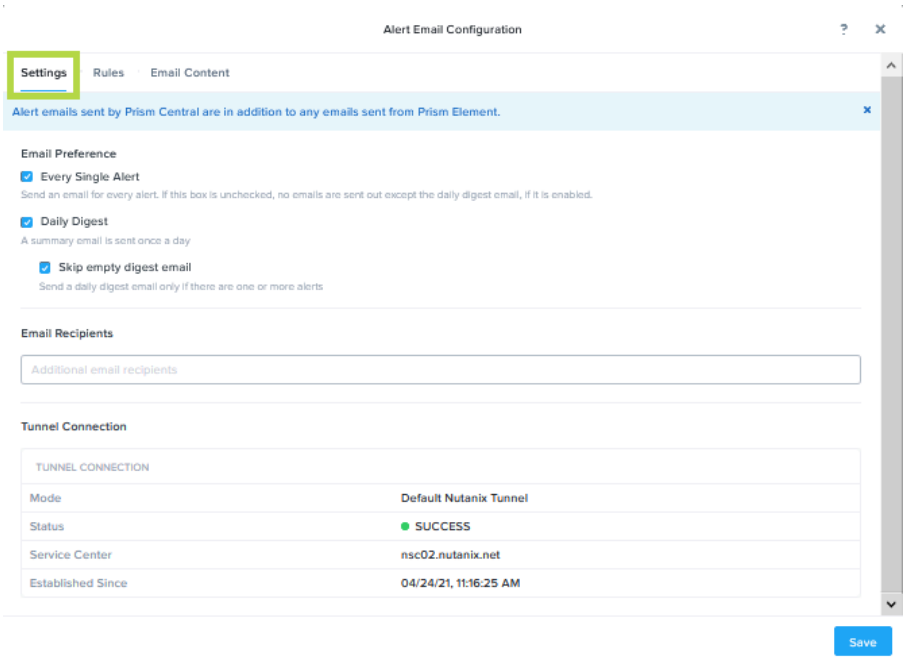

Figure 211: Alert Email Configuration Window (Settings tab)

3. To create a custom alert email rule, click the **Rules** tab and the **New Rule** button, and then configure the rule as follows:

   a. Specify the conditions for generating the alert:

      - **Alert Severity**: Select one or more of the severities from the pull-down list (Critical, Warning, Info, All).
      - **Impact Type**: Select one or more of the categories from the pull-down list (Availability, Capacity, Configuration, Performance, System Indicator, All).
      - **Cluster**: Select one or more of the clusters from the pull-down list (`cluster_name`, All).
      - **Alert Contains**: Enter a key phrase or word that should generate an email notification whenever the alert contains that phrase. For example, if you want to get an email notification when an alert contains the phrase "low memory", enter low memory in the field.

   b. Specify who should receive the alert email by entering recipient email addresses as a comma separated list in the **Send Email To** field.

   c. Click the **Save Rule** button to save the rule.

   d. Repeat these steps to create more custom rules.

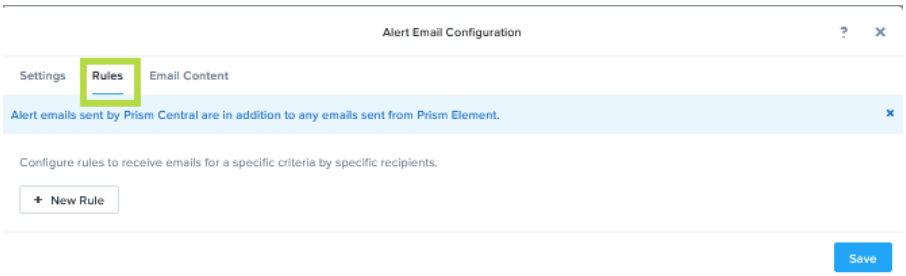

Figure 212: Alert Email Configuration Window (Rules tab)

4. To create a template for the email messages, click the **Email Content** tab and do the following:

   a. In the **Prepend Subject** field, enter the desired text.

      This text appears at the beginning of the subject field in each alert email. If the field is left blank, no prepended text appears in the subject.

   b. In the **Append Body** field, enter the desired text.

      This text appears at the end of the message body in each alert email. If the field is left blank, no appended text appears in the message body.

   c. Click the **Save** button to apply the changes and close the window.

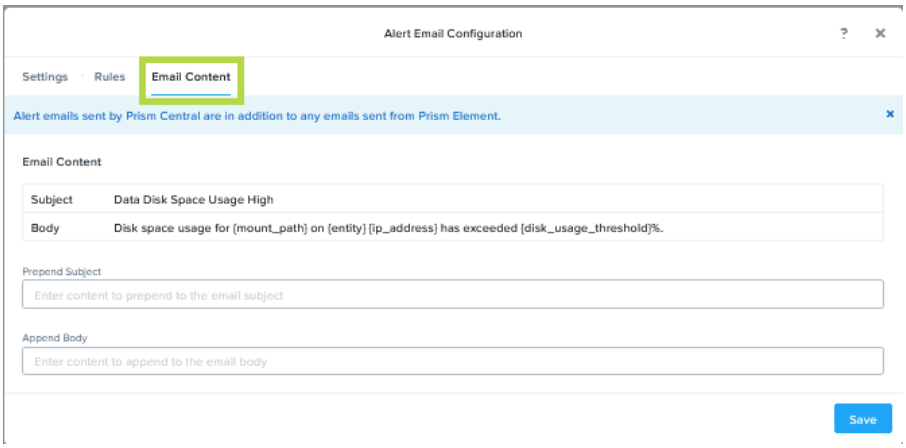

Figure 213: Alert Email Configuration Window (Template tab)

Configuring Alert Policies (Prism Central)

About this task

The system monitors a variety of conditions and sends an alert whenever one of the alert conditions is detected (when alerting is enabled). There are default policies for these alerts, but you have the option to modify the default policies and add new policies.

> **Note:** To help ensure that Prism Central and each managed cluster are taking advantage of NCC features, ensure that:
>
> - Each node in your cluster is running the same NCC version.
> - Prism Central and each cluster managed by Prism Central are all running the same NCC version.

There are three categories of alert policies:

Procedure

- System Defined Alert Policies. You can view and customize policies for any of the system alerts (see Modifying System Alert Policies on page 286).
- User Defined Alert Policies. You have the option to create custom alert policies (see Adding Custom Alert Policies on page 287).
- External Defined Alert Policies. There is an API that external sources can use to define alert policies. These are policies you can view (but not modify) through the alerts dashboard (see Viewing External Alert Policies on page 302).

Modifying System Alert Policies

About this task

There are default policies for system alerts, but you can customize which system events should trigger an alert. To modify the system alert policies, do the following:

Procedure

1. Go to the **System Defined** tab of the alerts dashboard and click the name of the target system alert policy (see Alerts Summary View (Prism Central) on page 269).
2. In the **Update Policy** window, do the following in the indicated fields:
   a. **Global Rule**: Uncheck (or check) the box next to the severity to disable (or re-enable) this as an alert.

      All the alerts are enabled (box checked) by default. In most cases this field includes just a single box with the word Critical, Warning, or Info indicating the severity level. Checking the box means this event will trigger an alert of that severity. Unchecking the box means an alert will not be issued when the event occurs. In some cases, such as in the example figure about disk space usage, the event can trigger two alerts, a warning alert when one threshold is reached (in this example 75%) and a critical alert when a second threshold is reached (90%). In these cases you can specify whether the alert should be triggered (check/uncheck the box) and at what threshold (enter a percentage in the box).

   b. **Auto Resolve These Alerts**: Uncheck (or check) the box to disable (or re-enable) automatic alert resolution.

      Automatic alert resolution is enabled for all alert types (where applicable) by default. When this is enabled, the system will automatically resolve alerts under certain conditions such as when the system recognizes that the error has been resolved or when the initiating event has not reoccurred for 48 hours. (Automatic resolution is not allowed for some alert types, and this is noted in the policy window for those types.)

   c. **Exceptions**: To exclude a cluster from this global rule, click the **Add Exceptions** link to open an exceptions line and then select the target cluster from the pull-down list.

      The same check boxes for severity and auto resolve appear. Create a custom rule for the cluster by checking (unchecking) the boxes that should apply to that cluster. Repeat for each cluster you want changed from the global rule.

d. Click the **Save** button to save the changes and return to the alerts dashboard.

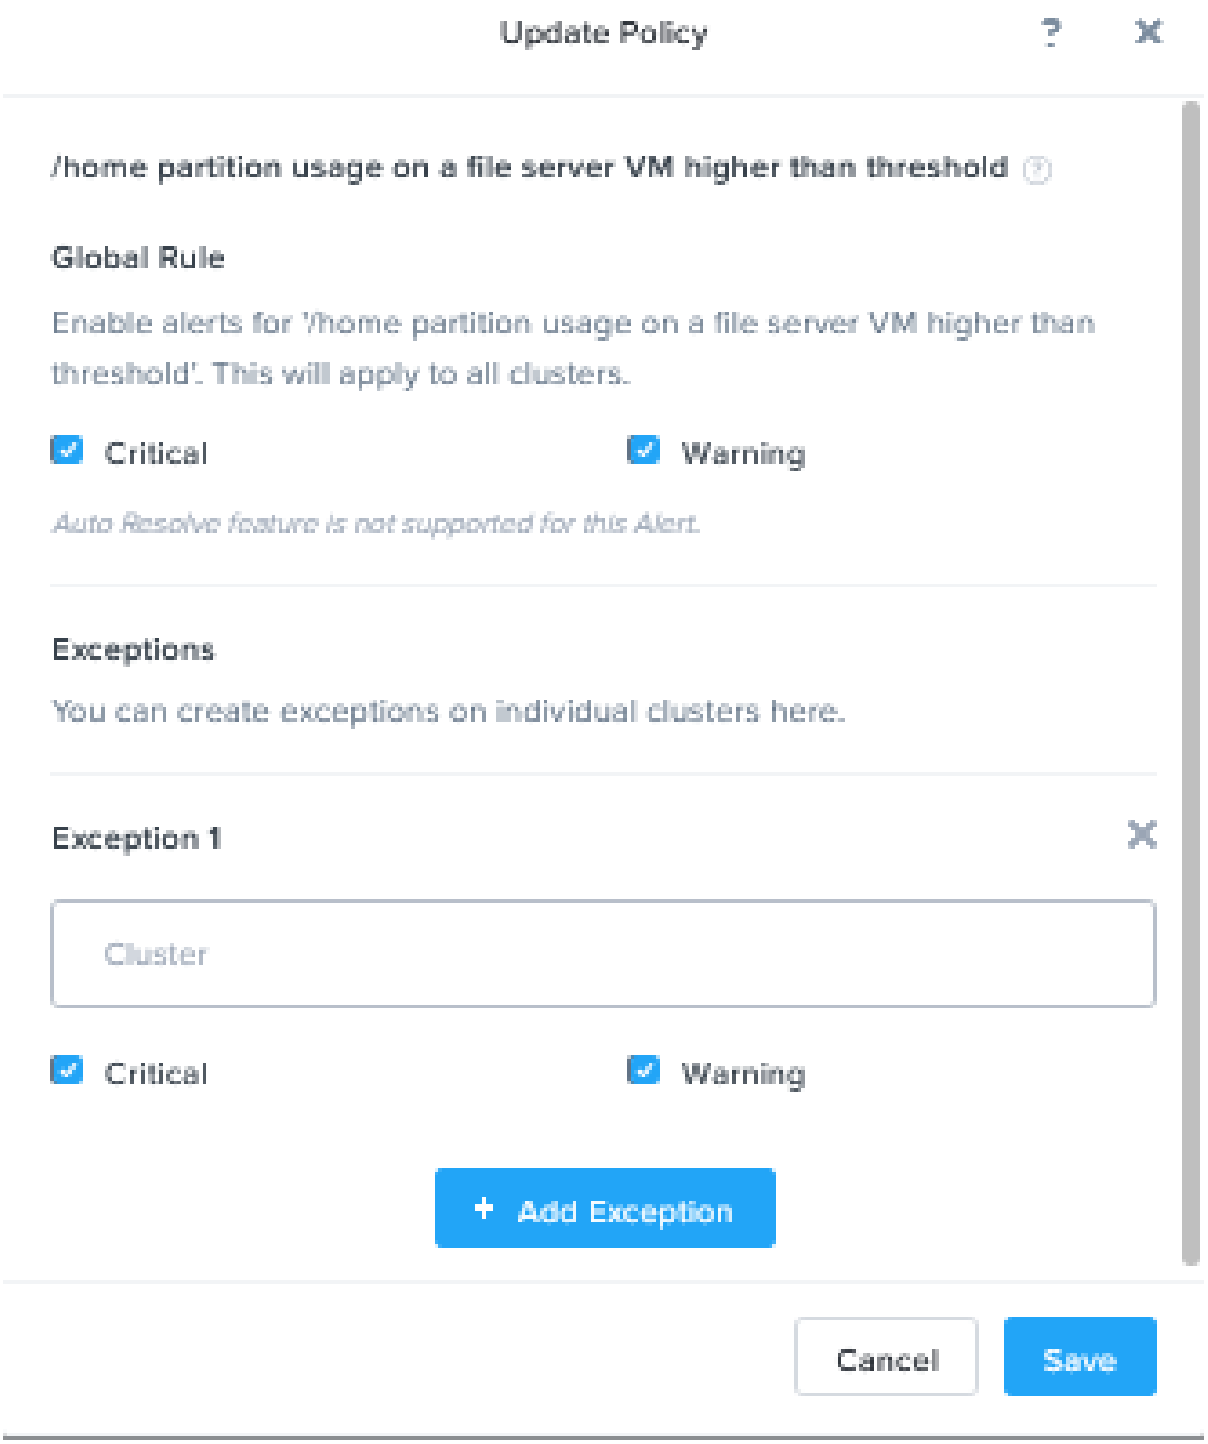

Figure 214: Update Policy Window (example)

Adding Custom Alert Policies

Prism Central provides a mechanism to define alert policies to monitor the entities with respect to the metrics that are critical to your virtual infrastructure. You can define alert policies to monitor entities such as VMs, hosts, and

clusters. You can also define a single alert policy for either all VMs, all hosts, or all clusters that share some common criteria. You can create, update, delete, enable, and disable the alert policies. Each enabled alert policy is executed every five minutes. You can also configure an alert policy so that Prism Central can automatically resolve the alert after the condition causing the alert is resolved.

For a complete description of the metrics available to create custom alert policies, see Alert Metrics on page 292.

Creating Custom Alert Policies

About this task

To create a custom alert policy, do the following.

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and click the **Create Alert Policy** button (see Alerts Summary View (Prism Central) on page 269).
   The **Create Alert Policy** page appears.

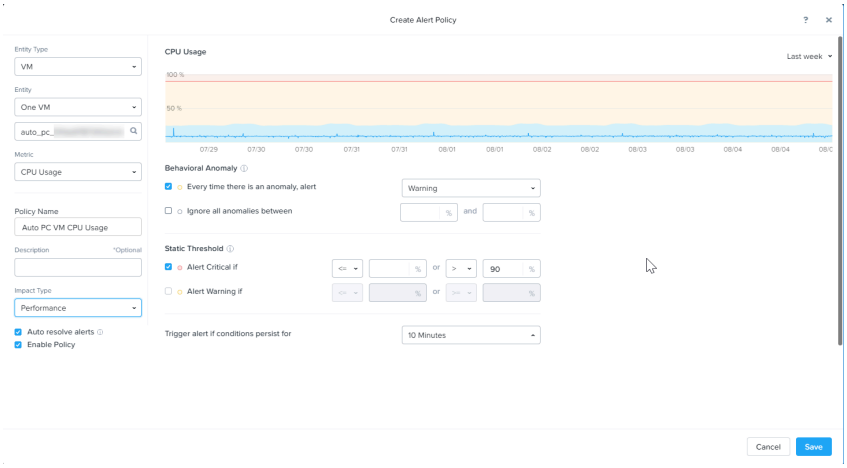

Figure 215: Create Alert Policy page (example)

2. In the left pane, do the following in the indicated fields:

   a. **Entity Type**: Select the entity type (**Host**, **VM**, or **Cluster**) from the pull-down list.

   b. **Entity**: Select the target entity from the pull-down list. The choices allow you to specify a single entity or a group of entity. For example, if the entity type is **Host**, you can specify **All Hosts**, **One Host**, or **All Hosts**

in a Cluster. For all options except **All Hosts**, a search field appears to select the target host, category, or cluster.

c. **Metric**: Select the desired metric from the pull-down list.

   The list of metrics varies depending on the entity type. When the metric is selected, more fields appear in the right pane to specify the conditions for the policy.

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

   The system automatically generates a name for this field based on the policy criteria you have selected, but you can edit the name as desired.

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

f. **Impact Type**: Select the type of impact the alert condition causes to the entity (**Performance**, **Capacity**, **Configuration**, or **Availability**) from the pull-down list. The default impact type is **Performance**.

g. **Auto resolve alerts**: Check this box to automatically resolve the alert generated by the policy after the condition causing the alert no longer persists (the condition must not occur again for 48 hours).

h. **Enable Policy**: Check this box to enable (or uncheck to disable) the policy.

3. In the right pane, do the following in the indicated fields:

a. **<Metric Name>**: The top field displays an anomalies chart from the (machine learned) expected behavior for the specified metric. Hovering over the data displays additional information. Select the chart duration (last 12 hours, last 24 hours, or last week) from the pull-down list on the right.

> **Note:** The following fields allow you to specify that an alert should be generated when the metric deviates from the expected behavior (**Behavioral Anomaly** fields), when the metric crosses a set limit (**Static Threshold**

fields), or both. See Behavioral Learning Tools on page 899 for more information about behavioral anomalies.

Note the following maximum limit of **Behavioral Anomaly** policies for VMs:

- For a small Prism Central VM (4 vCPUs, 16 GB memory) with more than 4000 VMs, you must not have more than 10 **Behavioral Anomaly** policies for VMs.
- For a large Prism Central VM (8 vCPUs, 32 GB memory) with more than 10,000 VMs, you must not have more than 5 **Behavioral Anomaly** policies for VMs.
- For a Prism Central VM in general, the total number of alert policies must not exceed 50.

If the number of policies exceeds the recommended maximum limit, the user-created alerts and behavioral anomaly alerts might be skipped.

b. **Every time there is an anomaly, alert**: Check this box to generate an alert when an anomaly occurs. Select **Critical** or **Warning** from the pull-down list to set this as a critical or warning level alert.

c. **Ignore all anomalies between**: Check this box to trigger the alert only when the anomaly occurs outside a certain range (percentages or units). Select the operator and value for the range.

d. **Alert critical if**: Check this box to generate a critical alert. Select the operator and value for the critical alert.

e. **Alert warning if**: Check this box to generate a warning alert. Select the operator and value for the warning alert. This option is only available if the behavioral anomaly threshold is not set.

f. **Trigger alert if condition persists for**: Select the duration from the pull-down list (0-240 minutes) for which the alert condition must persist before the system generates an alert.

Note: The UDA policies are executed every 5 minutes after checking the data generated during the last 5 minutes. Therefore, it might take up to 10 minutes for an alert generation, even though the **Trigger alert if condition persists for** is set to 0 minutes.

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

Overlapping Policies

If you are creating an alert policy specific to only an entity (a VM, host, or cluster), but you have earlier created a global policy with the same metric for all entities of that entity type (all VMs, all hosts, or all clusters), the policy that you are creating becomes an overlapping policy. For example, if you have created a policy (global-host-policy-1) that triggers a critical alert if the memory usage of any host exceeds 95%. However, you are now creating a policy (host1-policy) that triggers a critical alert if the memory usage (same metric) of host1 (a specific host) exceeds 90%, host1-policy is now an overlapping policy for global-host-policy-1.

When you are creating a new policy (see Creating Custom Alert Policies on page 288), if you select the same entity and metric as an existing policy but different values and you click **Save**, Prism Central displays a message that a similar policy exists as shown in the following screenshot:

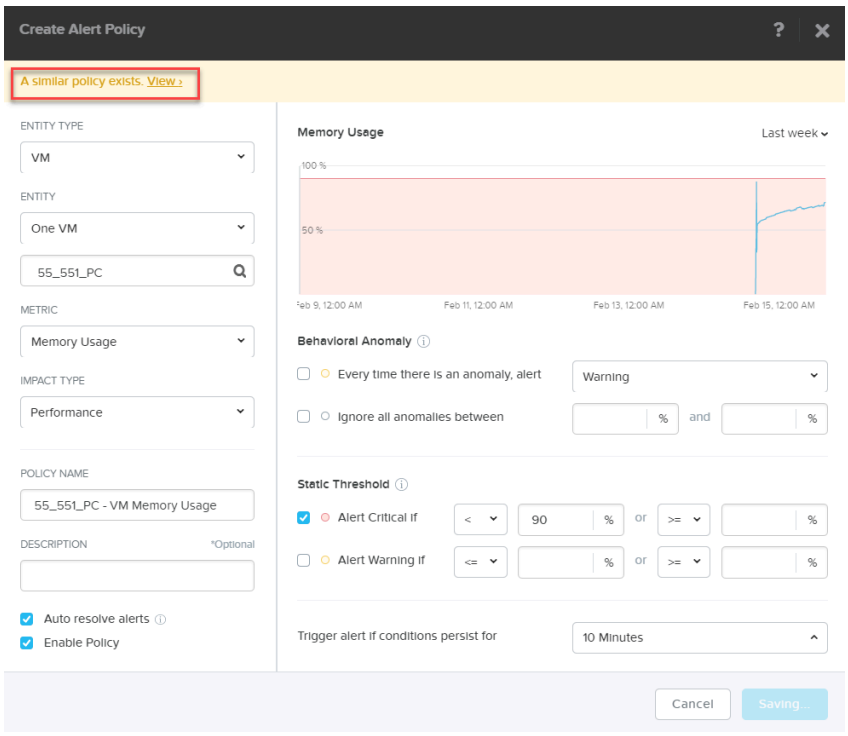

Figure 216: Overlapping Policy

Click **View** to display the overlapping policies and click **Save** to save the overlapping policy.

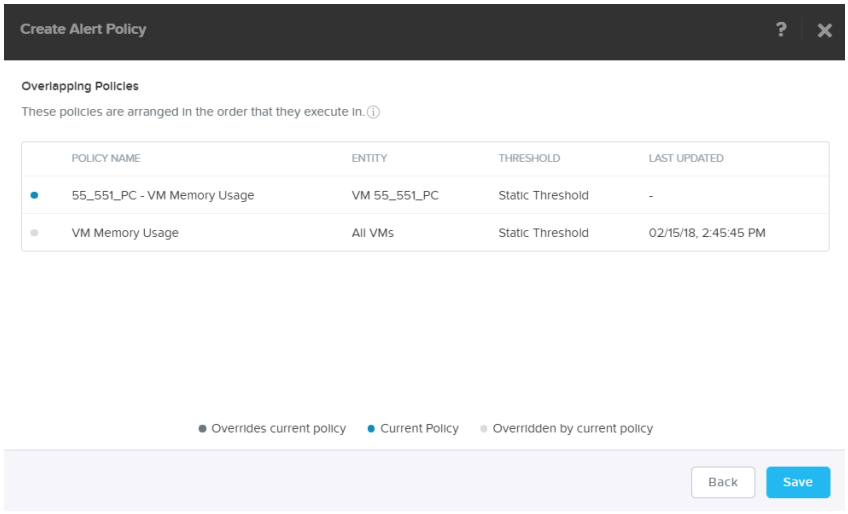

Figure 217: Overlapping Policies List

Precedence of Overlapping Polices

If an overlapping policy exists (policy that applies to a specific entity) for a global policy (policy that applies to all the entities of that entity type), the overlapping policy takes precedence over the global policy.

If multiple policies apply to an entity during policy execution (that is overlapping policies exist), the policy with the highest precedence for that entity is evaluated and all other policies that apply to that entity are ignored. Even if a policy is disabled, the policy might still overlap other policies. If the policy that has the highest precedence is disabled, that policy is not applied and all other policies that apply to that entity are also ignored.

Policies are evaluated in the following order of precedence (highest precedence policy at the top):

1. Policy applied to a specific entity
2. Policy applied to an entity type in a category
3. Policy applied to an entity type in a cluster
4. Policy applied to all entities of an entity type

Precedence of Overlapping Polices in a Category

A category policy (policy that applies to all VMs in a category) takes precedence over a global policy (a policy that applies to all VMs in the cluster). However, a VM can belong to multiple categories. In this case, the category policy that is most recently updated takes precedence over any other category policy that applies to a VM.

For example, consider that you have created a global policy (global-policy-1) for all VMs in the cluster. You have VM-1 that belongs to both the categories: category-1 and category-2. You have created the category-1-pol policy for all VMs in category-1 and the category-2-pol policy for all VMs in category-2. In this case, all three policies apply to VM-1, but only one policy can be applied to a single VM. Because category-1-pol and category-2-pol are category policies, category-1-pol and category-2-pol take precedence over global-policy-1. However, you still have two policies (category-1-pol and category-2-pol) that apply to the same VM (VM-1). In this case, the policy that is updated most recently takes precedence over the other.

Modifying Custom Alert Policies

About this task

To modify a custom alert policy, do the following:

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and click the name of the target custom alert policy (see Alerts Summary View (Prism Central) on page 269).
   The *policy_name* page appears. which includes the same fields as the **Create Alert Policy** page (see Creating Custom Alert Policies on page 288).
2. Make the desired changes in the appropriate fields and then click **Save**.

Deleting Custom Alert Policies

About this task

To delete a custom (user-created) alert policy, do the following:

Procedure

1. Go to the **User Defined** tab of the alerts dashboard and select the target custom alert policy in the list (see Alerts Summary View (Prism Central) on page 269).
2. Select **Delete** from the **Actions** pull-down menu.

Alert Metrics

The following table describes the VM metrics.

Table 128: VM Metrics

CPU Ready Time	

Metric	`hypervisor.cpu_ready_time_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the time a virtual machine waits to use the physical CPU out of the total CPU time allotted to the VM. AHV: Ratio of the VM wait time to the total CPU time allotted to the VM in percentage. ESXi: Value of `cpu.ready.summation` for a VM in percentage.
Supported Hypervisors	AHV and ESXi
CPU Usage	
Metric	`hypervisor_cpu_usage_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the total CPU cycles currently in use by the VM reported by the hypervisor. **Note:** This metric identical to the system created alert policy A3040. ESXi: Value of `cpu.usage.average` reported by ESXi for a VM. Hyper-V: Value of `AvgCpuUtilizationPercent` reported by Hyper-V for a VM.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Memory Usage	
Metric	`hypervisor_memory_usage_ppm`
Rollup	Average
Unit	Percentage (%)
Description	Percentage of the memory used by the VM out of the total configured memory reported by the hypervisor. **Note:** This metric is identical to the system created alert policy A3021. AHV: Memory usage is always 100% because AHV does not support overcommit. ESXi: Value of `mem.usage.average` reported by ESXi for a VM.
Supported Hypervisors	ESXi and AHV
Memory Balloon*	
Metric	`hypervisor_memory_balloon_reclaimed_bytes`
Rollup	Average
Unit	Bytes
Description	Total guest memory in bytes reclaimed through ballooning. **Note:** This metric is similar to the system created alert policy A3024.

Supported Hypervisors	ESXi
Memory Balloon Target*	
Metric	`hypervisor_memory_balloon_reclaim_target_bytes`
Rollup	Average
Unit	Bytes
Description	Target value in bytes for a VM’s memory balloon.
Supported Hypervisors	ESXi
Memory Swap In Rate*	
Metric	`hypervisor_swap_in_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the memory swapped in from the disk to active memory in kilobytes per second.
Supported Hypervisors	ESXi
Memory Swap Out Rate*	
Metric	`hypervisor_swap_out_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the memory swapped out from active memory to the disk in kilobytes per second.
Supported Hypervisors	ESXi
Controller AVG Read I/O Latency	
Metric	`controller_avg_read_io_latency_usecs`
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the read I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller AVG Write I/O Latency	
Metric	`controller_avg_write_io_latency_usecs`
Rollup	Derived from summary rollup type.
Unit	microseconds
Description	Average of the write I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller AVG I/O Latency	
Metric	`controller_avg_io_latency_usecs`
Rollup	Derived from summary rollup type.

Unit	microseconds
Description	Average of the read and write I/O latency of the VM in microseconds reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Read IOPS	
Metric	`controller_num_read_iops`
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM reads in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Write IOPS	
Metric	`controller_num_write_iops`
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM writes in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller IOPS	
Metric	`controller_num_iops`
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the VM read and writes in IOPS reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Read IO Bandwidth	
Metric	`controller_read_io_bandwidth_kBps`
Rollup	Derived from summary rollup type.
Unit	KBps
Description	Read IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller Write IO Bandwidth	
Metric	`controller_write_io_bandwidth_kBps`
Rollup	Derived from summary rollup type.
Unit	KBps
Description	Write IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Controller IO Bandwidth	
Metric	`controller_io_bandwidth_kBps`

Rollup	Derived from summary rollup type.
Unit	KBps
Description	Read and write IO bandwidth of the VM in kilobytes per second reported by the controller.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Network Transmit Packets Dropped	
Metric	`hypervisor_num_transmit_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network transmit packets dropped for the VM.
Supported Hypervisors	ESXi and AHV
Network Receive Packets Dropped	
Metric	`hypervisor_num_receive_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network receive packets dropped for the VM.
Supported Hypervisors	ESXi and AHV
Controller Disk Usage	
Metric	`controller_user_bytes`
Rollup	Average
Unit	bytes
Description	Disk usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
I/O working set size	
Metric	`controller.wss_3600s_union_MB`
Rollup	Average
Unit	MB
Description	Size of the I/O working set size in megabytes.
Supported Hypervisors	All hypervisors
Write I/O working set size	
Metric	`controller.wss_3600s_write_MB`
Rollup	Average
Unit	MB
Description	Size of the I/O working set size writes in megabytes.
Supported Hypervisors	All hypervisors
Read I/O working set size	
Metric	`controller.wss_3600s_read_MB`

Rollup	Average
Unit	MB
Description	Size of the I/O working set size reads in megabytes.
Supported Hypervisors	All hypervisors
Snapshot usage	
Metric	`controller.snapshot_usage_bytes`
Rollup	Average
Unit	bytes
Description	Snapshot usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
Shared Data	
Metric	`controller.shared_usage_bytes`
Rollup	Average
Unit	bytes
Description	Snapshot usage in bytes reported by the controller.
Supported Hypervisors	All hypervisors
Network Bytes Received	
Metric	`hypervisor_num_received_bytes`
Rollup	Summation
Unit	bytes
Description	Number of network bytes received reported by the hypervisor.
Supported Hypervisors	All hypervisors
Network Bytes Transmitted	
Metric	`hypervisor_num_transmitted_bytes`
Rollup	Summation
Unit	bytes
Description	Number of network bytes transmitted reported by the hypervisor.
Supported Hypervisors	All hypervisors

* Experimental only. Might not be supported in the future releases.

Table 129: Host Metrics

The following table describes the host metrics.

CPU Usage	
Metric	`hypervisor_cpu_usage_ppm`
Rollup	Average
Unit	Percentage (%)

Description	Percentage of the total CPU cycles currently in use by the host. **Note:** This metric is identical to the system created alert policy A6414. AHV: Ratio between the sum of kernal, user, and iowait and the sum of kernal, user, iowait, and idle over all the CPUs. For hosts with hyperthreads enabled, the host CPU usage is 10% more than the total value. ESXi: Value of `cpu.usage.average` for a host in percentage. Hyper-V: Value of `LogicalProcessorTotalRuntimePercent` from the output of the command Get-LocalHostPerformance | Convertto-Json -Depth 5 -Compress.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Memory Swap In Rate*	
Metric	`hypervisor_swap_in_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the swap-in memory in kilobytes per second. **Note:** This metric is similar to the system created alert policy A6413. ESXi: Value of `mem.swapinRate.average` for the host.
Supported Hypervisors	ESXi
Memory Swap Out Rate*	
Metric	`hypervisor_swap_out_rate_kBps`
Rollup	Average
Unit	KBps
Description	Rate of the swap-out memory in kilobytes per second. **Note:** This metric is similar to the system created alert policy A6413. ESXi: Value of `mem.swapoutRate.average` for the host.
Supported Hypervisors	ESXi
Memory Usage	
Metric	`hypervisor_memory_usage_ppm`
Rollup	Average
Unit	Percentage (%)

Description	Ratio of the total memory allocated to all VMs over the total memory size on the host. **Note:** This metric is identical to the system created alert policy A6414.
	AHV: Committed memory size from /proc/meminfo over the total memory size. Memory usage is always 100%.
	ESXi: Value of `mem.usage.average` for the host.
	Hyper-V: Amount of used memory over the total memory.
Supported Hypervisors	ESXi, AHV, and Hyper-V
Network Transmit Packets Dropped	
Metric	`hypervisor_num_receive_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network transmit packets dropped.
	ESXi: Value of `net.droppedTx.summation` for the host.
Supported Hypervisors	ESXi and AHV
Network Receive Packets Dropped	
Metric	`hypervisor_num_receive_packets_dropped`
Rollup	Summation
Unit	Packets
Description	Number of the network receive packets dropped.
	ESXi: Value of `net.droppedRx.summation` for the host.
Supported Hypervisors	ESXi and AHV
Controller AVG I/O Latency	
Metric	controller_avg_io_latency_usecs
Rollup	Average
Unit	Microseconds
Description	Average read and write I/O latency of the host in microseconds reported by the controller.
Controller AVG Read I/O Latency	
Metric	`controller_avg_read_io_latency_usecs`
Rollup	Average
Unit	Microseconds
Description	Average read I/O latency of the host in microseconds reported by the controller.

Controller AVG Write I/O Latency	
Metric	`controller_avg_write_io_latency_usecs`
Rollup	Average
Unit	Microseconds
Description	Average write I/O latency of the host in microseconds reported by the controller.
Controller Read IOPS	
Metric	`controller_num_read_iops`
Rollup	Average
Unit	IOPS
Description	Rate of the host reads in IOPS reported by the controller.
Controller Write IOPS	
Metric	`controller_num_write_iops`
Rollup	Average
Unit	IOPS
Description	Rate of the host writes in IOPS reported by the controller.
Controller IOPS	
Metric	`controller_num_iops`
Rollup	Average
Unit	IOPS
Description	Rate of the host reads and writes in IOPS reported by the controller.
* Experimental only. Might not be supported in the future releases.	

Table 130: Cluster Metrics

The following table describes the host metrics.

Controller Read IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller reads in IOPS.
Controller Write IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller writes in IOPS.

Controller IOPS	
Rollup	Derived from summary rollup type.
Unit	IOPS
Description	Rate of the controller reads and writes in IOPS.
Controller AVG Read I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the read I/O latency of the controller in microseconds.
Controller AVG Write I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the write I/O latency of the controller in microseconds.
Controller AVG I/O Latency	
Rollup	Derived from summary rollup type.
Unit	Microseconds
Description	Average of the read and write I/O latency of the controller in microseconds.
Controller I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Read and write IO bandwidth in kilobytes per second reported by the controller.
Controller Read I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Read IO bandwidth in kilobytes per second reported by the controller.
Controller Write I/O Bandwidth	
Rollup	Average
Unit	KBps
Description	Write IO bandwidth in kilobytes per second reported by the controller.
CPU Usage	
Rollup	Average

Unit	PPM
Description	CPU usage of the cluster in PPM reported by the hypervisor.
Memory Usage	
Rollup	Average
Unit	PPM
Description	Memory usage of the cluster in PPM reported by the hypervisor.

Viewing External Alert Policies

External alert policies are those created by applications.

Nutanix provides an API that applications can use to define alert policies. The **External Defined** tab displays a list of the application-defined policies (see Alerts Summary View (Prism Central) on page 269). You cannot modify an externally defined policy, but you can view the policy by clicking the policy name in the list.

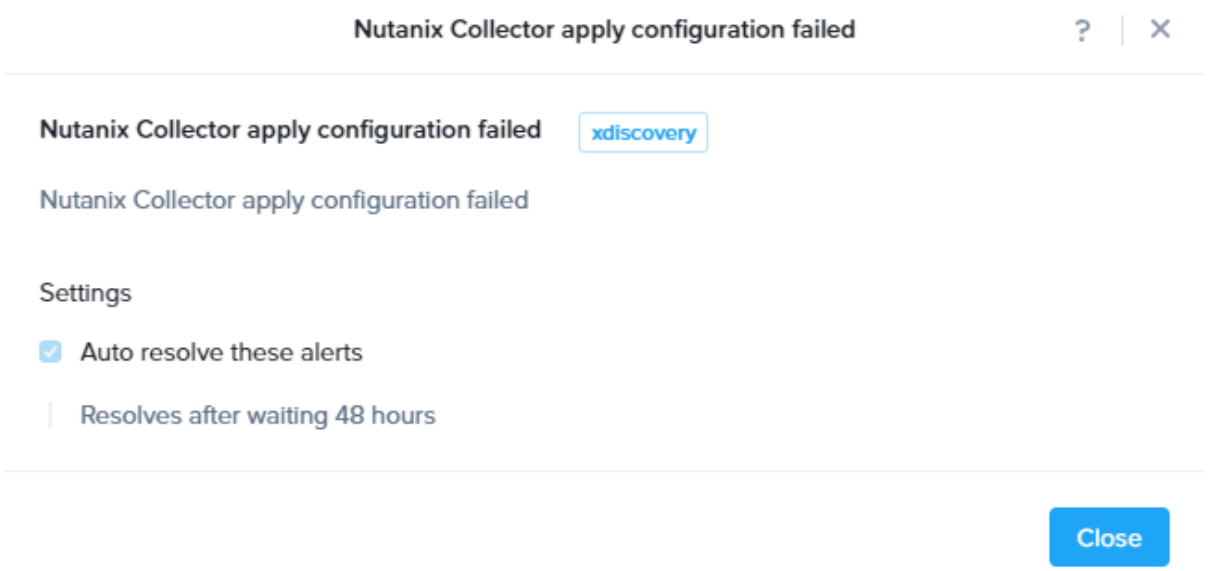

Figure 218: External Defined Policy Window (example 1)

Prism Central Logs

This section lists the logs that reside in a Prism Central VM. These logs record Prism Central-related events and actions. A Prism Central VM runs a subset of the services that run in a Controller VM plus a service about VMs (Metropolis) that is unique to Prism Central. Prism Central logs contain comparable information as the similarly named ones for a Controller VM. For more information, see the Logs section in the *Acropolis Advanced Administration Guide*.

Note:

- The timestamps for all Nutanix service logs are moved to UTC (in ISO 8601:2020-01-01 T00:00:00Z) from Prism version 5.18.

- All operating system logs will not be moved to UTC, hence Nutanix recommends that you set the server local time to UTC.

Table 131: Location: /home/nutanix/data/logs

Log	Contents	Frequency
alert_manager.[out, ERROR, FATAL, INFO, WARNING]	Alert manager process output	
aplos.out	Aplos V3 gateway logs	
aplos_engine.out	Aplos V3 engine logs	
athena.[out, ERROR, INFO, FATAL]	Authentication and identity management activity	
cassandra_monitor.[out, ERROR, FATAL, INFO]	Cassandra database monitor process output	
catalina.out	Catalina/Tomcat for Prism process output	
connection_splicer.[out, ERROR, FATAL, INFO, WARNING]	Internal process connection status	
consolidated_audit.log	Audit logs for VM operations, see Audit Log Events on page 305 for details.	
cron_avahi_monitor.log	Avahi process status	
cron_time_check.log	Check time difference across Prism Central VMs when Prism Central is a multi-VM instance.	every 1 min
disk_usage.log	Disk and inode usage of all partitions on the Prism Central VM	every 1 min
dynamic_ring_changer.[out, ERROR, FATAL]	Metadata migration across nodes activity	
genesis.out	Nutanix software start process output	
hyperint_monitor.[out, ERROR, FATAL, INFO, WARNING]	Hypervisor integration activity	
karbon_core.out (formerly acs_controller.out)	Karbon controller output. This file applies only when Karbon is enabled (see Enabling Karbon on page 993).	
karbon_ui.out (formerly acs_ui.out)	Karbon user interface output. This file applies only when Karbon is enabled.	

Log	Contents	Frequency
metropolis.out	Internal Prism Central service about VMs (Metropolis) to validate cluster choice with scheduler, check out images from the catalog service, and forward fanout requests to Prism Element	
prism_gateway.[out, ERROR, FATAL, INFO]	Prism leader activity	
prism_monitor.[out, ERROR, FATAL, INFO]	Prism (web console, nCLI, REST API) monitor process output	
scavenger.out	Log and core file clean-up status	
search.[out, FATAL]		
search_monitor.[ERROR, INFO, WARNING]		
send-email.log	E-mail alerts sent from the Controller VM	every 1 min
ssh_tunnel.log	Connect status to nsc.nutanix.com for the remote support tunnel	
stats_aggregator.[out, ERROR, FATAL, INFO]	Statistics aggregator process output	
support-info.log	Daily automated support (ASUP) alerts	
using-gflags.log	gflags status	
zeus_config_printer.INFO	Contents of cluster configuration database	
zookeeper_monitor.[out, ERROR, INFO]	Cluster configuration and cluster state activity	

Table 132: Location: /home/nutanix/data/logs/cassandra

Log	Contents
system.log	Cassandra system activity

Table 133: Location:/home/nutanix/data/logs/sysstats

Log	Contents	Frequency	Command
df.info	Mounted filesystems	every 5 sec	df -h
disk_usage.INFO	Disk usage across disks	every 1 hour	du
interrupts.INFO	CPU interrupts	every 5 sec	
iostat.INFO	I/O activity for each physical disk	every 5 sec	sudo iostat

Log	Contents	Frequency	Command
iotop.INFO	Current I/O in realtime	every 5 sec	sudo iotop
lsof.INFO	List of open files and processes that open them	every 1 min	sudo lsof
meminfo.INFO	Memory usage	every 5 sec	cat /proc/meminfo
metadata_disk_usage.INFO	Disk usage for metadata drives	every 5 sec	
mpstat.INFO	CPU activities per CPU	every 5 sec	mpstat
ntpq.INFO	NTP information	every 1 min	ntpq -pn
ping_gateway.INFO	Pings to the default gateway	every 5 sec	ping
ping_hosts.INFO	Pings to other Prism Central VMs (if Prism Central is a multi-VM instance)	every 1 min	ping
sar.INFO	Network bandwidth	every 5 sec	sar -n DEV, -n EDEV
top.INFO	Real-time CPU and memory activity	every 5 sec	top

Table 134: Location: /home/nutanix/data/serviceability/alerts

Log	Contents
num.processed	Alerts that have been processed

Table 135: Location: /var/log

Log	Contents
dmesg	OS start messages
kernel	OS kernel messages
messages	OS messages after starting

Audit Log Events

The following table provides the list of all the VM events that are audited and captured as part of the audit logs.

Table 136: Audit Event List

Event Type	Operation	Description
VmCreateAudit	Create	Created VM
VmReplicateAudit	Create	Replicated VM

Event Type	Operation	Description
VmCloneAudit	Create	Cloned VM
VmDeleteAudit	Delete	Deleted VM
VmRestoreAudit	Restore	Restored VM from snapshot
VmNicAddAudit	Update	Added NIC to VM
VmNicUpdateAudit	Update	Updated NIC of VM
VmNicDeleteAudit	Update	Deleted NIC of VM
VmDiskAddAudit	Update	Added disk to VM
VmDiskDeleteAudit	Update	Deleted disk of VM
VmPowerOnAudit	Power State Change	Powered on VM
VmPowerOffAudit	Power State Change	Powered off VM
VmResetAudit	Power State Change	Reset VM
VmPowerCycleAudit	Power State Change	Power cycled VM
VmPauseAudit	Power State Change	Paused VM
VmSuspendAudit	Power State Change	Suspended VM
VmResumeAudit	Power State Change	Resumed VM
VmGuestShutdownAudit	Power State Change	Guest shutdown on VM
VmGuestRebootAudit	Power State Change	Guest rebooted on VM
VmAcpiShutdownAudit	Power State Change	ACPI shutdown VM
VmAcpiRebootAudit	Power State Change	ACPI rebooted VM
VmPowerStateReconcileAudit	Power State Change	Reconciled VM power state
HAFailoverStartAudit	HA	Started HA failover on host
HAFailoverCompleteAudit	HA	Completed HA failover on host
EnteringMaintenanceModeAudit	Update	Entering maintenance mode on host
MaintenanceModeAudit	Update	Host in maintenance mode
ExitMaintenanceModeAudit	Update	Exited maintenance mode on host
VmHostAffinityUpdateAudit	Update	Updated host affinity of VM
VmDiskResizeAudit	Update	Resized disk of VM
VmDiskUpdateAudit	Update	Updated disk of VM
VmUpdateAudit	Update	Updated VM
VmExportAudit	Create	VM exported to OVA
OvaCreateAudit	Create	OVA created
OvaDeleteAudit	Delete	OVA deleted
OvaUpdateAudit	Update	OVA name updated
OvaConcatenateAudit	Update	OVA parts concatenated
VmMigrateAudit	Migrate	Migrated VM from one host to another host

Event Type	Operation	Description
VsCreateAudit	Create	Virtual switch created
VsUpdateAudit	Update	Virtual switch updated
VsDeleteAudit	Delete	Virtual switch deleted
MigrateBridgeToVirtualSwitchAudit	Migrate	Migrated bridge to virtual switch
IPAssignmentToVMAudit	Update	IP assigned to VM
IPUnassignmentToVMAudit	Update	IP unassigned from VM
SpanCreateAudit	Create	SPAN session created
SpanUpdateAudit	Update	SPAN session updated
SpanDeleteAudit	Delete	SPAN session deleted
VmHostAffinityPolicyCreateAudit	Create	VM Host Affinity Policy created
VmHostAffinityPolicyUpdateAudit	Update	VM Host Affinity Policy updated
VmHostAffinityPolicyDeleteAudit	Delete	VM Host Affinity Policy deleted
VmHostAffinityPolicyReEnforceAudit	Update	VM Host Affinity Policy Re-enforced
ImageCreateAudit	Create	Created Image
ImageDeleteAudit	Delete	Deleted Image
ImageUpdateAudit	Update	Updated Image

Alerts/Health checks

Cluster

Table 137: 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 138: 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 139: 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 140: 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 141: 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 142: 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 143: AHV Secondary IP Ping Check from PC [3071]

Name	AHV Secondary IP Ping Check from PC
Description	Checks if PC can ping secondary IP addresses of all the hosts of all the registered clusters.
Cause	The AHV secondary IP is down or there is a network connectivity issue.
Impact	If the AHV's secondary IP is not pingable from the PC, then advanced networking may encounter issues, if enabled and configured to use the corresponding virtual switch.
Resolution	Ensure that the AHV is running and that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	12446

Table 144: 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 145: 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 146: 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 147: 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 148: 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 149: 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 150: 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 151: 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 152: 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 153: 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 154: 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 155: 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 156: 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 157: 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 158: CVM IP address is unreachable [103088] [A103088]

Name	Inter-CVM connectivity check
Description	Check that all CVMs are reachable via ping.
Alert message	CVM `dest_ip` is unreachable
Cause	The Controller VM is down or there is a network connectivity issue.
Impact	Storage may be unavailable and/or performance issues may be observed.
Resolution	If the Controller VM does not respond to ping, turn it on. Ensure that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	1635
Severity	Warning

Table 159: 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 160: Checks if AOS version of the target site configured for synchronous replication is on AOS version >= 5.17 [110021] [A110021]

Name	Remote AOS Version Check
Description	Checks if the target site configured is on AOS version >= 5.17 to support Synchronous Replication
Alert message	For protection rule: `protection_rule_name`, target cluster(s): `remote_site_list` have version older than 5.17
Cause	Target site is on AOS version < 5.17
Impact	Synchronous Replication will not be supported.
Resolution	Upgrade the cluster on the target site to AOS version >= 5.17. It is recommended to have both the source and target sites at the same version.
KB Article	7364
Severity	Warning

Table 161: Round-trip travel (RTT) to target Prism Element(s) is greater than 5ms, the maximum value allowed for synchronous replications. [110022] [A110022]

Name	Remote Site Latency Check
Description	Check if the Round-trip travel (RTT) is less than 5ms, the maximum supported for Synchronous Replication
Alert message	Round-trip travel (RTT) to target Prism Element(s) `err_msg` is greater than `max_permitted_latency`
Cause	Target cluster is unreachable or Round-trip travel (RTT) to target cluster is greater than 5ms.
Impact	Synchronous Replication will be affected.
Resolution	Ensure that the target cluster is reachable and the RTT is less than 5ms or choose another target cluster.
KB Article	7366
Severity	Warning

Table 162: 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 163: Stale secondary synchronous replication configuration found. [110456] [A110456]

Name	Stale secondary Synchronous Replication configuration check
Description	Check if there are any stale secondary synchronous replication configurations present on Prism Central.
Alert message	Stale secondary synchronous replication configuration found for vms `err_msg`.
Cause	Stale secondary synchronous replication configuration left behind.

Impact	Stale entries could cause unwanted tasks to be generated and they cannot be cleaned up automatically on the secondary as it could affect the VM stretch state hence require a manual cleanup.
Resolution	Cleanup the stale secondary synchronous replication configuration left behind on Prism Central.
KB Article	9716
Severity	Warning

Table 164: Check if the Recovery Plan configured with Witness has asynchronously protected entities or has entities with Synchronous Replication paused [110458] [A110458]

Name	Witness configured Recovery Plan has asynchronously protected entities and entities with Synchronous Replication paused check
Description	Check if the Recovery Plan configured with Witness has asynchronously protected entities or has entities with Synchronous Replication paused
Alert message	`alert_msg`
Cause	Recovery Plan configured with Witness has asynchronously protected entities or has entities with Synchronous Replication paused
Impact	The entities will not be managed by Witness
Resolution	Ensure all the entities that are part of the Witness configured Recovery Plan are Synchronously protected
KB Article	11000
Severity	Warning

Table 165: Stale primary synchronous replication parameters found. [110460] [A110460]

Name	Stale primary Synchronous Replication configuration check
Description	Check if there are any stale primary synchronous replication configurations present on Prism Central.
Alert message	Stale primary synchronous replication parameters found for vms `err_msg`.
Cause	Stale primary synchronous replication configuration left behind probably because of a restore operation.
Impact	Stale configurations could cause unwanted tasks to be generated and they cannot be cleaned up automatically on the primary as it could affect the VM stretch state hence require a manual cleanup.
Resolution	Delete the stale VM(s) on primary site.
KB Article	11002
Severity	Warning

Table 166: 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 167: 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 168: 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 169: 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 170: 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 171: 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 172: 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 173: 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 174: 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 175: 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 176: Failure to copy Image to Cluster [111072] [A111072]

Name	Image copy to Cluster check
Description	Detect if image copy to cluster is failing
Alert message	Image could not be copied to Cluster
Cause	Cluster may be unreachable from PC or Catalog service may be down on Cluster.
Impact	Image will not be available on the Cluster.

Resolution	Check Catalog service on cluster and connectivity to PC.
Severity	Warning

Table 177: 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 178: 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 179: VM deployment from Template to Prism Element Cluster failed [111091] [A111091]

Name	VM Deployment from Template check
Description	Report on deploy VM to cluster status
Alert message	VM from Template could not be fully deployed to the Cluster
Cause	A PC service or Destination Cluster may be unreachable or undergoing maintenance.
Impact	Some VM(s) will not be present on the Cluster.
Resolution	Check connectivity between cluster, Prism Central and the destination Prism Element. Ensure PE-PC connection is healthy on the destination PE cluster and reinitiate the operation to make up for the partial failure.
KB Article	11356
Severity	Warning

Table 180: 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.
KB Article	2159
Severity	Critical
Severity	Warning

Table 181: 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 182: 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 183: 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 184: 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 185: 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 186: 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 187: 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 188: 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 189: 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 190: 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 191: 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 192: 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 193: 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 194: 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 195: 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 196: 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 197: 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 198: 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 199: 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 200: 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 201: 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 202: 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 203: 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 204: 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 205: 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 206: 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 207: 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 208: 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 209: 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 210: 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 211: 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 212: 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 213: 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 214: 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 215: 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 216: 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 217: 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`.
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 218: 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 219: 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 220: 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 221: Pulse is disabled [140002] [A140002]

Name	Pulse Disabled
Description	Pulse is disabled on the paired Prism Central.
Alert message	Pulse is disabled on the paired Prism Central `cluster_uuid`
Cause	Pulse enablement might be missed or disabled.
Impact	Monitoring Nutanix Service Health cannot be performed.
Resolution	Enable Pulse on current cluster and all of its registered clusters.
KB Article	6377
Severity	Info

Table 222: 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 223: 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 224: 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 225: 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 226: 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 227: 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 228: 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 229: 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 230: 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 231: 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 232: 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 233: 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 234: 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 235: 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 236: 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 237: 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
KB Article	11548
Severity	Critical
Severity	Warning
Severity	Info

Table 238: 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 239: 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 240: 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 241: 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 242: 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 243: 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 244: 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 245: 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 246: 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 247: 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 248: 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 249: 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 250: 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 251: 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 252: 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 253: 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 254: 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 255: 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 256: File server is in partial network state. [160148] [A160148]

Name	File Server Partial Network State
Description	File server is in partial network state.
Alert message	`message`
Cause	Scale-out operation performed on fileserver with multi-vlan configuration.
Impact	File server is in partial network state.
Resolution	Refer to KB article 12731. Contact Nutanix support if issue still persists or assistance needed.
KB Article	12731
Severity	Warning

Table 257: Cluster Connectivity Status [200000] [A200000]

Name	Cluster Connectivity Status
Description	Tests whether the cluster connectivity is fine
Alert message	`component` data from cluster `cluster_name` is not up-to-date.
Cause	The connection between Prism Central and Prism Element is down or Services on Prism Central are crashing.
Impact	Cluster data shown in the Prism Central is not up to date.
Resolution	Ensure that cluster network connectivity is up and all CVM services are up.
KB Article	3379
Severity	Warning

Table 258: 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 259: 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 260: Prism Central vCPU Availability Check [200301] [A200301]

Name	Prism Central vCPU Availability Check
Description	Checks if the number of vCPUs is sufficient for the number of VM entities in Prism Central.
Alert message	The Prism Central does not have enough vCPUs for the number of VM entities it has.
Cause	Prism Central doesn't have enough vCPUs to support the current number of VMs
Impact	Prism Central performance may be degraded.
Resolution	1) Reduce the number of VM entities in Prism Central or 2) If using 1 node small Prism Central scale up to 1 node large Prism Central or scaleout to 3 node small Prism Central. 3) If using 1 node large Prism Central scaleout it to 3 node large Prism Central.
KB Article	4381
Severity	Warning

Table 261: Insufficient Space for UVMs deployed on PC [200302] [A200302]

Name	PC Sufficient Disk Space Check
Description	Checks if the amount of storage is sufficient for the number of VM entities in Prism Central.
Alert message	The PC does not have enough storage for the number of VM entities it has.
Cause	Too many VMs in Prism Central for the amount of storage in Prism Central.
Impact	Prism Central may run out of disk space to store data.
Resolution	Reduce the number of VM entities in Prism Central or provide more storage for Prism Central.
KB Article	4340
Severity	Warning

Table 262: Prism Central Memory Availability Check [200303] [A200303]

Name	Prism Central Memory Availability Check
Description	Checks if the amount of memory is sufficient for the number of VM entities and the services enabled in Prism Central.
Alert message	The Prism Central does not have enough memory for the current number of VMs.
Cause	Prism Central doesn't have enough memory to support the current number of VMs
Impact	Services running in Prism Central may run out of memory and crash.
Resolution	1) Reduce the number of VM entities in Prism Central or 2) If using 1 node small Prism Central scale up to 1 node large Prism Central or scaleout to 3 node small Prism Central. 3) If using 1 node large Prism Central scaleout it to 3 node large Prism Central.
KB Article	4315
Severity	Warning

Table 263: Prism Central VM Limit Check [200304] [A200304]

Name	Prism Central VM Limit Check
Description	Checks if the number of VM entities is within the limit.
Alert message	The Prism Central cannot handle these many VM entities.
Cause	Too many VMs in Prism Central.
Impact	Prism Central performance may be degraded.
Impact	Prism Central may run out of disk space to store data.
Impact	Services running in Prism Central may run out of memory and crash.
Resolution	Reduce the number of VM entities in Prism Central.
KB Article	4296
Severity	Warning

Table 264: 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 265: Prism Central cluster version check [200310]

Name	Prism Central cluster version check
Description	Checks if Prism Central is running GA version.
Cause	Non GA version detected on Prism Central.
Impact	Prism Central is running a Non GA version.
Resolution	Nutanix recommends running officially released GA versions of Prism Central software. For assistance refer to KB 6325
KB Article	6325

Table 266: Prism Central Version EOL [200311] [A200311]

Name	Prism Central version EOL check
Description	Checks if Prism Central is running an EOL version.
Alert message	`info_msg`. Upgrade Prism Central at the earliest.
Cause	EOL version running on Prism Central
Impact	Support may not be able to resolve issues since this version will no longer be maintained.
Resolution	Move to a recent released version immediately, to avoid disruption in support. You can refer to the EOL Schedule on the Nutanix support portal.
KB Article	6325
Severity	Info

Table 267: Inadequate Resource Configuration for the Prism Central VM [200327] [A200327]

Name	Prism Central Resource Resize Check
Description	Configured resource for the Prism Central VM is inadequate.
Alert message	`inadequate_resources` for the `pc_vm_type` Prism Central VM is inadequate on `non_resized_pc_vm_ips`.

Cause	The current feature capabilities of Prism Central require resource on the Prism Central VM to be increased for optimum performance.
Impact	Prism Central performance may be degraded.
Resolution	Increase the resource of the Prism Central VM. Please refer KB-8932 for guidance.
KB Article	8932
Severity	Warning

Table 268: File system inconsistencies are detected. [200329] [A200329]

Name	Prism Central EXT4-fs Error Check
Description	Captures Prism Central EXT4-fs error messages
Alert message	EXT4 file system errors are detected on PCVM: `pcvm_ip`: `alert_msg`
Cause	File system inconsistencies are present on the Prism Central VM node.
Impact	Inability for the Prism Central VM to boot or for the upgrade pre-checks to run.
Resolution	Look for any problems in the file system. Review KB 9622.
KB Article	9622
Severity	Critical

Table 269: Prism Central Backup Sync has failed. [200331] [A200331]

Name	Prism Central Backup Sync check
Description	Checks if Prism Central backup sync is stopped for more than 3 hours
Alert message	Prism Central Backup Sync to these Prism Element(s) : `pe_names` is halted for the last : `delays`
Cause	Prism Central Backup Sync to designated Backup Prism Element(s) is halted for more than 3 hours.
Impact	Prism Central may not be restorable during disaster.
Resolution	Refer to KB 10788 for further details.
KB Article	10788
Severity	Critical

Table 270: Prism Central backup limit reached. [200332] [A200332]

Name	Prism Central backup limit check
Description	Checks if Prism Central backup limit reached.
Alert message	Prism Central backup limit reached on these Prism Elements : `pe_list`
Cause	Prism Central backup limit reached.
Impact	Prism Central backup to Prism Elements is paused.

Resolution	Refer to KB 12020 for further details.
KB Article	12020
Severity	Critical

Table 271: 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 272: 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 273: 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 274: 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 275: 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 276: 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 277: 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 278: 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 279: 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	Underlying issues with PE (Prism Element) - PC (Prism Central) connectivity. {reason}.
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	Validate if PE-PC is reachable from each other and if the required ports are open. Please refer to KB 6970 for additional details.
Resolution	Check that no services are crashing on the PC or PE. Please refer to KB 6970 for additional details.
Resolution	Contact Nutanix Support to troubleshoot this issue further. Please refer to KB 6970 for additional details.
KB Article	6970
Severity	Warning

Table 280: AZ Name Update Failure [200803] [A200803]

Name	Availability Zone Name Update Failure
Description	Raised if AZ Name update thread failed for one or more clusters
Alert message	AZ Name Update failed for the following clusters : `ip_list`
Cause	The Availability Zone name update might fail because of IDF or PC connectivity issues.
Impact	Remote PC has outdated PC name in AZ entity.
Resolution	Refer to KB 12211 for further details.
KB Article	12211
Severity	Critical

Table 281: 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 282: 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 283: 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 284: 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 285: 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 286: 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 287: 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 288: 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 289: 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 290: 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 291: 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 292: 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 293: 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 294: 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 295: Object Store deployment is not supported on this cluster. [805001] [A805001]

Name	Objects Store Deployment Check
Description	Checks if Object Store is deployed on an unsupported AOS cluster.
Alert message	The registered cluster(s) `affected_clusters_list` is/are on one of the versions between 5.18 and 5.18.0.4 which is unsupported for Object Store Deployment.
Cause	One or more Prism Element cluster(s) attached to this Prism Central is on version(s) between 5.18 and 5.18.0.4 which is unsupported for Object Store Deployment.
Impact	Object Store will not work as expected.
Resolution	Upgrade the cluster to 5.18.0.5 or above.
KB Article	10000

Severity	Critical

Table 296: Deployment failed in Foundation Central [805401] [A805401]

Name	Foundation Central deployment failure
Description	Deployment failure in Foundation Central will raise this alert
Alert message	`alert_msg`
Cause	Deployment process in Foundation Central has failed.
Impact	The remote nodes or cluster in the deployment request might not be ready for use.
Resolution	Please refer to the deployment history page in Foundation Central for more details.
KB Article	12863
Severity	Warning

Controller VM

Table 297: 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 298: 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 299: 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 300: 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 301: {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 302: 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 303: 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 304: Prism Central VM CPU Load High [6517] [A6517]

Name	PCVM CPU Load High
Description	Check that Prism Central VM CPU load is not high.
Alert message	CPU load in or Prism Central VM `ip_address` is high. Load average over 5 Minutes = #load_average_5min@, number of vCPUs in CVM or Prism Central VM = `num_procs`.
Cause	The CPU load on the Prism Central VM has been high.

Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix Support for diagnosis. CPU on the Prism Central VM may need to be increased.
KB Article	5227
Severity	Critical
Severity	Warning

Table 305: Prism Central VM disk usage high [101059] [A101059]

Name	Prism Central VM high disk space usage
Description	Check that disk space usage on the Prism Central VM is within limits.
Alert message	`alert_msg`
Cause	High disk usage in the Prism Central VM.
Impact	Prism Central VM may run out of storage space to store data.
Resolution	Refer to KB 5228 for further details.
Prism Central VM disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

Table 306: 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 307: 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 308: 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 309: 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 310: 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 311: 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 312: 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 313: 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 314: 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 315: 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 316: 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 317: 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 318: 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 319: 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 cluster and check the disk health. Contact the Hardware Vendor to replace the bad disk.
KB Article	8031
Severity	Critical

Table 320: Physical Disk Drive Might Have Failed. [130042] [A130042]

Name	Physical Disk Maybe Bad Status
Description	Physical drive might be failing.
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* might have failed. Node serial is *node_serial_number* and Node position *node_position*
Cause	The drive might be failing.
Impact	Cluster storage capacity is reduced.
Resolution	Inspect drive health. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 321: 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 322: 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 323: 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 324: 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 325: Prism Central VM platform disk space usage high [200328] [A200328]

Name	Prism Central VM platform disk space usage check
Description	Check that platform disk space usage on the Prism Central VM is within limits.
Alert message	Platform disk space usage in Prism Central VM `svm_ip` exceeds `percentage_exceed`% for disk(s): `disk_paths`.
Cause	High disk usage in the Prism Central VM.
Impact	Prism Central VM may run out of storage space to store data.
Resolution	Refer to KB 5228 for further details.

Disk Usage Percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

Table 326: 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 327: 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 328: 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 329: 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 330: 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 331: Atlas was unreachable to apply Flow rule [200610] [A200610]

Name	Flow rule failed to reach Atlas
Description	Programming a Flow rule failed to reach Atlas.
Alert message	Flow policy operation failed for `rule_name`. `reason`

Cause	Atlas service is unreachable in PC or the Microseg service in PC is unreachable
Impact	VMs will not be protected by the Flow rule.
Resolution	Check Atlas service in PC, Microseg service in PC, Microseg to Atlas connection.
Severity	Warning

Table 332: 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

DR

Table 333: 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 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 334: 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 335: Availability Zone configured in Protection Policy is not accessible [110401] [A110401]

Name	Availability Zone Accessibility Check
Description	Check if availability zones configured in protection policy are accessible.
Alert message	Availability Zone(s) `fault_azs` configured in the Protection Policy `protection_rule_name` is/are not accessible.
Cause	Availability zone configured in the protection rule is not accessible.
Impact	Availability zones of the backup site will not be reachable for taking backup.
Resolution	Check the health status of the configured availability zone.
KB Article	5685
Severity	Critical

Table 336: Protection Policy Max entities per Category Check Failed. [110402] [A110402]

Name	Protection Policy Max entities Per Category Exceeded
Description	Checks if the entity count for a category specified in Protection Policy exceeds the maximum allowed limit.
Alert message	Maximum number of entities for a category in a Protection Policy should not exceed `max_entity_count`. Following categories exceeds entities limit : `categories`.
Cause	Number of entities for the specified categories in the Protection Policy exceeds the limit.
Impact	Specified category will not be considered for the recovery as the Recovery Plan supports categories with limited number of entities.
Resolution	Reduce the protected entity count for the specified categories in the Protection Policy.
KB Article	6306
Severity	Warning

Table 337: Disconnected Availability Zones are affecting some entities. [110403] [A110403]

Name	Entities Affected By Disconnected Availability Zones
Description	Lists all entities which will be affected due to the disconnected Availability Zones.
Alert message	`entity_type` of name `name` has been affected by disconnected availability zones.
Cause	Availability zone configured in the protection rule has been disconnected.
Impact	Operations related to the affected entities will be affected.
Resolution	Check the health status of the configured availability zone.
KB Article	6283
Severity	Info

Table 338: 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 339: 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 340: 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 341: 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 342: 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 343: 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 344: 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 345: 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 346: 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 347: 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 348: 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 349: 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 350: 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 351: 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 352: 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 353: 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 354: 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 355: 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 356: 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 357: 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 358: 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 359: 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 360: 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 361: 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 362: 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 363: 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 364: 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.

Cause	Either the source or destination hypervisor is ESXi and one or more network name(s) has unicode characters.
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).
Resolution	Delete and recreate the network mappings associated with such network(s).
KB Article	10243
Severity	Warning

Table 365: 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 366: 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 367: 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 368: 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 369: 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 370: 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 371: 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 372: 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 373: 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 374: 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 375: 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 376: 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 377: 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 378: 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 379: 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 380: 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 381: 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 382: 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 383: 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 384: 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 385: 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 386: 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 387: 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 388: 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 389: 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 Element 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 390: 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 391: 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 392: 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 393: 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 394: 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 395: 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 396: 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 397: 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 398: 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 399: 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 400: 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 401: 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 402: 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 403: 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 404: 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 405: 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 406: 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 407: 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 408: 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 409: 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 410: 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 411: 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 412: 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 413: 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 414: 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.
KB Article	12108
Severity	Warning

Table 415: 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 416: 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 417: 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 418: 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 419: 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 420: 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 421: 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 422: 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 423: 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 424: 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 425: VM Recovery Point Replication Failed [130174] [A130174]

Name	Recovery Point Replication Failed
Description	VM Recovery Point replication failed.
Alert message	Failed to replicate VM 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	VM 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 VM 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	VM 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 VM 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 VM Recovery Point, retry replication again. For a scheduled VM Recovery Point replication, ensure all the remote clusters are healthy. Then wait for the scheduled Recovery Point replication.
Resolution	Retry the VM 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 426: 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 427: 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 428: 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 429: 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 430: Snapshot has Entities from the Storage Container that have Deduplication Enabled [130181] [A130181]

Name	Entity Centric 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 431: 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 432: 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 433: VM Protection Failed [130186] [A130186]

Name	Protection Policy Processing 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 (Nutanix Disaster Recovery).
Resolution	Upgrade cluster to version 5.19 or above.
Resolution	Resolve the stated reason for the failure.
KB Article	8472
Severity	Warning

Table 434: 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 435: 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 436: 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 437: 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 438: 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 439: 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 440: 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 441: VSS Snapshot is not supported for some VMs. [130200] [A130200]

Name	VSS Snapshot not supported for Protection Domain
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 442: 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 443: 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 444: 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 445: 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 446: 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 447: 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.
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 448: 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 449: 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 450: Failed to transition to a higher frequency snapshot schedule. [130213] [A130213]

Name	Protection Domain 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.
KB Article	11501
Severity	Warning

Table 451: Nearsync Replication is lagging for Protection Domain snapshot [130214] [A130214]

Name	Nearsync Replication is Lagging for the protection domain
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.
KB Article	11502
Severity	Warning

Table 452: VM failed to transition to a higher frequency snapshot schedule. [130215] [A130215]

Name	VM failed to transition to higher frequency snapshot schedule
Description	VM has failed to transition to a higher frequency snapshot schedule.
Alert message	Failed to enable nearsync for VM `vm_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	VM will automatically attempt to transition to a high frequency snapshot schedule after system constraints are resolved.
KB Article	11492
Severity	Warning

Table 453: Nearsync Replication is lagging for recovery point [130216] [A130216]

Name	Nearsync Replication is Lagging for the VM
Description	Replication of the nearsync recovery point is lagging.
Alert message	Nearsync replication for `vm_name` to availability zone `availability_zone_physical_name` is lagging by `replication_lagging_seconds` seconds.
Cause	Replication of the nearsync recovery point is slow.
Impact	Configured RPO to the remote site might get affected.
Resolution	Ensure that the number of entities are within the supported limit.
KB Article	11491
Severity	Warning

Table 454: 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 455: 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 456: 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 457: 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 458: 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 459: 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 460: 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 461: 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 462: 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 463: 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 464: Automatic Synchronous Replication pause is triggered by Witness [130353] [A130353]

Name	Automatic Synchronous Replication Pause is triggered by Witness
Description	Automatic Synchronous Replication pause is triggered by Witness for {vm_count} VMs protected by Protection Rule '{protection_rule_name}'
Alert message	Automatic Synchronous Replication pause has been triggered for *vm_count* VMs. '*reason*'.
Cause	Network connection to remote site or to Witness is lost probably because of a network partition or the remote site being unavailable.
Cause	One or more internal services are down or not working as expected.
Impact	Synchronous Replication for {vm_count} VMs will be paused.
Resolution	Ensure the remote site and Witness are reachable and resume Synchronous Replication.
Resolution	Contact Nutanix support.
Severity	Info

Table 465: Automatic Failover triggered by Witness [130354] [A130354]

Name	Automatic Failover triggered by Witness
Description	Automatic Failover triggered by Witness
Alert message	Unplanned Failover of the Recovery Plan *recovery_plan_name* has been triggered by Witness. '*reason*'.
Cause	Automatic Failover triggered by Witness probably because of a network partition or source site being down.
Cause	One or more internal services are down or not working as expected.
Impact	Effected entities are restored on remote site.
Resolution	Restore the connectivity between source and remote sites and the restored entities can be migrated back to source site.
Resolution	Contact Nutanix support.
Severity	Info

Table 466: 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 467: VM is part of multiple Recovery Plans configured with Witness [130356] [A130356]

Name	Same VM is in multiple Witness configured Recovery Plans
Description	VM is part of multiple Recovery Plans configured with Witness
Alert message	The VM '*vm_name*' is part of multiple Recovery Plans configured with Witness. *reason*
Cause	Same VM is part of multiple Recovery Plans configured with Witness
Impact	The VM will be managed only by one of the Recovery Plans configured with Witness
Resolution	Ensure the VM is part of only one Witness configured Recovery Plan
KB Article	11001
Severity	Warning

Table 468: 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 Nutanix Disaster Recovery.
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 469: 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 470: 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 471: 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 472: 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 473: Volume Group is part of multiple Recovery Plans configured with Witness [130364] [A130364]

Name	Same Volume Group is in multiple Witness configured Recovery Plans
Description	Volume Group is part of multiple Recovery Plans configured with Witness
Alert message	The Volume Group '*volume_group_config_name*' is part of multiple Recovery Plans configured with Witness. *reason*
Cause	Same Volume Group is part of multiple Recovery Plans configured with Witness
Impact	The Volume Group will be managed only by one of the Recovery Plans configured with Witness
Resolution	Ensure the Volume Group is part of only one Witness configured Recovery Plan
KB Article	11001
Severity	Warning

Table 474: Automatic Synchronous Replication pause is triggered by Witness [130365] [A130365]

Name	Automatic Pause Synchronous Replication is triggered by Witness
Description	Automatic Synchronous Replication pause is triggered by Witness for {entity_info_msg} protected by Protection Rule '{protection_rule_name}'
Alert message	Automatic Synchronous Replication pause has been triggered for *entity_info_msg*. '*reason*'.
Cause	Network connection to remote site or to Witness is lost probably because of a network partition or the remote site being unavailable.
Cause	One or more internal services are down or not working as expected.
Impact	Synchronous Replication for {entity_info_msg} will be paused.
Resolution	Ensure the remote site and Witness are reachable and resume Synchronous Replication.
Resolution	Contact Nutanix Support.
Severity	Info

Table 475: 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 476: 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 477: Subnet Creation Failure for the Recovery Plan [300405] [A300405]

Name	Subnet creation failure for Recovery Plan
Description	Subnet creation failed for the Recovery Plan.
Alert message	Subnet `network_cidr` creation in the virtual network `atlas_virtual_network_name` failed for the `operation_type` of the Recovery Plan `recovery_plan_name` executed through the recovery plan job `recovery_plan_job_name`.
Cause	Networks with conflicting CIDR might be present.
Impact	Some VMs will be recovered without vNICs.
Resolution	Resolve the issue. If you cannot resolve the issue, contact Nutanix support.
Severity	Critical

Table 478: Invalid Network Settings for the Recovery Plan [300407] [A300407]

Name	Recovery Plan Periodic Check for Invalid Network Settings
Description	Network settings have Availability Zones that are not in accordance with Availability Zone order list of Protection Policy applied to VMs in Recovery Plan.
Alert message	Invalid network settings found for the Recovery Plan `recovery_plan_name`. `reason`.
Cause	One or more Availability Zones listed in Protection Policy do not match the Availability Zones specified in the Recovery Plan network settings.
Impact	Recovery Plan will not be synced to Availability Zones defined in Protection Policy.
Resolution	Update Availability Zones in the Recovery Plan network settings to be in accordance with Availability Zone order in the Protection Policy.

Severity	Warning

Table 479: 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 480: 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 481: Recovery plans have conflicting network mappings [300410] [A300410]

Name	Recovery Plan network mapping conflict check
Description	Checks network mapping conflicts in multiple Recovery Plans.
Alert message	Recovery plans #rp_1@ and #rp_2@ have different source networks mapped to same destination network `target_network`
Cause	Multiple Recovery Plans have different source networks mapped to same destination network.
Impact	VMs isolated on source network might be able to reach each other on the destination after recovery.
Resolution	Avoid mapping multiple networks to same network in one or more Recovery Plans.
KB Article	7175
Severity	Info

Table 482: 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 483: Test Failover on Recovery Plan has not been executed recently [300412] [A300412]

Name	Recovery Plan Last Test Failover run status check
Description	Checks when last Test Failover was run on the Recovery Plan.
Alert message	`time` days have elapsed since a Test Failover was run on Recovery Plan `recovery_plan_name`.
Cause	Test Failover on the Recovery Plan has not been executed.
Impact	There might be issues in the recovery of VMs, which may be difficult to identify if the Test Failover is not run periodically.
Resolution	Run a Test failover on the Recovery Plan periodically, to ensure Recovery Plan has no errors.
KB Article	5793
Severity	Info

Table 484: Multiple Recovery Plans associated with a Category [300413] [A300413]

Name	Same Category in multiple Recovery Plans check.
Description	Checks if same category belongs to multiple Recovery Plans.
Alert message	A specific category should be associated with only one Recovery Plan. Recovery Plans : `rp_list` are associated with the same `entity_abbrv` category `category`.
Cause	Multiple Recovery Plans have Same Category
Impact	Execution of the following Recovery Plans will be affected and only one of those Recovery Plan will be able to recover VMs from specified category.
Resolution	Update the Recovery Plans such that the category is unique between them.
KB Article	5669
Severity	Warning

Table 485: Number of entities in Recovery Plan exceeds the threshold [300414] [A300414]

Name	Recovery Plan Entity Limit Check
Description	Checks if the entity count exceeds the threshold in Recovery Plan.
Alert message	Maximum number of entities in a recovery plan should not exceed `max_entity_count`. Recovery Plan `recovery_plan_name` have `entity_count` entities.
Cause	Number of entities in the Recovery Plan exceeds the limit.
Impact	Entity Recovery prone to failure.
Resolution	Reduce the number of entities in the Recovery Plan.
KB Article	5717
Severity	Warning

Table 486: VMs are part of multiple stages in Recovery Plan [300415] [A300415]

Name	VM in Multiple Stages Check
Description	Checks whether a VM is part of more than one stage in Recovery Plan.
Alert message	A VM should be part of a single stage in a Recovery Plan. `duplicate_entities`
Cause	VMs are part of multiple stages in the Recovery Plan.
Impact	VM specified in the first stage will be considered for recovery, later stages will be skipped.
Resolution	Update the Recovery Plan to ensure that the VMs are not part of multiple stages.
KB Article	5715
Severity	Info

Table 487: The same floating IP is associated with multiple VMs belonging to different Recovery Plans. [300416] [A300416]

Name	Conflicting Floating IP Check
Description	Checks if VMs which are part of different Recovery Plan have same Floating IPs.
Alert message	Same Floating IPs should not be assigned to multiple VMs and should not be part of multiple Recovery Plans. `alert_msg`
Cause	VMs belonging to different Recovery Plans are assigned the same Floating IP.
Impact	Floating IP assignment post VM recovery may fail.
Resolution	Update the Recovery Plans to ensure that a Floating IP address is to be assigned to only one VM.
KB Article	5797
Severity	Warning

Table 488: Some of the entities in the Recovery Plan are unprotected [300417] [A300417]

Name	Recovery Plan Entities Protection Policy Check
Description	Checks if entities configured in the Recovery Plan are protected.
Alert message	`entities` in the Recovery Plan `recovery_plan_name` are not protected.
Cause	Entities in the Recovery Plan are not protected.
Impact	Entities present in the Recovery Plan can't be recovered.
Resolution	Update associated entities in Recovery Plan to have Protection Policy.
KB Article	5795
Severity	Warning

Table 489: Entities not cleaned up following the Test Failover for Recovery Plan [300418] [A300418]

Name	Recovery Plan Checks for Cleanup of Entities after Test Failover
Description	Checks that entities recovered using Recovery Plan after Test Failover have been cleaned up.
Alert message	Entities recovered following the Test Failover of Recovery Plan `recovery_plan_name` have not been cleaned up. Manual clean up of the entities is required.
Cause	Entities recovered following the Test Failover of the Recovery Plan have not been cleaned up.
Impact	Subsequent recovery of the entities from Test Failover operation might get affected.
Resolution	Perform cleanup of the test entities for the Recovery Plan.
KB Article	5794
Severity	Warning

Table 490: Subnets Deletion Failure for the Recovery Plan [300419] [A300419]

Name	Dynamic Subnets deletion failure for Recovery Plan
Description	Deletion of dynamic subnets failed.
Alert message	Deletion of the subnets `networks_list` failed for the `operation_type` of the Recovery Plan `recovery_plan_name` executed through the `recovery_plan_job_name` due to `reason`.
Cause	Some of the VMs might be using the subnets.
Impact	Unused subnets will be billable.
Resolution	Resolve the issue. If you cannot resolve the issue, contact Nutanix Support.
Severity	Critical

Table 491: 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 492: 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 493: 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 494: Recovery Plan has multiple Availability Zone Orders [300424] [A300424]

Name	Recovery Plan Multiple Availability Zone Order Check
Description	Checks if the Recovery Plan has any entity protected to multiple Availability Zones.
Alert message	Recovery Plan `recovery_plan_name` has entities protected to multiple Availability Zones.
Cause	Recovery Plan contains entities protected to multiple Availability Zones.

Impact	The Recovery Plan update will not be allowed.
Resolution	Update associated entities in Recovery Plan to have single Availability Zone Order.
KB Article	6591
Severity	Warning

Table 495: Recovery Plan contains VMs with Unsupported CHDR VM configuration. [300425] [A300425]

Name	Recovery Plan CHDR VM Config Check
Description	Checks if Recovery Plan contains VMs with Unsupported CHDR VM configuration.
Alert message	Recovery Plan `recovery_plan_name` contains VMs `vm_list` with Unsupported CHDR VM configuration.
Cause	Recovery Plan contains VMs with unsupported configuration for Cross hypervisor disaster recovery.
Impact	The VMs will not be recovered.
Resolution	Install NGT on all the VMs in the Recovery Plan, ensure that you do not have UEFI boot, and have no SATA/PCI or delta disk attached.
KB Article	6835
Severity	Warning

Table 496: Incompatible Recovery Availability Zones for some VMs in the Recovery Plan [300426] [A300426]

Name	Recovery Plan Incompatible Recovery Availability Zone Check
Description	Checks if Recovery Plan contains VMware VMs and snapshots for these VMs are replicated to a recovery Availability Zone that doesn't support recovery of VMware VMs.
Alert message	Incompatible Recovery Availability Zones for Recovery Plan `recovery_plan_name`. Recovery of VMs `incompatibles_vms` on Recovery Availability Zones `incompatbile_target_availability_zone_names` will fail.
Cause	Target Availability Zone is running Prism Central version less than AOS 5.11 version and hence it does not support the recovery of VMware VMs.
Impact	The VM recovery will fail.
Resolution	Upgrade the Target Availability Zone to 5.11 or later version, or Remove the entity mentioned in the description of the alert from the Recovery plan.
KB Article	6872
Severity	Warning

Table 497: Floating IPs Deallocation Failed after Failback from Xi. [300427] [A300427]

Name	Unassociated Floating IPs Deallocation Failure after Failback from Xi.

Description	Failed to deallocate unassociated floating IPs after failback from Xi that were earlier dynamically created.
Alert message	Failed to deallocate unassociated floating IPs `failed_floating_ips` after failback from Xi that were earlier dynamically created, for `operation_type` on the Recovery Plan `recovery_plan_name`.
Cause	The floating IPs might be in use.
Impact	Floating IPs will continue to remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.
Severity	Warning

Table 498: Recovery Plan contains VMs with special configuration and no empty CDROM [300428] [A300428]

Name	Recovery Plan NGT Config Empty CD ROM Check
Description	Checks if VMs part of the recovery plan have an empty CDROM available to mount nutanix guest tools, so NGT reconfiguration tasks can be completed on restore.
Alert message	Recovery Plan `recovery_plan_name` contains VMs `vm_list` with special configuration that require NGT reconfiguration on restore. Nutanix Guest tools cannot be mounted as there is no empty CD ROM slot on the VM.
Cause	VMs on Source Availability Zone do not have a free CDROM slot.
Impact	VM will be recovered however post recovery VM reconfiguration tasks, including NGT reconfiguration, static IP reconfiguration will fail.
Resolution	Add an empty CDROM to the VM.
KB Article	7054
Severity	Warning

Table 499: Recovery Plan Execution Exceeded the Time Limit. [300429] [A300429]

Name	Recovery Plan Execution Exceeded the Maximum Expected Time Limit.
Description	Recovery Plan execution exceeded the time limit.
Alert message	Execution of Recovery Plan '`recovery_plan_name`' exceeded the estimated completion time by `excess_time_taken`.
Cause	Restore or Power On operation may have taken longer time.
Impact	Recovery Plan execution is prone to failure.
Resolution	Contact Nutanix customer service.
Severity	Warning

Table 500: Entities in category(s) present in multiple fault domains. [300432] [A300432]

Name	Category Protected entities Multiple Fault Domain Check

Description	Checks if entities protected by category(s) in Recovery Plan are present in multiple fault domains.
Alert message	`alert_message`
Cause	Protection Policy is applied on category(s) with Live entities on Local AZ and Remote AZ.
Cause	Entities from category(s) are getting replicated between clusters registered with Local Availability Zone and are configured for failover operations for different Availability Zones in Recovery Plan.
Impact	Selection of entities for recovery will be affected.
Resolution	Ensure category(s) present in Recovery Plan contain entities from only single Availability Zone.
Resolution	Ensure category(s) present in Recovery Plan contain entities from only one fault domain. Refer KB 9472 for further guidance. Contact Nutanix support if the issue persists.
KB Article	9472
Severity	Warning

License

Table 501: PC License Expiry [14008] [A14008]

Name	PCWideLicenseExpiry
Description	Check if licenses are about to expire.
Alert message	Detailed license expiry info: `license_expiry_msg`. `license_warning_msg`
Cause	The license is set to expire or has expired.
Impact	The Prism Central will be out of licensing compliance.
Resolution	Contact Nutanix to extend the expiration of the license, or apply a new license.
KB Article	2469
Severity	Warning

Table 502: Storage License Overshoot [14009] [A14009]

Name	StorageLicenseOvershoot
Description	Check if Storage licenses are in capacity compliance.
Alert message	Storage usage is out of compliance. Detailed info: `license_insufficiency_list`
Cause	Total storage usage exceeds the license applied.
Impact	The Prism Central is out of licensing compliance.
Resolution	Please upgrade the license with additional capacity. If you cannot resolve the issue, please contact Nutanix support.
KB Article	2469

Severity	Warning

Guest VM

Table 503: 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 504: 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 505: 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 506: 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 507: 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 508: 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 509: 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 510: 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 511: 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 512: 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 513: 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 514: 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 515: 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 516: 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 517: 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 518: 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 519: 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 520: 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 521: 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 522: 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 523: 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 524: 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 525: 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 526: 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 527: 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 528: 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 529: 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 530: 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 531: VM stored on RF1 container failed to start after host/CVM maintenance. [200405] [A200405]

Name	VM stored on RF1 container failed to start after host/CVM maintenance
Description	VM stored on RF1 container failed to start after host/CVM maintenance.
Alert message	VM *vm_name* (UUID: *vm_uuid*) failed to start with the following error: *message*
Cause	Insufficient resources available to start the VM.
Impact	VM which was running before the host/CVM maintenance will stay powered off until manual intervention.
Resolution	Free up resources and manually start the VM. Refer to KB-10804.
KB Article	10804
Severity	Critical

Hardware

Table 532: 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 533: 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 534: Disk Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Drive is marked offline.
Alert message	Disk mounted at *mount_path* on node *ip_address* with ID *service_vm_id* has been marked offline.
Cause	A drive on the node was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Investigate disk health. Please refer to KB 8453
KB Article	8453
Severity	Critical

Table 535: 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 536: 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 cluster and check the disk health. Contact the Hardware Vendor to replace the bad disk.
KB Article	8031
Severity	Critical

Table 537: Physical Disk Drive Might Have Failed. [130042] [A130042]

Name	Physical Disk Maybe Bad Status
Description	Physical drive might be failing.
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` might have failed. Node serial is `node_serial_number` and Node position `node_position`
Cause	The drive might be failing.
Impact	Cluster storage capacity is reduced.
Resolution	Inspect drive health. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 538: 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 539: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	Failed to evacuate VMs off of node `node_uuid`: `error`
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.
KB Article	8471
Severity	Critical

Node

Table 540: 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 541: 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 542: 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 543: 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 544: 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 545: 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 546: 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 547: 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 548: Host Network Uplink Configuration Failed. [130160] [A130160]

Name	Cluster Vswitch 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 549: 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 550: 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 551: 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 552: 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 553: Home partition is not expanded to 50GB. [200330] [A200330]

Name	Prism Central home partition expansion check
Description	Checks if home partition is expanded to 50GB after upgrade.
Alert message	Home partition is not expanded to 50GB on Prism Central `pc_ip` after upgrade.
Cause	During an upgrade, home partition extension has failed.
Impact	Prism Central performance may be degraded.
Resolution	Refer to KB 9525 for further details.
KB Article	9525
Severity	Warning

Table 554: 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 555: AHV Secondary IP Ping Check from PC [3071]

Name	AHV Secondary IP Ping Check from PC
Description	Checks if PC can ping secondary IP addresses of all the hosts of all the registered clusters.
Cause	The AHV secondary IP is down or there is a network connectivity issue.
Impact	If the AHV's secondary IP is not pingable from the PC, then advanced networking may encounter issues, if enabled and configured to use the corresponding virtual switch.
Resolution	Ensure that the AHV is running and that physical networking, VLANs, and virtual switches are configured correctly.
KB Article	12446

Table 556: 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 557: 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 558: 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 559: 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 560: 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 561: 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 562: 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 563: 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 564: Atlas was unreachable to apply Flow rule [200610] [A200610]

Name	Flow rule failed to reach Atlas
Description	Programming a Flow rule failed to reach Atlas.
Alert message	Flow policy operation failed for *`rule_name`*. *`reason`*
Cause	Atlas service is unreachable in PC or the Microseg service in PC is unreachable
Impact	VMs will not be protected by the Flow rule.
Resolution	Check Atlas service in PC, Microseg service in PC, Microseg to Atlas connection.
Severity	Warning

Table 565: 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 566: 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 567: 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 568: 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 569: 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 570: 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 571: 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 572: 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 573: 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 574: 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 575: 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 576: 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 577: 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 578: 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 579: 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 580: 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 581: 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 582: 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 583: 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 584: 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 585: 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 586: 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 587: 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 588: 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 589: 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 590: 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 591: 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 592: 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 593: 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 594: 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 595: 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 596: 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 597: 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 598: 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. `message`
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 599: 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 600: 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 601: 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 602: 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 603: 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 604: 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 605: 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 606: 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 607: 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 608: 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 609: 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 610: 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 611: 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 612: 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 613: 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 614: 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 615: 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 616: 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 617: 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 618: 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 619: 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 620: 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 621: 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 622: 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 623: 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 624: /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 625: 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 626: 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 627: 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`. `reason`
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 628: 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 629: 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 630: 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 631: 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 632: 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. `message`
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 633: 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 634: 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 635: 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 636: 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 637: 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 638: 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 639: 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 640: 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 641: 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 642: 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 643: 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 644: 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 645: 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 646: 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 647: 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 648: 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 649: 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 650: 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 651: 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 652: 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 653: 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 654: 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 655: 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 656: 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 657: 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 658: 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 659: 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 660: 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

Table 661: File Server Replication Job is stuck/idle [160143] [A160143]

Name	File Server Replication Job is stuck/idle
Description	File server replication job is stuck/idle - see the message for further details.
Alert message	`message`
Cause	See the message for further details.
Impact	See the message for further details.
Resolution	See the message for further details.
Severity	Warning

Other

Table 662: Detected recovery points on prism central managing clusters with high density nodes [110452] [A110452]

Name	High density nodes recovery points check

Description	Checks if recovery points are detected on prism central managing clusters with high density nodes
Alert message	Recovery points are detected on prism central managing clusters with high density nodes
Cause	Detected recovery points on prism central managing clusters with high density nodes
Impact	Recovery points are not currently supported on clusters with high density nodes. They may impact cluster operations
Resolution	Delete recovery points on prism central hosted by clusters with high density nodes
KB Article	7003
Severity	Info

Table 663: Detected protection policies on Prism Central managing clusters with high density nodes [110453] [A110453]

Name	High density nodes protection rules check
Description	Checks if protection policies are detected on Prism Central managing clusters with high density nodes
Alert message	Protection policies are detected on Prism Central managing clusters with high density nodes
Cause	Protection policies are detected on a Prism Central managing clusters with high density nodes
Impact	Data protection is currently not supported for clusters with high density nodes, this may impact cluster operations
Resolution	Delete protection policies configured on Prism Central managing clusters with high density nodes
KB Article	7003
Severity	Info

Table 664: 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 665: 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 666: 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 667: Memory configuration inconsistent. [200306] [A200306]

Name	Prism Central VM same memory level check
Description	Check all Prism Central VMs have the same memory level.
Alert message	The Prism Central VMs are not configured to have the same amount of memory.
Cause	Memory configuration among Prism Central VMs is inconsistent.
Impact	The Prism Central VM will not perform at the level necessary to manage the cluster.
Resolution	Fix Prism Central VM memory configuration by providing same amount of memory for all the Prism Central VMs
KB Article	5377
Severity	Warning
Severity	Info

Table 668: Prism Central VM type or annotation not set. [200307] [A200307]

Name	PCVM type and annotation check
Description	Check the type and annotation of all PCVMs are set.
Alert message	The type or annotation of Prism Central VMs is not set

Cause	Prism Central VMs do not have annotation set.
Impact	Guardrails to avoid potentially disruptive user actions on a Prism Central will not be applied.
Resolution	Set the annotation for the Prism Central VM. Refer to KB 6027.
KB Article	6027
Severity	Warning

Table 669: Unequal metadata partition size(s) across Prism Central VMs. [200308] [A200308]

Name	Prism Central VM same disk size check
Description	Checks if metadata partition size(s) are the same among Prism Central VMs.
Alert message	The metadata partitions across Prism Central VMs are not consistent due to mismatch in sizes: `disks_info`
Cause	Metadata partition size(s) of one or more Prism Central VMs is inconsistent with other Prism Central VMs.
Impact	Manageability of the cluster may be lost. Cluster may be significantly degraded and ultimately fail.
Resolution	Contact Nutanix support to resize the Prism Central VM disks.
KB Article	6028
Severity	Warning

Table 670: High time difference between Prism Central and registered Prism Elements. [200309] [A200309]

Name	PC-PE time drift check
Description	Checks ntp sync between the PC and the registered PEs.
Alert message	Time differs between Prism Central and registered Prism Elements.
Cause	Time is not synchronized between Prism Central and registered Prism Elements.
Impact	Stale statistics or configuration may appear in Prism Central.
Impact	Statistics may not be available in Prism Central.
Impact	Components that make use of the statistics or configuration in Prism Central may not function as expected.
Resolution	Ensure that NTP is configured properly and that the cluster has connectivity with the NTP server. Refer to KB 6121.
KB Article	6121
Severity	Warning

Table 671: PE-PC incompatible AOS versions. [200312] [A200312]

Name	PC-PE incompatible versions
Description	Checks if the Prism Element and Prism Central AOS versions are compatible.
Alert message	The AOS version of the registered cluster is incompatible with the Prism Central version.
Cause	The AOS version of the registered cluster is incompatible with the Prism Central version.
Impact	Prism Central features will not work as expected for this cluster.
Resolution	Upgrade the AOS version of the registered cluster to a version compatible with this Prism Central.
KB Article	6323
Severity	Warning

Table 672: Prism Central VM active upgrade check [200313]

Name	Prism Central VM active upgrade check
Description	Checks if there are any active upgrades
Cause	An upgrade is in process for this entity
Impact	Maintenance activities will have to wait until upgrade is finished.
Resolution	Wait until the upgrade is completed before doing any other actions.Contact Nutanix support for assistance.
KB Article	6644

Table 673: Prism Central Upgrades are disabled on {cvm_ip} [200314] [A200314]

Name	Prism Central VM upgrades are disabled check
Description	Checks if upgrades are disabled on Prism Central VM
Alert message	Prism Central Upgrades are disabled on `cvm_ip`
Cause	Prism Central Upgrade has been paused manually.
Impact	Upgrades from Prism won't proceed.
Resolution	Upgrades have been disabled and need to be re-enabled.Contact Nutanix support for assistance.
KB Article	6644
Severity	Info

Table 674: 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 675: 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 676: 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 677: Xi Payment Missed [200704] [A200704]

Name	Xi Invoice 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 678: Recovery Plan has conflicting IP mappings with other Recovery Plan(s) [300430] [A300430]

Name	Recovery Plan IP mappings check
Description	Checks for conflicting IP mappings across multiple Recovery Plans for the paired AZs.
Alert message	Recovery Plan `rp` has conflicting IP mappings with other Recovery Plans. `conflicting_ips` `conflicting_vms`
Cause	IPs mapped in Recovery Plans are conflicting. Same IP from a subnet is specified for multiple VMs.
Cause	IPs mapped in Recovery Plans are conflicting. Different IPs are specified for the same VM in different Recovery Plans.
Impact	IP mapping will fail if a VM has been recovered with the same IP from a different Recovery Plan.
Impact	VM will come up without the specified IP after recovery.
Resolution	Ensure that the same IP from a subnet is not specified for multiple VMs.
Resolution	Ensure that different IPs are not specified for the same VM in different Recovery Plans.
KB Article	7866
Severity	Warning

Table 679: Recovery Plan has errors in Consistency Group(s) configuration. [300434] [A300434]

Name	Recovery Plan consistency group configuration check
Description	Checks if Recovery Plans have errors in the configuration of consistency groups.
Alert message	Recovery Plan `recovery_plan_name` has errors in configuration of following Consistency Group(s): `partial_cg_names`.
Cause	Recovery Plan {recovery_plan_name} has the following partial Consistency Groups: {partial_cg_names}.
Impact	Partial Consistency Group might be recovered on failover.
Impact	Consistency Group will not be reformed post failover of entities in the Recovery Plan.
Impact	Recovered VM(s), Volume Group(s) might be in inconsistent state.
Resolution	Ensure that all VM(s) and Volume Group(s) belonging to a consistency group are present in Recovery Plan
KB Article	10532
Severity	Warning

Table 680: Recovery Plan has networks which are modified [300435] [A300435]

Name	Recovery Plan modified network check
Description	Checks if network names specified in the Recovery Plan are modified.
Alert message	Recovery Plan `recovery_plan_name` has following networks which are modified: `missing_networks`.
Cause	Network name(s) specified in the Recovery Plan are modified.
Cause	Network(s) specified in the Recovery Plan are deleted.
Impact	VMs will be recovered without vNIC upon failover
Resolution	Ensure that Network name(s) specified in the Recovery Plan are correct.
Resolution	Ensure that Network(s) specified in the Recovery Plan exists.
KB Article	11044
Severity	Warning

Table 681: Recovery Plan has VM(s) with iSCSI attached Volume Group(s) that do not have a compatible NGT installed/enabled. [300436] [A300436]

Name	Recovery Plan check for NGT version installed in VM(s) with iSCSI attached Volume Group(s)
Description	Checks if VM(s) with iSCSI attached Volume Group(s) specified in Recovery Plan have a compatible NGT installed.
Alert message	Recovery Plan `recovery_plan_name` has VM(s) `vm_list` with iSCSI attached Volume Group(s) that do not have a compatible NGT installed/enabled.
Cause	VM(s) with iSCSI attached Volume Group(s) do not have a compatible NGT installed.
Impact	Volume Group(s) attachment post recovery of a VM(s) would fail.
Resolution	Ensure that the VM(s) with iSCSI attached Volume Group(s) have NGT version of 2.3 or higher.
KB Article	11649
Severity	Warning

Table 682: Recovery Plan has VM to Volume Group(s) attachment(s) configuration that are different from the live VM to Volume Groups(s) attachment(s). [300437] [A300437]

Name	Recovery Plan VM to Volume Group(s) attachment(s) configuration check.
Description	Checks if the VM to Volume Group(s) attachment(s) configuration which are specified in the Recovery Plan are the same as the live VM to Volume Group(s) attachment(s).
Alert message	Recovery Plan `recovery_plan_name` has VM to Volume Group(s) attachment(s) `attachment_list` configuration specified that are not the same as the live VM to Volume Group(s) attachment(s).
Cause	VM to Volume Group(s) attachment(s) which are specified in the Recovery Plan are not present in the live state.

Cause	VM to Volume Group(s) attachment(s) present in the live state are not specified in the Recovery Plan.
Cause	VM to Volume Group(s) attachment(s) authentication configuration present in the live state are not the same as the configuration specified in the Recovery Plan.
Impact	VM to Volume Group(s) attachment(s) might not be done post recovery.
Resolution	Ensure that VM to Volume Group(s) attachment(s) which are specified in the Recovery Plan, are also present in the live state.
Resolution	Ensure that all VM to Volume Group(s) attachment(s) present in the live state are also specified in the Recovery Plan.
Resolution	Ensure that VM to Volume Group(s) attachment(s) authentication configuration present in the live state matches with the configuration specified in the Recovery Plan.
KB Article	11629
Severity	Warning

Table 683: Multiple Data Services IPs are configured on recovery cluster but valid Data Services IP address mapping(s) are not configured in the Recovery Plan. [300438] [A300438]

Name	Recovery Plan Data Services IP address mapping(s) configuration check.
Description	Checks if Data Services IP address mapping(s) are configured properly in a Recovery Plan.
Alert message	Multiple Data Services IPs are configured on recovery cluster but valid Data Services IP address mapping(s) are not configured in the Recovery Plan `recovery_plan_name`.
Cause	Multiple Data Services IPs are configured on recovery cluster but valid Data Services IP address mapping(s) are not configured in the Recovery Plan.
Impact	iSCSI attachments of Volume Group(s) to VM(s) will not be done on recovery cluster.
Resolution	Manual iSCSI reconfiguration will have to be performed from Guest VM(s).
KB Article	11858
Severity	Warning

Table 684: 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 685: 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 686: 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 687: Dataservice IP is unreachable [400105] [A400105]

Name	Data service IP connectivity check
Description	Data service IP is not reachable
Alert message	Data service ip '`data_service_ip`' is unreachable from prism central `node_name`
Cause	Invalid or empty data service ip
Impact	Nutanix Calm containers will be inaccessible
Resolution	Please provide correct data service ip. Review KB 5199
KB Article	5199
Severity	Critical

Table 688: 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 689: 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 690: 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 691: 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 692: 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 693: Calm containers health check [400112]

Name	Calm containers health check
Description	Check for Calm Container's state
Cause	Internal services of each calm containers are down, Docker plugin is not working properly
Impact	Nutanix Calm services may be inaccessible or performs incorrectly.
Resolution	Please check internal services of each calm containers.
KB Article	7343

Table 694: 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 695: Threshold set for the quota has been crossed [400116] [A400116]

Name	Quota threshold crossed
Description	Quota threshold crossed
Alert message	Quota utilisation has crossed threshold `threshold` for `quota_attributes` at `quota_level`
Cause	Number of virtual machine(s) deployed under quota exceeds the threshold
Impact	Quota consumption can go beyond the threshold value
Resolution	Either delete virtual machines or increase threshold value set for the quota
Severity	Warning

Prism Central VM

Table 696: XPilot Retry Approval Pending [111081] [A111081]

Name	XPilot Retry Approval Pending
Description	XPilot Playbook is waiting for Approval to retry
Alert message	XPilot Playbook `action_rule_name` is awaiting approval to begin retry no. `retry_count`
Cause	XPilot Playbook is waiting for approval to begin retry
Impact	Playbook will fail after wait period of 7 days
Resolution	Resume playbook to give approval. Refer to KB 9961 for further details
KB Article	9961
Severity	Info

Storage

Table 697: 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 698: 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 699: 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 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: 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 702: Disk Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Drive is marked offline.
Alert message	Disk mounted at *mount_path* on node *ip_address* with ID *service_vm_id* has been marked offline.
Cause	A drive on the node was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Investigate disk health. Please refer to KB 8453
KB Article	8453
Severity	Critical

Table 703: 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 704: 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 cluster and check the disk health. Contact the Hardware Vendor to replace the bad disk.
KB Article	8031
Severity	Critical

Table 705: Physical Disk Drive Might Have Failed. [130042] [A130042]

Name	Physical Disk Maybe Bad Status
Description	Physical drive might be failing.
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` might have failed. Node serial is `node_serial_number` and Node position `node_position`
Cause	The drive might be failing.
Impact	Cluster storage capacity is reduced.
Resolution	Inspect drive health. Please refer to KB 4158.
KB Article	4158
Severity	Critical

Table 706: 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 707: 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 708: 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 709: 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 710: 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 711: 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 712: 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 713: 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 714: 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 715: 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 716: 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 717: 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 718: 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 719: 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 720: 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 721: 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 722: 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 723: External datasource storage migration suspended due to unrecoverable errors [130347] [A130347]

Name	External datasource storage migration suspended due to unrecoverable errors
Description	Storage migration from an external datasource within an external repository suspended due to unrecoverable errors.
Alert message	Storage migration from the external datasource `datasource_uuid` with path `relative_uri` in the external repository `external_repository_name` suspended. Suspension reason: `message`.
Cause	Insufficient free storage space.
Cause	Reads to the external datasource failed with I/O error from the NFS server.
Cause	Network connectivity issues between the CVM subnet and the external repository. This can be either intermittent or permanent.
Cause	The datasource within the NFSv3 external repository does not have the correct access permission.
Cause	The datasource is not available in the NFS server.
Impact	Storage migration from a datasource within an external repository is suspended
Resolution	Make sure the destination container has enough free capacity.
Resolution	Make sure the NFS server is in a good state.
Resolution	Make sure that the CVM subnet has access to the NFS server.
Resolution	Make sure that the datasource is available for the CVM subnet and allows requests such as read and/or write.
Resolution	Make sure the datasource is available or recreate the VM disk with the correct datasource configuration.
KB Article	10864
Severity	Warning

Table 724: Storage Policies have become non-compliant [130366] [A130366]

Name	Storage Policies Non-Compliant
Description	Entities associated with storage policies have become non-compliant on PC
Alert message	Storage policies `storage_policy_name_list` have become non-compliant. Please refer the non-compliant reasons reported on VMs.
Cause	Storage policies {storage_policy_name_list} have become non-compliant. Please refer the non-compliant reasons reported on entities.
Impact	The storage properties applied on entities may not be in effect.
Resolution	Resolve the non-compliant reason(s) reported on entities.
Severity	Warning

Table 725: Prism Central VM System Root Partition Space Usage High [200316] [A200316]

Name	Prism Central VM System Root Partition Usage
Description	Checks if Prism Central system root partition usage is within threshold to ensure uninterrupted operations.
Alert message	Disk space usage for root partition `mount_path` on `entity-ip_address` has exceeded `threshold`%.
Cause	Increased Prism Central VM system root partition usage due to excessive logging or incomplete maintenance operation.
Impact	If Prism Central VM system root partition is highly utilised then certain maintenance operations, such as upgrades, may be impacted. If Prism Central VM system root partition is 100% utilized then services may stop and impact Prism Central cluster management functions.
Resolution	Reduce Prism Central VM system root partition usage by removing any known temporary or unneeded files. Refer to KB 5228 for further details.
Prism Central VM System Root Partition Usage Threshold	Critical Threshold: 95 %
KB Article	5228
Severity	Critical

Table 726: Prism Central VM home partition disk usage high [200317] [A200317]

Name	Prism Central VM home partition disk usage
Description	Checks if Prism Central home partition usage is within threshold to ensure uninterrupted operations.
Alert message	`alert_msg`
Cause	Increased Prism Central VM home partition usage due to excessive logging or incomplete maintenance operation.
Impact	If Prism Central VM home partition is highly utilised then certain maintenance operations, such as upgrades, may be be impacted. If Prism Central VM home partition is 100% utilised then services may stop and impact cluster storage availability.
Resolution	Reduce Prism Central VM home partition usage by removing any known temporary or unneeded files. Refer to KB 5228 for further details.
Prism Central VM home partition disk usage critical threshold percentage	Warning Threshold: 75 % Critical Threshold: 90 %
KB Article	5228
Severity	Critical
Severity	Warning

System Indicator (Prism Central)

Table 727: 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 728: 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

SYSTEM MANAGEMENT

You can configure various system settings for Prism Central.

- You can configure an HTTP proxy, which can be used to communicate with a Nutanix service center (see Configuring an HTTP Proxy on page 509
- You can specify one or more name servers (see Configuring Name Servers (Prism Central) on page 514).
- You can specify one or more NTP servers for setting the system clock (see Configuring NTP Servers (Prism Central) on page 515).
- You can configure SNMP (see Configuring SNMP (Prism Central) on page 517).
- You can configure a login banner page (see Configuring a Banner Page (Prism Central) on page 531).
- You can change the language setting for displayed text (see Changing the Language Settings (Prism Central) on page 532).

Configuring Prism Central When a Cluster Uses Proxy Servers

Note:

- Prism Central and its managed clusters are not supported in environments deploying Network Address Translation (NAT).
- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see Prism Central Proxy for Pulse Data on page 1005.

To communicate with the Nutanix service center directly, you need to configure the HTTP proxy. See Configuring an HTTP Proxy on page 509 for information on how to configure the HTTP proxy through Prism Central web console. After the HTTP proxy is configured, all the HTTP or HTTPS communication initiated by Prism Element or the Prism Central is routed through the proxy server.

The communication from Prism is routed through the proxy server until a whitelist target entry indicates otherwise. See Who Needs to Use the Whitelist Method on page 508 for information on who needs to whitelist the IP addresses. To overcome a communication failure between a Prism Element and the registered Prism Central, you need to manually whitelist the target entries from the Prism web console or nCLI. These options enable you to add Prism Central and its managed/registered clusters to a whitelist, where any HTTP proxy settings are ignored. This configuration allows network traffic between them, bypassing any proxy servers configured in the cluster. The whitelist also enables you to register new clusters with Prism Central successfully where clusters are using an HTTP proxy.

See *Step 3* of Configuring an HTTP Proxy on page 509 for information on how to configure the whitelist IP address through the Prism Central web console.

Alternatively, you can use the ncli http-proxy add-to-whitelist and ncli http-proxy delete-from-whitelist nCLI command options. See Whitelisting Prism Central and Its Managed Clusters (nCLI) on page 513 for information on how to configure the whitelist IP address through the nCLI.

- You can add or delete one whitelist entry at a time.
- Each whitelist entry cannot exceed 253 characters.
- A maximum of 1000 whitelist entries are supported.
- When deleting an entry from a whitelist, delete the target, not the target type.

- The commands do not support the IPv4 network mask network prefix or * (asterisk) prefix notation. You can specify a subnet and netmask, such as 10.0.0.0/255.0.0.0 or 192.168.1.0/255.255.255.0. This results in adding the specified subnet to the whitelist.
- When applying a whitelist, domain names like contoso.com are not processed the same as www.contoso.com and are treated as separate, distinct entities.
- Use fully qualified domain names to whitelist.

Who Needs to Use the Whitelist Method

SSL port 9440 needs to be open in both directions between the Prism Central VM and any registered clusters or clusters to be registered. For the complete list of required ports, see Port Reference.

If you are implementing a proxy server in your cluster environment with this port open as shown in this simple graphic, you do not need to allow Prism Central and its managed/registered clusters.

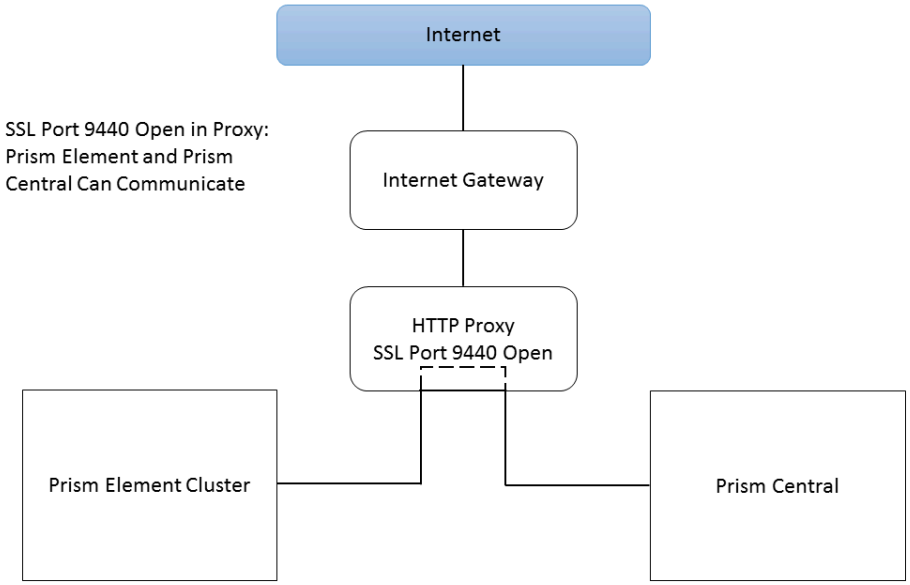

Figure 219: Prism Central, Proxy Server with SSL Port 9440 Open

If you are implementing a proxy server in your cluster environment with this port closed as shown in this simple graphic, you must allow direct communication between Prism Central and its managed/registered clusters as described in *Step 3* of Configuring an HTTP Proxy on page 509 or alternatively in Whitelisting Prism Central and Its Managed Clusters (nCLI) on page 513.

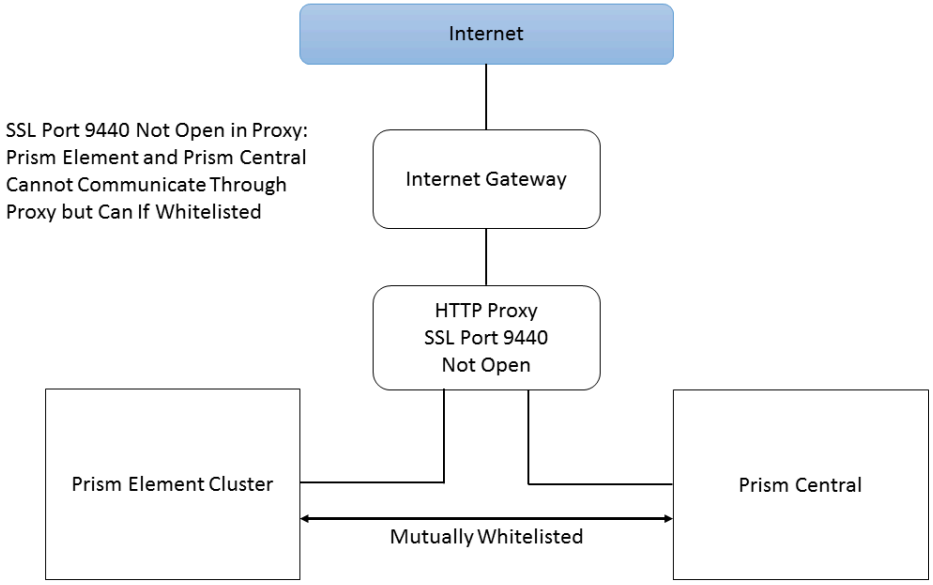

Figure 220: Prism Central, Proxy Server with SSL Port 9440 Closed and with Whitelisting

Configuring an HTTP Proxy

About this task

Note:

- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see Prism Central Proxy for Pulse Data on page 1005.

About this task

If Prism Central cannot send traffic to a Nutanix Service center directly, an HTTP proxy is required. To configure an HTTP proxy, do the following:

Procedure

1. Click the gear icon in the main menu and then select **HTTP Proxy** in the **Settings** page. The **HTTP Proxies** dialog box appears.

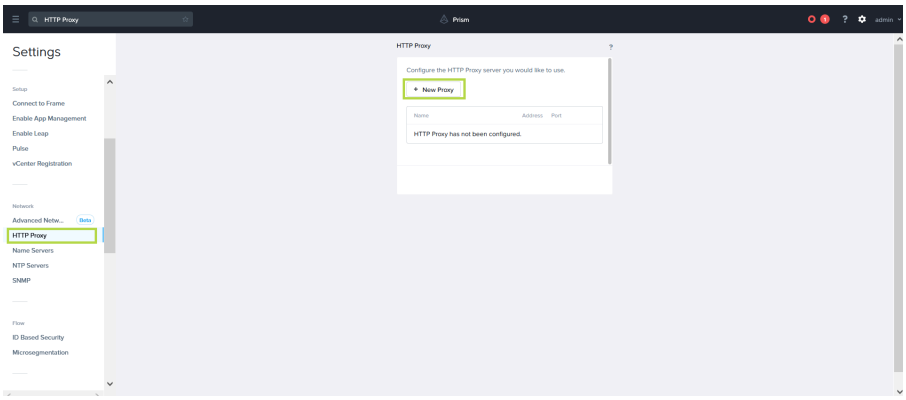

Figure 221: HTTP Proxy Window

2. To add an HTTP proxy, click the **New Proxy** button and do the following in the displayed fields:

   > **Note:** You can configure only one HTTP proxy at a time. If one exists currently, you must first delete it before creating a new one.

   a. **Name**: Enter a proxy server name.

   b. **Address**: Enter an IP address or host name for the proxy server.

   c. **Port**: Enter the port number to use.

   d. **Username**: Enter a user name.

   e. **Password**: Enter a password.

   f. **Protocols**: Select (check) the protocol to use, **HTTP**, **HTTPS**, or both.

      Select **HTTP** to route all HTTP requests through the proxy.

      Select **HTTPS** to route all HTTPs requests through the proxy.

      Select **HTTP** and **HTTPS** to route both HTTP and HTTPS requests through the proxy.

   g. When all the fields are correct, click the **Save** button (lower right).

      This saves the configuration and redisplays the dialog box with the new HTTP proxy entry appearing in the list.

      > **Note:** To return to the **HTTP Proxy** window without saving, click the **Cancel** button.

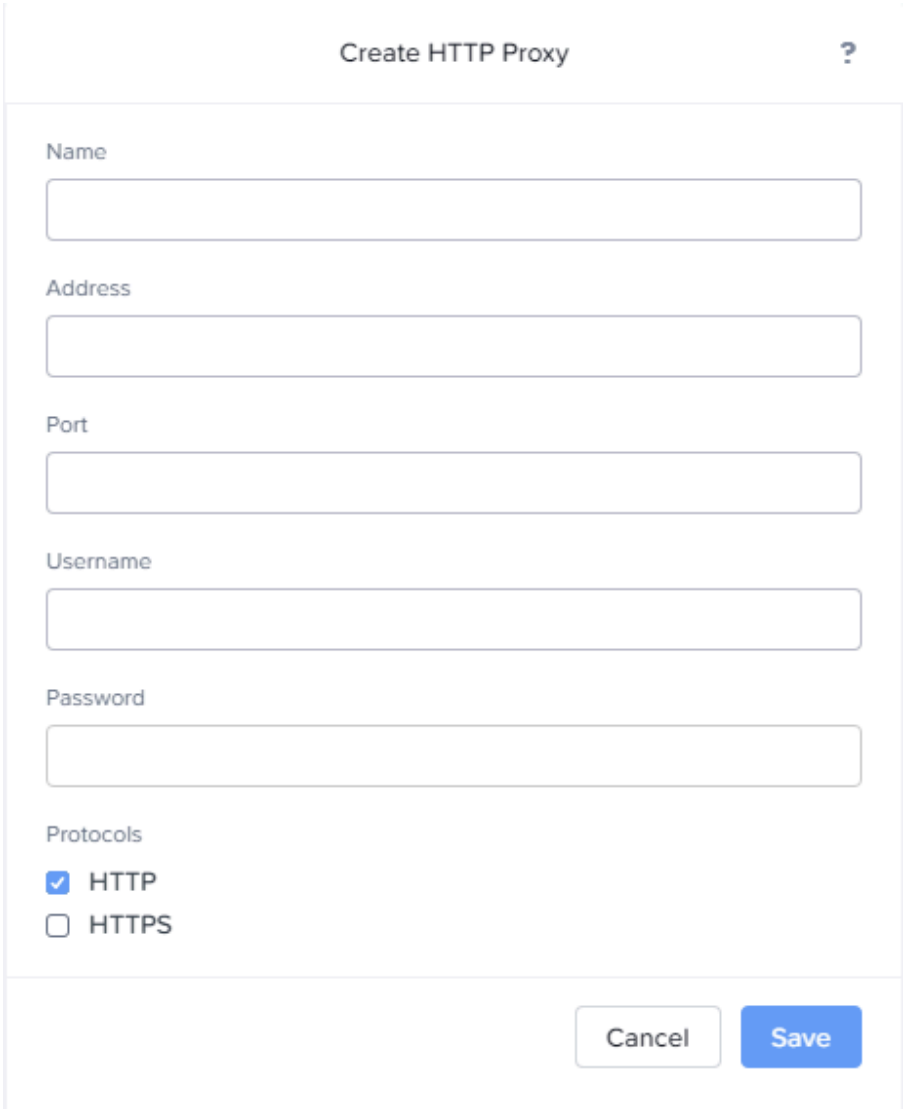

Figure 222: Create HTTP Proxy Window

3. To edit an HTTP proxy entry, click the pencil icon on the line for that entry, update one or more of displayed field entries as desired, and then click the **Save** button.

   The **Update HTTP Proxy** dialog box appears with the same fields as the **Create HTTP Proxy** dialog box plus the option (below the protocol check boxes) to add whitelist entries.

   To configure HTTP proxy whitelist entries, do the following:

   - To add a whitelist target, click the **+ Create** link. This step opens a line to enter a target address. Enter the target IP address or the network address and the subnet mask and then click the **Save** link in that field.

     Prism Central sends traffic to the whitelist IP addresses directly rather than through the HTTP or HTTPS proxy.
   - To edit a whitelist target, click the pencil icon for that target and update as needed.
   - To delete a whitelist target, click the X icon for that target.

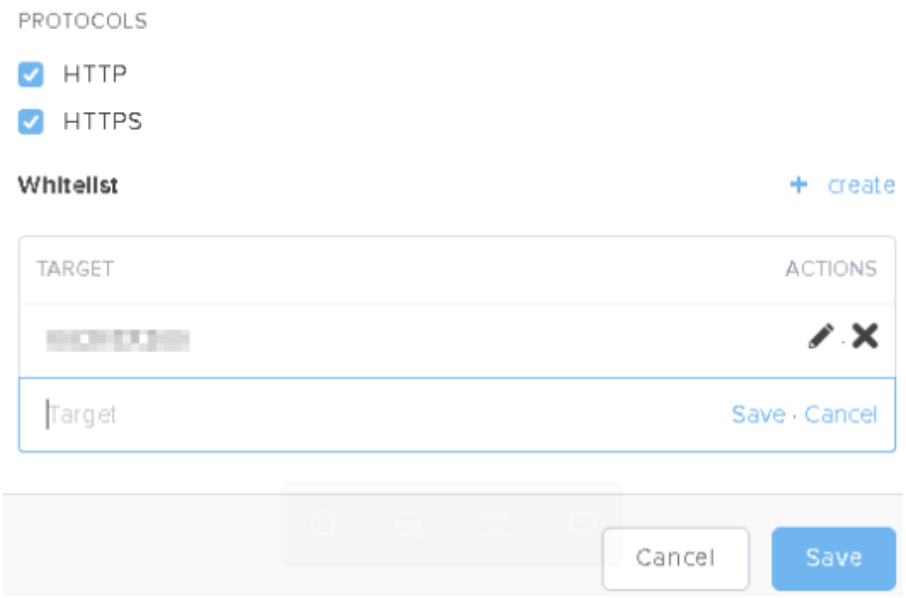

Figure 223: Whitelist Targets

> **Note:** Where proxy whitelist configuration is needed, it is recommended that you add the cluster virtual IP and all external CVM IPs of each registered cluster to the proxy whitelist for Prism Central. The virtual IP of Prism Central and the IP of each Prism Central VM (PCVM) should be added to the proxy whitelist on the registered Prism Element clusters. For easier management, you can add the subnets instead of adding the individual IPs.

4. To delete an HTTP proxy entry, click the X icon for that entry.

   A window prompt appears to verify the action; click the **OK** button. The entry is removed from the HTTP proxy list.

Whitelisting Prism Central and Its Managed Clusters (nCLI)

About this task

In this example, bypass a proxy server used by a managed Prism Element cluster and allow network traffic between Prism Central and the cluster. Previously, if you attempted to register a cluster that implemented a proxy server, the registration failed.

Procedure

1. Open a SSH session to any Controller VM in the cluster to be managed by Prism Central.
2. In this example, add the Prism Central VM IP address to the whitelist, then ensure the Prism Central VM IP address was added to the whitelist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address
 target=10.4.52.40
```

```
nutanix@cvm$ ncli http-proxy get-whitelist
Target Type : IPV4_ADDRESS
Target : 10.4.52.40
```

> **Note:** Repeat this step for additional Prism Central VM IP addresses and the virtual IP address in the case of scaleout Prism Central.

3. Open a SSH session to the Prism Central VM managing the cluster where you just modified the HTTP whitelist.
4. Add the cluster virtual IP address to the whitelist, then ensure the IP address was added to the whitelist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address
 target=10.4.52.10
```

```
nutanix@cvm$ ncli http-proxy get-whitelist
Target Type : IPV4_ADDRESS
Target : 10.4.52.10
```

> **Note:**
>
> - Repeat this step for individual CVM IP addresses.
> - For large clusters or for multiple clusters in a network it can be more efficient to add a network to the http-proxy whitelist. For example,
>
>   The following is a sample command to add a network to the http-proxy whitelist.
>
>   ```
>   nutanix@PCVM:~$ ncli http-proxy add-to-whitelist target-
>   type=ipv4_network_mask target=10.4.52.0/255.255.255.0
>   ```
>
>   The following is sample output of the get-whitelist command, which displays the whitelisted network added above.
>
>   ```
>   nutanix@CVM:~$ ncli http-proxy get-whitelist
>   ```
>
>   Target Type : IPV4_NETWORK_MASK Target : 10.4.52.0/255.255.255.0

In this case, Prism Central and its managed cluster can communicate, with network traffic bypassing any proxy servers configured in the cluster.

Configuring Name Servers (Prism Central)

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.

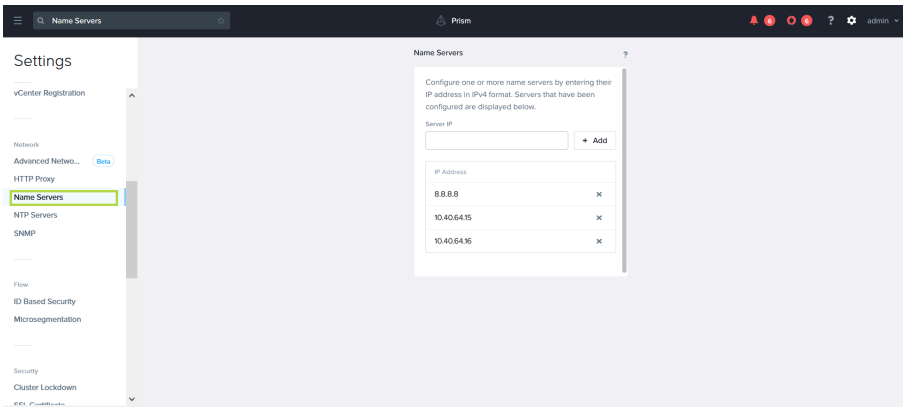

Figure 224: Name Servers Window

2. To add a name server, enter the server IP address in the **Server** field and then click the **Add** button to the right of that field.

   The server is added to the **IP Address** list (below the **Server** field).

   > **Note:** Changes in name server configuration may take up to 5 minutes to take effect. Functions that rely on DNS may not work properly during this time. You can configure a maximum of three name servers.

3. To delete a name server entry, click the X icon for that server in the **Host Name or IP Address** list.

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

4. Click the **Close** button to close the **Name Servers** window.

Configuring NTP Servers (Prism Central)

About this task

Network Time Protocol (NTP) is a protocol for clock synchronization between computers, and Prism Central must connect to an NTP server to synchronize the system clock. To add (or delete) an NTP server entry, do the following:

> **Note:** If Prism Central is running on Hyper-V, you must specific the IP address of the Active Directory Domain Controller server, not the hostname. Do not use DNS hostnames or external NTP servers.

Procedure

1. Click the gear icon in the main menu and then select **NTP Servers** in the **Settings** page.

   The **NTP Servers** dialog box appears.

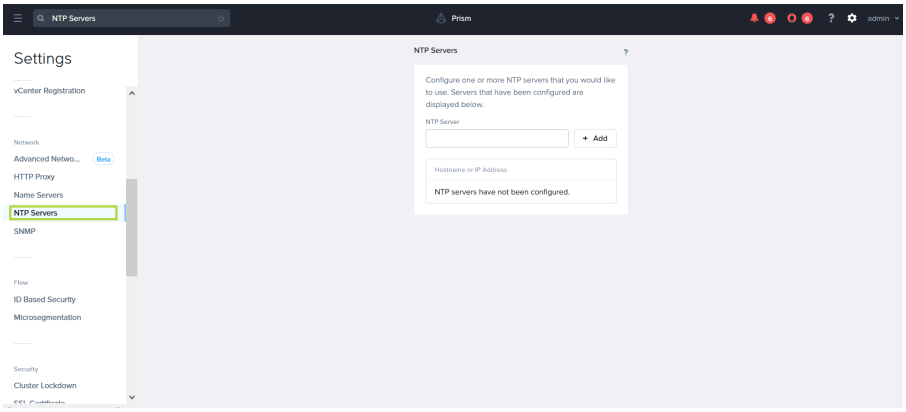

Figure 225: NTP Servers Window

2. To add an NTP server entry, enter the server IP address or fully qualified host name in the **NTP Server** field and then click the **Add** button to the right of that field.

   The name or address is added to the **Server** list (below the **NTP Server** field).

3. To delete an NTP server entry, click the delete icon for that server in the **Servers** list.

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

4. Click the **Close** button to close the **NTP Servers** window.

Configuring an SMTP Server (Prism Central)

About this task

Simple Mail Transport Protocol (SMTP) is an Internet standard protocol for electronic mail transmission across Internet Protocol (IP) networks, and Prism Central uses SMTP to send alert emails and to exchange emails with Nutanix customer support. To configure an SMTP server entry, do the following:

Procedure

1. Click the gear icon in the main menu and then select **SMTP Server** in the **Settings** page.

   The **SMTP Server Settings** dialog box appears.

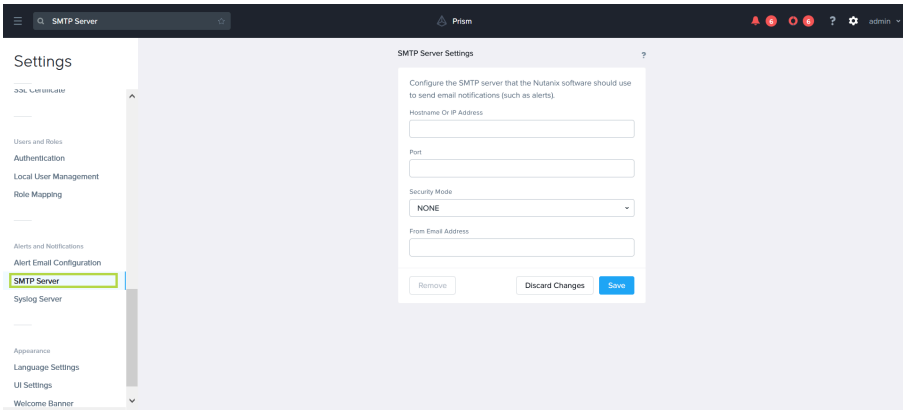

Figure 226: SMTP Server Settings Window

2. Do the following in the indicated fields:

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

   b. **Port**: Enter the port number to use.

      The standard SMTP ports are 25 (unencrypted), 587 (TLS), and 465 (SSL). For the complete list of required ports, see Port Reference.

   c. **Security Mode**: Enter the desired security mode from the pull-down list.

      The options are **NONE** (unencrypted), **STARTTLS** (use TLS encryption), and **SSL** (use SSL encryption).

   d. **User**: Enter a user name.

      The **User** and **Password** fields apply only when a secure option (**STARTTLS** or **SSL**) is selected. The user name might need to include the domain (`user@domain`) depending on the authentication process.

   e. **Password**: Enter the user password.

   a. **From Email Address** (optional): Enter an e-mail address that appears as the sender address.

      By default, alert and status information e-mails display "cluster@nutanix.com" as the sender address. You have the option to replace that address with a custom address by entering a sender address in this field.

3. When all the fields are correct, click the **Save** button.

Configuring SNMP (Prism Central)

About this task

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. To configure SNMP on Prism Central, do the following:

> **Note:** Prism Element (individual cluster) supports both the SNMP service (agent) and SNMP traps, but Prism Central only supports SNMP traps. Prism Central does not support the SNMP service (unlike Prism Element), so its SNMP capability is limited to just sending traps. For information on configuring SNMP for an individual cluster and for details about the Nutanix MIB, see Prism Element Web Console Guide.

Procedure

1. Click the gear icon in the main menu and then select **SNMP** in the **Settings** page.

   The **SNMP Configuration** dialog box appears.

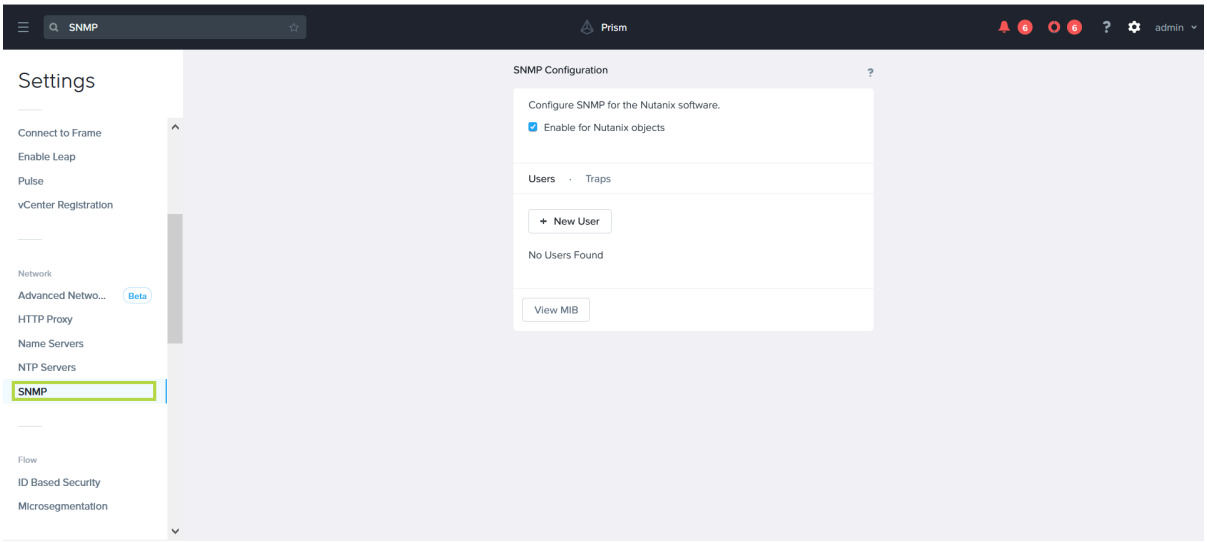

Figure 227: SNMP Configuration Window

2. To enable SNMP for Nutanix Objects, check the **Enable for Nutanix Objects** box.

   For more information about Nutanix Objects, see Enabling Objects on page 993.

3. To view the Nutanix MIB (NUTANIX-MIB.txt), click the **View MIB** link. To download NUTANIX-MIB.txt, right-click and select the appropriate download action for your browser.

4. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:

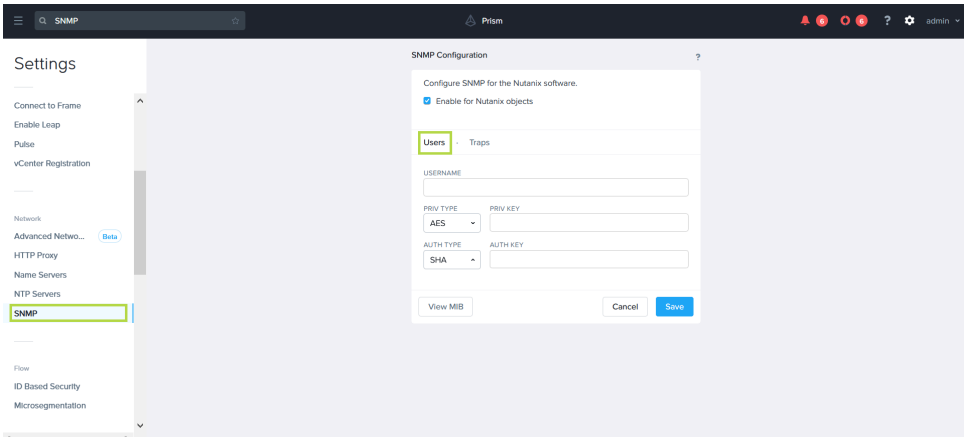

Figure 228: SNMP Configuration: Users Tab

a. **Username**: Enter a user name.

b. **Priv Type**: Select the privacy encryption type from the pull-down list.

   The only option is **AES** (Advanced Encryption Standard).

c. **Priv Key**: Enter a privacy key phrase (password) into this field.

   The key phrase is AES encrypted when the user is created.

d. **Auth Type**: Select the authentication hash function type from the pull-down list.

   The only option is **SHA** (Secure Hash Algorithm).

e. **Auth Key**: Enter an authentication key phrase (password) into this field.

   The key phrase is SHA-1 encrypted when the user is created.

f. When all the fields are correct, click the **Save** button (lower right).

   This saves the configuration and redisplays the dialog box with the new user entry appearing in the list.

5. To add an SNMP trap receiver, click the **Traps** tab and the **New Trap Receiver** button, and then do the following in the indicated fields:

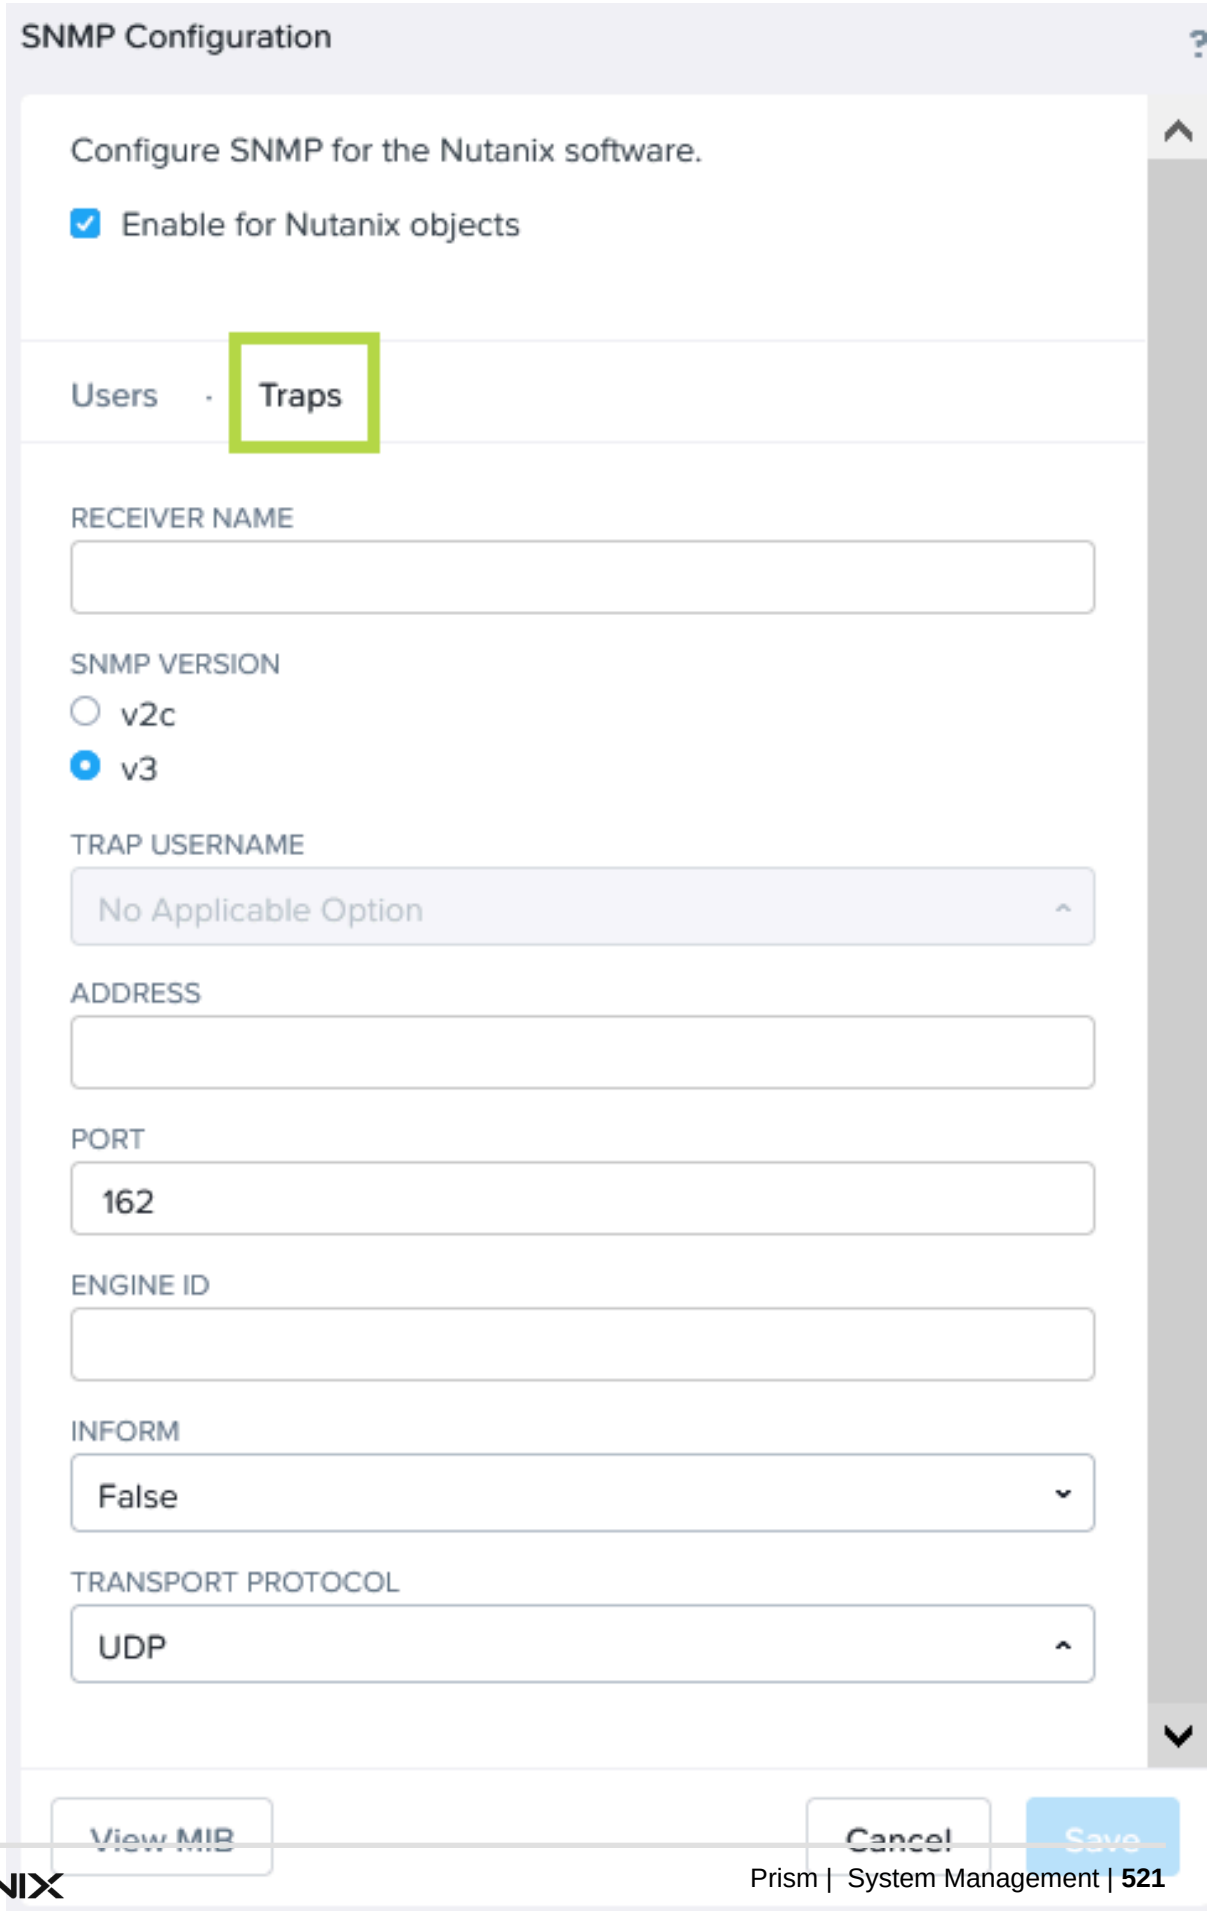

Figure 229: SNMP Configuration: Traps Tab

a. **Receiver Name**: Enter the receiver name.

b. **SNMP Version**: Select (click the radio button for) the SNMP version, either v3 or v2c. For SNMP v2c, Nutanix supports only SNMP TRAP and not SNMP GET.

c. This field is displayed based on your selection in the **SNMP Version**:

- **Trap Username**: This field is displayed if you select v3 in the **SNMP Version**. Select a user from the pull-down list.
- **Community**: This field is displayed if you select v2c in the **SNMP Version**. The default value for v2c trap community is public, or you can enter any other name of your choice.

All users added previously (see step 4) appear in the pull-down list. You cannot add a trap receiver entry until at least one user has been added.

d. **Address**: Enter the target address.

An SNMP target address specifies the destination and user that receives outgoing notifications, such as trap messages. SNMP target address names must be unique within the managed device.

e. **Port**: Enter the port number to use.

The SNMP trap receiver uses UDP port number 162. For the complete list of required ports, see Port Reference.

f. **Engine ID** (optional): Enter an engine identifier value, which must be a hexadecimal string between 5 and 32 characters long.

If you do not specify an engine ID, an engine ID is generated for you for use with the receiver. Every SNMP v3 agent has an engine ID that serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

g. **Inform**: Select **True** from the pull-down list to use inform requests as the SNMP notification method; select **False** to use traps as the SNMP notification method.

SNMP notifications can be sent as traps or inform requests. Traps are one-way transmissions; they do not require an acknowledgment from the receiver. Informs expect a response. If the sender never receives a response, the inform request can be sent again. Therefore, informs are more reliable than traps. However, informs consume more resources. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform may be retried several times. The retries increase traffic and add overhead on the network. Thus, traps and inform requests provide a trade-off between reliability and resources.

h. **Transport Protocol**: Select the protocol to use from the pull-down list.

The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.

i. When all the fields are correct, click the **Save** button (lower right).

This saves the configuration and redisplays the dialog box with the new trap entry appearing in the list.

j. To test all configured SNMP traps, click the **Traps** tab, and then click **Test All**.

The Nutanix cluster sends test alerts to all the SNMP trap receivers configured on the cluster.

6. To edit a user or trap receiver entry, click the appropriate tab (**Users** or **Traps**) and then click the pencil icon for that entry in the list.

An edit window appears for that user or trap receiver entry with the same fields as the add window. (Transport entries cannot be edited.) Enter the new information in the appropriate fields and then click the **Save** button.

7. To delete an SNMP entry, click the appropriate tab (**Users** or **Traps**) and then click the X icon for that entry in the list.

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

8. Click the **Close** button to close the **SNMP Configuration** window.

Configuring Syslog Monitoring

Using Prism Central, you can configure syslog monitoring to forward system logs (API Audit, Audit, Security Policy Hitlogs, and Flow Service Logs) of the registered clusters to an external syslog server. Prism Central enables you to configure multiple remote syslog servers. Additionally, you can configure separate log modules to be sent to each of the rsyslog servers.

About this task

> **Note:** The Prism Central method of syslog monitoring configuration propagates the configuration to the Prism Element clusters. If you do not want the configuration to be propagated to the clusters, you must use Nutanix command-line interface (nCLI) for syslog monitoring configuration.

To configure syslog monitoring through Prism Central, do the following.

Before you begin

- You must have the IP address of the syslog server that is deployed in your environment.
- For forwarding Flow logs, the Flow feature must be enabled.

Procedure

1. Click the gear icon in the main menu and then select **Syslog Server** in the **Settings** page.
2. Click **Add Syslog Server**.
3. In **Server Name**, enter a descriptive name for the server.
4. Enter the **IP Address** and **Port**.
5. Select the **Transport Protocol** (TCP or UDP).
6. Optionally, click the check box to enable RELP (Reliable Logging Protocol).
7. Click **Edit** against **Data Sources**.
8. Select one or more log modules from the following log types. See Syslog Modules on page 524 for details.

   - **API Audit**
   - **Audit**
   - **Security Policy Hitlogs** (policy hitlog files logs)
   - **Flow Service Logs** (Flow processes logs)

   > **Caution:** Users recommend that you only configure rsyslog modules listed in Syslog Modules on page 524. Configuration of other modules using ncli commands will be over-written with any modifications done on the Prism Central web console.

9. Select the **Severity Level** from the following available options. See Syslog Modules on page 524 for details.

   Click **Save** to complete.

Syslog Modules

Refer to the following table to understand the information that is sent to the syslog server based on the selected log modules and severity levels.

Table 729: Log Modules and Severity Levels

Module Name	Severity Level	Description
API Audit	0-7	API Audit logs send information about REST API endpoints that are called and who called them from both Prism Central and Prism Element. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent.
Audit	0-7	Audit logs send information about VM, Category, and Security Policy creation, update, and delete operations. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent. See the *Syslog Module: AUDIT Fields* table below for information on Syslog Audit fields.
Security Policy Hit Logs	All Levels	Security Policy Hit Logs send information about network flows that are acted on by a security policy. This information includes the source and destination IP address, protocol, and port. It also includes the action taken by the policy such as allow, monitor, or drop. Network statistics are also included in each message for bytes sent and received in the flow. Messages are always sent to the syslog server as level 6 - informational.
Flow Service Logs		Flow service logs contain detailed debug information used for troubleshooting the backend processes used by Flow. **Note:** Do not enable these unless instructed to do so by Nutanix support for troubleshooting purposes. Unlike the other data sources, the syslog level selected for Flow Service Logs has an impact on the types of data sent to the syslog server.
	0 - 5	Critical service failure messages as well as log rotation indications. Sent at the selected log level to syslog.
	6 - 7	Detailed log messages from the various components that make up the Flow product on both Prism Central and Prism Element.

The following table provides information on the syslog severity levels.

Table 730: Syslog Severity Levels

Value	Severity	Keyword	Description
0	Emergency	emerg	System is unusable.
1	Alert	alert	Should be corrected immediately.
2	Critical	crit	Critical conditions.
3	Error	err	Error Conditions.
4	Warning	warning	May indicate that an error will an action is not taken.
5	Notice	notice	Events that are unusual, but not error conditions.
6	Informational	info	Normal operational messages that require no action.
7	Debug	debug	Information useful to developers for debugging the application.

The following table provides information on the AUDIT fields of the Syslog Module.

Table 731: Syslog Module: AUDIT Fields

AUDIT Field	Description
entityType	Indicates the resource type of the entity that is associated with the audit record. Example. The entity type is `vm` for operations on virtual machines.
name	Indicates the name of the entity that is associated with the audit record. Example. Name of the virtual machine for audits or operations on virtual machines.
uuid	Indicates the unique identifier of the entity that is associated with the audit. uuid is also used to fetch the details of the entity instance.
alertUid	Indicates the name of the audit record. Example. `VmMigrateAudit` and `VmCreateAudit`.
creationTimestampUsecs	Indicates the time (in microseconds) from Epoch when the audit record was created.
defaultMsg	Provides a brief description of the current state of the operation that is captured by the audit.
opStartTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit begins.
opEndTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit is completed.

AUDIT Field	Description
operationType	Indicates the category of the audit record. Example. `Create`, `Delete`, and `Migrate`.
originatingClusterUuid	Indicates the cluster (if applicable) where the operation of is performed.
params	Additional key value pairs that capture information pertaining to the operation that is captured by the audit record.
userName	Indicates the name of the user who triggered the operation.
userUuid	Indicates the unique identifier of the user who triggered the operation.

The following table provides information on the different fields of the Security Policy Hit Logs.

Table 732: Security Policy Hit Logs Fields

Field	Description
<Timestamp>	Indicates the time of the hit log capture.
<Policy UUID>	Indicates the unique identifier of the policy that is associated with the hit log.
<Policy Name>	Indicates the name of the policy as configured by the user.
<Session Information> Create, Update, or Destroy	Indicates the session information. • Create - New session started. TCP handshake or UDP bidirectional packets. • Update - Sent every 30 seconds for active sessions. • Destroy - Session ended by TCP handshake or UDP timeout.
SRC	Indicates the source IP address.
DST	Indicates the destination IP address.
PROTO	Indicates the transport protocol.
ACTION	Indicates the action taken on traffic by the applied policy. • Allow - traffic allowed by policy • Drop - traffic dropped by policy • Monitor - traffic allowed by policy in monitor mode
<Packet Stats>	Indicates the packet statistics.

Modifying UI Settings (Prism Central)

About this task

The Prism Central login page includes background animation by default, and users are logged out automatically after being idle for 15 minutes. You can change one or both of these settings.

Procedure

1. Click the gear icon and select **UI Settings** from the **Settings** > **Appearance** menu (see Settings Menu (Prism Central) on page 16).
   The **UI Settings** window appears.

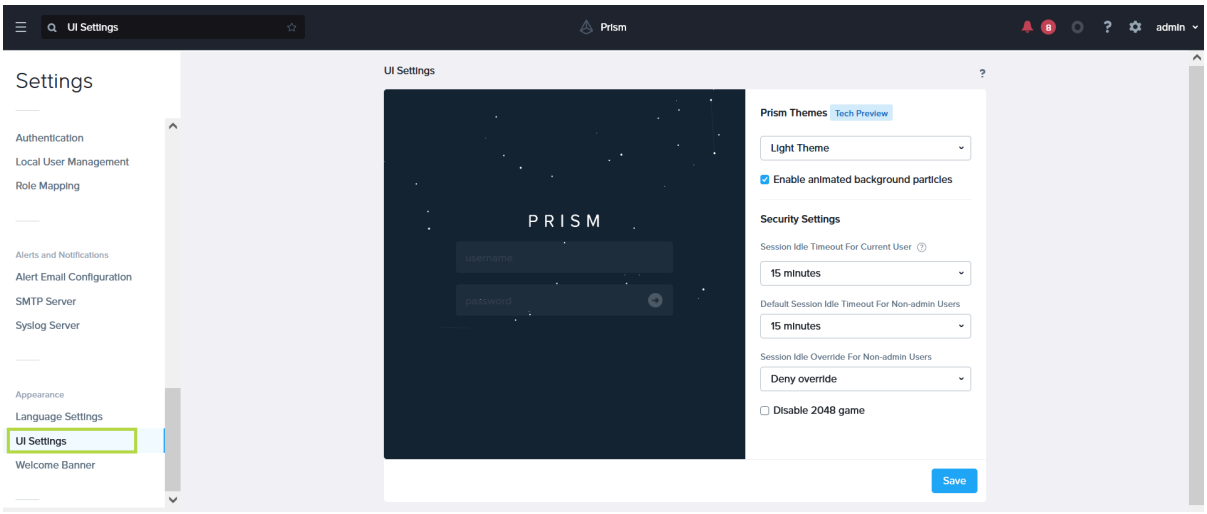

Figure 230: UI Settings Window

2. To set different variations of the Prism UI background themes, select any of the following options from the **Prism Themes** drop-down menu. The UI changes cosmetically to reflect the selection.

Important:

- The Prism themes feature is currently in technical preview. You may encounter visual anomalies in few settings or views where the prism themes are not applied. For example, Licensing Settings. Nutanix recommends that you register a case on the support portal to report any anomalies.

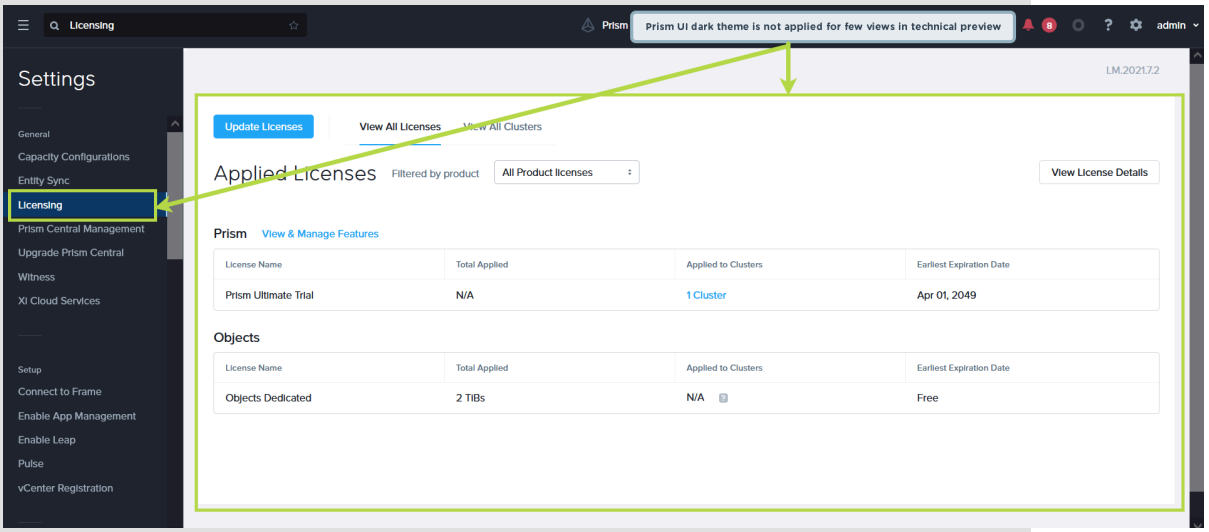

Figure 231: Example: Licensing View (Dark Theme Exception)

- 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.

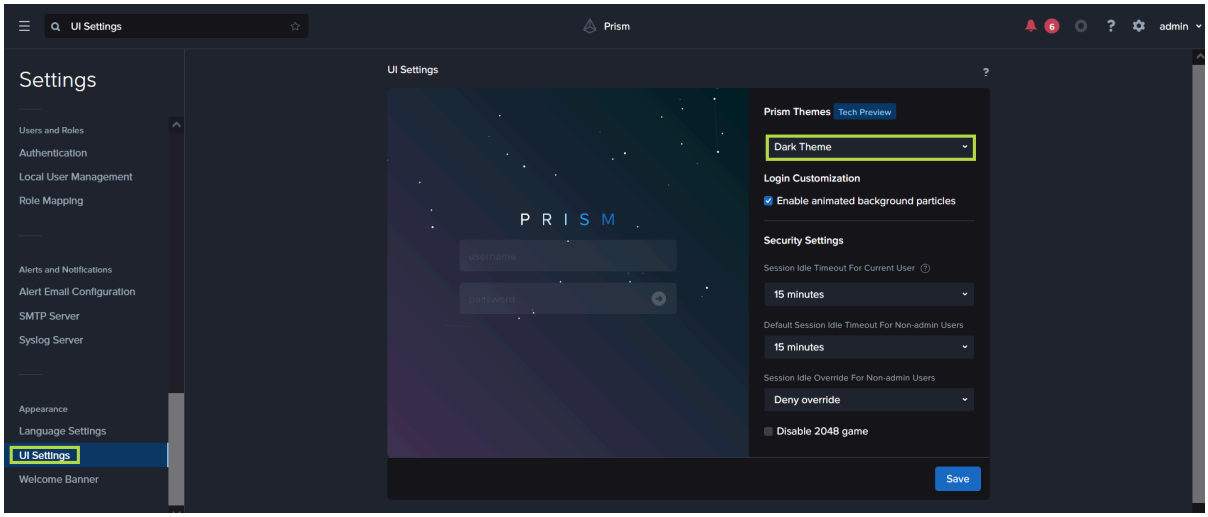

Figure 232: Prism Theme (Dark Mode)

- 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.

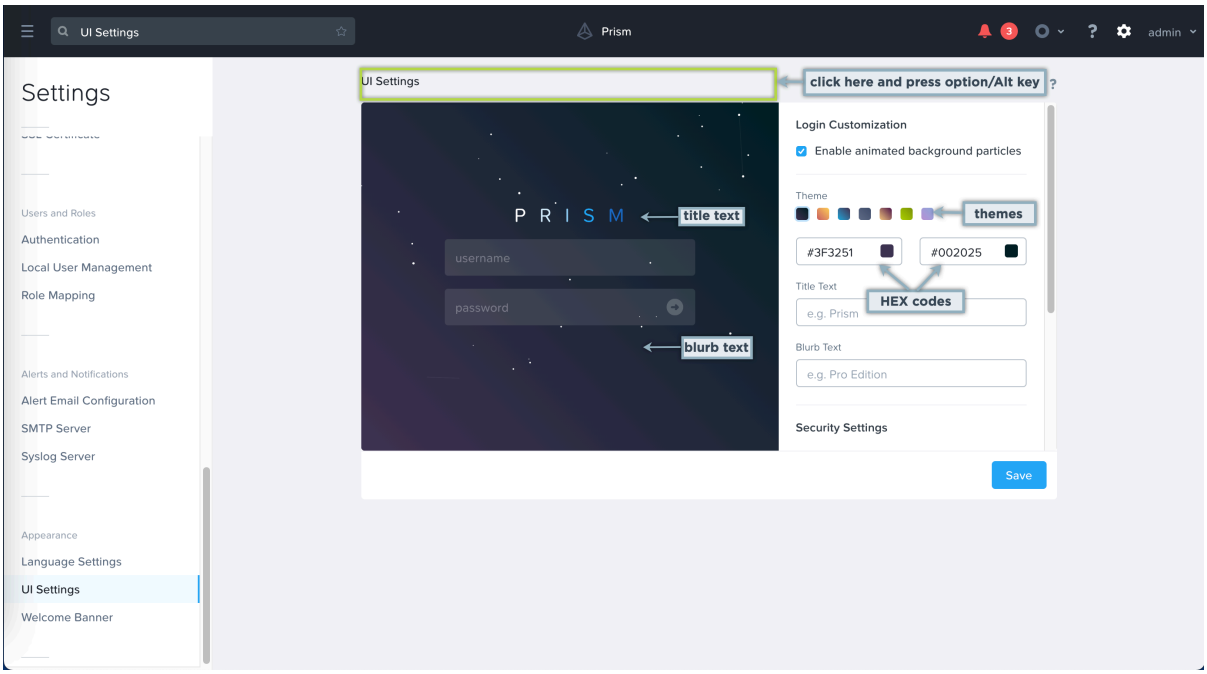

Figure 233: UI Settings Window for customizing the theme, title text, and blurb text

5. To configure session timeout, do the following under **Security Settings**:

   - Select the session timeout for the current user from the **SESSION TIMEOUT FOR CURRENT USER** drop-down list.
   - Select the default session timeout for all non-administrative users from the **DEFAULT SESSION TIMEOUT FOR NON-ADMIN USERS** drop-down list.
   - Select the appropriate option from the **SESSION TIMEOUT OVERRIDE FOR NON-ADMIN USERS** drop-down list to override the session timeout for non-administrative users.

   > **Note:** The timeout interval for an administrator cannot be set for longer than 1 hour.

6. Clear the **Disable 2048 game** option to disable the 2048 game.
7. Click **Save** to save your changes and close the window.

Configuring a Banner Page (Prism Central)

You have the option to create a welcome banner, which will be the first screen that appears when a user attempts to log into Prism Central. The content of the banner page is configurable, and it can include a custom message and graphics.

About this task

To configure a banner page, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Welcome Banner** in the **Settings** page.

   The **Edit Welcome Banner** dialog box appears.

2. Enter (paste) the desired content in HTML format in the pane on the left.

   Only "safe" HTML tags are supported. Inline event handlers, scripts, and externally-sourced graphics are not allowed.

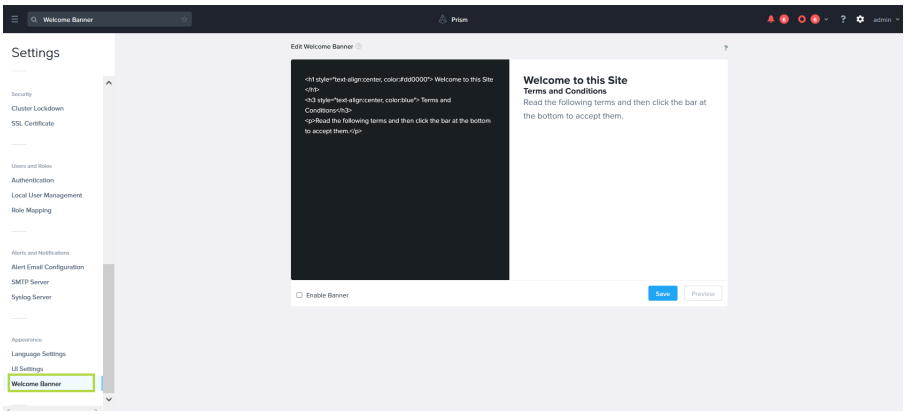

Figure 234: Welcome Banner Window

3. Click the **Preview** button to display the banner in the pane on the right.
4. If the banner is not correct, update the HTML code as needed until the preview pane displays the desired message.
5. When the preview is correct, check the **Enable Banner** box (lower left) and then the **Save** button.

   A live banner page includes an "Accept terms and conditions" bar at the bottom. Clicking on this bar sends the user to the login page.

   You can disable (or enable) the banner at any time by unchecking (checking) the **Enable Banner** box.

Internationalization (i18n) (Prism Central)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 733: Internationalization Support for Prism Central

Supported Entities	Unsupported Entities
User-defined dashboard name	Password fields

Supported Entities	Unsupported Entities
Custom widget name	Static dashboard name
First and last name under **Update Profile**	Static widget name
User name, first name, and last name under **User Management**	
Chart name	

Localization (L10n) (Prism Central)

Nutanix localizes the user interface in Simplified Chinese and Japanese language. All the static screens are translated to the selected locale language.

You have an option to change the language settings of the cluster from English (default) to Simplified Chinese or Japanese. For more information, see Changing the Language Settings (Prism Central) on page 532.

If the Prism Element instance is launched from Prism Central, language settings of Prism Central takes precedence over Prism Element.

The dashboards (including tool tips) and menus of Prism Central are localized.

Guidelines and Limitations

- Logical entities that do not have a contextual translation available in the localized language are not translated.
- The AOS generated alerts and events are not localized to the selected locale language.
- Following strings are not localized: VM, CPU, vCPU, Language Settings, licensing details page, hardware names, storage denominations (GB, TB), About Nutanix page, EULA, service names (SNMP, SMTP), hypervisor types.

Changing the Language Settings (Prism Central)

Perform the following procedure to change the language settings in Prism Central. You can change the language setting to Simplified Chinese or Japanese.

Procedure

1. Click the gear icon in the main menu and then select **Language Settings** in the **Settings** page.

   The Language Setting window appears. The **English** language is selected by default.

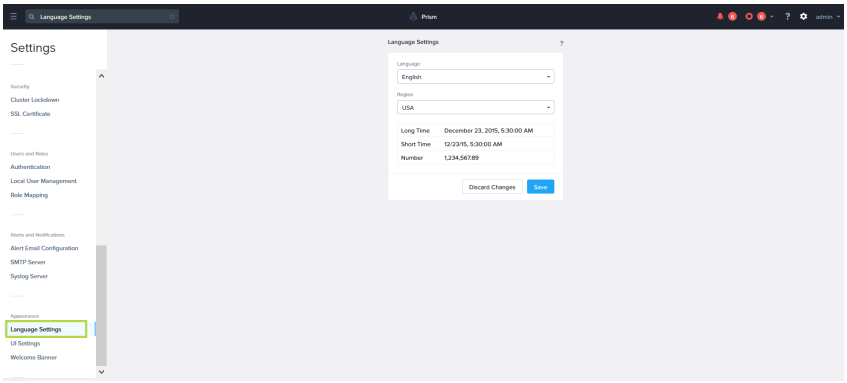

Figure 235: Language Settings Window

2. To change the setting, do the following:

   » To change the language setting of the cluster to Simplified Chinese, select **Simplified Chinese** from the drop-down menu.

   » To change the language setting of the cluster to Japanese, select **Japanese** from the drop-down menu.

   » To change the locale settings (date, time, calendar), select the appropriate region from the **Region** drop-down menu.

   By default, the locale is set to the language setting that you have set in the **Language** drop-down menu. However, you can change the **Region** to display the date, time, or calendar in some other format. This format for date, time, and calendar is applied for the entire cluster.

3. Click **Save**.

   The language and locale settings (date, time, calendar) is changed according to the selection. For example, in the below screen shot, once you click **Save** the language setting for the cluster is changed to Chinese and locale setting is changed to Russian. For more information on the entities that are supported in Simplified Chinese, see Internationalization (i18n) (Prism Central) on page 531. Also, the user interface is localized according to the selection. For more information about localization, see Localization (L10n) (Prism Central) on page 532.

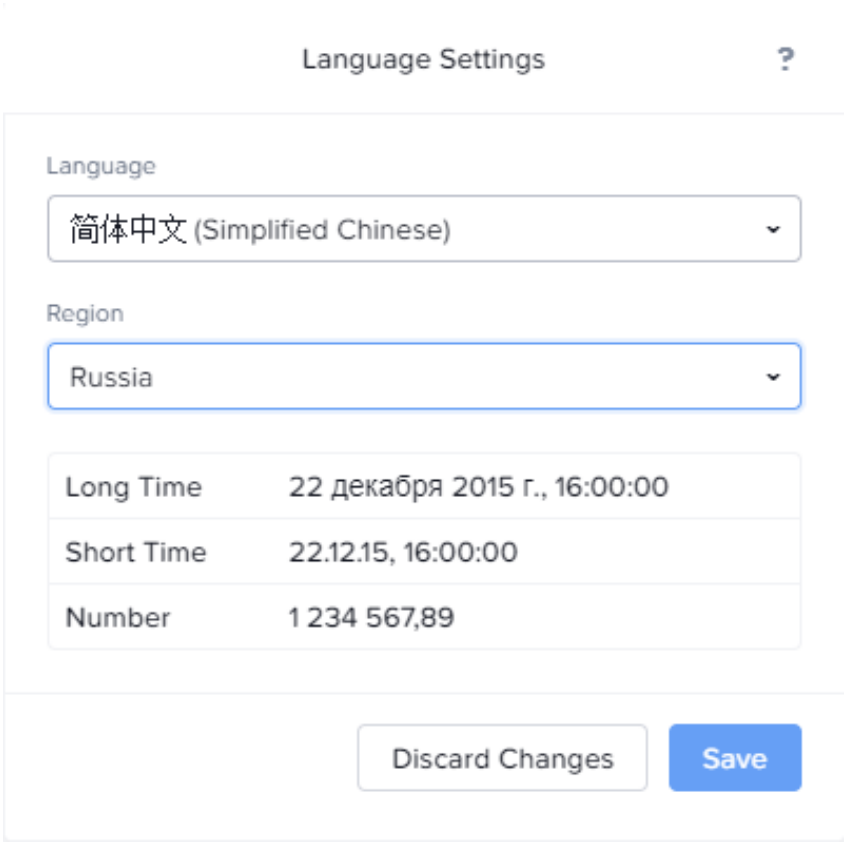

Figure 236: Localized Settings (Chinese/Russian)

Note: If you are logged on as a Domain administrator and the language settings changes are not saved, you might need to update the *User logon name* for the Domain administrator account. For more information, see KB10166.

SECURITY AND USER MANAGEMENT (PRISM CENTRAL)

Prism Central provides several mechanisms to manage security and control user access. For more information, see the following topics in the Security Management using Prism Central (PC) chapter of *Nutanix Security Guide*.

- Set the user authentication method to local, directory service, or both (see Configuring Authentication)
- Add, edit, or delete local user accounts (see Managing Local User Accounts)
  - Update your account (see Updating My Account)
- Assign roles to users (see Controlling User Access (RBAC))

  Note: For recovery plan RBAC, see Nutanix Disaster Recovery Guide.
- Install SSL certificates (see Installing an SSL Certificate)
- Control SSH access to Prism Central (see Controlling Remote (SSH) Access)
- Create and apply security-related policies (see Security Policies in the *Flow Microsegmentation Guide*).

COMPUTE AND STORAGE (CLUSTER) ADMINISTRATION

You can view hardware and infrastructure information about registered (individual) clusters through the Entities menu (see Hardware Entities on page 201 and Compute and Storage Entities on page 93). In addition, you can configure certain components on a registered cluster directly through Prism Central.

- To create or manage a VM on an AHV or ESXi cluster, see
  - AHV: Creating a VM through Prism Central (AHV) on page 535, Creating a VM (Self Service) on page 545, Managing a VM through Prism Central (AHV) on page 552, and Managing a VM (Self Service) on page 566
  - ESXi: Creating a VM through Prism Central (ESXi) on page 580 and Managing a VM through Prism Central (ESXi) on page 585
- To upload disk or ISO images that can by applied when creating a VM, see Image Management on page 670.
- To add images or VMs to a catalog that can be applied by a self-service portal administrator to project users, see Catalog Management on page 666.
- To configure affinity polices for an AHV cluster, see Affinity Policies Defined in Prism Central on page 608
- To perform other administrative tasks for a cluster, see Performing Other Administrative Tasks on page 717.
- To register (or unregister) an ESXi cluster with vCenter, see vCenter Server Integration on page 718.

VM Management

You can create and manage VMs directly from Prism Central when the hypervisor is either ESXi or AHV.

- ESXi
  - To create a VM, see Creating a VM through Prism Central (ESXi) on page 580.
  - To manage or modify a VM, see Managing a VM through Prism Central (ESXi) on page 585.
- AHV
  - To create a VM, see Creating a VM through Prism Central (AHV) on page 535.
  - To manage or modify a VM, see Managing a VM through Prism Central (AHV) on page 552.
- Prism Self Service (AHV)
  - To create a VM, see Creating a VM (Self Service) on page 545.
  - To manage or modify a VM, see Managing a VM through Prism Central (AHV) on page 552.

Creating a VM through Prism Central (AHV)

You can create virtual machines (VMs) in Acropolis managed clusters through Prism Central.

About this task

Note: For instructions about how to install Nutanix VirtIO on Windows VMs, see Windows VM Provisioning in the *AHV Administration Guide*.

To create a VM, do the following:

> **Note:** If you are logged in as the self-service administrator or a project member, see Creating a VM (Self Service) on page 545.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 94) and click the **Create VM** button.
   The **Create VM** wizard appears.

2. In the **Configuration** step, do the following in the indicated fields:

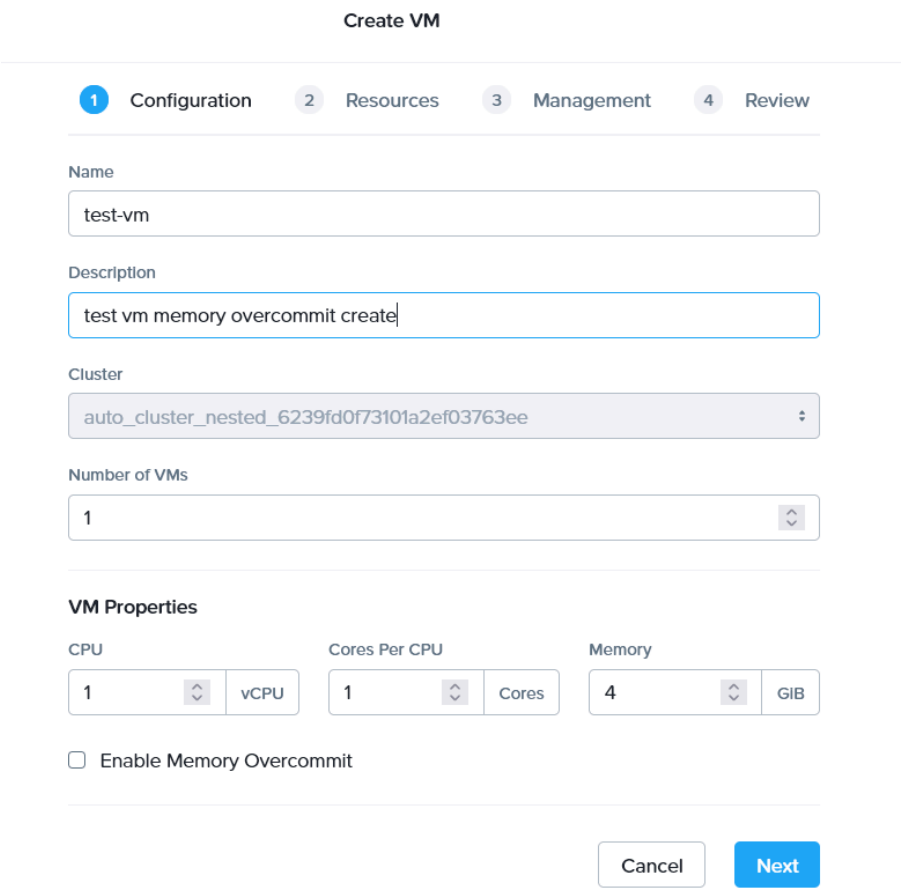

Figure 237: Create VM - Configuration

a. **Name**: Enter a name for the VM.

b. **Description** (optional): Enter a description for the VM.

c. **Cluster**: Select the target cluster from the drop-down list on which you intend to create the VM.

d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.

e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.

f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.

g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

h. **Enable Memory Overcommit**: Select this checkbox to enable memory overcommit for the VM.

Before you enable or disable memory overcommit, see Memory Overcommit in the *AHV Administration Guide*. For more information about configuring memory overcommit, see Memory Overcommit Management on page 604.

> **Note:** Ensure that the status of the **Memory Overcommit** property under **Configuration** in the **Review** tab is set as **Enabled**.

3. In the **Resources** step, do the following.

   Disks: To attach a disk to the VM, click the **Attach Disk** button. The **Add Disks** dialog box appears. Do the following in the indicated fields:

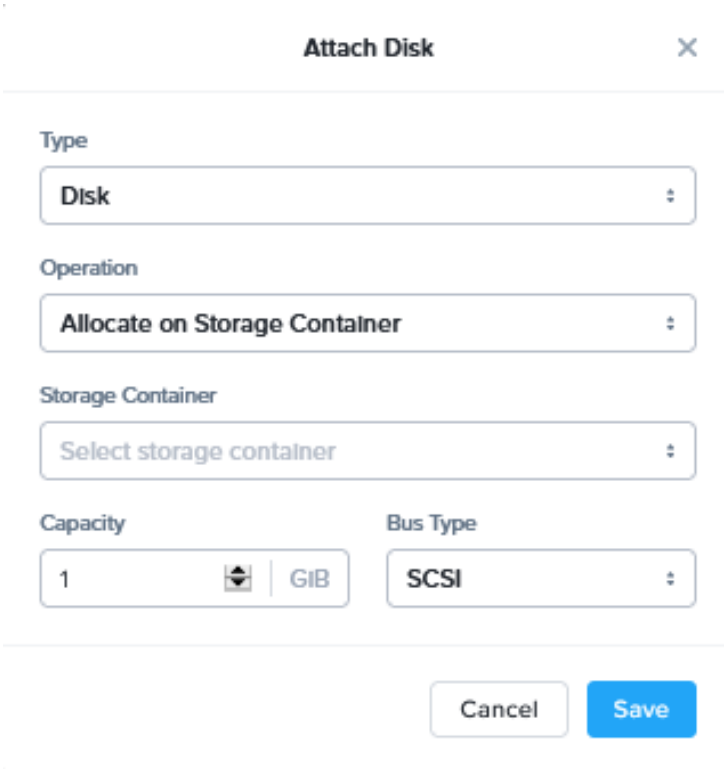

Figure 238: Add Disk Dialog Box

   a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the drop-down list.

   b. **Operation**: Specify the device contents from the drop-down list.

      - [CD-ROM only] Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
      - [Disk only] Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
      - Select **Clone from Image** to copy an image that you have imported by using image service feature onto the disk.

   c. **Bus Type**: Select the bus type from the pull-down list.

      The options displayed in the pull-down list varies based on the storage **Type** selected in Step a.

      - For type **Disk** the choices are **SCSI**, **SATA**, **PCI**, or **SATA**.
      - For type **CD-ROM** the choices are **IDE**, or **SATA**.

   d. **Clone from Image**: Select the image that you have created by using the image service feature.

      This field appears only when **Clone from Image Service** is selected. It specifies the image to copy.

> **Note:** If the image you created does not appear in the list, see KB-4892.

> **Note:** The image transfer can trigger image bandwidth throttling if a bandwidth throttling policy is associated with the image. For more information, see Bandwidth Throttling Policies on page 695.

e. **Storage Container**: Select the storage container to use from the drop-down list.

   This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.

f. **Capacity**: Enter the disk size in GiB.

g. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the **Create VM** dialog box.

h. Repeat this step to attach additional devices to the VM.

4. In the **Resources** step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The **Attach to Subnet** dialog box appears.

   Do the following in the indicated fields:

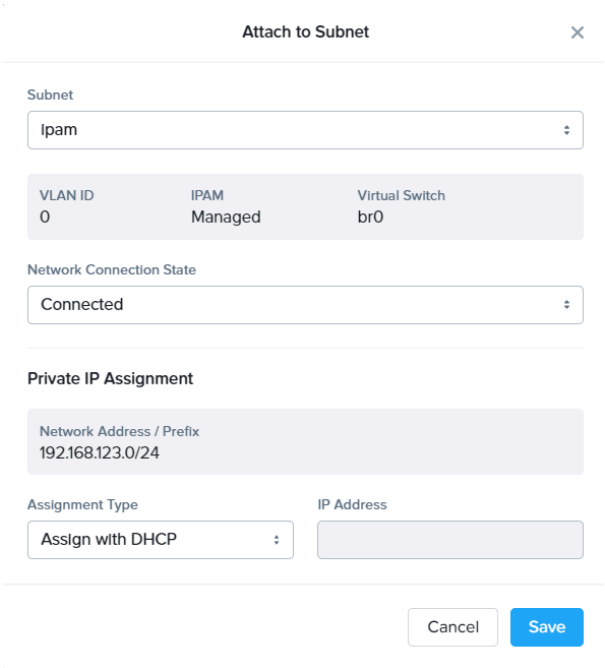

Figure 239: Attach to Subnet Dialog Box

   a. **Subnet Name**: Select the target subnet from the drop-down list.

      The list includes all defined networks (see Network Configuration for VM Interfaces in *Prism Element Web Console Guide*).

   b. **VLAN ID**: This is a read-only field that displays the VLAN ID.

   c. **IPAM**: This is a read-only field that informs you if the subnet is IPAM managed or not.

   d. **Virtual Switch**: This is a read-only field that displays the name of the virtual switch associated with the subnet.

   e. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are `Connected` or `Disconnected`.

   f. **Network Address/Prefix**: This is a read-only field that displays the network IP address and prefix.

   g. **Assignment Type**: Select the IP assignment type from the drop-down list. The options are `Assign with DHCP`, `Assign Static IP`, and `No Private IP`.

   h. **IP Address**: Enter an IP address for the VLAN.

      This field appears only if the NIC is placed in a managed network. Entering an IP address in this field is optional when the network configuration provides an IP pool. If the field is left blank, the NIC is assigned an IP address from the pool.

   i. When all the field entries are correct, click the **Save** button to create a network interface for the VM and return to the **Create VM** dialog box.

j. Repeat this step to create additional network interfaces for the VM.

5. In the **Resources** step, do the following.

   Boot Configuration: Select one of the following firmware to boot the VM.

   » **Legacy BIOS**: Select legacy BIOS to boot the VM with legacy BIOS firmware.

   » **UEFI**: Select UEFI to boot the VM with UEFI firmware. UEFI firmware supports larger hard drives, faster boot time, and provides more security features. For more information about UEFI firmware, see UEFI Support for VM section in the *AHV Administration Guide*.

6. In the **Resources** step, do the following.

   GPUs: (For GPU-enabled AHV clusters only) To add a GPU, click the **Add GPU** button. The **Add GPU** dialog box appears. Do the following in the indicated fields:

   a. **GPU Mode**: Click the radio button for the desired mode, either **vGPU** or **Passthrough**.

   b. If you selected **vGPU**, do the following:

   - **NVIDIA Virtual GPU License**: Select a license type from the drop-down list. This sets (filters the list of) available profiles. Click the help link (circled question mark) for information about the license types.

     Note:

     You can add multiple vGPUs to the VM only if you select the license for NVIDIA Virtual GPU software version 10.1 (440.53) or later.

     Before you add multiple vGPUs to the VM, see Multiple Virtual GPU Support and Restrictions for Multiple vGPU Support in the *AHV Admin Guide*.

   - **vGPU Profile**: Click the radio button for the desired profile. Click the help links (Virtual Slice column and end of line) for more information about the profiles.

     Note:

     Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.

   - After you add the first vGPU, to add multiple vGPUs, see Adding Multiple vGPUs to the Same VM on page 573.

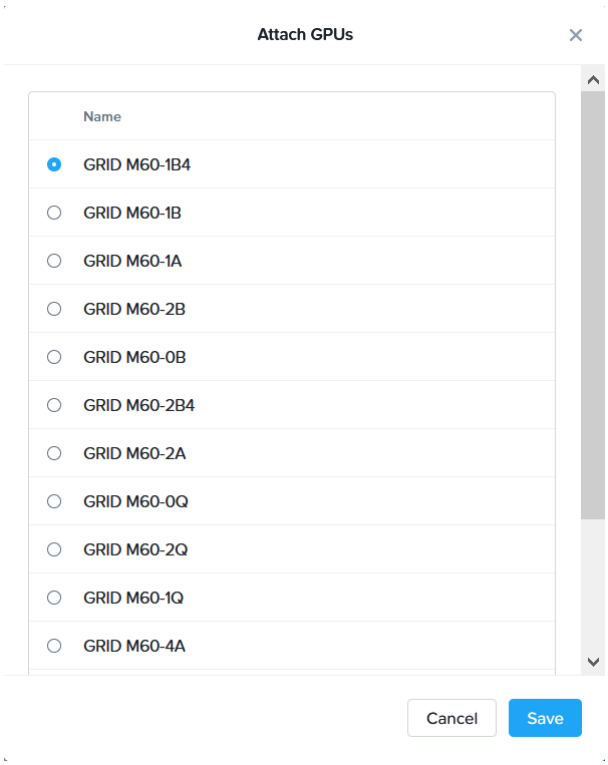

Figure 240: Add GPU Dialog Box: vGPU

c. If you selected **Passthrough**, click the radio button for the desired type.

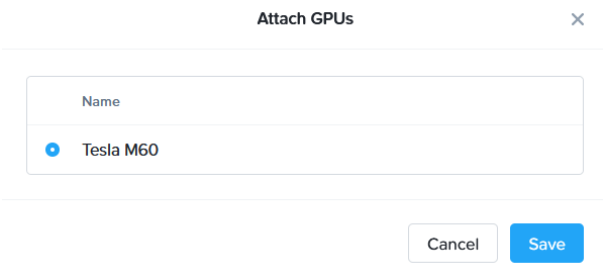

Figure 241: Add GPU Dialog Box: Passthrough

d. Click the **Add** button.

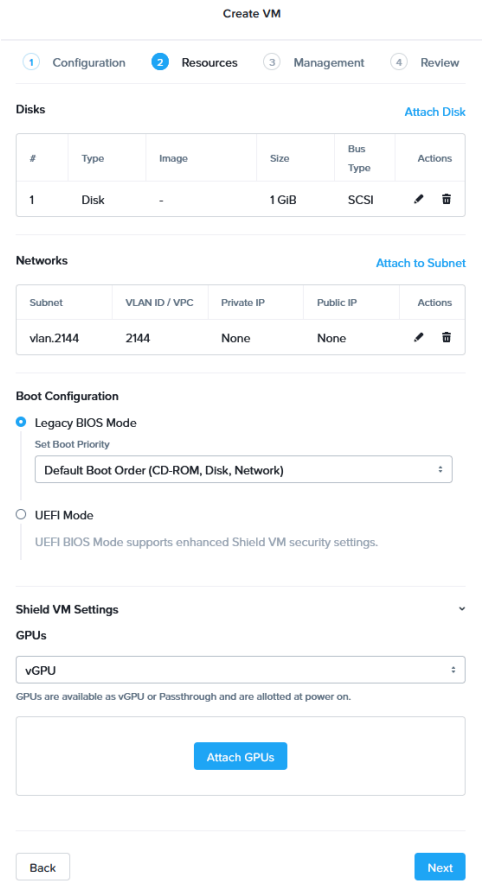

Figure 242: Create VM Window (Resources)

7. In the **Management** step, do the following.

   a. **Categories**: Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

   b. **Timezone**: Select the local timezone to use from the drop-down list.

   c. **Use this VM as an agent VM**: Select this option to make this VM as an agent VM.

      You can use this option for the VMs that must be powered on before the rest of the VMs (for example, to provide network functions before rest of VMs are powered on the host) and must be powered off and migrated after rest of the VMs (for example, during maintenance mode operations).

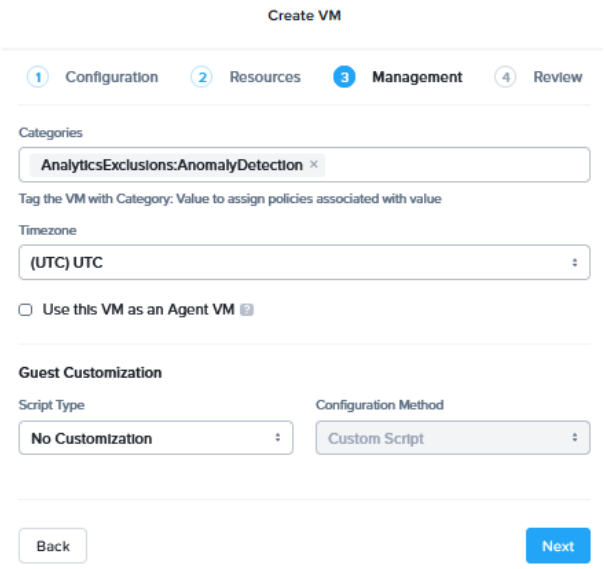

Figure 243: Create VM Window (Management)

8. In the **Management** step, do the following.

   Guest Customization: To customize the VM, select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs).

   Script Type: Select the VM script customization type from the drop-down list. The options are `No Customization`, `Sysprep(Windows)`, or `Cloud-init (Linux)`.

   Configuration Method: Select the VM script customization method from the drop-down list. This field is activated when the **Script Type** is either `Sysprep(Windows)` or `Cloud-init (Linux)`. The options are `Custom Script` or `Guided Script`.

   Field required for configuring Cloud-init and Sysprep, such as options for specifying a configuration script and option to upload script appears.

9. To specify a user data file (Linux VMs) or answer file (Windows VMs) the script for unattended provisioning, do one of the following.

   » If the file is available on your local computer, click the **Upload Script** button, choose and upload the file.

   » Create or paste the contents of the file in the text box below the **Upload Script** button.

   Note: The script type supports the following file formats.

   - Sysprep: XML
   - Cloud-init (Linux): YAML, JSON, or Shell.

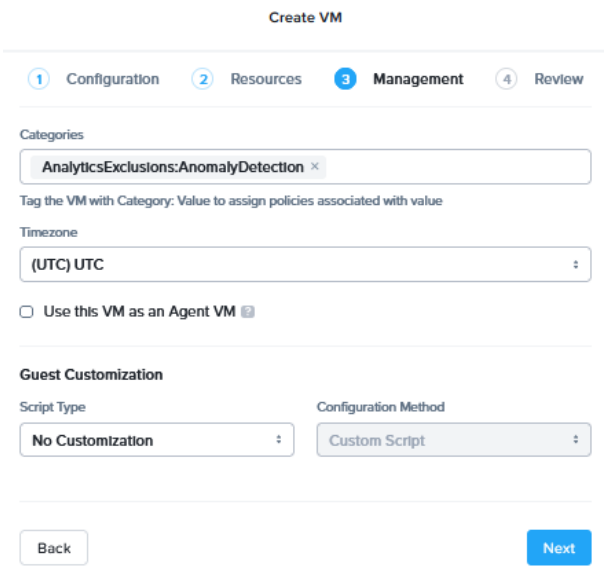

Figure 244: Create VM Window (Management)

10. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the **Create VM** dialog box.

    The new VM appears in the VMs entity page list.

Creating a VM (Self Service)

About this task

When you are logged in as a self-service administrator or a project member with permission to create a VM (see Prism Self Service Overview on page 723), you can create a VM in two ways:

- You can create a VM through the usual workflow that an admin user follows to create a VM (see Creating a VM through Prism Central (AHV) on page 535)
- You can also create a VM based on a source file stored in the Prism Central catalog (see Creating a VM from Catalog Items (Self Service) on page 546).

Creating a VM from Catalog Items (Self Service)

About this task

When you are logged on as a self-service administrator or a project member with permission to create a VM (see Prism Self Service Overview on page 723), you can create a VM based on a source file stored in the Prism Central catalog (see Catalog Management on page 666). To create a VM in this case, do the following:

> **Note:** If you are not logged on as the self-service administrator or a project member or you are creating a VM not based on a source file stored in the Prism Central catalog, see Creating a VM through Prism Central (AHV) on page 535.

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 94).
2. Click the **Create VM from Catalog Item** button.

   The **Create VM from Catalog Items** page appears.
3. In **Catalog Type**, select the source image for the VM.

   - Click **VM Template** if you want to create a VM from a template in the catalog.

     A VM template includes all the configuration information required to create a VM. You can use a template to quickly deploy a VM without having to specify all configuration information.
   - Click **Image** if you want to create a VM from a mounted disk image.

     Disk images can be CD-ROM images such as installer ISO images or images of hard drives that contain preinstalled applications and data. Disk images enable you to share data with other VMs, but you need to specify configuration information to create the VM that will use the data.

4. In **Catalog Item**, do one of the following:

   - For **VM Template** select the target VM template. The available VM templates from the catalog are listed. If the target VM image does not appear in the list, you can search for it by name in the search field.

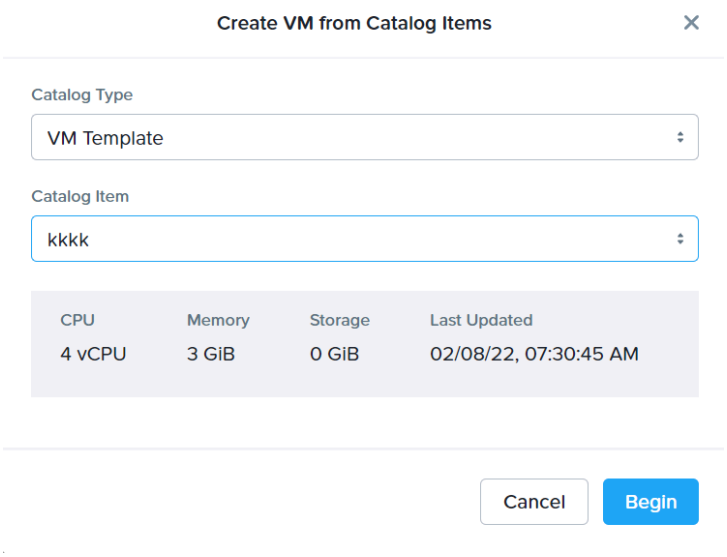

Figure 245: Create VM Window (Configuration)

   - For **Image** select one or more of the disk images. The available disk images from the catalog are listed. If the target disk image does not appear in the list, you can search for it by name in the search field.

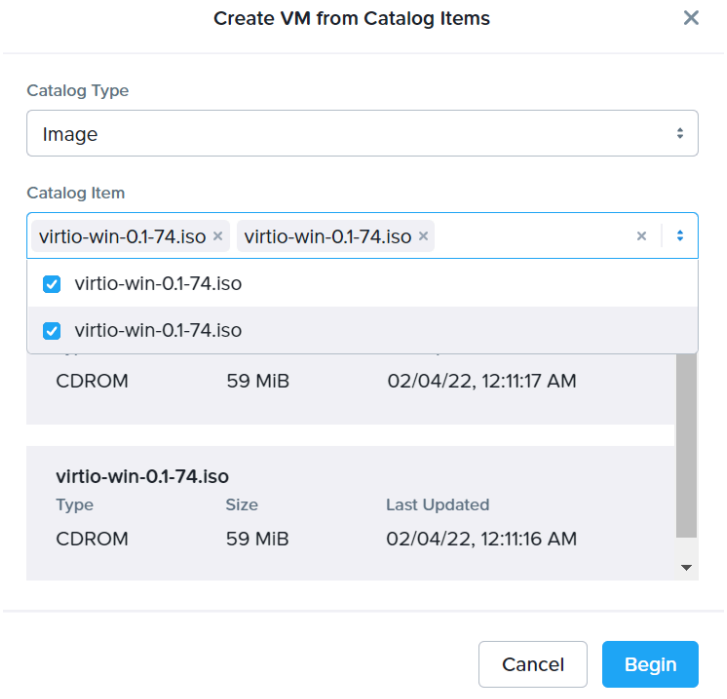

Figure 246: Create VM Window (Configuration)

5. Click **Begin**.

   The **Create VM** page appears.

6. In the **Configuration** step, do the following in the indicated fields:

   a. **Name**: Enter a name for the VM.

   b. **Project**: Select the project associated with this VM from the drop-down list.

   c. **Cluster**: Select the target cluster from the drop-down list on which you intend to create the VM.

   d. **Number of VMs**: Enter the number of VMs you intend to create. The VM names are suffixed with sequential numbers (1 to 5).

   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 GB) to allocate to this VM.

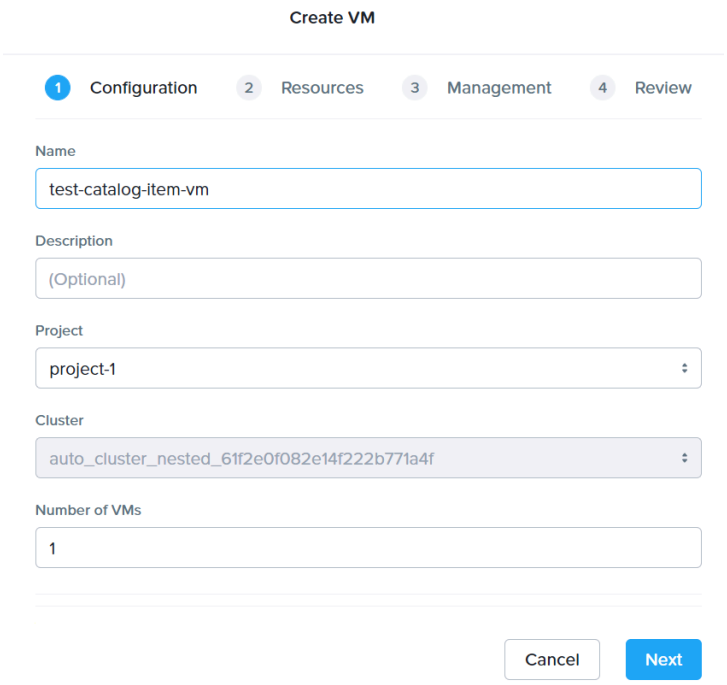

Figure 247: Create VM Window (Configuration)

7. In the **Resources** step, do the following.

   Disks: Displays the disks and CD-ROMs attached to this VM.

   - For **VM Template** the list is fixed. Click the radio button for the device to boot from.
   - For **Image** there are **+ New Disk** and **+ New CDROM** links above the list.
     - Click **+ New Disk** to add a disk. This displays a new line at the bottom of the list. Specify the disk size and then click **Save** at the end of line to add the disk. (The name and type field values are entered automatically.)
     - Click **+ New CDROM** to add a CD-ROM. This displays a new line at the bottom of the list. Click **Save** at the end of line to add the CD-ROM. (All field values are entered automatically.)

   Repeat this step until you have added all the desired disks and CD-ROMs. When the list is complete, click the radio button for the device to boot from.

8. In the **Resources** step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The **Attach to Subnet** dialog box appears.

   Do the following in the indicated fields:

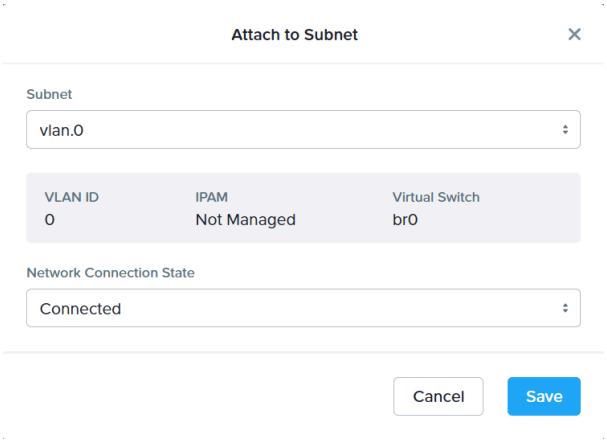

Figure 248: Attach to Subnet Dialog Box

   a. **Subnet Name**: Select the target subnet from the drop-down list.

      The list includes all defined networks (see Network Configuration for VM Interfaces in *Prism Element Web Console Guide*).

   b. **VLAN ID**: This is a read-only field that displays the VLAN ID.

   c. **IPAM**: This is a read-only field that informs you if the subnet is IPAM managed or not.

   d. **Virtual Switch**: This is a read-only field that displays the name of the virtual switch associated with the subnet.

   e. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are `Connected` or `Disconnected`.

   f. When all the field entries are correct, click the **Save** button to create a network interface for the VM and return to the **Create VM** dialog box.

   g. Repeat this step to create additional network interfaces for the VM.

9. In the **Resources** step, do the following.

   Boot Configuration: Select one of the following firmware to boot the VM.

   » **Legacy BIOS Mode**: Select legacy BIOS to boot the VM with legacy BIOS firmware.

   » **UEFI BIOS Mode**: 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 the UEFI Support for VM section in the *AHV Administration Guide*.

10. In the **Management** step, do the following.

    a. **Categories** (optional): To assign the VM a category value, enter the category name or select one from the drop-down list. The policies associated with the category value are assigned to the VM.

    This field acts like a search field; it provides a list of matching categories as you enter a string. Select the desired category value when you see it in the list. You can repeat this step for as many categories as desired.

11. In the **Management** step, do the following.

    Guest Customization: To customize the VM, select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs).

    Script Type: Select the VM script customization type from the drop-down list. The options are `No Customization`, `Sysprep(Windows)`, or `Cloud-init (Linux)`.

    Configuration Method: Select the VM script customization method from the drop-down list. This field is activated when the **Script Type** is either `Sysprep(Windows)` or `Cloud-init (Linux)`. The options are `Custom Script` or `Guided Script`.

    Field required for configuring Cloud-init and Sysprep, such as options for specifying a configuration script and option to upload script appears.

12. To specify a user data file (Linux VMs) or answer file (Windows VMs) the script for unattended provisioning, do one of the following.

    » If the file is available on your local computer, click the **Upload Script** button, choose and upload the file.

    » Create or paste the contents of the file in the text box below the **Upload Script** button.

    Note: The script type supports the following file formats.

    - Sysprep: XML
    - Cloud-init (Linux): YAML, JSON, or Shell.

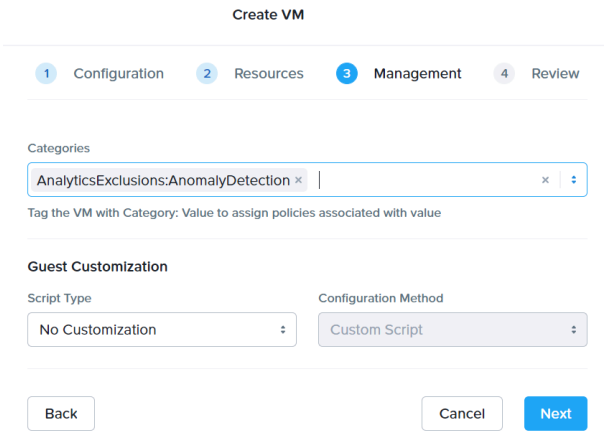

Figure 249: Create VM Window (Management)

13. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the **Create VM** dialog box.

    The new VM appears in the VMs entity page list.

Managing a VM through Prism Central (AHV)

About this task

After creating a VM (see Creating a VM through Prism Central (AHV) on page 535), you can use Prism Central to perform operations on the VM.

You can perform these operations by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 94) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the dropdown list.
- Go to the details page of a selected VM (see VM Details View on page 105) and select the desired action.

You can perform some of the available actions on a single VM at a time, while others can be performed on multiple VMs simultaneously. The available actions appear in bold; the unavailable actions are grayed out. The available actions depend on the current state of the VM and your permissions.

> **Note:** You can perform only those operations for which you have permissions from the admin.

Procedure

1. Select the target VM and choose the action that you want to perform. See the following table for the actions and their descriptions.

Table 734: VM Actions

Action	Description	Applicable to Multiple VMs
Update	Update the VM configuration. For more information, (see Updating a VM through Prism Central (AHV) on page 555.	No
Delete	Delete the VM.	Yes
Clone	Clone a VM. For more information, see Cloning a VM through Prism Central (AHV) on page 559.	No
Create VM Template	Create a VM template. For more information, see Creating a VM Template on page 614.	No
Launch Console	Launch the VM console. For more information, see Launching a VM Console through Prism Central (AHV) on page 560.	No
Power On	Power on the VM.	Yes
Power Off	Power off the VM.	Yes
Power Cycle	power off the VM followed by a power on.	Yes
Reset	Perform an ACPI reset action through the BIOS of the VM.	No
Guest Shutdown	Perform a graceful shutdown of the operating system of the VM.	No
Guest Reboot	Perform a graceful restart of the operating system of the VM.	No
Pause/Suspend	Pause or suspend the VM. This option is available only when the VM is powered on.	Yes
Resume	Resume a paused or suspended VM.	No

Action	Description	Applicable to Multiple VMs
Disable Efficiency Measurement	Disable efficiency measurement for the VM. You can also create a category of VMs and assign the category in **Inefficiency Measurement Exclusion** policy to control this function through operations policies. For more information, see Operations Policy Management on page 873. **Note:** By default, the efficiency measurement function is enabled for all VMs.	Yes
Enable Efficiency Measurement	Enable efficiency measurement for the VM. **Note:** If the VM is part of the VM category that is mapped to the **Inefficiency Measurement Exclusion** operations policy, the policy takes precedence and the system does not allow you to enable efficiency measurement. For information about how to assign or remove the VM category in **Inefficiency Measurement Exclusion** policy, see Modifying an Operations Policy on page 874. You cannot enable efficiency measurement on individual VMs if the efficiency measurement is disabled on the cluster level.	Yes
Disable Anomaly Detection	Disable Anomaly Detection for the VM. You can also create a category of VMs and assign the category in **Anomaly Detection Exclusion** policy to control this function through operations policies. For details, see Operations Policy Management on page 873. **Note:** By default, the anomaly detection function is enabled for all VMs.	Yes
Enable Anomaly Detection	Enable Anomaly Detection for the VM. **Note:** If the VM is part of the VM category that is mapped to the **Anomaly Detection Exclusion** operations policy, the policy takes precedence and the system does not allow you to enable efficiency measurement. For information about how to assign or remove the VM category in **Anomaly Detection Exclusion** policy, see Modifying an Operations Policy on page 874.	Yes
Protect	Assign the VM to a protection policy. Selecting **Protect** opens a dialog box to specify the protection policy to which this VM should be assigned. You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Nutanix Disaster Recovery and configuring the Availability Zone.	Yes
Unprotect	Remove the VM from a protection policy.	Yes

Action	Description	Applicable to Multiple VMs
Create Recovery Point	Create VM recovery point. The VM can be restored or replicated from a recovery point either locally or remotely in a state of a chosen recovery point. In the **Create VM Recovery Point** dialog box, enter a name for the VM recovery point. You can choose to create an app consistent VM recovery point by enabling the **App Consistent** check-box. For more information about app consistent recovery point, see Terminology section in the *Data Protection and Recovery with Prism Element* guide.	No
Migrate	Migrate the VM to another host. In the **Migrate VM** dialog box, select the target host from the pull-down list (or select the System will automatically select a host option to let the system choose the host) and then click the Migrate button to start the migration. **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. For more information about VM Migration, see VM Migration Specifications on page 561.	No
Add to Recovery Plan	Add the VM to a recovery plan you created previously. For more information, see the Adding Guest VMs Individually to a Recovery Plan section in the *Nutanix Disaster Recovery Guide*.	Yes
Run Playbook	Run a playbook you created previously. For more information, see Running a Playbook (Manual Trigger) on page 932.	Yes
Manage Categories	Assign the VM a category value. For more information, see Assigning a Category on page 797.	Yes
Install NGT	Install Nutanix Guest Tools (NGT). For more information, see Installing NGT on page 594.	Yes
Manage NGT Applications	Enable or disable NGT. For more information, see Managing NGT Applications on page 596.	Yes
Upgrade NGT	Upgrade NGT. For more information, see Upgrading NGT on page 597.	Yes
Configure VM Host Affinity	This option is disabled and cannot be used to define VM-Host affinity in Prism Central. The VM-Host affinity policies can be created using VM categories and host categories. For more information about affinity policies Prism Central, see Affinity Policies Defined in Prism Central on page 608. For information about how to create an affinity policy, see Creating an Affinity Policy on page 610. It is recommended to create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure.	No
Add to Catalog	Add the VM to the catalog. For more information, see Adding a Catalog Item on page 667.	No

Action	Description	Applicable to Multiple VMs
Manage Ownership	Specify a project and user who owns the VM. In the **Manage VM Ownership** dialog box, select the target project from the pull-down list and enter a user name as owner of the VM.	No
Set QoS Attributes	Configure quality of service (QoS) settings. For more information, see Setting QoS for an Individual VM on page 602.	Yes
Export as OVA	Export the VM as OVA. For more information, see OVA Management on page 698 and Uploading an OVA on page 700.	No

2. You can check the status of the VM actions under tasks widget.

Updating a VM through Prism Central (AHV)

You can use Prism Central to update the VM configuration.

Procedure

1. Follow any one of the navigation paths mentioned in Managing a VM through Prism Central (AHV) on page 552 for selecting a VM to perform an operation.

2. To modify the VM configuration, select **Update**.

   The **Update VM** dialog box appears, which includes the same fields as the **Create VM** dialog box. Make the desired changes and then click the **Save** button in the **Review** step.

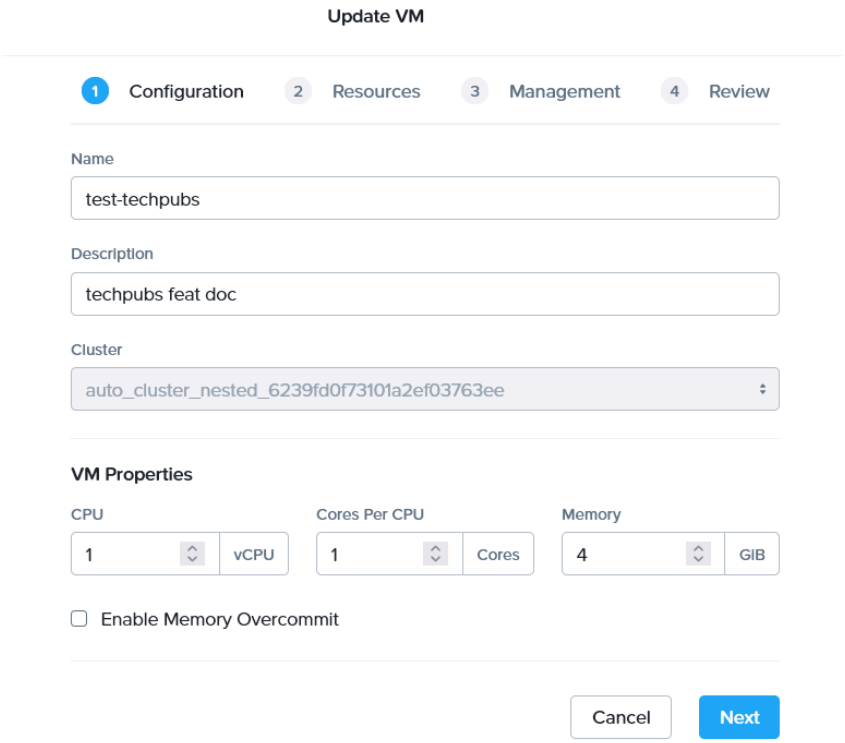

Figure 250: Update VM Window (Configuration)

> **Caution:** See the recommendations specified in KB-9832 before making changes to the VM properties such as memory, vCPU, or storage. Any values other than the recommended memory, vCPU, or storage values for a VM might cause discrepancies in the Prism Central behavior.

3. **Enable Memory Overcommit**: Select this checkbox to enable memory overcommit for the VM.

   Before you enable or disable memory overcommit, see Memory Overcommit in the *AHV Administration Guide*. For more information about configuring memory overcommit, see Memory Overcommit Management on page 604.

   > **Note:** Verify that the status of the **Memory Overcommit** property under **Configuration** in the **Review** tab is set as **Enabled**.

4. **Disks**: You can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM through Prism Central

(AHV) on page 535 before you modify the disk or create a new disk for a VM. You can enable or disable the flash mode settings for the VM and VM disks.

a. To enable flash mode on the VM, click the **Enable Flash Mode** check box.

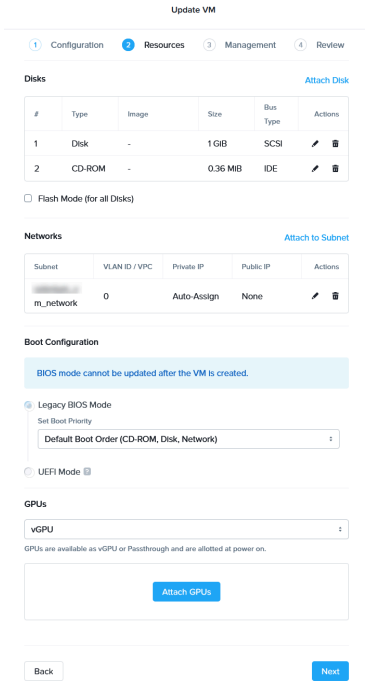

Figure 251: Update VM Window (Resources)

» 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.

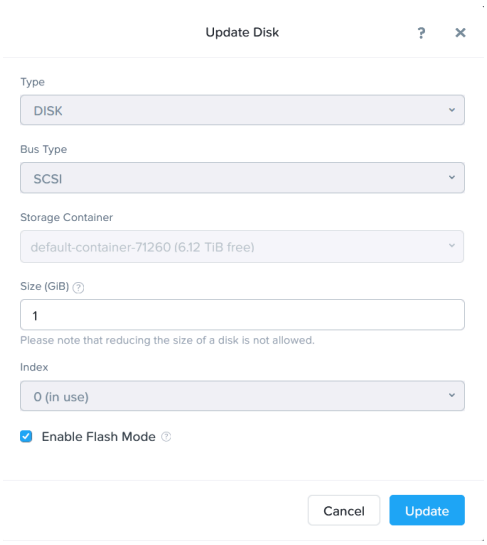

Figure 252: Update Disk Window (VM Disk Flash Mode)

Note: The flash mode setting is not supported for **CD-ROM** storage device.

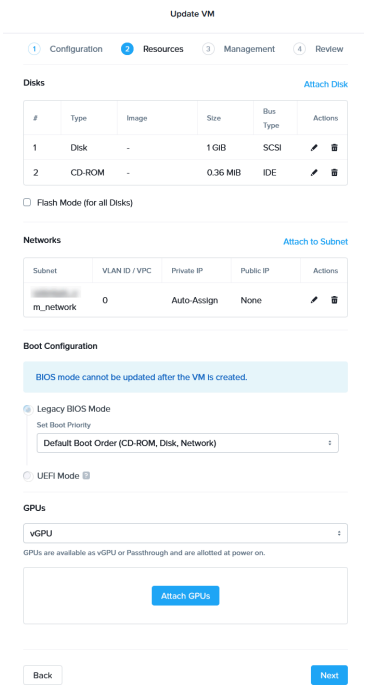

Figure 253: Update VM Window (Resources)

5. **Networks:** you can attach new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM through Prism Central (AHV) on page 535 and Limitation for VNIC Hot-Unplugging topic in *AHV Admin Guide* before you modify NIC network or create a new NIC for a VM.

6. **Boot Configuration:** You cannot modify the BIOS mode after the VM is created. However, if you have enabled the Secure Boot option under UEFI Mode while creating the VM, you can modify this setting during the VM update.
7. **GPUs:** (For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the specified vGPU license. However, you must shut down the VM before you perform these operations.

   - Before you add multiple vGPUs to the VM, see Multiple Virtual GPU Support and Restrictions for Multiple vGPU Support in the *AHV Admin Guide*.
   - Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
   - After you add the first vGPU, to add multiple vGPUs, see Adding Multiple vGPUs to the Same VM on page 573.

Cloning a VM through Prism Central (AHV)

You can use Prism Central to clone a VM.

Procedure

1. Follow any one of the navigation paths mentioned in Managing a VM through Prism Central (AHV) on page 552 for selecting a VM to perform an operation.
2. To clone the VM, select **Clone**.

   This displays the **Clone VM** dialog box, which includes the same fields as the **Create VM** dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking

the **Save** button. For example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

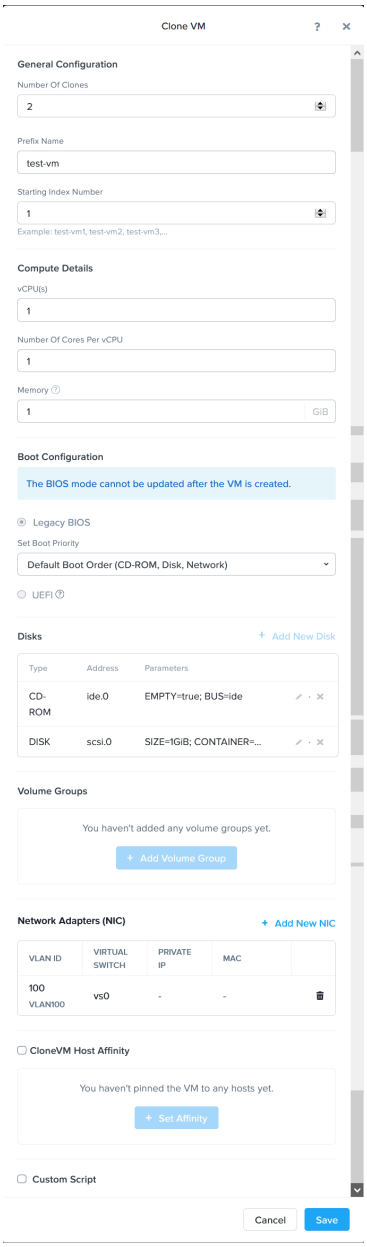

Figure 254: Clone VM Window

Launching a VM Console through Prism Central (AHV)

You can use Prism Central to launch a VM console.

Procedure

1. Follow any one of the navigation paths mentioned in Managing a VM through Prism Central (AHV) on page 552 for selecting a VM to perform an operation.
2. To launch a console window, select **Launch Console**.

   This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

   Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

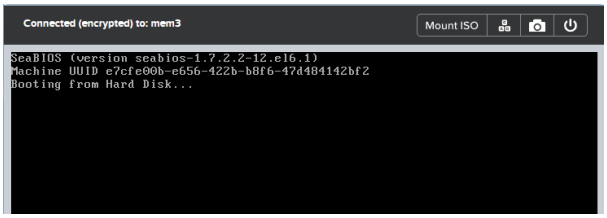

Figure 255: Console Window (VNC)

Note: For enabling or disabling console support for a VM with one vGPU configured, see Enabling or Disabling Console Support for vGPU VMs.

VM Migration Specifications

VM migration can be performed in the following ways:

- *Live Migration* – This method is used to transfer the VM from one physical host to another host without affecting the normal functions and operations of the VM, and with minimal and often no interruption to the users of the VM.

  The live migration is supported for the following scenarios:

  - Intra-cluster live migration from one host to another host. For more information, see:
    - Migrating within the Cluster.
    - Migrating Live a vGPU-enabled VM Within the Cluster section in *Prism Element Web Console Guide* to perform Intra-cluster live migration Prism Element web console.
  - Cross-cluster (Inter-cluster) live migration from one AHV cluster to another cluster using Prism Central. For information about how to perform VM live migration between AHV to AHV cluster using Prism Central, see Migrating outside the cluster.
- *Cold Migration* - This method is used to transfer the VM from one physical host to another host with minimal downtime. You can use **Nutanix Move** for Inter-cluster VM migration <Third-Party Hypervisor> to AHV cluster. For details, see Move User Guide.

  Note: The <Third-Party-Hypervisors> are Xen (Amazon EC2), Microsoft Hyper-V, and VMware ESXi.

Live Migration Cases

VM live migration is applicable for the following cases:

- Maintenance activity required on the cluster - This scenario is applicable for a planned event (for example, scheduled maintenance of guest VMs) at the primary AZ when you put node into maintenance mode. For information about how to configure a planned failover, see Cross-Cluster Live Migration section in *Nutanix Disaster Recovery Guide*.
- AHV or Firmware Upgrades - When you upgrade the AHV or Firmware, the VMs need to be live migrated to another AHV cluster.
- Load balancing or to isolate specific VMs on hosts.
- Changes in Host affinity policy.
- Acropolis Dynamic Scheduling (ADS)- If ADS detects a resource contention in the cluster, it creates a migration plan to eliminate the hotspots in the cluster by migrating VMs from one host to another. For more information about ADS, see Acropolis Dynamic Scheduling in AHV section in *AHV Administration Guide*.
- Defragmentation activity to create resources for VM power-on operation. This activity is performed when the cluster can accommodate resources for VM power-on operation after moving some VMs to another host.

Live Migration Restrictions

The following VM condition (or properties) impacts the VM Live migration.

- VM is powered off.
- VM is configured as an agent VM.

  You can select or clear the **Use this VM as an agent VM** checkbox to enable or disable this function when you create or update a VM. For more information, see Creating a VM through Prism Central (AHV).
- Windows Defender Credential Guard is enabled. For more information about Windows Defender Credential Guard, see Windows Defender Credential Guard Support section in *AHV Administration Guide*.
- WSL2 is enabled. For more information about WSL2, see Windows Subsystem for Linux (WSL2) Support on AHV section in *AHV Administration Guide*.
- GPU passthrough is enabled.

  For information about how to verify if GPU passthrough is enabled for the VM, see Checking Live Migration Status of a VM on page 562.
- CPU passthrough is enabled.
- VM-Host affinity is set from the Prism Element web console with one host.

  For information about how to verify the VM-Host affinity policy from Prism Element web console, see Verifying Affinity Policy Association on page 564.
- VM Affinity policies are defined from Prism Central using VM Categories and Host categories.

  For information about how to verify the affinity policies mapped to a VM from Prism Central, see Checking Affinity Policies of a VM From Prism Central on page 565.
- Defragmentation activity to create resources for VM-power-on is in progress.

Checking Live Migration Status of a VM

This section describes how to check whether the live migration is allowed for the VM.

Procedure

To check if a VM can be live migrated, perform the following steps:

1. Log on to the CVM as a nutanix user using SSH.
2. Run the following command:

   ```
   nutanix@cvm$ acli vm.get <VM_NAME> | grep 'allow_live_migrate
   ```

   The following attribute value confirms the live migration is allowed:

   ```
   allow_live_migrate: True
   ```

What to do next

If `allow_live_migrate: False` is returned for the above command, then check the status of the VM properties that restricts live migration.

The following table provides the information about the verification procedure to check the status of VM properties:

VM Properties	Verification Procedure
Windows Defender Credential Guard	1. Log in to Prism Central. 2. Navigate to **Compute & Storage** > **VMs**. The system displays the VMs **Lists** tab. 3. Select the target VM and click **Update**. 4. Observe the **Shield VM Security Settings** to confirm if Windows Credential Guard is enabled. 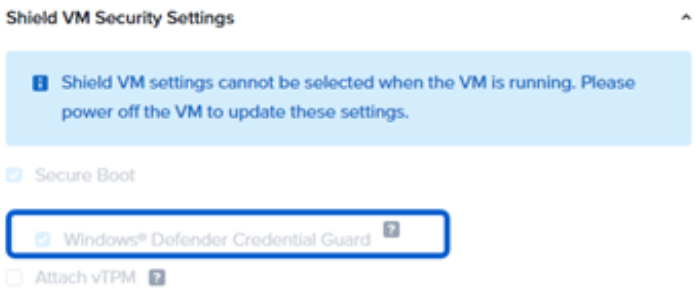 **Figure 256: Shield VM Security Settings**

WSL2	Log on to the CVM as a Nutanix user using SSH, and run the following command:
	`nutanix@cvm$ acli vm.get <VM_NAME> | grep 'hardware_virtualization`
	The following attribute value confirms WSL2 is enabled:
	`hardware_virtualization: True`

VM Properties	Verification Procedure
GPU	1. Log in to Prism Central. 2. Navigate to **Compute & Storage** > **VMs**. The system displays the **Lists** tab by default. 3. Apply the **GPU TYPE** filter from the **Modify Filter** option. The system displays the VMs based on the selected GPU criteria. 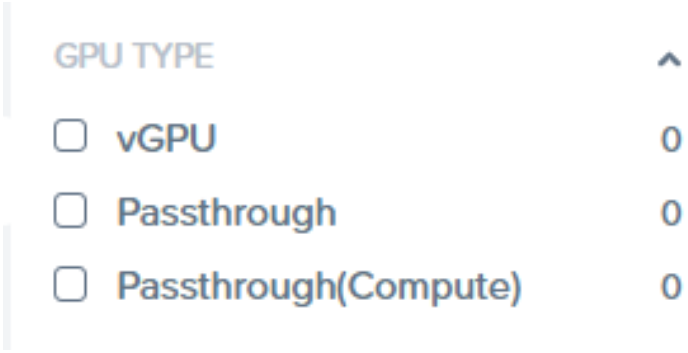

Figure 257: GPU TYPE Filter

Verifying Affinity Policy Association

This section describes how to verify the VM-Host affinity from Prism Element web console and affinity policies mapped to the VM from Prism Central.

Procedure

To verify affinity policy association for a VM, perform the following checks:

- Checking VM-Host Affinity From Prism Element on page 564.
- Checking Affinity Policies of a VM From Prism Central on page 565.

Checking VM-Host Affinity From Prism Element

About this task

Using Prism Element, you can define the VM-host affinity policy that 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.

Procedure

To verify the VM-Host affinity, perform the following steps:

1. Log in to Prism element web console.
2. Navigate to **VM** > **Table**, and double-click the target VM name. The system displays the **Update VM** window.

3. Under **VM Host Affinity**, check the hosts specified for the VM.

> **Note:**
>
> VMs with Host affinity policies can only be migrated to the hosts specified in the affinity policy during an HA event. If only one host is specified, the VMs cannot be migrated.

Checking Affinity Policies of a VM From Prism Central

About this task

Using Prism Central, you can define the affinity policy based on VM Categories and Host categories.

Procedure

To verify the affinity policy associated with a VM, perform the following steps:

1. Log in to Prism Central.
2. Navigate to **Compute & Storage** > **VMs**

   The system displays the **Lists** tab by default.
3. Click the target VM name.
   The system displays the details view of the target VM.
4. Click the **Categories** tab in the target VM's details view.
   The system displays the Categories and Affinity Policy associated with the VM.

Figure 258: Categories Tab

Defining Behaviour for Non-migratable VMs (GPU/CPU/PCI/Host Affinity configured VMs)

About this task

When an attempt is made for a host to enter maintenance mode, the VMs with GPU passthrough, CPU passthrough, PCI passthrough, and host affinity policies are not migrated to the other hosts in the cluster and therefore may block the attempt to enter maintenance mode.

To overcome this situation, Nutanix recommends you define either of the following behavior for these non-migratable VMs when host enters the maintenance mode:

- Automatically shutdown these non-migratable VMs.

- Set a block migration indication for these non-migratable VM.

Procedure

To define the above behaviors for non-migratable VMs, perform the following steps:

1. Log on to CVM with SSH.
2. Run the following command to put the host into maintenance mode with default action set for non-migratable VMs:

```
nutanix@cvm$ acli host.enter_maintenance_mode <Hypervisor-IP-address> [wait="{ true | false }" ]
 [non_migratable_vm_action="{ acpi_shutdown | block }"
```

   Replace:

   - `<Hypervisor-IP-address>` - with the actual host name at your site.
   - `wait`: Set the wait parameter to true to wait for the host evacuation attempt to finish.
   - `non_migratable_vm_action`: By default, this parameter is set to *block*, which means VMs with GPU, CPU passthrough, PCI passthrough, and host affinity policies are not migrated or shut down when you put a node into maintenance mode. If you want to automatically shut down such VMs for the duration of the maintenance mode, set this parameter to *acpi_shutdown*.

Managing a VM (Self Service)

About this task

After creating a VM (Creating a VM (Self Service) on page 545), you can use Prism Central to update the VM configuration, delete the VM, clone the VM, launch a console window, power on (or off) the VM, disable efficiency measurement, disable anomaly detection, set flash mode settings for a VM and VM disks, assign the VM to a protection policy, migrate the VM, add the VM to a recovery plan, run a playbook, manage categories, install and manage Nutanix Guest Tools (NGT), configure VM host affinity, add the VM image to the catalog, manage the VM ownership, configure QoS settings, or export VM as OVA.

> **Note:** You can perform only those operations for which you have permission.

You can perform these tasks by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VMs Summary View on page 94) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
- Go to the details page of a selected VM (see VM Details View on page 105) and select the desired action.

> **Note:** The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM and your permissions.

Procedure

- To modify the VM configuration, select **Update**.

  The **Edit VM** dialog box appears, which includes the same fields as the **Create VM** dialog box. Make the desired changes and then click the **Save** button.

  (For GPU-enabled AHV clusters only) You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through,

change the vGPU profile, add more vGPUs, and change the specified vGPU license. However, you need to power off the VM before you perform these operations.

- Before you add multiple vGPUs to the VM, see Multiple Virtual GPU Support and Restrictions for Multiple vGPU Support in the *AHV Admin Guide*.
- Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
- After you add the first vGPU, to add multiple vGPUs, see Adding Multiple vGPUs to the Same VM on page 573.

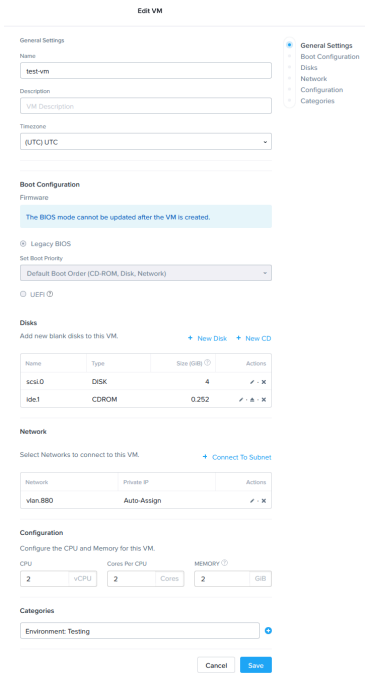

Figure 259: Edit VM Window

- **Disks:** You can add new disks to the VM using the **+ New Disk** option. You can also add new CDs to the VM using the **+ New CD** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM (Self Service) on page 545 before you modify the disk or create a new disk for a VM.
- **Network:** You can attach new network to the VM using the **Connect to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM (Self Service) on page 545 before you modify NIC network or create a new NIC for a VM.
- **Configuration:**You can update the CPU and memory configuration of the VM.
- **Categories:**You can update the VM category value, enter the category name and then click the plus sign (+) to the right of the field. See Creating a VM (Self Service) on page 545 before you modify or assign new categories for a VM.
- To delete the VM, select **Delete**. A window prompt appears; click the **OK** button to delete the VM.
- To clone the VM, select **Clone**.

  This displays the **Clone VM** dialog box, which includes the same fields as the **Create VM** dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking

the **Save** button. For example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot enabled.

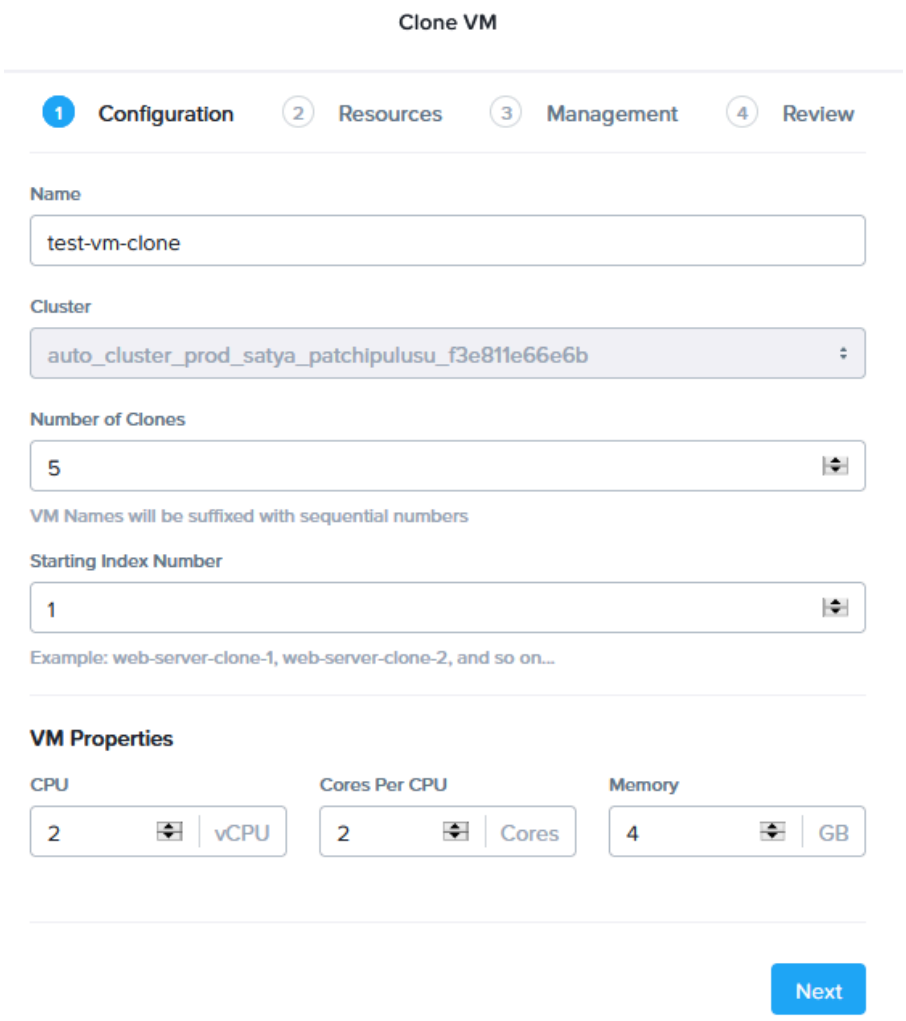

Figure 260: Clone VM Window (Configuration)

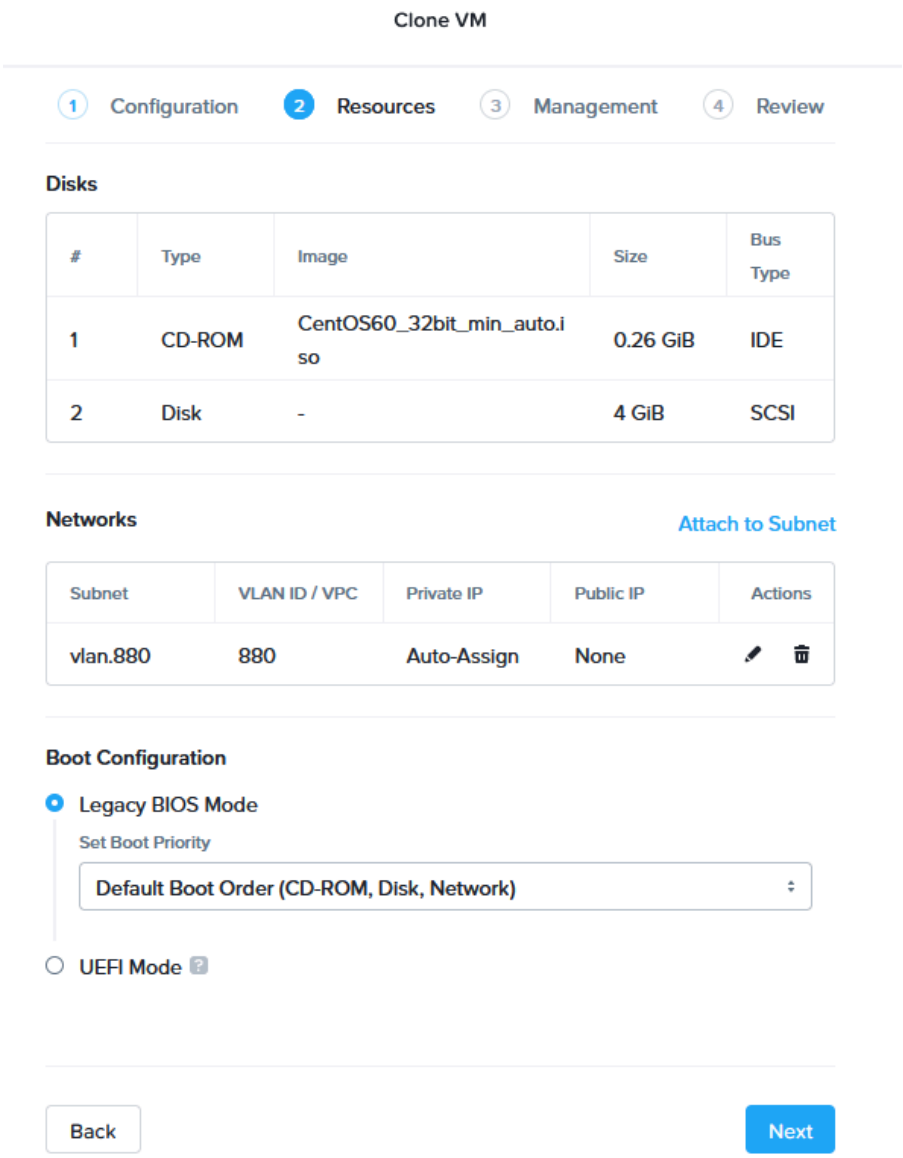

Figure 261: Clone VM Window (Resources)

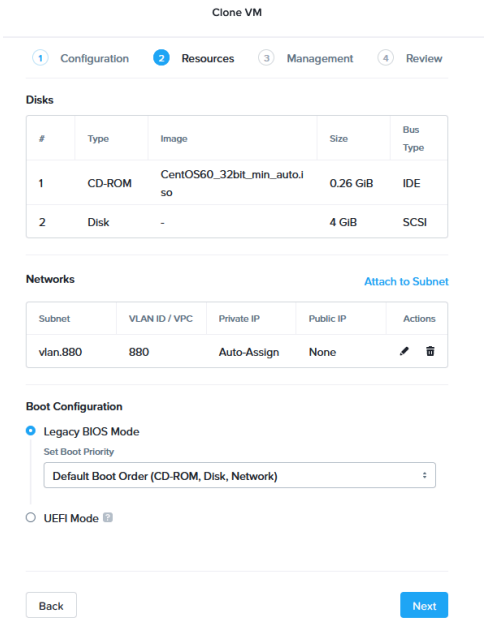

Figure 262: Clone VM Window (Management)

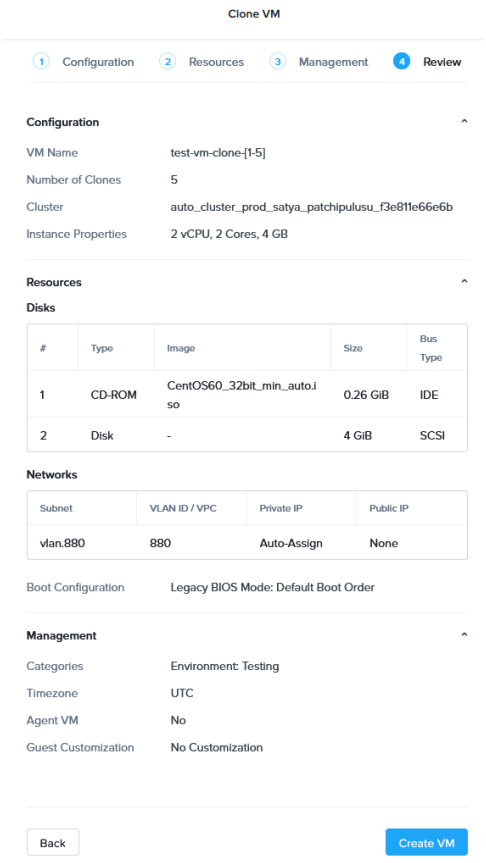

Figure 263: Clone VM Window (Review)

- To launch a console window, select **Launch Console**.

  This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

  > **Note:** A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

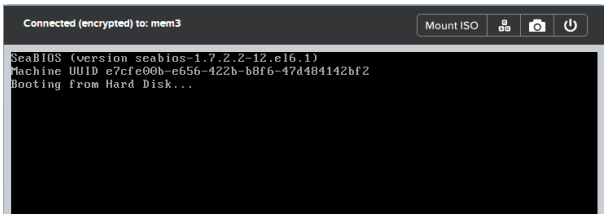

Figure 264: Console Window (VNC)

- To power on (or off) the VM, select **Power on** (or **Power off**).
- To disable (or enable) efficiency measurement for the VM, select **Disable Efficiency Measurement** (or **Enable Efficiency Measurement**).
- To disable (or enable) anomaly detection the VM, select **Disable Anomaly Detection** (or **Enable Anomaly Detection**).
- To create VM recovery point, select **Create Recovery Point**.

  This displays the **Create VM Recovery Point** dialog box. Enter a name of the recovery action for the VM. You can choose to create an **App Consistent** VM recovery point by enabling the check-box. The VM can be restored or replicated from a Recovery Point either locally or remotely in a state of a chosen recovery point.

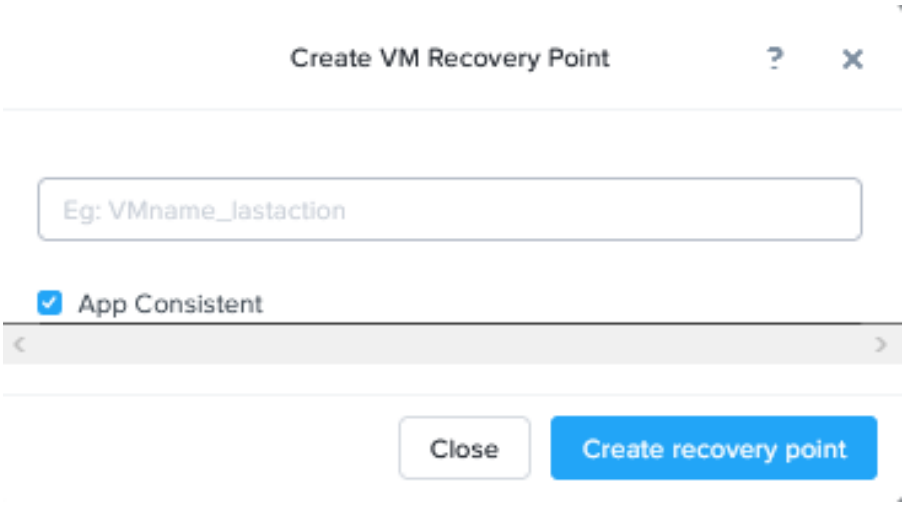

Figure 265: Create VM Recovery Point Window

- To add this VM to a recovery plan you created previously, select **Add to Recovery Plan** (see the Adding Guest VMs Individually to a Recovery Plan section in the *Nutanix Disaster Recovery Guide*).

  - **Adding Guest VMs Individually to a Recovery Plan**

- To run a playbook you created previously, select **Run Playbook** (see Running a Playbook (Manual Trigger) on page 932).
- To assign the VM a category value, select **Manage Categories**.

  This displays the **Manage VM Categories** page (see Assigning a Category on page 797).
- To install Nutanix Guest Tools (NGT), select **Install NGT** (see Installing NGT on page 594).
- To enable (or disable) NGT, select **Manage NGT Applications** (see Managing NGT Applications on page 596).

  The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

  > **Note:** If you clone a VM, by default NGT is not enabled on the cloned VM. You need to again enable and mount NGT on the cloned VM. If you want to enable NGT on multiple VMs simultaneously, see the Prism Web Console Guide.

  If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

  ```
  ncli> ngt mount vm-id=virtual_machine_id
  ```

  For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

  ```
  ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987
  ```

  > **Note:**
  >
  > - Self-service restore feature is not enabled by default on a VM. You need to manually enable the self-service restore feature.
  > - If you have created the NGT ISO CD-ROMs on AOS 4.6 or earlier releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to a later AOS version.
- To upgrade NGT, select **Upgrade NGT** (see Upgrading NGT on page 597).

- To specify a project and user who own this VM, select **Manage Ownership**.
  In the **Manage VM Ownership** window, do the following in the indicated fields:

  a. **Project**: Select the target project from the pull-down list.

  b. **User**: Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.

  c. Click the **Save** button.

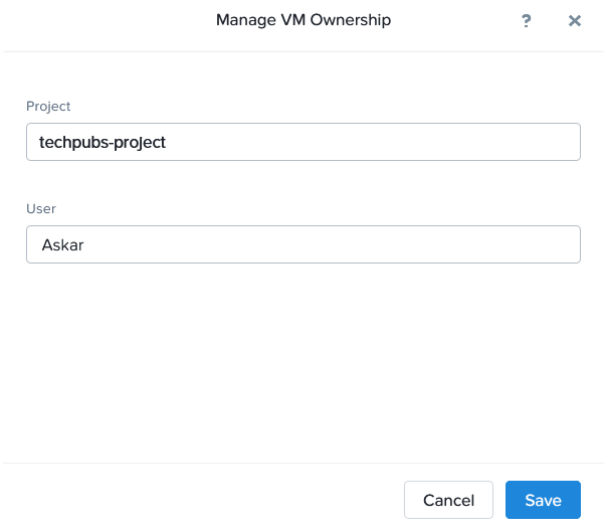

Figure 266: VM Ownership Window

- To export as OVA, select **Export as OVA** (see OVA Management on page 698 and Uploading an OVA on page 700).

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.

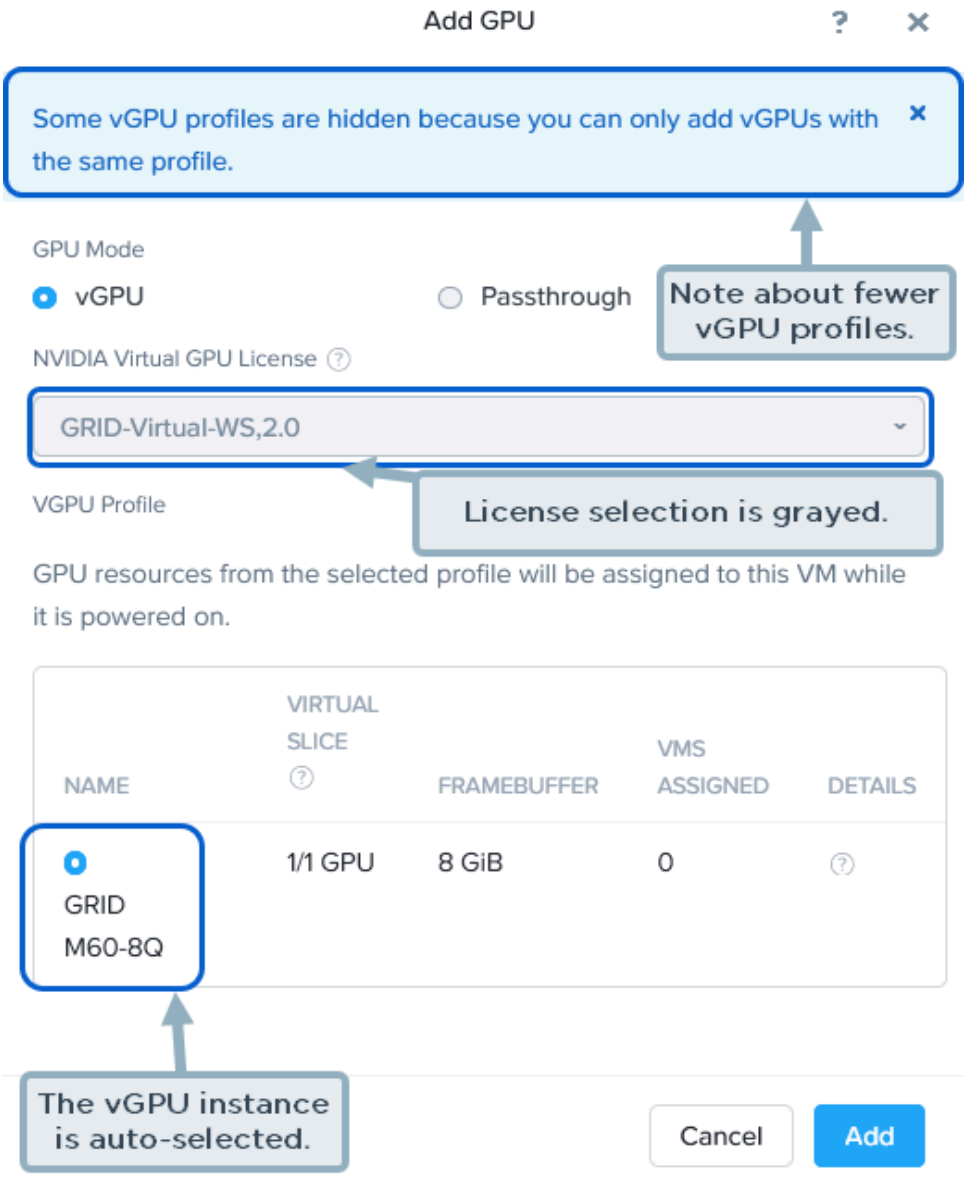

Figure 267: Add GPU for multiple vGPUs

3. In the main dialog box, you see the newly added vGPU.

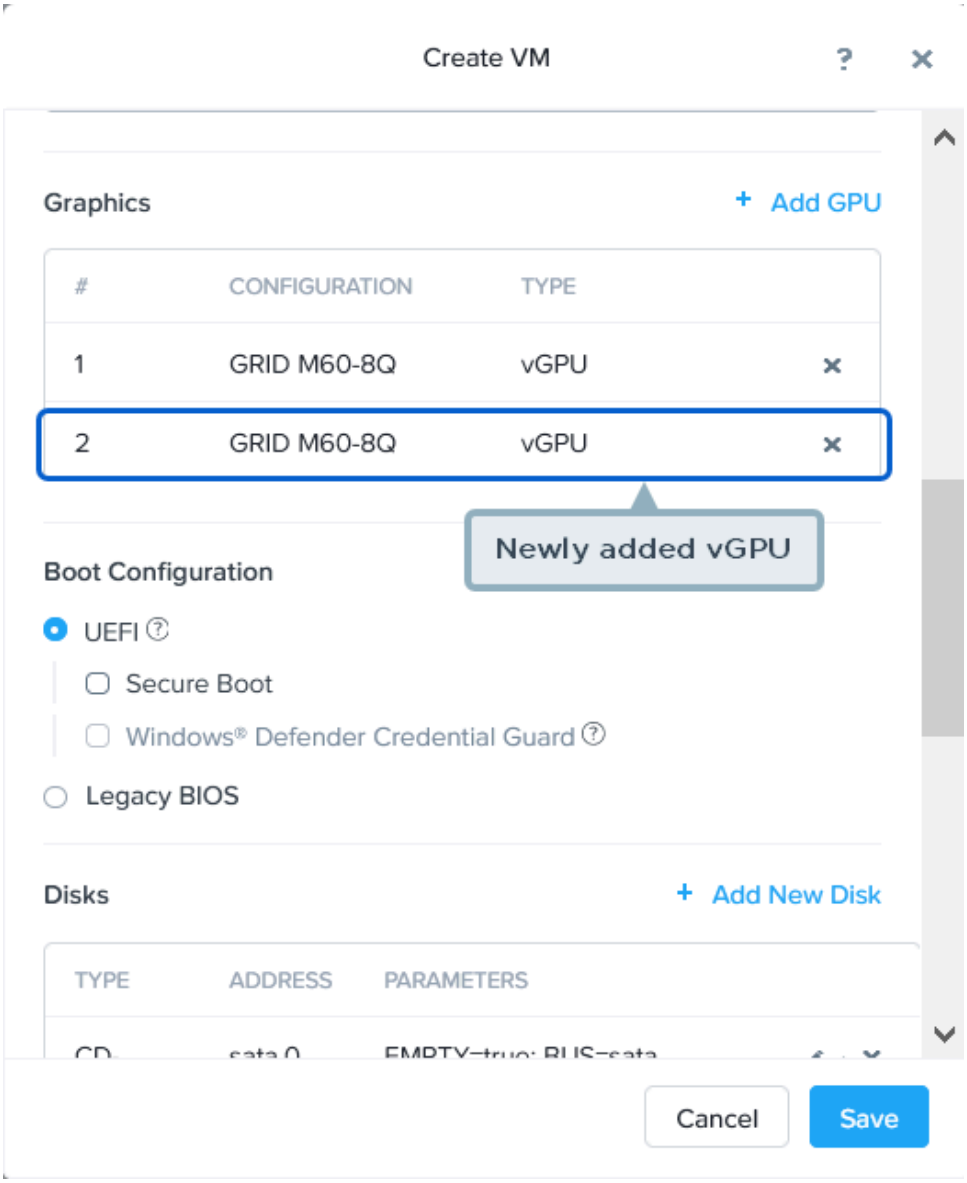

Figure 268: New vGPUs Added

4. Repeat the steps for each vGPU addition you want to make.

Live Migration of vGPU-enabled VMs

You can live-migrate a vGPU-enabled VM to another host within the cluster. You can also migrate the vGPU-enabled VM to a host in another cluster.

Note: Live migration of VMs with vGPUs is supported:

- For vGPUs created with minimum NVIDIA Virtual GPU software version 10.1 (440.53).
- Minimum AOS version that supports live migration within the same cluster is 5.18.1.
- Minimum AOS version that supports live migration across cluster is 6.1.

You can perform live migration of VMs enabled with virtual GPUs (vGPU-enabled VMs) only on best effort basis, if the destination node is equipped to provide enough resources to the vGPU- enabled VMs. However, if the destination node is not equipped with the enough resources, the vGPU-enabled VMs are shut down and you might experience a downtime.

In a successful migration case, the vGPUs can continue to run while the VMs that are running the vGPUs are seamlessly migrated in the background.

When you perform the LCM update, the vGPU-enabled VMs are listed as *Non-HA-protected VMs*. LCM also migrates the non-HA-protected VMs on best effort basis to the destination node if the following requirements are met:

- Destination node is equipped with the required resources for the VM.
- The VM GPU drivers are compatible with the AHV host GPU drivers.

If the destination node is not equipped with the enough resources or there is any compatibility issue between the VM GPU drivers and AHV host GPU drivers, the LCM forcibly shuts down the non-HA-protected VMs.

Prerequisites

Ensure the following prerequisites are met before you live migrate the vGPU-enabled VMs:

- The VM is not powered off.
- The host affinity is not set for the VM.

  If the host affinity of the VM is set to only one host, you cannot migrate the VM. However, based on the GPU resources required, if the host affinity is set to multiple hosts with the same or similar GPU resources, you can migrate the VM among the hosts with which the affinity is set.
- You have another host in the same cluster that supports the resources necessary to support the vGPU-enabled VM that you want to live migrate.

Limitations

For limitations applicable to live migration of vGPU-enabled VMs, see Limitations for Live Migration in *AHV Administration Guide*.

Migrating Within the Cluster

You can live-migrate a vGPU-enabled VM to another host within the cluster.

About this task

The **Migrate Within Cluster** option in the **Migrate VM** dialog box allows you to migrate vGPU-enabled VMs to other hosts within the clusters.

For limitations applicable to live migration of vGPU-enabled VMs, see Limitations for Live Migration in *AHV Administration Guide*.

To live migrate a vGPU-enabled VM to another host in the same cluster, do the following.

Procedure

1. Click **Compute & Storage** > **VMs** > **List**.

2. Select the VM you want to migrate live. Click **Actions** > **Migrate**

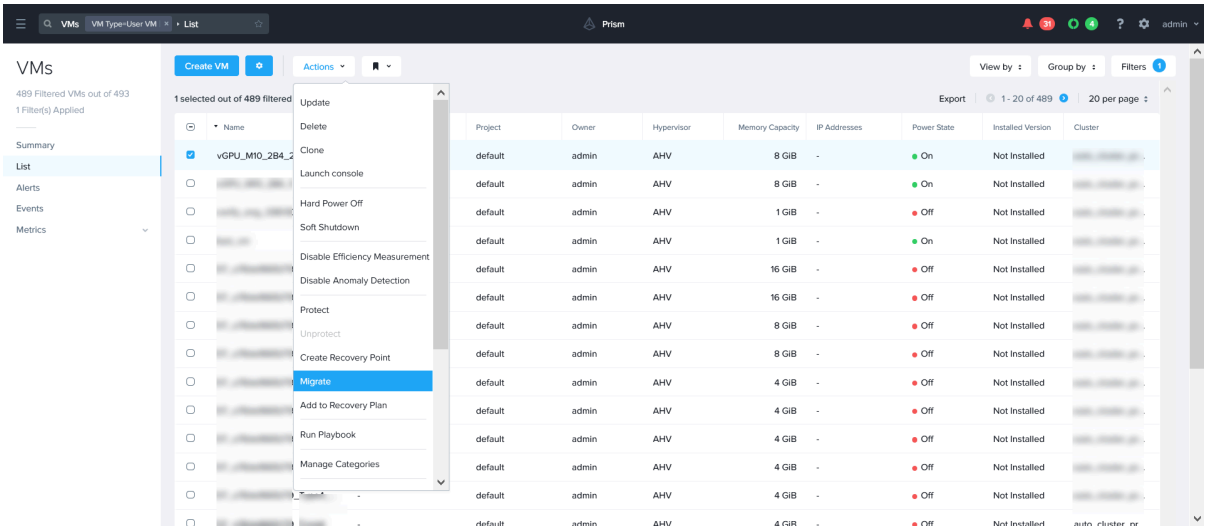

Figure 269: VM Selection

3. In the **Migrate VM** dialog box, the **Migrate Within Cluster** option is default selection.

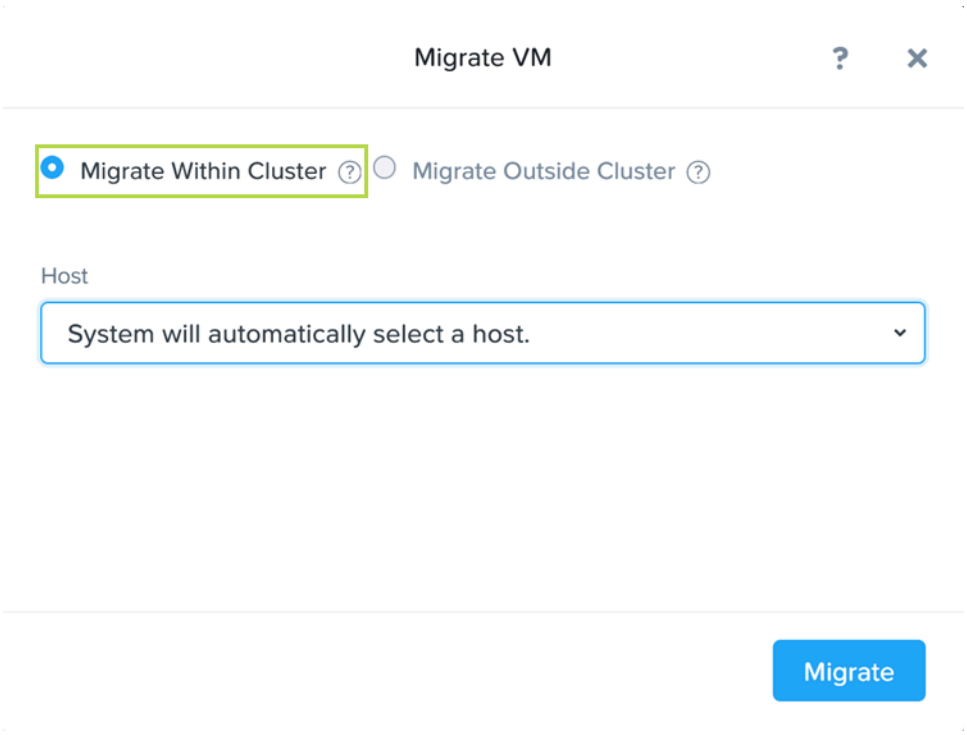

Figure 270: Migrate VM dialog box

4. In the **Host** drop-down list, do one of the following:

   - Retain the `System will automatically select a host` default option if you want to migrate the VM to a host selected by the system.

     The system selects a host based on the GPU resources available with the host as appropriate for the VM to be migrated live.

   - Select the host listed in the drop-down list that you want to migrate the VM to.

5. Click **Migrate**.

   Prism submits the task and displays the following message:

   ```
   Successfully submitted migrate operation.
   Task details
   ```

   Task details is a link to the **Tasks** page. Click the **Task details** link to monitor the migration task in the **Tasks** page.

   The host name of the VM in the **List** view changes to the host name to which you migrated the VM.

Migrating Outside the Cluster

You can migrate the vGPU-enabled VM to a host outside the cluster, that is, a host in another cluster.

About this task

The **Migrate Outside Cluster** option in the **Migrate VM** dialog box allows you to migrate vGPU-enabled VMs across clusters.

> **Note:** Live migration to a host in another cluster is supported only if the VM is protected by protection policy with Synchronous replication schedule.

For limitations applicable to live migration of vGPU-enabled VMs, see Limitations for Live Migration in *AHV Administration Guide*.

To migrate a vGPU-enabled VM to a host in another cluster, do the following.

Procedure

1. Click **Compute & Storage** > **VMs** > **List**.

2. Select the VM you want to migrate live. Click **Actions** > **Migrate**

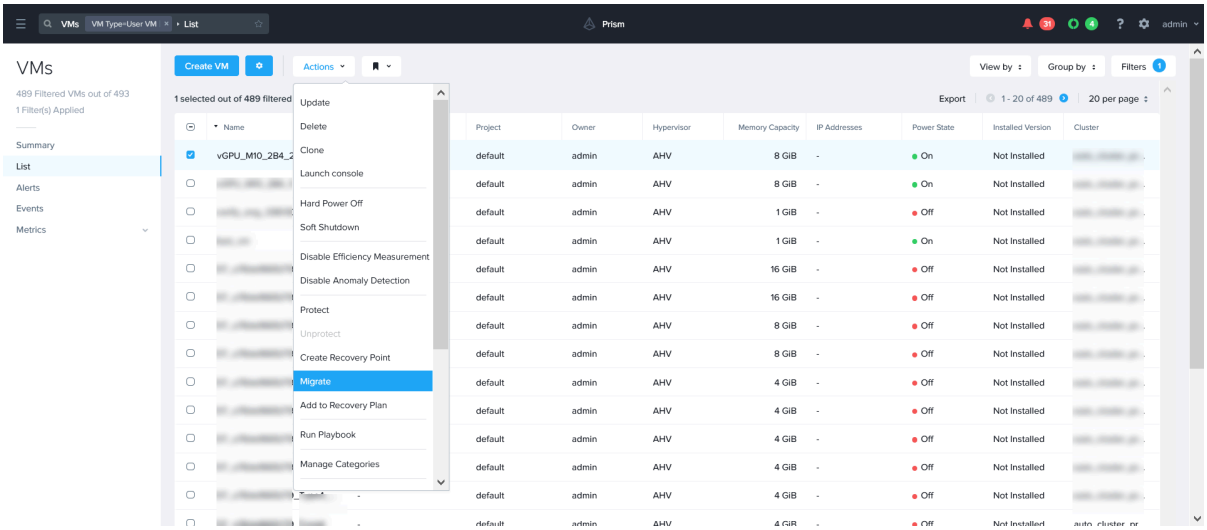

Figure 271: VM Selection

3. In the **Migrate VM** dialog box, select the **Migrate Outside Cluster** option.

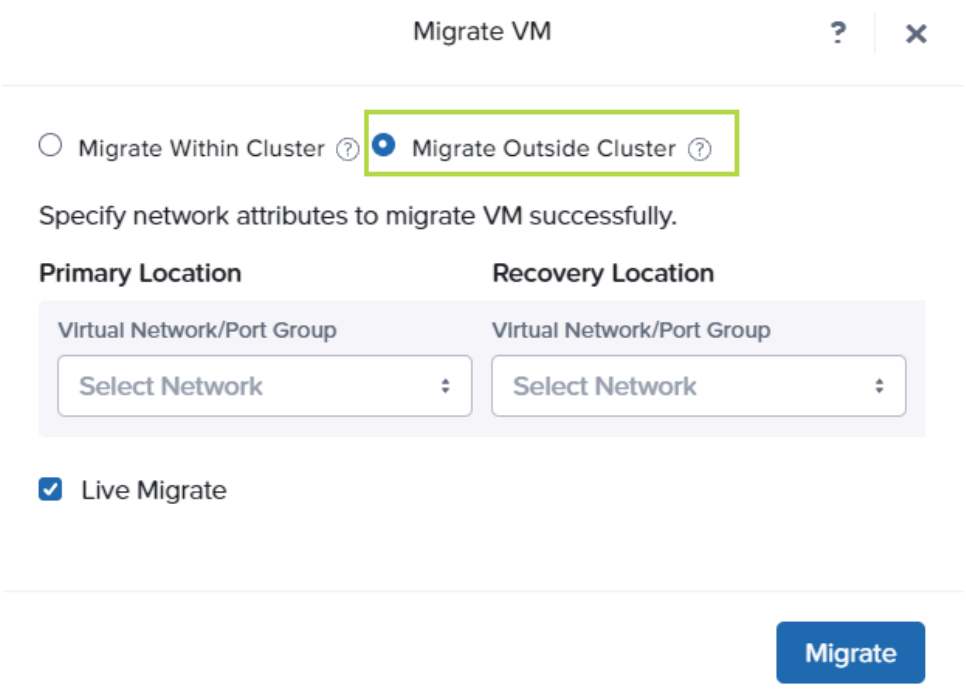

Figure 272: Migrate VM dialog box

4. In the **Select network attributes to migrate VM successfully** section, select the parameters for the following:

   > **Note:** If the VM is already configured in a recovery plan, then the network attribute fields are already populated and in-actionable. You cannot modify or add any values in these fields.

   - **Primary Location**: Virtual Network/Port Group—Select the network in the Primary AZ to be used for the live migration.
   - **Recovery Location**: Virtual Network/Port Group—Select the network in the Recovery AZ to be used for live the migration.

5. Check the **Live Migrate** box.

6. Click **Migrate**.

   Prism submits the task and displays the following message:

   ```
   Successfully submitted migrate operation.
   Task details
   ```

   Task details is a link to the **Tasks** page. Click the **Task details** link to monitor the migration task in the **Tasks** page.

   The host name of the VM in the **List** view changes to the host name to which you migrated the VM.

Creating a VM through Prism Central (ESXi)

In ESXi clusters, you can create a new virtual machine (VM) through Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration on page 718 before proceeding.
- Register the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central) on page 719.

About this task

To create a VM, do the following:

Procedure

1. Go to the **List** tab of the VMs dashboard (see VMs Summary View on page 94) and click the **Create VM** button.
   The **Create VM** wizard appears.

2. In the **Configuration** step, do the following in the indicated fields:

   a. **Name**: Enter a name for the VM.

   b. **Description** (optional): Enter a description for the VM.

   c. **Cluster**: Select the target cluster from the pull-down list on which you intend to create the VM.

   d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.

   e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.

   f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.

   g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.

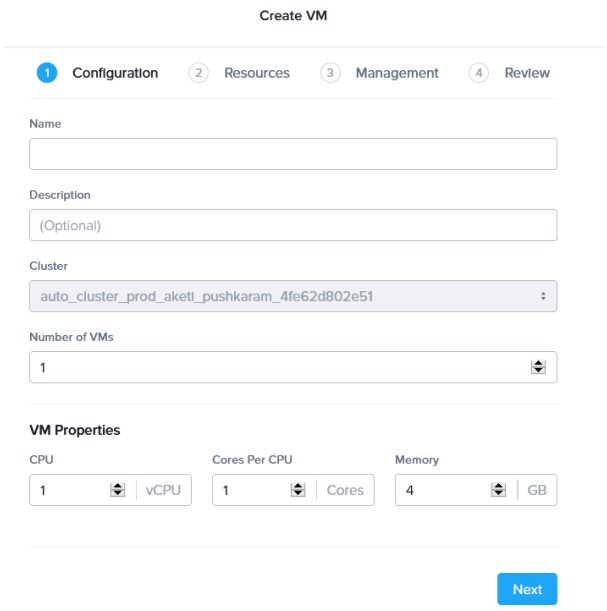

Figure 273: Create VM Window (Configuration)

3. In the **Resources** step, do the following.

   Disks: To attach a disk to the VM, click the **Attach Disk** button. The **Add Disks** dialog box appears. Do the following in the indicated fields:

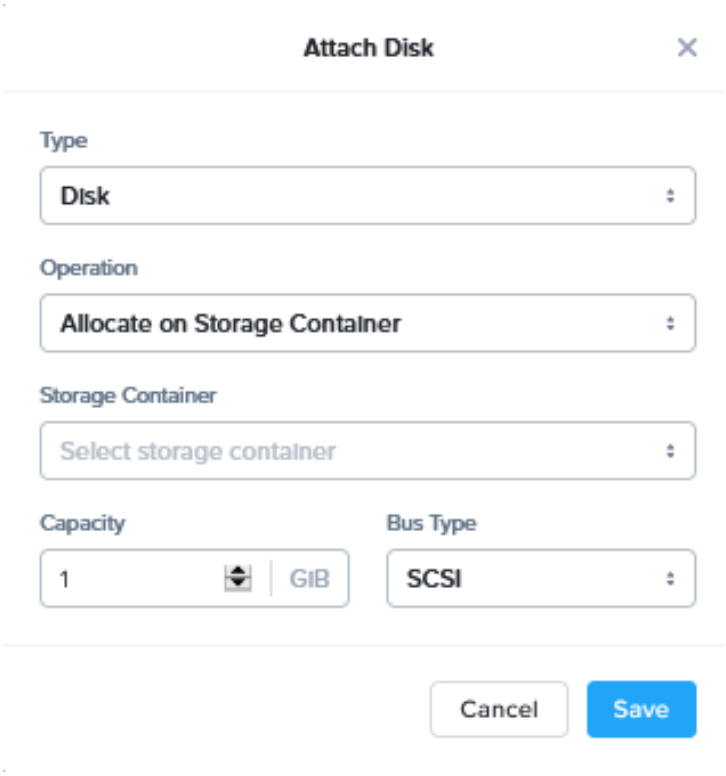

Figure 274: Add Disk Dialog Box

a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the dropdown list.

b. **Operation**: Specify the device contents from the dropdown list.

   - Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM. The **Empty CD-ROM** option is available only when *CD-ROM* is selected as the storage device in the **Type** field.
   - Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or

other source. The **Allocate on Storage Container** option is available only when *Disk* is selected as the storage device in the **Type** field.

c. If you select **Allocate on Storage Container** in the **Operation** field, the system prompts you to specify the **Storage Container**. Select the appropriate storage container.

d. **Bus Type**: Select the bus type from the dropdown list.

   The options displayed in the dropdown list varies based on the storage device **Type** selected in the **Type** field.

   If the storage device **Type** is:

   - **Disk** - The available choices are **SCSI**, **SATA**, **PCI**, or **SATA**.
   - **CD-ROM** - The available choices are **IDE**, or **SATA**

e. **Path**: Enter the path to the desired system image.

f. **Capacity**: Enter the disk size in GiB.

g. When all the field entries are correct, click **Save** to attach the disk to the VM and return to the **Create VM** window.

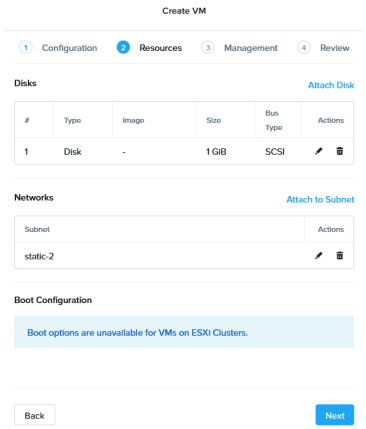

Figure 275: Create VM Window (Resources)

4. In the **Resources** step, do the following.

   Networks: To create a network interface for the VM, click the **Attach to Subnet** button. The **Attach to Subnet** dialog box appears.

   Do the following in the indicated fields:

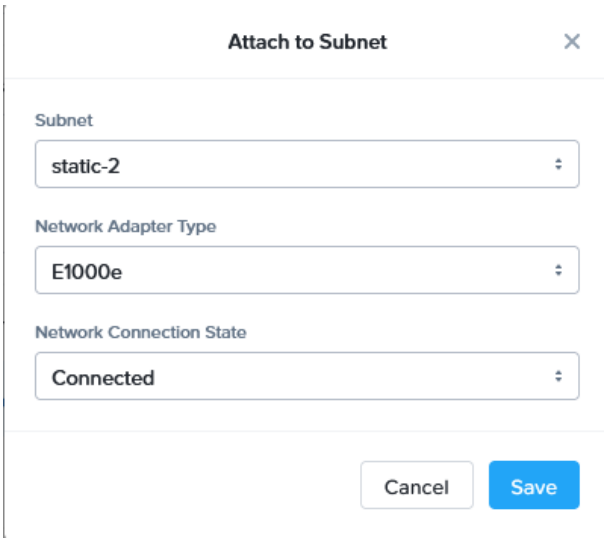

Figure 276: Attach to Subnet Dialog Box

   a. **Subnet**: Select the target virtual LAN from the pull-down list.

      The list includes all defined networks (see Configuring Network Connections on page 784).

   b. **Network Adapter Type**: Select the network adapter type from the pull-down list.

      See vCenter Server Integration on page 718 for the list of supported adapter types.

   c. **Network Connection State**: Select the state for the network that you want it to operate in after VM creation. The options are `Connected` or `Disconnected`.

   d. When all the field entries are correct, click the **Add** button to create a network interface for the VM and return to the **Create VM** dialog box.

   e. Repeat this step to create more network interfaces for the VM.

5. In the **Management** step, do the following.

   a. **Categories**: Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

   b. **Guest OS**: Type and select the guest operating system.

      The guest operating system that you select affects the supported devices and number of virtual CPUs available for the virtual machine. The Create VM wizard does not install the guest operating system. See the list of supported operating systems in vCenter Server Integration on page 718 topic.

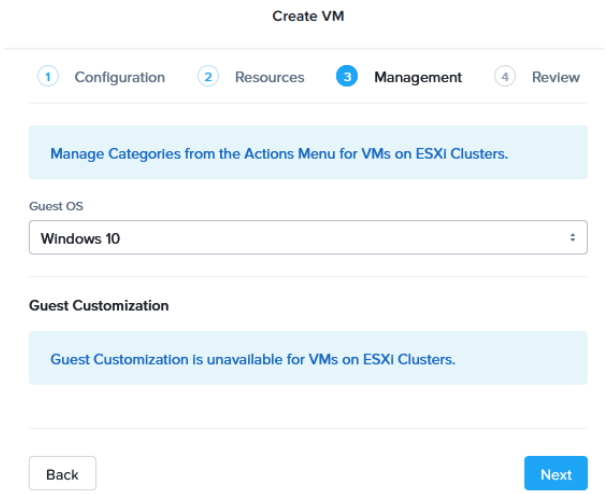

Figure 277: Create VM Window (Management)

6. In the **Review** step, when all the field entries are correct, click the **Create VM** button to create the VM and close the **Create VM** dialog box.

   The new VM appears in the VMs entity page list.

Managing a VM through Prism Central (ESXi)

You can manage virtual machines (VMs) in an ESXi cluster through Prism Central.

Before you begin

- See the requirements and limitations section in the vCenter Server Integration before proceeding.
- Ensure that you have registered the vCenter Server with your cluster. For more information, see Registering vCenter Server (Prism Central).

About this task

After creating a VM (see Creating a VM through Prism Central (ESXi) on page 580), you can use Prism Central to perform operations on the VM.

You can perform these operations by using any of the following methods:

- Select the target VM in the **List** tab of the VMs dashboard (see VM Summary View) and choose the required action from the **Actions** menu.
- Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.

- Go to the details page of a selected VM (see VM Details View) and select the desired action.

The available actions appear in bold; the unavailable actions are grayed out. You can perform some of the available actions on a single VM at a time, while others can be performed on multiple VMs simultaneously. The available actions depend on the current state of the VM and your permissions.

> **Note:** You can perform only those operations for which you have permissions from the admin.

Procedure

- Select the target VM and choose the action that you want to perform. See the following table for the actions and their descriptions.

Table 735: VM Actions

Action	Description	Applicable to Multiple VMs
Update	Update the VM configuration. For more information, (see Updating a VM through Prism Central (ESXi) on page 588.	No
Delete	Delete the VM.	Yes
Clone	Clone the VM. For more information, see Cloning a VM through Prism Central (ESXi) on page 589.	No
Create VM Template	Create a VM template. For more information, see Creating a VM Template on page 614.	No
Launch Console	Launch the VM console. For more information, see Launching a VM Console through Prism Central (ESXi) on page 590.	No
Power On	Power on the VM.	Yes
Power Off	Power off the VM.	Yes
Power Cycle	power off the VM followed by a power on.	Yes
Reset	Perform an ACPI reset action through the BIOS of the VM.	No
Guest Shutdown	Perform a graceful shutdown of the operating system of the VM.	No
Guest Reboot	Perform a graceful restart of the operating system of the VM.	No
Pause/Suspend	Pause or suspend the VM. This option is available only when the VM is powered on.	Yes
Resume	Resume a paused or suspended VM.	No
Disable Efficiency Measurement	Disable efficiency measurement for the VM. For more information, see Operations Policy Management on page 873.	Yes
Enable Efficiency Measurement	Enable efficiency measurement for the VM. For more information, see Operations Policy Management on page 873.	Yes
Disable Anomaly Detection	Disable Anomaly Detection for the VM. For more information, see Operations Policy Management on page 873.	Yes
Enable Anomaly Detection	Enable Anomaly Detection for the VM. For more information, see Operations Policy Management on page 873.	Yes

Action	Description	Applicable to Multiple VMs
Protect	Assign the VM to a protection policy. Selecting **Protect** opens a dialog box to specify the protection policy to which this VM should be assigned. You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Nutanix Disaster Recovery and configuring the Availability Zone.	Yes
Unprotect	Remove the VM from a protection policy.	Yes
Create Recovery Point	Create VM recovery point. The VM can be restored or replicated from a recovery point either locally or remotely in a state of a chosen recovery point. In the **Create VM Recovery Point** dialog box, enter a name for the VM recovery point. You can choose to create an app consistent VM recovery point by enabling the **App Consistent** check-box. For more information about app consistent recovery point, see Terminology section in the *Data Protection and Recovery with Prism Element* guide.	No
Migrate	Migrate the VM to another host. In the **Migrate VM** dialog box, select the target host from the pull-down list (or select the System will automatically select a host option to let the system choose the host) and then click the Migrate button to start the migration. **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.	No
Add to Recovery Plan	Add the VM to a recovery plan you created previously. For more information, see the Adding Guest VMs Individually to a Recovery Plan section in the *Nutanix Disaster Recovery Guide*.	Yes
Run Playbook	Run a playbook you created previously. For more information, see Running a Playbook (Manual Trigger) on page 932.	Yes
Manage Categories	Assign the VM a category value. For more information, see Assigning a Category on page 797.	Yes
Install NGT	Install Nutanix Guest Tools (NGT). For more information, see Installing NGT on page 594.	Yes
Manage NGT Applications	Enable or disable NGT. For more information, see Managing NGT Applications on page 596.	Yes
Upgrade NGT	Upgrade NGT. For more information, see Upgrading NGT on page 597.	Yes
Configure VM Host Affinity	Establish VM host affinity. In the **Configure VM Host Affinity** dialog box, select (click the icon for) the hosts and then click the **Save** button. This creates an affinity between the VM and the selected hosts. For more information, see Affinity Policies Defined in Prism Central on page 608. It is recommended to create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure.	No

Action	Description	Applicable to Multiple VMs
Set QoS Attributes	Configure quality of service (QoS) settings. For more information, see Setting QoS for an Individual VM on page 602.	Yes
Export as OVA	Export the VM as OVA. For more information, see OVA Management on page 698 and Uploading an OVA on page 700.	No

•

Updating a VM through Prism Central (ESXi)

You can use Prism Central to update the VM configuration.

Procedure

1. Follow any one of the navigation paths mentioned in Managing a VM through Prism Central (ESXi) on page 585 for selecting a VM to perform an operation.
2. To modify the VM configuration, select **Update**.

   The **Update VM** dialog box appears, which includes the same fields as the **Create VM** dialog box. Make the desired changes and then click the **Save** button in the **Review** step.

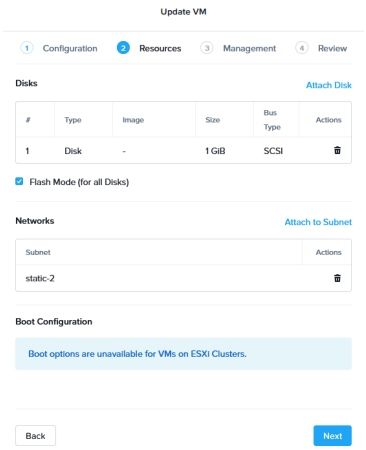

Figure 278: Update VM Window (Resources)

> **Caution:** See the recommendations specified in KB-9832 before making changes to the VM properties such as memory, vCPU, or storage. Any values other than the recommended memory, vCPU, or storage values for a VM might cause discrepancies in the Prism Central behavior.

3. **Disks:** You can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column. See Creating a VM through Prism Central (ESXi) on page 580 before you create a new disk for a VM. You can enable or disable the flash mode settings for the VM. To enable flash mode on the VM, click the **Enable Flash Mode** check box. After you enable this feature on the VM, the status is updated in the VM table view.
4. **Networks:** You can attach new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet attached to the VM. See Creating a VM through Prism Central (ESXi) on page 580 before you modify NIC network or create a new NIC for a VM.

Cloning a VM through Prism Central (ESXi)

You can use Prism Central to clone a VM.

Procedure

1. Follow any one of the navigation paths mentioned in Managing a VM through Prism Central (ESXi) on page 585 for selecting a VM to perform an operation.
2. To clone the VM, select **Clone**.

   This displays the **Clone VM** dialog box, which includes the same fields as the **Create VM** dialog box. A cloned VM inherits most the configurations (except the name) of the source VM. Enter a name for the clone and then click the **Save** button to create the clone. You can optionally override some of the configurations before clicking the **Save** button. For example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

   Note:

   - You can clone up to 250 VMs at a time.
   - You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

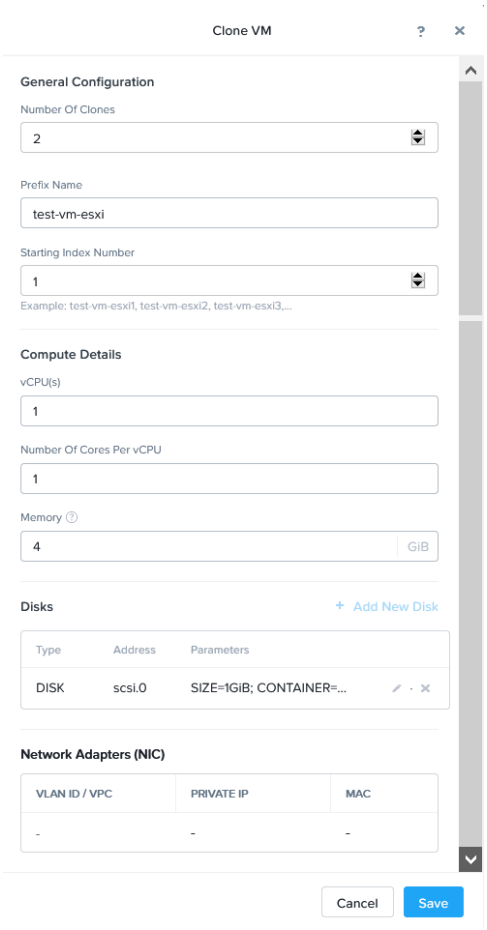

Figure 279: Clone VM Window

Launching a VM Console through Prism Central (ESXi)

You can use Prism Central to launch a VM console.

Procedure

1. Follow any one of the navigation paths mentioned in Managing a VM through Prism Central (ESXi) on page 585 for selecting a VM to perform an operation.
2. To launch a console window, select **Launch Console**.

   This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

   Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

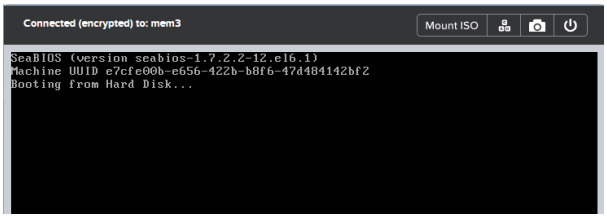

Figure 280: Console Window (VNC)

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 (VM) to enable advanced VM management functionalities provided by Nutanix.

Nutanix Guest Tools Overview

The NGT software package for Linux and Windows includes the following components:

- **NGT Installer**: Allows you to install NGT in a guest VM.
- **Nutanix Guest Agent (NGA) Service**: Maintains a communication channel between the Controller VM (CVM) and the guest VMs.
- **Nutanix VirtIO Package**: Includes the drivers required to run the VM on the AHV hypervisor. For more information about VirtIO drivers, see Nutanix VirtIO for Windows in *AHV Administration Guide*.
- **Python**: The NGA service is written in Python and is shipped with a dedicated installation of Python so it does not interact with the other components in the system.
- **Nutanix VSS package**: Enables Nutanix native in-guest Volume Snapshot Service (VSS) agent to take application-consistent snapshots for all the VMs that support VSS.

In addition, the NGT software package for Windows also includes the following components:

- A custom module that performs checks and fulfills all the prerequisites.
- Microsoft Visual C++ 2008/2012 x64 redistributable component
- Integration with Windows Performance Monitor for hypervisor statistics.

Installing NGT in a guest VM allows you to use the following advanced VM management features:

- **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. 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.

- **Application-consistent snapshots for Windows VMs**

  Enables Nutanix native in-guest Volume Snapshot Service (VSS) agent 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 about Nutanix VSS-based snapshots for the Windows VMs, see the *Application-consistent Snapshots with Microsoft Volume Shadow Copy Service (VSS)* section in the Conditions for Application-consistent Snapshots topic of the *Data Protection and Recovery with Prism Element* guide.

- **Application-consistent snapshots for Linux VMs**

  Supports application-consistent snapshots for Linux VMs by running specific scripts on VM quiesce. For more information, see the Conditions for Application-consistent Snapshots topic of the *Data Protection and Recovery with Prism Element* guide.

- **Static IP address preservation support after failover for Nutanix Disaster Recovery (DR)**

  This feature allows the preservation of the IP address of a guest VM (with a static IP address) for its failover (DR) to an IPAM network. For more information, see the *Networking Requirements* section in the Nutanix Disaster Recovery Requirements topic of the *Nutanix Disaster Recovery Guide*.

- **In-guest scripts execution support for Nutanix Disaster Recovery (DR)**

  In-guest scripts automate various task executions upon recovery of the VMs. For example, you can automate the task of changing the domain controller or performing the reassigning of the DNS IP address and reconnection to the database server at the recovery AZ. For more information about the tasks that can be automated using In-guest scripts, see Creating a Recovery Plan in the *Nutanix Disaster Recovery Guide*.

Nutanix Guest Tools Requirements

Ensure that the following requirements are met so that you can successfully install NGT, and use all the NGT features.

General Requirements

- For information about the supported OS for specific NGT features, see the NGT section in *Compatibility and Interoperability Matrix*.
- A 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. For more information, see Virtual IP Address Impact in the *Prism Web Console Guide*.
- Network access from the virtual IP address of the cluster to the guest VMs where Prism Central-based management of NGT is required.
- Network access from the guest VM to the virtual IP address of the cluster where IP-based connectivity is used for NGT.
- VMs must have at least one empty IDE or SATA CD-ROM to attach the ISO required for the configuration and installation of NGT.

- The following ports must be accessible to use the specific NGT functionality:
  - If IP-based communication is enabled, the TCP port 2074 of the cluster virtual IP address must be accessible from the guest VMs so that the guest VMs can communicate with the CVM. For more information, see Nutanix Guest Agent and Controller VM Communication.
  - The TCP port 23578 in the guest VM must be accessible if you want to use the VSS service.

  For the complete list of required ports, see Port Reference.

Windows Requirements

In addition to the General Requirements, ensure that the following requirements are met so that you can successfully install NGT in a Windows VM:

- To install NGT, you need a local user account in the VM with administrative privileges. Domain accounts are currently not supported.
- For Windows Server Edition VMs, ensure that Microsoft Volume Shadow Copy Service (VSS) is enabled before you install NGT with the VSS application enabled. For more information about how to verify if Microsoft VSS is enabled on your server, see Microsoft documentation.

  Note:
  - Before you enable the VSS feature, upgrade the NGT version on the Windows Server Edition VMs to 2.3.2 or later. If not, the VSS snapshots will be marked as crash-consistent.
  - The VSS feature is not supported on Windows OS VMs and Linux VMs.
- During NGT installation a few PowerShell commands are executed. To run these commands, the PowerShell executable must be present in the PATH system variable of the environment. Otherwise, the installation of NGT might fail. For more information, see KB-7284.
- The PowerShell version on the Windows VMs must be 3.0 or later.
- Antivirus scanners might need exclusions added for the C:\Program Files\Nutanix directory for successful NGT installation.
- Windows Remote Manager Service (winrm) must be running and configured to allow SSL-based connections. The following commands are one of the many methods to configure winrm with SSL. These commands must be executed from an administrative Windows PowerShell console.

```
> $certificate = New-SelfSignedCertificate -DnsName $env:computername -
CertStoreLocation cert:\LocalMachine\My
```

```
> winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$env:computername`";CertificateThumbprint=`"$($certificate.ThumbPrint)`"}"
```

```
> cmd /c 'winrm set winrm/config/service/auth @{Basic="true"}'
```

```
> netsh advfirewall firewall add rule name=\"WinRM-HTTPS\" dir=in localport=5986
 protocol=TCP action=allow
```

  Note: Windows Remote Manager Service (winrm) is required to install NGT from Prism Central only. It is not required when you upgrade NGT or install NGT manually by logging in to the guest VM.

Linux Requirements

In addition to the General Requirements, ensure that the following requirements are met so that you can successfully install NGT in a Linux VM:

- Ensure that the VM has Python 2.6 or a later version installed.
- You must have a guest VM user account with password-based SSH, and root or passwordless SUDO access enabled.
- Ensure that the user has write permissions for the /usr/local directory as NGT is installed at this location.
- Ensure that the user has write permissions for the /mnt/nutanix/ngt directory to mount the NGT ISO at this location.
- Ensure that a minimum available disk space of 550 MB is available in the /mnt/nutanix/dir directory for copying the NGT installer when upgrading NGT.
- Ensure that the user has write access to /tmp directory for the duration of the installation.

General Notes

- NGT includes VirtIO drivers when you install it on a Windows VM. If these drivers are already installed in your guest VM, they might be replaced during the installation of NGT, depending on the version that is bundled with NGT. If the drivers were downgraded, you might need to manually reinstall VirtIO and restart the VM if prompted. For more information about VirtIO drivers, see Nutanix VirtIO for Windows in *AHV Administration Guide*. For information about the compatible VirtIO driver versions, see the *AOS - NGT - VM Mobility Matrix* section in the *NGT* tab of the Compatibility and Interoperability Matrix.
- VirtIO driver installation updates the NIC driver in the guest VM. During the update, the guest VM might lose network connectivity. The network connectivity is restored after the update is complete.
- If a VM is connected to a volume group (VG), NGT captures the iSCSI Qualified Name (IQN) of the VM and stores the information. If you change the IQN of the VM and take a snapshot of the VM before the NGT refresh cycle (currently at 5 minutes) occurs, NGT does not provide the auto restore capability because the snapshot operation cannot capture the VM-VG connection.

  Workaround:

  - **Linux VM**: Restart the Nutanix Guest Agent (NGA) service by running the $ sudo service ngt_guest_agent restart command.
  - **Windows VM**: In the command prompt, run the net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent" command.
- By default, NGT client certificates expire every 1,000 days. Regenerate the certificates before they expire. For more information, see Regenerating NGT Certificates for Guest VMs in the *Prism Element Web Console Guide*.

NGT Management in Prism Central

You can use Prism Central to perform the following operations on a guest VM or multiple guest VMs:

- Enable NGT and mount the NGT installer
- Install NGT
- Upgrade NGT
- Manage NGT applications

In addition, NGT installation through Prism Central provides the following capabilities:

- Enable Nutanix Volume Shadow Copy Service (VSS) and Self Service Restore (SSR) applications as part of the install workflow
- Select multiple VMs in Prism Central from the VM Entity browser and upgrade NGT on these VMs

- Defer restarting of the VMs to a specified later time and date after installing or upgrading NGT
- Define policies that allow you to defer restarting of the VMs to a specified later time and date

Installing NGT

The installation of NGT using Prism Central is fully automated. You do not need to manually log in to the VM to install NGT. For large-scale deployments, Nutanix recommends using Prism Central for installing NGT.

Before you begin

Ensure that all the requirements specified in Nutanix Guest Tools Requirements are met.

About this task

To install NGT, perform the following steps on a guest VM or multiple guest VMs at the same time in Prism Central:

> **Note:**
>
> - The guest VM must be powered on to install NGT in the VM.
> - You cannot install NGT on VMs created on storage containers with replication factor 1.

Procedure

1. Log in to Prism Central.
2. Navigate to **Compute & Storage** > **VMs** from the Navigation Bar.
3. From the VMs dashboard, click the **List** tab.

   > **Tip:** The **NGT** status column associated with a VM displays whether the NGT is installed in the VM, whether it is the latest version, and whether an upgrade is available or not.

4. Select the checkbox associated with the VM in which you want to install NGT.

5. Click **Actions**, and from the dropdown menu that appears click **Install NGT**.

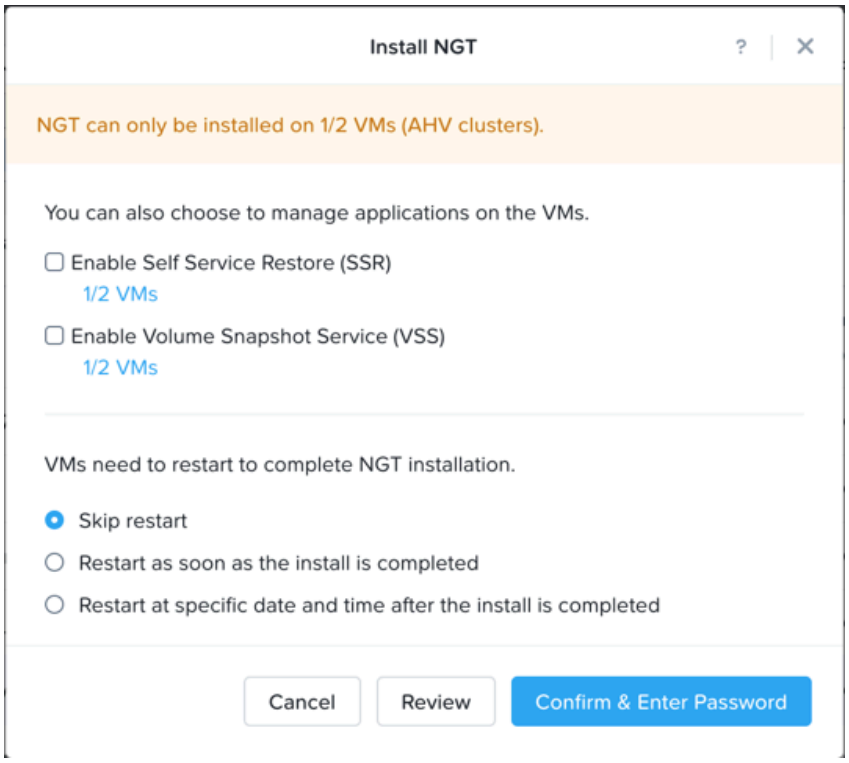

Figure 281: Installing NGT

> **Note:** If you select multiple VMs, this window displays the number of VMs on which you can install NGT. For example, if you select a VM on which NGT is already installed, it is skipped during this operation.

6. (Optional) To enable SSR and VSS applications, select the **Enable Self Service Restore (SSR)** and **Enable Volume Snapshot Service (VSS)** checkboxes.

   If you choose not to enable these features during the installation of NGT, you can do it later by following the instructions in Managing NGT Applications.

7. Select the restart schedule of the VM by doing one of the following:

   - **Skip restart**: The VM does not restart after the installation completes.
   - **Restart as soon as the install is completed**: The VM restarts after the installation completes.
   - **Restart at specific date and time after the install is completed**: The VM restarts at the specific date and time that you configure in the **Date** and **Time** fields.

   > **Note:** Nutanix recommends that you restart the VM after the NGT installation completes.

8. (Optional) Click **Review** to view your configurations.

9. Click **Confirm & Enter Password**.

10. In the **Username** and **Password** fields, enter a username and password of an account that meets the requirements specified in Nutanix Guest Tools Requirements, and is accessible to all the VMs on which NGT is being installed.

    > **Note:** If you choose to install NGT on multiple VMs and the VMs do not have the same username and password, you can add them to a JSON file and paste the file in the box provided. Click **Download sample json** to download a sample of the JSON file.

11. Click **Done** to start the installation process.

    When the installation completes successfully on a VM, the **NGT** status column of the VM in the **List** tab displays the status as **Latest**. It takes a few minutes for the status to reflect in Prism Central because this is an asynchronous operation.

    > **Note:** If the NGT requirements are not met or you experience any issue with the installation process, the NGT installation fails and an alert message displays. Ensure that all the requirements are met and perform the installation process again to install NGT in the VMs.

12. (Optional) If you do not want to install NGT automatically or cannot provide the username and password of an account that meets the requirements, click **Skip and Mount** to complete the installation manually.

    Prism Central enables the NGT feature in the VM, mounts the NGT installer, and attaches an ISO to the virtual CD drive with the volume label NUTANIX_TOOLS to the selected VMs. You can then install NGT by logging in to the VM. For more information, see Installing NGT in a Windows VM or Installing NGT in a Linux VM in the *Prism Element Web Console Guide*.

    > **Note:** If NGT is already installed in the VM, it detects the NUTANIX_TOOLS CD and reconfigures itself. Once NGT is installed and the reconfiguration is successful, the system automatically unmounts NUTANIX_TOOLS. For more information, see CD-ROM Eject Functionality of NGT.

Enabling NGT and Mounting the NGT Installer on Cloned VMs

If you have cloned a VM or multiple VMs from a VM that had NGT installed, perform the following steps to re-enable and configure NGT on the cloned VMs.

About this task

Perform the following steps to enable NGT and mount the NGT installer on cloned VMs.

> **Note:** After you perform the following steps, you do not need to separately install NGT on the cloned VMs.

Procedure

1. Navigate to the **Install NGT** window by following the instructions mentioned in Steps 1 through 5 of Installing NGT.
2. Enable NGT and mount the NGT installer on the cloned VMs by following the instructions in Step 12 of Installing NGT (choosing **Skip and Mount**).

   NGT automatically identifies the ISO, update the configuration, and unmount the CD.

Managing NGT Applications

Prism Central allows you to enable or disable the Nutanix Volume Shadow Copy Service (VSS) and Self-Service Restore (SSR) applications even after NGT is installed in a guest VM.

About this task

Perform the following steps to enable the SSR and VSS applications if you have not enabled them during the NGT installation workflow. You can also disable these applications (if already enabled) after NGT is installed using this procedure.

Procedure

1. Log in to Prism Central.
2. Navigate to **Compute & Storage** > **VMs** from the Navigation Bar.
3. From the VMs dashboard, select the **List** tab.
4. Select the checkbox associated with the VM on which you want to enable or disable the SSR and VSS applications.
5. Click **Actions**, and from the dropdown menu that appears click **Manage NGT Applications**.
6. From the **Manage Applications** window, enable or disable applications by selecting the corresponding option.

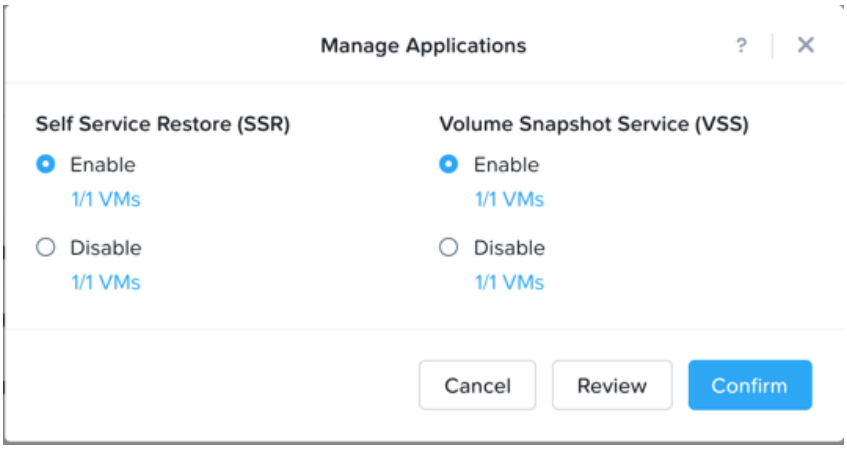

Figure 282: Manage NGT Applications

7. Click **Confirm** to save the changes.

   When you enable an application, the **Services Enabled** field in the **Properties** widget of the **Summary** tab displays the services enabled for the VM. For example, if you enable the SSR application the **Services Enabled** field displays **SSR**.

Upgrading NGT

Prism Central allows you to upgrade NGT on the guest VMs.

Before you begin

- Ensure that all the requirements specified in Nutanix Guest Tools Requirements are met.
- Unless you upgrade AOS, you cannot upgrade NGT.
- Upgrading NGT using Prism Central is supported on VMs that have NGT version 1.2.3 or later. If the VM has an earlier version of NGT, upgrade NGT using the procedure in the Upgrading NGT topic in the *Prism Element Web Console Guide*.

About this task

Perform the following steps to upgrade NGT on a guest VM.

> **Note:**
>
> - You can upgrade NGT on a maximum of 60 VMs at the same time.
> - For Linux VMs, upgrading a cluster running any NGT version prior to NGT 4.0 to NGT 4.0 or later versions, might fail. For more information, see KB-15396.

Procedure

1. Log in to Prism Central.
2. Navigate to **Compute & Storage** > **VMs** from the Navigation Bar.
3. From the VMs dashboard that is displayed, select the **List** tab.
4. Select the checkbox associated with the VMs for which you want to upgrade NGT.
5. Click **Actions**, and from the dropdown menu that appears click **Upgrade NGT**.

6. From the **Upgrade VMs** window, select the restart schedule of the VMs by doing one of the following:

   - **Skip restart**: The VM does not restart after the installation completes.
   - **Restart as soon as the install is completed**: The VM restarts after the installation completes.
   - **Restart at specific date and time after the install is completed**: The VM restarts at the specific date and time that you configure in the **Date** and **Time** fields.

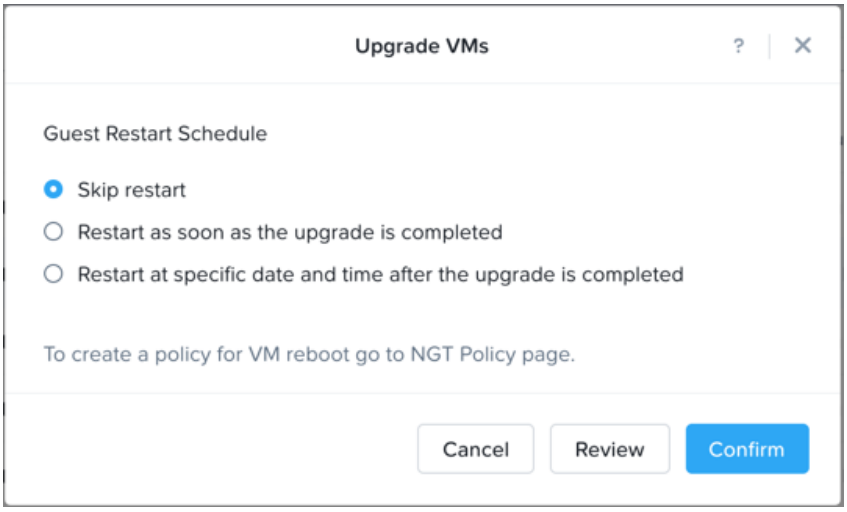

Figure 283: Upgrading NGT

> **Note:**
>
> - If you select multiple VMs to upgrade, the Upgrade VMs window displays the number of VMs on which you can upgrade NGT. For example, if you select a VM on which the latest version of NGT is already installed, that VM is skipped during this operation.
> - Nutanix recommends that you restart the VM after upgrading NGT.

7. (Optional) Click **Review** to view your configurations.
8. Click **Confirm** to start the upgrade process.

   When the upgrade is complete, the **NGT** status column associated with the VMs displays the status as **Latest**. It takes a few minutes for the status to reflect in Prism Central because this is an asynchronous operation.

> **Note:** If the Nutanix Guest Tools Requirements are not met or if you experience any issue with the upgrade process, the NGT upgrade fails and an alert message displays. Ensure that all the prerequisites are met and perform the upgrade process again to upgrade NGT.

Uninstalling NGT

Prism Central does not support automatic uninstallation of NGT.

You must log in to a guest VM to uninstall NGT in that VM. For more information, see Uninstalling and Removing Nutanix Guest Tools in the *Prism Element Web Console Guide*.

Note: If the NGT Status in Prism Central displays **Latest** after uninstalling NGT, follow the instructions mentioned in Step 3 of Uninstalling and Removing Nutanix Guest Tools in the *Prism Element Web Console Guide* to remove the NGT information from Prism Central.

Nutanix Guest Agent and Controller VM Communication

You can set up connectivity between the controller virtual machine (CVM) and the guest VMs using the Nutanix Guest Tools (NGT) service that runs on the CVM and the Nutanix Guest Agent (NGA) service that runs on the guest VMs in the Nutanix cluster. After you install NGT in a guest VM, the NGA service in the guest VM starts periodic communication with the CVM over SSL connections.

Each Nutanix cluster is configured as a certificate authority (CA). When NGT is enabled in a guest VM, a certificate pair is generated specifically for that 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 inside the guest VM initiates an SSL connection to port 2074 of the cluster virtual IP address to communicate with the CVM. Any firewall must be configured to allow the guest VM to reach port 2074 on the cluster virtual IP address. For the complete list of required ports, see Port Reference.

NGT provides three levels of security for VM communication.

- SSL certificates that ensure a secure TCP connection between the CVM and the guest VM
- Capability-based authorization that ensures only NGT features that are supported by the cluster are enabled on the guest VM
- NGT ensures that the communication is established only if the system or BIOS UUID of the guest VM and the UUID of the guest VM (provided by the hypervisor) stored in the CVM are the same.

Communication is successful only if all the three conditions are met. For example, if a guest VM that has NGT installed is cloned, the new VM cannot communicate with the CVM. You must separately enable NGT, which includes mounting the NGT installation media on the cloned VM. The installation media includes configuration specific to the VM, and when the VM is re-configured and is able to communicate with the CVM, it automatically unmounts the installation media. You do not need to separately install NGT again on the cloned VMs.

NGA also publishes information about the guest VM to the CVM. For example, guest OS type, status of VM mobility, and VSS services.

Automatic CD-ROM Ejection

After NGT is installed or re-configured in a guest VM, the CD-ROM is auto-ejected from the VM in the following stages.

1. CD-ROM is automatically ejected from the VM immediately after the installation of NGT is successfully completed.
2. NGA indicates to the controller VM (CVM) that the installation is complete and the CD-ROM is detached from the VM through the hypervisor. This operation depends on the communication between the guest VM and the CVM and may take up to 10 minutes.

Note:

- These operations can occur at different times. It is possible that the CD-ROM shows up as empty inside a VM, but the status in the NGT-Controller VM service Tools Mounted might display as true. If this situation occurs, unmount the CD-ROM using the Prism Element web console.
- If the NGT software version in the ISO is more recent than the installed version in the guest VM, the CD-ROM does not eject automatically. This functionality enables the upgrade of the NGT software inside the VM.

NGT Usage in Disaster Recovery

If NGT is enabled on a guest VM that is protected by Nutanix Disaster Recovery, the snapshots of the VM include NGT relevant information along with its capabilities.

The NGT information is preserved for a restored guest VM in the following scenarios:

- Migrating the VM to the remote cluster (Planned Failover)
- Performing an in-place restore of the VM on the source cluster
- Creating a clone of the VM on the source cluster

In these scenarios, a new NGT ISO image containing only the relevant configuration information (SSL certificates, Controller VM IP address, ans so on) is created for the recovered guest 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 NGT ISO CD-ROM automatically.

The NGT information is not preserved for a restored guest VM in the following scenarios:

- Performing an out-of-place restore of the VM on the remote cluster.
- Retrieving and restoring the snapshot from the remote cluster to the source cluster.

In these scenarios, you must enable NGT again in the restored VM.

Storage Quality of Service (QoS)

Storage QoS provides administrators granular control to manage the performance of virtual machines and ensure that the system delivers consistent performance for all workloads. You can use a controllable knob to limit the IOPS that the storage layer would serve for individual virtual machines. IOPS is the number of requests the storage layer can serve in a second. You can set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources.

Note:

- Storage QoS is applicable only if you have an AOS Pro or above license. See the License Manager Guide for more information on AOS licenses.
- Setting storage QoS is not allowed while the AOS upgrade is in progress.

Terminology

Throttled IOPS

The maximum IOPS that the storage layer tries to admit over a sustained period. Applications that are running on the VM are not permitted to exceed this level. This prevents VMs from affecting a system beyond the set limits. However, if the system is in a state of flux, your application might not be throttled. To prevent application crashing due to misconfiguration, the throttled IOPS setting must not be less than one hundred. Otherwise, an error is raised.

Setting maximum IOPS on a VM may affect the cluster latency for that VM and may result in higher latency being displayed in the cluster latency charts. Performance of any other VM is not affected by this setting. You can view individual VM latency graph to ensure that the other VMs are not affected. You can also create a custom focus view to display the throttled IOPS for all the VMs.

Throughput

Throughput is the amount of data that an application sees from the storage layer underneath per second. The maximum throughput is calculated by multiplying the number of IOPS with the block size.

Calculation of the IOPS and Throughput

The default block size used for IOPS calculation is 32 KiB. If the block size is less than or equal to 32 KiB, then the system is rate limited by the number of IOPS and the bandwidth is calculated by multiplying the number of IOPS with the block size of the application. Otherwise, if the block size is greater than 32 KiB, the system is rate limited by the throughput and the number of IOPS is calculated by dividing the throughput by block size of the application.

For example, if you have set a throttle limit of 800 IOPS on a VM and the application is admitting IOPS with 8 KiB block size, then the application would see 800 IOPS and throughput of 6.25 MBps. However, if application is admitting IOPS with 64 KiB block size, then the application would see 400 IOPS and a throughput of 25 MBps. For more information about the block size, IOPS, and throughput relationship when throttle limit it set to 800 IOPS, see *Relationship between Block Size, IOPS, and Throughput* table.

Table 736: Relationship between Block Size, IOPS, and Throughput

Block size	IOPS	Throughput
8 KiB	800	6.25 MBps
16 KiB	800	12.50 MBps
32 KiB	800	25 MBps
64 KiB	400	25 MBps
128 KiB	200	25 MBps

Limitations

Storage QoS has the following limitations:

- Storage QoS is not supported for an in-place restore, out-of-place restore, and snapshot operation on VM.
- While creating a clone of a VM, any QoS attributes throttle limit set on the original VM might not be applied to the new cloned VM. QoS throttle limit has to be set on the new cloned VM separately.
- For linked clones, when you are sharing the vDisks across multiple VMs, shared disks are not a part of QoS policy and shared disks would not be throttled. Each VM might exceed the maximum limit depending on the shared vDisk usage. For example, during bootstorm in a VDI setup.
- If a VM has volume groups attached, then QoS is not applicable.
- Storage QoS is not supported for Metro Availability, synchronous replication, and Nutanix Files virtual machines.

Setting QoS for an Individual VM

About this task

Storage QoS allows you to set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources. Perform the following procedure to set the QoS for an individual VM.

Figure 284: Setting QoS for an Individual VM

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275779955001 on Brightcove.

Procedure

1. Go to the **List** tab of the VM dashboard (see VMs Summary View on page 94) and select the target VM.
2. Do one of the following:
   - » From the **Actions** drop-down list, select **Set QoS Attributes**.
   - » Click the VM and then select **Set QoS Attributes** from the **More** drop-down list.
3. Under **Input Type**, select one of the following:
   - » IOPS (For more information about IOPS, see *Throttled IOPS* in *Storage Quality of Service (QoS)*.)
   - » Throughput (For more information about throughput, see *Throughput* in *Storage Quality of Service (QoS)*.)
4. If you selected **IOPS**, enter a value for the throttled IOPS in the **Throttled IOPS**.
5. If you selected **Throughput**, enter a value for the throttled throughput in the **Throttled Throughput**.

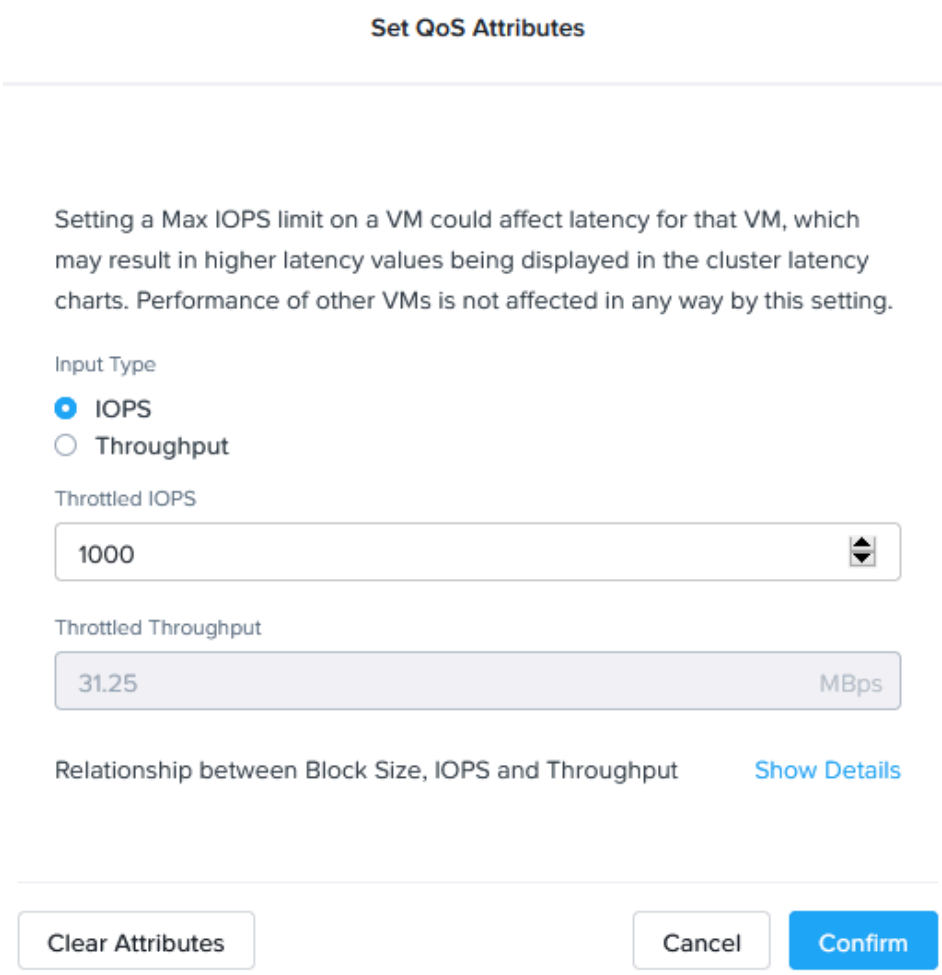

Figure 285: Setting QoS Attributes

6. Click **Show Details** to populate a table that describes the relationship between block size, IOPS, and throughput.
7. Click **Clear Attributes** to delete all the values.

8. To save details, click **Confirm**.
   This step stores the number of IOPS or throughput that would be served by a storage layer. To edit the storage QoS setting, follow the previous process and enter a new value for IOPS or throughput.

Memory Overcommit Management

Enabling Memory Overcommit While Creating a VM

This procedure helps you enable memory overcommit in a VM that you create in Prism Central. You need a minimum Prism Central version of pc.2022.4 for this procedure.

Before you begin

You cannot enable the memory overcommit feature using the Prism Element Web Console. Enable memory overcommit on a VM that you create by using Prism Central. If you create a VM using Prism Web Console, then you can enable memory overcommit on that VM using Prism Central. For the procedure to enable memory overcommit on a VM that you have already created, see Enabling Memory Overcommit for Existing VMs on page 606.

About this task

When you are creating a VM, you can enable memory overcommit for it. Follow the procedure to create a VM.

Procedure

1. On the **VMs** dashboard, click **Create VM**.
   The **Create VM** page opens.

2. On the **Configuration** tab, provide the other details necessary to create the VM. Select the **Enable Memory Overcommit** checkbox.

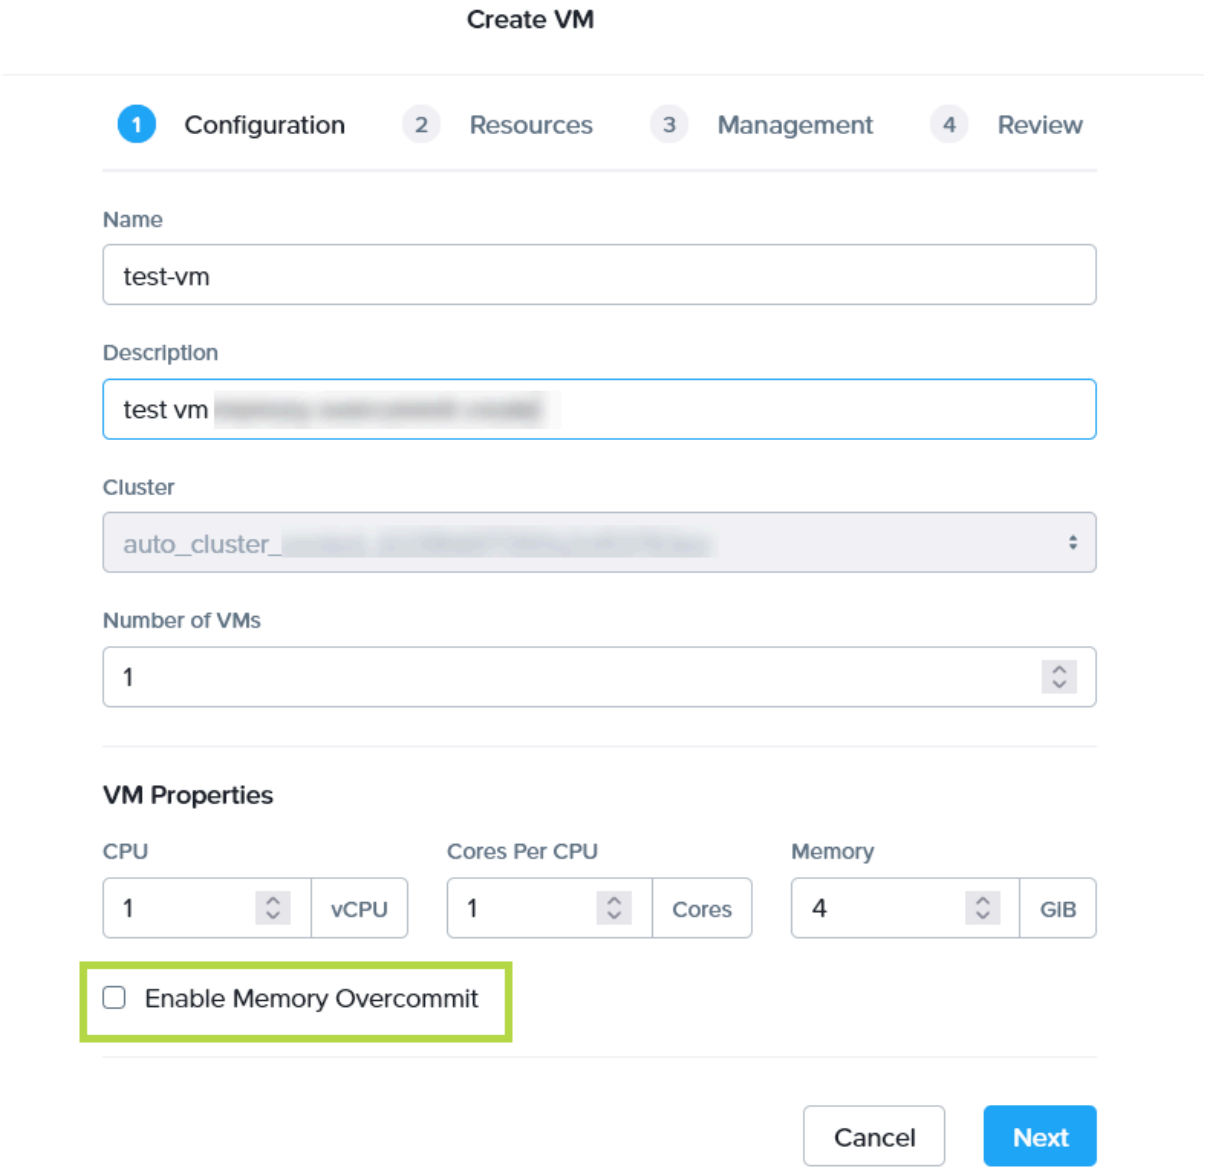

Figure 286: Create VM - Enable Memory Overcommit

3. Click **Next**.
4. Provide other details necessary for the VM creation in the **Resources** and **Management** tabs, and click **Next**.

   For more information, see Creating a VM through Prism Central (AHV).
5. In the **Review** tab, under **Configuration** ensure that the **Memory Overcommit** configuration is displayed as **Enabled**.

   If **Memory Overcommit** configuration is not displayed as **Enabled**, click **Edit** to go back to the **Configuration** tab and select the **Enable Memory Overcommit** checkbox.

6. Click **Save**.

What to do next

The **Tasks** page displays the VM creation task. After it is successfully completed, check the **VMs** dashboard to verify that the VM is created and the value in the **Memory Overcommit** column displays **Enabled**.

> **Note:** The default **General** list view does not provide the **Memory Overcommit** column. Create your own customized view and add the **Memory Overcommit** column to that view. You can also add other columns to your customized view.

Click the newly created VM to open the VM details page. In the **Summary** tab, the **Properties** widget displays **Enabled** for the **Memory Overcommit** property.

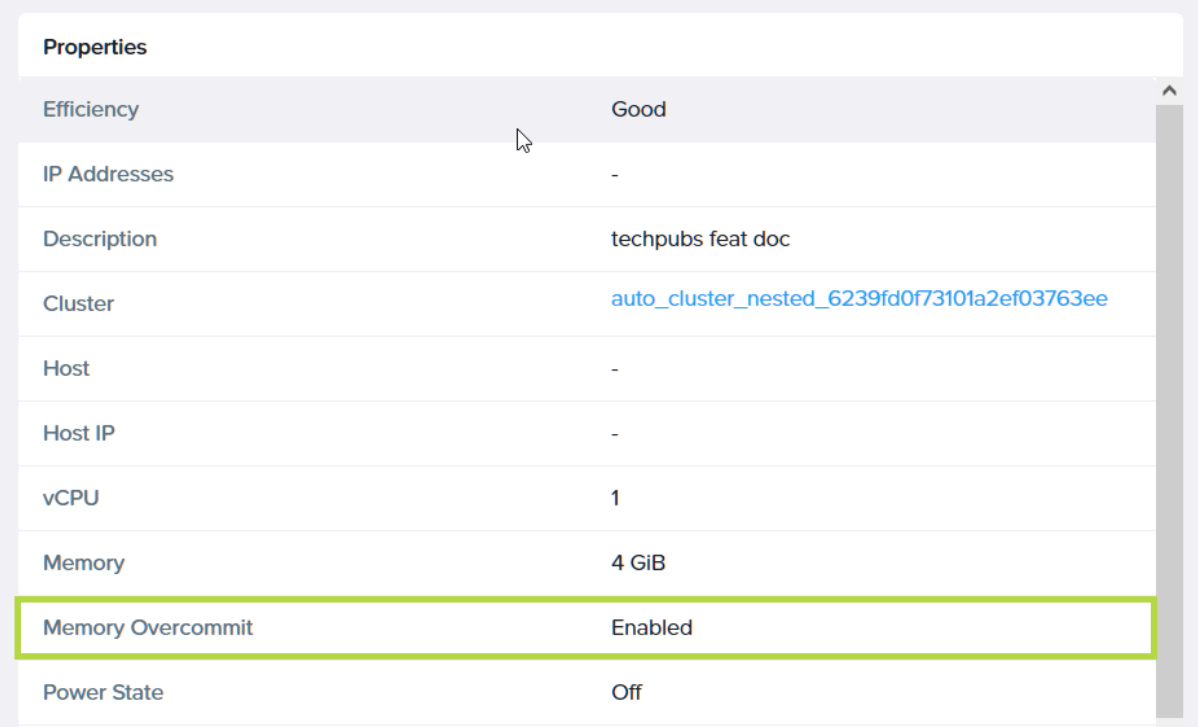

Figure 287: VM Details Page - Memory Overcommit Enabled

Enabling Memory Overcommit for Existing VMs

The procedures in this section enable you to enable memory overcommit in one or more existing VMs. You need a minimum Prism Central version of pc.2022.4 for this procedure.

Before you begin

Ensure that the VM on which you want to enable memory overcommit is powered off. If the VM is powered on or is in **Soft Shutdown** state, the **Enable Memory Overcommit** action is not available in the **Actions** dropdown list.

About this task

To enable memory overcommit in one or more existing VM by updating it, follow this procedure.

Procedure

1. On the **VMs** dashboard (**List** tab), select the VM or VMs for which you want to enable memory overcommit.

   You can select one VM and enable memory overcommit. You could also select multiple VMs and enable memory overcommit in all the selected VMs as a bulk update.

2. In the **Actions** dropdown list, click the **Enable Memory Overcommit** action.

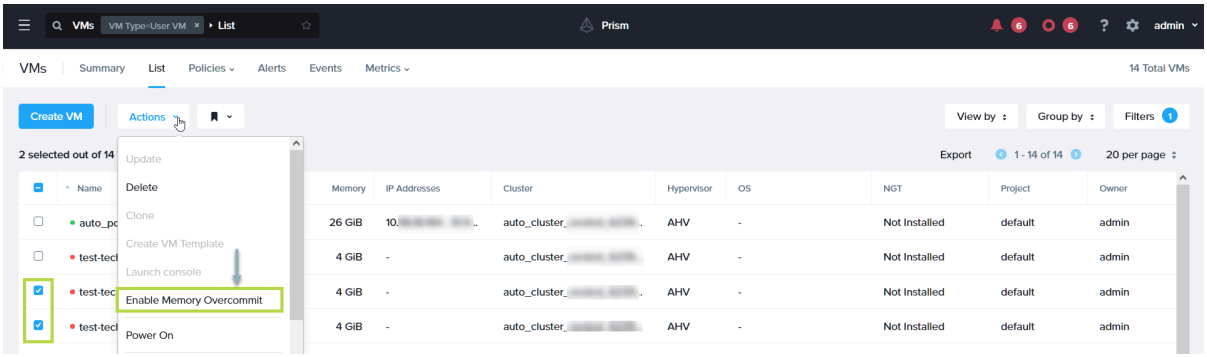

Figure 288: VM(s) Selection to Enable Memory Overcommit

Note: For an individual VM, you can also select the VM, click **Actions** > **Update** and in the **Configuration** tab of the **Update VM** page, select the **Enable Memory Overcommit** checkbox.

Click **Next** on the **Configuration**, **Resources** and the **Management** tabs, and **Save** on the **Review** tab.

What to do next

You can check the update tasks in the **Tasks** page. If you selected multiple VMs to enable memory overcommit on, the **Task** page displays the update for each VM as a separate **Update VM** task.

You could verify the status as **Enabled** in the individual VM details page of each VM. You could also verify the status in the **List** tab with customized view having the **Memory Overcommit** column.

Disabling Memory Overcommit on Existing VMs

This procedure helps you to disable memory overcommit in one or more existing VMs. You need a minimum Prism Central version of pc.2022.4 for this procedure.

Before you begin

Ensure that the VM on which you want to disable memory overcommit is powered off. If the VM is powered on or is in **Soft Shutdown** state, the **Disable Memory Overcommit** action is not available in the **Actions** dropdown list.

About this task

To disable memory overcommit in one or more existing VM by updating it, follow this procedure.

Procedure

1. On the **VMs** dashboard (**List** tab), select the VM or VMs for which you want to disable memory overcommit.

   You can select one VM and disable memory overcommit. You could also select multiple VMs and disable memory overcommit in all the selected VMs as a bulk update.

2. In the **Actions** dropdown list, click the **Disable Memory Overcommit** action.

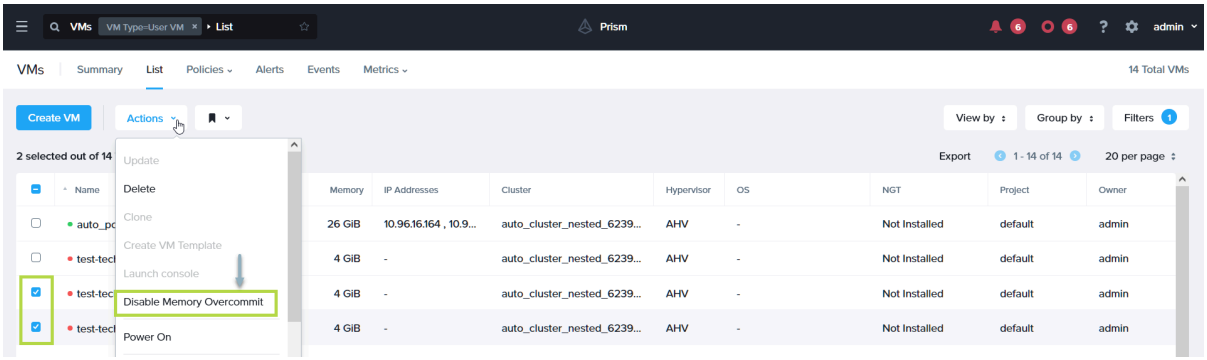

Figure 289: VM(s) Selection to Disable Memory Overcommit

> **Note:** For an individual VM, you can also select the VM, click **Actions** > **Update** and in the **Configuration** tab of the **Update VM** page, select the **Disable Memory Overcommit** checkbox.
>
> Click **Next** on the **Configuration**, **Resources** and the **Management** tabs, and **Save** on the **Review** tab.

What to do next

You can check the update tasks in the **Tasks** page. If you selected multiple VMs to disable memory overcommit on, the **Task** page displays the update for each VM as a separate **Update VM** task.

You could verify the status as **Disabled** in the individual VM details page of each VM. You could also verify the status in the **List** tab with customized view having the **Memory Overcommit** column.

VM Policy Management

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

- Affinity Policies (see Affinity Policies Defined in Prism Central on page 608)
- NGT Policies (see NGT Policies on page 611)

Affinity Policies Defined in Prism Central

In Prism Central, you can define category-based VM-Host affinity policies, where a set of VMs can be affined to run only on a particular set of hosts. Category-based affinity policy enables you to easily manage affinities for a large number of VMs. In case of any changes to the affined hosts, you only need to update the category of the host, and it updates the affinity policy for all the affected VMs.

This policy checks and enforces where a VM can be hosted when you start or migrate the VM. If there are no resources available on any of the affined hosts, the VM does not get started.

> **Note:**
>
> If you create a VM-Host affinity policy for a VM that is configured for asynchronous replication, you must create similar categories and corresponding policies on the remote site as well. If you define similar categories and policies on the remote site, affinity policies will be applied when the VMs are migrated to the remote site.

Limitations of Affinity Policies

Affinity policies created in Prism Central have the following limitations:

- Only a super admin can create, modify, or delete affinity policies.
- The minimum supported versions for VM-Host affinity policies are version 6.1 for Prism Element and version 2022.1 for Prism Central.
- Host category attach or detach takes around five minutes to get reflected in the applicable affinity policies.
- You cannot apply VM-Host affinity policy on a VM that is enabled for synchronous replication. Also, you cannot enable synchronous replication on a VM that is associated with a VM-Host affinity policy.

Affinity Policy Configuration Workflow

About this task

To set up an affinity policy, do the following:

Procedure

1. Create categories for the following entities:
   a. VMs
   b. Hosts
   For information about creating a category, see Creating a Category.
2. Apply the VM categories to the VMs and host categories to the hosts.
   For information about associating categories with VMs, see Associating VMs with Categories. For information about associating categories with hosts, see Associating hosts with Categories.
3. Create the affinity policy. See Creating an Affinity Policy.

Associating VMs with Categories

About this task

To associate categories with VMs, do the following:

Procedure

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

Associating Hosts with Categories

About this task

To associate categories with hosts, do the following:

Procedure

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

Creating an Affinity Policy

About this task

To create an affinity policy, do the following:

Before you begin

Complete the following tasks:

- Configure the categories that you need to associate with the VMs. Associate this category with all the relevant VMs. Alternatively, you can associate the category with the VMs after creation of the affinity policy.
- Configure the categories that you need to associate with the hosts. Associate this category with all the relevant hosts. Alternatively, you can associate the category with the hosts after creation of the affinity policy.

> **Note:** If you have configured any legacy affinity policy (non-category-based affinity policy) associated with the VMs, you must first remove those legacy affinity policies to allow the creation of category-based affinity policies associated with the same VMs.

Procedure

1. In the VMs Summary View, go to **Policies** > **Affinity Policies**, and then click **Create**.
2. In the **Create Affinity Policy** page, enter a name and description (optional) for the policy.

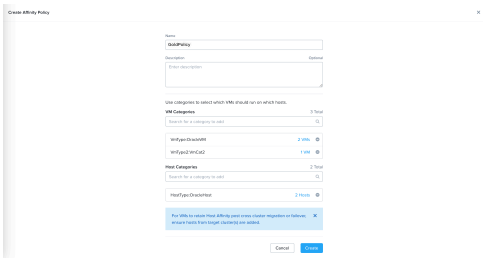

Figure 290: Create Affinity Policy

3. Click inside the **VM Categories** search field and select the category you want to associate with the VMs.
4. Click inside the **Host Categories** search field and select the category you associated with the hosts.
5. Click **Create**.

Managing an Affinity Policy

You can update, delete, and re-enforce an affinity policy. Updates to an affinity policy can result in a policy violation. Prism Central attempts to correct the violation by executing a series of actions.

About this task

To manage an affinity policy, do the following:

Procedure

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

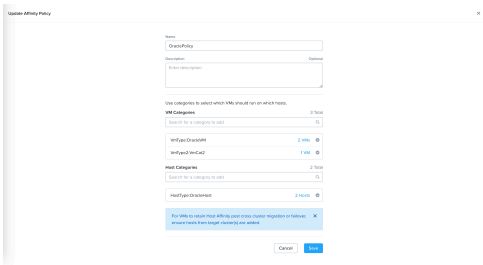

Figure 291: Update Affinity Policy

4. Update the settings that you want, and then click **Save**.

   To delete the affinity policy, select the target affinity policy checkbox and choose **Delete** from the **Actions** dropdown menu. The system prompts you to confirm the delete action. Click **Delete** to confirm the delete affinity policy action.

   To re-enforce the affinity policy, select the target affinity policy checkbox and choose **Re-Enforce** from the **Actions** dropdown menu. The system prompts you to confirm the action. Click **Re-enforce** to apply the updated affinity policy.

NGT Policies

If a restart is required post installation or upgrade, you can define a policy on when the restart should happen. You can use these policies when you have different set of VMs which you would like to be restarted at different times, for example if they are in different time zones or if their down times are expected to be different. Policy workflow are defined on the categories. You have to attach the category to the VM and then you can create a policy on the VM.

Creating a New Restart Policy

Perform the following procedure to create a new restart policy.

Procedure

1. Navigate to **Compute & Storage** > **VMs** from the Navigation bar. The VMs dashboard page opens.
2. Select **Policies** > **NGT Policies** from the dropdown menu.
   The NGT dashboard displays a list of current policies that includes the name and type for each policy.
3. Click **New VM Reboot Policy**.

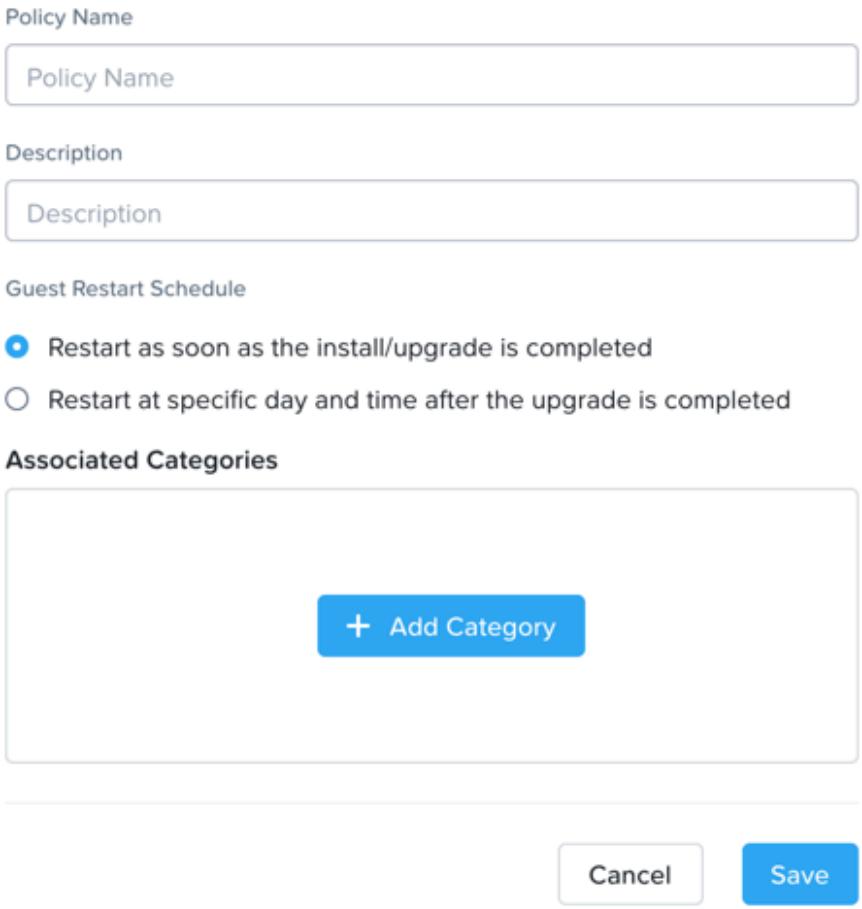

Figure 292: NGT Policy

4. In the **Policy Name** text box, type a policy name.
5. In the **Policy Description** text box, type a description of the policy.
6. Configure the guest restart schedule.

   - **Restart as soon as the install/upgrade is completed**: Select this option if you want to restart the VMs of the selected category right after the install process is completed.
   - **Restart at specific day and time after the upgrade is completed**: Select this option and choose the date and time on which you want to the restart to happen.

     Note: The VMs will be restarted depending on the Prism Element timezone settings.

7. Click **Add Category** and select the category that you want to apply this policy.

8. Click **Save**.

> **Note:** Any VM restart policy created using this procedure will be overridden if you choose to restart the VM during installation or upgradation of NGT by using a different option compared to the option configured in the policy for VM restart. If you choose not to restart the VM during installation or upgradation of NGT, the existing restart policy for the VM continues. For example, if you select the **Restart as soon as the install/upgrade is completed** option as part of a newly created policy, and if you select the **Restart at specific date and time after the install is completed** option during installation or upgradation of NGT the latter take precedence and the VM will be restarted at the scheduled date and time. However, if you select the **Skip restart** option during installation or upgradation of NGT, the VM restart policy takes precedence and the VM will restart immediately after an installation or upgrade.

The policy is created and saved in the NGT Policies pane.

Modifying a Restart Policy

You can update and delete the restart policy that you have created by following this procedure.

Procedure

1. Go to the NGT policies dashboard (see NGT Policies View on page 192).
2. Check the box for the policy to update and then select **Update** from the **Actions** drop-down menu.
   a. Make the necessary changes to the policy and click **Save**.
3. To delete a policy, select the policy and from the **Actions** drop-down menu click **Delete**.
   a. Click Delete in the confirmation dialog box to delete the policy.

VM Template Management

In Prism Central, you can create VM templates to manage the golden image of a VM. A VM template can be considered as a master copy of a virtual machine. It captures the virtual machine configuration and the contents of the VM including the guest operating system and the applications installed on the VM. You can use this template to deploy multiple VMs across clusters.

> **Note:** You can create or manage a VM template only as an admin user.

Limitations of VM Template Feature

The current implementation of the VM template feature has the following limitations:

- You cannot create a VM template if any of the following conditions is applicable at your site:
  - VM is not on AHV.
  - VM is an agent or a PC VM.
  - Any volume group is attached to the VM.
  - VM is undergoing vDisk migration.
  - VM has disks located on RF1 containers.
  - VM is protected by PD-based DR.

- VM templates do not copy the following attributes from the source VMs:
  - Host affinity attributes
  - HA priority attributes
  - Nutanix Guest Tools (NGT) installation
  - Quality of service (QoS) configuration

  You must reconfigure the above-listed attributes at the deployed VMs.

Creating a VM Template

Before you begin

- Before selecting a VM for creating a template, ensure that the VM is powered off.
- If you plan to use Sysprep for customizing the OS of the VMs that will be deployed using the template, you must first generalize the VM using Sysprep and then power off the VM before creating the template.

To create a VM template, do the following:

Procedure

1. Log in to Prism Central.
2. Select the VM in **List** tab under the VM dashboard.
3. In the **Actions** list, select **Create VM Template**.

   Alternatively, on the details page of the VM, select **More** and click **Create VM Template**.
4. This displays the **Create template from VM** dialog box. In the **Create Template from VM** window, do the following in the indicated fields:
   a. In the **VM Template Name** field, enter a name for the template.
   b. In the **Description** field, enter a description for the template. Description is an optional field.
   c. In the **Guest Customization** field, select your options for guest operating system (OS) customization.

      You can choose to provide a customization option for the OS of the VMs that will be deployed using this template. In the **Script Type** field, select **Sysprep (Windows)** for customizing the Windows OS, and **Cloud-init (Linux)** for customizing the Linux OS.

      For each of these script types, you can choose to either upload a custom script or opt for a guided setup in the **Configuration Method** field.

      - If you select **Custom Script**, you can either upload a script to customize the guest OS of the VMs, or you can copy-paste the script in the text box.
      - If you select **Guided Setup** , you must enter the authentication information such as username, password, locale, hostname, domain join, and license key. The information that you enter is used for customizing the OS of the VMs that are deployed using this template.

      You can also specify if you want to allow the template users to override the guest customization settings of the template while deploying the VM. If you select a script type and then allow the users to override the settings, the users can change the settings only for the configuration method. For example, the users can change the

authentication information at the time of deploying a VM from template, or they can change from a guided setup to a custom script.

If you select **No Customization** at the time of creating the template but allow the users to override the guest customization settings, it gives the maximum customization control to the users. They can customize the script type and the configuration method.

> **Note:** If you opt for a guest customization script, ensure that the script is in a valid format. The guest customization scripts are not validated, hence even though the VM deployment might succeed, the guest customization script may not work as expected and it will be apparent only by observing the VM after it gets deployed.

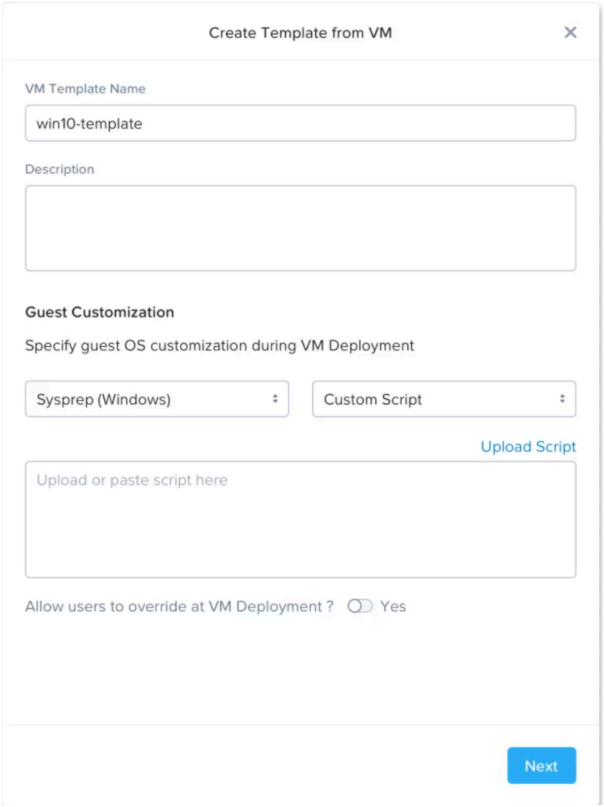

Figure 293: Create VM Template

5. Click **Next**.

   On the next page, you can review the configuration details, resource details, network details, and management details. Except the name, description and guest customization options that you chose on the previous page, you cannot modify any other setting while creating the template.

6. Click **Save** to save the inputs and create a template.

   The new template appears in the templates entity page list.

Deploying VM from a Template

About this task

After creating a VM template, you can use the template to deploy any number of VMs across clusters.

> **Note:** VM Template versions are replicated inline during VM deployment on a cluster. The first VM deployment takes longer due to the version replication. Subsequent deployments from that template version take lesser time.

To deploy a VM from a template, do the following:

Procedure

1. You can use any of the following methods to deploy VMs:

   - Select the target template in the templates dashboard (see VM Template Summary View on page 118) and click the **Deploy VMs** button. By default, the active version of the template is used for deployment.
   - Go to the details page of a selected template (see VM Template Details View on page 120) and click the **Deploy VMs** button. By default, the active version of the template is used for deployment.
   - Go to the details page of a selected template (see VM Template Details View) and then go to **Versions** view. Select the desired version that you want to use for VM deployment, and click the **Deploy VMs** button. Here, you can choose any active or non-active version for the deployment.

2. The **Deploy VM from Template** page appears. By default, you see the **Quick Deploy** page. You can click**Advanced Deploy** button to access the Advanced Deploy page. In **Advanced Deploy** option, you can view and modify some VM properties and network settings.

   - **Quick Deploy**: To deploy VMs using the quick deploy method, provide inputs in the indicated fields:

     - **Name**: Enter a name for the VM.
     - **Cluster**: Select the cluster where you want to deploy the VM.
     - **Number of VMs**: Enter the number of VMs that you want to deploy.
     - **Starting Index Number**: Enter the starting index number for the VMs when you are deploying multiple VMs simultaneously. These index numbers are used in the VM names. For example, if you are deploying two VMs and specify the starting index number as 5, the VMs are named as `vm_name`-5 and `vm_name`-6.
     - **Guest Customization**: The template can have any one of the following options for guest OS customization: **No Customization**, **Sysprep (Windows)**, or **Cloud-init (Linux)**. For **Sysprep (Windows)**, or **Cloud-init (Linux)**, you can choose to either upload a custom script or opt for a guided setup. The fields are enabled for modification only if the template allows you to override its guest

customization settings while deploying the VM. If you are not allowed to override the guest customization settings, the settings that are already provided in the template are used for VM deployment.

- Click **Next** to verify the configuration details of the VMs to be deployed.
- Click **Deploy** to deploy the VMs.

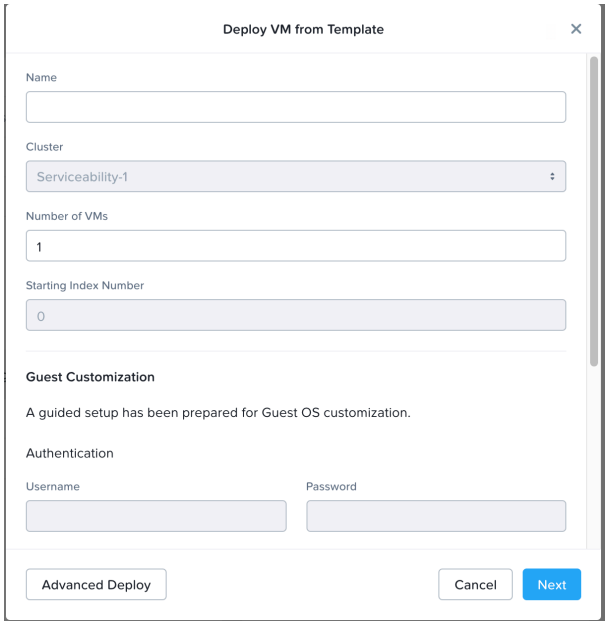

Figure 294: Deploy VM from Template using Quick Deploy Method

- **Advanced Deploy**: To deploy VMs using the advanced deploy method, provide inputs in the indicated fields and review the information displayed under various tabs:
  - **Configuration**: Provide inputs for name and description (optional) of the VM, cluster where you want to deploy the VM, Number of VMs to be deployed, and starting index number (only if deploying multiple VMs). In this tab, you can also view and modify the VM properties such as CPU, core per CPU, and memory.
  - **Resources**: Review the configuration settings for the VM resources such as disks, networks, and boot configuration. Here, you can modify the network settings but cannot modify any other settings.
  - **Management**: If the template allows you to modify the settings for guest OS customization, provide inputs for the same.

- **Review**: Review the information displayed on this tab. Click **Deploy** to deploy the VMs.

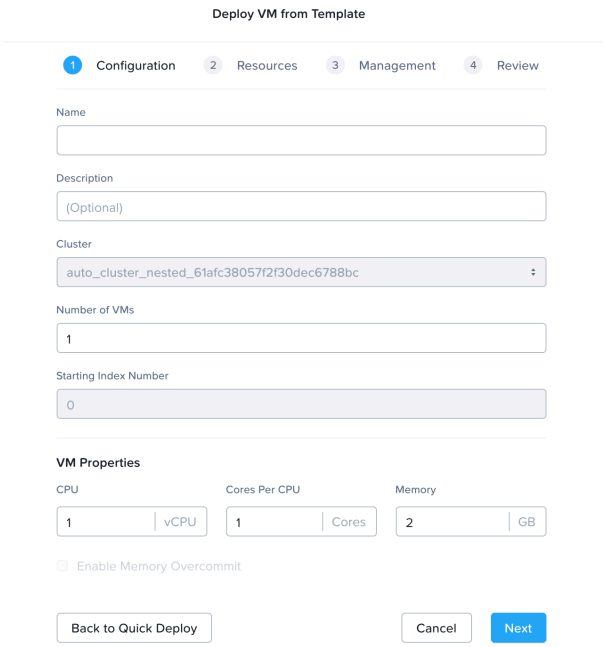

Figure 295: Deploy VM from Template using Advanced Deploy Method

Managing a VM Template

About this task

After creating a template (see Creating a VM Template on page 614), you can update the guest OS of the source VM of the template, complete guest OS update, cancel guest OS update, update the configuration of the template to create a version, or delete the template.

You can perform these tasks by using any of the following methods:

- Select the target template in the templates dashboard (see VM Template Summary View on page 118) and choose the required action from the **Actions** menu.
- Go to the details page of a selected template (see VM Template Details View on page 120) and select the desired action.

> **Note:** The available actions appear in bold; other actions are not available. The available actions depend on the current state of the template and your permissions.

Procedure

- To update the guest OS of the template, select **Update Guest OS**. Click **Proceed**.

  This step deploys a temporary VM and gives you access to that VM. You can also access the new VM from VM dashboard. You must start the VM, log on to the VM, and update the guest OS of that VM.

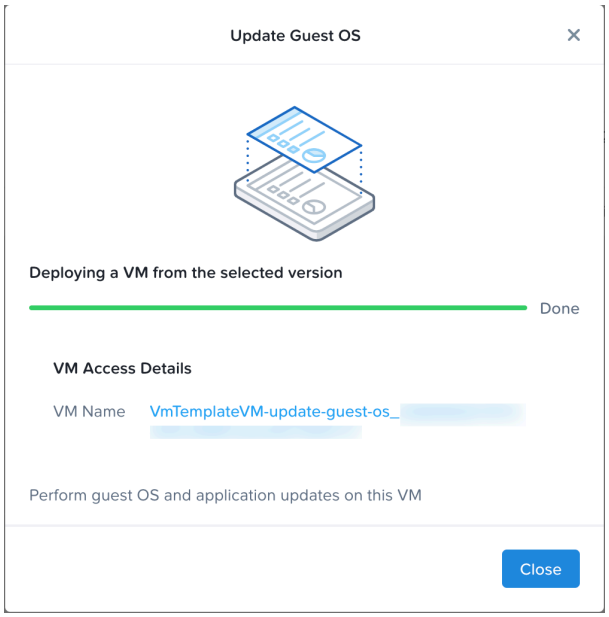

Figure 296: Update Guest OS

  Once you have updated the guest OS of the temporary VM, you must complete the guest OS update from Prism Central UI. Alternatively, you can choose to cancel the guest OS update, if you choose to discontinue the guest OS update process.

  The temporary VM automatically gets deleted after completion or cancellation of the guest OS upgrade process.

- To complete the process of guest OS update that you had earlier initiated, select **Complete Guest OS Update**. You must select this option only after successful update of the guest OS of the temporary VM.
- To cancel the process of guest OS update that you had earlier initiated, select **Cancel Guest OS Update**.
- To modify the template configuration, select **Update Configuration**. You cannot modify the configuration of the initial version of the template, but you can update the configuration settings to create a new version of the template. In **Select a version to Upgrade** window, select the version of the template that you want to modify and create a version on top of it.

  The **Update Template Configuration** window displays the following sections:

  - **Configuration**: You can view the name of the base version that you want to update, change notes for that version, cluster name, VM properties (CPU, cores per CPU), and memory). In this section, you can modify only VM properties.
  - **Resources**: You can view the information about disks, networks, and boot configuration. In this section, you can modify only network resources.
  - **Management**: In this section, you can modify the guest customization settings.
  - **Review**: In this section, you can review and modify the configuration settings that you are allowed to modify. You must provide a name and change notes for the new version. You can also choose to set this new version

as active version. An active version is the version of the template that by default gets deployed when you click the **Deploy VMs** button.

Click **Save** to save the settings and create a new version of the template.

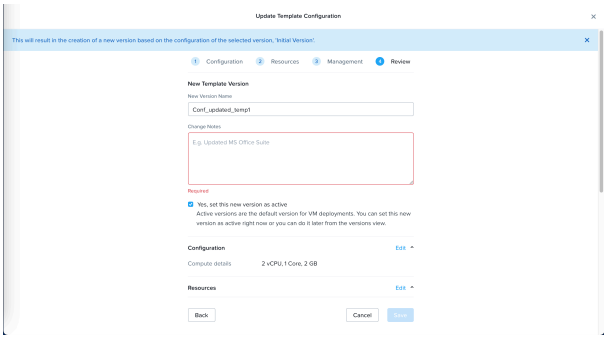

Figure 297: Update Template Configuration

- To delete a template, select **Delete Template**. A window prompt appears; click the **OK** button to delete the template.

Storage Management (Prism Central)

Storage in a Nutanix cluster is organized into several components that allow you to manage capacity and performance. Nutanix clusters include storage pool, storage container, volume group, and virtual disk components. Prism Element allows you to manage all aspects of storage for a cluster, but you can also manage some storage components directly from Prism Central.

- The storage containers dashboard displays information about storage containers across the registered clusters (see Storage Containers Summary View on page 130).
- The volume groups dashboard displays information about volume groups across the registered clusters (see Volume Groups Summary View on page 147). Using Prism Central you can do the following:

  Note: Some actions are unavailable for clusters running a pre-6.0 AOS version.

  - Create a storage container (see Creating a Storage Container (AHV) on page 629, Creating a Storage Container (ESXi) on page 634, or Creating a Storage Container (Hyper-V) on page 637)
  - Update or modify a storage container (see Modifying a Storage Container on page 640)
  - Delete a storage container (see Deleting a Storage Container on page 644)
  - Create a volume group (see Creating a Volume Group on page 644)
  - Update or delete a volume group (see Modifying a Volume Group on page 648)
  - Add, update, or remove connections to external clients for a volume group (see Managing Volume Group Connections on page 649)
  - Add, update, or delete virtual disks for a volume group (see Managing Volume Group Virtual Disks on page 653)
- The storage policy dashboard displays information about the storage policies that are configured to manage the govern the storage behaviour of VMs or entities (see Storage Policies Summary View on page 141 and Storage Policy Management on page 657).
- For information about storage management tasks that are performed through Prism Element, see the Storage Management chapter in the *Prism Web Console Guide*.

Storage Components

Storage in a Nutanix cluster is organized into the following components.

Storage Tiers

Each type of storage hardware (SSD-PCIe (NVMe), SSD (SATA SSD), and HDD) is placed in a storage tier. You can determine the tier breakdown for disks in a storage pool through the web console (see Storage Table View 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.

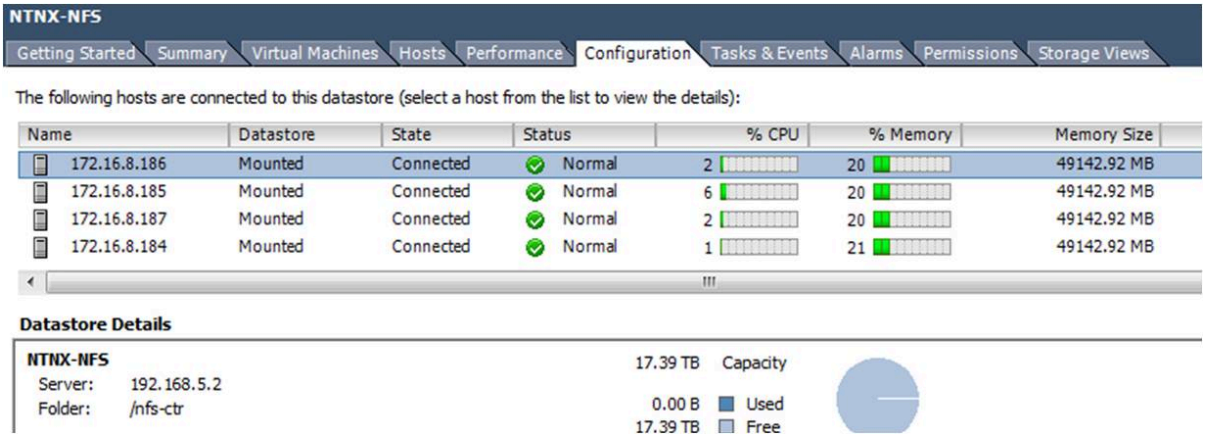

Figure 298: vSphere Configuration of NFS Datastore

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.

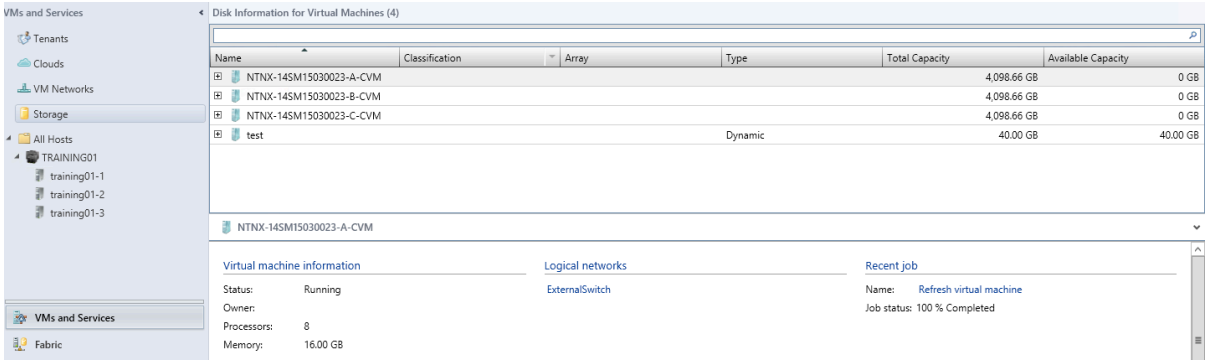

Figure 299: Hyper-V Configuration of an SMB Share

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.

Data Reduction Ratios and Data Reduction Savings

You can view data reduction ratios and data reduction savings in Prism Central for any of the following entities:

- *For the entire cluster*

  Navigate to **Summary** tab in the **Clusters** page, and observe the **Data Reduction** column entry for the target cluster in **Storage Usage** widget. For information about how to access the **Clusters** page, see Clusters Summary View on page 202.

  > **Note:** The **Data Reduction Savings** for the entire cluster is visible only in the Storage dashboard in Prism Element. For details, see Storage Overview View information in *Prism Element Web Console Guide*.

- *For an individual Storage Container*

  Navigate to **Summary** tab of an individual storage container, and observe the **Data Reduction Ratio** and **Data Reduction Savings** fields in **Optimization & Performance** widget. For information about how to navigate to the **Summary** tab of an individual storage container, see Storage Container Details View on page 136.

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. Capacity 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 in Prism Central.

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 (AHV) for AHV, Creating a Storage Container (ESXi) for ESXi or Creating a Storage Container (Hyper-V) for Hyper-V.

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.

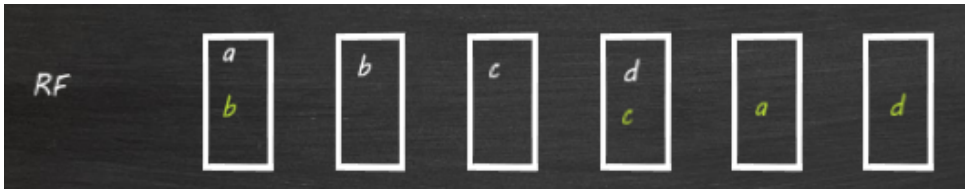

Figure 300: Data copies before Erasure Coding

When the data becomes cold, the erasure code engine computes parity “P” for the data by performing an exclusive OR operation.

Figure 301: 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.

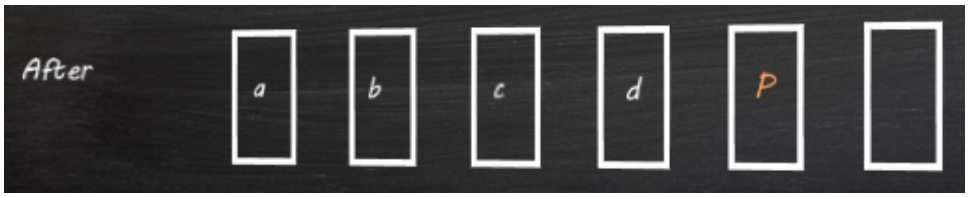

Figure 302: After Computation of Parity

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.

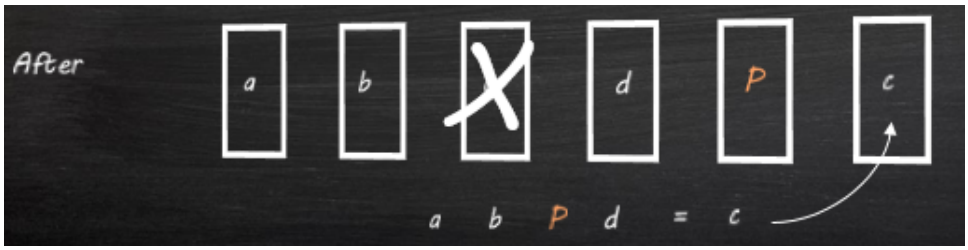

Figure 303: Post Node Failure

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.

You can view the data reduction savings from Erasure coding using any of the following methods:

- Access the Storage Container Details View on page 136, and observe the following fields:
  - **Erasure Coding** value as *On* in the **Properties** widget.
  - **Data Reduction Savings** in the **Optimization & Performance** widget.
- Access the Storage Containers Summary View on page 130, and observe the **Data Reduction Savings** column for the target storage container using any of the following methods:
  - Click the **Lists** tab, and filter the storage containers based on the **Erasure Coding** value as *On*. The system displays the results in the **Lists** page.
  - Select **Optimization** from the **View by** option and **Erasure Coding** from the **Group by** option. The system displays the results in the **Lists** page with **Erasure Coding: Off** and **Erasure Coding: On.**

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.

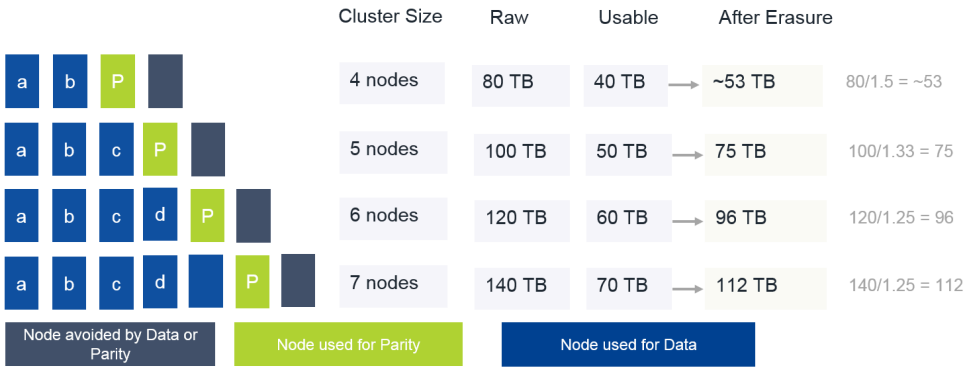

Figure 304: Example of Space Saving from Erasure Coding on 20 TiB Nodes

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 737: 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.

Before you begin

Caution:

- Nutanix recommends that you enable inline erasure coding for Object storage containers only. To enable inline erasure coding for any other type of storage container, contact Nutanix Support.
- Erasure coding must be enabled on the container to enable inline erasure coding. For information about how to enable erasure coding, see Creating a Storage Container (AHV) on page 629, Creating a Storage Container (ESXi) on page 634, or Creating a Storage Container (Hyper-V) on page 637.

Procedure

To enable inline erasure coding, perform the following actions:

- Run the following nCLI command:

  ```
  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 inline erasure coding.

  » To explicitly configure inline erasure coding type as same vDisk strips:

  ```
  ncli> container create name=container_name sp-id=storage_pool_id erasure-code=on
   inline-ec-enabled=true inline-ec-type=same-vdisk-strips
  ```

  » To explicitly configure inline erasure coding type as cross vDisk strips:

  ```
  ncli> container create name=container_name sp-id=storage_pool_id erasure-code=on
   inline-ec-enabled=true inline-ec-type=cross-vdisk-strips
  ```

  To change an existing inline erasure coding type, run the following nCLI commands.

  » To change an existing inline erasure coding type to same vDisk strips:

  ```
  ncli> container edit inline-ec-enabled=true inline-ec-type=same-vdisk-strips
   id=container_id
  ```

  » To change an existing inline erasure coding type to cross vDisk strips:

  ```
  ncli> container edit inline-ec-enabled=true inline-ec-type=cross-vdisk-strips
   id=container_id
  ```

  Replace *container_id* with the ID of the storage container.

Capacity Reservation Best Practices

By default, each storage container has access to all of the unused storage in the storage pool. If a storage pool has multiple storage containers, one storage container may take all the remaining storage space and leave others with no available space. To make sure that there is space for a storage container, you can enable capacity reservation.

Capacity reservation allows you to guarantee that a storage container has a minimum amount reserved that is unavailable to other storage containers.

Following are best practices for capacity reservation.

- Reserve capacity for a storage container only if the storage pool has multiple storage containers. Unless there is a specific reason to have multiple storage containers, Nutanix recommends having a single storage pool with a single storage container.
- In total, reserve no more than 90% of the space in the storage pool.
- When setting an advertised capacity for a storage container, be aware that some extra space should be allocated beyond the projected size of any VMs placed in the container to allow room for data that has not yet been garbage collected, which can be substantial depending on the workload (10% or more of the storage capacity in some cases).

Storage Container Management

A storage pool is created during cluster configuration. There is one storage pool per cluster. Storage efficiency features including compression, deduplication, erasure coding, and replication factor are enabled at the storage container level. To access the storage containers dashboard, select **Compute & Storage** > **Storage Containers** from the entities menu (see Entities Menu on page 13).

The storage containers dashboard displays information about storage containers across the registered clusters (see Storage Containers Summary View on page 130). Using Prism Central you can do the following:

-
  - Create a storage container (see Creating a Storage Container (AHV) on page 629, Creating a Storage Container (ESXi) on page 634, or Creating a Storage Container (Hyper-V) on page 637)
  - Update or modify a storage container (see Modifying a Storage Container on page 640)
  - Delete a storage container (see Deleting a Storage Container on page 644)
- For information about storage management tasks that are performed through Prism Element, see the Storage Management chapter in the *Prism Web Console Guide*.

Note:

- A storage pool and three storage container are created automatically when the cluster is created.
- A storage container is not created if you have not configured the Controller VMs with enough memory. Controller VM memory allocation requirements differ depending on the models and features that are being used. For more information, see *Controller VM Memory Configurations* in *Acropolis Advanced Administration Guide*.
- The NutanixManagementShare, SelfServiceContainer, and a default-container are created by default.
  - The NutanixManagementShare storage container is a built-in storage container for Nutanix clusters for use with the Nutanix Files and Self-Service Portal (SSP) features. This storage container is used by Nutanix Files and SSP for file storage, feature upgrades, and other feature operations. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.
  - SelfServiceContainer storage container is also a built-in storage container for Nutanix clusters for use with Image service features.

  To ensure proper operation of these features, do not delete these storage containers. Nutanix also recommends that you do not delete this storage container even if you are not using these features.

Known Issues and Limitations

RF1 containers are not supported in Prism Central. You can create RF1 containers only from Prism Element. However, you cannot update or delete Prism Element created RF1 containers from Prism Central.

Creating a Storage Container (AHV)

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers (Prism Central) on page 515 for information on how to configure the NTP server) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor. The storage container is accessible transparently for AHV.

To create a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage containers dashboard.

4. Click the **Create Storage Container** button.
   The **Create Storage Container** page appears.

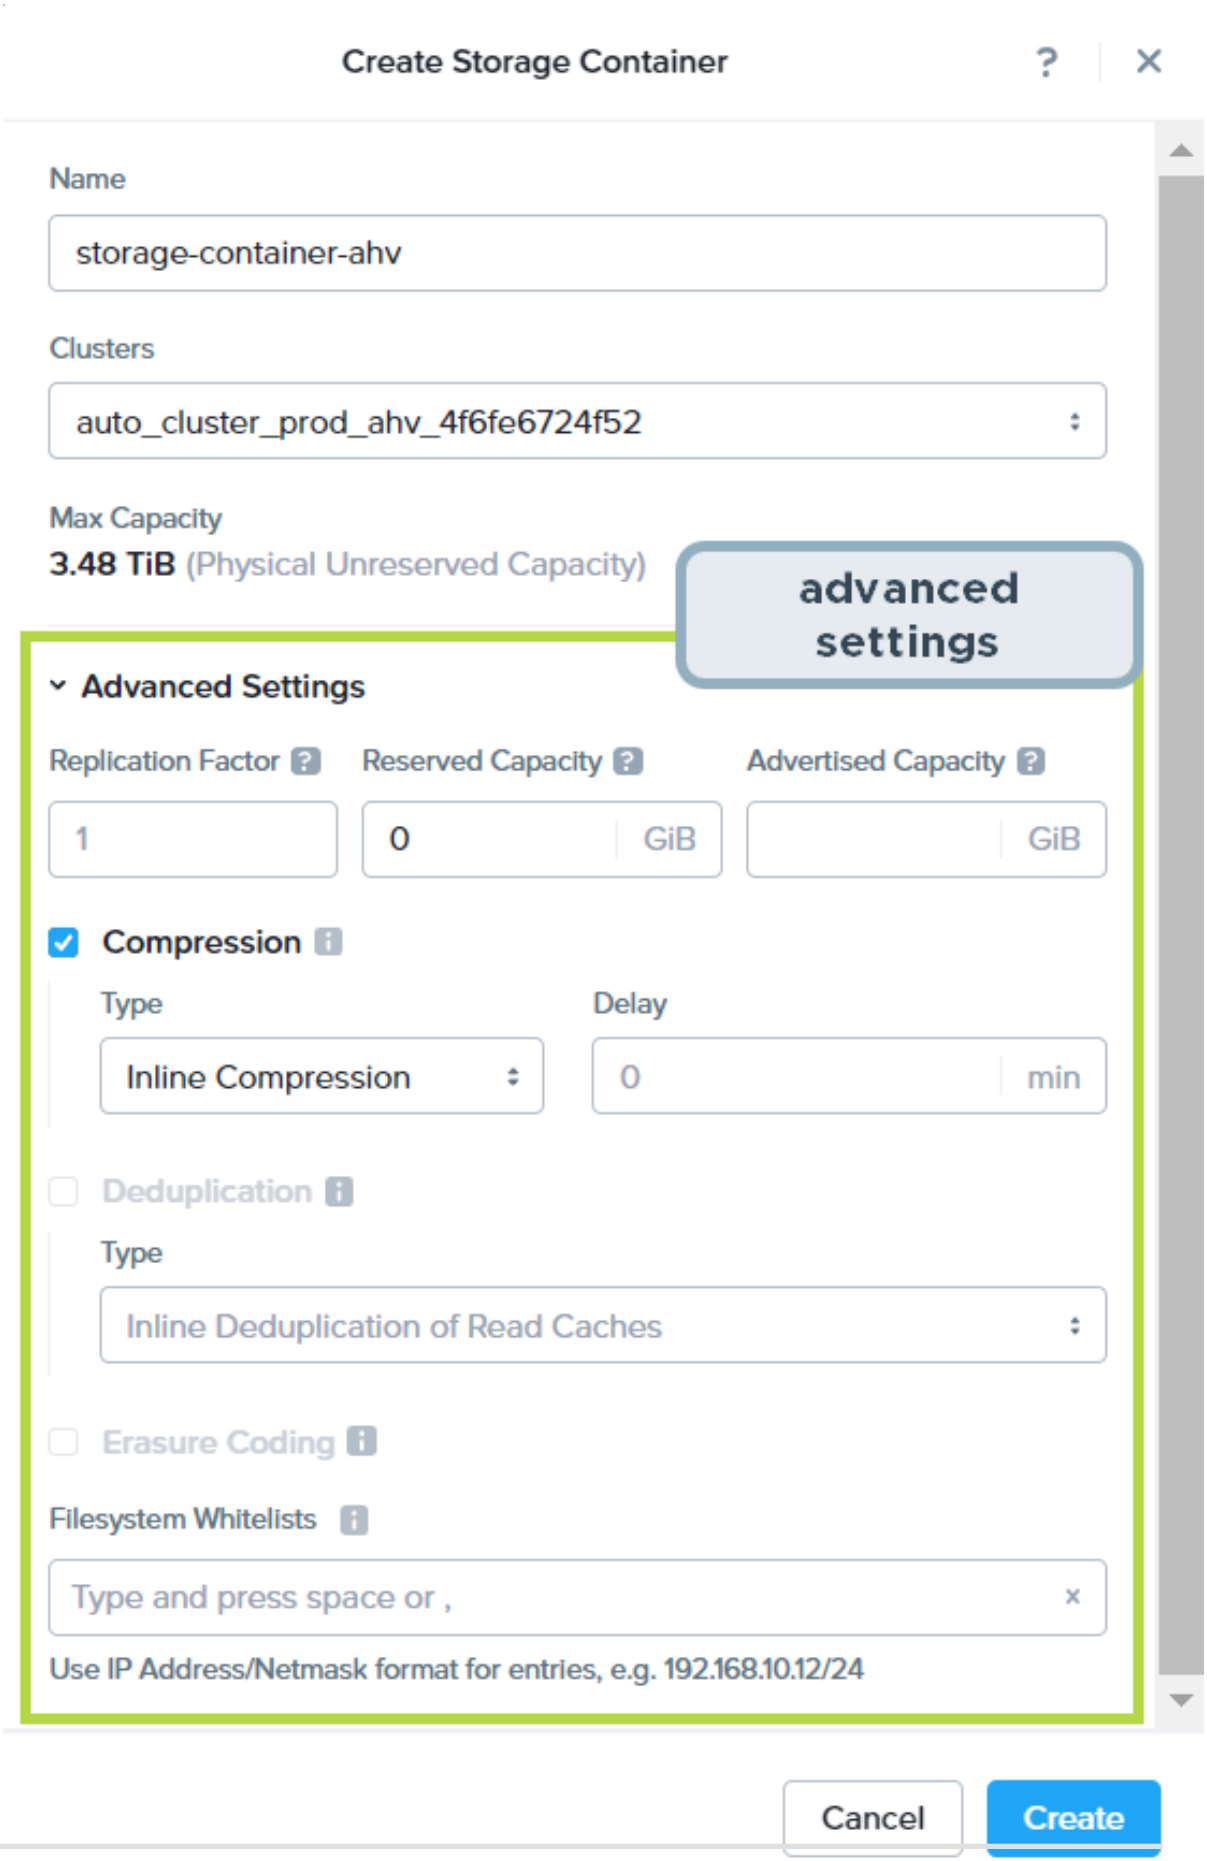

Figure 305: Create Storage Container (AHV)

5. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

   > **Note:** This entity has the following naming restrictions.
   >
   > - Container Name Length: Maximum length is 75 characters.
   > - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   > - Case Sensitivity: Container names are case sensitive.

   b. **Clusters**: Select a cluster from the drop-down list.

   c. **Max Capacity** displays the amount of free space available in the selected cluster.

6. To configure additional parameters, click the **Advanced Settings**.

   Do the following in the indicated fields:

   a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

   Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > **Note:** To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See the *Increasing the Cluster Fault Tolerance Level* in the *Prism Web Console Guide*.

   b. **Reserved Capacity**: To reserve storage space for this storage container, enter the capacity (in GiB).

   You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 628 for more information.

   c. **Advertised Capacity**: To reserve the maximum storage space for this storage container, enter the capacity (in GiB).

   This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity.

   d. **Compression**: By default, **Inline compression** is selected in the **Type** field and **0** selected in the **Delay field (In Minutes)** which indicates that the data is compressed immediately as it is written. To configure the delay time between the write and compression, under **Type**, select **Post Process Compression** from the drop-down menu. For post-process compression, where data is compressed after it is written, Nutanix recommends settings a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

   All data in the storage container is compressed when you select **Compression**. See Compression on page 623 for information about using compression.

   e. **Deduplication**: Select the **Deduplication** checkbox, and choose any of the following options in the **Type** dropdown menu to optimize performance:

   - **Inline Deduplication of Read Caches** - Used to perform deduplication on the data in memory. This option is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. For details, see *Cache Deduplication* information in Deduplication on page 623.
   - **Inline Deduplication of Read Caches + Post Process Deduplication** - Used to perform post-process deduplication of persistent data. This option is primarily recommended for full clone,

persistent desktops, and physical to virtual migration use cases that need both storage capacity savings and performance savings from deduplication. For details, see *Capacity Deduplication* information in Deduplication on page 623.

> **Important:** Deduplication is only supported on clusters with a minimum of three nodes.

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 624.

> **Note:** Erasure coding is supported on clusters with a minimum of 4 nodes when using RF2 and a minimum of 6 nodes when using RF3.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form `ip_address/ netmask`).

A whitelist (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. If you set an allowlist at storage container level, the system overrides any global whitelist for this storage container.

Setting an allowlist helps you provide access to the container via NFS. Some manual data migration workflows might require the allowlist to be configured temporally, while some third-party backup vendors might require the allowlist to be configured permanently to access the container via NFS.

> **Caution:**
>
> - There is no user authentication for NFS access, and the IP address in the allowlist has full read or write access to the data on the container.
> - It is recommended to allow single IP addresses (with net mask such as 255.255.255.255) instead of allowing subnets (with netmask such as 255.255.255.0).

7. When all the field entries are correct, click the **Create** button.

Creating a Storage Container (ESXi)

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers (Prism Central) on page 515 for information on how to configure the NTP server) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor. For ESXi, the storage container is accessible as an NFS datastore. This requires access to the vSphere APIs. Ensure that you have appropriate license of vSphere to access the APIs.

To create a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage containers dashboard.

4. Click the **Create Storage Container** button.
   The **Create Storage Container** page appears.

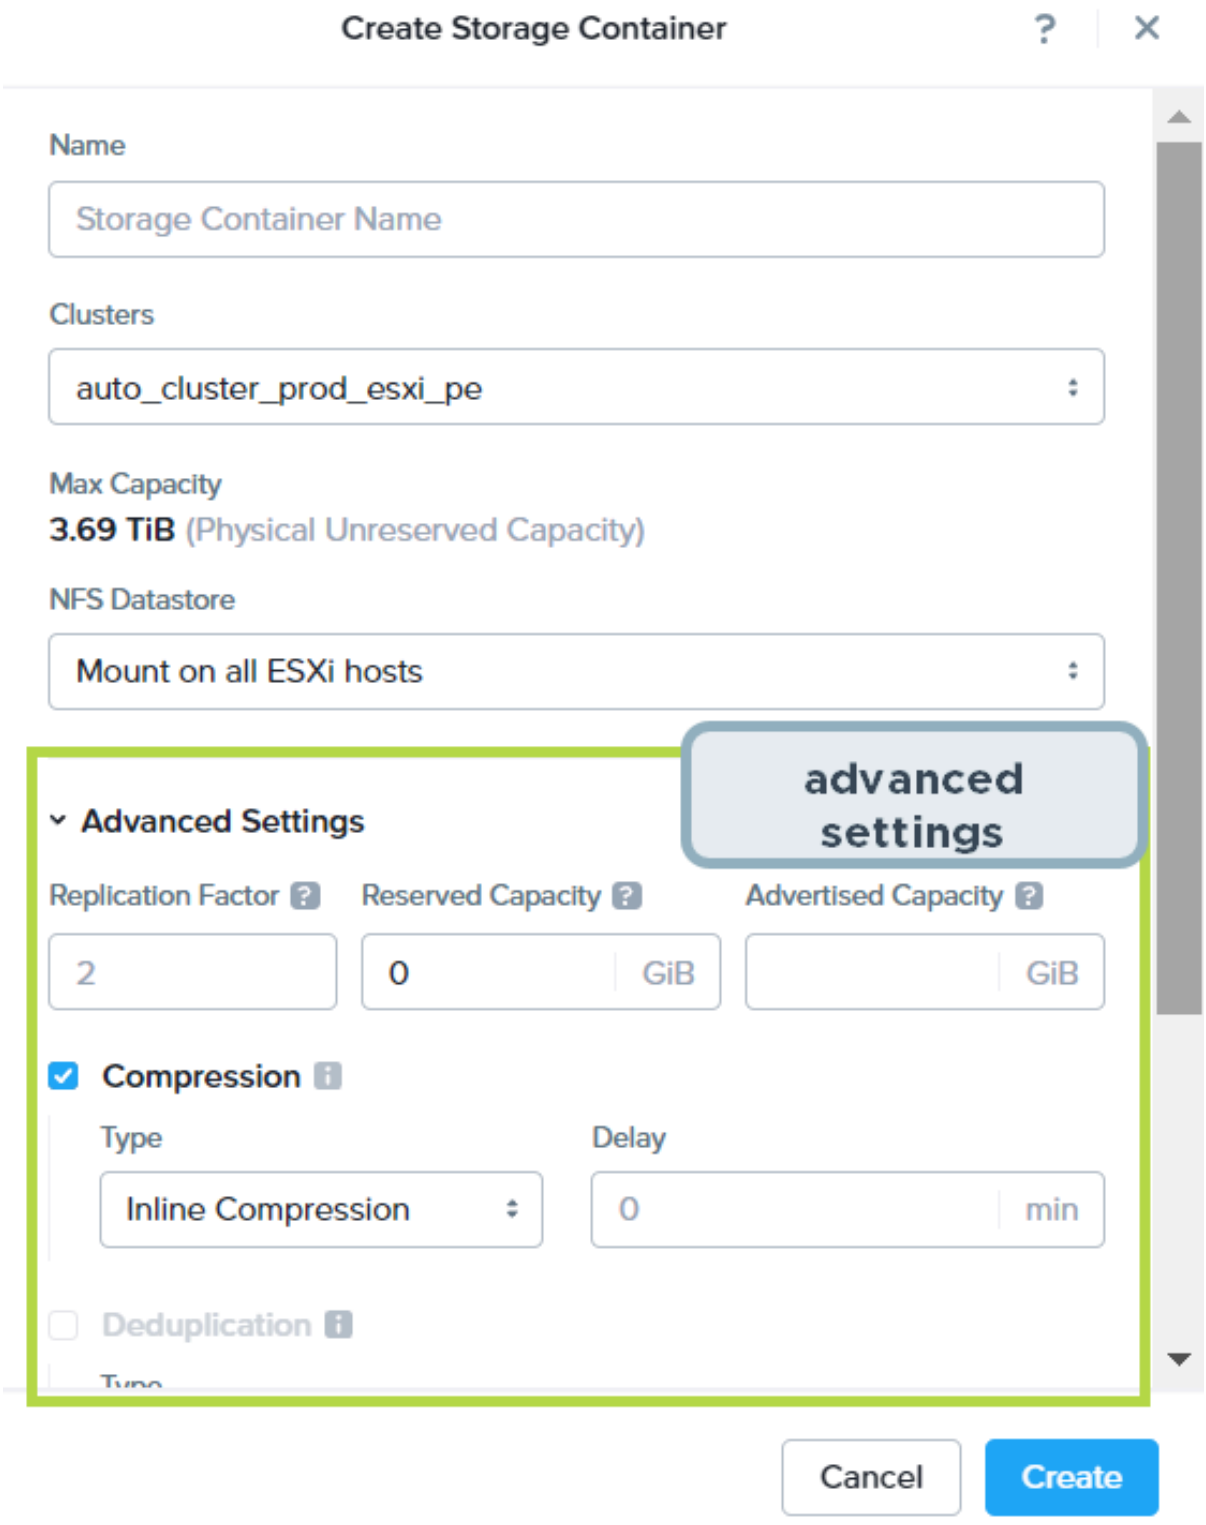

Figure 306: Create Storage Container (AHV)

5. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

   > **Note:** This entity has the following naming restrictions.
   >
   > - Container Name Length: Maximum length is 42 characters.
   > - Supported Characters: Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-) , hash (#), and underscores (_).
   > - Case Sensitivity: Container names are case sensitive.

   b. **Clusters**: Select a cluster from the drop-down list.

   c. **Max Capacity** displays the amount of free space available in the selected cluster.

   d. **NFS Datastore**: Select the **Mount on all ESXi hosts** button to mount the storage container on all hosts. Select the **Mount on the following ESXi hosts** button to mount the storage container on a subset of hosts, which displays a list of host IP addresses below this field. Check the boxes of the hosts to be included.

6. To configure additional parameters, click the **Advanced Settings**.

   Do the following in the indicated fields:

   a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

   Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > **Note:** To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See the *Increasing the Cluster Fault Tolerance Level* in the *Prism Web Console Guide*.

   b. **Reserved Capacity**: To reserve storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 628 for more information.

   c. **Advertised Capacity**: Sets a maximum storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

   d. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is

configurable. For post-process compression, where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

All data in the storage container is compressed when you select **Compression**. See Compression on page 623 for information about using compression.

e. **Deduplication**: Select the **Deduplication** checkbox, and choose any of the following options in the **Type** dropdown menu to optimize performance:

- **Inline Deduplication of Read Caches** - Used to perform deduplication on the data in memory. This option is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. For details, see *Cache Deduplication* information in Deduplication on page 623.
- **Inline Deduplication of Read Caches + Post Process Deduplication** - Used to perform post-process deduplication of persistent data. This option is primarily recommended for full clone, persistent desktops, and physical to virtual migration use cases that need both storage capacity savings and performance savings from deduplication. For details, see *Capacity Deduplication* information in Deduplication on page 623.

> **Important:** Deduplication is only supported on clusters with a minimum of three nodes.

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 624.

> **Note:** Erasure coding is supported on clusters with a minimum of 4 nodes when using RF2 and a minimum of 6 nodes when using RF3.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form `ip_address/ netmask`).

A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

> **Note:** Setting a storage container level whitelist overrides any global whitelist for this storage container.

7. When all the field entries are correct, click the **Create** button.

Creating a Storage Container (Hyper-V)

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see Configuring NTP Servers (Prism Central) on page 515 for information on how to configure the NTP server) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor. The storage container is accessible as an SMB share for Hyper-V.

To create a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.

3. Go to the **List** tab in the storage containers dashboard.
4. Click the **Create Storage Container** button.
   The **Create Storage Container** page appears.

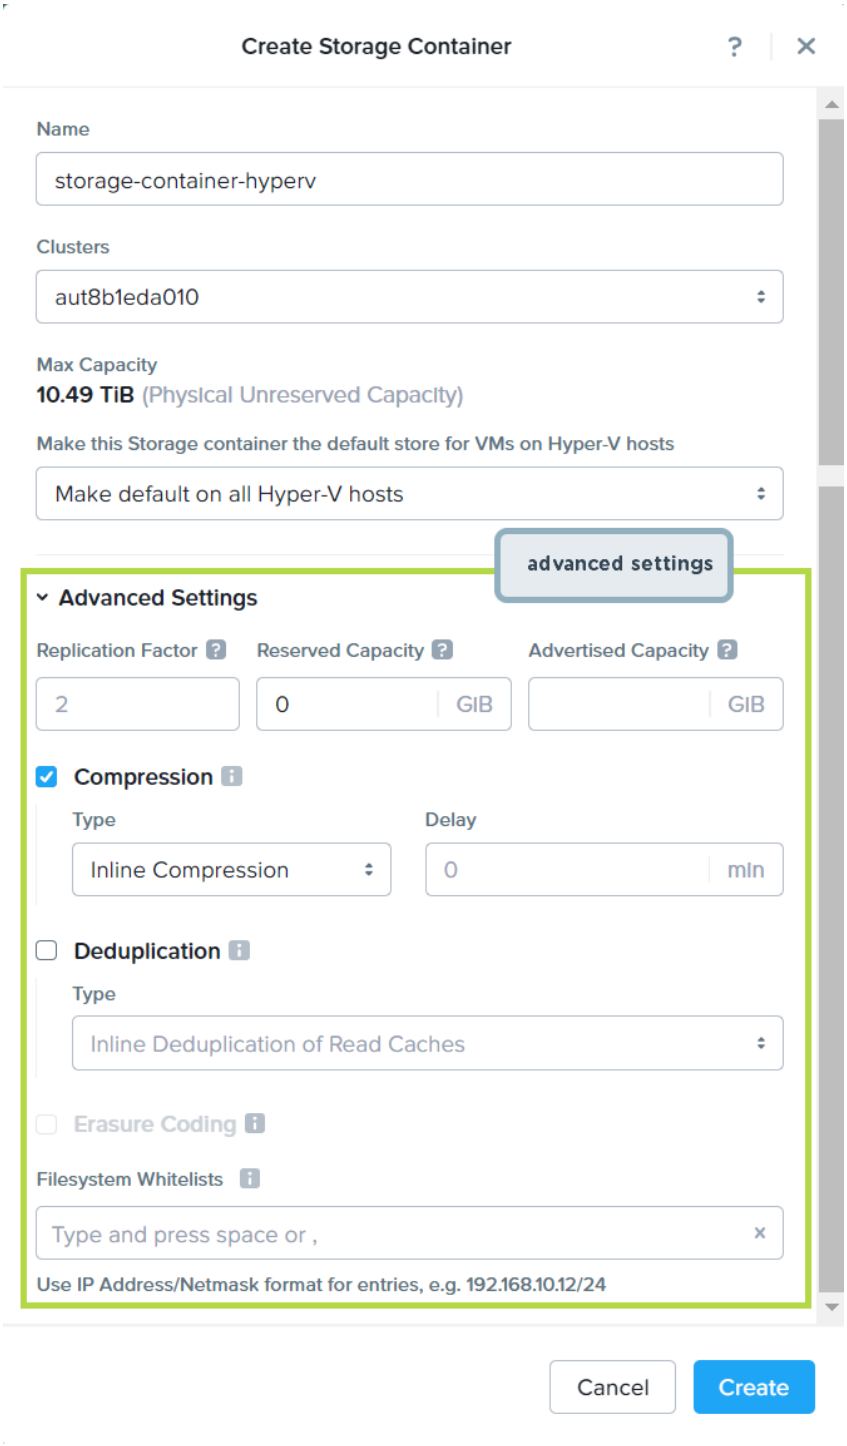

Figure 307: Create Storage Container (AHV)

5. Do the following in the indicated fields:

   a. **Name**: Enter a name for the storage container.

   > **Note:** This entity has the following naming restrictions.
   >
   > - Container Name Length: Maximum length is 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. **Clusters**: Select a cluster from the drop-down list.

   c. **Max Capacity**: Displays the amount of free space available in the selected cluster.

   d. **Make this storage container a default store for VMs on Hyper-V Hosts**: Select from the following options:

Option	Description
Make default on all Hyper-V hosts	Makes this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.
Make default on particular Hyper-V hosts	Provides you with an option to select the hosts on which you want to make this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.

   Depending on your selection in the Create Virtual Machine Wizard of Hyper-V, the system automatically populates the storage location with the relevant storage container.

6. To configure additional parameters, click the **Advanced Settings**.

   Do the following in the indicated fields:

   a. **Replication Factor**: Displays the number of data copies to maintain in the cluster.

   Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

   > **Note:** To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See the Capacity Reservation Best Practices on page 628 for more information.

   b. **Reserved Capacity**: To reserve storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer available to other storage containers even if the reserved space is unused. See Capacity Reservation Best Practices on page 628 for more information.

   c. **Advertised Capacity**: Sets a maximum storage space for this storage container, enter the amount (in GiB) to reserve in this field.

   This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go

beyond the advertised capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

d. **Compression**: Inline compression is enabled by default with the **Delay (In Minutes)** field set to 0. A value of 0 means data is compressed immediately as it is written. The delay time between write and compression is configurable. For post-process compression, where data is compressed after it is written, Nutanix recommends a delay of 60 minutes. Compression is delayed for 60 minutes after the initial write operation.

   All data in the storage container is compressed when you select **Compression**. See Compression on page 623 for information about using compression.

e. **Deduplication**: Select the **Deduplication** checkbox, and choose any of the following options in the **Type** dropdown menu to optimize performance:

   - **Inline Deduplication of Read Caches** - Used to perform deduplication on the data in memory. This option is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. For details, see *Cache Deduplication* information in Deduplication on page 623.
   - **Inline Deduplication of Read Caches + Post Process Deduplication** - Used to perform post-process deduplication of persistent data. This option is primarily recommended for full clone, persistent desktops, and physical to virtual migration use cases that need both storage capacity savings and performance savings from deduplication. For details, see *Capacity Deduplication* information in Deduplication on page 623.

     Important: Deduplication is only supported on clusters with a minimum of three nodes.

f. **Erasure Coding**: Select the check box to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see Erasure Coding on page 624.

   Note: Erasure coding is supported on clusters with a minimum of 4 nodes when using RF2 and a minimum of 6 nodes when using RF3.

g. **Filesystem Whitelists**: Enter the comma-separated IP address and netmask value (in the form `ip_address/ netmask`).

   A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

   Note: Setting a storage container level whitelist overrides any global whitelist for this storage container.

7. When all the field entries are correct, click the **Create** button.

Modifying a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be modified as conditions change.

About this task

Storage Containers can be modified to change how the data in that storage container is handled, for example to apply compression.

Note: You cannot rename a storage container in an AHV cluster when modifying container details through the **Update Storage Container** dialog box. You cannot rename a storage container if it contains vDisks.

To modify a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.

2. Click **Storage Containers**.
3. Go to the **List** tab in the storage container dashboard.
4. Select the target storage container checkbox, and choose **Update** from the **Actions** dropdown menu.
   The **Update Storage Container** window appears, which includes the same fields as the **Create Storage Container** window. Through this window you can specify NFS datastore mounts, reserve storage capacity, enable

(or disable) compression, enable or disable erasure coding, select or deselect fingerprinting on writes which enables deduplication on reads, and configure filesystem whitelist.

For example,

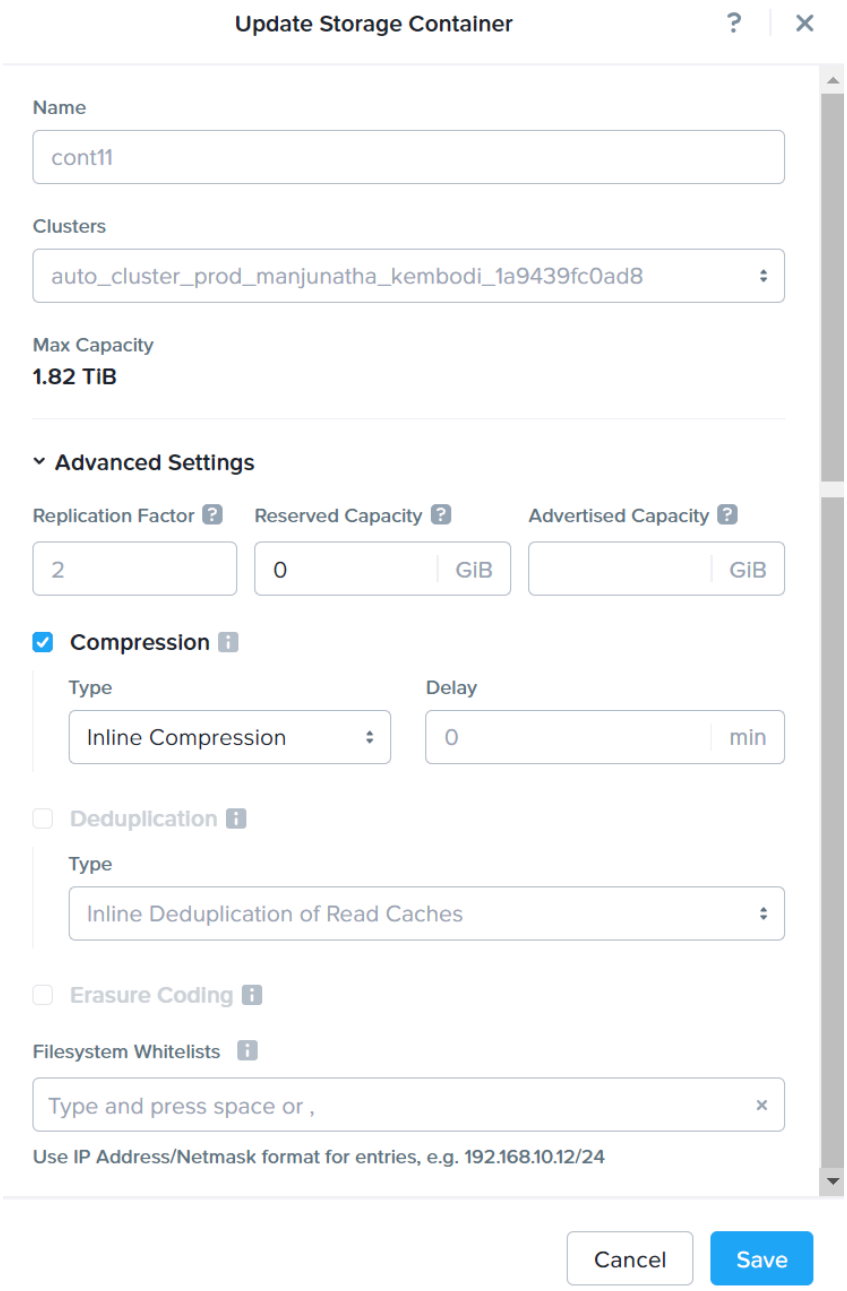

Figure 308: Update Storage Container (AHV)

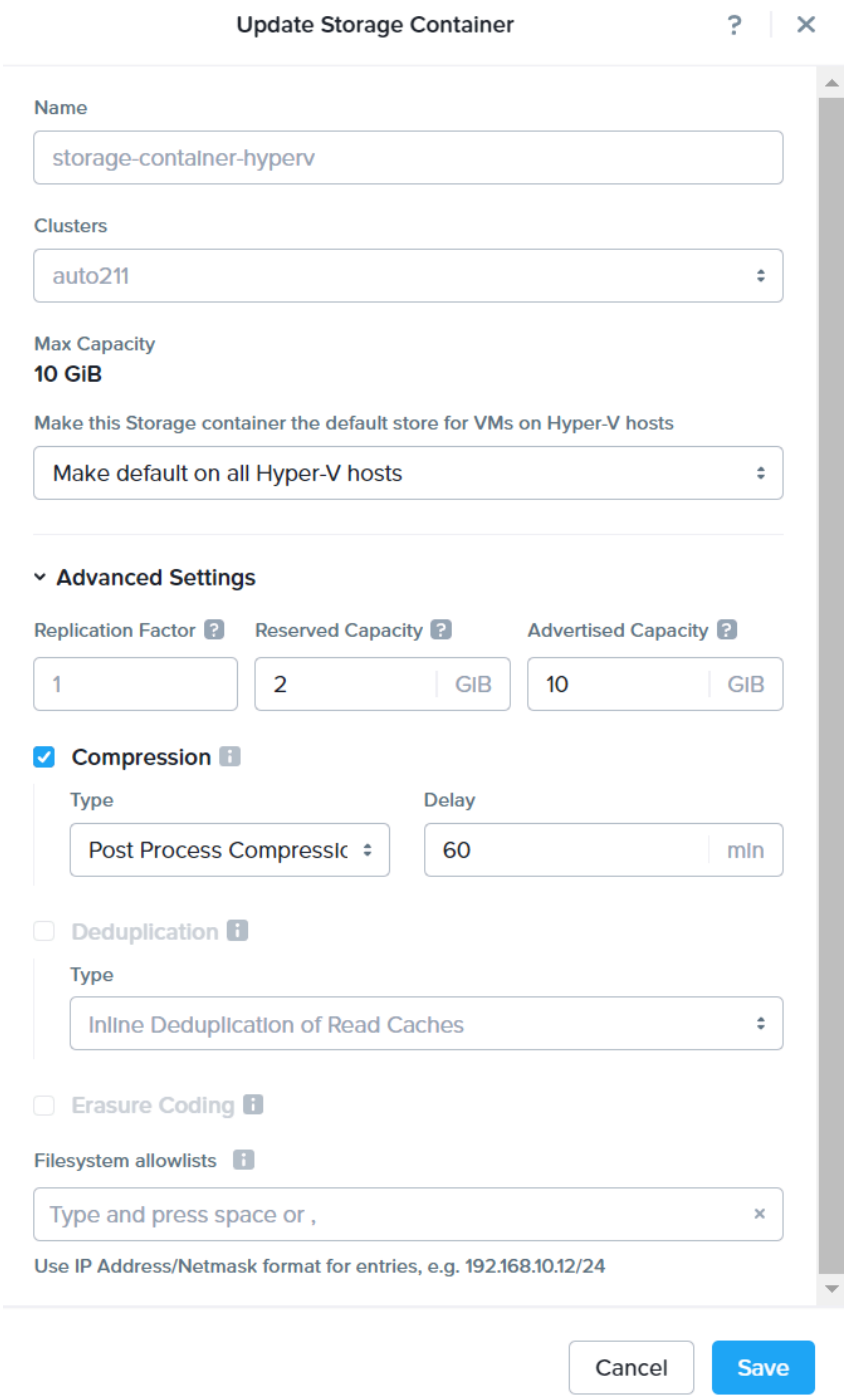

Figure 309: Update Storage Container (Hyper-V)

Note:

- If the compression policy is changed from compressed to uncompressed (or vice versa), the existing compressed (uncompressed) data in the storage container will be uncompressed (compressed) as a background process when the next data scan detects the data that needs this change.

- The Prism Central does not provide an option to change the container replication factor. That can be done only through the nCLI (see Increasing the Cluster Fault Tolerance Level section in the *Prism Web Console Guide*).

Deleting a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be deleted.

About this task

To delete a storage container, do the following:

Procedure

1. In the **Entities** menu, click **Compute and Storage**.
2. Click **Storage Containers**.
3. Go to the **List** tab in the storage container dashboard.
4. To delete the storage container, select the target storage container.
5. Click the **Delete** link.
   The **Deleting Storage Container** window appears which requests you to confirm deletion of the selected storage container.
6. Click the **Delete** button.

Creating a Volume Group

About this task

To create a volume group through Prism Central, do the following:

Figure 310: Creating a Volume Group

A video on this topic is available https://players.brightcove.net/5850956868001/xJM5EiUJJ_default/index.html?videoId=6275775660001 on Brightcove.

Procedure

1. Go to the **List** tab in the volume groups dashboard (see Volume Groups Summary View on page 147) and click the **Create Volume Group** button.
   The **Create Volume Group** page appears.

2. In the **Configuration** screen, do the following in the indicated fields:

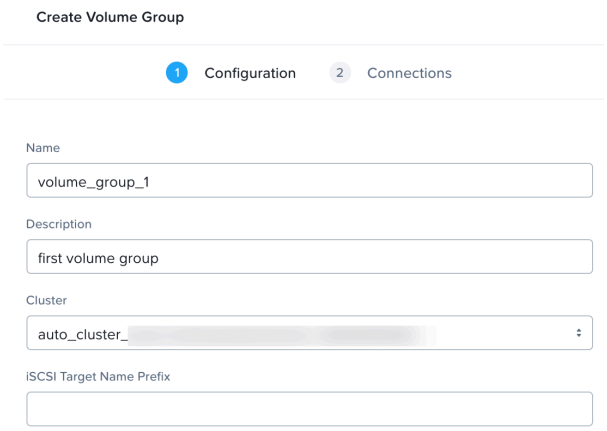

Figure 311: Create Volume Group: Configuration

a. **Name**: Enter a name for the volume group.

b. **Description**: Enter an optional description for the volume group.

c. **Cluster**: Select the cluster in which to create the volume group from the pull-down list.

   The list includes all clusters registered with this Prism Central.

d. **iSCSI Target Name Prefix**: Enter the iSCSI target name prefix.

e. **Virtual Disks**: Click the **Add Disk** button. In the **Add Virtual Disk** window, select the target storage container from the pull-down list, enter the disk size (in GiBs), and then click the **Add** button.

   The **Configuration** screen reappears with a table of virtual disks. The table includes index (assigned number), storage container (name), size, and action fields with a line for each disk in the volume group.

   - To add another disk to the volume group, click the **Add Disk** link above the table. This action redisplays the **Add Virtual Disk** window. Repeat this step for the new disk.
   - To change the size of a disk, click the pencil mark icon for that disk. An **Edit Virtual Disk** window appears in which you can change the size (but not the storage container).
   - To delete a disk, click the garbage can icon for that disk.

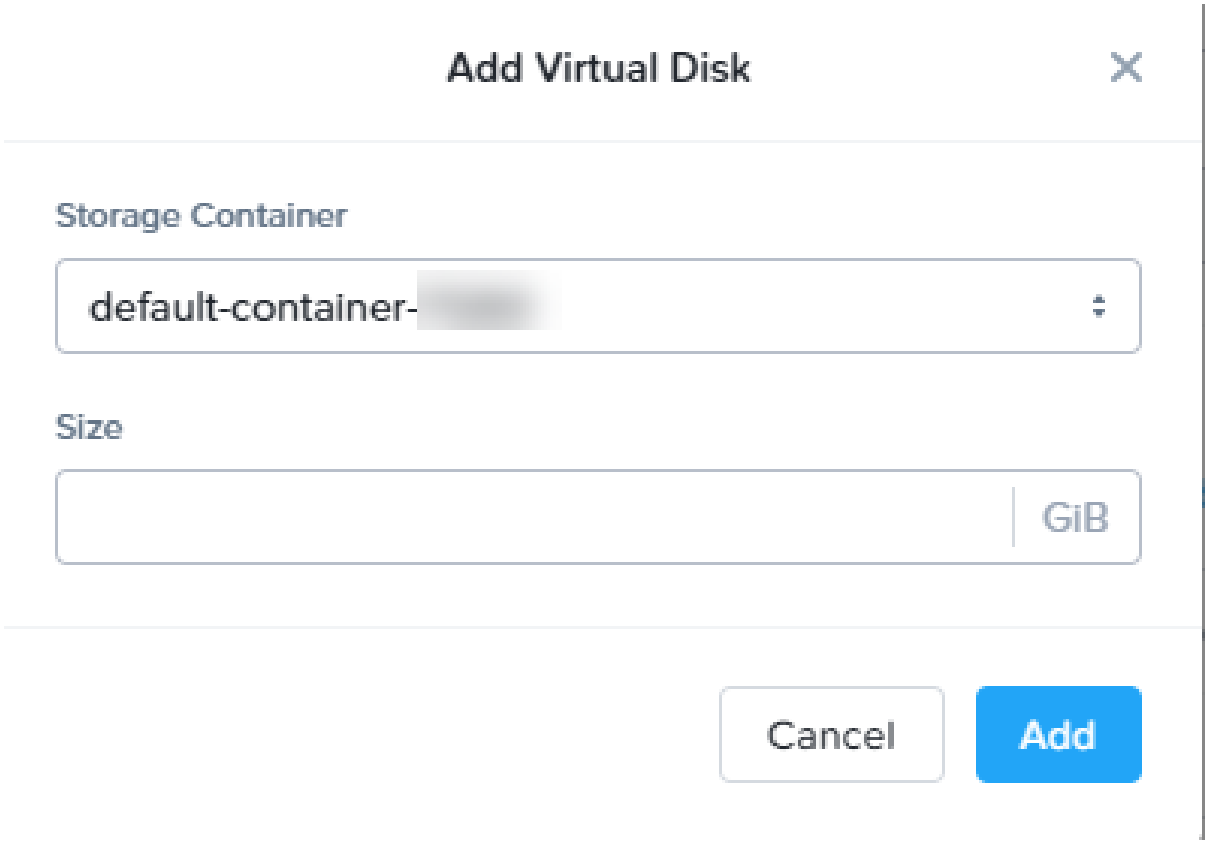

Figure 312: Add Virtual Disk Window

f. **Advanced Settings** (optional): Check the **Flash Mode** box to ensure no down migration of data occurs from the flash tier.

   Clicking **Advanced Settings** switches (display/hide) the **Flash Mode** box.

g. Click the **Next** button.

3. In the **Connections** screen, do the following in the indicated fields:

   » **Virtual Machines** - Allows you to configure direct attachment to VMs on the same cluster.

   » **External Clients** - Allows you to configure access to external clients or clients not residing on the same cluster.

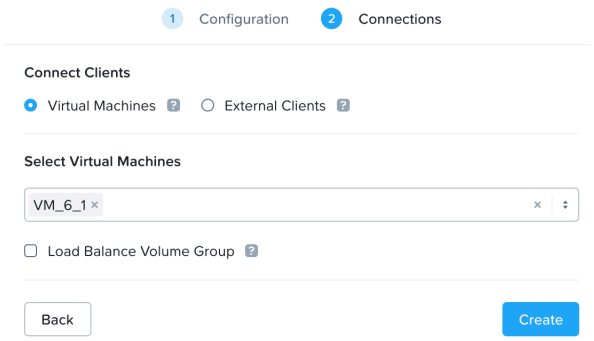

Figure 313: Create Volume Group: Connections - Virtual Machines

Configure the **Virtual Machines** settings as indicated.

- **Select Virtual Machines**: Select one or more VMs using the drop-down menu.
- (Optional) **Load Balance Volume Group**: Check to enable load balancing between the selected VMs. Virtual disks in the volume group are evenly distributed between the CVMs instead of being hosted locally.

> **Note:** You cannot change the load balance configuration once you attach the VMs.

Configure the **External Clients** configuration as indicated.

a. **CHAP Authentication**: Check the box to enable the Challenge Handshake Authentication Protocol (CHAP) and then enter the password in the indicated field.

b. **Attached Client**: Configure the clients as desired. The configured clients appear in the table.

   - To attach (unattach) a client, check (uncheck) the box for that client in the table.
   - To add a client, click the **Add Client** link. In the **Add External Client** window, enter the client IP address or IQN designator name in the indicated field, check the **CHAP Authentication** box and

password to enable CHAP authentication, and then click the **Add** button. The added client appears in the table.

- To update a client, click the pencil icon for that client in the table. The **Edit External Client** window appears. Check (uncheck) the box to enable (disable) CHAP authentication and then click the **Save** button. (You cannot change the client address.)

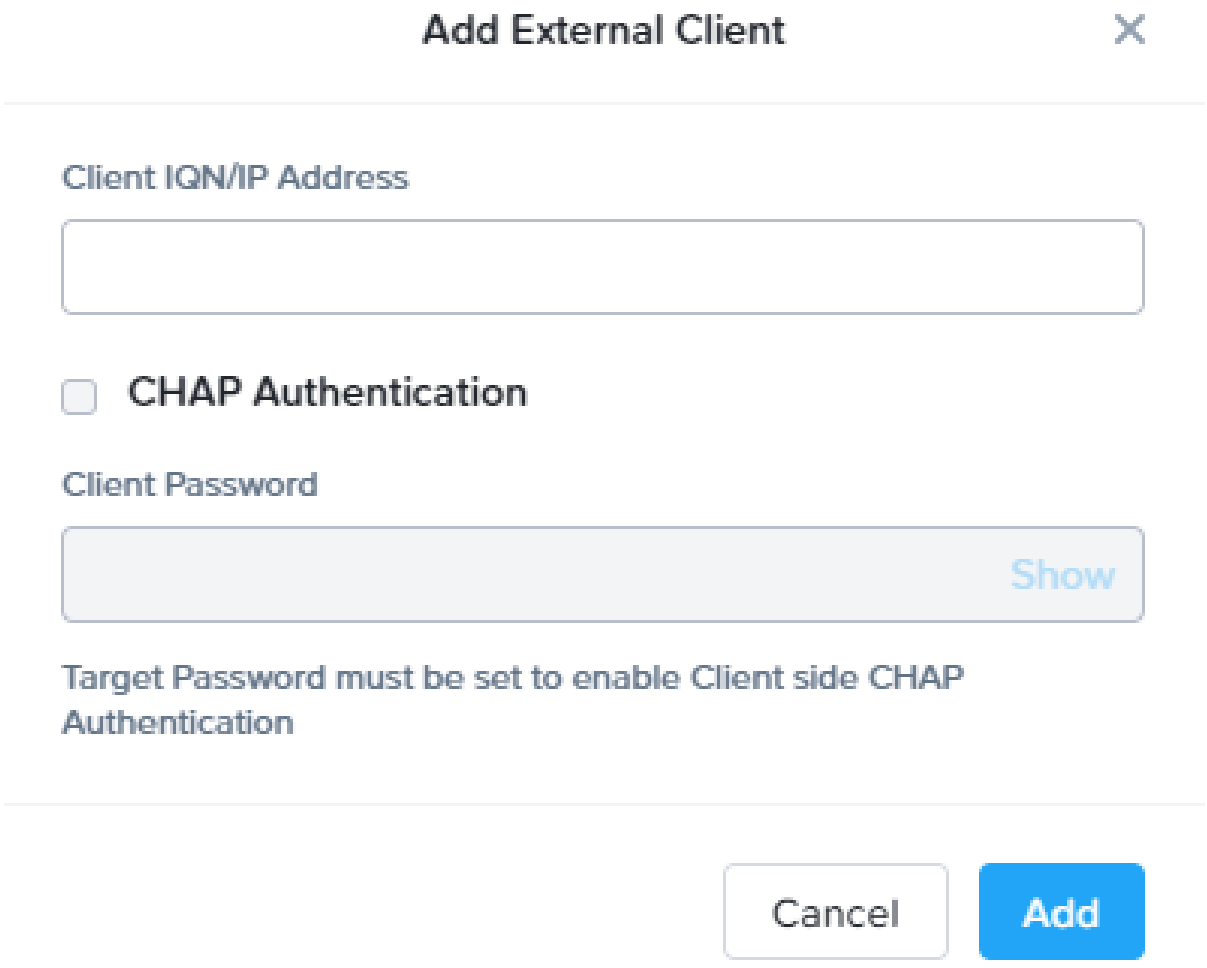

Figure 314: Add External Client Window

c. Click the **Create** button.

The volume group is created, and the new volume group appears in the list of volume groups.

Modifying a Volume Group

About this task

You can modify a volume group (update or delete) at any time after creating it.

Procedure

- To update a volume group, do the following:

  a. Do one of the following:

     » Go to the **Summary** tab in the details page of the target volume group (see Volume Group Details View on page 151) and click the **Update** button.

     » Select the target volume group in the **List** tab of the volume groups dashboard and then click the **Update** button (see Volume Groups Summary View on page 147).

  b. In the **Update Volume Group** page, update the target parameters and then click the **Save** button.

     The fields are the same as when creating the volume group (see Creating a Volume Group on page 644).

- To delete a volume group, either select the volume group in the **List** tab of the dashboard or go to the **Summary** tab in the details page and then click the **Delete** button. (In the dashboard **Delete** is an option in the **Action** drop-down menu.)

  If there are active connections, you must first remove the connections before deleting the volume group. for more information about how to manage connections, see Managing Volume Group Connections on page 649.

Managing Volume Group Connections

About this task

Connections to external clients are configured when the volume group is created (see Creating a Volume Group on page 644), but you configure authentication and add or remove connections to a volume group at any time.

Procedure

- To manage connections as a whole, do the following:

  a. Do one of the following:

     » Go to the **Summary** tab in the details page of the target volume group (see Volume Group Details View on page 151) and click the **Manage Connections** button.

     » Select the target volume group in the **List** tab of the volume groups dashboard and then select **Manage Connections** from the **Actions** pull-down menu (see Volume Groups Summary View on page 147).

     The **Manage Connections** page appears.

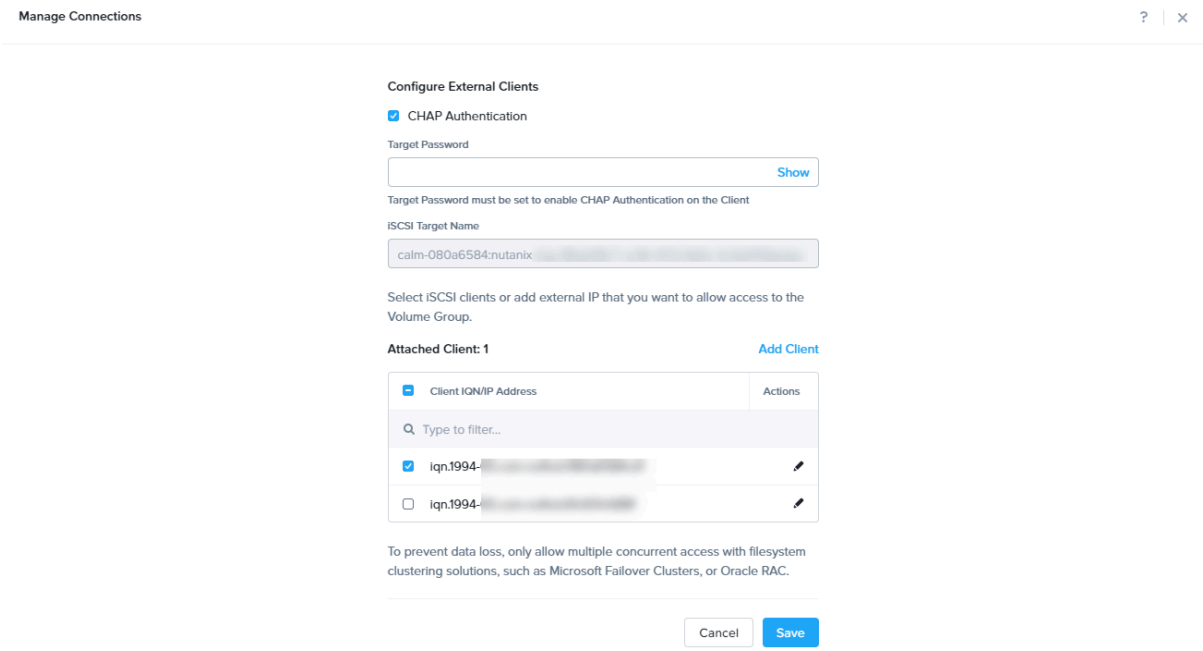

Figure 315: Manage Connections Page

  b. Check the **CHAP Authentication** box to enable the Challenge Handshake Authentication Protocol (CHAP) and then enter the password in the indicated field.

  c. The iSCSI target name appears in the **iSCSI Target Name** field. This is a read-only field because you cannot change the name.

  d. Configure the clients as desired in the **Attached Clients** section. The configured clients appear in the table.

     - To attach (unattach) a client, check (uncheck) the box for that client in the table.
     - To add a client, click the **Add Client** link. In the **Add External Client** box, enter the client IP address or IQN designator name in the indicated field, check the **CHAP Authentication** box and enter the password to enable CHAP authentication, and then click the **Add** button. The added client appears in the list.
     - To change the CHAP authentication for a client, click the pencil icon for that client in the table. The **Edit External Client** box appears. Check (uncheck) the box to enable (disable) CHAP authentication and then click the **Save** button. (You cannot change the client address.)

  e. Click the **Save** button.

- To add an external client connection to a volume group, do the following:

  a. Go to the **Connections** tab in the details page of the target volume group (see Volume Group Details View on page 151) and then click the **Add Connection** button.

  b. In the **Add External Client** window, do the following in the indicated fields:

  - **Add External Client**: Enter the client IP address or IQN designator name
  - **CHAP Authentication**: Check the box to enable CHAP authentication and enter the client password in the indicated field.
  - Click the **Add** button.

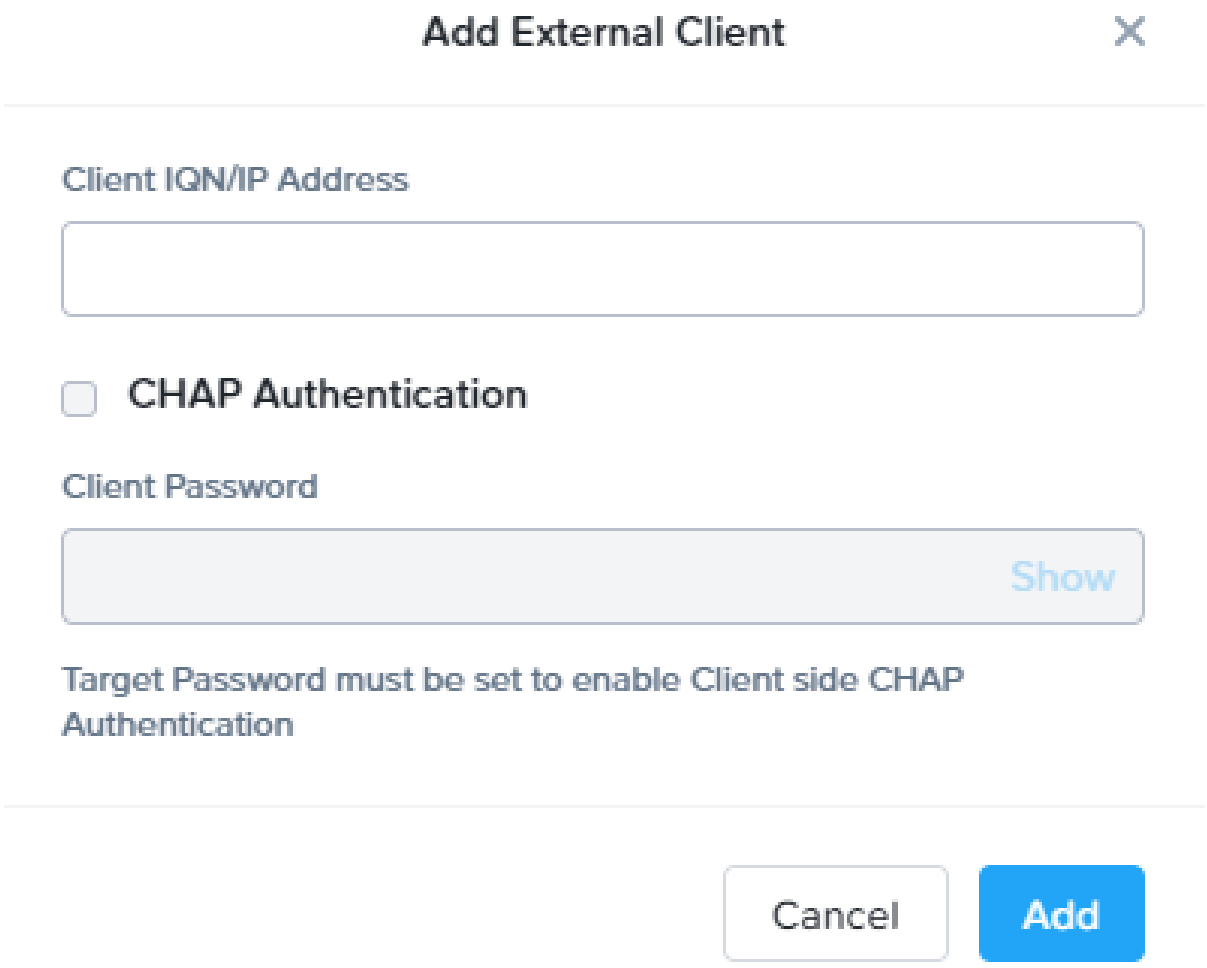

Figure 316: Add External Client Window

The connection is added, and the new client appears in the list of connections for the volume group.

- To update a client connection, do the following:

  Note: A separate CHAP authentication setting for an individual client is not allowed when CHAP authentication is not enabled for the volume group, and you cannot change the client address.

  a. Go to the **Connections** tab in the details page of the target volume group, select the target connection, and then click the **Update** button.

  b. In the **Edit External Client** window, uncheck the **CHAP Authentication** box to disable or check to enable. If you check the box, enter the client password in the indicated field.

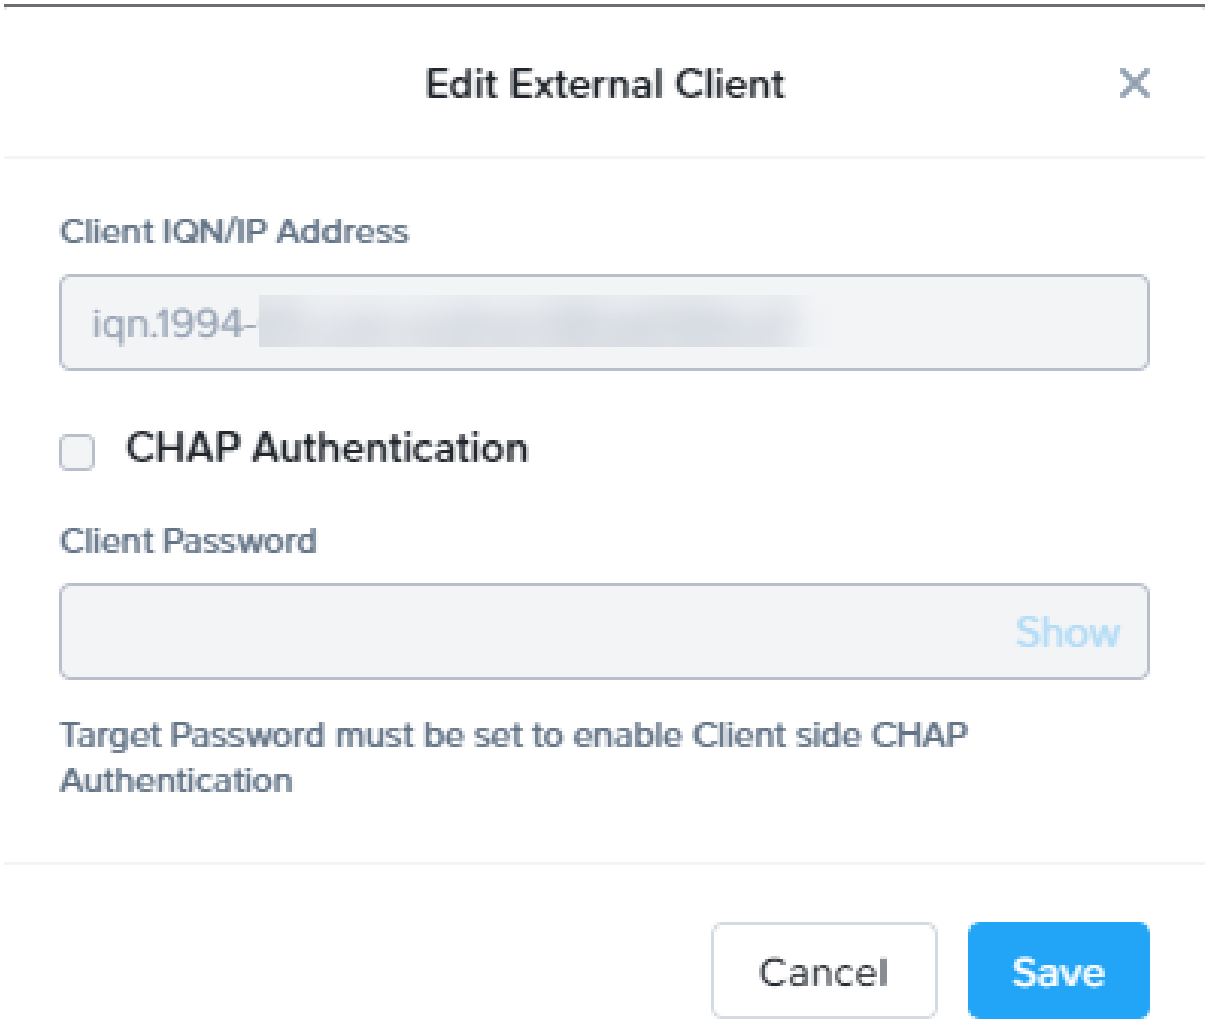

Figure 317: Edit External Client Window

- To remove a client connection, select the **Connections** tab in the details page of the target volume group, select the target connection, and then click the **Remove** button.

  The connection is removed, and the entry is deleted from the list of connections.

Managing Volume Group Virtual Disks

About this task

Virtual disks are added when the volume group is created (see Creating a Volume Group on page 644), but you add, update, or delete disks for a volume group at any time.

Procedure

- To add a virtual disk to a volume group, do the following:
  a. Go to the **Virtual Disks** tab in the details page of the target volume group (see Volume Group Details View on page 151) and then click the **Add Virtual Disk** button.
  b. In the **Add Virtual Disk** window, do the following in the indicated fields:
     - **Storage Container**: Select the target storage container from the pull-down list.
     - **Size**: Enter the desired size (in GiBs) for the virtual disk.
     - Click the **Add** button.

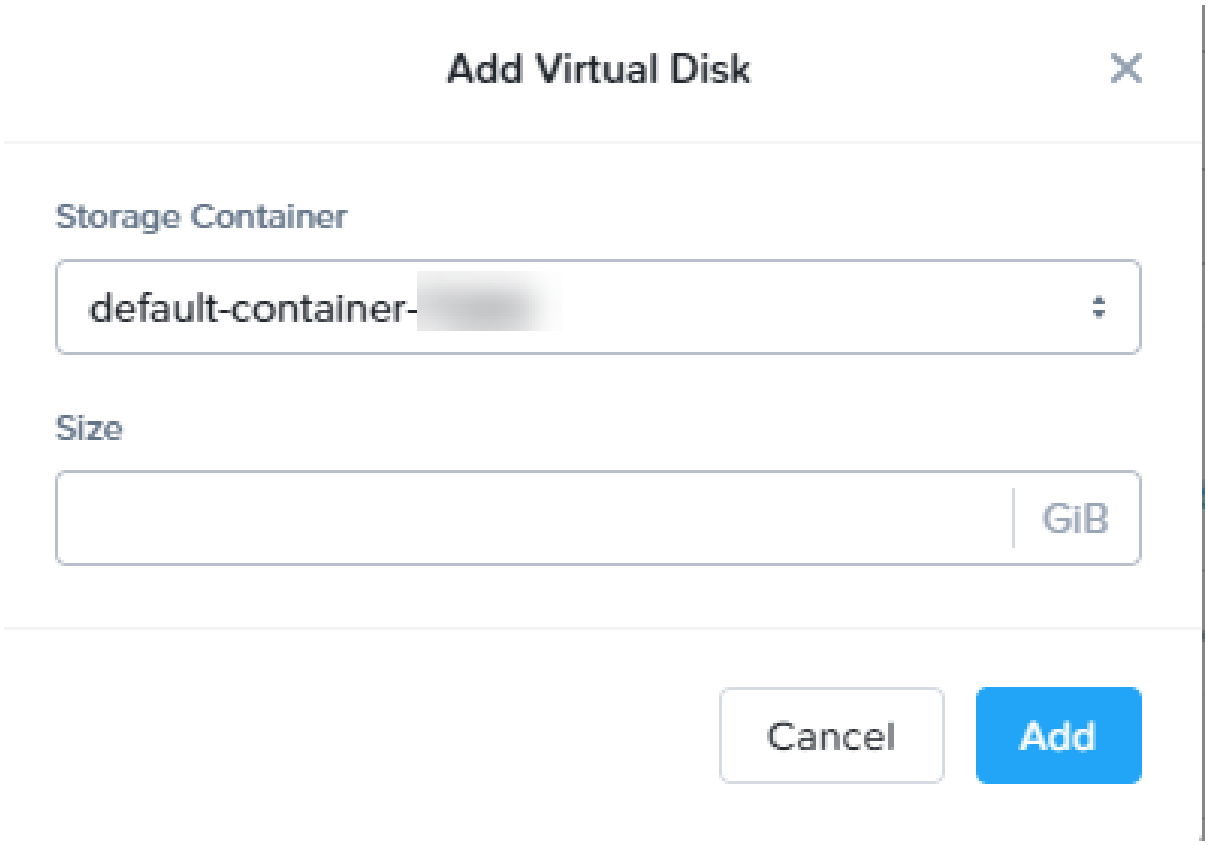

Figure 318: Add Virtual Disk Window

The disk is created, and the new disk appears in the list of virtual disks for the volume group.

- To update a virtual disk, do the following:

  Note: You can only change the disk size, not the storage container.

  a. Go to the **Virtual Disks** tab in the details page of the target volume group, select the target disk, and then click the **Update** button.

  b. In the **Edit Virtual Disk** window, update the disk size (in GiBs) as desired and then click the **Save** button.

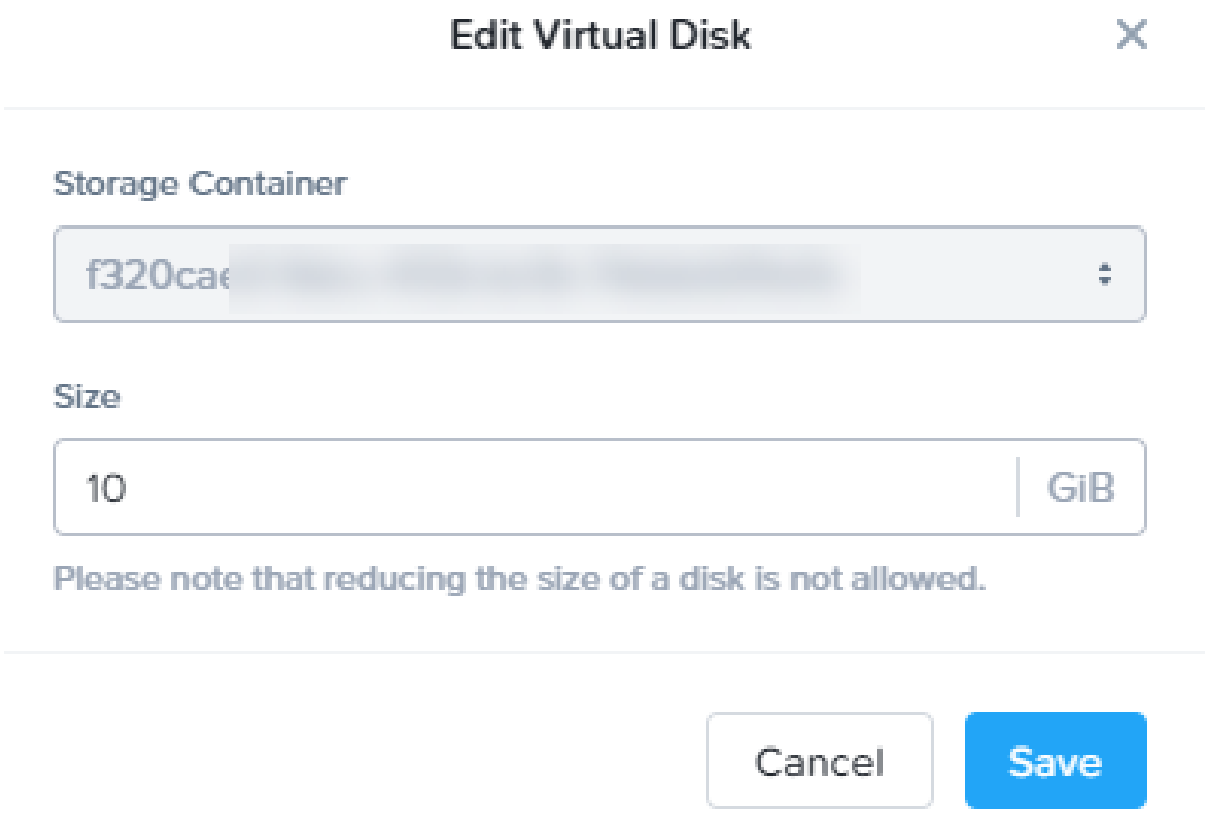

Figure 319: Edit Virtual Disk Window

- To delete a virtual disk, select the **Virtual Disks** tab in the details page of the target volume group, select the target disk, and then click the **Delete** button.

Cluster RBAC for Volume Group

Cluster role-based access control (RBAC) for Volume Group feature enables a super-admin user to provide Prism Admin and Prism Viewer roles access to one or more clusters registered with Prism Central. A user with Prism Admin role can view and update the entities like volume groups, virtual disks, and storage containers from the allowed clusters. However, a user with Prism Viewer role can only view the entities.

Cluster RBAC is currently supported on an on-prem Prism Central instance hosted in a Prism Element cluster running AHV. After you enable the Micro Services Infrastructure feature on Prism Central, the Cluster RBAC feature is then automatically enabled.

Cluster RBAC for Volume Group feature is supported on AHV and ESXi clusters.

Table 738: List of Permissions for Prism Admin and Prism Viewer Roles

Role	Privileges
Prism Admin	Full administrator privileges except for creating or modifying the user accounts
Prism Viewer	View-only privileges

Configuring Cluster RBAC for Volume Group

About this task

To configure Cluster RBAC for Volume Group in Prism Central for users or user groups, do the following.

Procedure

1. Log on to Prism Central as an admin user or any user with super admin access.
2. Configure Active Directory settings.

   > **Note:** You can skip this step if an Active Directory is already configured.

   Go to **Prism Central Settings** > **Authentication**, click **+ New Directory** and add your preferred Active Directory.
3. Click the hamburger menu and go to **Administration** > **Roles**.
   The page displays system defined and custom roles.
4. Select **Prism Admin** or **Prism Viewer** role, then click **Actions** > **Manage Assignment**.

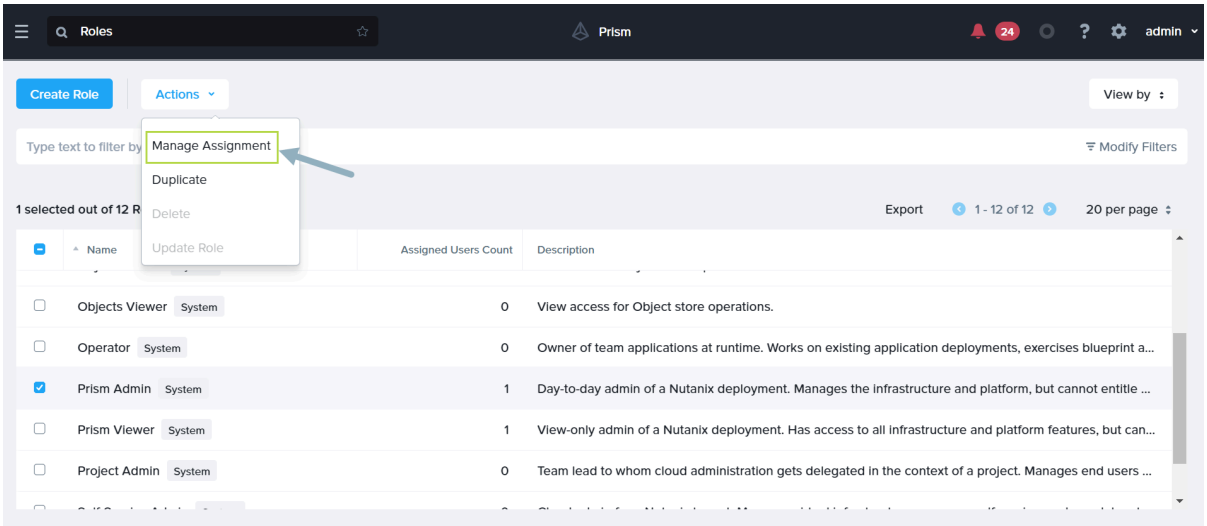

Figure 320: Prism Central Roles

For illustration purpose, the Prism Admin role is selected in this step.

5. Click **Add New** to add a new user or user groups to this role.

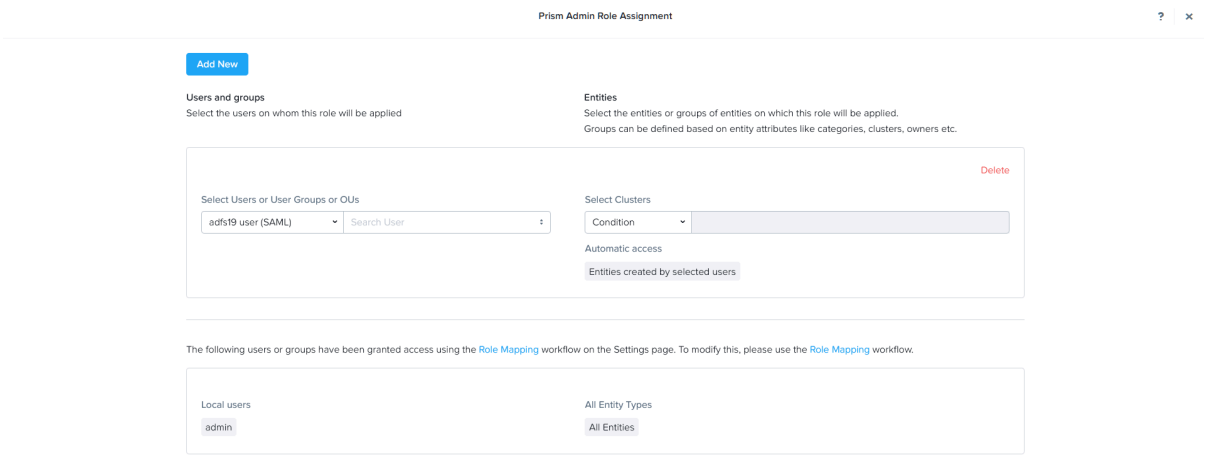

Figure 321: Role Assignment

You will add users or user groups and assign clusters to the Prism Admin or Prism Viewer role in the upcoming steps.

6. In the **Select Users or Groups** field, do the following:

   a. Select the configured active directory (AD) from the drop-down.

   b. Search and add the users or user groups.

   To search a user or user group, start typing few letters and the system will automatically suggest the names.

7. In the **Select Clusters** field, you can provide cluster access to AD users or User Groups using the **Individual entity** option (one or more registered clusters) or **ALL Clusters** option.

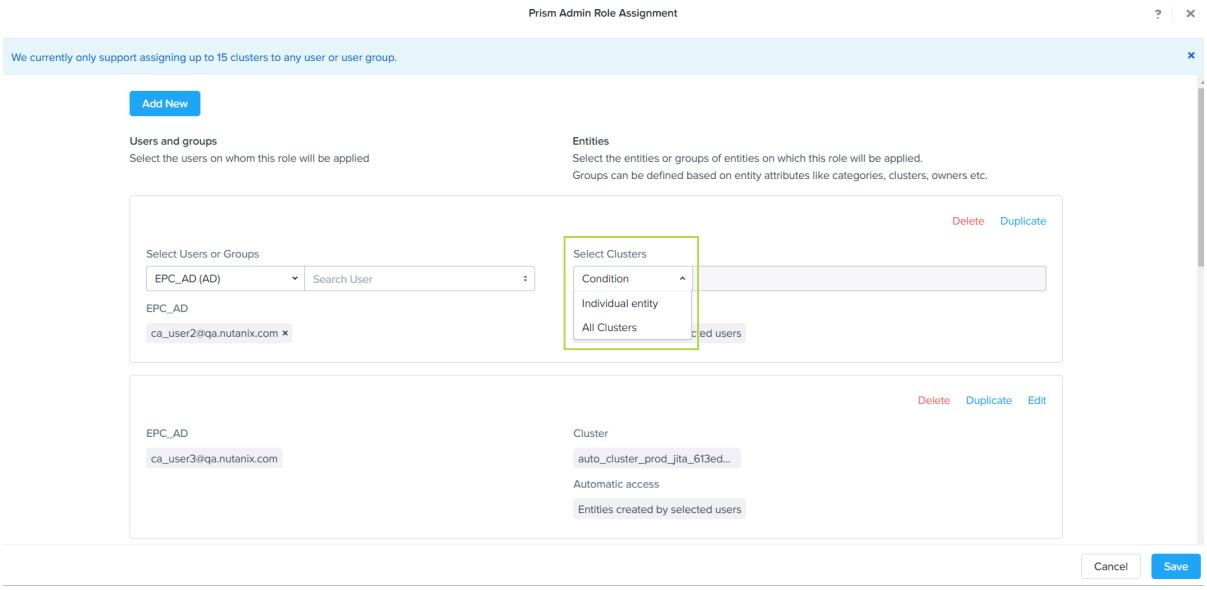

Figure 322: Select Clusters

8. Click **Save**.
   AD users or User Groups can log on and access Prism Central as a Prism Admin or Prism Viewer. They can view or act on the available entities in the configured clusters such as volume groups, virtual disks, and storage containers.

Storage Policy Management

Storage Policies in Prism Central let you manage the storage attributes like encryption, compression and QoS of entities like VMs. A single storage policy can manage the attributes of several entities that are associated with various Categories.

You can create a storage policy in Prism Central. A storage policy uses Categories to associate with entities such as VMs. Using a storage policy, you can manage parameters of entities (VMs) such as encryption, type of values for data compression, and IOPS or Throughput throttling to be applied to the entities.

To create and manage storage policies, open the **Storage Policies** dashboard (see Storage Policies Summary View on page 141).

Requirements

Prism Central supports storage policies only on AHV. Prism Central supports storage policies only for VMs. It does not support storage policies for Volume Groups.

To successfully complete the inbuilt pre-checks and create a storage policy, ensure that:

1. You provide a unique name for every storage policy. In other words, no two storage policies can have the same name. If you use the name of an existing policy, Prism Central displays an error.
2. You provide a *defined* (non-default) value for at least one storage policy attribute or property. For example, for **Compression** the default value is **Inherit from Cluster**, so you may instead select **On** with **Inline** as the value. You could select or provide a defined value for any one ore more properties of the storage policy.

   See Creating or Updating a Storage Policy on page 660 for list of properties that you can configure during creation or updating of a storage policy.
3. The throttled IOPS value is greater than 99 (to enable rate limiting or throttling) or equal to -1 (to clear the throttling) in the storage policies.

> **Note:** The storage policy engine completes the pre-checks in the same order as stated above. For example, if you repeat an existing storage policy name for a new storage policy, the policy engine fails the storage policy creation, generates the failure message stating Duplicate name and stops other checks. The engine triggers the storage policy creation only after all the pre-checks are validated.

Limitations

> **Important:**
>
> The storage policy engine takes 30 minutes or more to apply a newly created storage policy to the associated entities like VMs. The time taken depends on the number of VMs associated with the storage policy.
>
> Any changes you make to values for the storage attributes, like compression, encryption or throttling, for any storage policy that is already applied, take effect immediately.

- Prism Central does not support storage policies for entities running on ESXi and Hyper-V.
- Prism Central does not support storage policies for Nutanix Disaster Recovery with on-prem recovery AZ (on-prem recovery cluster managed by a Prism Central or Nutanix Disaster Recovery with multiple Prism Central deployments, see Nutanix Disaster Recovery Guide for more information about on-prem AZ).

- Nutanix supports storage policies for Nutanix Disaster Recovery with single Prism Central. However, the storage properties defined for the container are applied to the snapshots replicated to the replication cluster. In other words, the snapshot data replicated to the remote site is transformed based on the settings defined for the remote site containers. The storage policies are applied only after the VM failover.
- Storage policies use categories to associate with entities like VMs. A storage policy cannot associate directly with an entity like a VM. Add an appropriate category to the VM and then associate the category with the storage policy.

  > **Note:** Prism Central does not support storage policies for volume groups (VGs).

  When you associate a category with a storage policy, all the VMs that are associated with that category are associated with the storage policy. The storage policy, thus, applies to all the VMs associated with that category.
- You can associate multiple categories with one storage policy. However, you can associate one category with only one storage policy. After associating a category with one storage policy, you cannot associate the same category with another storage policy. When you associate the same category with another storage policy, Prism Central adds the category to the **Associated Categories** list but displays an alert (an exclamation mark) and disables the **Save** button until you remove the category. The following message pops up when you hover on or click the alert (exclamation mark):

  ```
  This category has an existing storage policy (<name of storage policy>)
   associated with it.
  ```

- You can create a maximum of 100 storage policies irrespective of the type of Prism Central deployed (whether XLarge, Large or Small PC with Scale-out or Single Instance).
- You cannot create a storage policy for a snapshot. A storage policy applies to a live vDisk and the entire chain of snapshots of that vDisk.
- The snapshots of the vDisk that are in the storage are not encrypted if the live vDisk is deleted before AOS applies the storage policy to the vDisk.
- Storage policies support entity-centric encryption only if you did not already enable encryption on the underlying clusters in the deployment.
- After you enable encryption using storage policies, you cannot disable it.

  You can, however, delete the storage policy that has encryption enabled. In such cases, any new data written into the entities is not encrypted. The data that was written prior to the deletion of the storage policy remains encrypted.
- If two VMs share data and one or both the VMs have a storage policy with encryption enabled, then the shared data is encrypted.
- If two VMs share data and one of the VMs have a storage policy with compression enabled, even then the shared data may not be compressed.
- If you associate a VM with multiple storage policies associated, the policy engine sorts the storage policy list using a method based on the policy name and applies the highest ranking policy from the sorted list to the VM. The highest ranking storage policy is the one with the lexicographically smallest name in the list. If you add another storage policy later and this newly added policy becomes the highest ranking policy in the newly sorted list, the engine continues to apply the previously applied storage policy. To apply the newly added storage policy to the VM, remove the association of the VM with the previously applied storage policy.
- Storage policy inheritance is supported for operations on VMs performed using API v3. This support is not available for operations performed using other API versions like API v2.

For information about Storage Policy summary and detailed views, see Storage Policies Summary View on page 141.

Storage Policy Based Encryption

Storage policies allow you to make Encryption-at-Rest choices at the VM level.

Deployments can continue to have Nutanix Encryption-at-Rest capabilities scoped for the entire cluster. Storage policy provides the additional option to control the encryption scopes decisions at the VM.

> **Note:** Prism Central does not support Policy based encryption storage containers on which Erasure Coding or Deduplication is enabled.

The Key Manager System (KMS) choice is configured as a cluster-wide setting, leveraging either Nutanix Built-In Native Key manager or leveraging a third party Key Manager. As part of the Key manager configuration, the scope of encryption-at-rest is defined either as Cluster-wide or specific entity-scoping. Entity-scoping means that the encryption-at-rest scope is defined at the container or the data policy level.

> **Important:**
>
> - When you select the cluster, you can manage the KMS type by clicking **Actions** > **Manage KMS Type**.
> - From Prism Central, if you try to configure storage policy based encryption for any single VM or Category entity without enabling **Entity Encryption** on that cluster, Prism Central does not transform the entity data to be encrypted.
>
>   Log on to Prism Central and do the following:
>
>   - Go to **Hardware** > **Clusters**, select the cluster and click **Actions** > **Enable Data-at-Rest Encryption**.
>   - In the **Data-at-Rest Encryption** window, under **Select Encryption Type**, select **Entity Encryption** and click **Save Encryption Type**.
>   - In the **Data-at-Rest Encryption** window, under **Select Key Management Server (KMS)**, select **Native KMS (Local)** and click **Enable Encryption**.
>   - In the **Data-at-Rest Encryption** window, under **Type SET to Confirm**, enter **SET** and click **Set Encryption**.

When Storage policy is used, AOS data for the VM is encrypted on the disk for all the vDisks attached to the VM. The main features include:

- Once encryption is enabled in a policy, it cannot be disabled. If the VM moves out of the policy or the policy is deleted, then the new writes to the VM are unencrypted and existing data remains encrypted.
- The system automatically generates a unique key for each storage container.
- Multi-tenant requirements having a unique key per tenant can be fulfilled by ensuring that the set of VMs belonging to each tenant is mapped to its own container. Nutanix recommends having only 50 tenants or less per cluster.
- Encryption can be enabled for new and existing VMs.
- Storage policies do not support Volume Group. For any VM with a Policy with encryption enabled, the data stored in the VG is not encrypted.
- Unencrypted VM needs to remain unencrypted during the cloning process.
- Shared data between the encrypted VMs and unencrypted VMs is encrypted.

For information about containers, see the Storage Management section in the *Prism Web Console Guide*.

For information about configuring key manager, see the Security Guide.

Creating or Updating a Storage Policy

You can create or update a storage policy in Prism Central only.

About this task

The **Create Storage Policy** has two tabs - **Configuration** and **Association**. The **Configuration** tab lets you configure the Data Security (Encryption), data Reduction (Compression) and QoS (IOPS or Throughput throttling) aspects of the storage policy. The **Association** tab lets you associate the storage policy with categories (and thence, entities).

The **Create Storage Policy** and the **Update Storage Policy** pages have the same storage policy attribute or property configuration fields.

Do the following on the **Storage Policies** dashboard:

Procedure

1. Click **Create Storage Policy**.
   On the **Create Storage Policy** page, the **Configuration** tab opens.

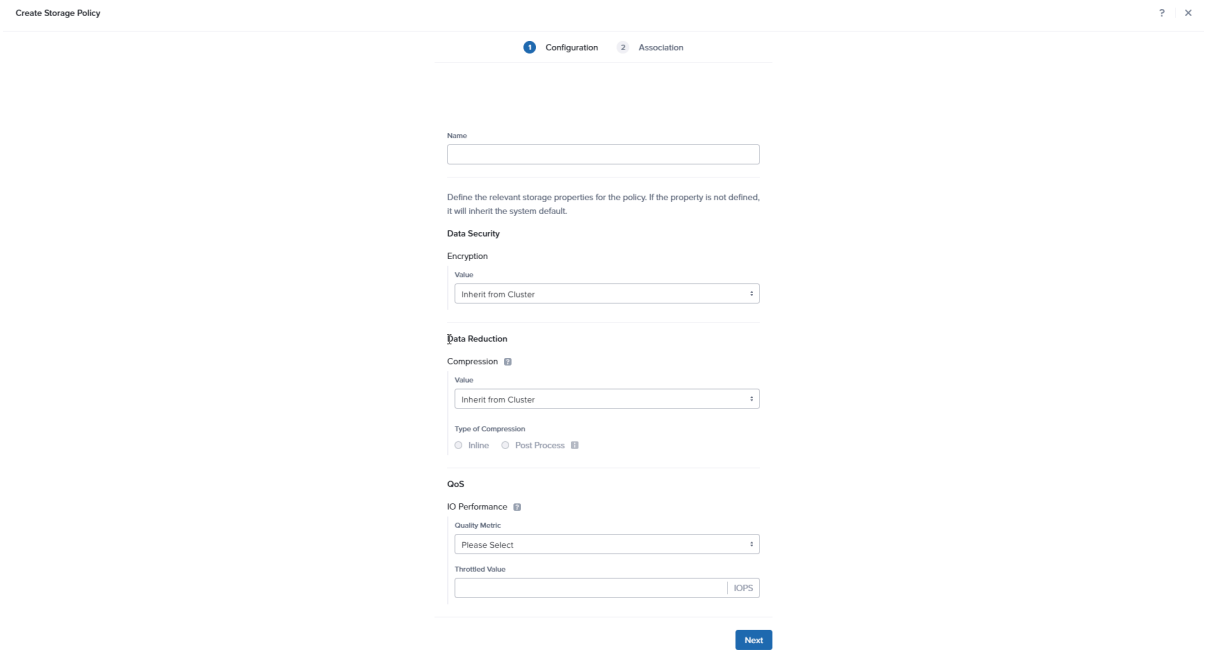

Figure 323: Create Storage Policy - Configuration Tab

Configure the following properties in the **Configuration** tab. Select a defined (non-default) value for a least one of the properties.

Table 739: Configuration Tab Fields

Field	Description and Values
Name	Enter a unique name for the storage policy. **Note:** Do not repeat a name previously used for another storage policy. The built-in pre-check fails if you enter a non-unique name.

Field	Description and Values
Data Security	
Encryption Value	Select one of the following values from the drop-down list. • **Enabled**—to enable encryption through the storage policy. **Note:** After you enable encryption using storage policies, you cannot disable it. • **Inherit from Cluster** (default)—the category or entity inherits the encryption configuration from the cluster. This is the default value for Encryption.
Data Reduction	
Compression Value	Select one of the following values from the drop-down list. • **On**—to enable compression. If you select **On** then you must select the type of compression that you want to apply. You can select **Inline** compression (default) or **Post Process** compression. If you select **Post Process**, the data is compressed with some delay (up to 3600 seconds) that is set by the system default. • **Off**—to disable compression. • **Inherit from Cluster** (default)—the category or entity inherits the compression configuration from the cluster. This is the default value for compression.
IO Performance (Also see the *Relationship between Block Size, IOPS and Throughput* table for information about how Throttle IOPS or Throughput is calculated based on block size)	
Quality Metric	Select one of the following values from the dropdown list. • **Throughput** • **IOPS**
Throttled Value	Provide a numeric value. **Note:** For **IOPS** the value must be greater than 99 or equal to -1 (minus 1).

Throttle IOPS or Throughput Relationship: Set the Throttle IOPS or Throughput based on the block size. Nutanix uses a fixed block size of 32k. For example, with a block size of 32K, if you set a Throttle IOPS value of 1000 IOPS, the Throttle Throughput would be 31.25 MBps. See the table.

Table 740: Relationship between Block Size, IOPS and Throughput

Block Size	IOPS	Throughput (MBps)
8k	1000	7.81
16K	1000	15.63
32K	1000	31.25

Block Size	IOPS	Throughput (MBps)
64K	500	31.25
128K	250	31.25

2. Click **Next**.

   If you did not select or enter an appropriate value and the pre-checks failed, the relevant error is displayed below the field that failed the pre-check. Select or enter an appropriate value and click **Next** again.

3. On the **Association** tab, in the **Add Category** field, click the category or categories you want to associate with the storage policy.

   Enter the first 2-3 characters of the category name you want to associate with in the **Add Category** field. This opens a drop-down list of categories that you can select from. Click 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.

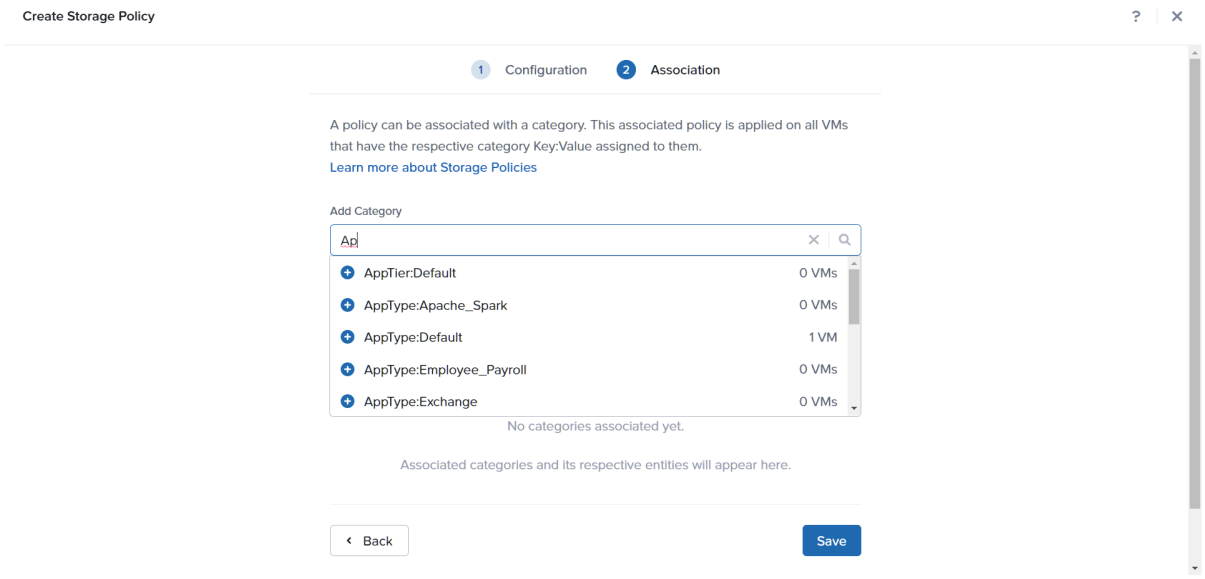

Figure 324: Create Storage Policy - Association Tab

   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 662.

   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.

Before you begin

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. For more information, see *Limitations* section in Storage Policy Management on page 657.

> **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.

About this task

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

Procedure

1. Log in to Prism Central.
2. Navigate to **Compute & Storage** > **Storage Policies**.
   The system displays the **List** tab by default.
3. Click the name of the storage policy. In the storage policy details page, select the **Entities** > **Realized** tab.
4. On the **Entities** page, click **Manage Associations**.
   In the **Manage Associations** page, you can add or remove category associations.
5. 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.

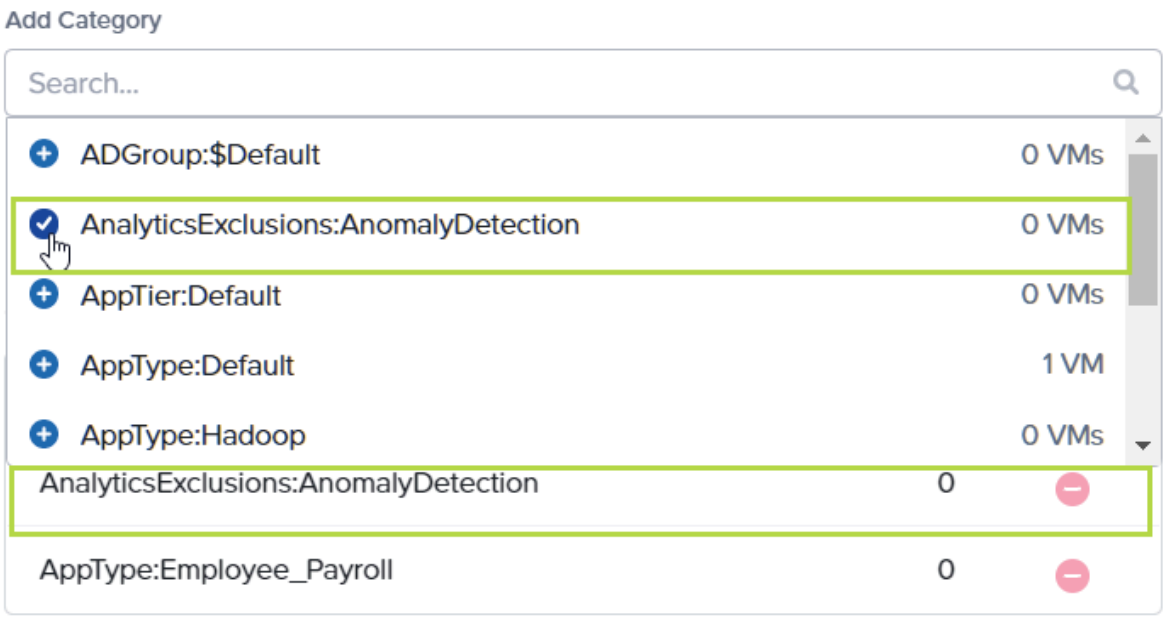

Figure 325: Add Category

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

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

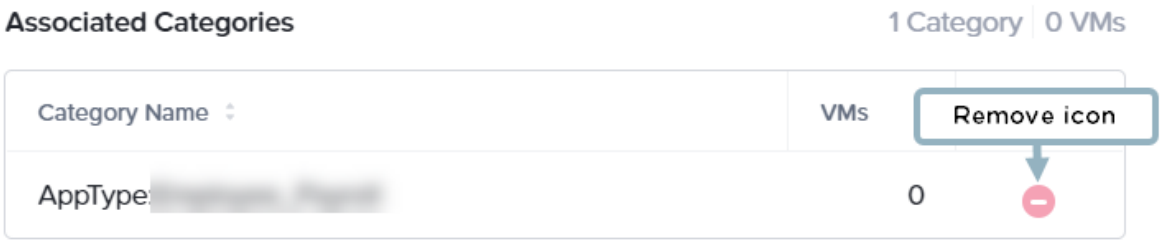

Figure 326: Remove Categories

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 662 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 662 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:

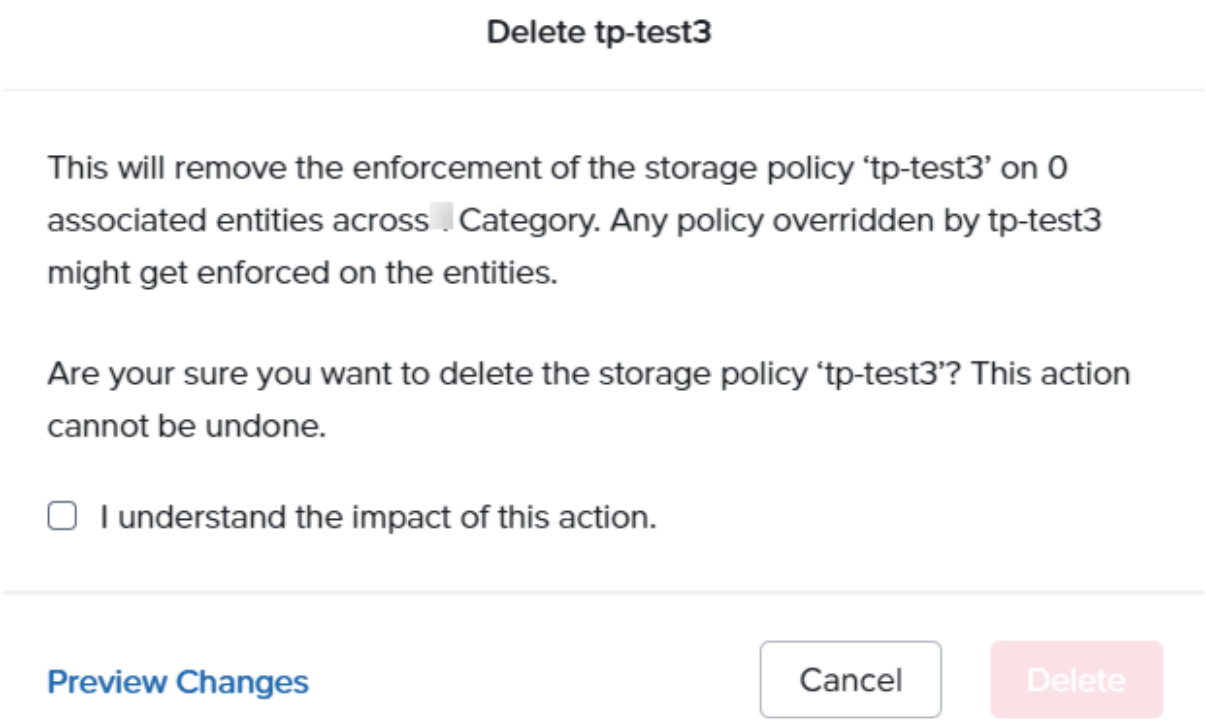

Figure 327: Delete Storage Policy

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.

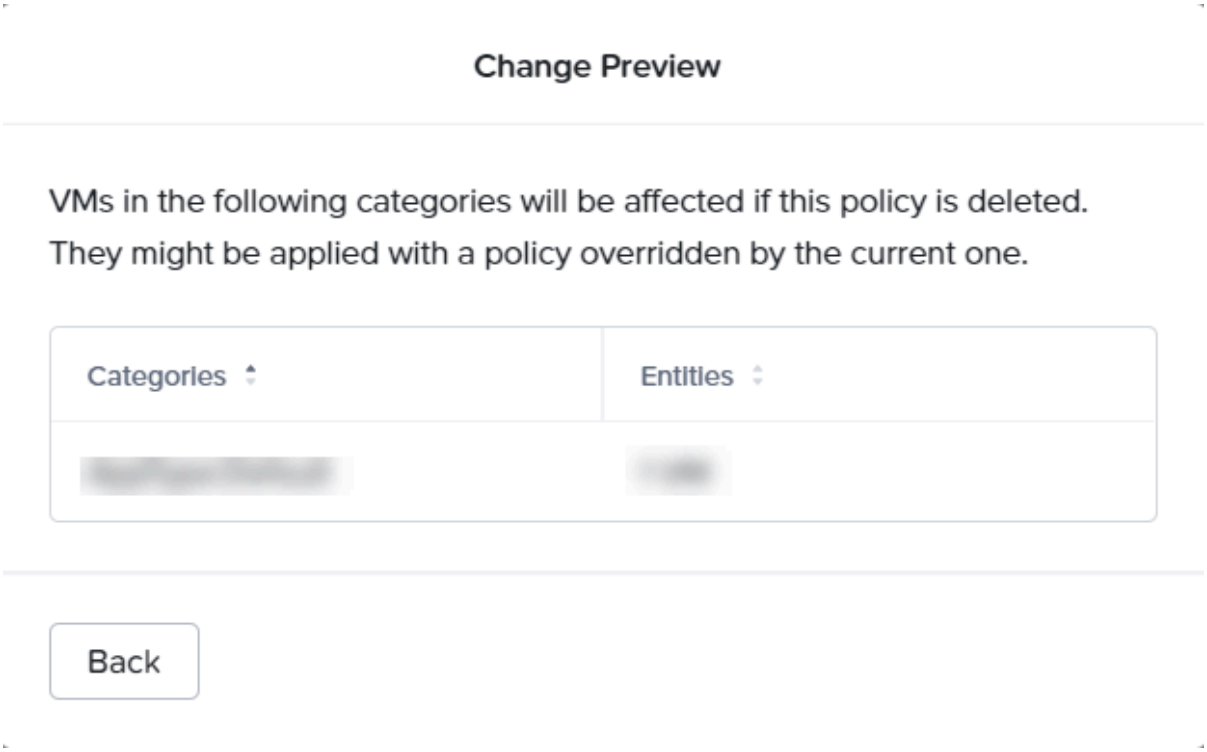

Figure 328: Change Preview

   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 723).
- 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 94).
   - » Go to the images dashboard and select the image that you want to add to the catalog (see Images Summary View on page 125).

2. To add an image, select **Add Image to Catalog** from the **Actions** menu and do the following in the **Add Image to Catalog** window in the indicated fields:

   a. **Name**: Enter a name for the image.

   b. **Description** (optional): Enter a description of the image.

   c. Click the **Save** button.

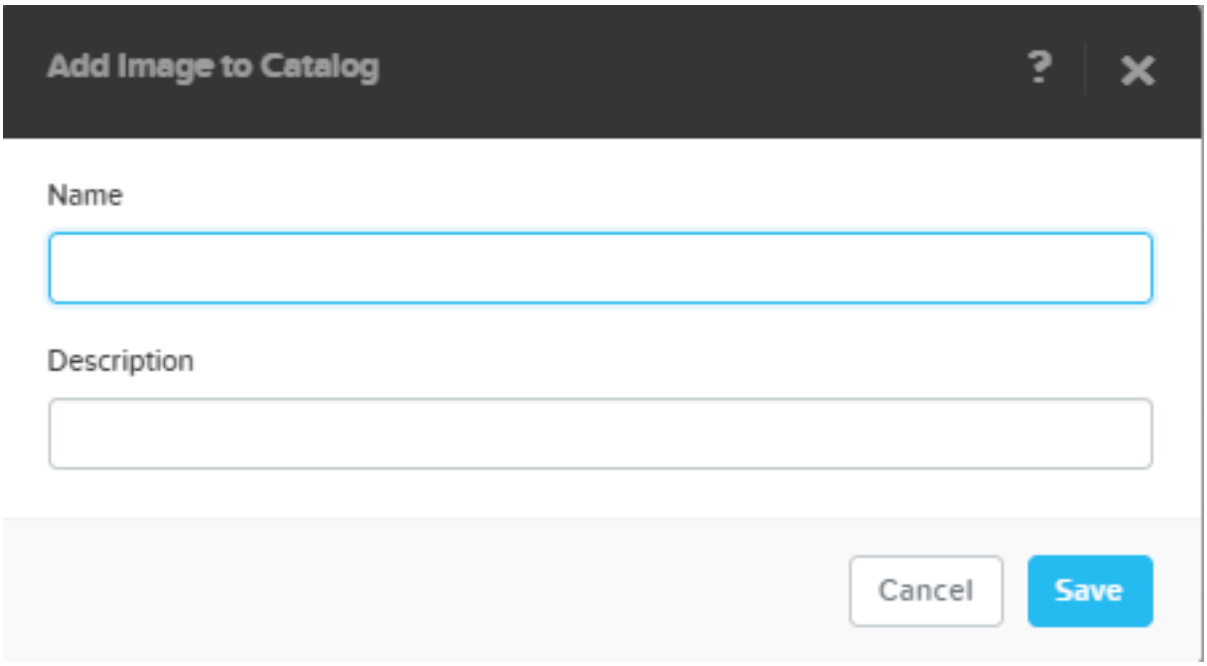

Figure 329: Add Image to Catalog Window

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 (Linux) 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.

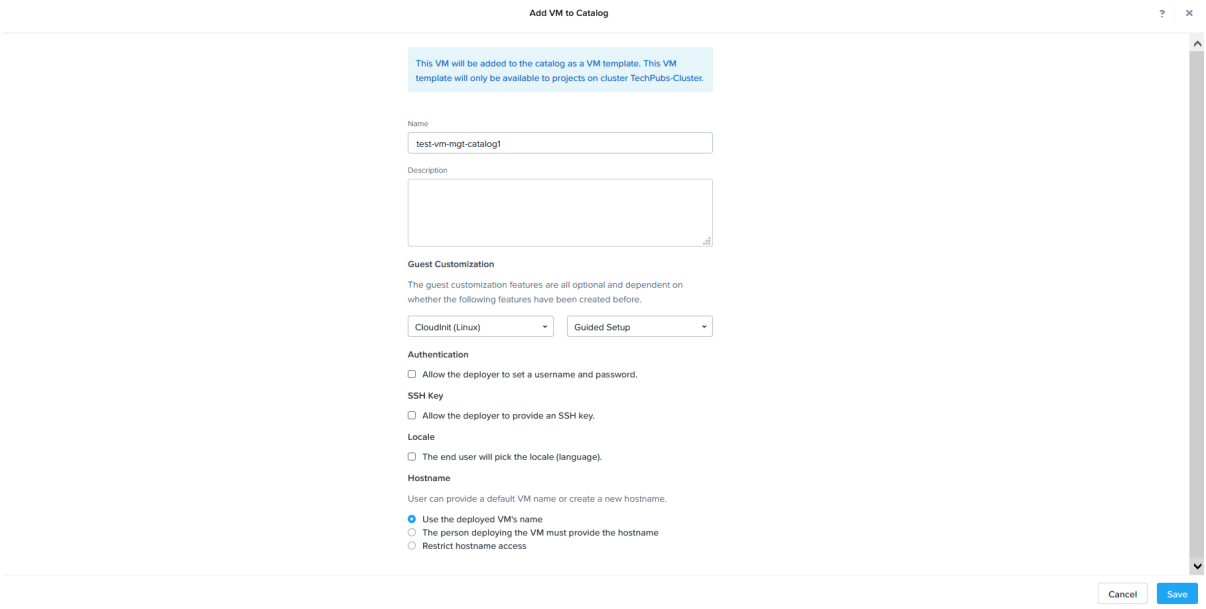

Figure 330: Add VM to Catalog Page

Deleting a Catalog Item

About this task

VM snapshots and images deleted from a catalog makes those items unavailable to project members.

Procedure

To delete a VM snapshot or image from the catalog, do the following:

1. Log in to Prism Central.
2. Navigate to **Compute & Storage** > **Catalog Items** from the entities menu.

3. Select the target VM snapshot or Image catalog item, and click **Delete Catalog Item**.
   The selected catalog item is deleted from all the projects in the cluster.

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 690).

Figure 331: 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 671.

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 690.

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 672 and Uploading Images from a Remote Server on page 677, 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 689.

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 you upload images from a workstation and place the images on multiple clusters during the initial image upload, the uploaded image becomes active (in Prism Element) on only one of the selected clusters and remains inactive on the other clusters registered to Prism Central.
- If you add a new cluster that contains images to an existing Prism Central, you must first import the images to Prism Central and then use the image placement policies to push the images to any other cluster already registered with 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 672.
- Uploading Images from a Remote Server. See Uploading Images from a Remote Server on page 677.
- Uploading Images from a VM disk. See Uploading Images from a VM Disk on page 683.

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 125).
   The **Add Images** page appears.

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

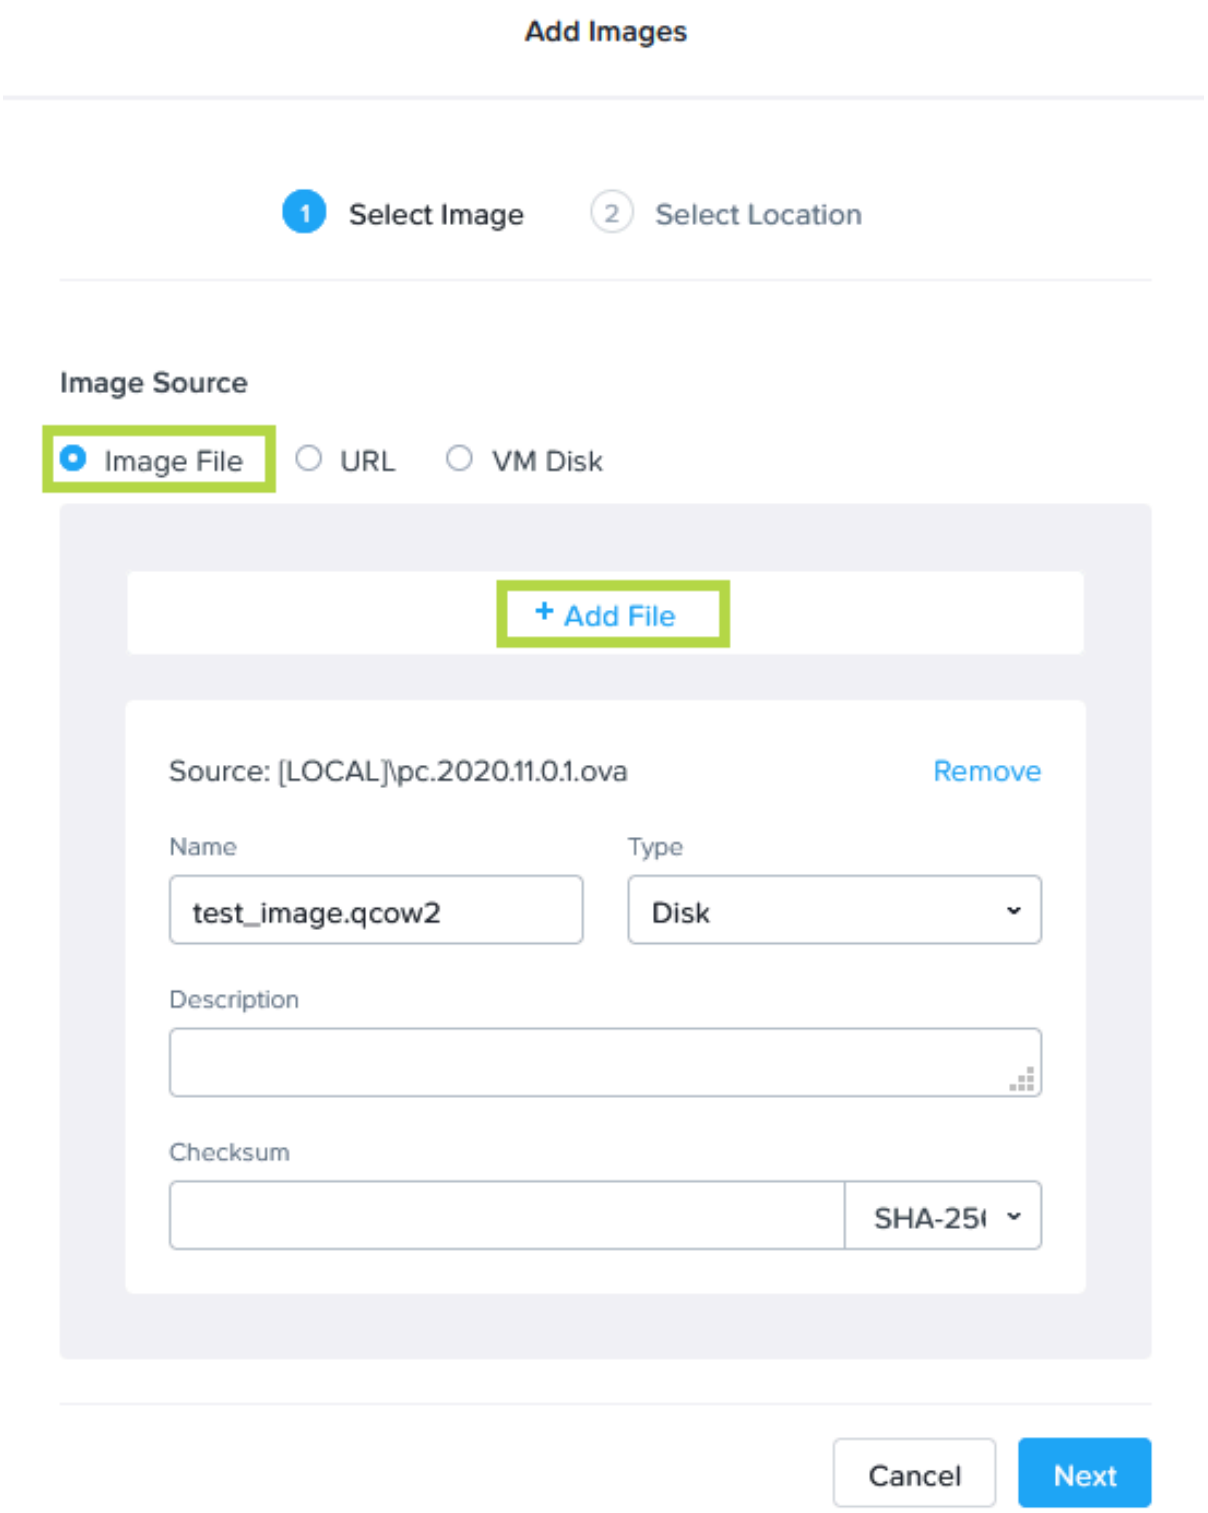

Figure 332: Add Images: File

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.

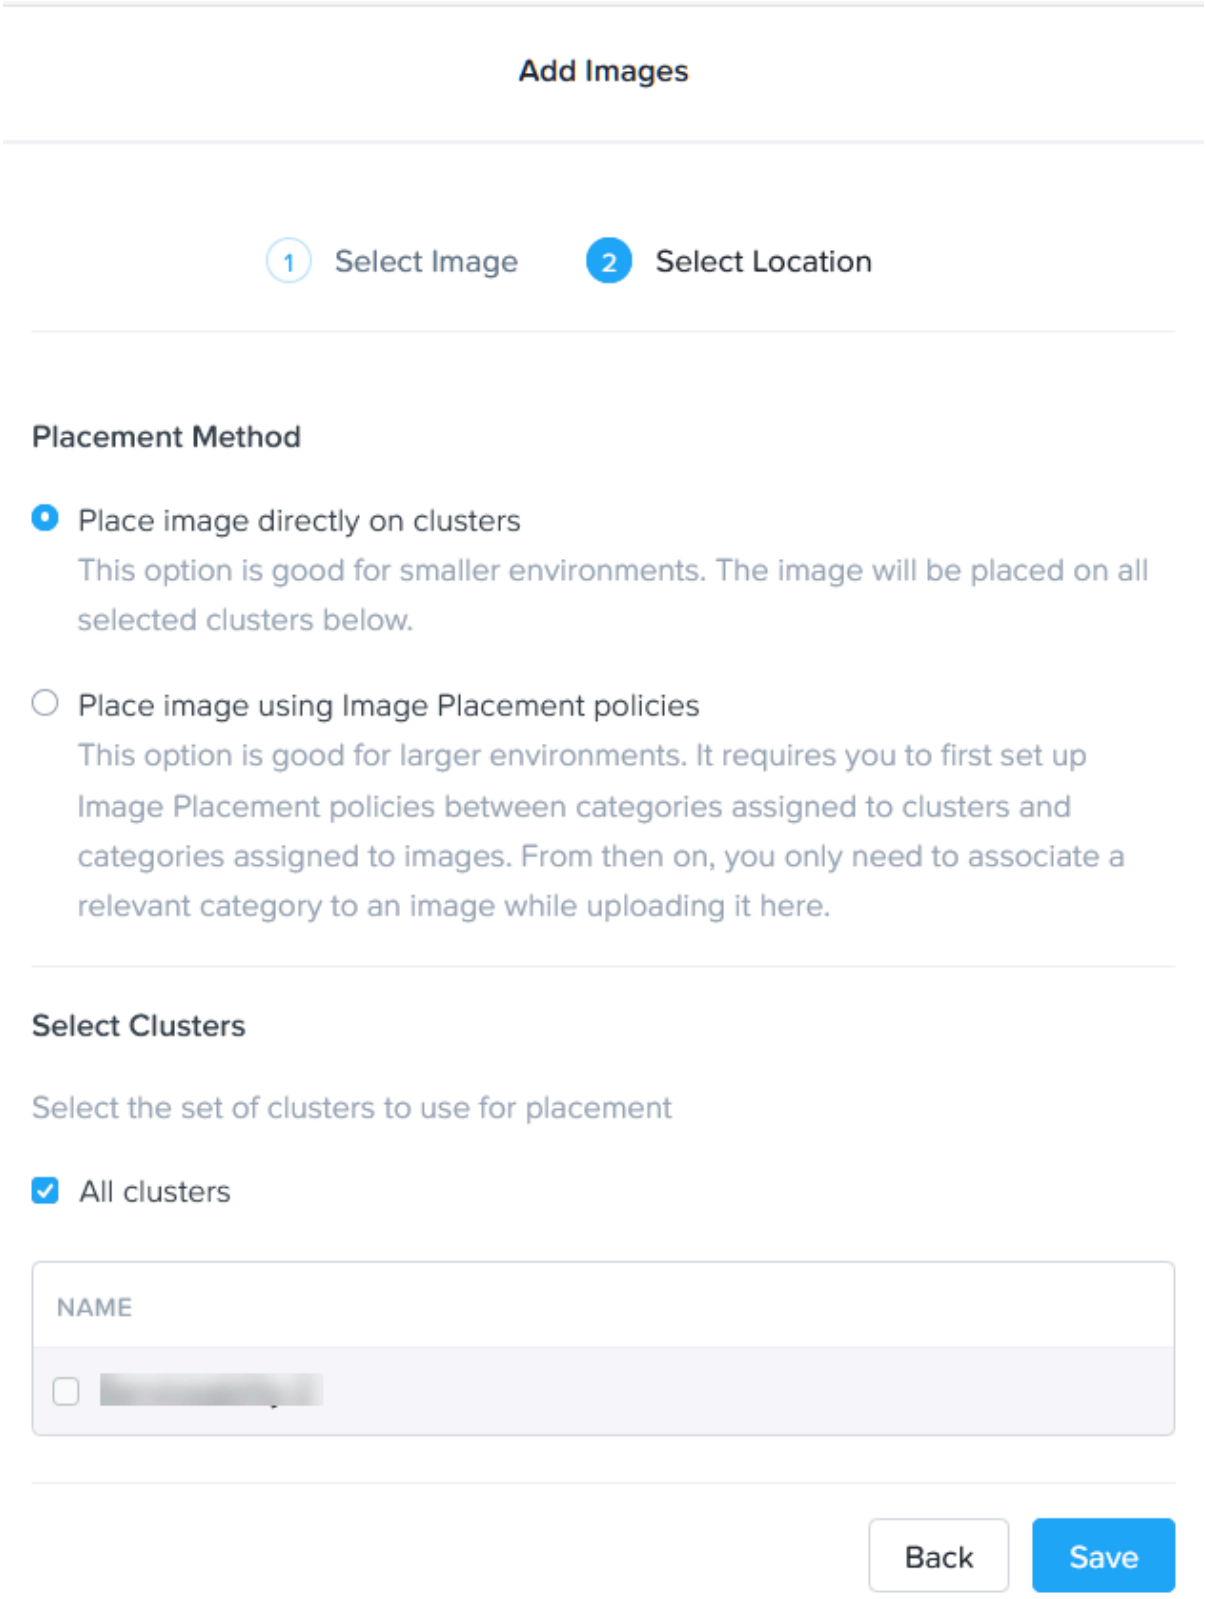

Figure 333: Add Images: Select Location

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 125).
   The **Add Images** page appears.

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

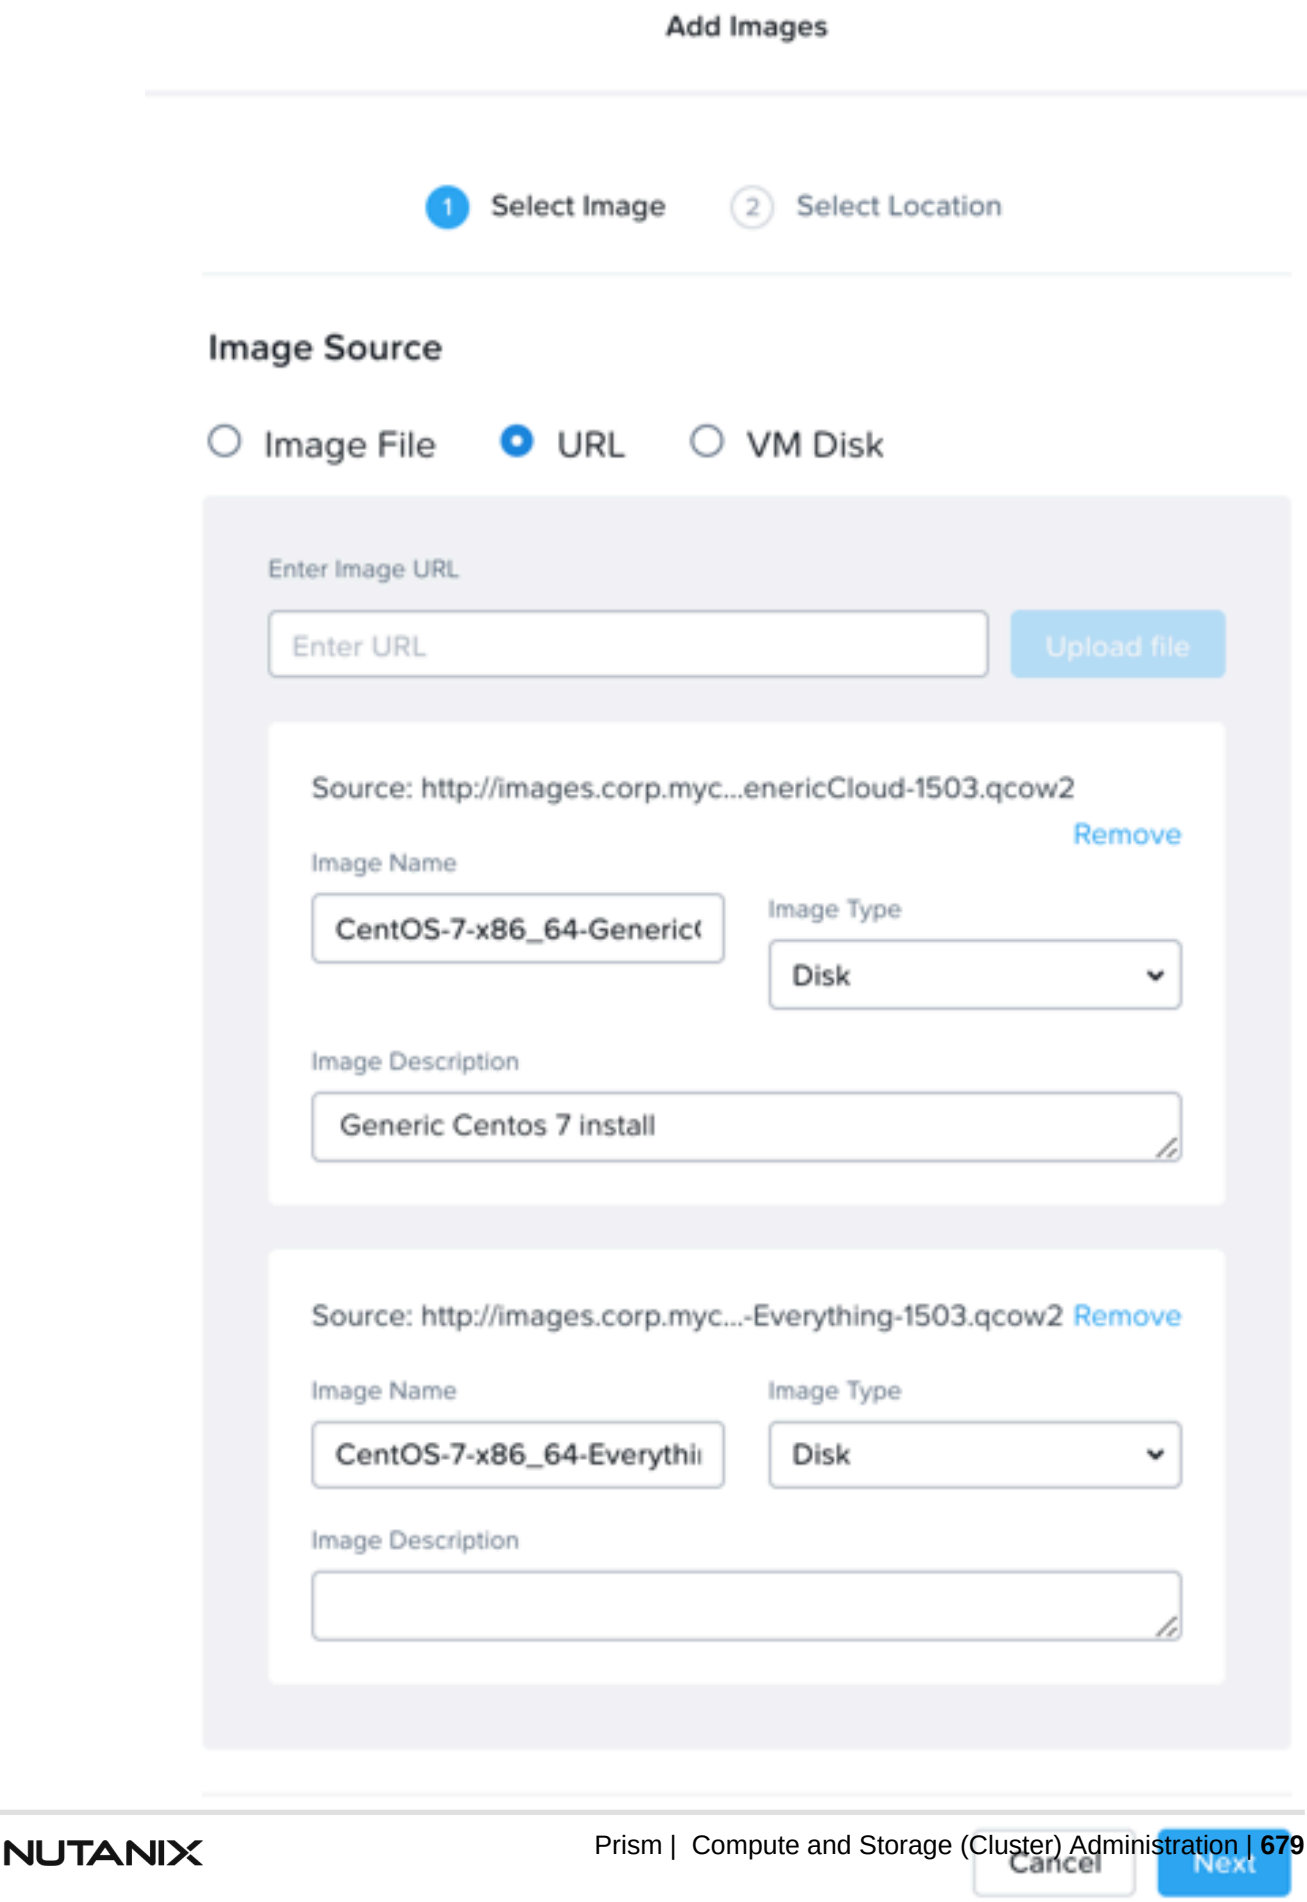
Add Images
1 Select Image
2 Select Location
Image Source
Image File
URL
VM Disk
Enter Image URL
Enter URL
Upload file
Source: http://images.corp.myc...enericCloud-1503.qcow2
Remove
Image Name
CentOS-7-x86_64-Generic(
Image Type
Disk
Image Description
Generic Centos 7 install
Source: http://images.corp.myc...-Everything-1503.qcow2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

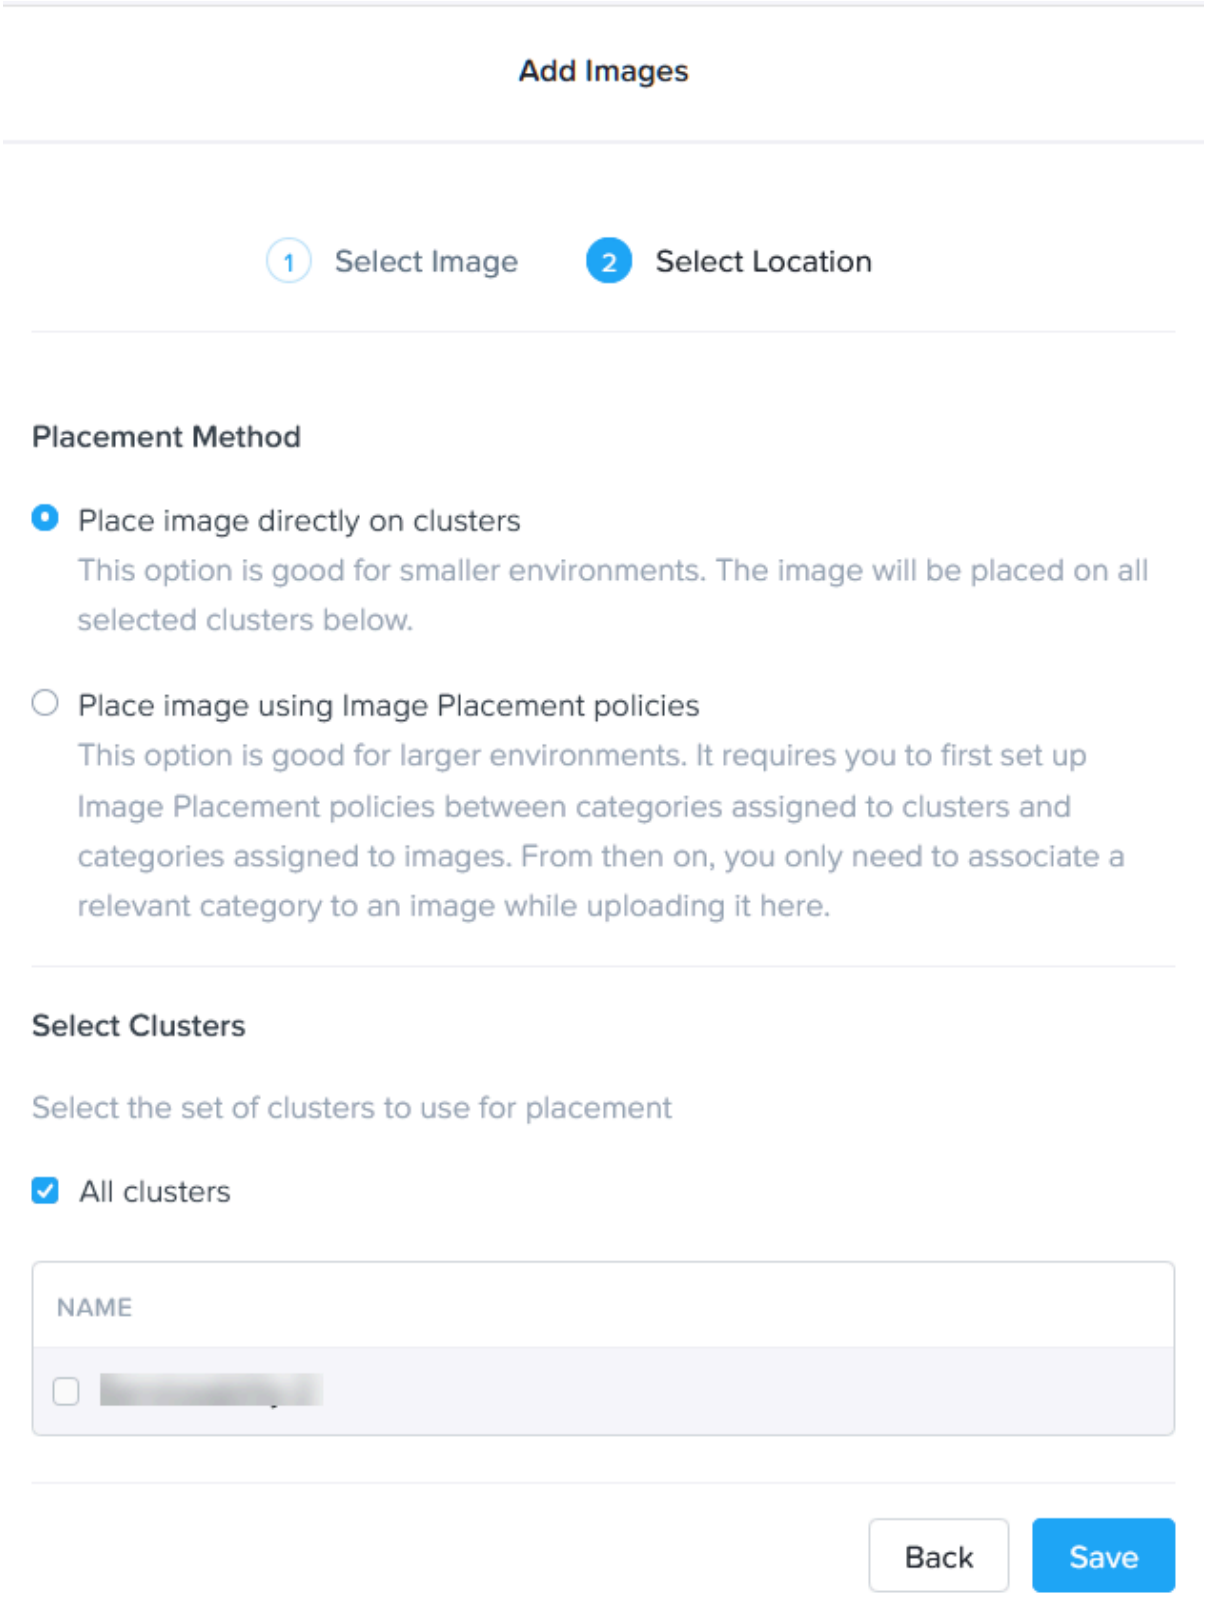

Figure 335: Add Images: Select Location

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 336: 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 795 and Assigning a Category on page 797 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 125).
   The **Add Images** page appears.
2. Under **Image Source**, click **VM Disk**.

Figure 337: 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.

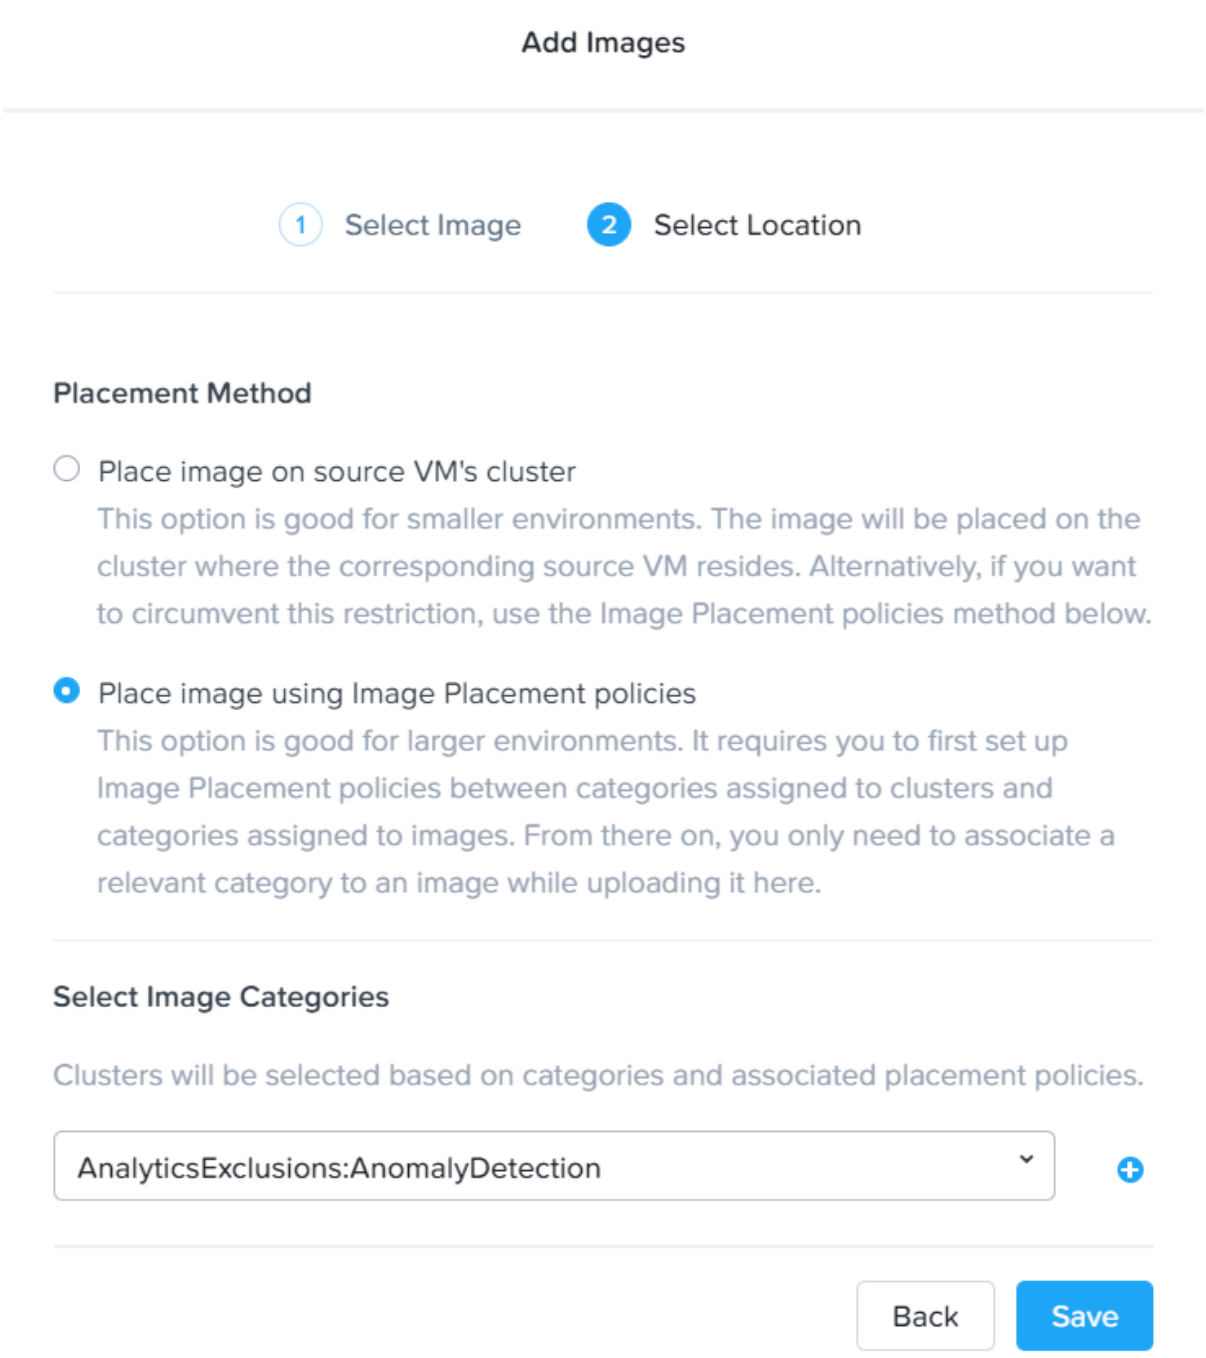

Figure 338: Add Images: Select Location

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 671 and Associating Images with Categories on page 691 for more details about adding an image.

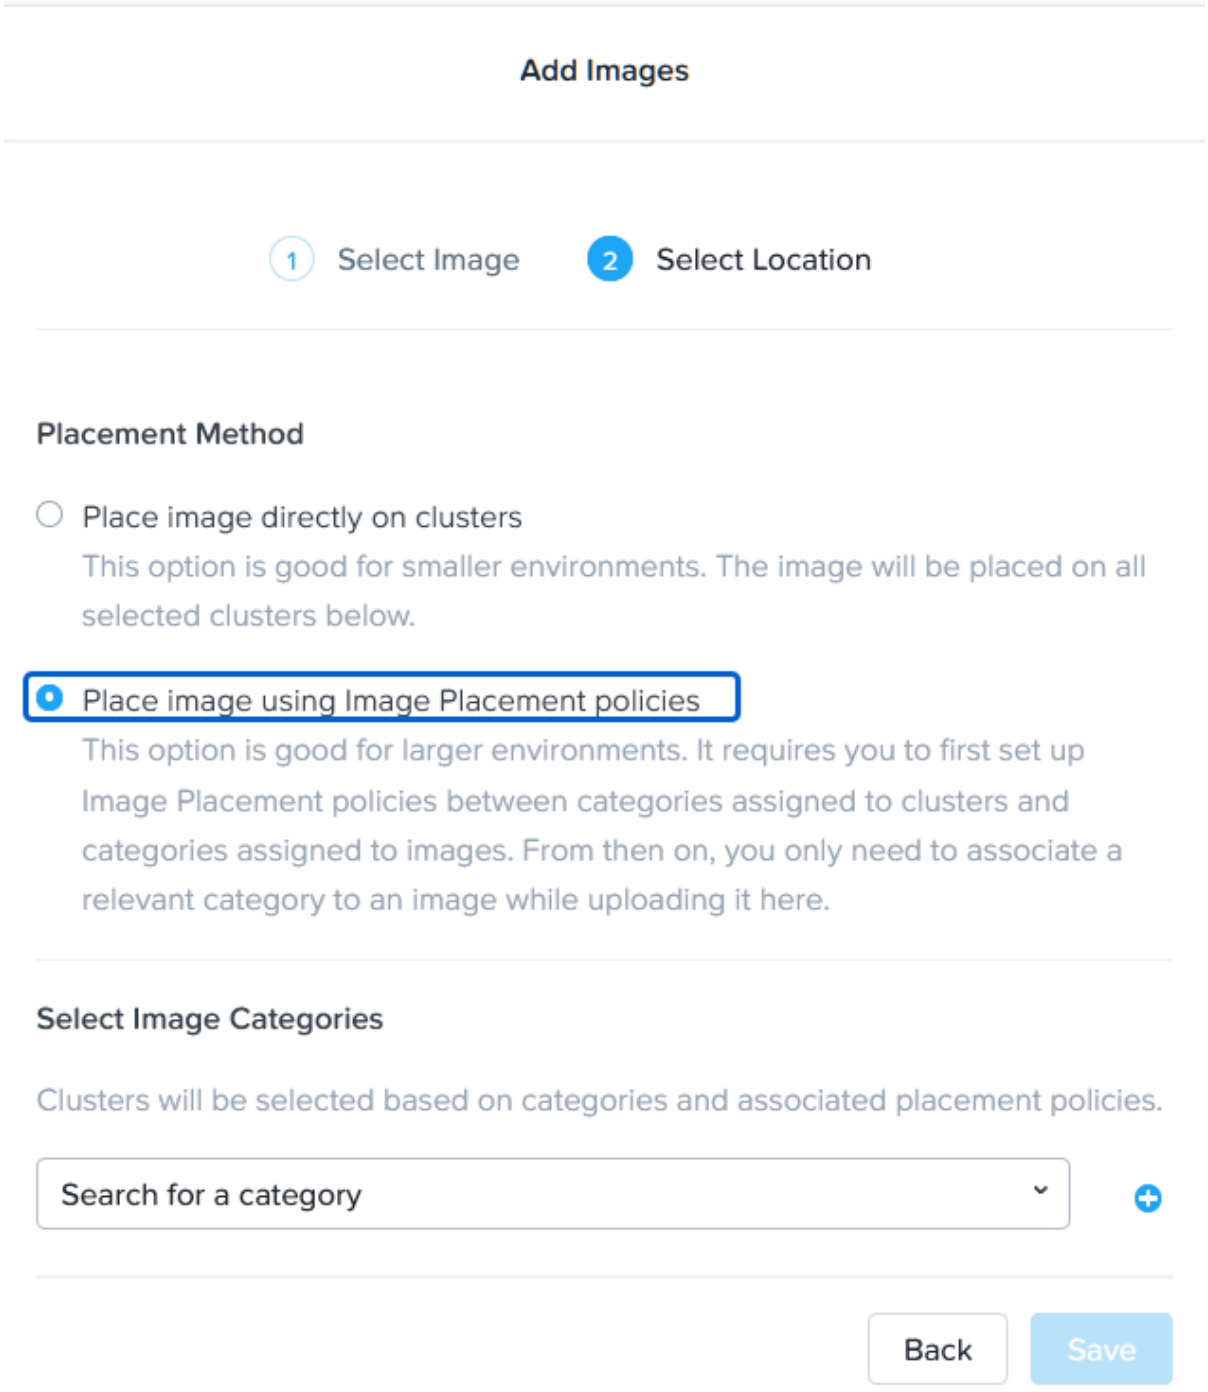

Figure 339: Image Placement Method selection for multiple clusters

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 125), and then do the following:

Procedure

- To delete an image, select (check the box for) the target image and then select **Delete** from the **Action** pull-down menu.
- To update an image, select the target image and then select **Update Image** from the **Action** pull-down menu. In the **Edit Image** page, update the name, description, or type as desired, and then click the **Save** button.

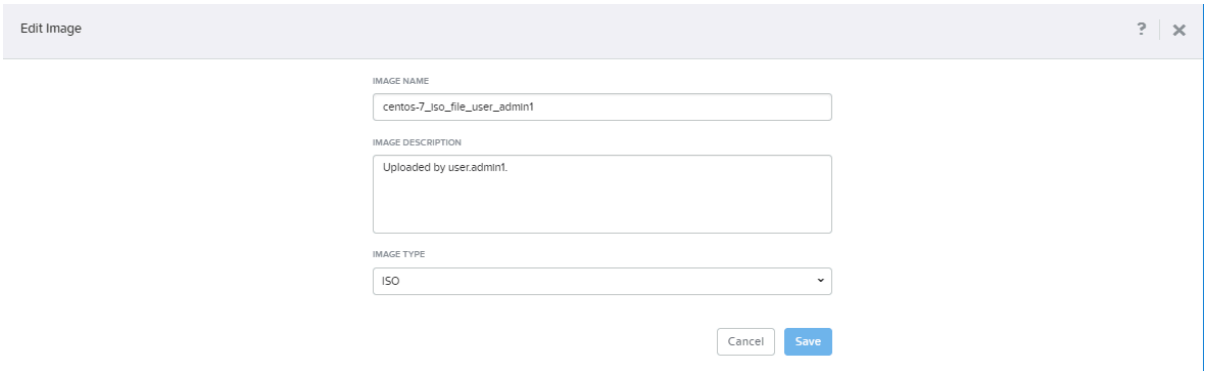

Figure 340: Update Image

- To add an image to the catalog (see Catalog Management on page 666), 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 125).

2. In the **Import Images** dialog box, under **Select Import Type**, do one of the following:

   » To import all images from all registered clusters, click **All Images**.

   » To import all images from a selection of registered clusters, click **Images On a Cluster**, and then, in the list of clusters, select the clusters.

   » To import specific images from a given cluster, click the **Select Images** link provided for the cluster. In the dialog box that is displayed, select the images that you want to import, and then click **Done**.

   Repeat this step for all the clusters from which you want to import images.

3. To begin the import, click **Save**.
   Prism Central imports the metadata of the selected images and marks the images as read-only entities on the clusters.

   The image continues to stay in the container that hosted the disk before importing the image.

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 690)
- Bandwidth Throttling Policies (see Bandwidth Throttling Policies on page 695)

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 and Configurations on page 693.

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 and Configurations on page 693.

To view the location of, the category and Image Placement policy associated with an image, see Images Summary View on page 125 and Image Details View on page 127.

Configuration Workflow

About this task

To set up an image placement policy, do the following:

Procedure

1. Create categories for the following entities:
   a. Clusters (for example, create categories based on cluster location or region)
   b. Images (for example, create categories based on operating system and size)

   For information about creating a category, see Creating a Category.
2. Apply the cluster categories to clusters and image categories to images.

   For information about associating categories with a Nutanix cluster, see Associating a Registered Nutanix Cluster with Categories. For information about associating categories with an image, see Associating an Image with Categories.
3. Configure the image placement policy. See Configuring an Image Placement Policy.

Associating Clusters 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 Images 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 With All Of 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 Clusters With All Of 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 and Configurations on page 693.
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 741: 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 677 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 availability for the other cluster operations.

Requirements for Bandwidth Throttling Policy

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

- Prism Central release 2021.9 or later and AOS version 6.0 or later are deployed at your site.
- The clusters are registered with Prism Central.
- A cluster category is created and associated with the cluster.

  For information about how to create a cluster category, see Creating a Category on page 795.

  For information about how to associate a category to a cluster, see Assigning a Category on page 797.
- AHV Hypervisor is deployed at your site.

  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**.

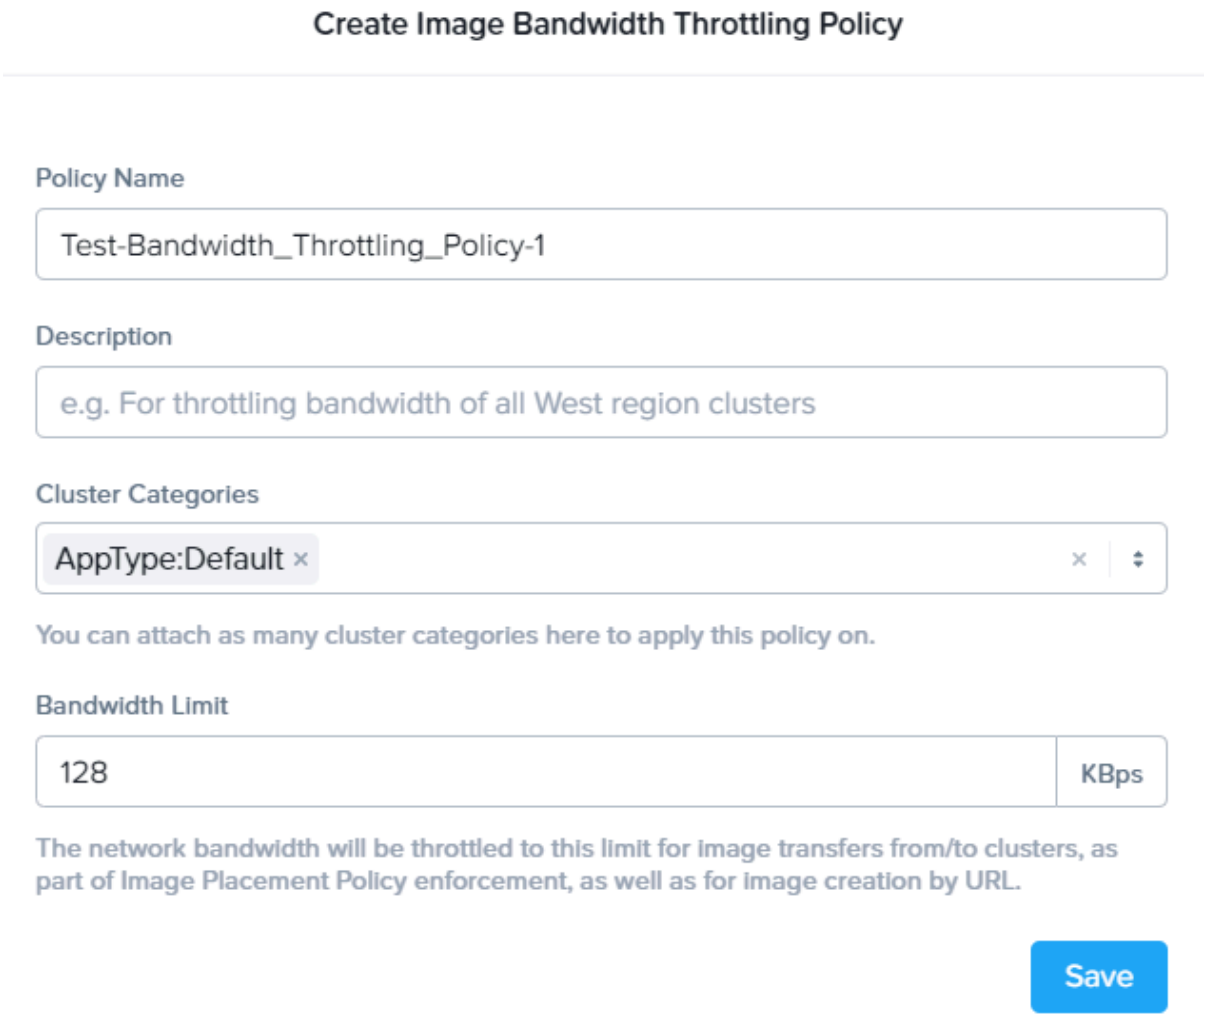

Figure 341: Bandwidth Throttling Policy

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 795 .

   > **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 Images with Categories on page 691 .

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 configured.
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.

The minimum supported versions for performing OVA operations are AOS 5.18, Prism Central 2020.8, and AHV-20190916.253.

Of the following OVA operations, you can export a VM only from the **VMs** dashboard (see VMs Summary View on page 94). You can perform the other actions from the **OVAs** dashboard (see OVAs View on page 122).

- To export a VM in the VM dashboard as an OVA, see Exporting a VM as an OVA on page 699.
- To upload an OVA, see Uploading an OVA on page 700.
- To deploy an OVA file as a VM, see Deploying an OVA as VM on page 709.
- To download an OVA file to your local machine, see Downloading an OVA on page 715.
- To rename an OVA file, see Renaming an OVA on page 716.
- To delete an OVA file, see Deleting an OVAs on page 716.

To perform any action on 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.

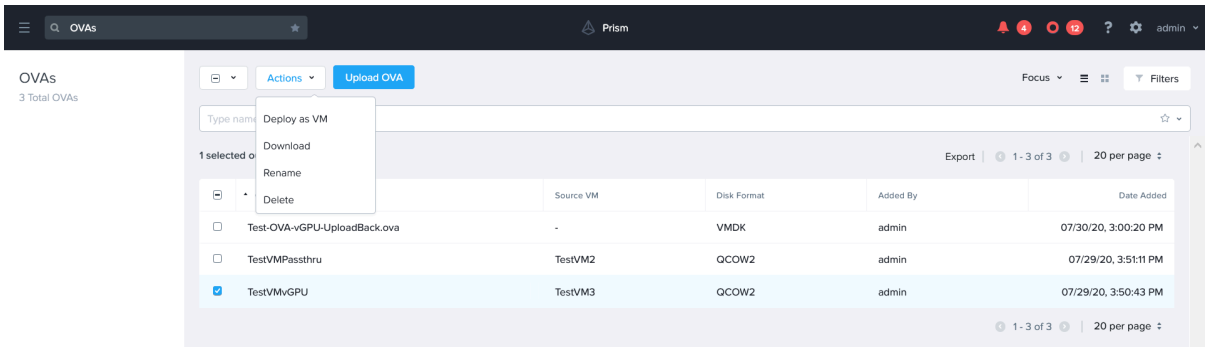

Figure 342: Actions menu

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**.

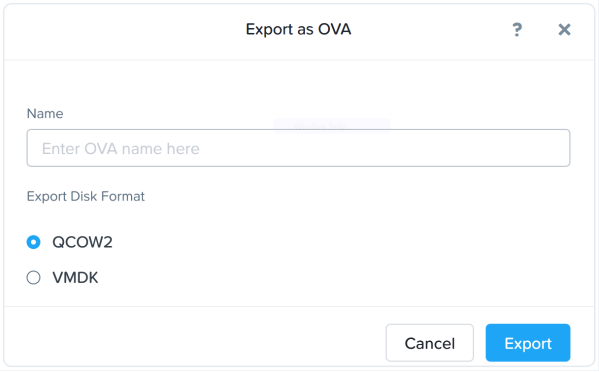

Figure 343: Export as OVA dialog box

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.
>
> There 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. If the validation operation is interrupted due to a concatenation error, you must run the concatenate API to resume concatenation and validation. See Concatenating Upload using APIs on page 706.
- You can also check the status of the upload in **Tasks**. Prism Central displays the following three subtasks 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 122.

To upload an OVA, do the following:

Procedure

1. In the **Entities** menu, go to **Compute & Storage** > **OVA**, and then click the **Upload OVA** button.

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.

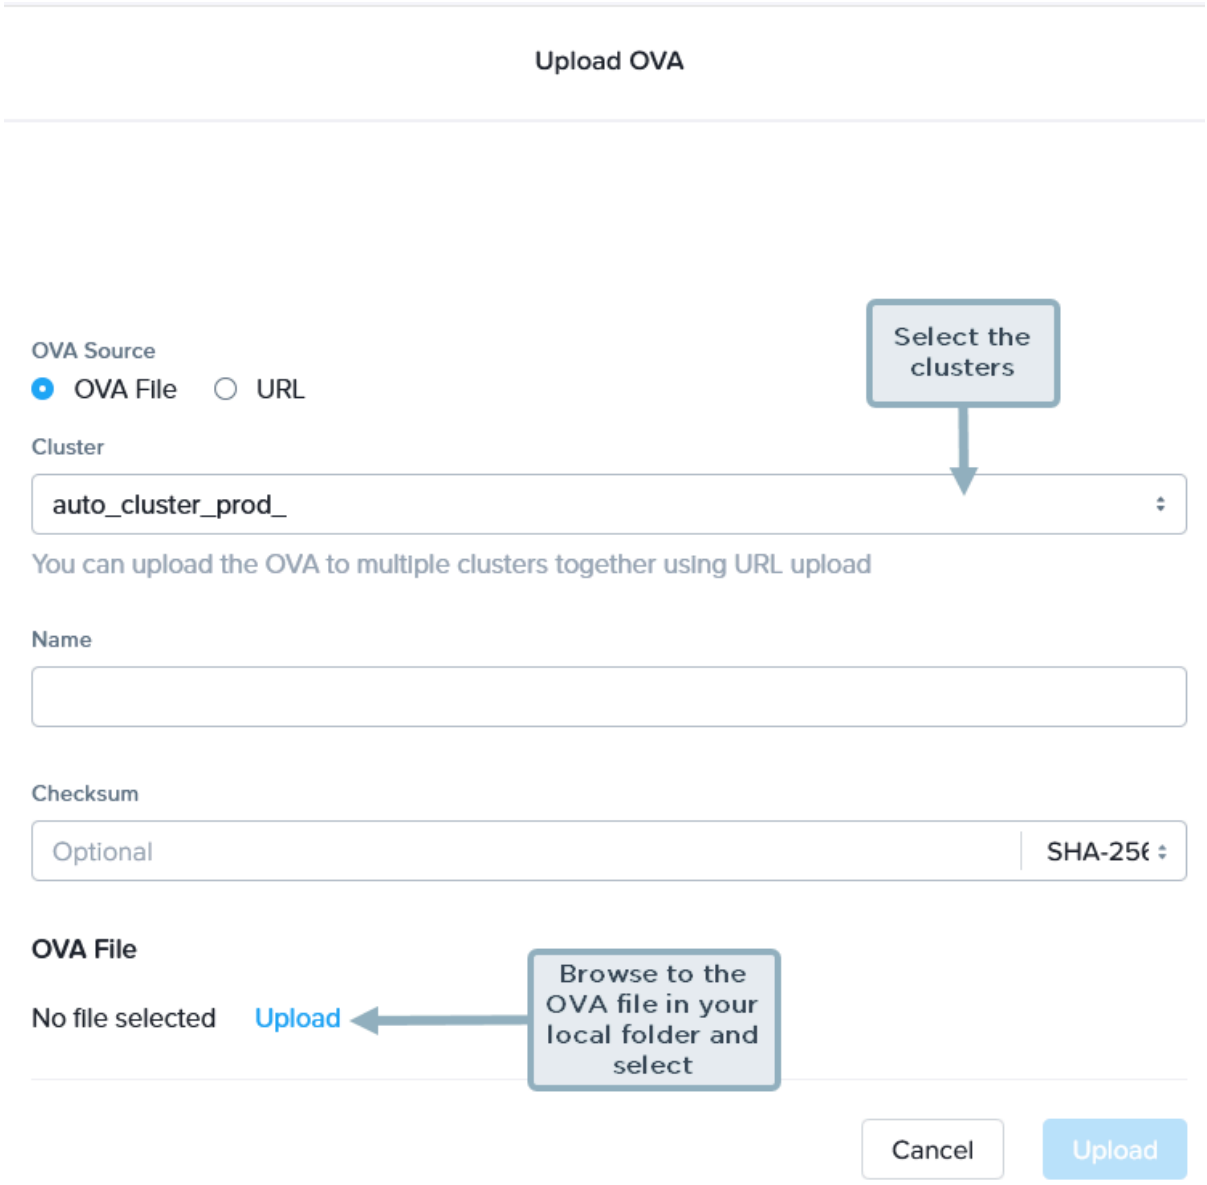

Figure 344: Upload OVA dialog box

   b. If you choose *URL* as **OVA Source**, you can select multiple clusters from the **Cluster** dropdown list that is displayed. One cluster is already selected in the **Cluster** field.

   c. Enter a **Name** for the OVA that 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. To specify one or more URLs, 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 (Prism Central) on page 514). A file uploaded through NFS must have 644 permissions.

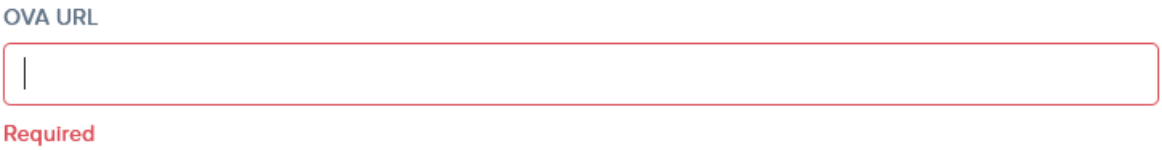

Figure 345: OVA URL field

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

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

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.

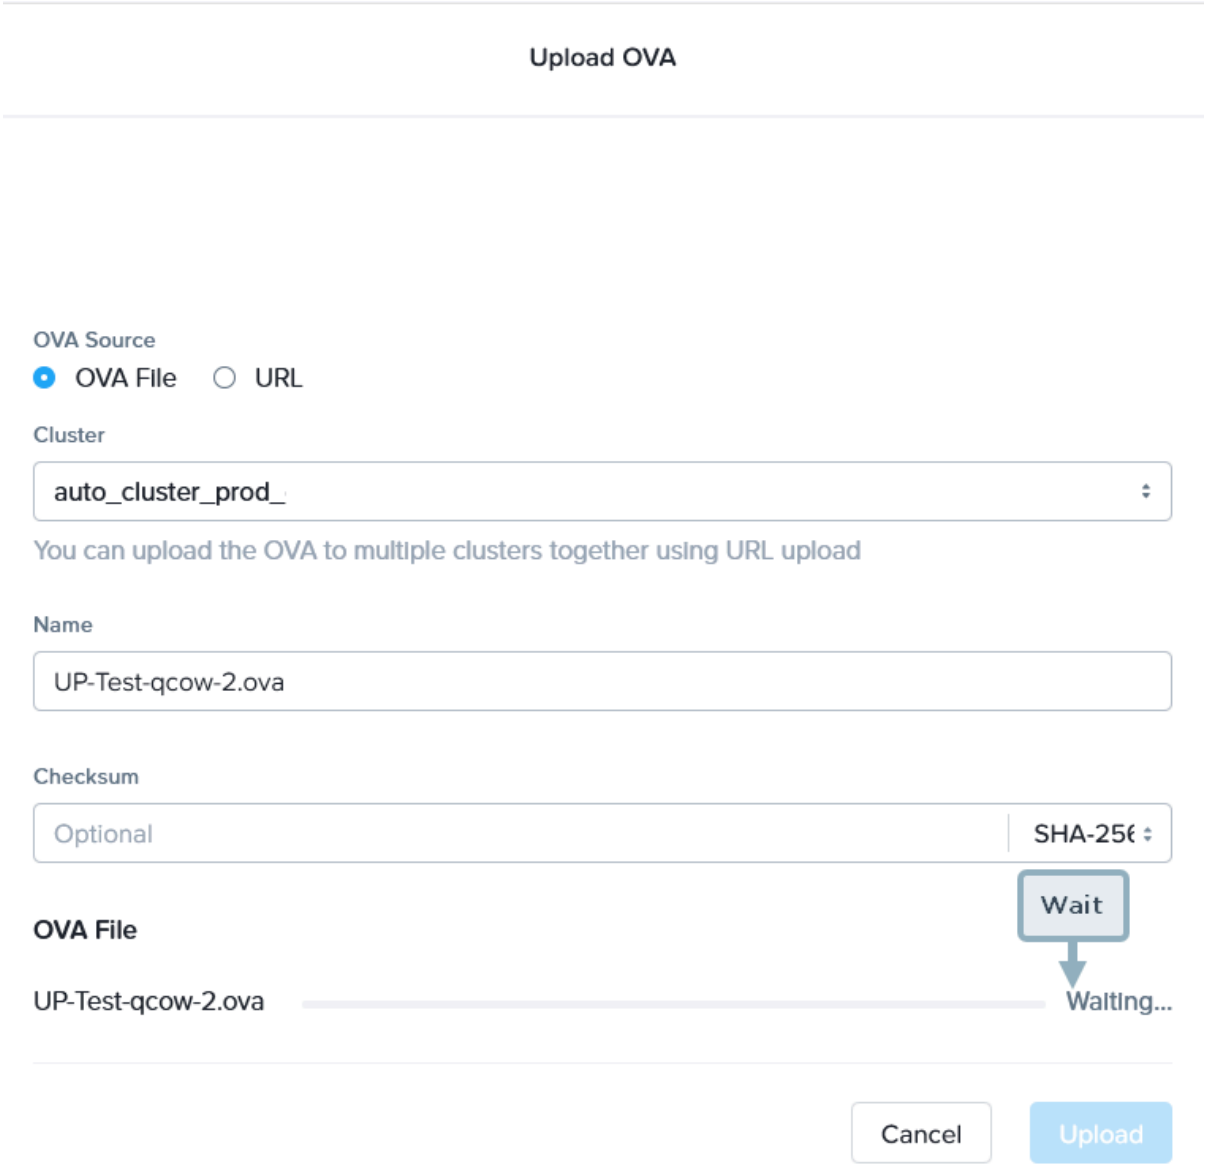

Figure 346: Upload OVA - Waiting

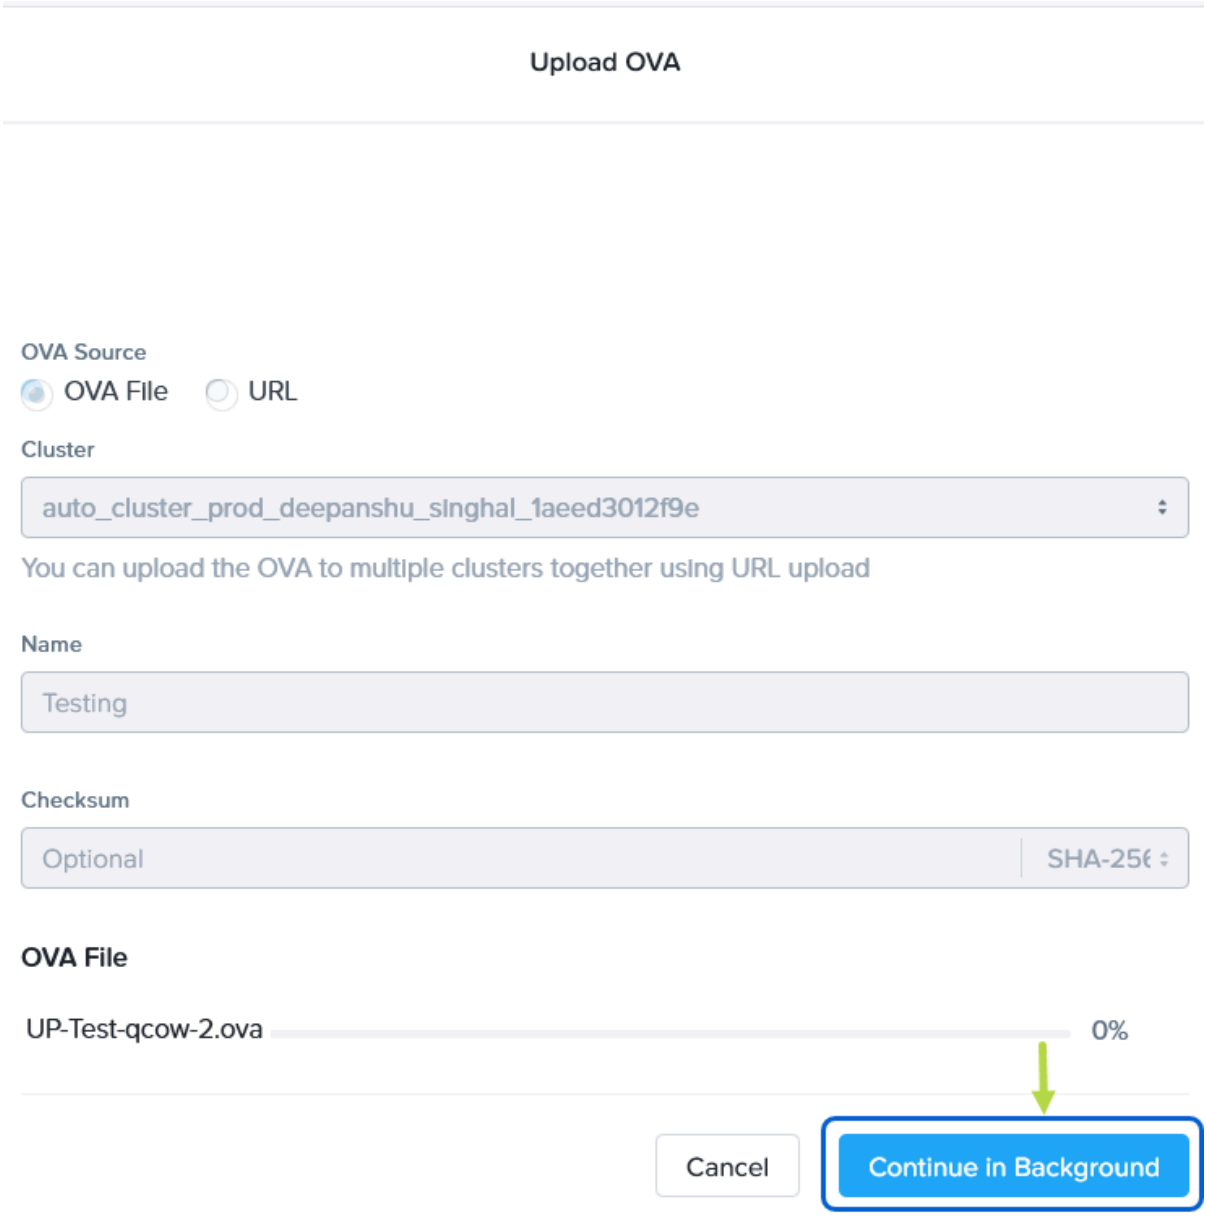

Figure 347: Continue in Background button

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 to 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

This section describes how to run the concatenation API using the REST API Explorer in Prism Central.

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 due to the concatenation error, you need to run the concatenation API to resume the upload process.

To run the concatenation API successfully, you need to perform the following actions:

- Derive the UUID of the OVA for which the upload is interrupted. The UUID is derived using the name of the OVA. For details, see *Step 1* to *Step 4* in the following procedure.
- Run the concatenation API. For details, see *Step 5* to *Step 6* in the following procedure.

Procedure

To run the concatenation API, perform the following steps:

1. Log in to Prism Central.
2. Select **REST API Explorer** from the `<user>` dropdown menu in the upper-right corner.
   The REST API Explorer opens in a new browser tab and displays a list of the objects that can be managed by the API.

3. In the REST API Explorer browser tab, navigate to **ovas** > **Post /ovas/list**, and click **Try it Out**. The **Post /ovas/list** helps you filter the OVAs and get the details of the details of the interrupted OVA.

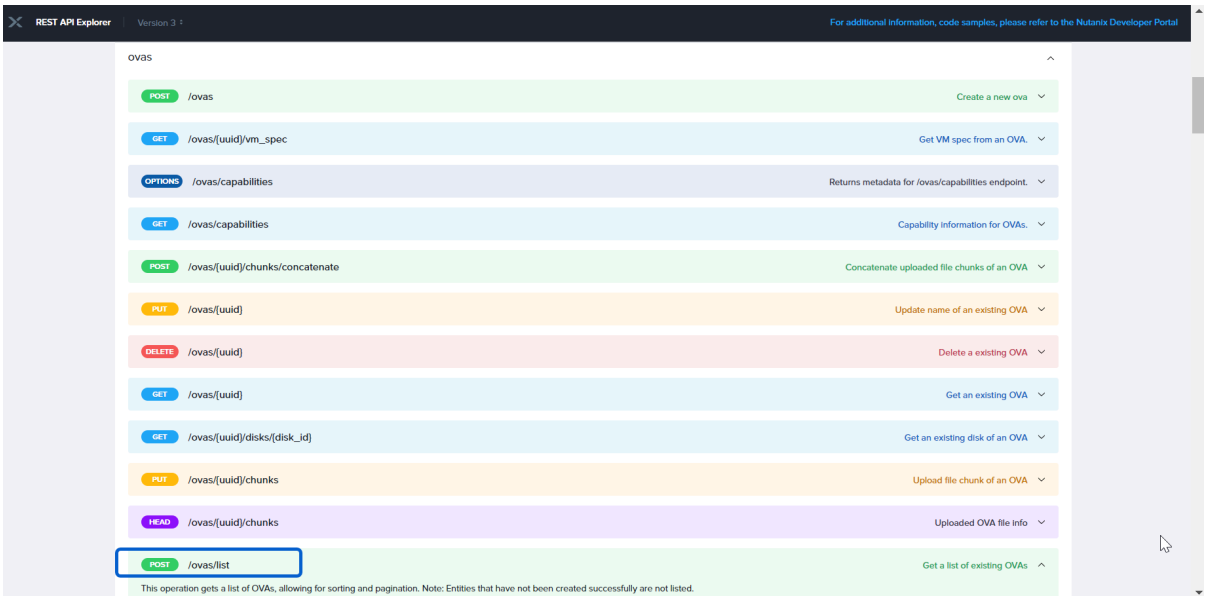

Figure 348: REST API Explorer

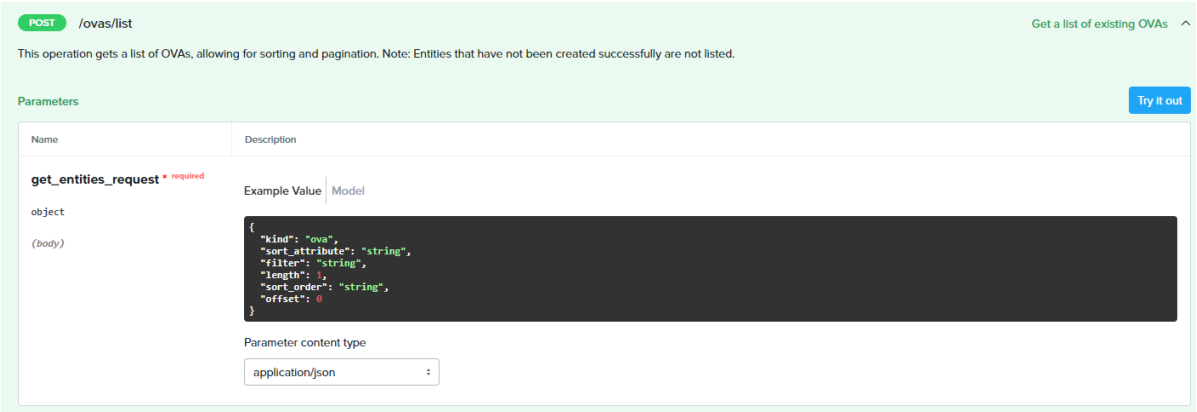

Figure 349: REST API Explorer -Try it Out

4. Enter the name of the OVA in the **"filter"** field of **get_entities_request** API in the following format, and click **Execute**.

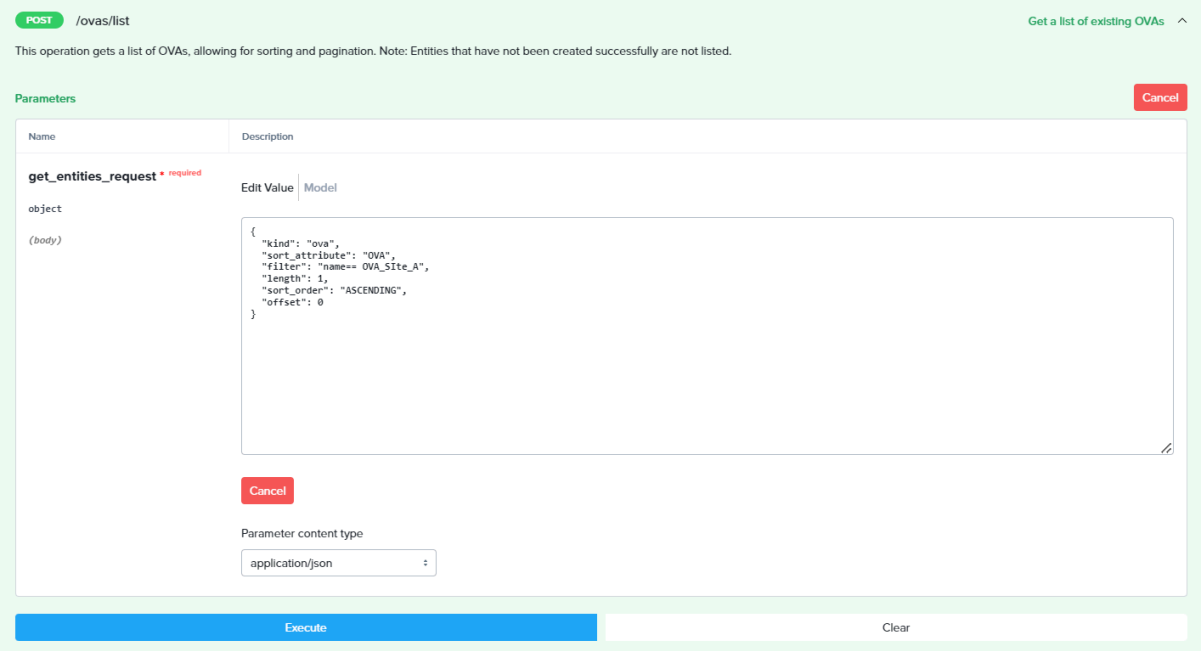

Figure 350: Parameter Filter

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

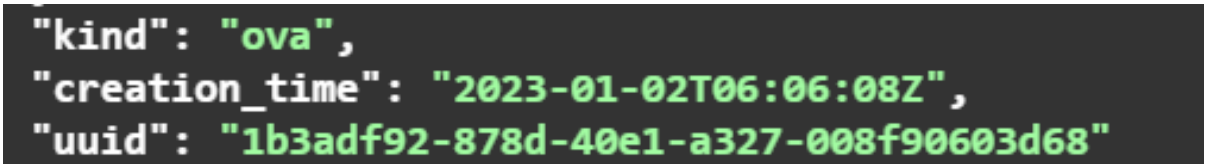

Figure 351: OVA Upload UUID

5. In the REST API Explorer browser tab, navigate to **ovas > Post /ovas/{uuid}/chunks/concatenate**, and click **Try it Out**.

6. Enter the UUID of the OVA in the **uuid** field, and click **Execute**.

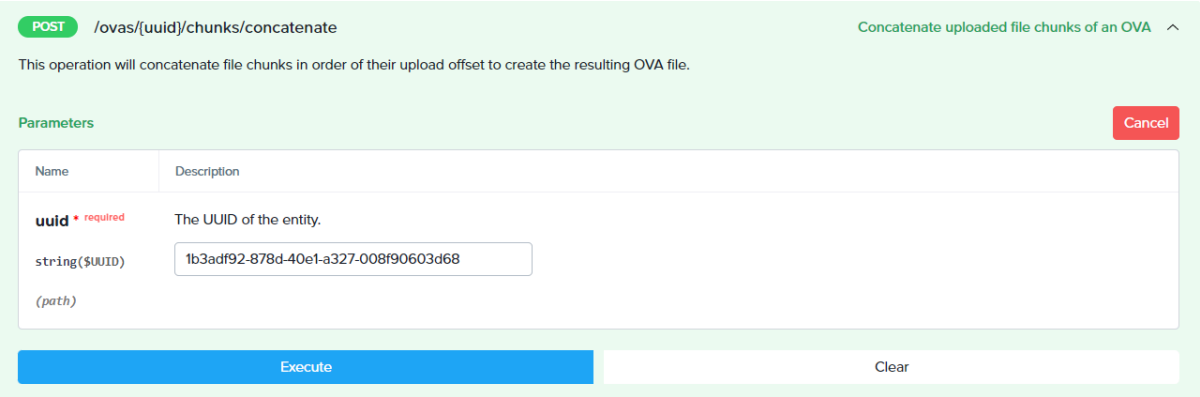

Figure 352: OVA Upload UUID - Concatenation

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 through Prism Central (AHV) on page 535. 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:**
>
> For the OVAs that are not created from AHV, ensure that the pre-installed VirtIO drivers are available for them. For more information about VirtIO drivers, see Nutanix VirtIO for Windows in *AHV Administration Guide*.

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:

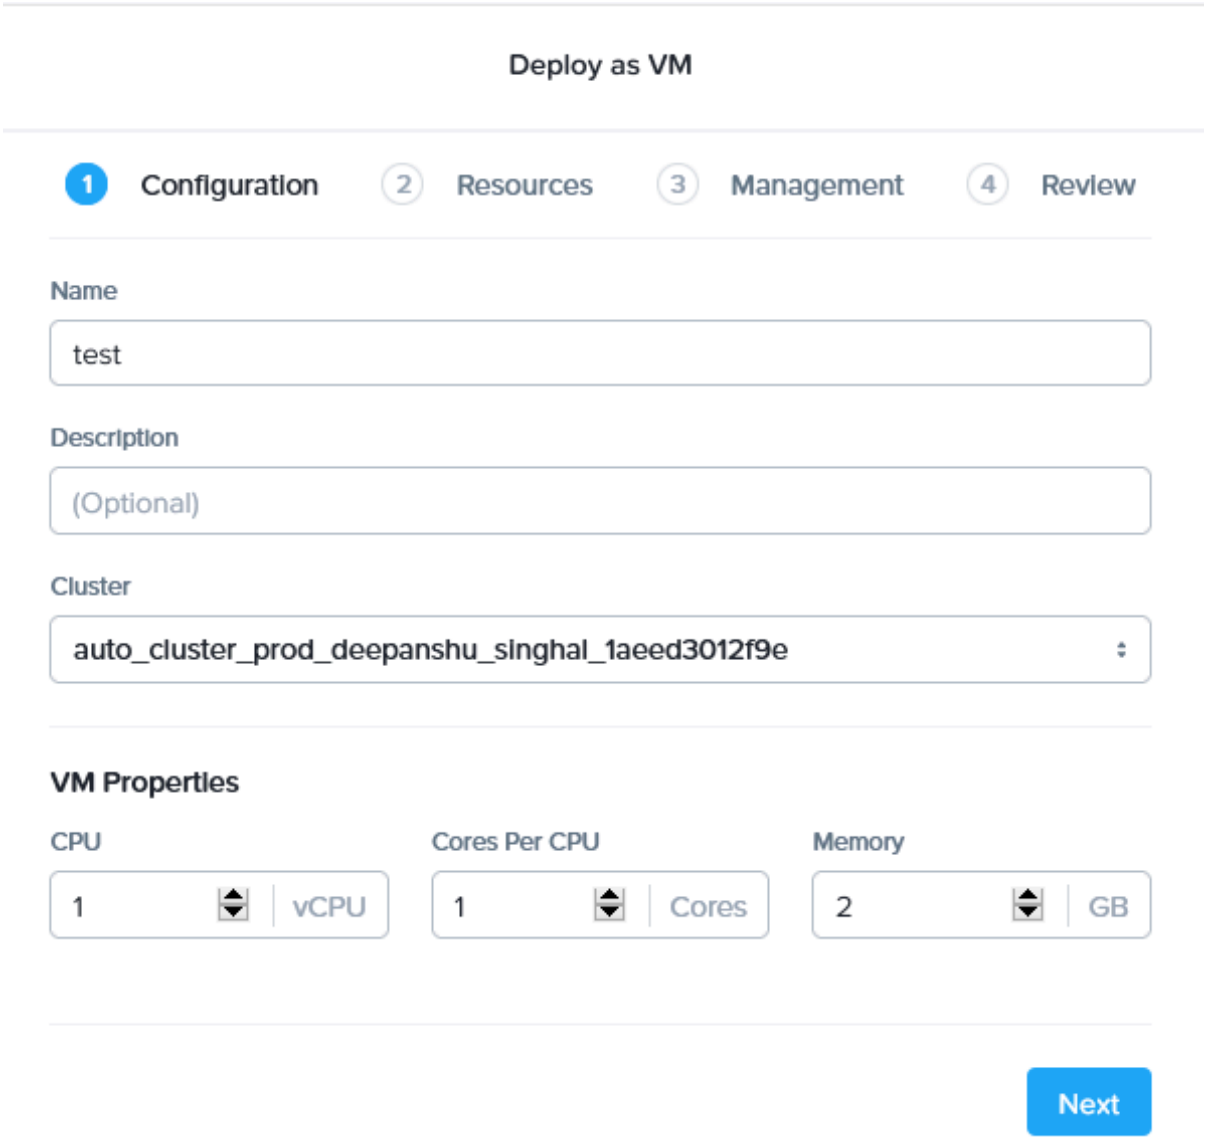

Figure 353: Deploy as VM dialog box

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**.

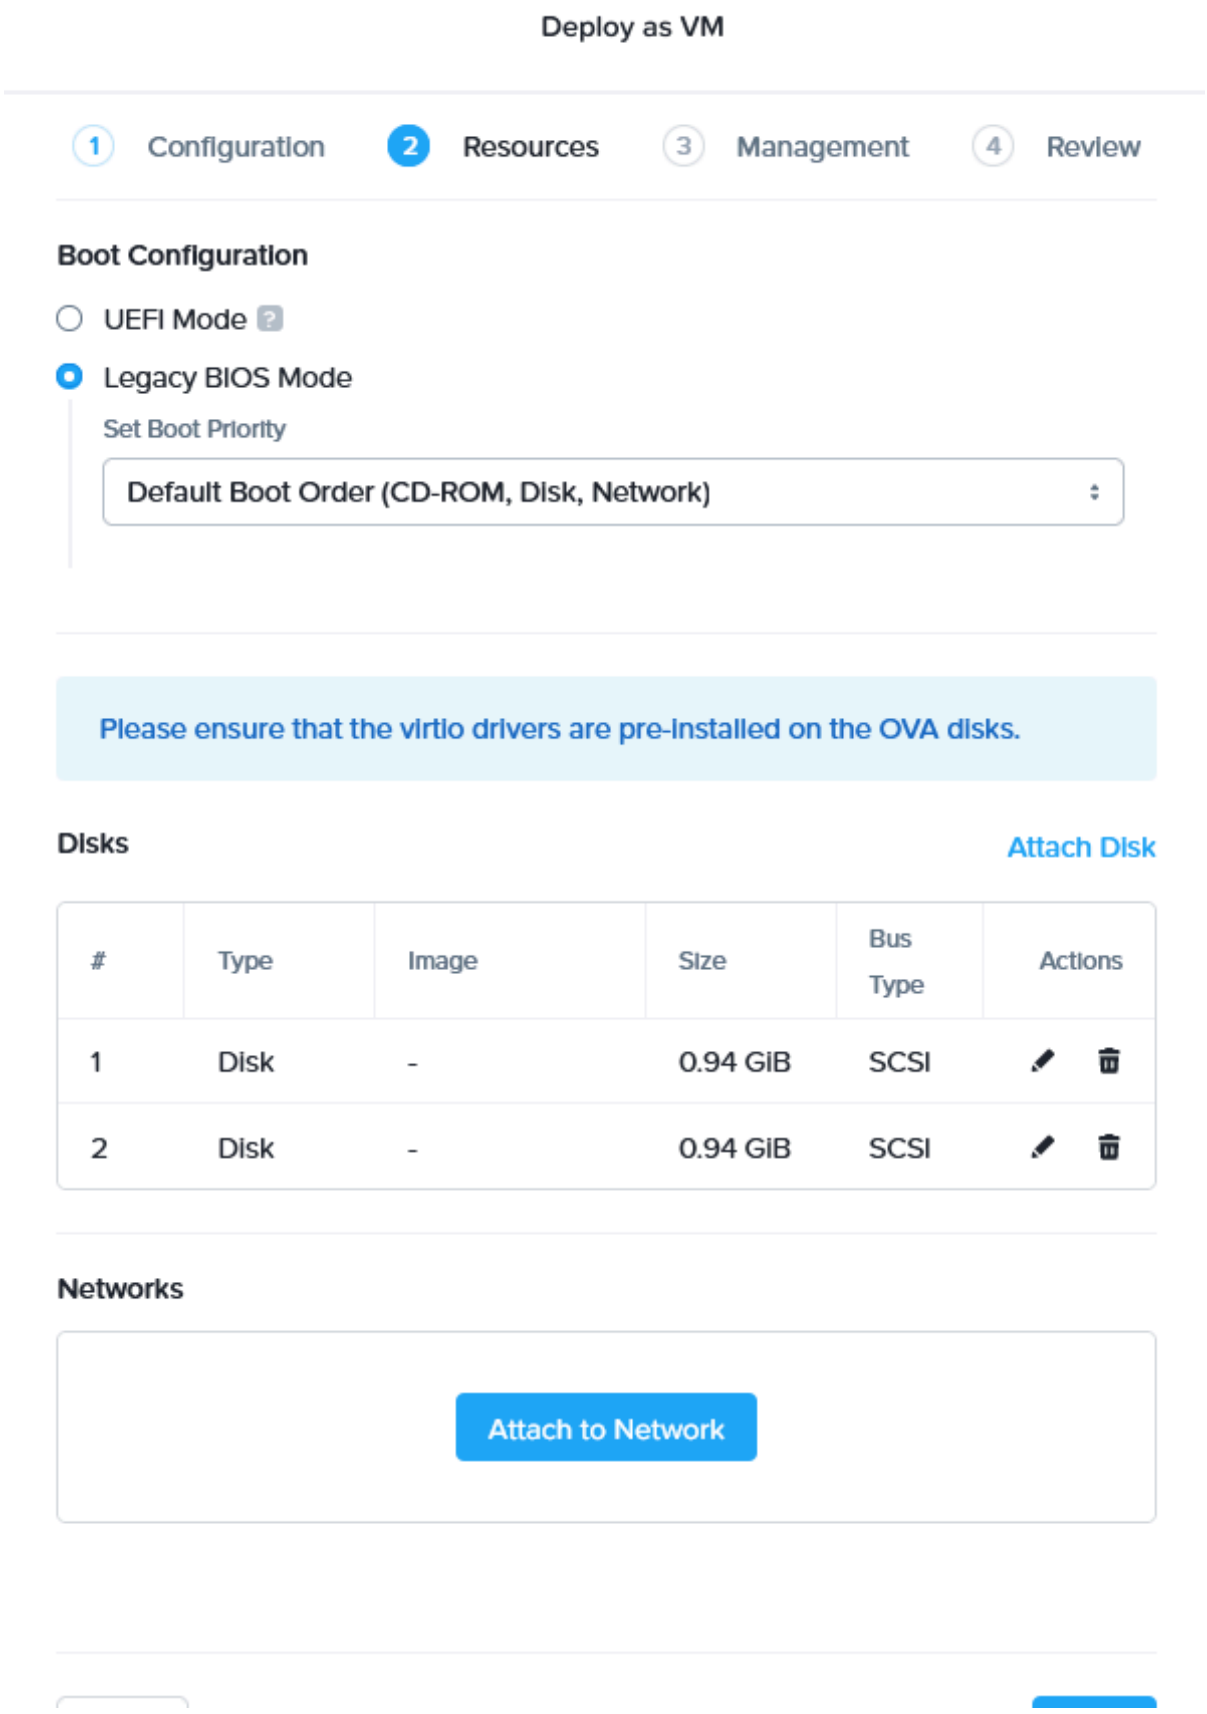

Figure 354: Deploy A VM - Resources tab

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 through Prism Central (AHV) on page 535 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 through Prism Central (AHV) on page 535 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 through Prism Central (AHV) on page 535 and Adding Multiple vGPUs to the Same VM on page 573 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**.

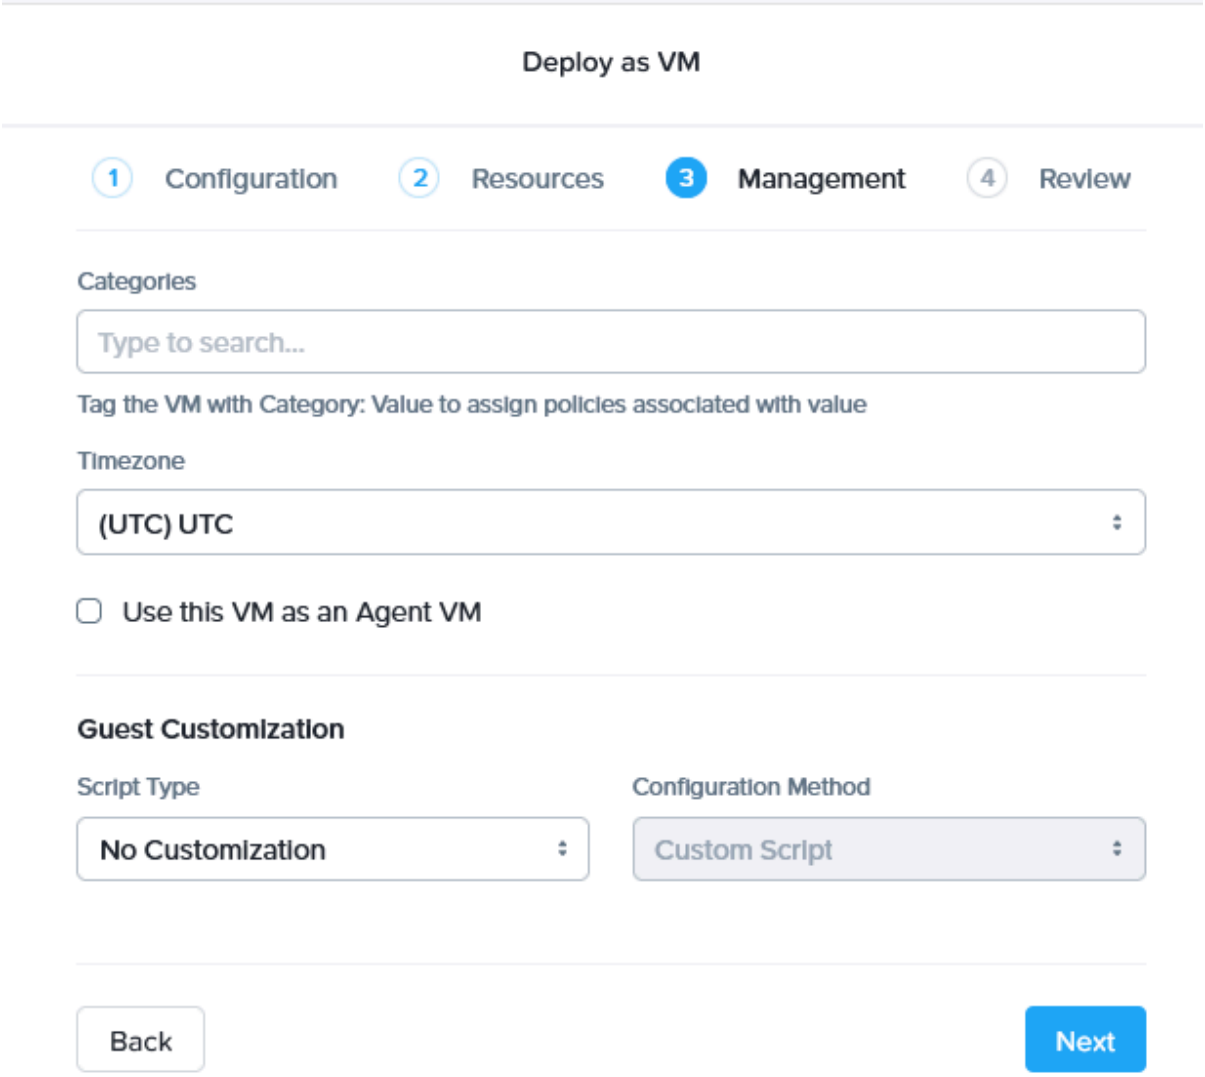

Figure 355: Deploy A VM - Management tab

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 795.

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 up to 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.

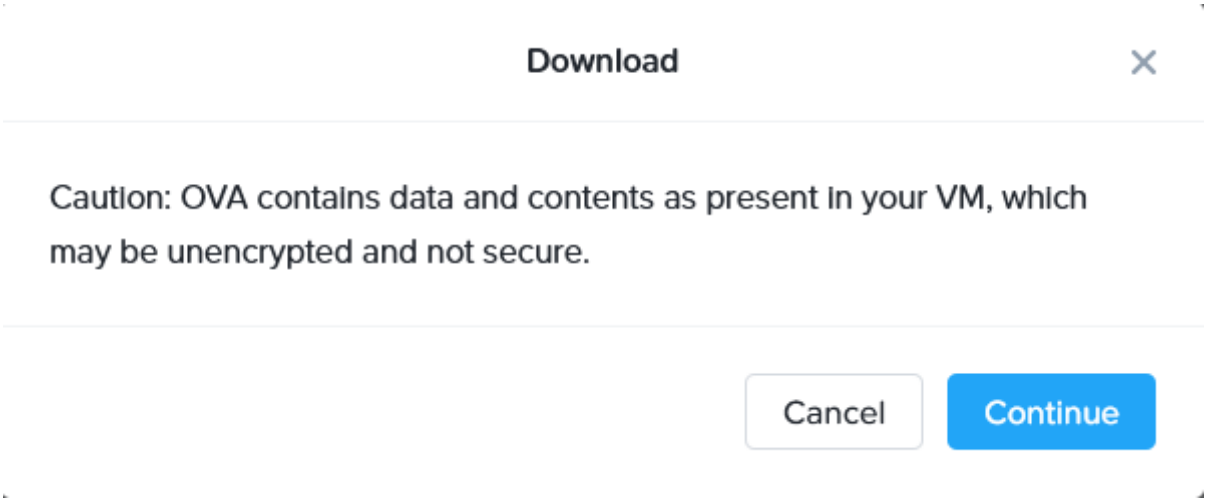

Figure 357: Download message box

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 your browser's download settings.

   Once the download starts, it continues even after you sign out from the console, as the TCP connection does not break. However, if you close the browser or the internet connection is interrupted, the download fails as the TCP connection breaks.

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.

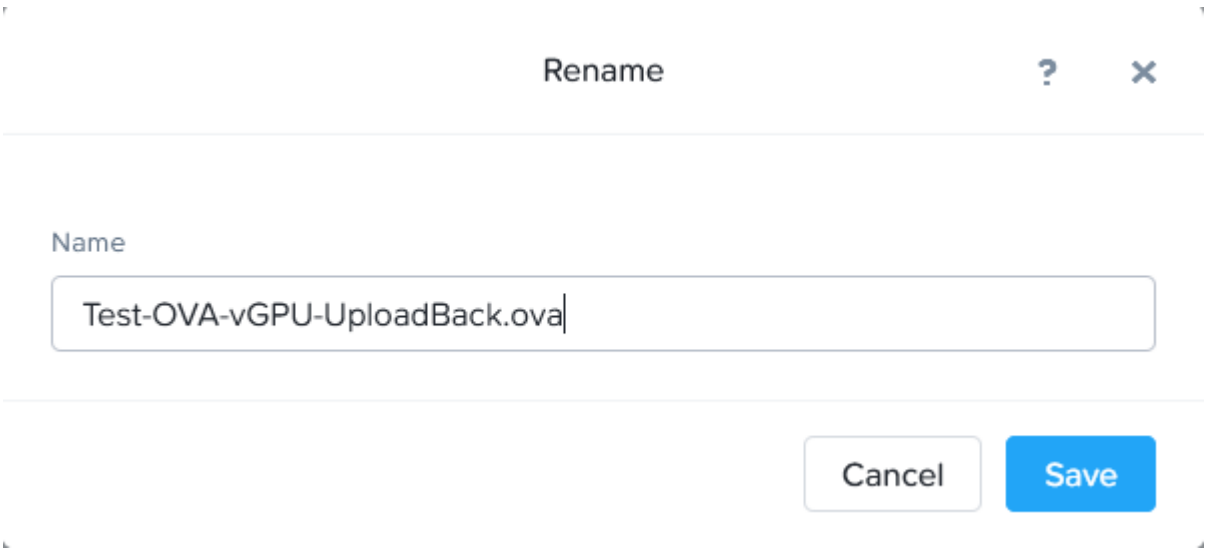

Figure 358: Rename dialog box

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>.

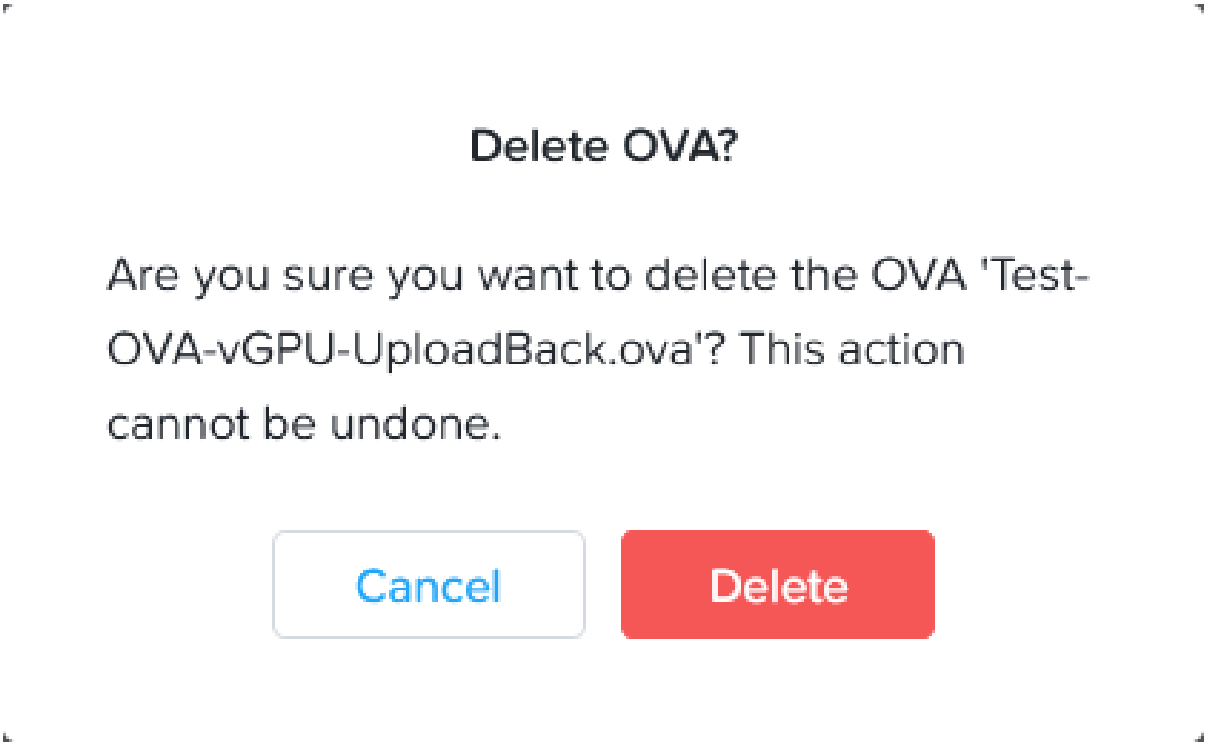

Figure 359: Delete OVA dialog box

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 202).

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 Element 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 742: 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 719. 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 722.

  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 (Prism Central) 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.

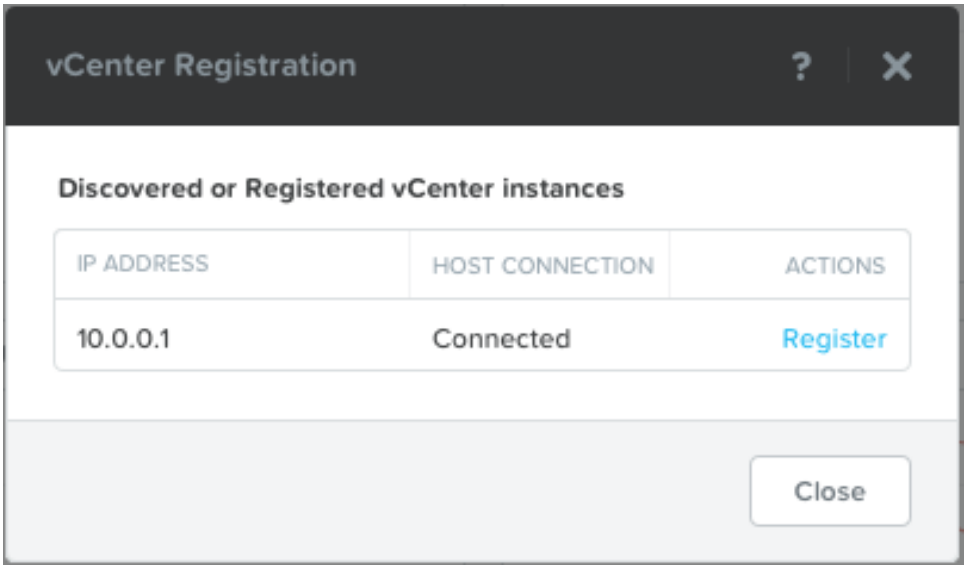

Figure 360: vCenter Registration

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**.

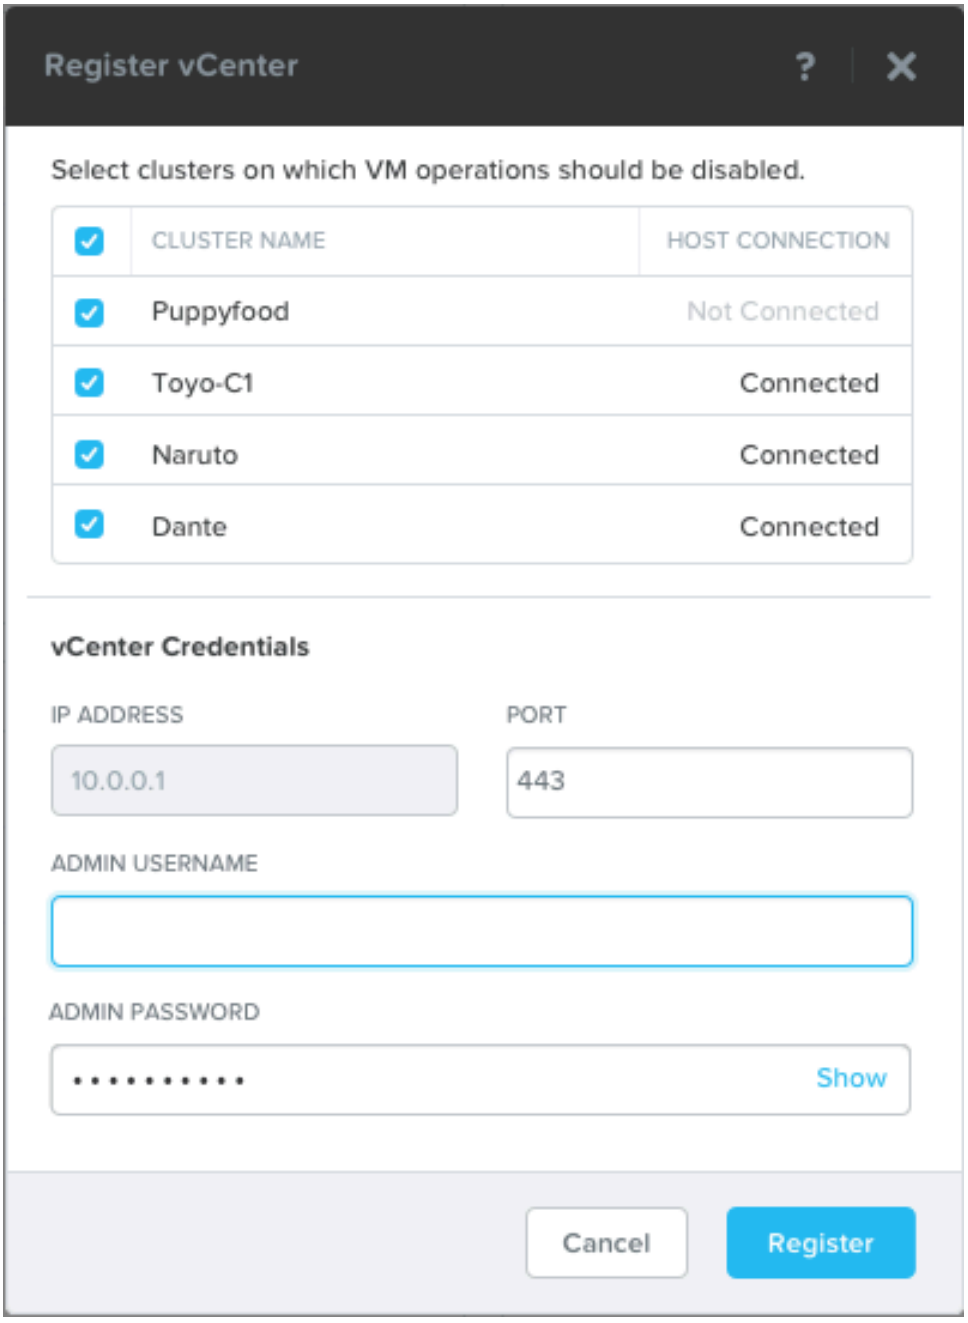

Figure 361: Cluster Selection

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 722.

Procedure

1. Click the gear icon in the main menu and select **vCenter Registration** from the **Settings** menu (see Settings Menu (Prism Central) 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 (Prism Central) 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 721.
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 719.

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 130) 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 724).
- Create a project for each team that needs self service and add users and groups to the projects (see Projects Overview on page 726).
- Configure roles for project members (see Controlling User Access (RBAC) in *Security Guide*).

  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 666).
- 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 256 and Project Details View on page 258).
- 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.

Procedure

1. Configure the Directory Services (Active Directory or Open LDAP) that includes the pool of self-service users if one is not already configured.

   For more information, See Configuring Authentication in *Security Guide*.

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).

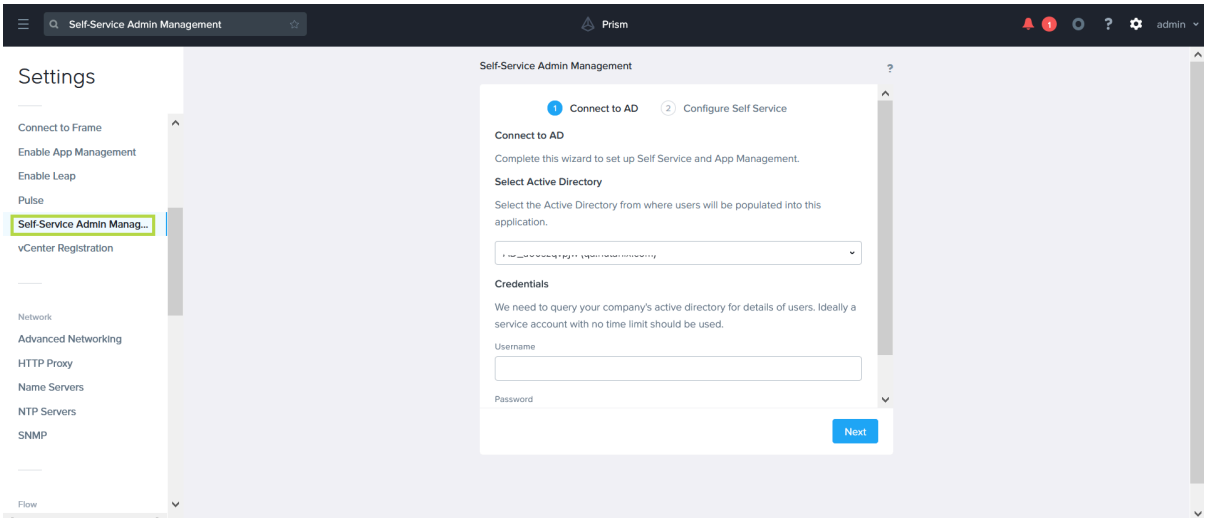

Figure 362: Self-Service Admin Management Window

3. Do the following in the **Connect to AD** tab:

   a. Select the desired Directory Services (Active Directory or Open LDAP) from the pull-down list.

   b. Enter the user name and password of the Directory Services (Active Directory or Open LDAP) 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.

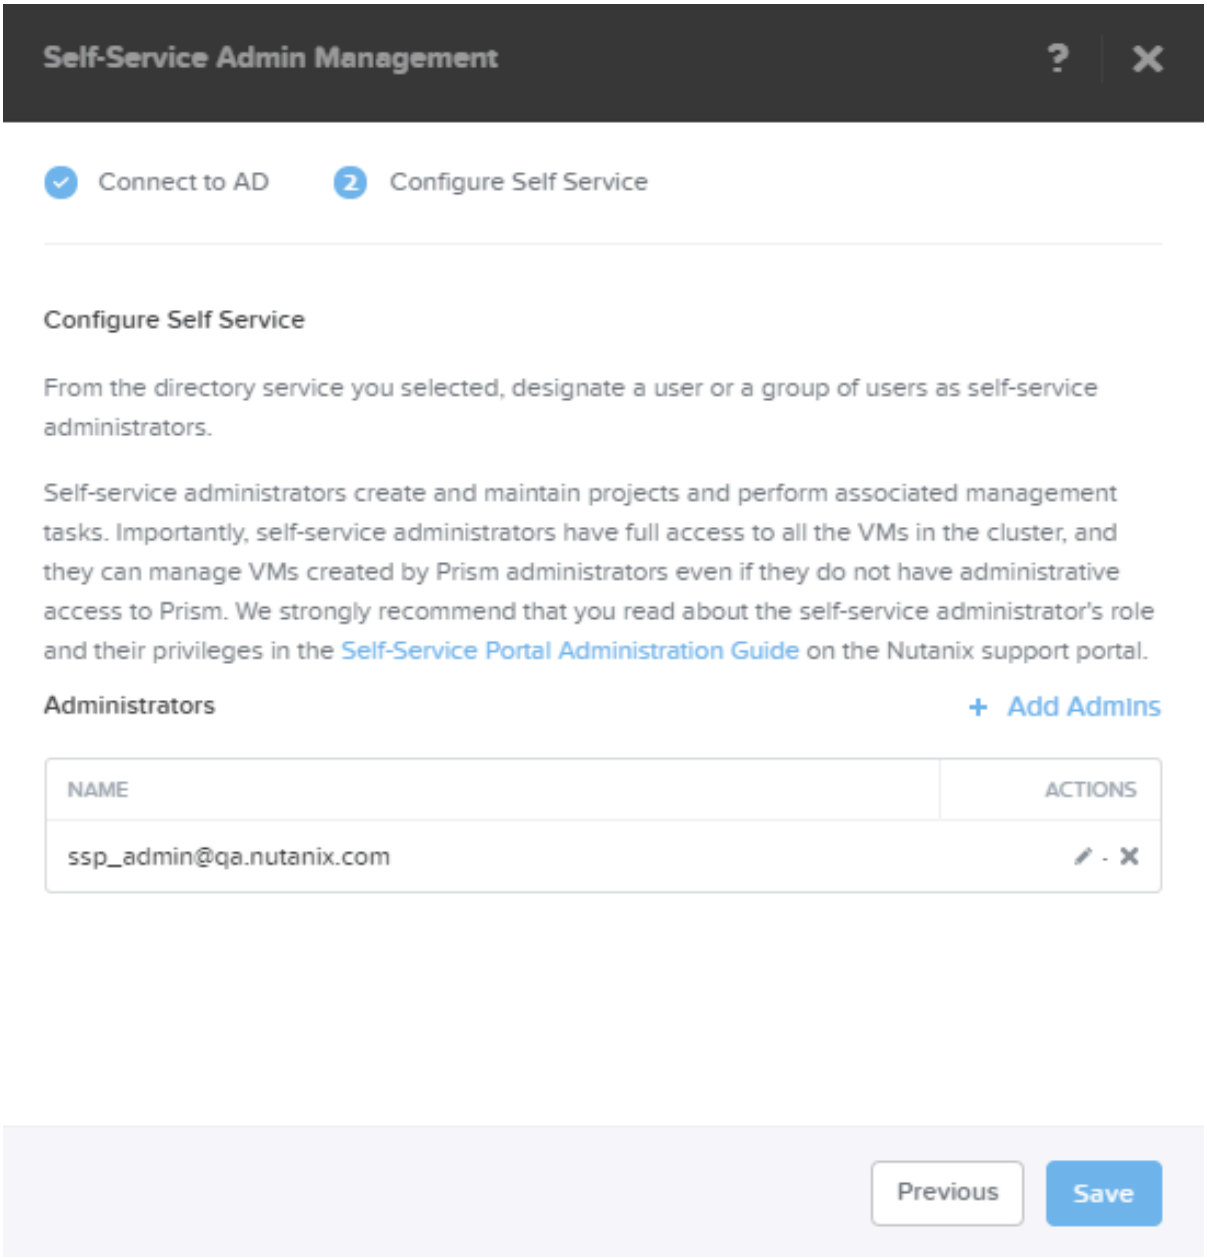

Figure 363: Configure Self Service tab

Projects Overview

A project is a set of users with a common set of requirements or a common structure and function, such as a team of engineers collaborating on an engineering project.

Projects provide logical groupings of user roles for managing resource usage within your organization. You use a directory or identity provider configured in your Prism Central to assign different roles to users or groups in a project.

The project construct helps you define various settings, such as:

- Permissions, such as the user accounts and groups who can deploy a marketplace application.

- The networks to use while deploying an application.
- Default VM specifications and deployment options, such as vCPUs, vRAM, storage, base images, Cloud-init, or Sysprep specs.
- The quota and snapshot policy definition.
- Credentials.

A project has user roles. Each role has predefined functions that a user to whom the role is assigned can perform in a project. For information on the user roles, see Roles Summary View on page 252 and Role-Based Access Control in Calm.

Project Feature Availability in Prism Central

The refactored projects provide a consistent experience when you access it from Prism Central or from Calm. However, the project features and configurations differ depending on whether you have enabled Calm in your Prism Central or not.

When Calm is disabled, the project dashboard typically appears as follows:

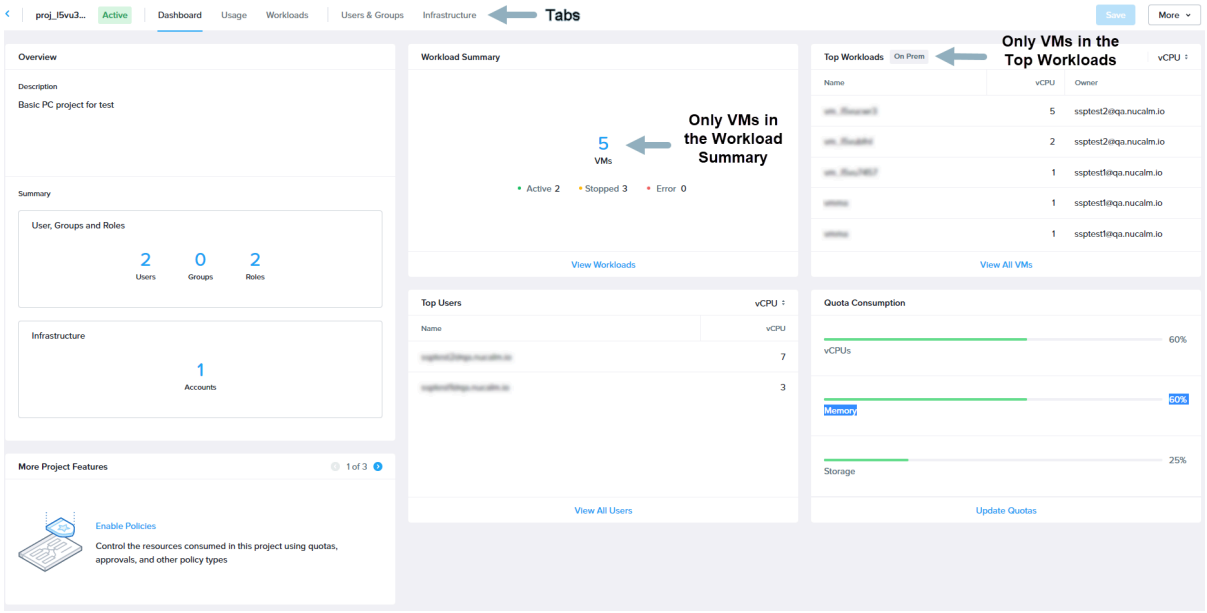

Figure 364: Projects Dashboard when Calm is Disabled

When Calm is enabled, the project dashboard typically appears as follows:

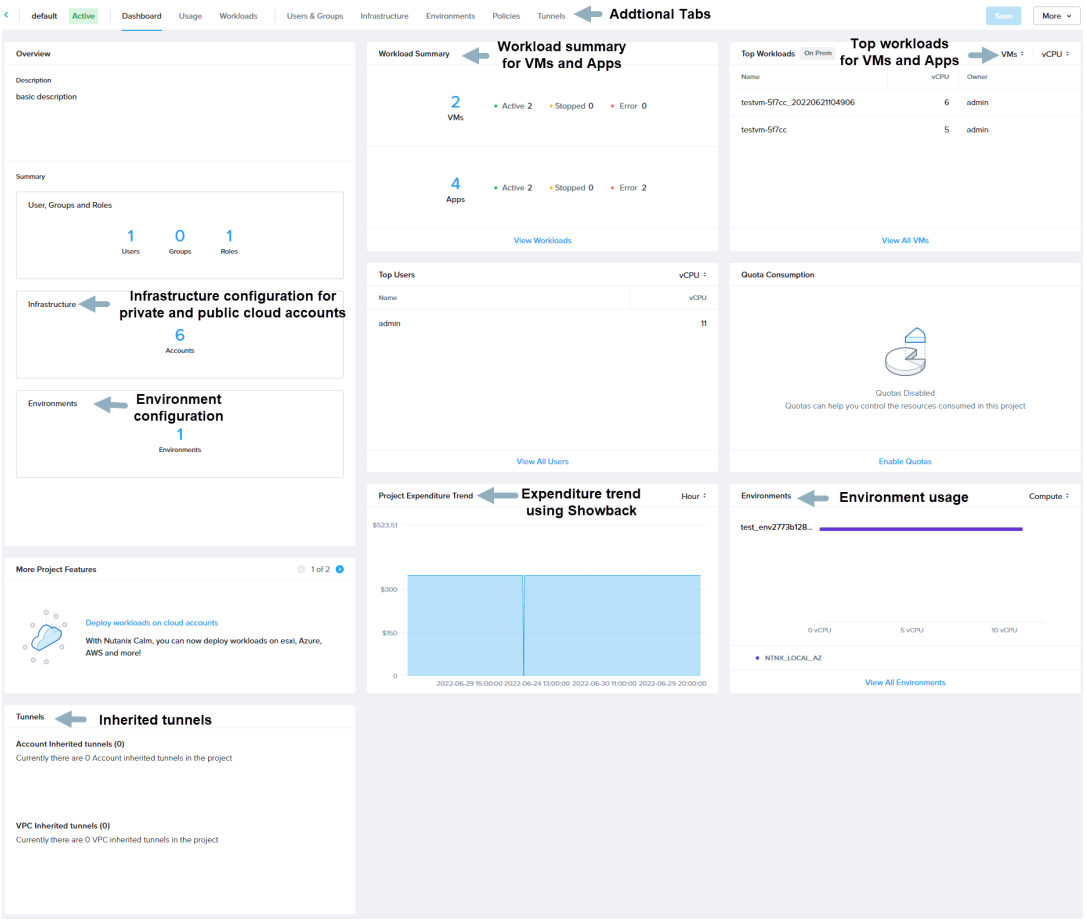

Figure 365: Projects Dashboard when Calm is Enabled

The following table lists the project features that are available with Prism Central when Calm is disabled and when Calm is enabled.

Table 743: Project Feature Availability in Prism Central

Project features	When Calm is disabled	When Calm is enabled
Allowing multiple AHV clusters and subnets	Yes	Yes
Viewing workloads per project	Yes (for VMs)	Yes (for VMs and Applications)
Viewing usage statistics (vCPU, Memory, Usage)	Yes	Yes
Internal project	Yes	Yes
Managing project-level quotas	Yes	Yes
Adding other private or public cloud accounts (ESXi, AWS, Azure or GCP)	No	Yes
Allowing VPC Overlay subnets	No	Yes

Project features	When Calm is disabled	When Calm is enabled
Managing account or cluster quotas & quotas for ESXi (new quotas with policy engine)	No	Yes
Managing environments	No	Yes
Managing tunnels for VPC or SaaS in Projects	No	Yes
Managing snapshot policies	No	Yes

Upgrade to the Refactored Project

You upgrade to the refactored project when you upgrade to Prism Central version 2022.6 or later. With this upgrade, all existing entities in Prism Central that are assigned to the default project are migrated to an Internal Project. The Internal Project is a system-defined project that encompasses all underlying infrastructure.

After the upgrade when you create VMs as an admin, the Internal Project is automatically assigned to the VMs. You can later assign the VM to a different project using the **Manage Ownership** option (see VM Management on page 535).

> **Note:** With the refactored project, Calm no longer provides an out-of-the-box Default Project when you enable it. The Default project is visible only when you upgrade your Prism Central from a previous version to version 2022.6 or later to retain existing Calm-specific entities (if any) using the Default project. The Default Project does not appear for any new deployments of Prism Central version 2022.6 or later.

Creating a Project

You create a project to provide logical groupings of user roles for managing resource usage within your organization.

About this task

Use this procedure to define the basic setup of your project. After you define the basic setup, you can further add users to your project or configure other application management features such as accounts, environments, and policies.

> **Note:** You must enable Calm in your Prism Central instance to configure application management features.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » If Calm is enabled, then in the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.

2. Click the **+ Create Project** button on the Project Summary page to create a project. The **Create Project** window appears.

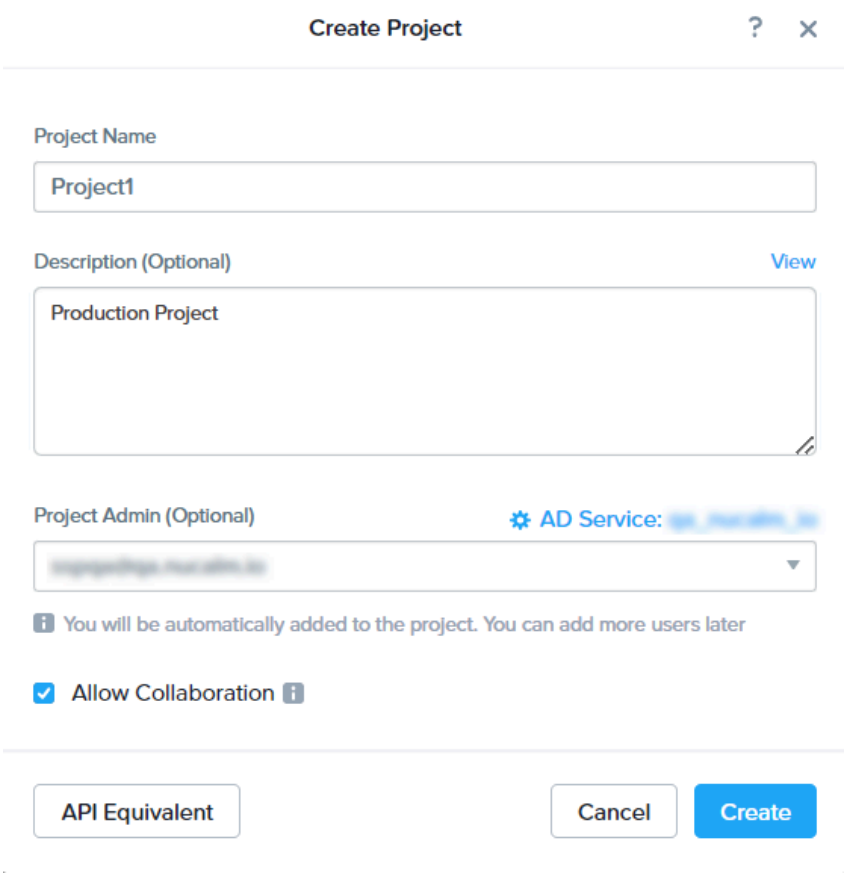

Figure 366: Create Project window

3. Type a name for the project in the **Project Name** field.
4. Type a description for the project in the **Description** field.

5. (Optional) Select an admin for the project in the **Project Admin** field.

   Type the first few letters of the project admin to get the list of matching results in the drop-down list.

   > **Note:** The project automatically adds you as a Project Admin when you create it. You can also add other users after you configure the basic setup of the project.

   If you have multiple directories or identity providers configured, ensure that the directory or identity provider from which you want to add the project admin is selected. Do the following:

   - Click the gear icon to open the Search Directories window.
   - Select the option for the directory service or identity provider that you want to use to add an admin.

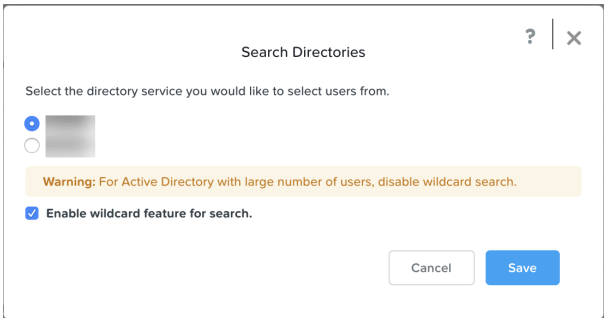

Figure 367: Search Directories Window

   - Click the **Save** button.

   > **Note:** Local users are not supported in a project. You can only add an admin from your configured directory or identity provider.

6. Check the **Allow Collaboration** check box to allow project admins and users to collaboratively manage VMs and applications within the project.

   The **Allow Collaboration** check box appears when you add your first user to the project. By default, the **Allow Collaboration** check box is checked and it enables a project admin or user to view and manage VMs and applications of other users in the same project. If you uncheck the **Allow Collaboration** check box, project admins or users can manage only the VMs and applications that they create. You cannot change this configuration after you add the first user and save the project. To change the configuration, you need to remove all users from the project.

7. Click the **Create** button.
   The **Dashboard** tab for the project appears.

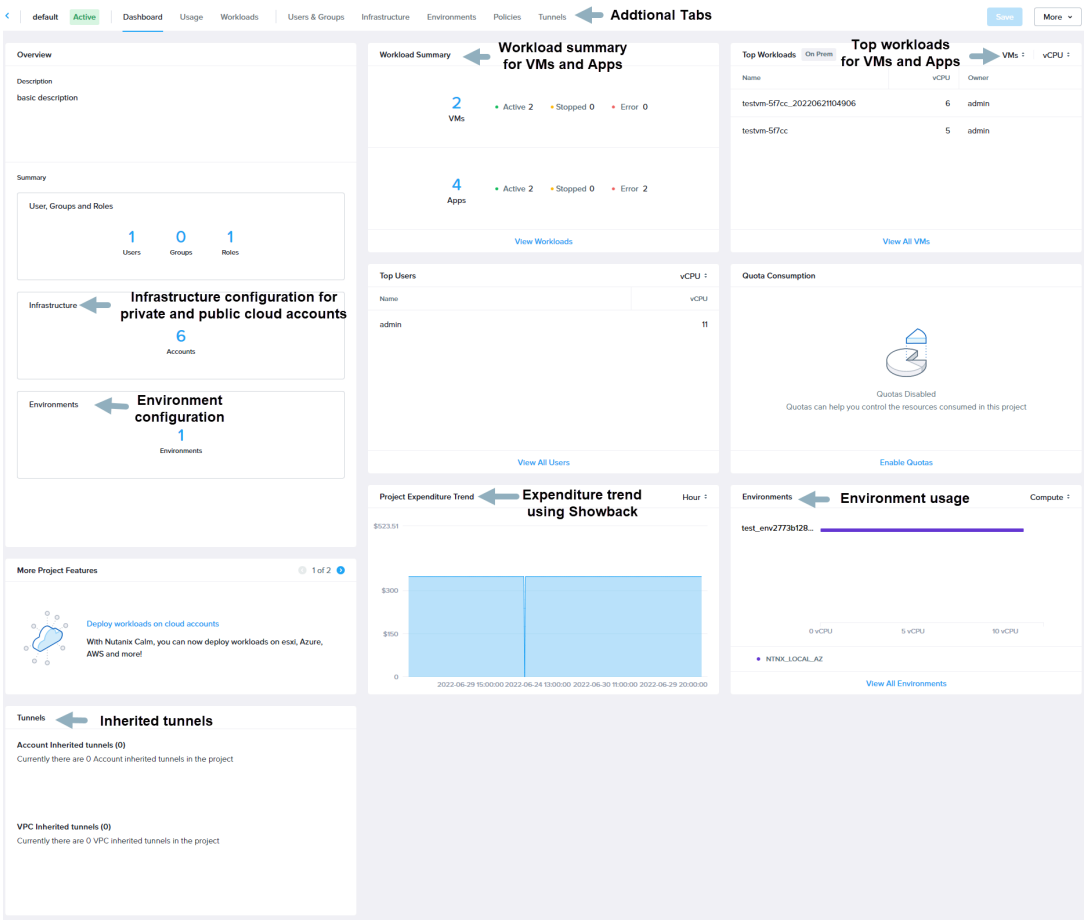

Figure 368: Project Dashboard

8. Add users or groups to your project on the **Users & Groups** tab. For more information, see Adding Users to a Project on page 733.

9. Add infrastructure to your project on the **Infrastructure** tab. For more information, see Adding Infrastructure in a Calm-Disabled Prism Central on page 737.

10. If you have enabled Calm in your Prism Central instance, you can configure infrastructure, environments, and policies in your project. Do the following:

    a. Use the **Infrastructure** tab to add different accounts to your project. For more information, see Adding Infrastructure in a Calm-Enabled Prism Central on page 739.

    b. Use the **Environments** tab to add credentials and configure VMs for the provider accounts that you selected for your project. For more information, see Configuring Environments in a Project on page 746.

    c. Use the **Policies** tab to define your quota and snapshot policies. For more information, see Quota Policy Overview on page 773 and Creating a Snapshot Policy on page 777.

11. Click the **Save** button on the Project Details page.

Adding Users to a Project

Use projects to assign roles to a particular user or group. Based on the role assigned, users are allowed to configure and manage resources within the organization.

About this task

For more information about the user roles, see Roles Summary View on page 252 and Role-Based Access Control in Calm.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » If Calm is enabled, then in the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.
2. Do one of the following:
   - » Click the **+ Create Project** button to create a new project and add users to the project. For more information about creating a project, see Creating a Project on page 729.
   - » Click a project name in the list of existing projects to add users to that project.
3. Click the **Users & Groups** tab.
4. Click the **Add/Edit Users & Groups** button.

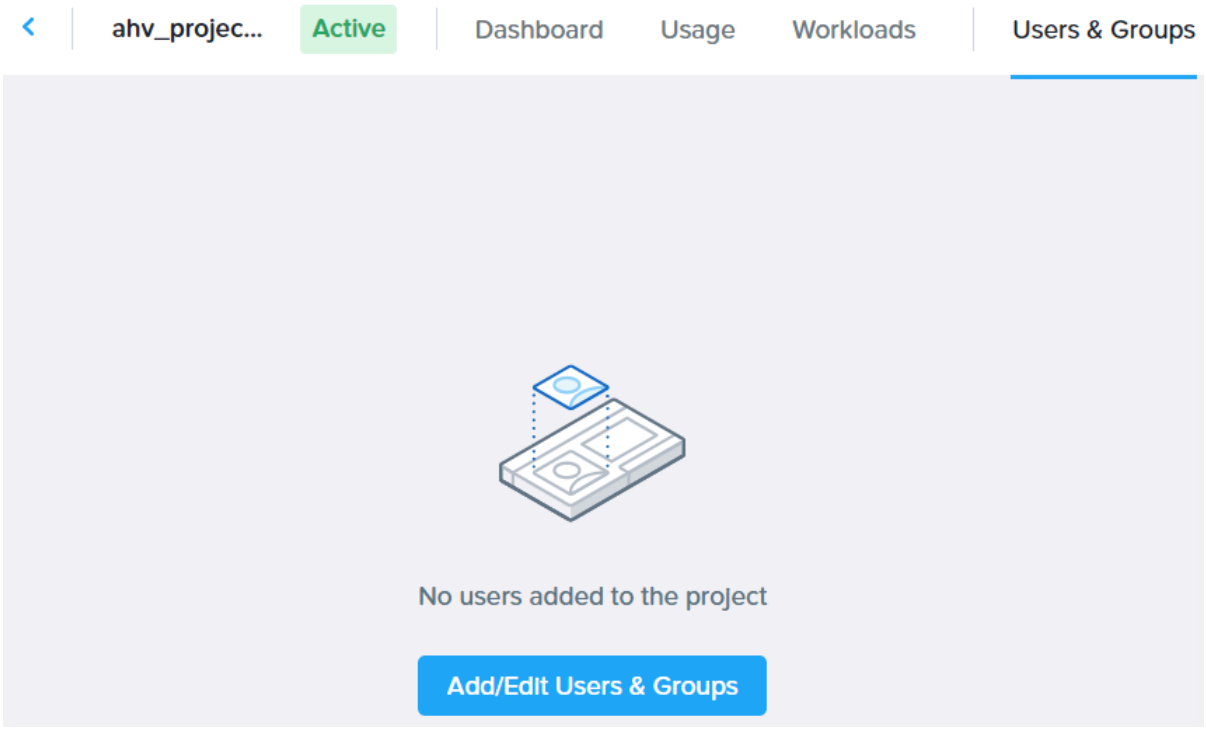

Figure 369: Add/Edit Users & Groups

5. If you have multiple directories or identity providers configured, ensure that the directory or identity provider from which you want to add users and groups is selected. Do the following:

   a. Click the gear icon to open the Search Directories window.

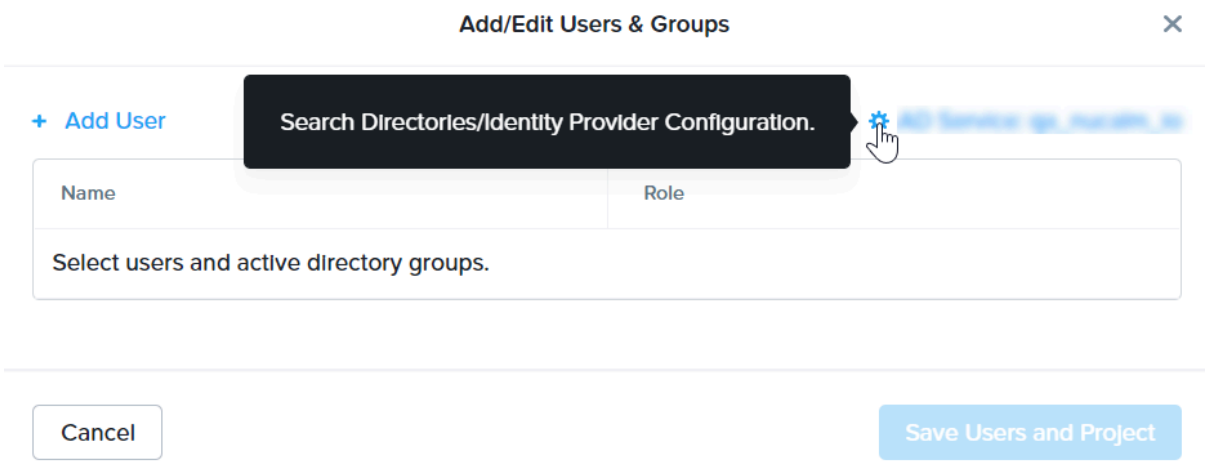

Figure 370: Search Directories

   b. Select the radio button for the directory service or identity provider that you want to use to add users and groups.

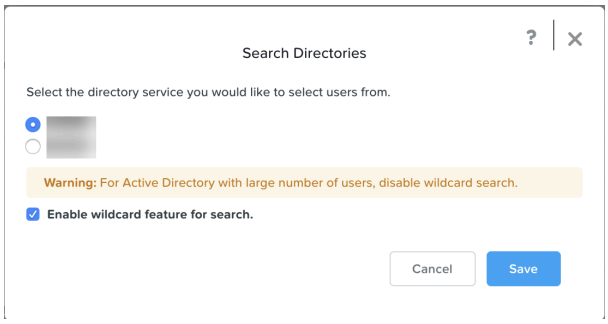

Figure 371: Search Directories Window

   c. Click the **Save** button.

> **Note:** Local users are not supported in a project. You can only add users from your configured directory or identity provider.

6. To add a user or group with a role in the project, do the following.

   a. Click **+ Add User**.
      A blank row is added with the **Name** and **Role** columns.

   b. In the **Name** column, enter the Active Directory name of a user or a group (typically in the form of `name@domain`).

   c. In the **Role** column, select a user role from the list.

      The default value in the **Role** column is **Project Admin**. You can select a value from the list to change the user role.

      The **Allow Collaboration** check box appears when you add the first user to your project. By default, the **Allow Collaboration** check box is checked and enables a project admin or user to view and manage VMs and applications of other users in the same project. If you uncheck the **Allow Collaboration** check box, project admins or users can manage only the VMs and applications that they create. You cannot change

this configuration after you add the first user and save the project. To change the configuration, you need to remove all users from the project.

d. Click the **Save Users and Project** button.
The **Users, Groups & Roles** section on the Overview tile displays the number of users you added to the project.

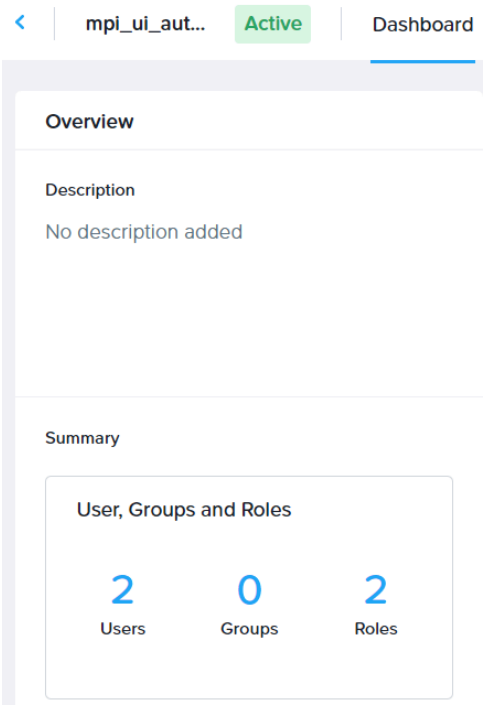

Figure 372: Users, Groups & Roles

Note:

- If you add a group to a project, users in the group might not appear in the project members list until they log in.
- Nested groups (groups within a group) are not supported. For example, if a selected group (Group1) includes a nested group (Group1.1) along with individual names, only individual names are added to the project. The group members of Group 1.1 are not added to the project.

7. To add multiple users and groups to the project, repeat step 6.

Infrastructure in Projects

Infrastructure allows you to add accounts to your project. When Calm is disabled, you can add only the local Nutanix account to your project (see Adding Infrastructure in a Calm-Disabled Prism Central on page 737). However, you can enable Calm and add multiple accounts of the same provider or different providers you configured in Calm to your projects (see Adding Infrastructure in a Calm-Enabled Prism Central on page 739).

Infrastructure supports accounts of the following providers:

- Nutanix
- VMware

- AWS
- Azure
- GCP
- Kubernetes
- Xi

When you add Nutanix accounts to your project, you can allow clusters and their corresponding VLAN subnets. A VLAN subnet is bound to a Prism Element cluster. When you use the VLAN subnet to provision a VM, the VM is placed on that Prism Element.

When Calm is enabled, you can also allow clusters and their corresponding overlay subnets. An overlay subnet can span from a few clusters to all the clusters of your Prism Central. Therefore, when you configure your Nutanix account in your project, you must allow clusters before allowing the subnets.

> **Note:**
>
> - Allowing clusters before their subnets enables you to have projects where you can allow VPCs and their corresponding subnets without allowing any VLAN subnets.
> - You can also define resource quota limits for the Nutanix and VMware accounts within your project. However, to define resource quota limits for the VMware account, you must enable policy engine in Calm. For more information, see Quota Policy Overview on page 773.

Adding Infrastructure in a Calm-Disabled Prism Central

When Calm is disabled in your Prism Central instance, you can add only the local Nutanix account and allow clusters and subnets in the project. You can also define resource quota limits for quota checks.

Procedure

1. In the Prism Central entities menu, click **Administration** > **Projects**.
2. Do one of the following:
   - » Click the **+Create Project** button to create a new project and add the Nutanix account to the project. For more information about creating a project, see Creating a Project on page 729.
   - » Click a project name in the list of existing projects to add the Nutanix account to that project.
3. Click the **Infrastructure** tab.

4. Click the **Add Infrastructure** button and select the local Nutanix account.

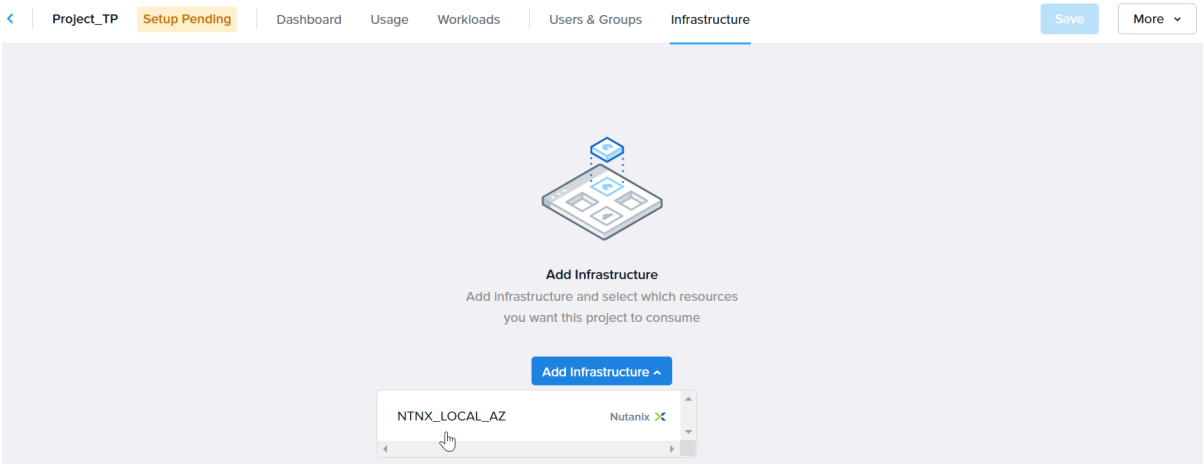

Figure 373: Add Infrastructure

5. Do the following.

   a. Click **Configure Resources**.

   b. On the **Select Clusters and VLANs** tab, select the cluster that you want to allow in the project from the **Select clusters to be added to this project** list.

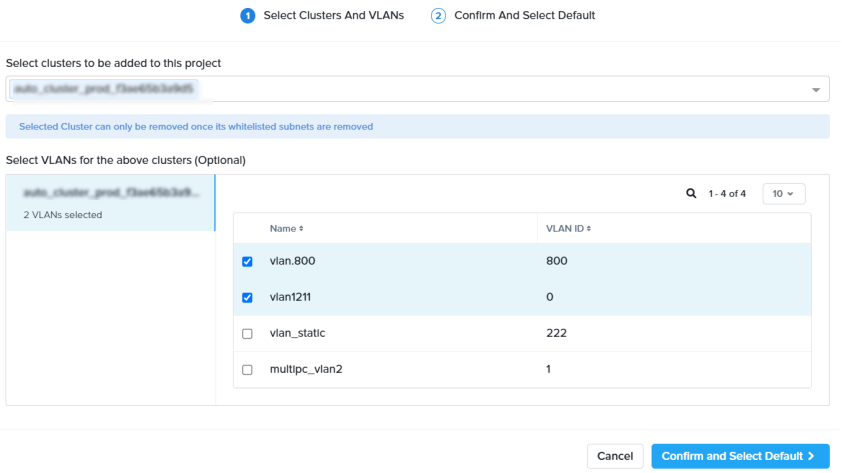

Figure 374: Select Clusters and VLANs

   Selecting a cluster for the account is mandatory. You must select a cluster irrespective of whether you want to allow VLAN subnets or not.

   c. Under Select VLANs for the above clusters section, click the cluster and then **Select VLANs** to view and select the VLANs that you want to allow in the project.

   d. Click **Confirm and Select Default**.

   e. Select a default subnet.

   Default subnets are used when you create a VM from Prism Central.

   f. Click **Confirm**.

6. To define resource quota limits for the Nutanix account, under the Quotas section, enter quota values for **vCPU**, **Storage**, and **Memory**.
7. Click **Save**.
   The Overview tile on the **Dashboard** tab displays the number of accounts you added to the project.

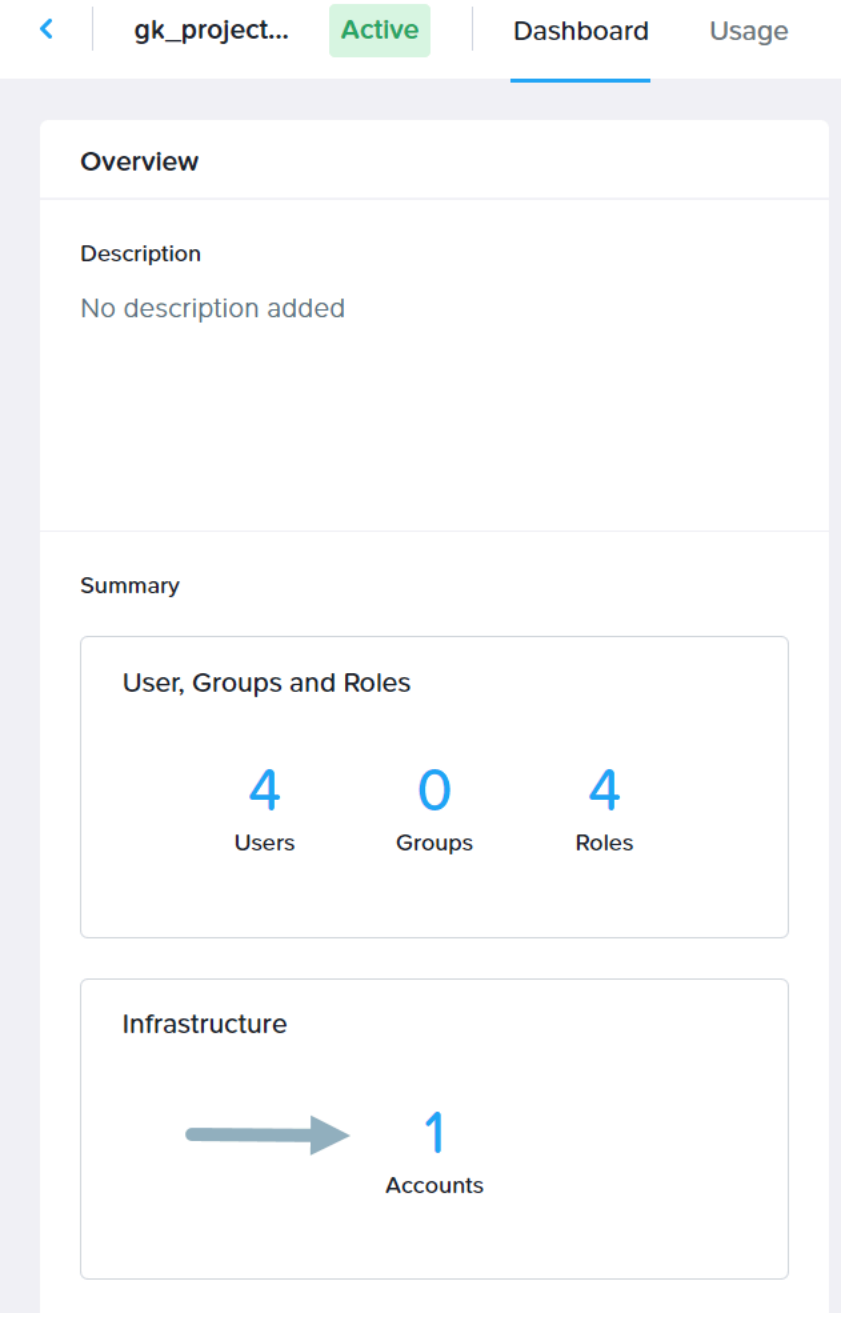

Figure 375: Accounts in a Project

Adding Infrastructure in a Calm-Enabled Prism Central

When Calm is enabled in your Prism Central, you can add multiple provider accounts or credential provider accounts that you configured in Calm to your project. You can also define resource quota limits for quota checks for Nutanix and VMware accounts within the project.

About this task

Use this procedure to add provider accounts or credential provider accounts to your project.

Before you begin

- Ensure that you have configured your provider accounts or credential provider accounts on the **Settings** tab. For information on configuring a provider account, see Provider Account Settings in Calm. For information on configuring a credential provider account, see Credential Settings in Calm.
- For resource quota limit definition, ensure that you have enabled the policy engine on the **Settings** tab. For more details about enabling the policy engine, see Enabling Policy Engine.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » In the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.
2. Do one of the following:
   - » Click the **+Create Project** button to create a new project and add providers to the project. For more information about creating a project, see Creating a Project on page 729.
   - » Click a project name in the list of existing projects to add providers to that project.
3. Click the **Infrastructure** tab.
4. Click the **Add Infrastructure** button and select a provider account or a credential provider account from the list.

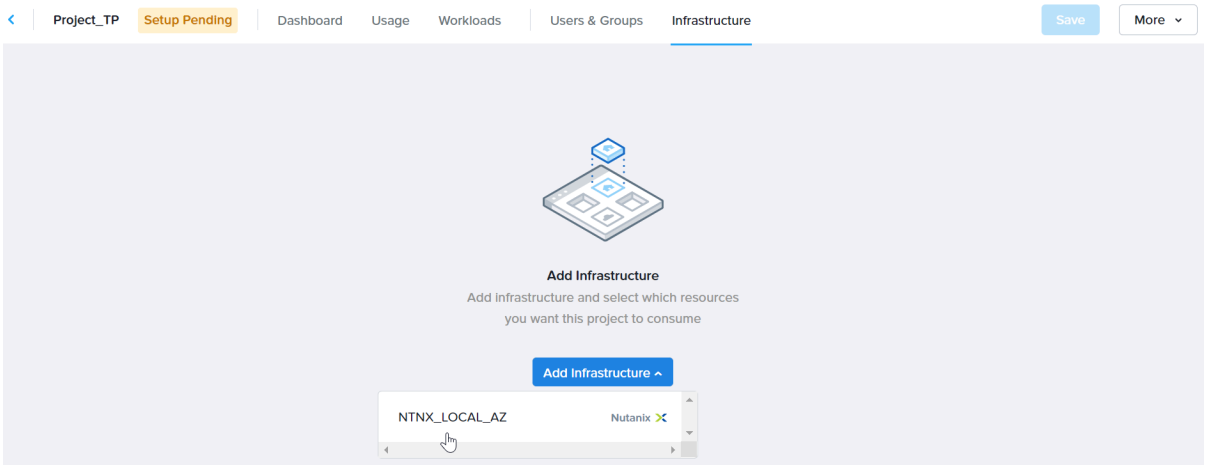

Figure 376: Add Infrastructure

5. If you selected a credential provider account, then go to step 8.

6. If you selected a **Nutanix** account, do the following.

   a. Click **Configure Resources**.

   b. On the **Select Clusters and VLANs** tab, select the cluster that you want to allow in the project from the **Select clusters to be added to this project** list.

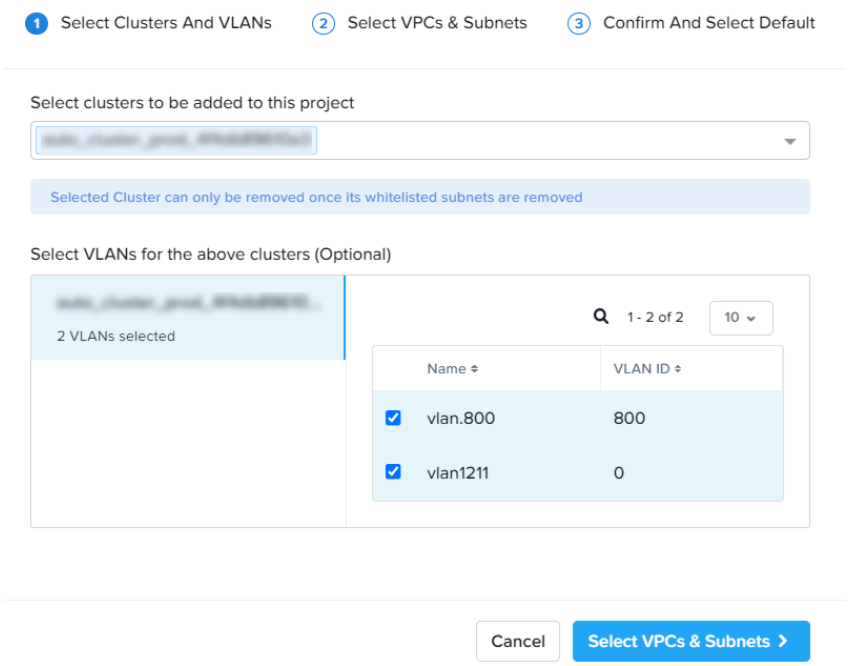

Figure 377: Select Clusters and VLANs

   Selecting a cluster for the account is mandatory. You must select a cluster irrespective of whether you want to allow VLAN subnets or not.

   c. Under Select VLANs for the above clusters section, click **Select VLANs** to view and select the VLANs that you want to allow in the project. This step is optional.

   d. Click the **Select VPCs & Subnets** button.

   The **Select VPCs & Subnets** button or tab appears only on VPC-enabled setups.

   e. On the **Select VPCs & Subnets** tab, select the VPC from the **Select VPCs to view overlay subnets below** list to view the associated subnets. This step is optional.

   The Select overlay subnets section displays the overlay subnets associated with the VPC you selected.

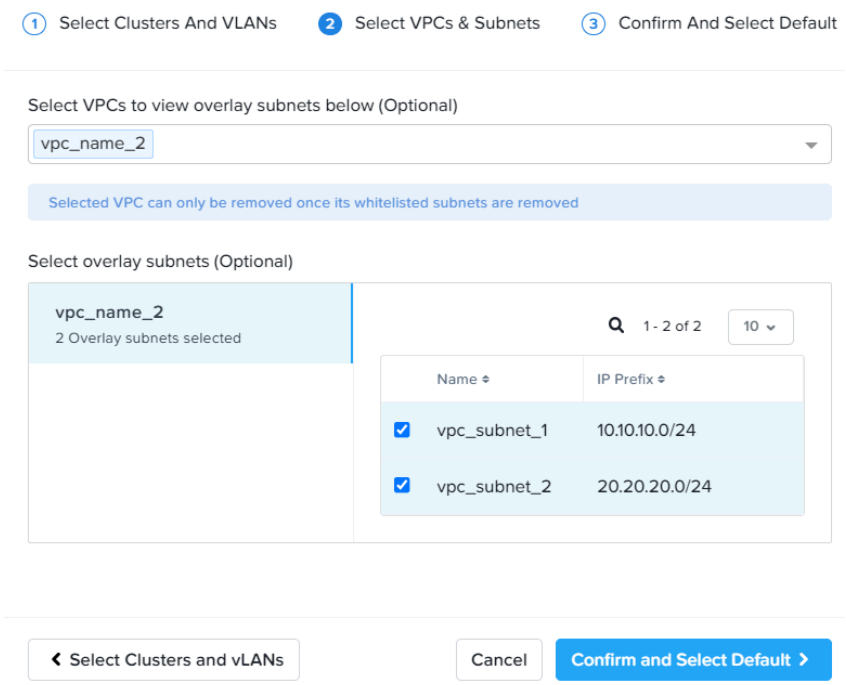

Figure 378: Select VPCs and Subnets

If tunnel is not configured for the selected VPC, you can perform only basic operations, such as VM provisioning, on the VPC. To perform check log-in and orchestration, ensure that you create a tunnel for the VPC. For more information, see Creating VPC Tunnels.

f. Under the Select overlay subnets section, select the overlay subnets that you want to allow in the project. This step is optional.

g. Do one of the following:

   » For the local Nutanix account, click **Confirm and Select Default**, select a default VLAN subnet, view the configuration summary, and then click **Confirm**.

   » For a remote PC account, click **Confirm**, view the configuration summary, and then click **Confirm**.

   The default VLAN subnet is used when you create a virtual machine in Prism Central.

> **Note:** You can select one or more AHV clusters. You can allow one or more subnets per cluster. When you configure a blueprint, only the allowed networks and clusters appear for the user to select during network configuration. If the network selection is a runtime attribute, only the allowed networks and clusters are available to update while launching a blueprint.

7. If you selected a **Nutanix** or **VMware** account, you can define resource quota limits.

   a. Select the **Quotas** check box.
   The **vCPU**, **Memory**, and **Disk** fields are enabled.

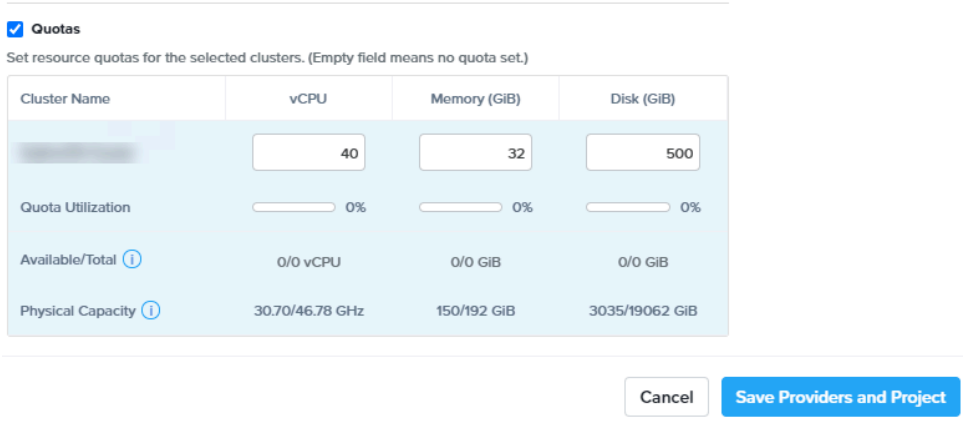

Figure 379: Quota Definition

   b. Enter quota values for **vCPU**, **Memory**, and **Disk** for the selected clusters.

   The **Available/Total** row shows the available resources quota and the total quota allocated to the provider. The **Physical Capacity** row shows the used and total physical capacity. Use these details while defining resource quota limits.

8. Click **Save Accounts and Project**.
   The Overview tile on the **Dashboard** tab displays the number of accounts you added to the project.

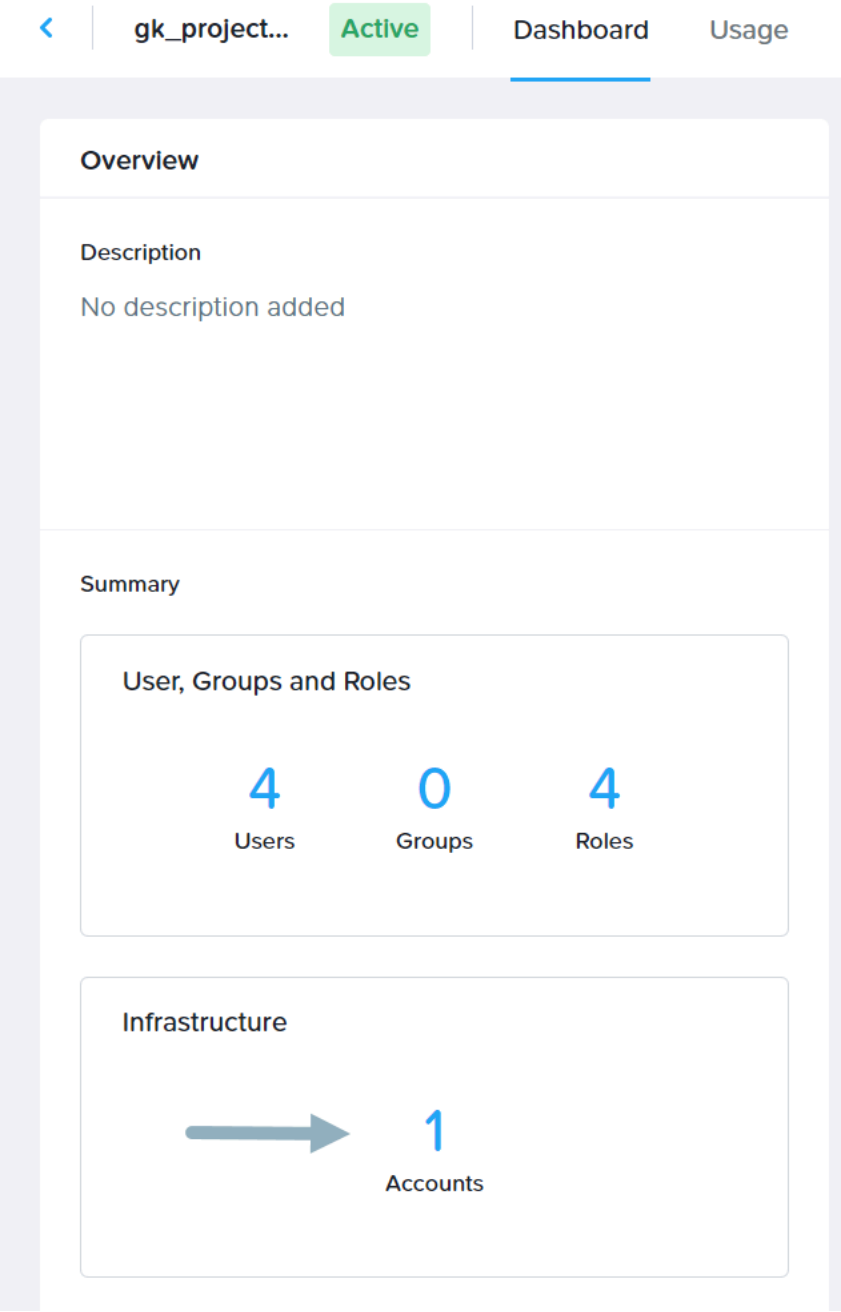

Figure 380: Accounts in a Project

Modifying a Project

You can modify the description, users, infrastructure, or any application management specific details of your project.

About this task

Use this procedure to modify an existing project.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » If Calm is enabled, then in the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.
2. On the Project Summary page, click the project that you want to modify.
   The Project Details page appears. The page includes the same tabs and fields that appeared while creating the project you want to modify.
3. Update the field values as required and then click the **Save** button.

Deleting a Project

You can delete a project that is not associated with any application or blueprint. If the project is already used to create any applications or blueprints, you cannot delete the project. In such cases, a dialog box appears displaying the association of the project with different applications or blueprints.

About this task

Use this procedure to delete a project.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » If Calm is enabled, then in the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.
2. On the Project Summary page, select the check box adjacent to the project that you want to delete.
3. From the **Action** list, select **Delete**.
   Calm verifies the association of the project with any application or blueprint. After successful verification, a confirmation message appears.
4. Click **Delete**.
   The project is deleted from the **Projects** tab.

Environments in Projects

If Calm is enabled in your Prism Central instance, you can configure environments as the subset of your project.

To configure environments, you must add the accounts that you configured in Calm to your project and then configure environments using those accounts. Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project. See Configuring Environments in a Project on page 746.

You use a configured environment either during your blueprint creation or during an application launch. When you select a configured environment while launching an application from the marketplace, the values for the application launch are picked up from the selected environment. See Launching a Blueprint from the Marketplace.

You can add multiple accounts of the same provider or accounts of different providers to your project. You can also configure multiple environments in a project and set one of the environments as the default environment for the project. You can optionally select one environment for each application profile in the blueprint.

Note: Environment is not supported for Kubernetes.

Configuring Environments in a Project

You configure environments as a part of your project creation so that you can use the configured environments when you create blueprints or launch marketplace application blueprints. You can configure multiple environments in your project.

Before you begin

- Ensure that you have enabled Calm in your Prism Central instance. See Enabling and Accessing Calm.
- Ensure that you have configured a project and have added the required accounts to your project. For more information, see Creating a Project on page 729 and Infrastructure in Projects on page 736.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » In the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.
2. On the Projects Summary page, click the project in which you want to configure the environment.
   The **Dashboard** tab appears.
3. Click the **Environments** tab.
4. Click **Create Environment**.
   The Create Environment page appears.

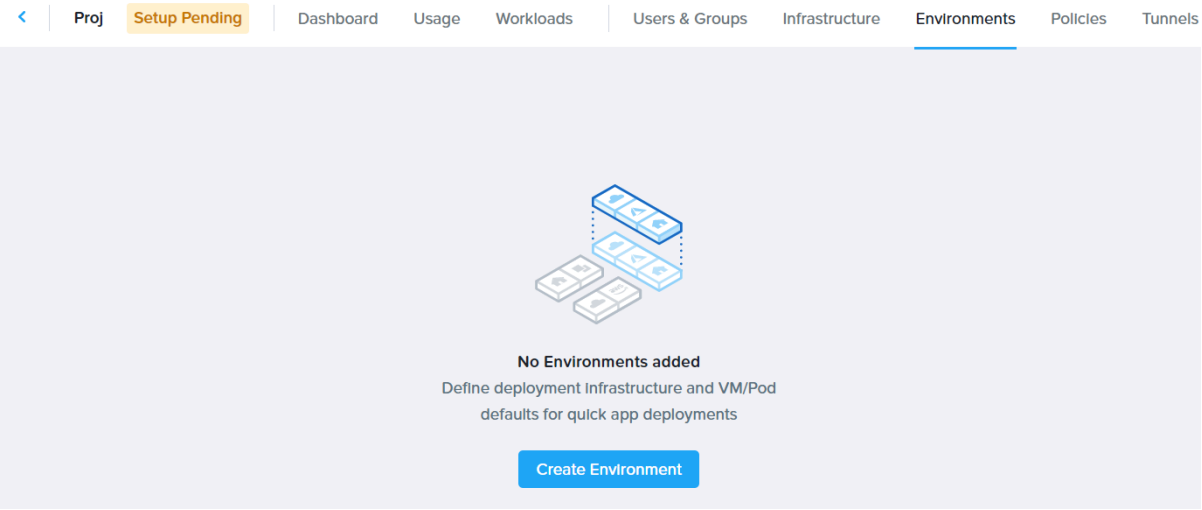

Figure 381: Create Environment

5. On the **General** tab, do the following:

   a. Enter a name for the environment in the **Name** field.

   b. Provide a description for the environment in the **Description** field.

   c. Select **Set as default environment** check box to use the environment as a default environment to launch applications for the project.

      The **Set as default environment** check box is selected by default for the first environment you configure in your project.

   d. Click **Next**.

6. On the **Infrastructure** tab, do the following:

   a. Click the **Select Infrastructure** list, and select an account.

      The **Select Infrastructure** list shows the accounts that you added to your project. For more information, see Infrastructure in Projects on page 736.

      You can click **+ Add Infrastructure** to select and add multiple provider accounts to the environment.

   b. Expand the **VM Configuration** section to configure the virtual machine details for the environment.

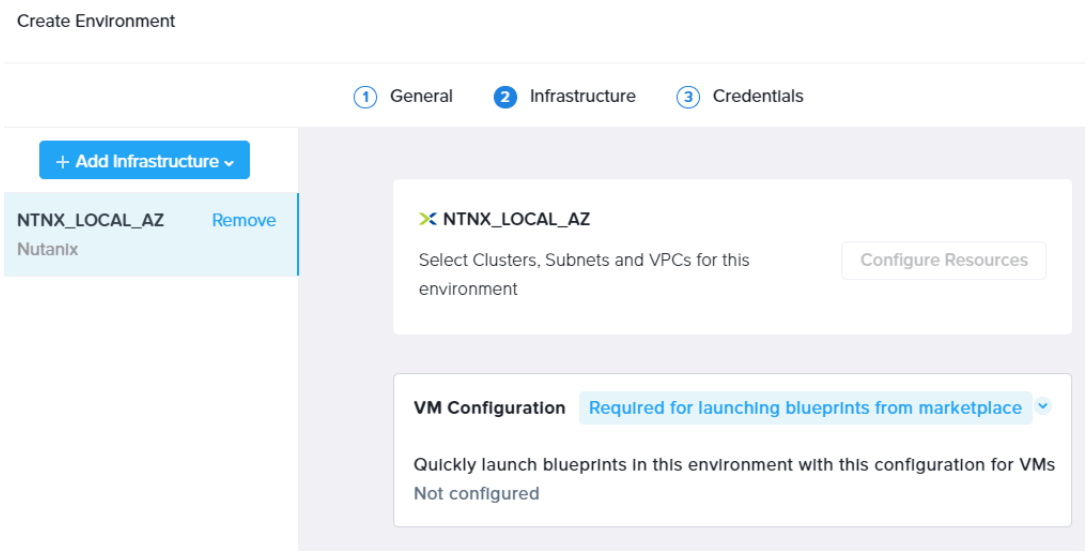

Figure 382: Infrastructure Configuration

   - For more information on how to configure a VM for Nutanix, see Configuring Nutanix Environment on page 751.
   - For more information on how to configure a VM for AWS, see Configuring AWS Environment on page 755.
   - For more information on how to configure a VM for VMware, see Configuring VMware Environment on page 759.
   - For more information on how to configure a VM for GCP, see Configuring GCP Environment on page 763.
   - For more information on how to configure a VM for Azure, see Configuring Azure Environment on page 766.

   c. Click **Next**.

7. On the **Credentials** tab, add credentials for your environment. For more information, see Adding Credentials to the Environment on page 772.

8. Click **Save Environment & Project**.
   The **Overview** tile on the **Dashboard** tab displays the number of environments you configured for the project. You can go to the **Environments** tab to view the details of each environment you configured for your project.

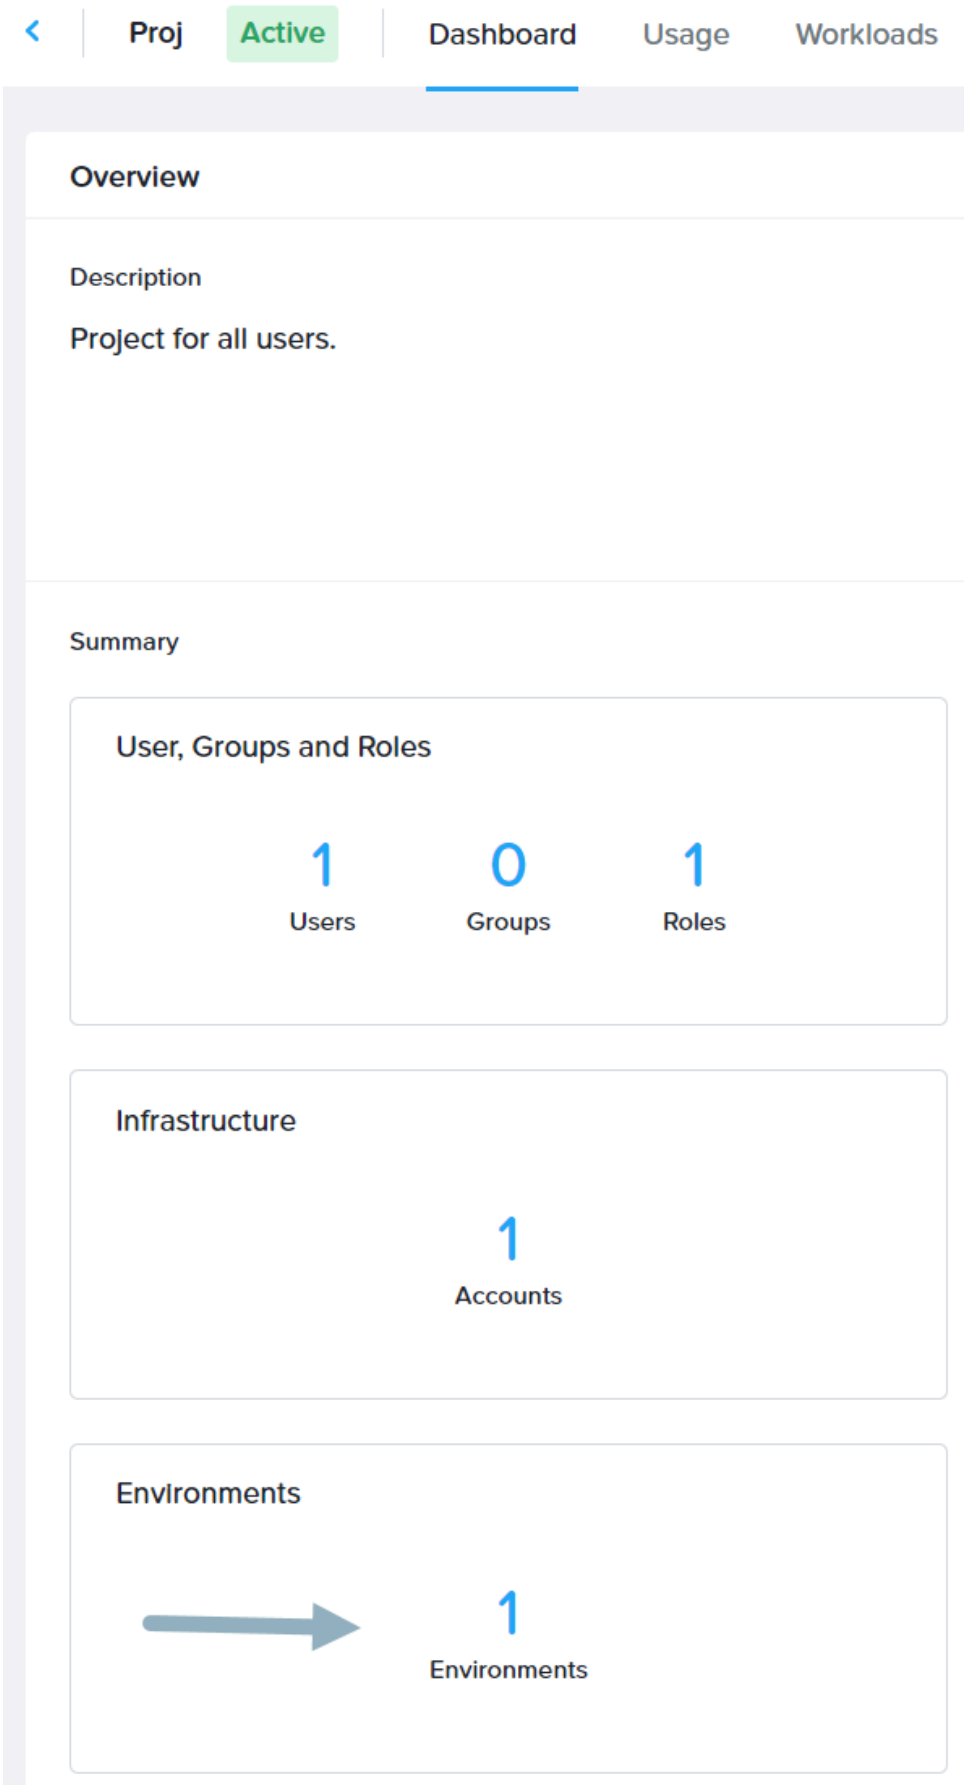

Figure 383: Environments on Dashboard

Configuring Nutanix Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables with a Nutanix account.

Before you begin

Ensure that you have configured a project and selected a Nutanix account. For more information, see Configuring Environments in a Project on page 746.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** list, and add a Nutanix account.
2. Select the account you added in the left pane.

   The Resource Configuration section displays the cluster, VLAN subnets, and overlay subnets you selected in the project for the account.

   When you configure the environment for the first time, the cluster and subnets that you selected for the project is selected for the environment by default. However, you can click the **Configure Resources** button to configure resources specific to the environment. You can select additional subnets for the environment or remove any subnets that you have already selected for the project.

3. In the right pane, expand the **VM Configuration** section and select either **Windows** or **Linux** as the operating system for the VM.

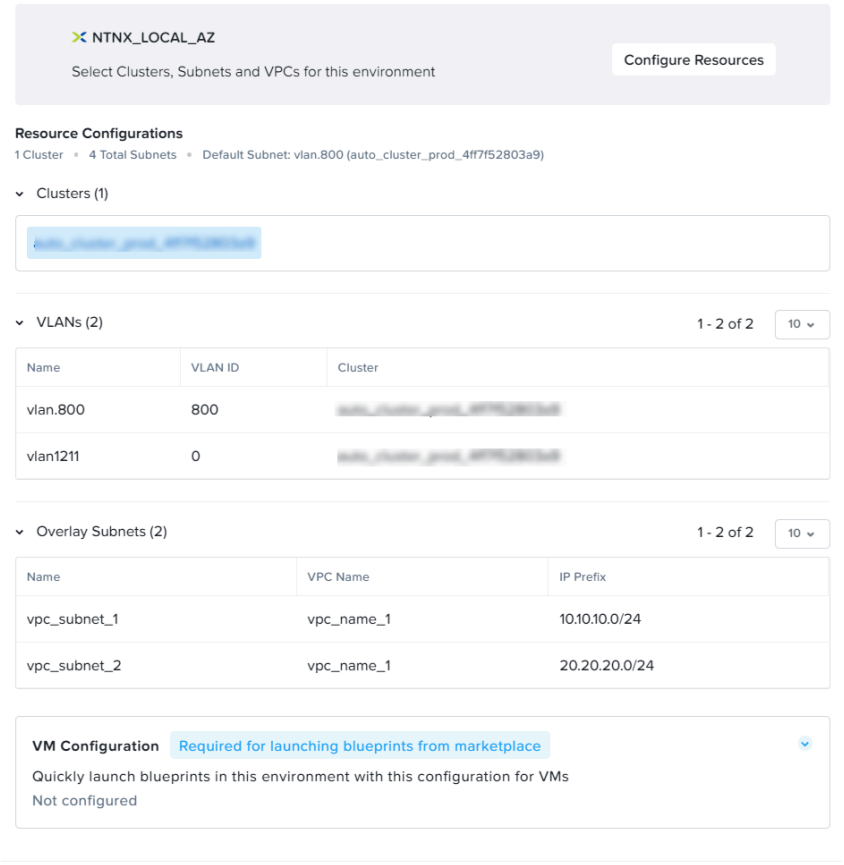

Figure 384: VM Configuration

4. In the **Cluster** list, select the cluster where you want to place the VM.

   The **Cluster** list displays the clusters that you allowed in the project.

   The VLAN subnets have direct association with the cluster. When you select a VLAN subnet under the Network Adapters section, the associated cluster is auto-populated in the **Cluster** list. However, if you intend to use overlay subnets, you must select the cluster in the list.

5. Enter a name of the VM in the **VM Name** field.

   You can use Calm macros to provide a unique name to the VM. For example, `vm-@@{calm_time}@@`. For more information on Calm macros, see Macros Overview.

6. Configure the processing unit of the VM by entering the number of vCPU, cores of each vCPU, and total memory in GB of the VM in the **vCPU**, **cores per vCPU**, and **Memory (GiB)** fields.

7. If you want to customize the default OS properties of the VM, select the **Guest Customization** check box.

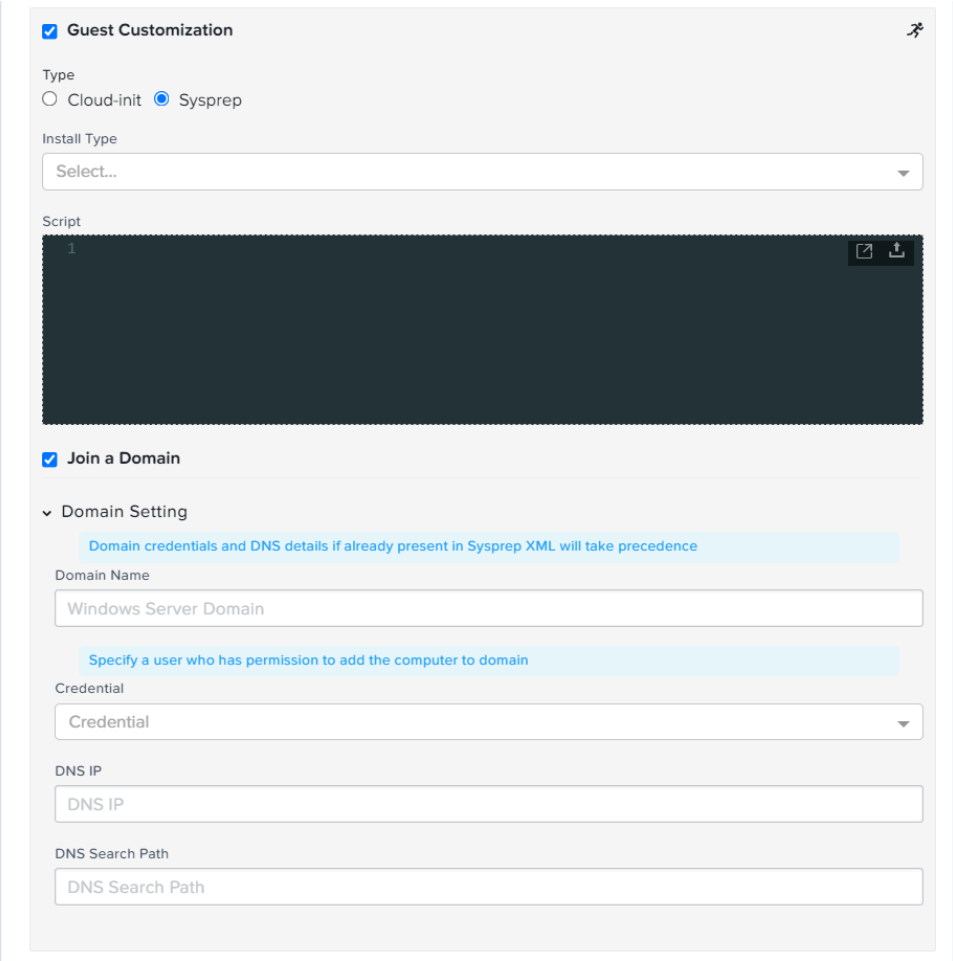

Figure 385: Guest Customization

Guest customization allows you to modify the properties of the VM operating system. You can prevent conflicts that might result due to the deployment of virtual machines with identical settings, such as duplicate VM names or same SID. You can also change the computer name or network settings by using a custom script.

a. Select **Cloud-init** for Linux or **SysPrep** for Windows, and enter or upload the script in the **Script** panel. For Sysprep, you must use double back slash for all escape characters . For example, \\v.

b. For Sysprep script, click **Join a Domain** check box and configure the following fields.

- Enter the domain name of the Windows server in the **Domain Name** field.
- Select a credential for the Windows VM in the **Credentials** list. You can also add new credentials.
- Enter the IP address of the DNS server in the **DNS IP** field.
- Enter the DNS search path for the domain in the **DNS Search Path** field.

8. To add a virtual disk to the VM, click the **+** icon next to the **DISKS** section and do the following.

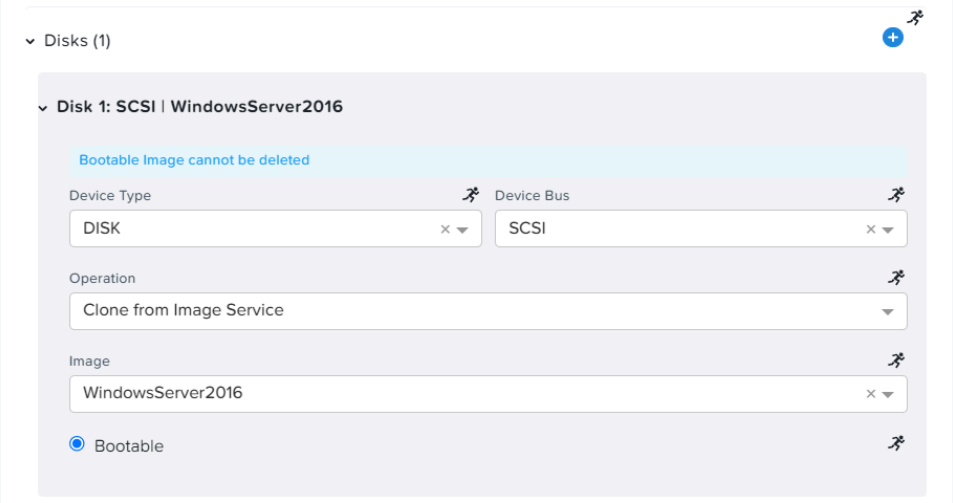

Figure 386: Disks

a. Select the device for the image from the **Device Type** list.
   You can select either **CD-ROM** or **Disk**.

b. Select the device bus from the **Device Bus** list.
   You can select **IDE** or **SATA** for CD-ROM and **SCSI**, **IDE**, **PCI**, or **SATA** for DISK.

c. From the **Operations** list, select one of the following.

   » To allocate the disk memory from the storage container, select **Allocate on Storage Container**.

   » To clone an image from the disk, select **Clone from Image Service**.

d. If you selected **Allocate on Storage Container**, enter the disk size in GB in the **Size (GiB)** field.

e. If you selected **Clone from Image Service**, select the image you want to add to the disk in the **Image** field.
   All the images that you uploaded to Prism Central are available for selection. For more information about image configuration, see *Image Management* section in the *Prism Central* guide.

f. Select the **Bootable** check box for the image that you want to use to start the VM.

> **Note:** You can add more than one disk and select the disk with which you want to boot up the VM.

9. Under the **Boot Configuration** section, select a firmware type to boot the VM.

   » To boot the VM with legacy BIOS firmware, select **Legacy BIOS**.

   » To boot the VM with UEFI firmware, select **UEFI**. UEFI firmware supports larger hard drives, faster boot time, and provides more security features.

10. (For GPU-enabled clusters only) To configure a vGPU, click the **+** icon under the **vGPUs** section and do the following:

    a. From the **Vendor** list, select the GPU vendor.

    b. From the **Device ID** list, select the device ID of the GPU.

    c. From the **Mode** list, select the GPU mode.

11. Under the **Categories** section, select a category in the **Key: Value** list.

    Use this option to tag your VM to a defined category in Prism Central. The list options are available based on your Prism Central configuration. If you want to protect your application by a protection policy, select the category defined for the policy in your Prism Central. Categories list is available only for Nutanix.

12. To add a network adapter, click the **+** icon next to the Network Adapters (NICS) field.

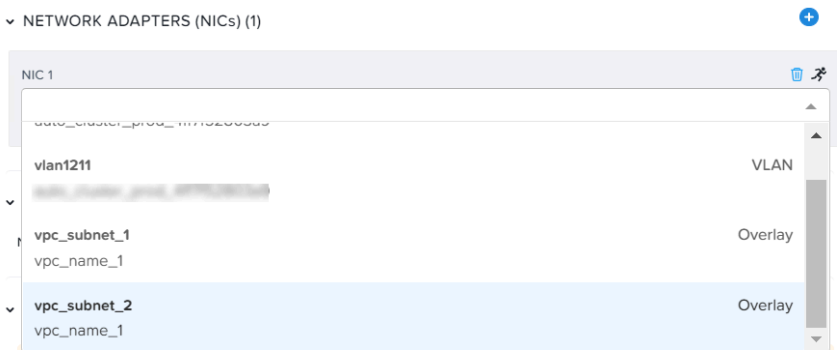

Figure 387: NIC

The **NIC** list shows all the VLAN and overlay subnets. The VLAN subnets have direct association with the cluster. Therefore, when you select a VLAN subnet, the associated cluster is auto-populated in the **Cluster** list.

The NICs of a VM can either use VLAN subnets or overlay subnets. For example, if you select an overlay subnet in NIC 1 and then add NIC 2, the NIC 2 list displays only the overlay subnets.

If you select a VLAN subnet in NIC 1, any subsequent VLAN subnets belong to the same cluster. Similarly, if you select an overlay subnet, all subsequent overlay subnets belong to the same VPC.

13. Configure the connection in your environment. For more information, see Configuring Check Log-In for your Environment on page 771.

14. Click **Next**.

15. Add credentials for the environment. For more information, see Adding Credentials to the Environment on page 772. This step is optional.

16. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for Nutanix or launching a blueprint.

Configuring AWS Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables with an AWS account.

Before you begin

Ensure that you have configured a project and selected an AWS account. For more information, see Configuring Environments in a Project on page 746.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** list, and add an AWS account.
2. Select the account you added in the left pane.

3. In the right pane, expand the **VM Configuration** section and select either **Windows** or **Linux** as the operating system for the VM.

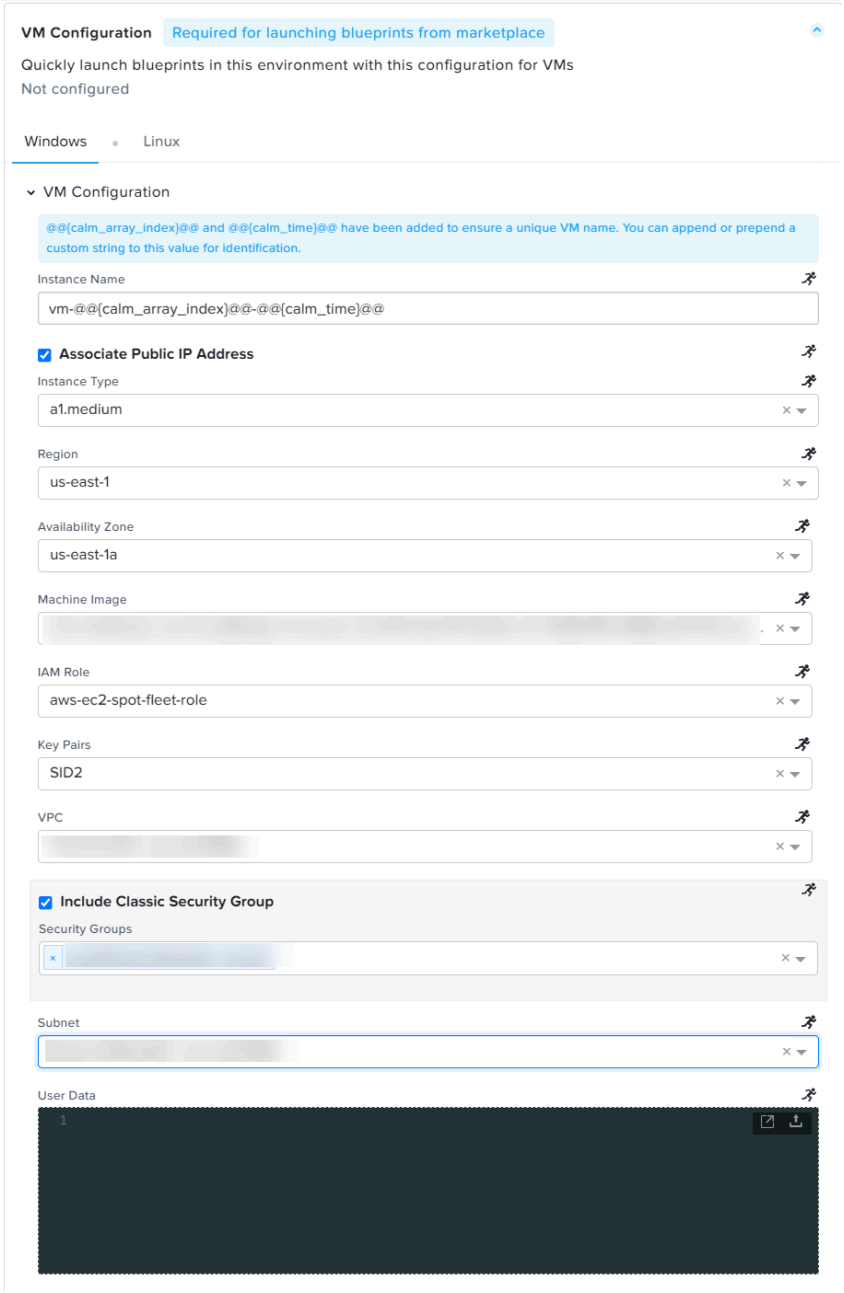

Figure 388: VM Configuration

4. Enter the name of the instance in the **Instance Name** field.

   This field is pre-populated with a macro as suffix to ensure name uniqueness. The service provider uses this name as the VM name.

5. Select the **Associate Public IP Address** check box to associate a public IP address with your AWS instance.

   If you do not select the **Associate Public IP Address** check box, ensure that the AWS account and Calm are on the same network for the scripts to run.

6. Select an AWS instance type from the **Instance Type** list.

   Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload.

   The list displays the instances that are available in the AWS account. For more information, see AWS documentation.

7. Select the region from the **Region** list and configure the following.

   Note: The list displays the regions which are selected while configuring the AWS setting.

   a. Select the availability zone from the **Availability Zone** list.

      An Availability Zone is one or more discrete data centers with redundant power, networking, and connectivity in an AWS region. Availability Zones allow you to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.

   b. Select the machine image from the **Machine Image** list.

      An Amazon Machine Image is a special type of virtual appliance that is used to create a virtual machine within the Amazon Elastic Compute Cloud. It serves as the basic unit of deployment for services delivered using EC2.

   c. Select the IAM role from the **IAM Role** list.

      An IAM role is an AWS Identity and Access Management entity with permissions to make AWS service requests.

   d. Select the key pairs from the **Key Pairs** list.

      A key pair (consisting of a private key and a public key) is a set of security credentials that you use to prove your identity when connecting to an instance.

   e. Select the VPC from the **VPC** list.

      Amazon Virtual Private Cloud (Amazon VPC) allows you to provision a logically isolated section of the AWS cloud where you can launch AWS resources in your defined virtual network.

      - Select the **Include Classic Security Group** check box to enable security group rules.
      - Select security groups from the **Security Groups** list.

8. Enter or upload the AWS user data in the **User Data** field.

9. Enter the AWS tags in the **AWS Tags** field.

   AWS tags are key and value pair to manage, identify, organize, search for, and filter resources. You can create tags to categorize resources by purpose, owner, environment, or other criteria.

10. Under the **Storage** section, configure the following to boot the AWS instance with the selected image.

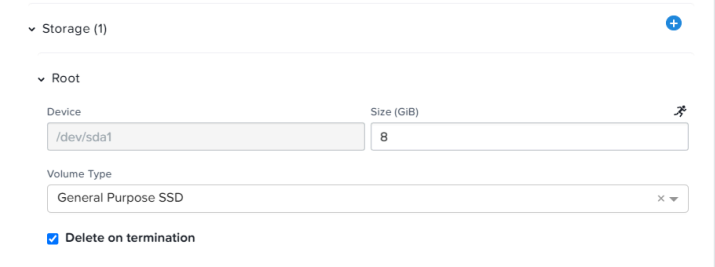

Figure 389: Storage

a. In the **Device** field, select the device to boot the AWS instance. The available options are based on the image you have selected.

b. In the **Size (GiB)** field, enter the required size for the bootable device.

c. In the **Volume Type** list, select the volume type. You can select either **General Purpose SSD**, **Provisioned IOPS SSD**, and **EBS Magnetic HDD**.

   For more information on the volume types, see AWS documentation.

d. (Optional) Select the **Delete on termination** check box to delete the storage when the instance is terminated.

You can also add more secondary storages by clicking the **+** icon next to the **Storage** section.

11. Configure the connection in your environment. For more information, see Configuring Check Log-In for your Environment on page 771.

12. Click **Next**.

13. (Optional) Add credentials for the environment. For more information, see Adding Credentials to the Environment on page 772.

14. Click **Save Environment & Project**.
    The **Environments** tile on the **Overview** tab displays the number of environments you configured for the project. You can go to the **Environments** tab to view the details of the environment you configured for your project.

What to do next

You can use the environment details while configuring a blueprint for AWS or launching a blueprint.

Configuring VMware Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables for VMware.

Before you begin

Ensure that you have configured a project and selected a VMware account. For more information, see Configuring Environments in a Project on page 746.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** list, and add a VMware account.
2. Select the account you added in the left pane.
3. In the right pane, expand the **VM Configuration** section, and select either **Windows** or **Linux** as the operating system for the VM.

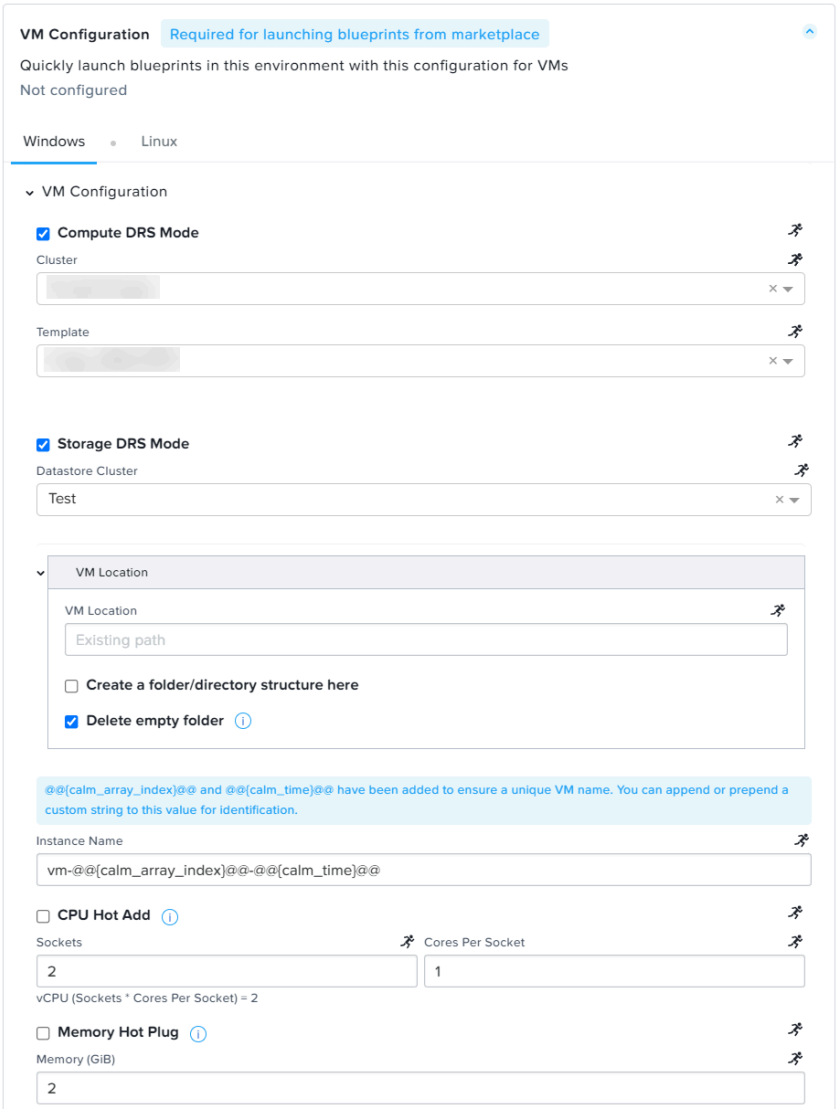

Figure 390: VM Configuration

4. Select the **Compute DRS Mode** check box to enable load sharing and automatic VM placement.

   Distributed Resource Scheduler (DRS) is a utility that balances computing workloads with available resources in a virtualized environment. For more information about DRS mode, see the *VMware documentation*.

   » If you selected **Compute DRS Mode**, then select the cluster where you want to host your VM from the **Cluster** list.

   » If you have not selected **Compute DRS Mode**, then select the host name of the VM from the **Host** list.

5. Do one of the following:

   » Select the **VM Templates** radio button and then select a template from the **Template** list.

   Templates allow you to create multiple virtual machines with the same characteristics, such as resources allocated to CPU and memory or the type of virtual hardware. Templates save time and avoid errors when configuring settings and other parameters to create VMs. The VM template retrieves the list options from the configured vCenter.

   Note:

   - Install the VMware Tools on the Windows templates. For Linux VMs, install *Open-vm-tools* or *VMware-tools* and configure the *Vmtoolsd* service for automatic start-up.
   - Support for *Open-vm-tools* is available. When using *Open-vm-tools*, install Perl for the template.
   - Do not use SysPrepped as the Windows template image.
   - If you select a template that has unsupported version of VMware Tools, then a warning appears stating VMware tool or version is unsupported and could lead to VM issues.
   - You can also edit the NIC type when you use a template.

   For more information, refer to VMware KB articles.

   » Select the **Content Library** radio button, a content library in the **Content Library** list, and then select an OVF template or VM template from the content library.

   A content library stores and manages content (VMs, vApp templates, and other types of files) in the form of library items. A single library item can consist of one file or multiple files. For more information about the vCenter content library, see the *VMware Documentation*.

   Caution: Content Library support is currently a technical preview feature in Calm. Do not use any technical preview features in a production environment.

6. If you want to use the storage DRS mode, then select the **Storage DRS Mode** check box and a datastore cluster from the **Datastore Cluster** list.

   The datastore clusters are referred as storage pod in vCenter. A datastore cluster is a collection of datastores with shared resources and a shared management interface.

7. If you do not want to use storage DRS mode, then do not select the **Storage DRS Mode** check box, and select a datastore from the **Datastore** list.

8. In the **VM Location** field, specify the location of the folder in which the VM must be created when you deploy the blueprint. Ensure that you specify a valid folder name already created in your VMware account.

   To create a subfolder in the location you specified, select the **Create a folder/directory structure here** check box and specify a folder name in the **Folder/Directory Name** field.

   Note: Calm gives preference to the VM location specified in the environment you select while launching an application. For example, you specify a subfolder structure as the VM location in the blueprint and the top-level folder in the environment. When you select this environment while launching your application, Calm considers the VM location you specified in the environment and creates the VM at the top-level folder.

   Select the **Delete empty folder** check box to delete the subfolder you create within the specified location, in case the folder does not contain any VM resources. This option helps you to keep a clean folder structure.

9. Enter the instance name of the VM in the **Instance Name** field.
   This field is pre-populated with a macro as suffix to ensure name uniqueness. The service provider uses this name as the VM name.
10. Select the **CPU Hot Add** check box if you want to increase the VCPU count of a running VM.
    Support for **CPU Hot Add** depends on the Guest OS of the VM.
11. Update the **vCPUs** and **Core per vCPU** count.
12. Select the **Memory Hot Plug** check box if you want to increase the memory of a running VM.
    Support for **Memory Hot Plug** depends on the Guest OS of the VM.
13. Update the memory in the **Memory** field.
14. Under **Controller**, click **+** to add the type of controller.
    You can select either SCSI or SATA controller. You can add up to three SCSI and four SATA controllers.
15. Under the **Disks** section, click the **+** icon to add vDisks and do the following:
    a. Select the device type from the **Device Type** list.
       You can either select **CD-ROM** or **DISK**.
    b. Select the adapter type from the **Adapter Type** list.
       You can select **IDE** for CD-ROM.
       You can select **SCSI**, **IDE**, or **SATA** for DISK.
    c. Enter the size of the disk in GiB.
    d. In the **Location** field, select the disk location.
    e. If you want to add a controller to the vDisk, select the type of controller in the **Controller** list to attach to the disk.

       Note: You can add either SCSI or SATA controllers. The available options depend on the adapter type.

    f. In the **Disk mode** list, select the type of the disk mode. Your options are:
       - **Dependent**: Dependent disk mode is the default disk mode for the vDisk.
       - **Independent - Persistent**: Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
       - **Independent - Nonpersistent**: Changes to disks in nonpersistent mode are discarded when you shut down or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you shut down or reset.
16. Under the **Tags** section, select tags from the **Category: Tag pairs** field.
    You can assign tags to your VMs so you can view the objects associated with your VMs in your VMware account. For example, you can create a tag for a specific environment and assign the tag to multiple VMs. You can then view all the VMs that are associated with the tag.
17. (Optional) If you want to customize the default OS properties of the VM, then click the **Enable** check box under **VM Guest Customization** and select a customization from the **Predefined Guest Customization** list.

18. If you do not have any predefined customization available, select **None** and do the following.

    a. Select **Cloud-init** or **Custom Spec**.

    b. If you selected **Cloud-init**, enter or upload the script in the **Script** field.

    c. If you have selected **Custom Spec**, enter the network details for the VM in the following fields:

    - Enter the hostname in the **Hostname** field.
    - Enter the domain in the **Domain** field.
    - Select timezone from the **Timezone** list.
    - Select **Hardware clock UTC** check box to enable hardware clock UTC.
    - Click the **+** icon to add network settings.
    - To automatically configure DHCP server, enable the **Use DHCP** check box and then skip to the **DNS Setting** section.
    - Enter a name for the network configuration you are adding to the VM in the **Setting name** field. Settings name is the saved configuration of your network that you want to connect to your VM.
    - Enter values in the **IP Address**, **Subnet Mask**, **Default Gateway**, and **Alternative Gateway** fields.
    - Under the **DNS Settings** section, enter values in the **DNS Primary**, **DNS Secondary**, **DNS Tertiary**, and **DNS Search Path**.

19. Configure the connection in your environment. For more information, see Configuring Check Log-In for your Environment on page 771.

20. Click **Next**.

21. (Optional) Add credentials for the environment. For more information, see Adding Credentials to the Environment on page 772.

22. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for VMware or launching a blueprint.

Configuring GCP Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables for GCP.

Before you begin

Ensure that you have configured a project and selected a GCP account. For more information, see Configuring Environments in a Project on page 746.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** list, and add a GCP account.

2. Select the account you added in the left pane.
3. In the right pane, expand the **VM Configuration** section, and select either **Windows** or **Linux** as the operating system for the VM.

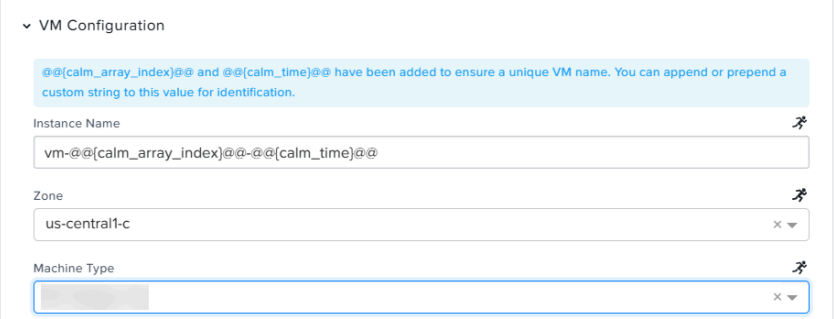

Figure 391: VM Configuration

4. Under **VM Configuration**, enter the instance name of the VM in the **Instance Name** field. This field is pre-populated with macro as suffix to ensure name uniqueness.
5. Select the zone from the **Zone** list.

   Zone is a physical location where you can host the VM.
6. Select machine type from the **Machine Type** list.

   The machine types are available based on your zone. A machine type is a set of virtualized hardware resources available to a virtual machine (VM) instance, including the system memory size, virtual CPU (vCPU) count, and persistent disk limits. In Compute Engine, machine types are grouped and curated by families for different workloads.
7. Under the **Disks** section, click the **+** icon to add a disk.

   You can also mark the added vDisks runtime editable so you can add, delete, or edit the vDisks while launching the blueprint. For more information about runtime editable attributes, see Runtime Variable Overview.
8. To use an existing disk configuration, select the **Use existing disk** check box, and then select the persistent disk from the **Disk** list.

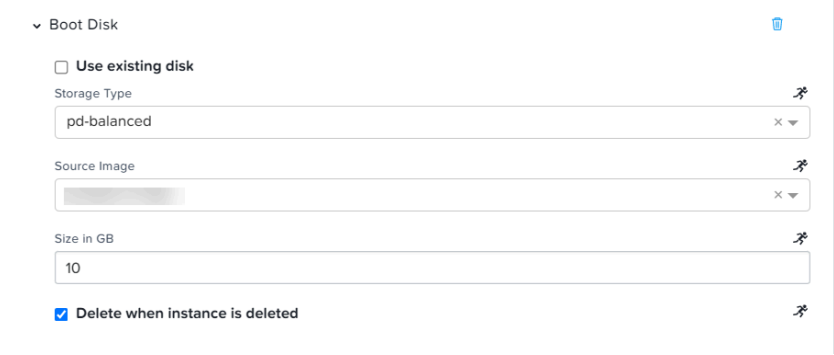

Figure 392: Disks

9. If you have not selected the **Use existing disk** check box, then do the following:

   a. Select the type of storage from the **Storage Type** list. The available options are as follows.

      » **pd-balanced**: Use this option as an alternative to SSD persistent disks with a balanced performance and cost.

      » **pd-extreme**: Use this option to use SSD drives for high-end database workloads. This option has higher maximum IOPS and throughput and allows you to provision IOPS and capacity separately.

      » **pd-ssd**: Use this option to use SSD drives as your persistent disk.

      » **pd-standard**: Use this option to use HDD drives as your persistent disk.

      The persistent disk types are durable network storage devices that your instances can access like physical disks in a desktop or a server. The data on each disk is distributed across several physical disks.

   b. Select the image source from the **Source Image** list.

      The images available for your selection are based on the selected zone.

   c. Enter the size of the disk in GB in the **Size in GB** field.

   d. To delete the disk configuration after the instance is deleted, select the **Delete when instance is deleted** check box under the **Disks** section.

10. To add a blank disk, click the **+** icon under the **Blank Disks** section and configure the blank disk.

11. To add networking details to the VM, click the **+** icon under the **Networking** section.

12. To configure a public IP address, select the **Associate Public IP address** check box and configure the following fields.

    a. Select the network from the **Network** list and the sub network from the **Subnetwork** list.

    b. Enter a name of the network in the **Access configuration Name** field and select the access configuration type from the **Access configuration type** list.

       These fields appear when you select the **Associate public IP Address** check box.

13. Under the **SSH Key** section, click the **+** icon and enter or upload the username key data in the **Username** field.

14. Select **Block project-wide SSH Keys** to enable blocking project-wide SSH keys.

15. Under the **Management** section, do the following:

    a. Enter the metadata in the **Metadata** field.

    b. Select the security group from the **Network Tags** list.

       Network tags are text attributes you can add to VM instances. These tags allow you to make firewall rules and routes applicable to specific VM instances.

    c. Enter the key-value pair in the **Labels** field.

       A label is a key-value pair that helps you organize the VMs created with GCP as the provider. You can attach a label to each resource, then filter the resources based on their labels.

16. Under the **API Access** section, do the following:

    a. Specify the service account in the **Service Account** field.

    b. Under Scopes, select **Default Access** or **Full Access**.

17. Configure the connection in your environment. For more information, see Configuring Check Log-In for your Environment on page 771.

18. Click **Next**.

19. (Optional) Add credentials for the environment. For more information, see Adding Credentials to the Environment on page 772.

20. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for GCP or launching a blueprint.

Configuring Azure Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables for Azure.

Before you begin

- Ensure that the following entities are already configured in the Azure account.
  - Resource group
  - Availability set
  - Network security group
  - Virtual network
  - Vault certificates
- Ensure that you have configured a project and selected an Azure account. For more information, see Configuring Environments in a Project on page 746.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** list, and add a Azure account.

2. Select the account you added in the left pane.

3. In the right pane, expand the **VM Configuration** section, and select either **Windows** or **Linux** as the operating system for the VM.

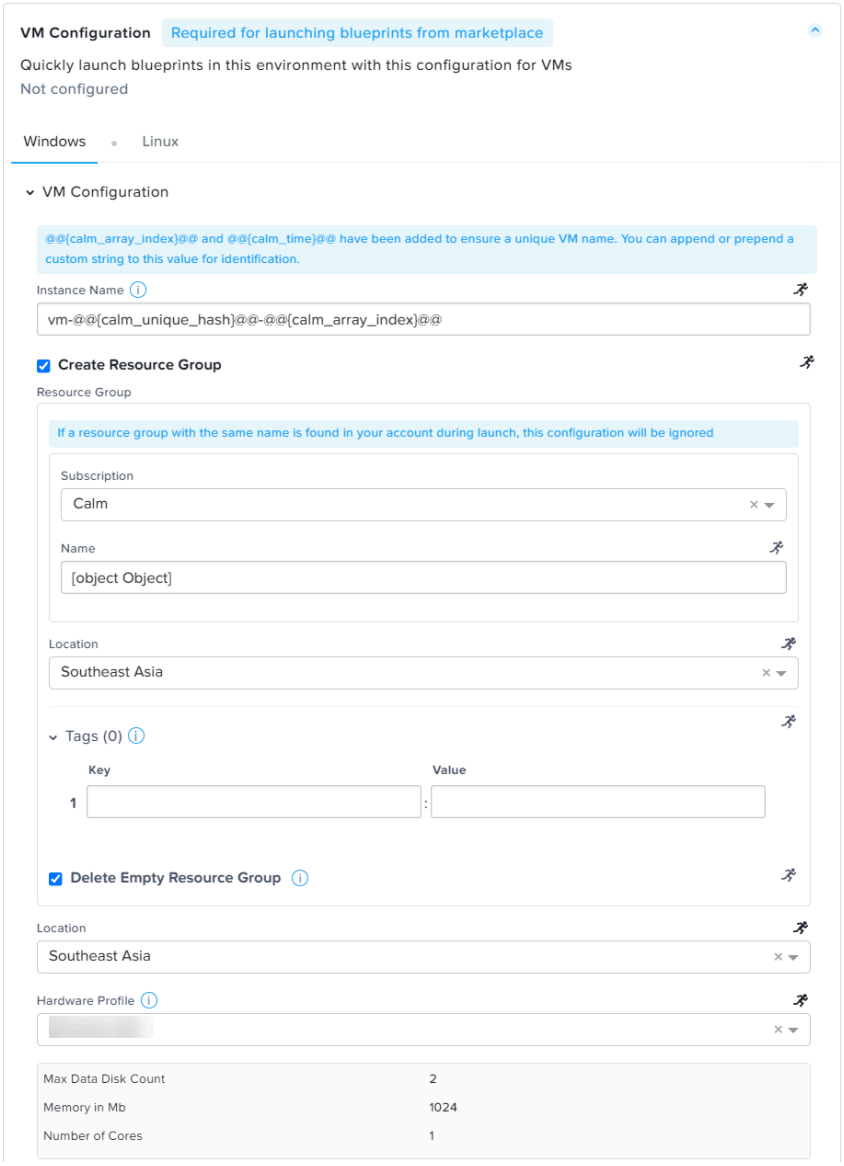

Figure 393: VM Configuration

4. Under **VM Configuration**, enter the instance name of the VM in the **Instance Name** field.

   This field is pre-populated with a macro as suffix to ensure name uniqueness. The service provider uses this name as the VM name.

5. Select a resource group from the **Resource Group** list or select the **Create Resource Group** check box to create a resource group.

   Each resource in Azure must belong to a resource group. A resource group is simply a logical construct that groups multiple resources together so you can manage the resources as a single entity. For example, you can create or delete resources as a group that share a similar life cycle, such as the resources for an n-tier application.

   The **Resource Group** list displays the resource groups that are associated with the subscriptions you selected in your Azure account. In case you have not selected any subscriptions, Calm considers all the subscriptions that

are available in the Azure service principal to display the resource groups. Each resource group in the list also displays the associated subscription.

6. If you selected a resource group from the **Resource Group** list, then do the following:

   a. Select the geographical location of the datacenter from the **Location** list.

   b. Select **Availability Sets** or **Availability Zones** from the **Availability Option** list.

      You can then select an availability set or availability zone. An availability set is a logical grouping capability to ensure that the VM resources are isolated from each other to provide High Availability if deployed within an Azure datacenter. An availability zone allows you to deploy your VM into different datacenters within the same region.

   c. Select the hardware profile as per your hardware requirements from the **Hardware Profile** list.

      The number of data disks and NICs depends upon the selected hardware profile. For information about the sizes of Windows and Linux VMs, see Windows and Linux Documentation.

7. If you selected the **Create Resource Group** check box to create a resource group, then do the following:

   a. Select a subscription associated to your Azure account in the **Subscription** field.

   b. Enter a unique name for the resource group in the **Name** field.

   c. Select the geographical location of the datacenter that you want to add to the resource group in the **Location** list.

   d. Under **Tags**, enter a key and value pair in the **Key** and **Value** fields respectively.

      Tags are key and value pairs that enable you to categorize resources. You can apply a tag to multiple resource groups.

   e. If you want to automatically delete a resource group that has empty resources while deleting an application, click the **Delete Empty Resource Group** check box.

   f. Specify the location and hardware profile.

8. Under the **Secrets** section, click the **+** icon and do the following:

   a. Enter a unique vault ID in the **Vault ID** field.

      These certificates are installed on the VM.

   b. Under **Certificates**, click the **+** icon.

   c. Enter the URL of the configuration certificate in the **URL** field.

      The URL of the certificate is uploaded to key vault as a secret.

   d. Enter the certificate store for the VM in the **Store** field.

      - For Windows VMs, specify the certificate store on the virtual machine to which the certificate is added. The specified certificate store is implicitly created in the LocalMachine account.
      - For Linux VMs, the certificate file is placed under the /var/lib/waagent directory with the file name <UppercaseThumbprint>.crt for the X509 certificate file and <UppercaseThumbpring>.prv for private key. Both of these files are .pem formatted.

9. (For Windows) Select the **Provision Windows Guest Agent** check box.

   This option indicates whether or not to provision the virtual machine agent on the virtual machine. When this property is not specified in the request body, the default behavior is to set it to true. This ensures that the VM Agent is installed on the VM, and the extensions can be added to the VM later.

10. (For Windows) To indicate that the VM is enabled for automatic updates, select the **Automatic OS Upgrades** check box.

11. Under the **Additional Unattended Content** section, click the **+** icon and do the following:

    a. Select a setting from the **Setting Name** list.
    You can select **Auto Logon** or **First Logon Commands**.

    > **Note:** Guest customization is applicable only on images that allows or support guest customization.

    b. Enter or upload the xml content. See Sample Auto Logon and First Logon Scripts.

12. Under the **WinRM Listeners** section, click the **+** icon and do the following:

    a. Select the protocol from the **Protocol** list.
    You can select **HTTP** or **HTTPS**.

    b. If you selected HTTPS, then select the certificate URL from the **Certificate URL** list.

13. Under the **Storage Profile** section, select the **Use Custom Image** check box to use a custom VM image created in your subscription.
    You can then select a custom image or publisher-offer-SKU-version from the **Custom Image** list.

14. Under the **VM Image Details** section, select an image type in the **Source Image Type** list.
    You can select **Marketplace**, **Subscription**, or **Shared Image Gallery**.
    Do one of the following:

    » If you selected **Marketplace**, then specify the publisher, offer, SKU, and version for the image.

    » If you selected **Subscription**, then select the custom image.

    » If you selected **Shared Image Gallery**, then select the gallery and the image.

15. Under the **OS Disk Details** section, do the following:

Figure 394: OS Disk Details

a. Select the storage type from the **Storage Type** list.
   You can select **Standard HDD**, **Standard SSD**, or **Premium SSD**.

b. Select a disk storage account from the **Disk Storage** list.
   This field is available only when the **Use Custom Image** check box is enabled.

c. Select disk caching type from the **Disk Caching Type** list.
   You can select **None**, **Read-only**, or **Read write**.

d. Select disk create option from the **Disk Create Option** list.
   You can select **Attach**, **Empty**, or **From Image**.

16. Under the **Network Profile** section, add NICs as per your requirement and do the following for each NIC:

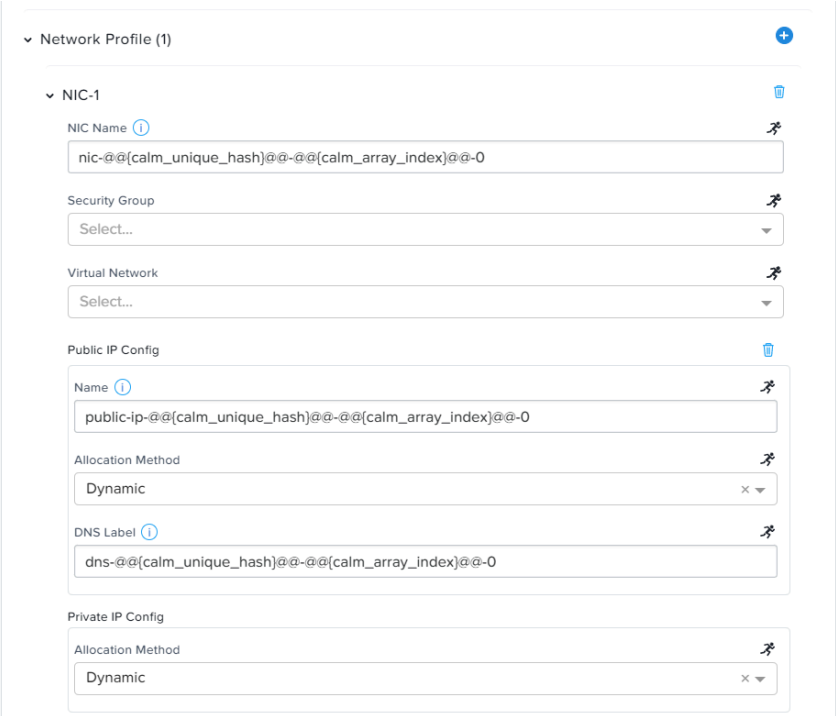

Figure 395: Network Profile

   a. Select a security group from the **Security Group** list.

   b. Select a virtual network from the **Virtual Network** list.

   c. Under **Public IP Config**, enter a name and select an allocation method.

   d. Under **Private IP Config**, select an allocation method.
   If you selected **Static** as the allocation method, then enter the private IP address in the **IP Address** field.

17. (Optional) Enter tags in the **Tags** field.

18. Configure the connection in your environment. For more information, see Configuring Check Log-In for your Environment on page 771.

19. Click **Next**.

20. (Optional) Add credentials for the environment. For more information, see Adding Credentials to the Environment on page 772.

21. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for Azure or launching a blueprint.

Configuring Check Log-In for your Environment

You configure a check log-in task to check whether you are able to SSH into the VM provisioned. Perform the following steps to configure check log-in for your environment.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, under **VM Configuration**, expand the **Connection** section.
2. Click the **Check log-in upon create** check box.
3. In the **Credential** list, select **Add New Credential** to add a new credential and do the following:
   a. Enter a name of the credential in the **Credential Name** field.
   b. Enter user name in the **Username** field.
   c. Select the secret type from the **Secret Type** list.
      You can either select **Password** or **SSH Private Key**.
   d. Do one of the following.
      » If you selected password, enter the password in the **Password** field.
      » If you selected SSH Private Key, enter or upload the SSH private key in the **SSH Private Key** field.

      (Optional) If the private key is password protected, click **+Add Passphrase** to provide the passphrase.
   e. If you want this credential as your default credential, select the **Use as default** check box.
   f. Click **Done**.
4. Select address from the **Address** list.

   You can either select the public IP address or private IP address of a NIC.
5. Select the connection from the **Connection Type** list.
   Select **SSH** for Linux or **Windows (Powershell)** for Windows.
   The **Connection Port** field is automatically populated depending upon the selected **Connection Type**. For SSH, the connection port is 22 and for PowerShell the connection port is 5985 for HTTP and 5986 for HTTPS.
6. If you selected **Windows (Powershell)**, then select the protocol from the **Connection Protocol** list. You can select **HTTP** or **HTTPS**.
7. Enter the delay in seconds in the **Delay** field.

   Delay timer defines the time period when the check login script is run after the VM starts. It allows you to configure the delay time to allow guest customization script, IP, and all other services to come up before running the check login script.
8. In the **Retries** field, enter the number of log-on attempts the system must perform after each log on failure.

Support for Multiple Credential

Credentials help in abstracting identity settings while connecting to an external system. You can configure multiple credentials of the same type (either SSH key or password) under the **Environments** tab. You can use the configured credentials during the launch of an application blueprint.

Adding Credentials to the Environment

Credentials are used to authenticate a user to access various services in Calm. Calm supports key-based and password-based authentication method.

About this task

Use this procedure to add credentials.

Before you begin

Ensure that you have configured a project and created environments for the project. For more information, see Configuring Environments in a Project on page 746.

Procedure

1. On the **Credentials** tab of the Create Environment page, click **+ Add Credentials**.
2. Enter a name of the credential in the **Credential Name** field.
3. Enter a username in the **Username** field.
4. Select a secret type from the **Secret Type** list.
   You can either select **Password** or **SSH Private Key**.
5. Do one of the following.
   - » If you selected password, enter the password in the **Password** field.
   - » If you selected SSH Private Key, enter or upload the SSH private key in the **SSH Private Key** field.

   Optionally, if the private key is password protected, click **+Add Passphrase** to provide the passphrase.
   The type of SSH key supported is RSA. For information on how to generate a private key, see Generating SSH Key on a Linux VM or Generating SSH Key on a Windows VM.
6. Click **Save Environment & Project**.

Quota Policy Overview

Quota policies enforce a usage limit on an infrastructure resource for projects and restrict project members to use more than the specified quota limits. Quotas ensure that a single project or a few projects do not overrun the infrastructures. If the cluster runs out of a resource, project members cannot use the resource even if the project has not reached its specified limit.

Quota policies also enforce a usage limit on an infrastructure resource at the provider account level to ensure that the resource consumption is within the specified quota limits across all projects of that provider account.

> **Note:** Quotas do not reserve any specific amount of infrastructure resources.

Quota Allocation

Quotas are allocated at the account and project levels. Enforcement of resource quota depends on the following factors:

- The status of the policy engine.
  You must enable the policy engine to enforce resource quota policies. For more information, see Enabling policy Engine.
- The resource quotas you allocate to the Nutanix and VMware provider accounts. For more information, see Allocating Resource Quota to an Account.
- The resource quotas you allocate for a project at the project level. For more information, see Managing Quota Limits for Projects on page 774.

- The resource quotas you allocate to the Nutanix and VMware accounts within a project. For more information, see Adding Infrastructure in a Calm-Enabled Prism Central on page 739.

Project-Level Quota Allocation

You can define resource quota limits at the project level for the projects that you create from Prism Central or Calm. The **Policies** tab of a project provides a unified view of all the resource quota limits that you defined for the project and the accounts within the project. You can manage all your project-specific quota definition on the **Policies** tab.

You must consider these conditions while allocating quotas at the project level.

- You can define resource quota limits for a project without defining quota limits for the associated Nutanix or VMware accounts.
- The quota limits you define at the project level cannot be less than the sum of quotas you allocated to different accounts within the project. You can, however, increase the project-level quota limits.
- The project-level quota limits cannot be more than the sum of quotas allocated at the account level to the associated providers. For example, if your project has a Nutanix account and a VMware account, the project-level quota limits cannot be more than the sum of quotas allocated globally to the associated Nutanix and the VMware accounts.
- When you disable quota at the project level, the quota utilization checks are not performed for any actions that you perform within the project from the time it is disabled.

Quota Checks

Quota checks are performed for every resource provisioning request. Quota check happens for multi-VM applications, single-VM applications, and the VMs that are created from Prism Central within a project.

- For multi-VM applications, quota check happens when you launch a blueprint, update an application, or perform a scale-out action to increase the number of replicas of a service deployment.
- For single VM applications, quota check happens when you launch a blueprint or update an application.
- For successful launch of a blueprint, application update, or scale-out, the requested resource must be within the quota limit allocated to the associated provider and project.
- To get the quota consumption of running applications after enabling the policy engine, you can wait for the platform sync to happen or run the platform sync manually. After the first update, all future updates will happen instantly. For more information on how to run platform sync, see Synchronizing Platform Configuration Changes.
- In case of quota violation, appropriate notification is displayed with details such as the associated project, associated account if applicable, and the reasons for violation.
- In case of a scale down or application delete action, the consumed resources are released back as available quotas and added to the total available quota.

Reporting

The Prism Admin and Project Admin can view project-wise usage of infrastructure resources for each cluster. For more information, see Viewing Quota Utilization Report and Managing Quota Limits for Projects on page 774.

Managing Quota Limits for Projects

The **Policies** tab of a project provides you the options to define and manage quota limits for the project and its associated Nutanix and VMware accounts.

About this task

You can do the following as part of the quota limit management at the project level:

- Enable or disable project-level quota checks.
- Assign quota limits for the project.
- Edit quota limits for the associated Nutanix and VMware accounts within the project. For information on allocating resource quota limits to the associated accounts, see Adding Infrastructure in a Calm-Enabled Prism Central on page 739.
- View quota utilization report.

Procedure

1. Do one of the following:
   - » In the Prism Central entities menu, click **Administration** > **Projects**.
   - » In the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.
2. Do one of the following:
   - » Click the **+Create Project** button to create a new project and define quota limits. For more information about creating a project, see Creating a Project on page 729.
   - » Click a project name in the list of existing projects to define quota limits for that project.
3. Configure your project with users, accounts, and environments. For more information, see Projects Overview on page 726.
4. Click the **Policies** tab of the project.

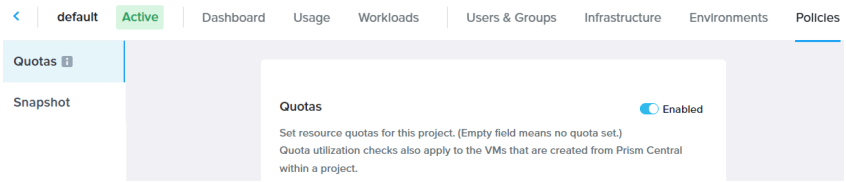

Figure 396: Policies Tab

5. Ensure that the **Quotas** tab is selected in the left pane.
6. To enable quota checks at the project level, enable the **Quotas** toggle button in the right pane.
7. To assign resource quota limits at the project level, enter quota values for **vCPU**, **Memory**, and **Disk** for the project.
   - The quota limits you define at the project level for vCPU, Memory, and Disk must be equal to or more than the sum of quotas you allocated to different accounts within the project.
   - The **Project/Global Quota** row shows the total quota limit defined for the different associated accounts within the project and globally at the account levels. The project-level quota limits must be equal to or less than the sum of the quota allocated at the account level to the associated providers globally.
   - If you hover your mouse over the status bar of a resource in the **Quota Utilization** row, you can view the resources consumed and the resources allocated to the project.

8. To manage the provider quota limits within the project, expand the provider account in the **Provider Quotas** section and do the following:

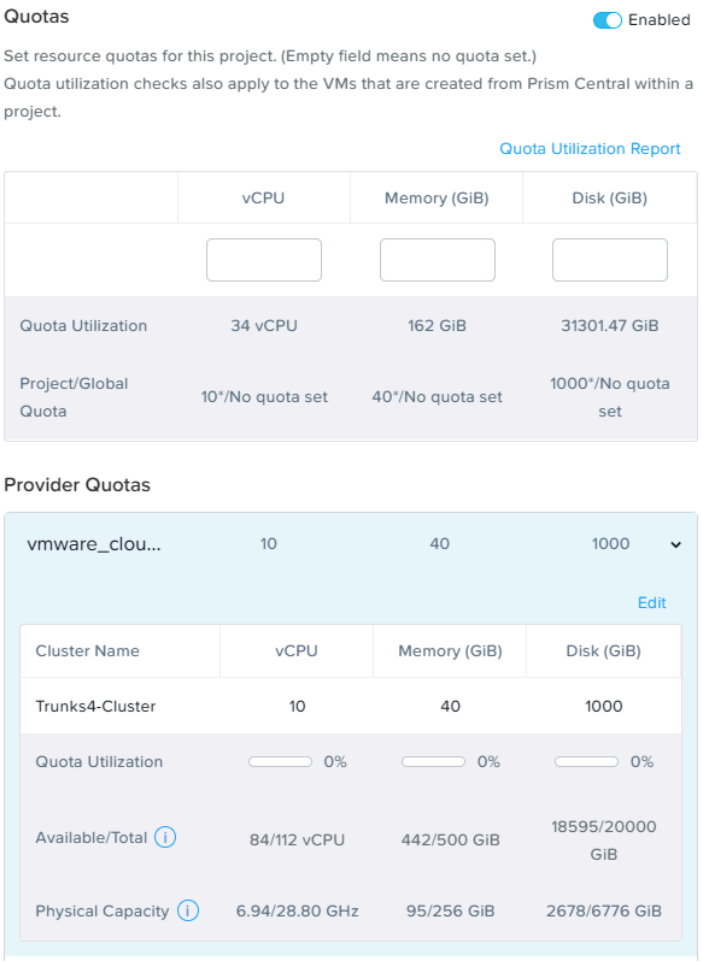

Figure 397: Quota Definition

a. View the quota limits (vCPU, memory, and disk) allocated to the provider account within the project.

b. Click **Edit** to enable and add the quota limit or modify the existing quota limit.

   The **Edit Account** window opens where you can enable quotas for the account, add quota limits, or modify the existing limits.

The **Available/Total** row shows the available resources quota and the total quota allocated to the provider. The **Physical Capacity** row shows the used and total physical capacity. Use these details while allocating resource quotas to the cluster.

9. To view the resource utilization, quota utilization, and application quota utilization at the project level, click **Quota Utilization Report**.

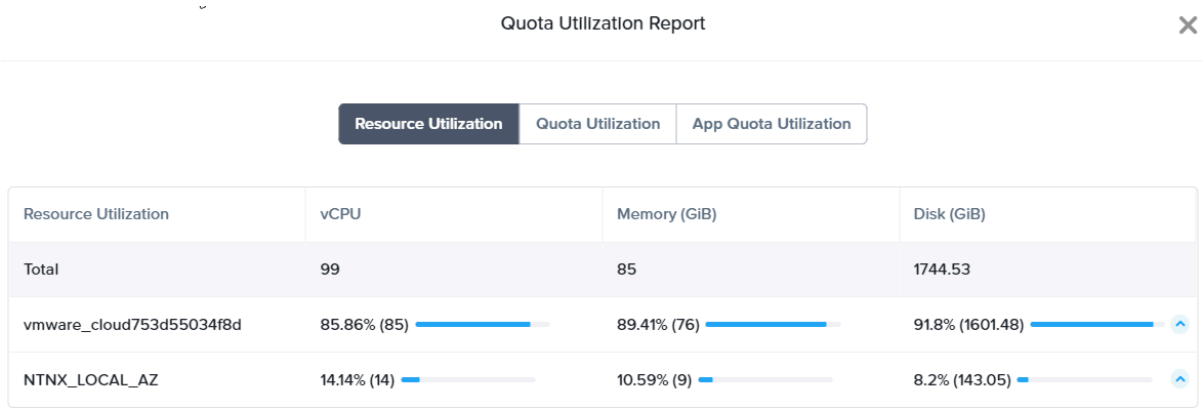

Figure 398: Quota Utilization Report

- Use the **Resource Utilization** tab to view the total utilization of vCPU, Memory, and Disk at the project level. You can also view the utilization of infrastructure resources by each cluster of the accounts associated with the project.
- Use the **Quota Utilization** tab to view detailed utilization data at the level of each cluster of the associated accounts. You can view the quota allocated and quota used by each account. You can also view the quota allocated for each cluster and the percentage of quota utilization at the cluster level.
- Use the **App Quota Utilization** tab to review the resource quota utilization in terms of applications in Calm.

10. To disable quota checks for the actions that are performed within the project, disable the **Quotas** toggle button.

Creating a Snapshot Policy

A snapshot policy allows you to define rules to create and manage snapshots of application VMs that run on a Nutanix platform. The policy determines the overall intent of the snapshot creation process and the snapshot expiration. You can create rules in your snapshot policy to manage your snapshots on a local cluster, on a remote cluster, or both.

About this task

Perform the following steps to create your snapshot policy. For information on the snapshot configuration and creation, see Snapshot and Restore for Nutanix Platform.

Procedure

1. Do one of the following:

   » In the Prism Central entities menu, click **Administration** > **Projects**.

   » In the Prism Central entities menu, click **Services** > **Calm** and then click the **Projects** icon in the left pane.

2. Do one of the following:

   » Click the **+Create Project** button to create a new project and create snapshot policies within the project. For more information about creating a project, see Creating a Project on page 729.

   » Click a project name in the list of existing projects to create snapshot policies within the project.

3. Ensure that you have your project configured with users, accounts, and environments. For more information, see Projects Overview on page 726.

4. On the **Policies** tab, click **Snapshot** in the left pane.

5. Click **+Create Snapshot Policy**.
The Create Snapshot Policy page appears.

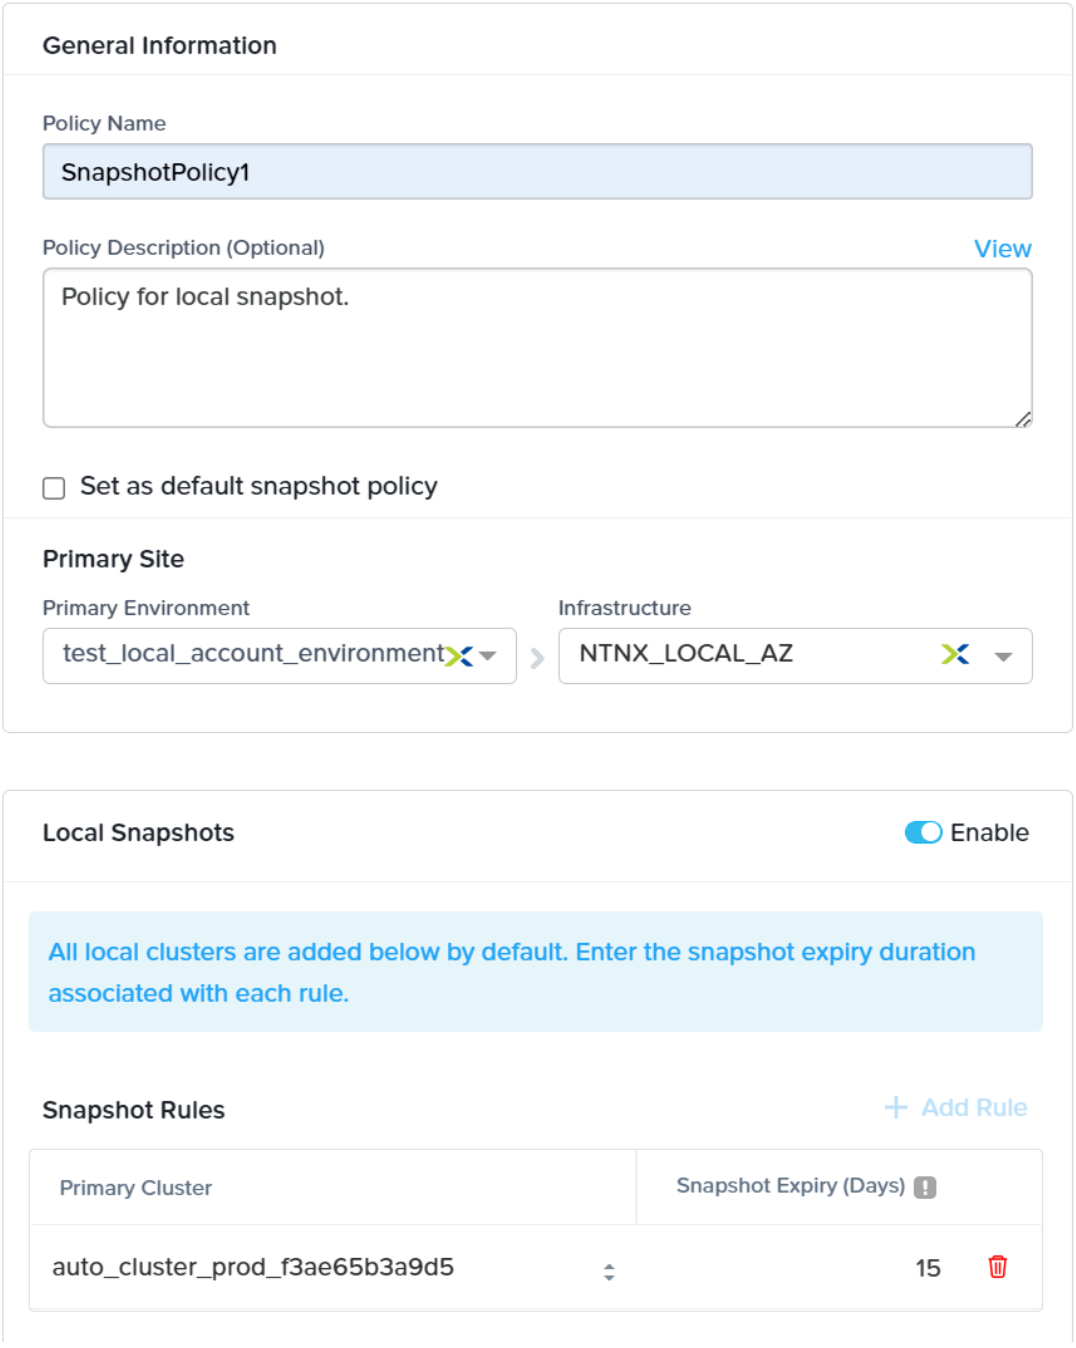

Figure 399: Create Snapshot Policy

6. In the **Policy Name** field, enter a name for the snapshot policy.
7. In the **Policy Description** field, enter a description for the snapshot policy.
8. Select the **Set as default snapshot policy** check box to make this your default snapshot policy.
9. Under Primary Site, select a primary environment.

10. Select an account in the primary environment to which you want to associate the snapshot policy.

    Selecting an account in the **Infrastructure** list enables **Local Snapshots**. You can view all the clusters that you allowed for the account in the **Primary Cluster** column under Snapshot Rules.

11. Under Snapshot Rules, in the **Snapshot Expiry** field for each cluster, specify the number of days after which the snapshot should expire.

    > **Note:** The storage cost of the snapshot depends on the days of expiration you specify for the cluster. The longer the days of expiration, the higher the storage cost. The default value is zero days, which indicates that the snapshot will never expire.

    If you do not want to include an allowed cluster in the policy, click the **Delete** icon next to the cluster. You can use the **+ Add Rule** option to include a deleted cluster to the policy.

12. To enable remote snapshots in your policy, do the following:

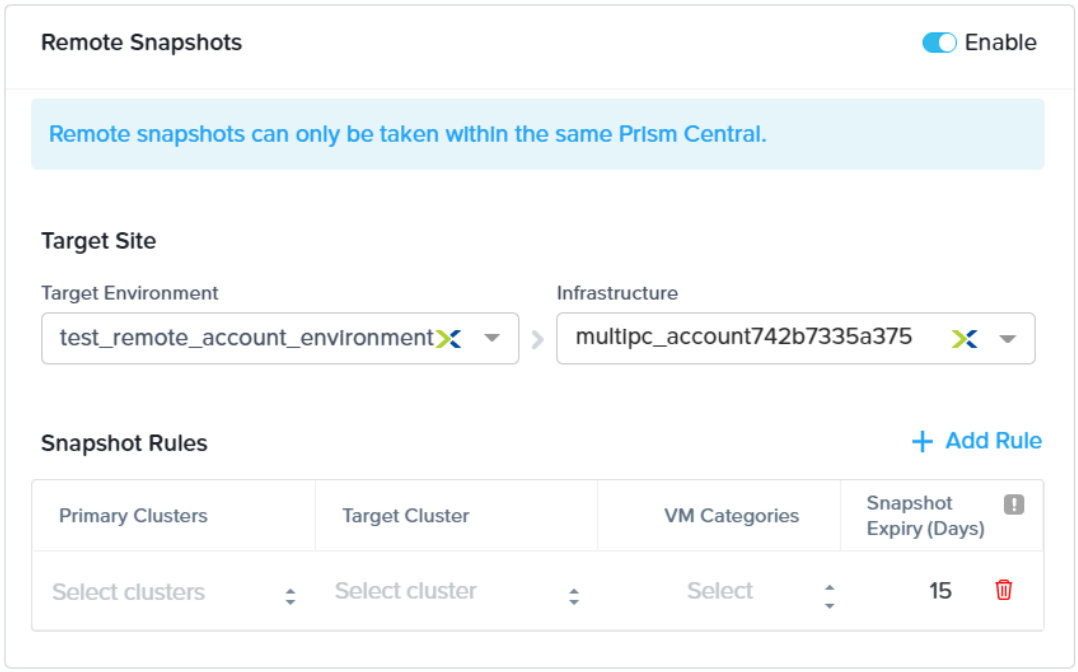

Figure 400: Remote Snapshots

a. Enable the toggle button next to **Remote Snapshots**.

   You can enable remote snapshots if your target environment and the associated account has multiple allowed clusters, and you want to use one of the clusters to store snapshots. Remote snapshots are particularly useful

when your Prism Central has a computer-intensive cluster managing workloads and a storage-intensive cluster managing your data, snapshots, and so on.

You can anytime use the toggle button to enable or disable remote snapshots in your policy.

b. Click **+ Add Rule**.

c. From the **Primary Cluster** list, select the primary cluster where the snapshots of the VMs are taken.

d. From the **Target Cluster** list, select the target cluster where you want to store the snapshots.

e. From the **VM Categories** list, select the VM category.

f. For each cluster, in the **Snapshot Expiry** field for each cluster, specify the number of days after which the snapshot should expire.

g. Click **+ Add Rule** and repeat the steps to add more primary and target clusters to the rule.

13. Click **Save Snapshot Policy**.

NETWORK AND SECURITY MANAGEMENT

This section provides the information on how to set up the network and security infrastructure for a cluster.

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 94 or Subnets on page 170).

The **Network Configuration** page includes three tabs.

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

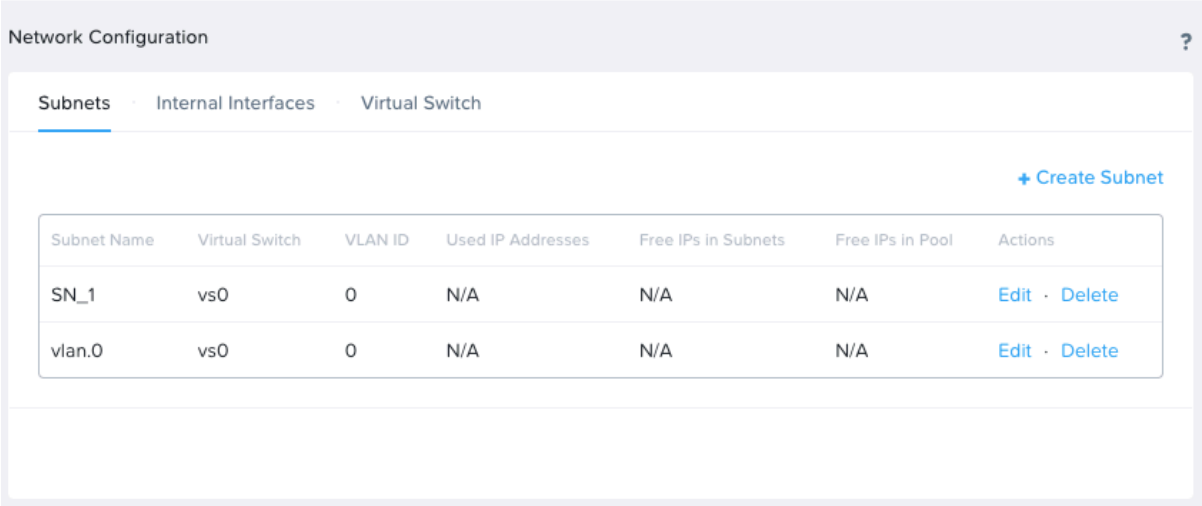

Figure 401: Subnets Tab

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

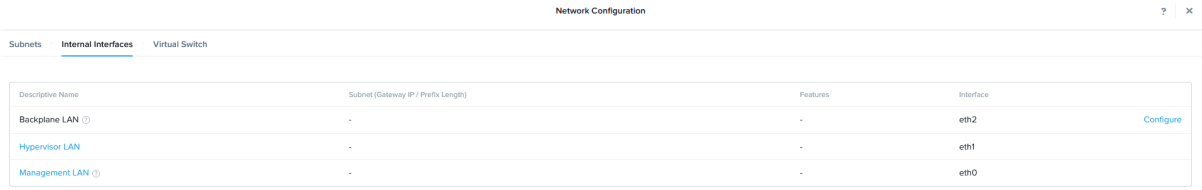

Figure 402: Internal Interfaces Tab

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

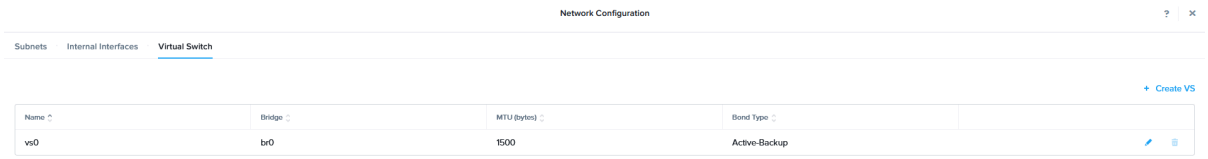

Figure 403: Virtual Switch Tab

The following table describes the fields in each tab.

Table 744: 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	Displays the number of IP addresses in the subnet that are used, for example IP address of a VM or any other entity. This parameter is applicable only when you have configured a managed network or subnet.	(number of IP addresses)
Free IPs in Subnets	Displays the number of free or unused IP addresses in the subnet. This parameter is applicable only when you have configured a managed network or subnet.	(number of IP addresses)
Free IPs in Pool	Displays the number of free or unused IP addresses in the configured pool. This parameter is applicable only when you have configured a managed network or subnet.	(number of IP addresses)
Actions	Action link for editing or deleting a network configuration.	(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)

Parameter	Description	Values
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 LACP and LAG (AHV Only) on page 791.

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 94) and click the **Network Config** button.
   - » Go to the Network and Security dashboard (see Subnets on page 170 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.

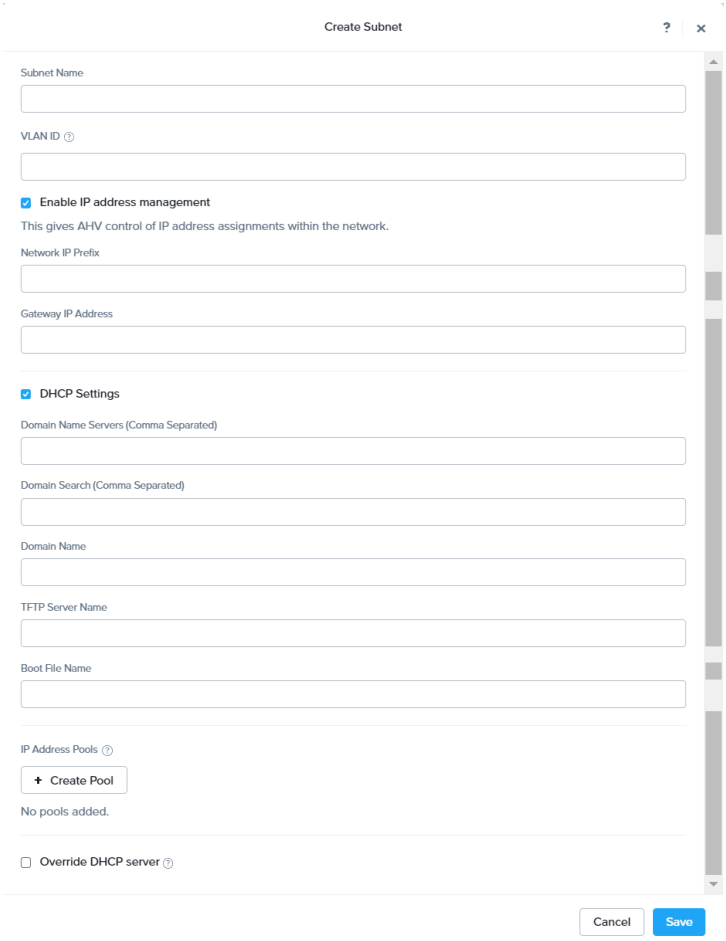

Figure 404: Create Subnet Page (expanded)

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 787.

      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.

   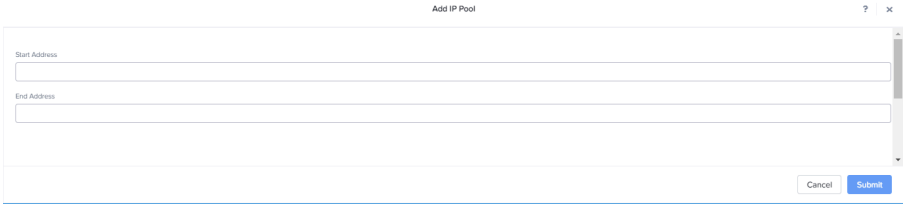

   Figure 405: Add IP Pool Page

   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.

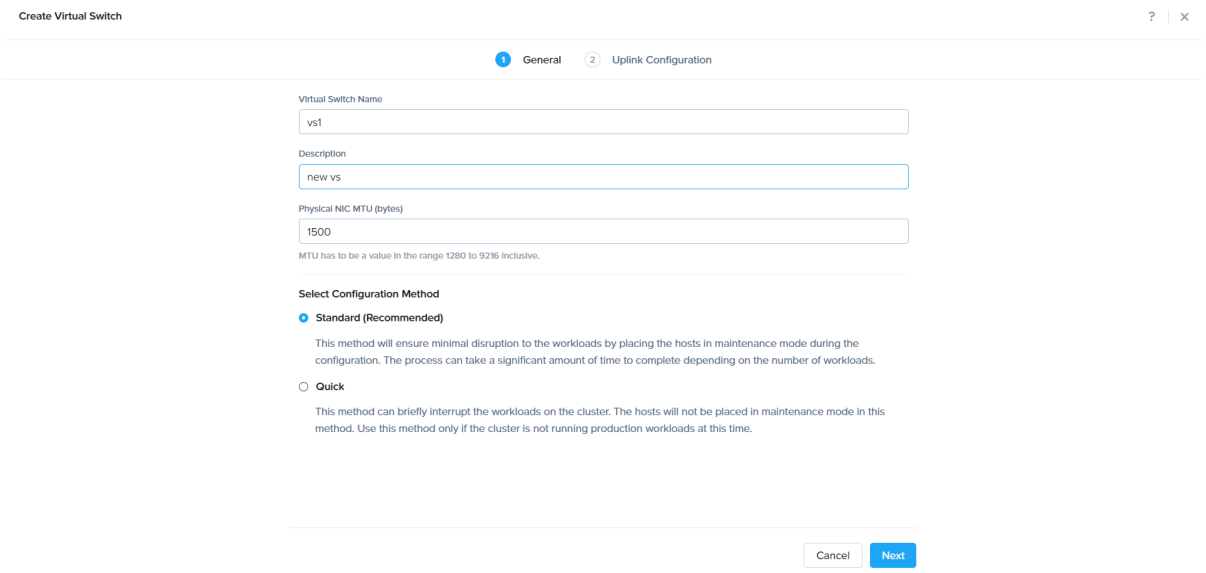

Figure 406: Create Virtual Switch - General tab

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.
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:

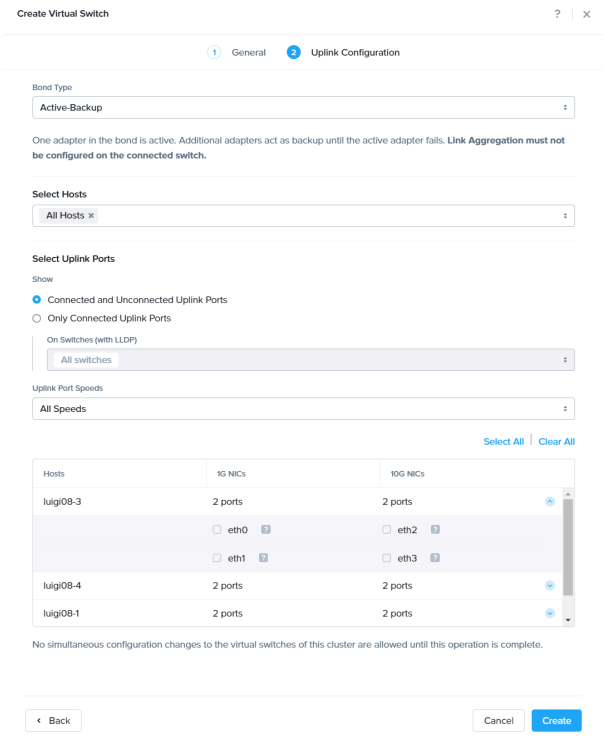

Figure 407: Create Virtual Switch - Uplink Configuration tab

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.

Field	Description and Value
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 745: 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
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

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
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 Virtual Switch Workflow 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 Element 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 LACP and LAG (AHV Only) on page 791

Enabling LACP and LAG (AHV Only)

If you select the **Active-Active** bond type, you must enable LACP and LAG on the corresponding ToR switch for each node in the cluster one after the other. This section describes the procedure to enable LAG and LACP in AHV nodes and the connected ToR switch.

About this task

Procedure

1. Change the uplink **Bond Type** for the virtual switch.

   a. Open the **Edit Virtual Switch** window.

      » In Prism Central, open **Network & Security** > **Subnets** > **Network Configuration** > **Virtual Switch**.

      » In Prism Element or Web Console, open **Settings** > **Network Configuration** > **Virtual Switch**

   b. Click the *Edit* icon of the virtual switch you want to configure LAG and LACP.

   c. On the **Edit Virtual Switch** page, in the **General** tab, ensure that the **Standard** option is selected for the **Select Configuration Method** parameter. Click **Next**.

      The **Standard** configuration method puts each node in maintenance mode before applying the updated settings. After applying the updated settings, the node exits from maintenance mode. See Virtual Switch Workflow.

   d. On the **Uplink Configuration** tab, in **Bond Type**, select **Active-Active**.

   e. Click **Save**.

   The **Active-Active** bond type configures all AHV hosts with the *fast* setting for LACP speed, causing the AHV host to request LACP control packets at the rate of one per second from the physical switch. In addition, the **Active-Active** bond type configuration sets LACP fallback to **Active-Backup** on all AHV hosts. You cannot modify these default settings after you have configured them in Prism, even by using the CLI.

   This completes the LAG and LACP configuration on the cluster.

Perform the following steps on each node, one at a time.

2. Put the node and the Controller VM into maintenance mode.

   Before you put a node in maintenance mode, see Verifying the Cluster Health and carry out the necessary checks.

   See Putting a Node into Maintenance Mode. Step 6 in this procedure puts the Controller VM in maintenance mode.

3. Change the settings for the interface on the ToR switch that the node connects to, to match the LACP and LAG setting made on the cluster in step 1 above.

   This is an important step. See the documentation provided by the ToR switch vendor for more information about changing the LACP settings of the switch interface that the node is physically connected to.

   - Nutanix recommends that you enable LACP fallback.
   - Consider the LACP time options (*slow* and *fast*). If the switch has a *fast* configuration, set the LACP time to *fast*. This is to prevent an outage due to a mismatch on LACP speeds of the cluster and the ToR switch. Keep in mind that the **Active-Active** bond type configuration set the LACP of cluster to *fast*.

   Verify that LACP negotiation status is negotiated.

4. Remove the node and Controller VM from maintenance mode.

   See Exiting a Node from the Maintenance Mode (AHV). The Controller VM exits maintenance mode during the same process.

What to do next

Do the following after completing the procedure to enable LAG and LACP in all the AHV nodes the connected ToR switches:

- Verify that the status of all services on all the CVMs are Up. Run the following command and check if the status of the services is displayed as **Up** in the output:

```
nutanix@cvm$ cluster status
```

- Log on to the Prism Element of the node and check the **Data Resiliency Status** widget displays **OK**.

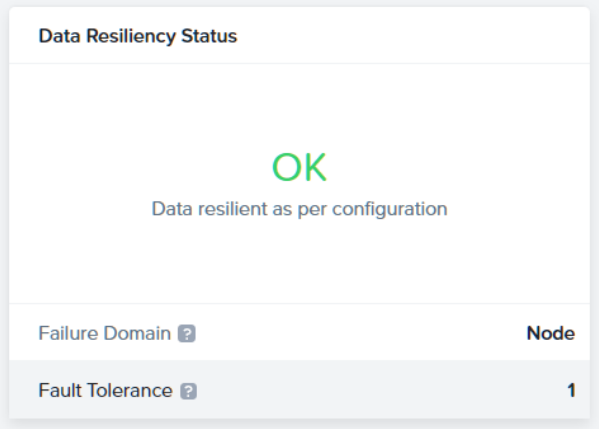

Figure 408: Data Resiliency Status

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 94) and click the **Network Config** button.

   » Go to the Networks dashboard (see Subnets on page 170 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 784).

   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.

Security Policies

For information on Security Policies, see Flow Microsegmentation Guide.

Flow Networking

For information on how to configure Virtual Private Cloud, see Virtual Private Cloud topic in *Flow Virtual Networking Guide*.

For information on how to set up the Network Gateway, VPN connection, and Subnet extension, see Connections Management topic in *Flow Virtual Networking Guide*.

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.

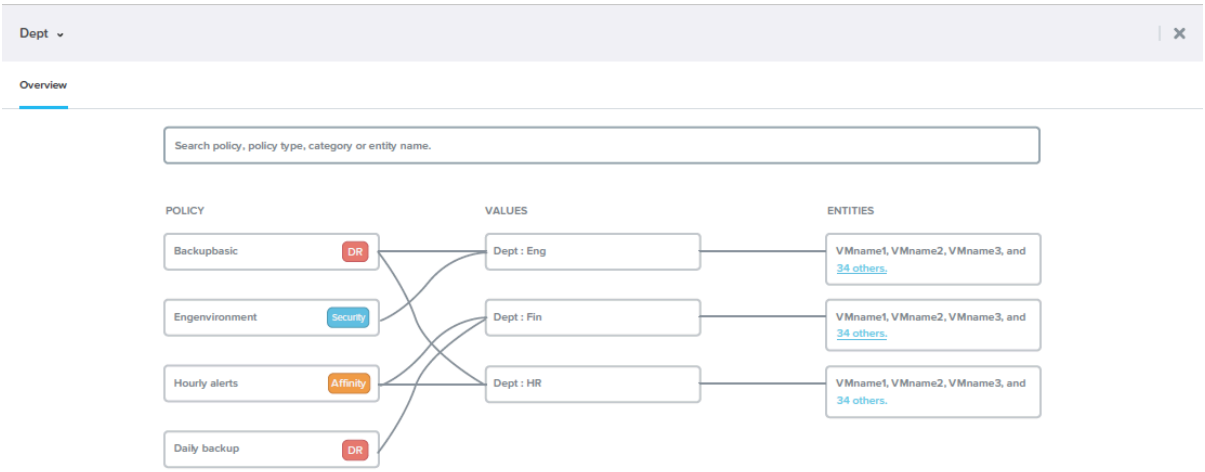

Figure 409: Dept Details View (example)

Creating a Category

About this task

To create a category, do the following:

Figure 410: 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 164) and click the **Create Category** button.
   The **Create Category** page appears.

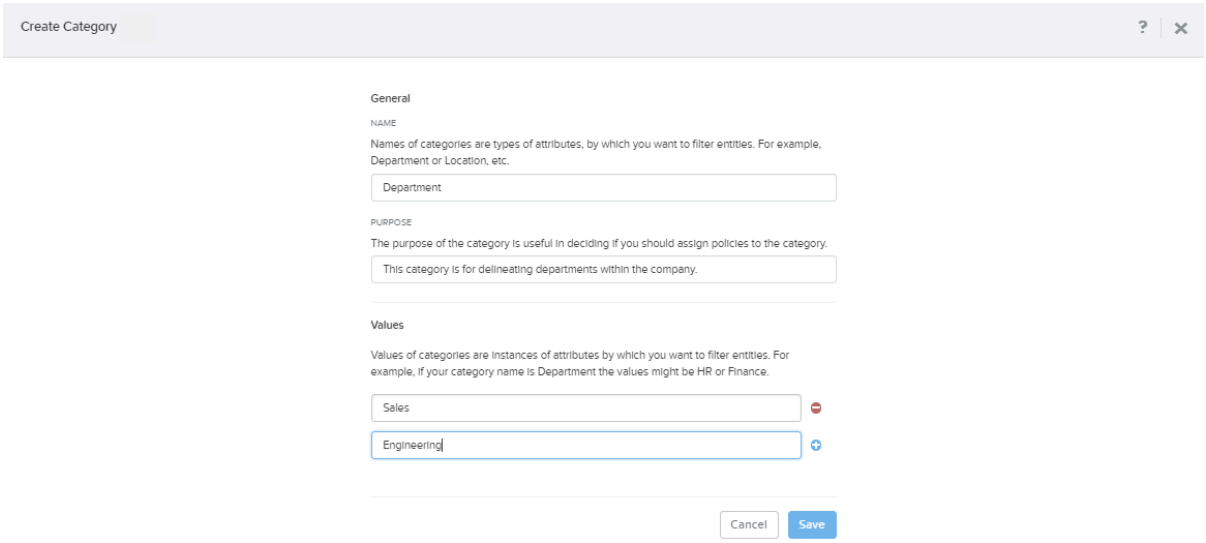

Figure 411: Create Category Page

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 164) 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 795). 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, host, volume group, catalog, image, report, subnet, blue print, app, and marketplace item. 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 202) 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 94) and select the target VMs.
   - Host: Go to the **List** tab in the hosts dashboard (see Hosts Summary View on page 221) and select the target hosts.
   - Volume group: Go to the **List** tab in the volume groups dashboard (see Volume Groups Summary View on page 147) and select the target volume groups.
   - Catalog: Go to the catalogue items dashboard (see Catalog Items View on page 129) and select the target catalog items.
   - Image: Go to the **List** tab in the Images dashboard (see Images Summary View on page 125) and select the target images.
   - Report: Go to the reports dashboard (see Reports View on page 246) and select the target reports.
   - Subnet: Go to the Subnets dashboard (see Subnet Summary View on page 173) and select the target subnets.
   - Blue Print, App, or Marketplace Item: See the Marketplace Manager section of the *Calm Administration and Operations Guide*.

2. Select **Manage Categories** from the **Actions** pull-down menu.
   This displays the **Manage [Cluster|VM|Image|Subnet] Categories** page.

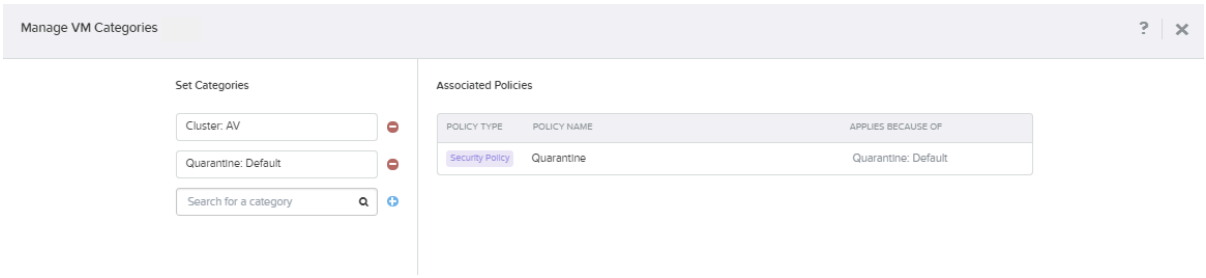

Figure 412: Manage VM Categories Page

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.

OPERATIONS MANAGEMENT

- Analysis: Analyze cluster activity (see Analysis Dashboard (Prism Central) on page 799)
- App Discovery: Discover applications running in a cluster (see Application Discovery on page 831)
- Cost Management: Analyze and manage costs (see Cost Management (Xi Beam) on page 855)
- Integrations: Monitor select applications in a cluster (see Application Monitoring on page 856)
- Ops Policies: Manage operations policies (see Operations Policy Management on page 873)
- Planning: Plan for resource capacity and usage (see Resource Planning on page 876)
- Playbooks: Create playbooks for task automation (see Task Automation on page 902)
- Reports: Generate resource and activity reports (see Reports View on page 246)

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 (Prism Central) on page 799).
- You can create custom entity and metric performance charts (see Chart Management on page 809).

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

Analysis Dashboard (Prism Central)

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 824.

> **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.

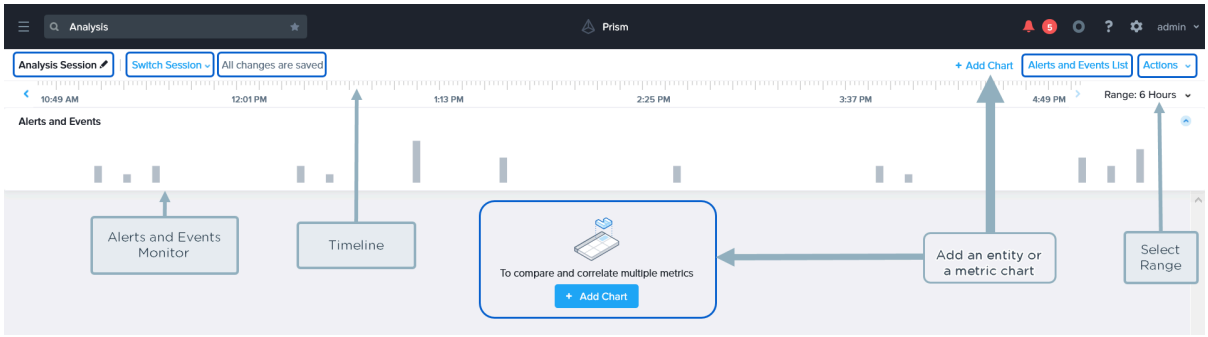

Figure 413: Analysis Dashboard - Prism-generated Session

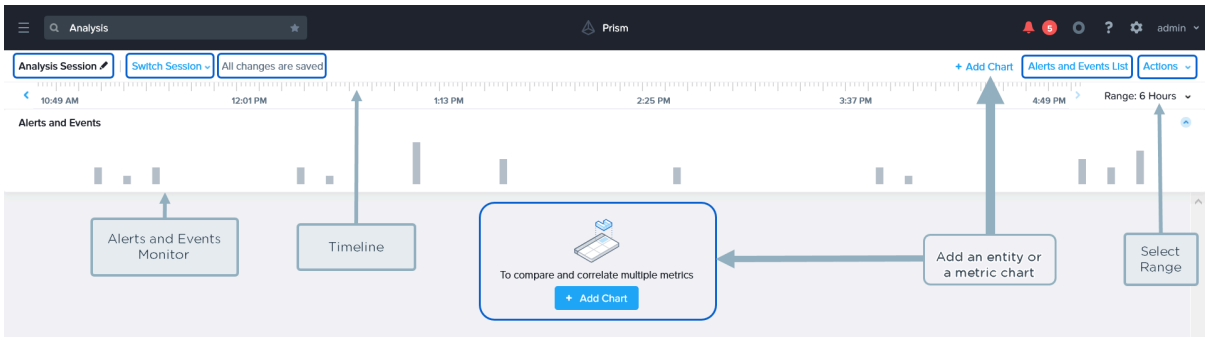

Figure 414: Analysis Dashboard - User-defined Session

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 746: 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.

Field	Description
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.
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.

Field	Description
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.
More	Provides the following options for the current session: 1. **Edit Session Details** 2. **Share** 3. **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. 4. **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 747: 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**
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**.

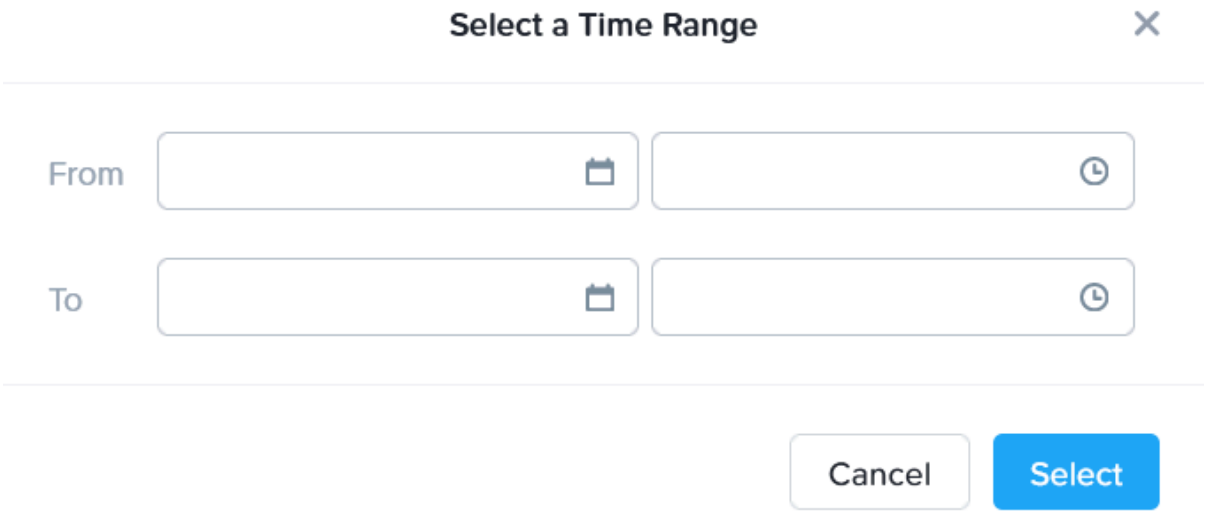

Figure 415:

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.

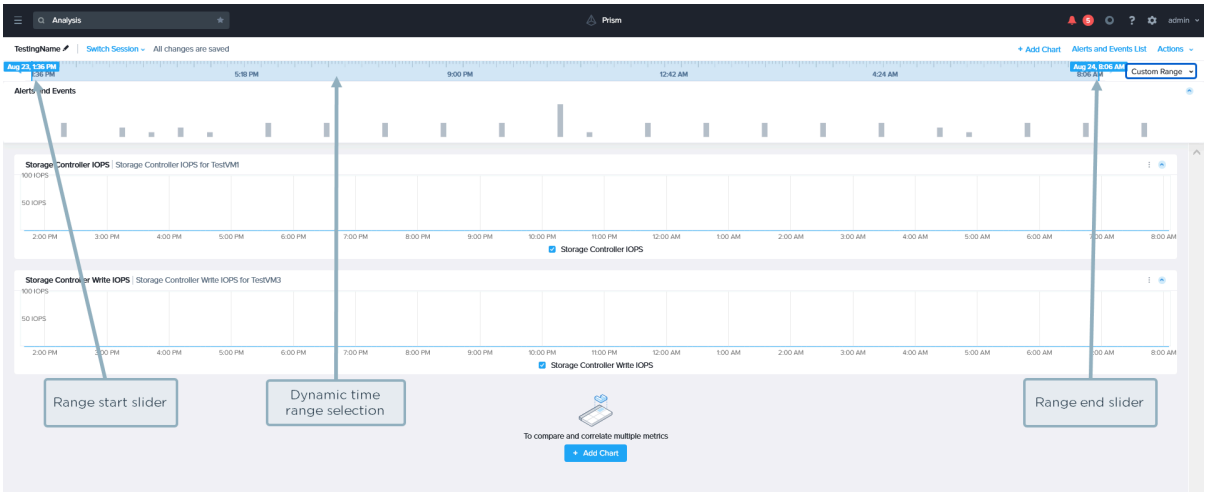

Figure 416: Time Range

Viewing Alerts and Events

About this task

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

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

Procedure

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

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

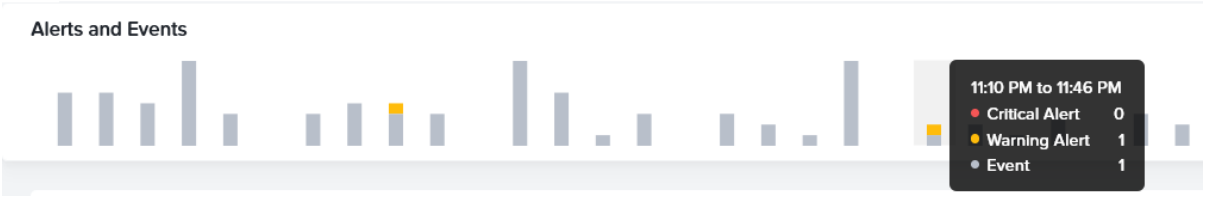

Figure 417: Stacked bar - Alerts and Events Information

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 747: Alerts and Events List Panel Elements on page 803 table for more information.

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

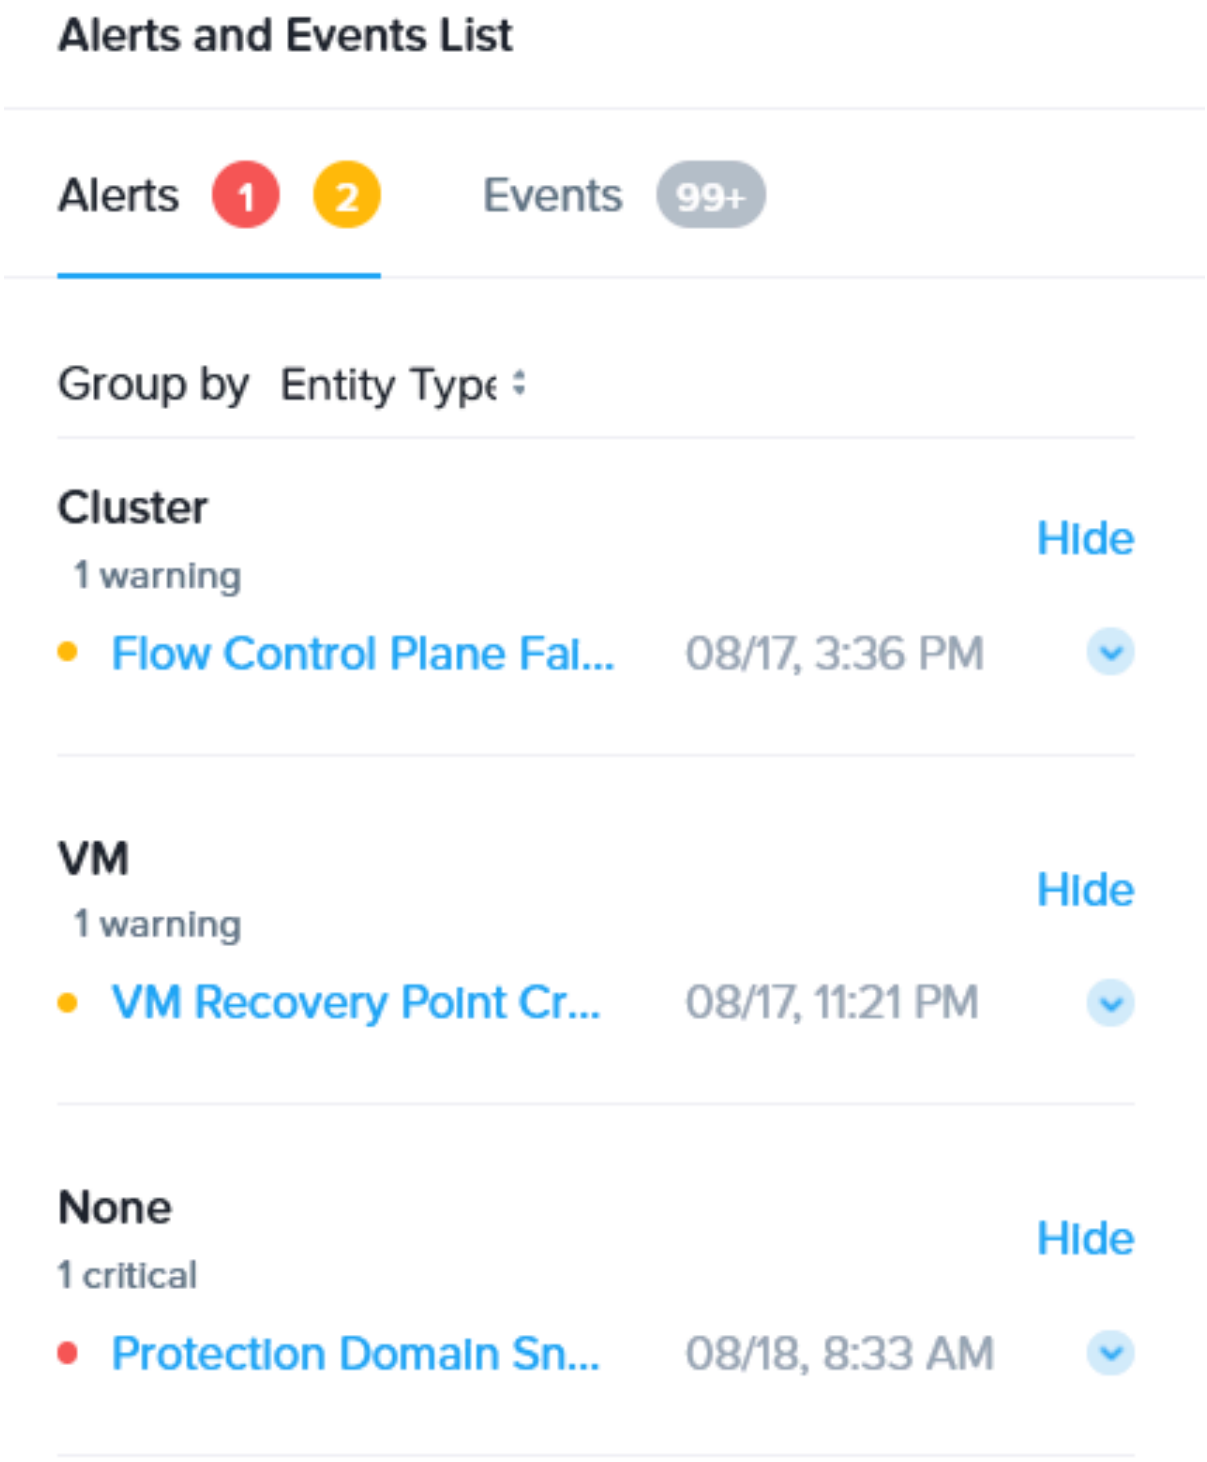

Figure 418: Alerts tab

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.

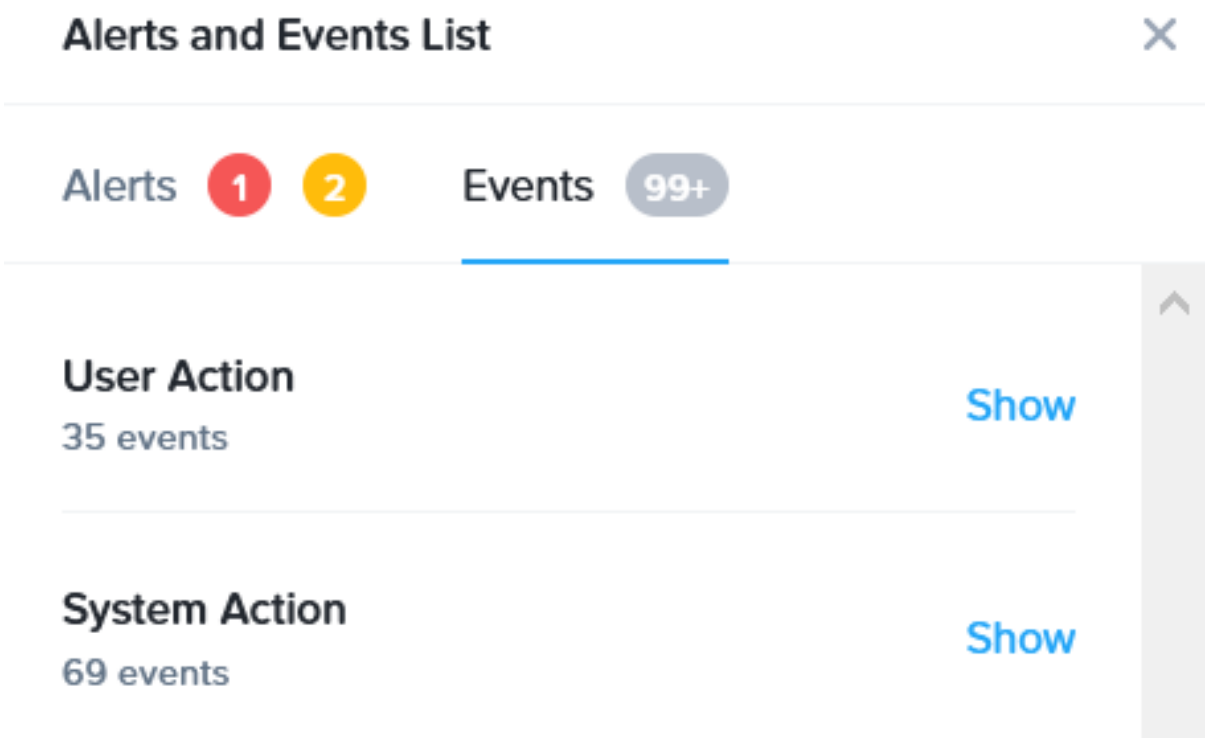

Figure 419: Events tab

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 810
Edit a chart	Editing a Chart on page 812
Delete a chart	Deleting a Chart on page 813
Clone a chart	Cloning a Chart on page 814
Export a chart as a JSON or CSV file	Exporting a Chart on page 814

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.

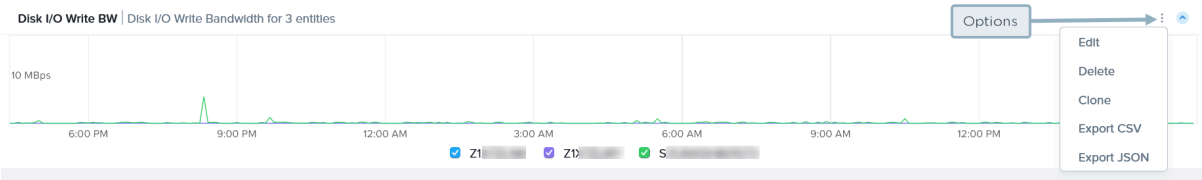

Figure 420: Sample Chart - Chart Options

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.

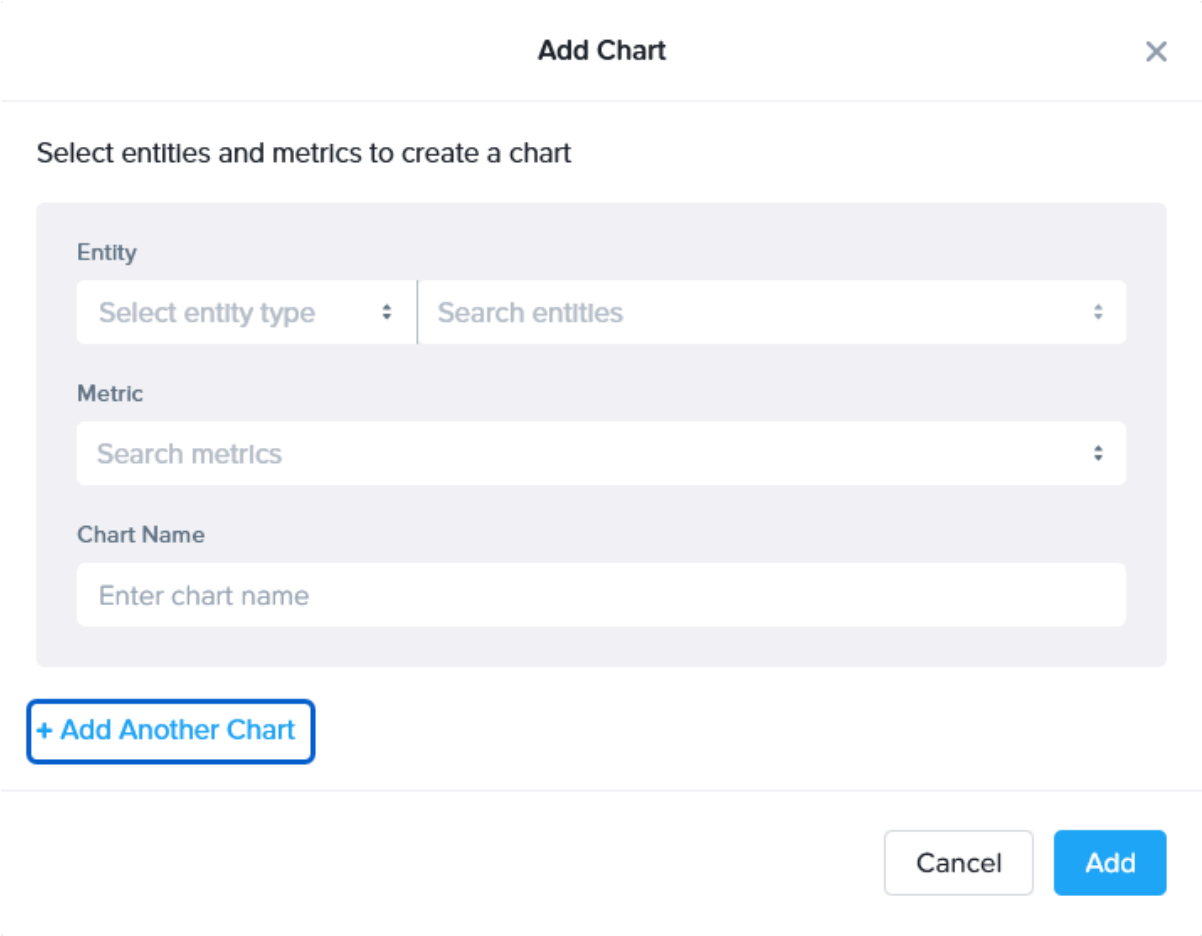

Figure 421: Add Chart dialog box

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 815 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 422: 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**.

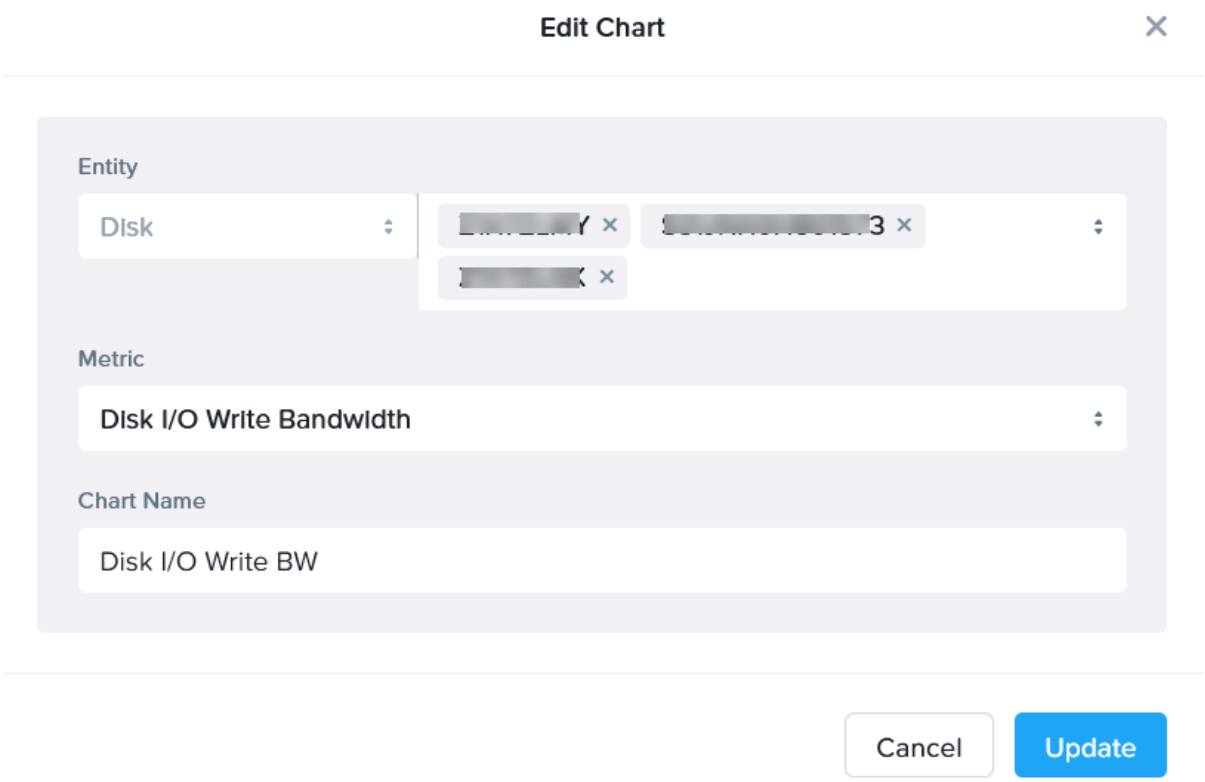

Figure 423: Edit Chart dialog box

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**.

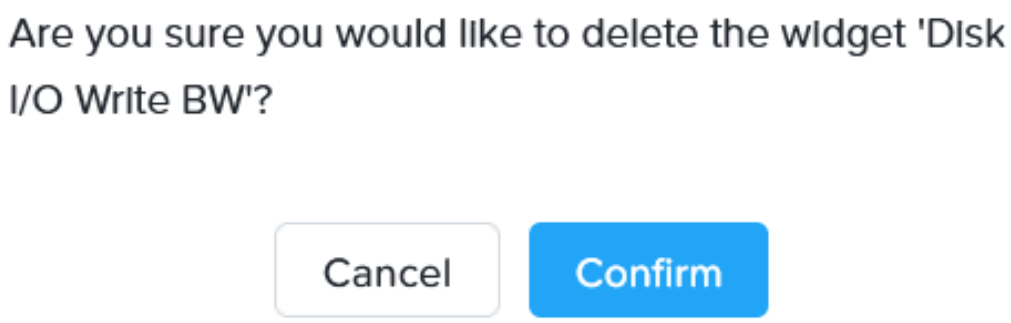

Figure 424: Delete confirmation dialog box

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`
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`

Metric	Entity Type(s)	Description
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`
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`

Metric	Entity Type(s)	Description
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`
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

Metric	Entity Type(s)	Description
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.
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	Network transmitted bytes reported by the hypervisor. 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.
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.

Metric	Entity Type(s)	Description
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`
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.
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.

Metric	Entity Type(s)	Description
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 **Sessions**dashboard 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 **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. 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.

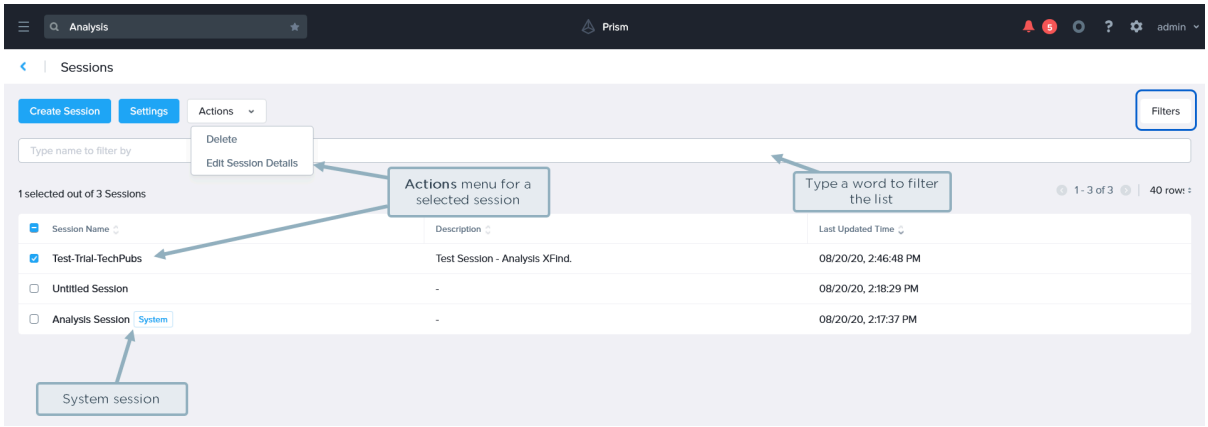

Figure 425: Sessions dashboard

Table 748: 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. • **Share** 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.
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 828.

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.

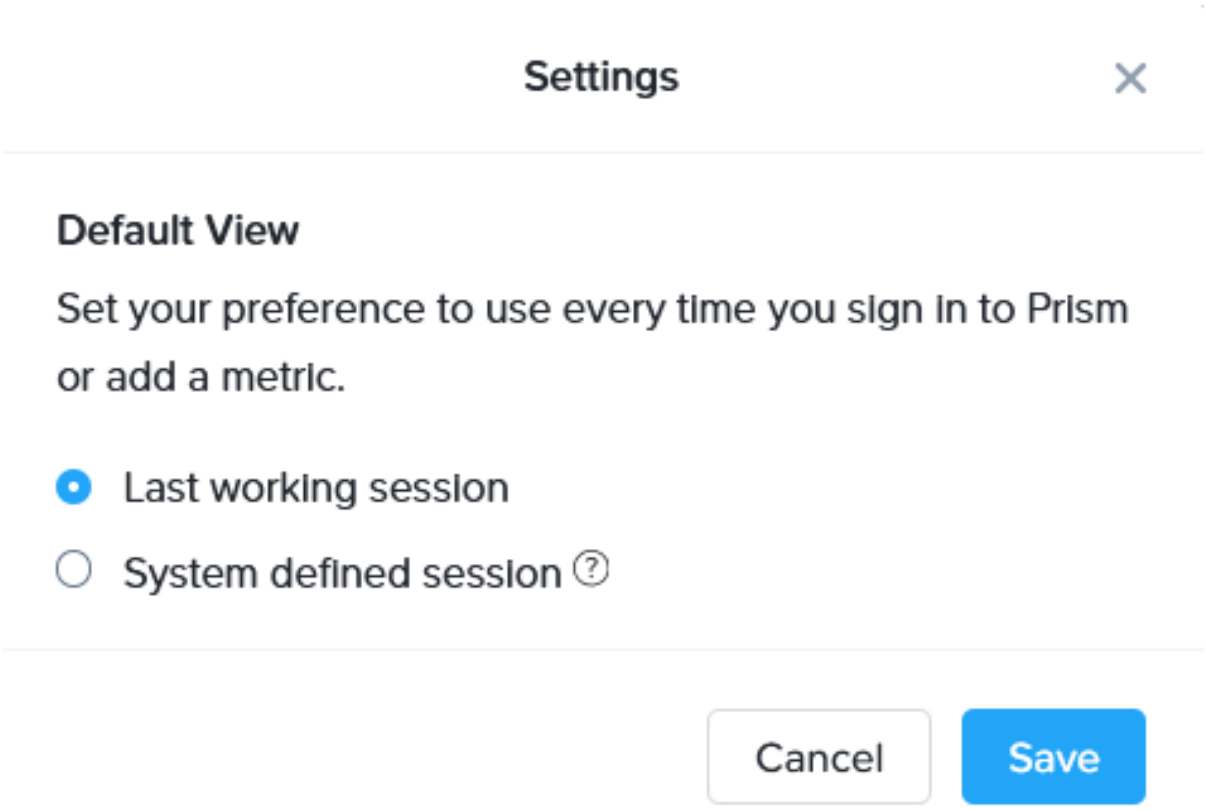

Figure 426: Settings dialog box

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**.

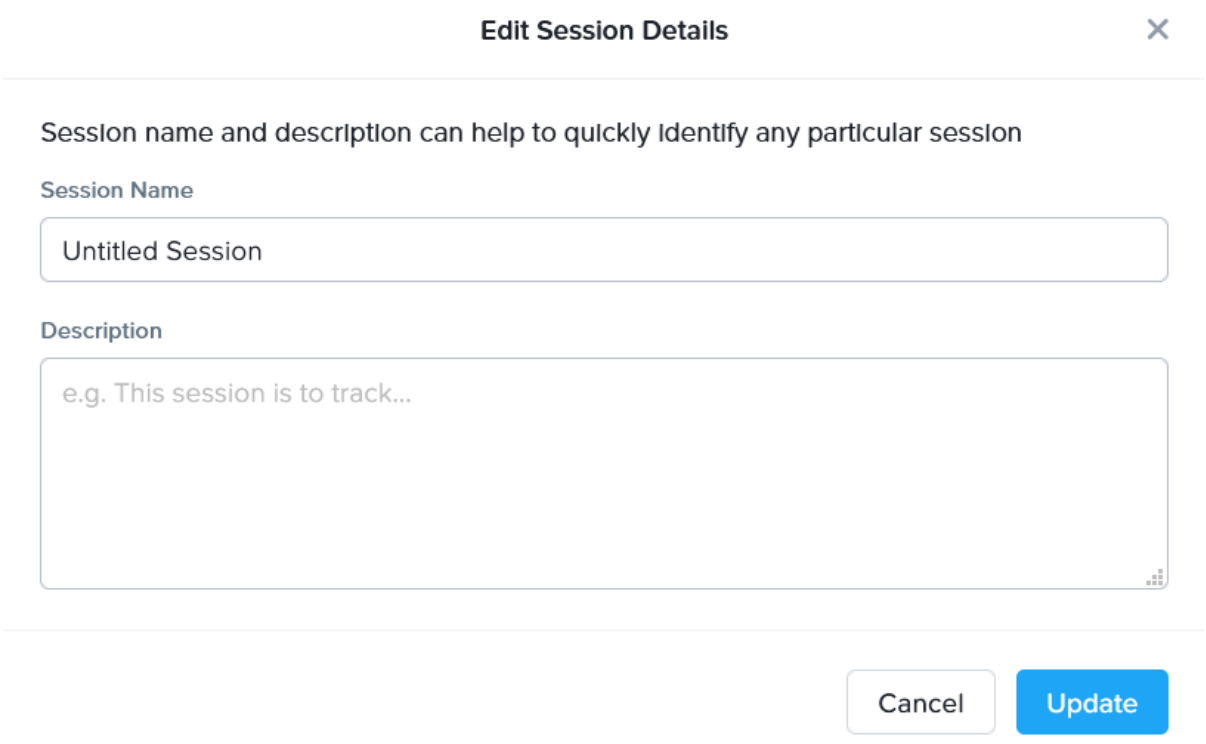

Figure 427: Edit Session Details dialog box

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

Sharing a Session

About this task

Any user with edit permission can share a session with other users and provide view or edit access to them.

> **Note:** Enable CMSP to use this feature.

Procedure

1. In the **Analysis** dashboard, click **Switch Session > View All Sessions**.
   This switches the view to the **Sessions** dashboard.

2. Select a session, go to **Actions** and click **Share**.

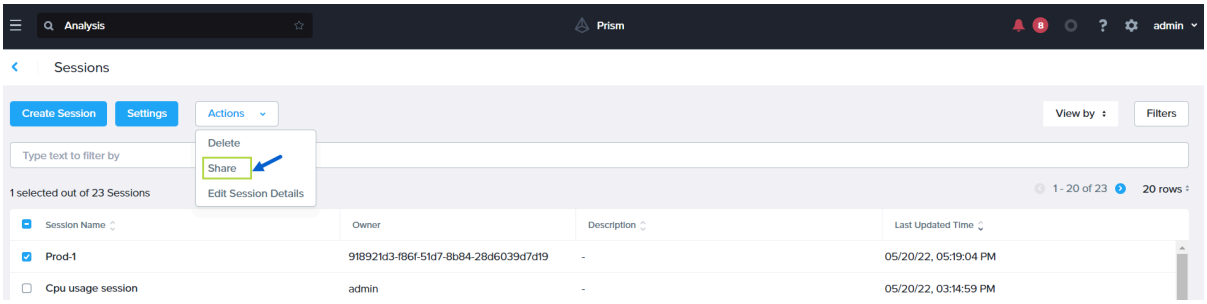

Figure 428: Sessions Dashboard

The **Sharing Session** window appears.

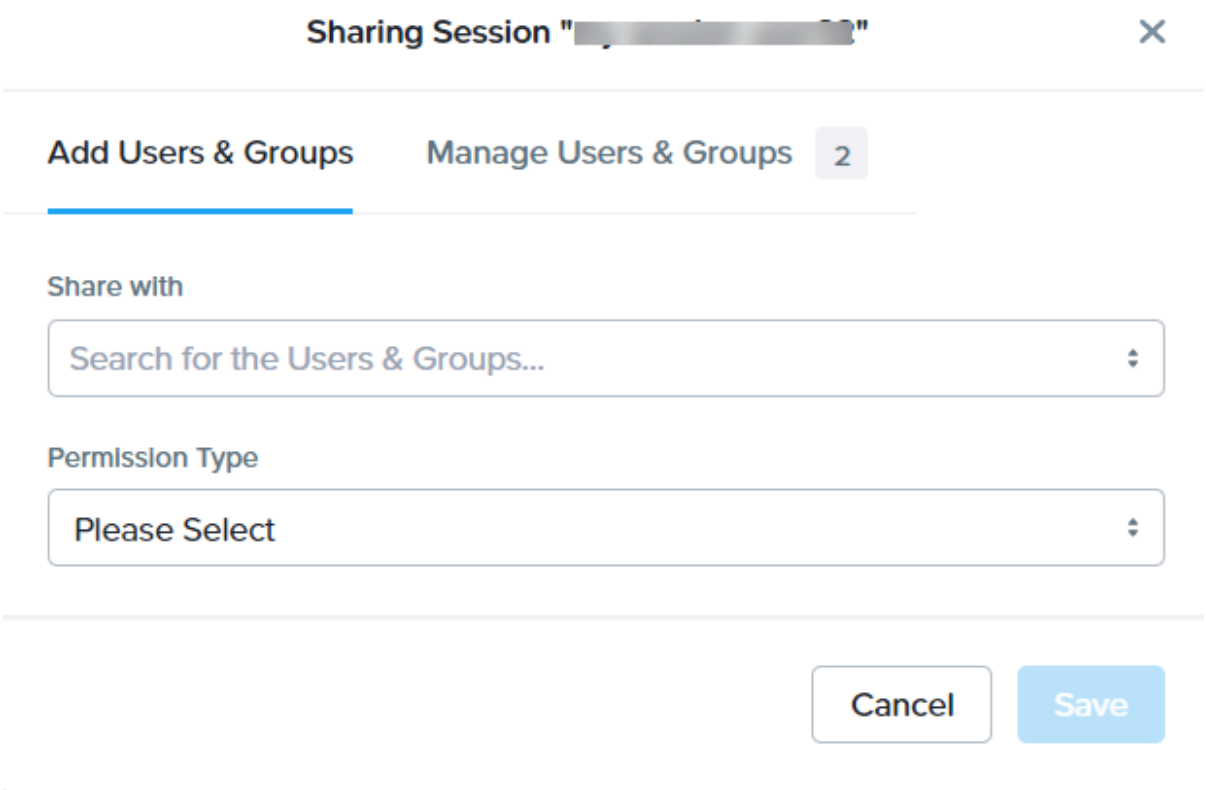

Figure 429: Sharing Session

3. In the **Add Users & Groups** tab, do the following:

   a. **Share with**. Add the users or groups to share the session. You can type first few characters of a user or group, as you type a pull-down list appears. Click the username or group name to add. You can add more than one user or a group.

   b. **Permission Type**. You can provide View only or Edit permissions to the user or group. User with Edit permission can perform these tasks: View, Edit or Share the session.

   The **Manage Users & Groups** tab allows you to change the permission type for users or groups with whom you have shared the session earlier. You can also remove the users or groups.

4. Click **Save**.
   A message that the session is shared successfully appears.

   After sharing the session, wait for few minutes for changes to take effect.

   Users will not be able to view any information about entities in the analysis chart without entity permissions.

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 841)
  - connect to the cloud-based discovery service (see Connecting to the Cloud on page 847)
  - add clusters to monitor (see Adding (Removing) Clusters to Monitor on page 848)
  - authenticate with vCenter for ESXi cluster access (see Configuring vCenter Authentication on page 850)

- To generate current findings at any time:
  - start an application discovery (see Discovering Applications on page 851)
  - publish data for identified applications (see Publishing Applications on page 852)
  - create policies for unidentified applications (see Creating Discovery Policies on page 852)
- To view discovery findings, see the application discovery dashboard (see Application Discovery View on page 832).

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. For instructions on how to view and organize that information in a variety of ways, see Entity Exploring on page 89 .
- If application discovery is not enabled, an enablement page appears instead of the dashboard. See Enabling Application Discovery on page 841 .

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 850) or click **Manage** to see the cloud connection details (see Connecting to the Cloud on page 847). 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 848 .
- **Top Discovered Apps**: Lists the applications with the most instances.
- **Top Discovered Ports**: Lists the ports with the most instances.

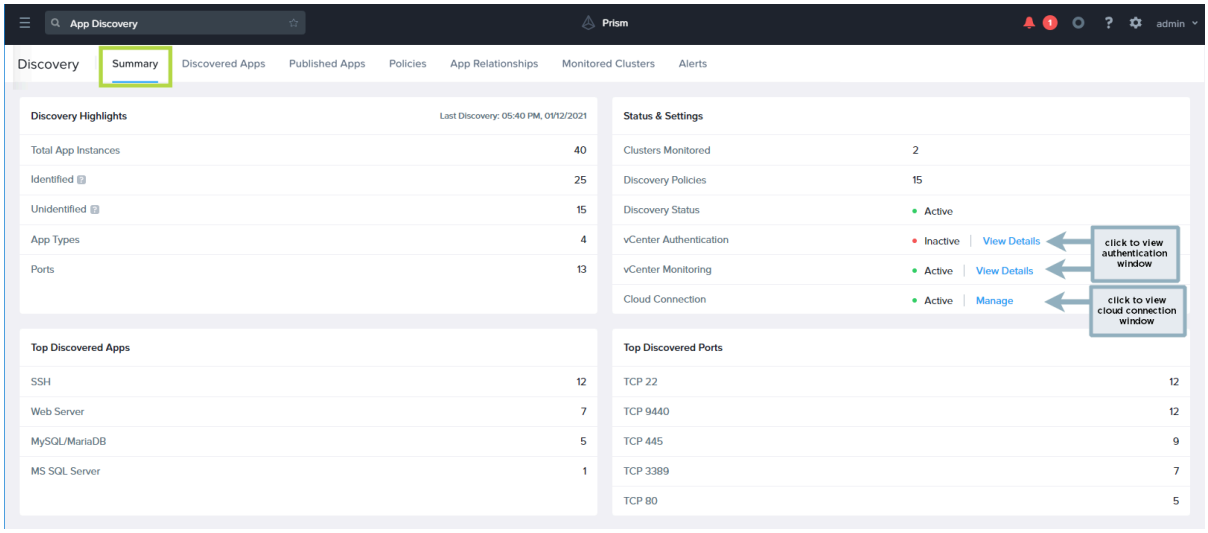

Figure 430: App Discovery Summary Tab

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.

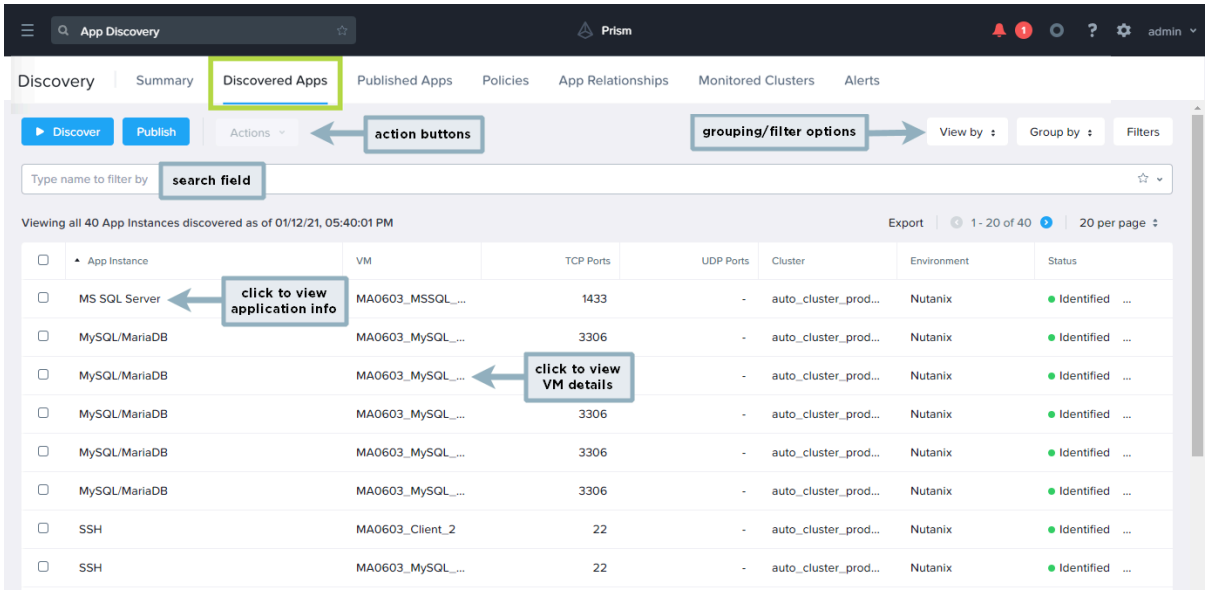

Figure 431: Discovered Apps Tab

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 749: 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 105 .	(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 209). See Configuring vCenter Authentication on page 850 .	(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 852 .	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 750: 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
Ignored	Filters for ignored applications.	Show Ignored Apps (check box)

Parameter	Description	Values
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 851 .
- Click the **Publish** button to publish application information. See Publishing Applications on page 852 .
- 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.

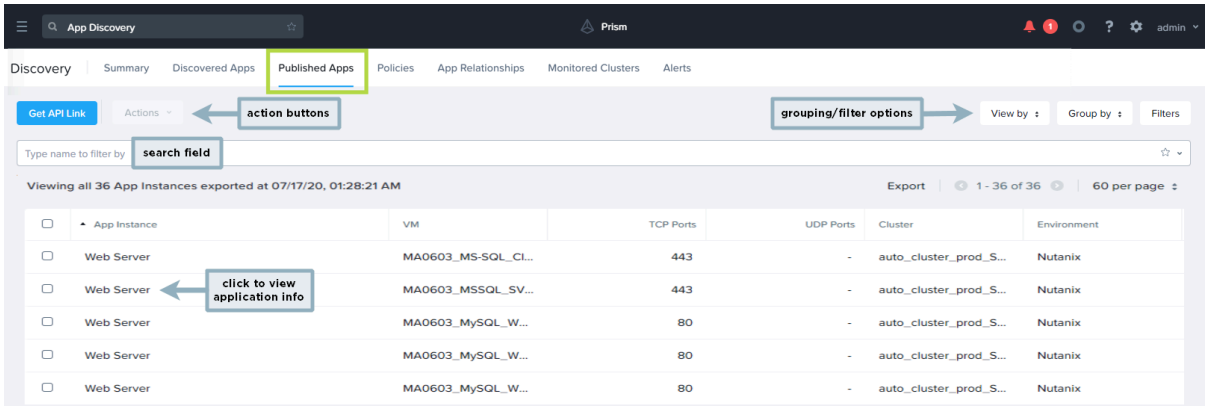

Figure 432: Published Apps Tab

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.

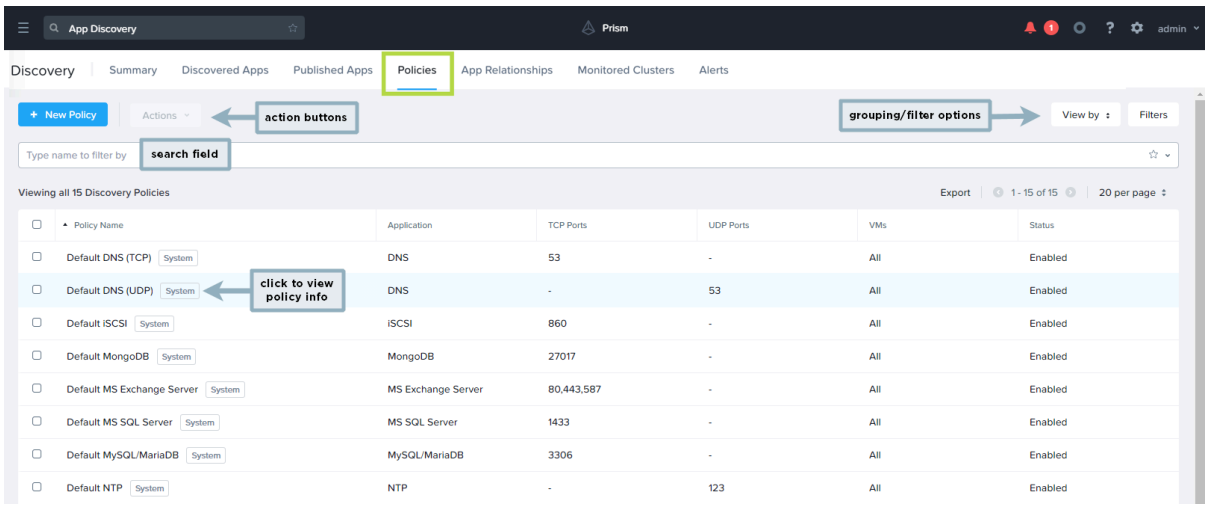

Figure 433: Policies Tab

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 751: 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 854 .	(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 852 .	(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 852 .
- 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.

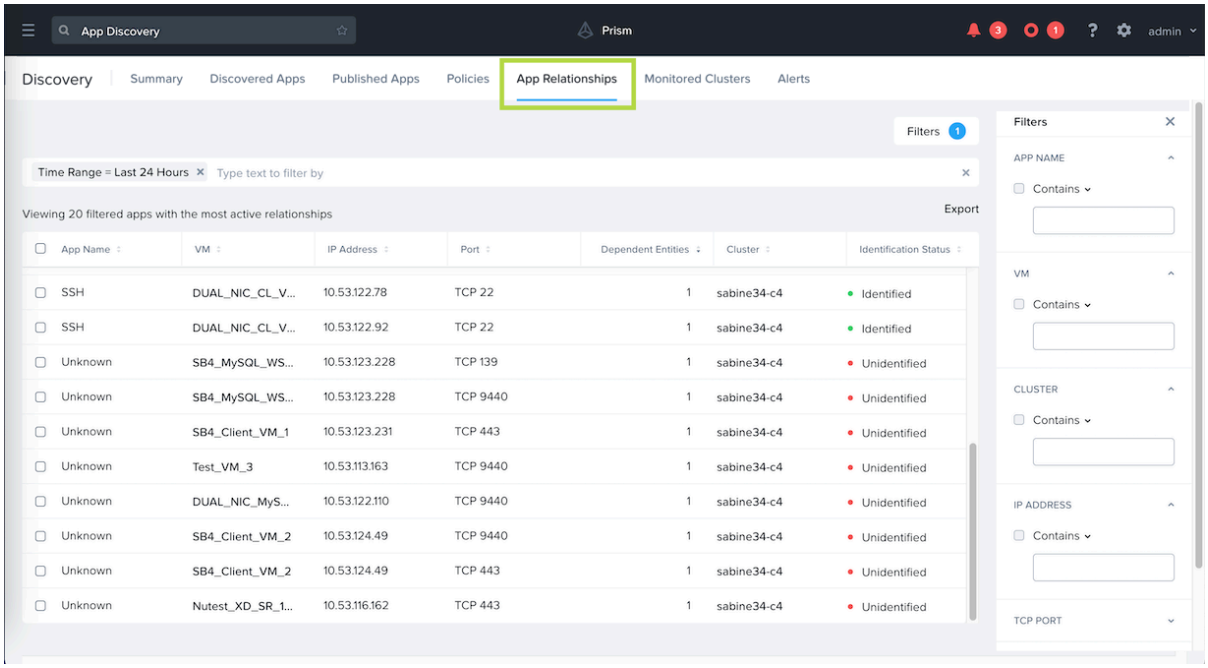

Figure 434: App Relationships Tab

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

Table 752: 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 105 .	(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 209). See Configuring vCenter Authentication on page 850 .	(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 753: 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.

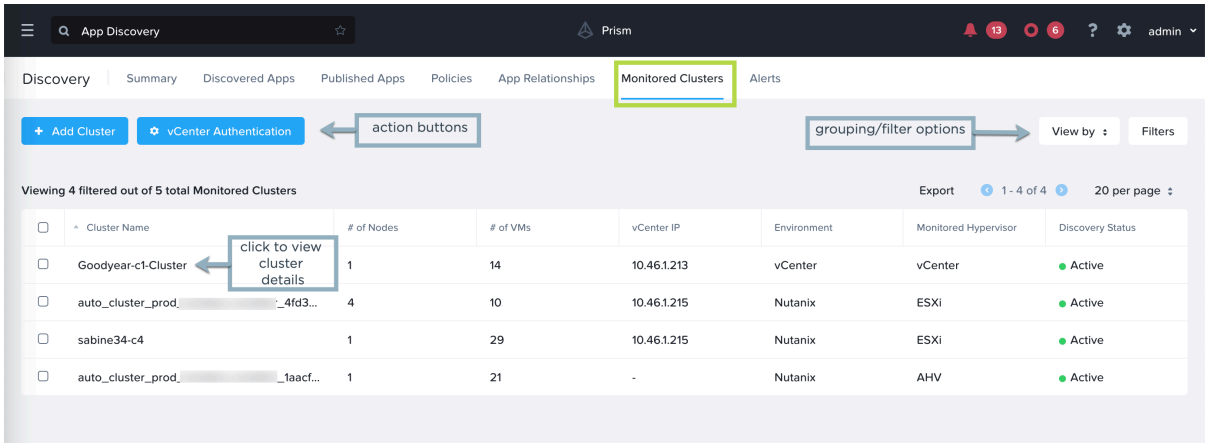

Figure 435: Monitored Clusters Tab

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 754: 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 209 .	(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 755: 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 848 .
- Click the **vCenter Authentication** button to authenticate application discovery on the monitored clusters with vCenter. See Configuring vCenter Authentication on page 850 .
- 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 (Prism Central) on page 269 .

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 846.

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.

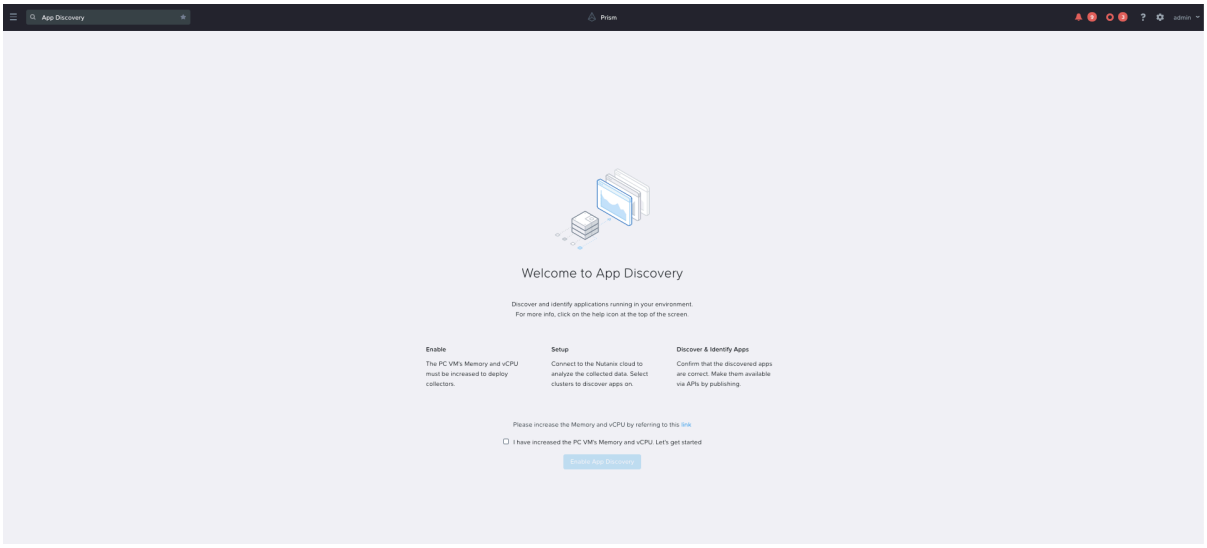

Figure 436: Welcome to App Discovery screen

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.

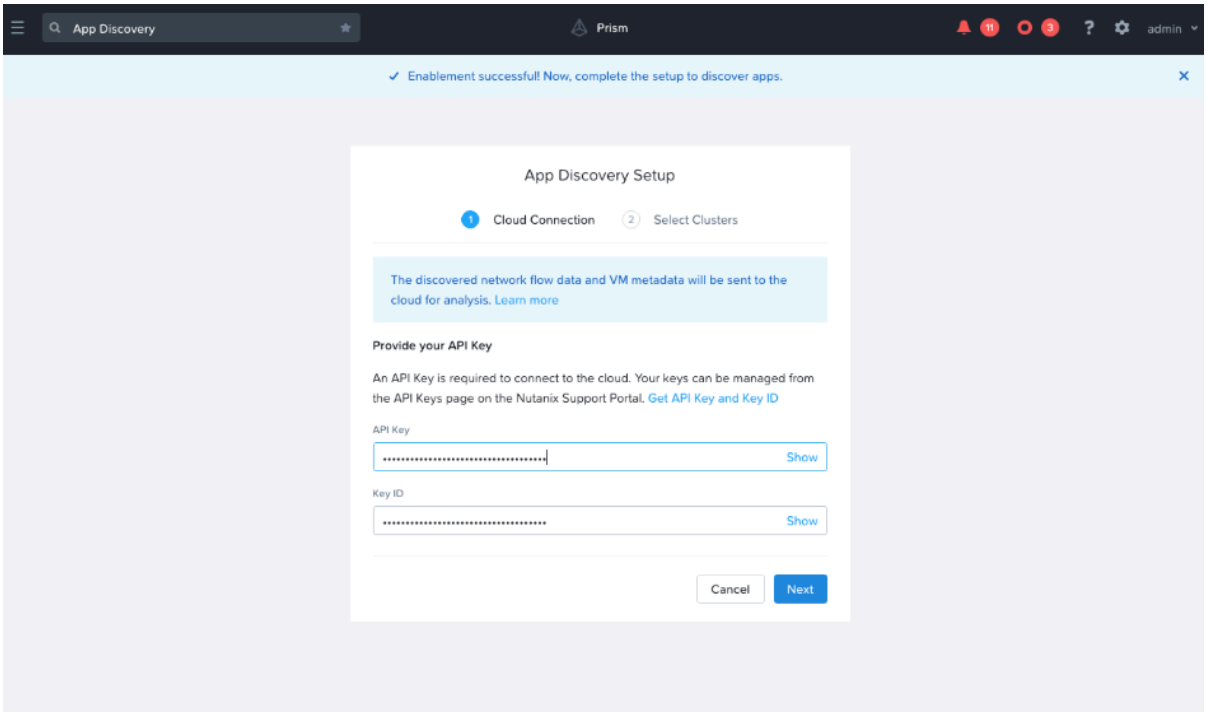

Figure 437: App Discovery Setup: Cloud Connection screen

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.

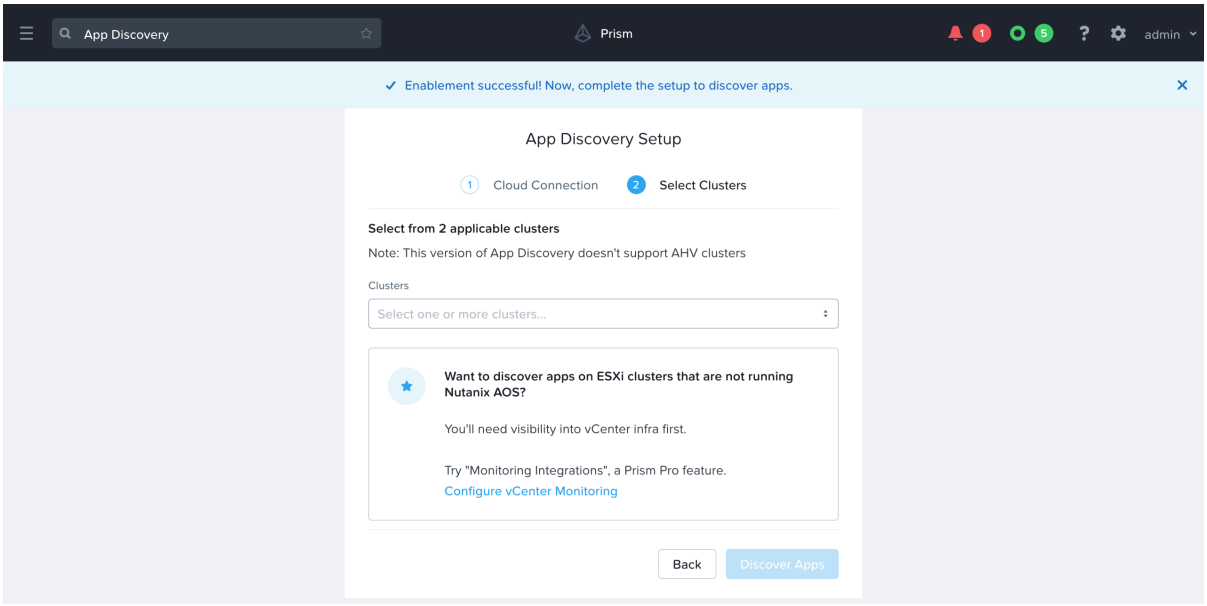

Figure 438: App Discovery Setup: Select Clusters (initial)

   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 719.

   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 832). This completes the setup for a Nutanix environment.

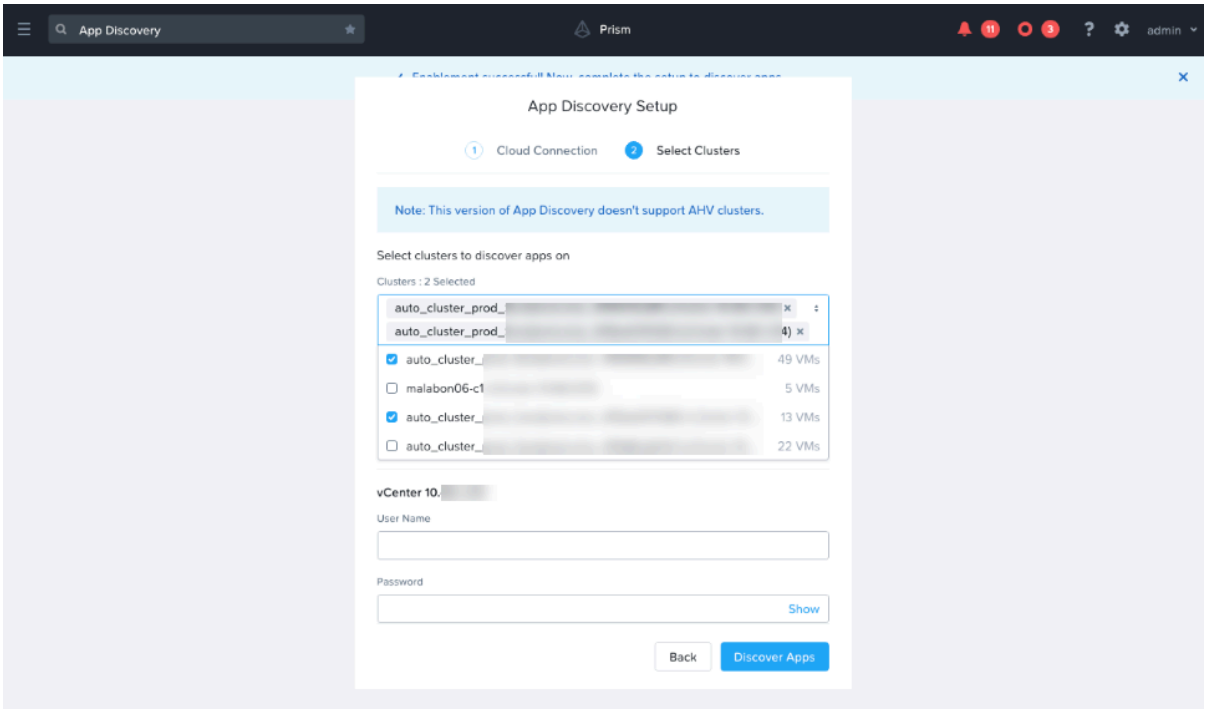

Figure 439: App Discovery Setup: Select Clusters

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

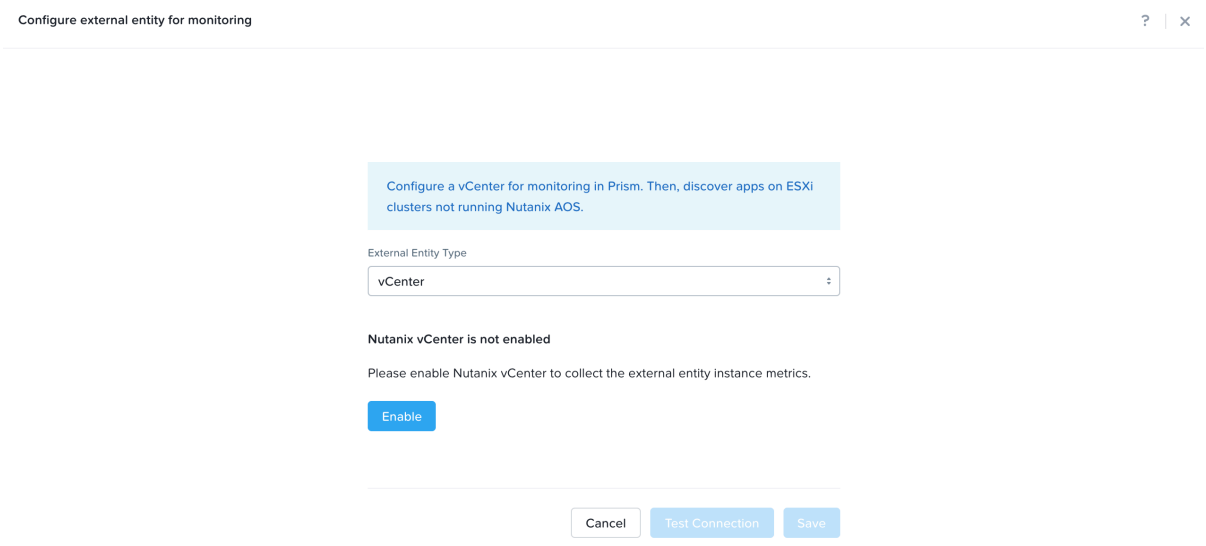

Figure 440: Enable vCenter Discovery

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.

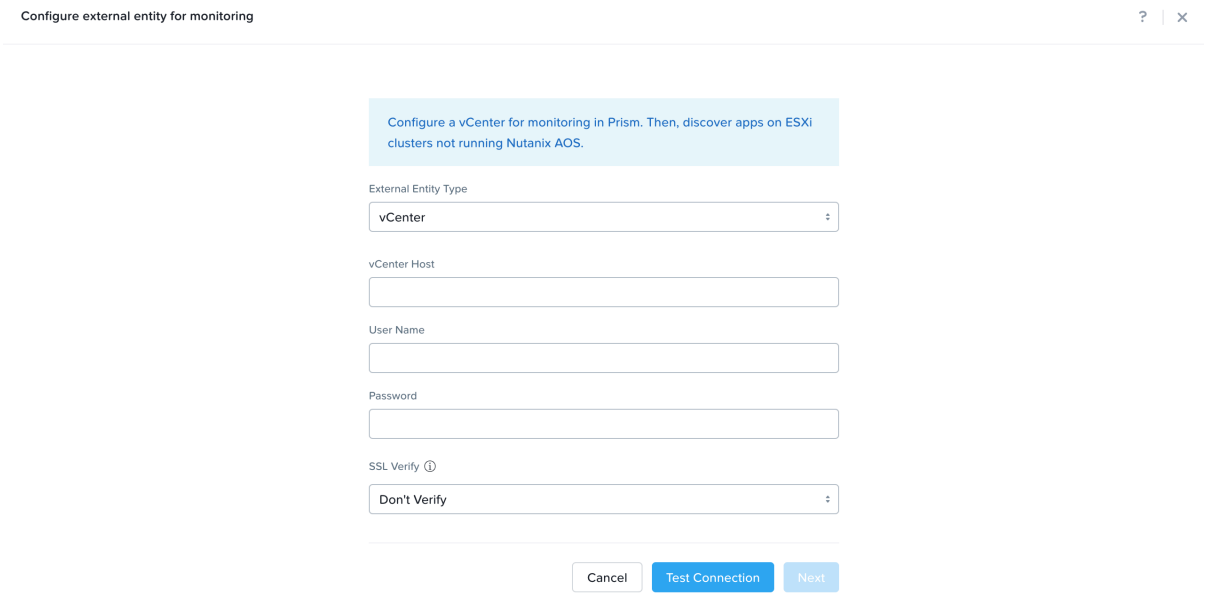

Figure 441: Configure vCenter Instance

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.

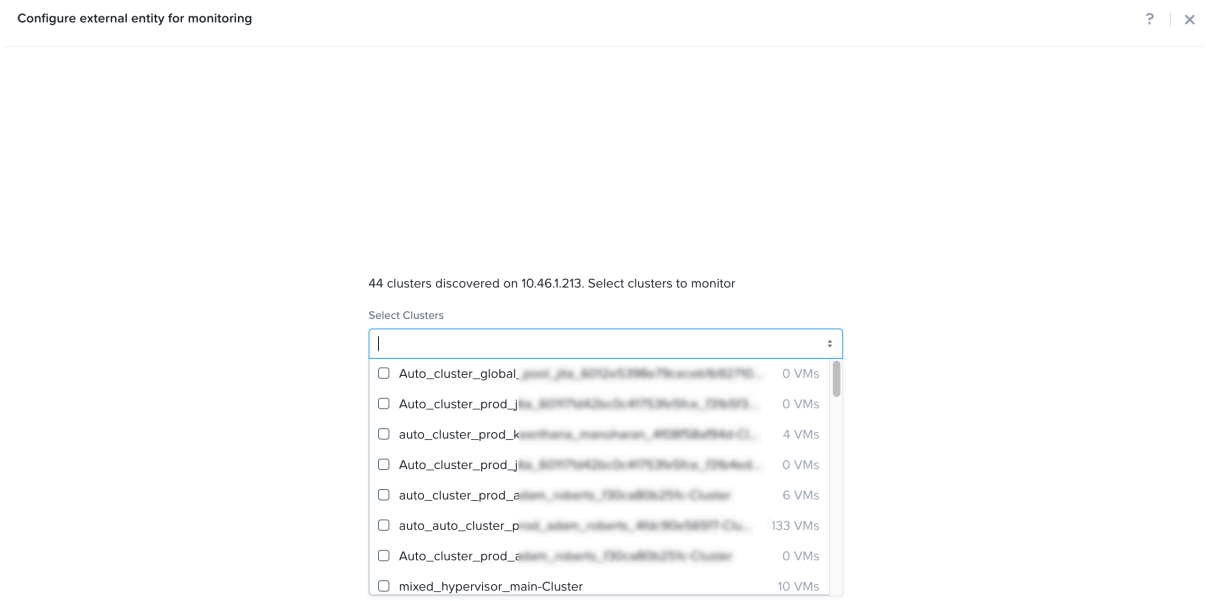

Figure 442: Select Clusters (vCenter)

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.

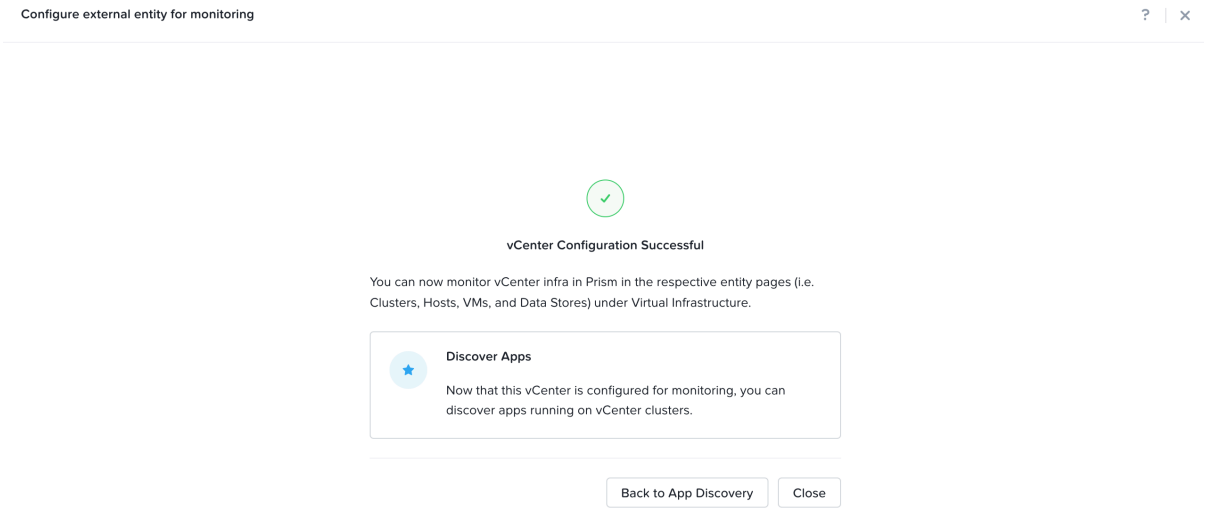

Figure 443: vCenter Configuration Complete

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 License Manager 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.

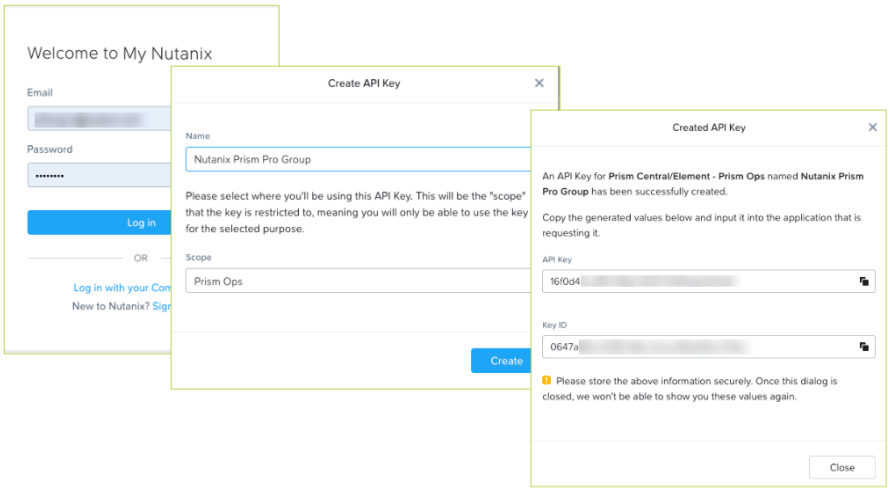

Figure 444: Create API Key Flow (example)

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 841), 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 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 832).

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.

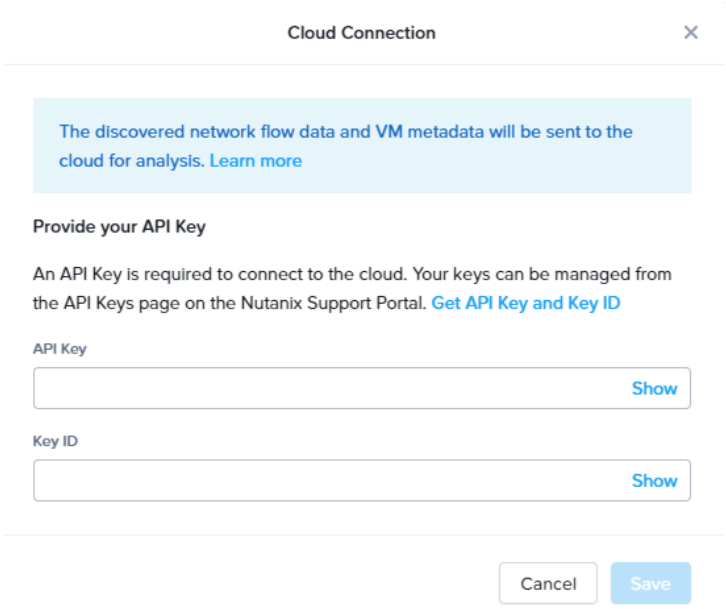

Figure 445: Cloud Connection Window

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 841), 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 832).

  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.

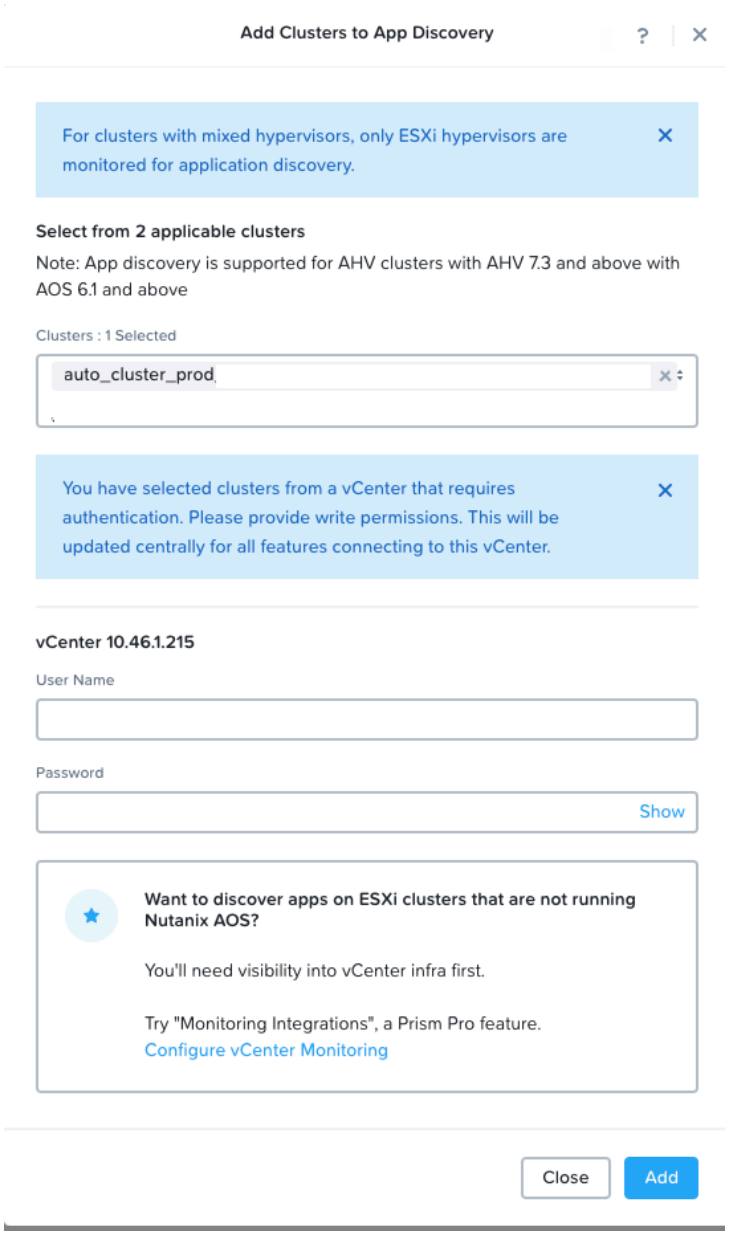

Figure 446: Add Clusters to App Discovery Window

  c. Run discovery to identify applications on the new clusters (see Discovering Applications on page 851).

- 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 841), 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 832).
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 719.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 868).

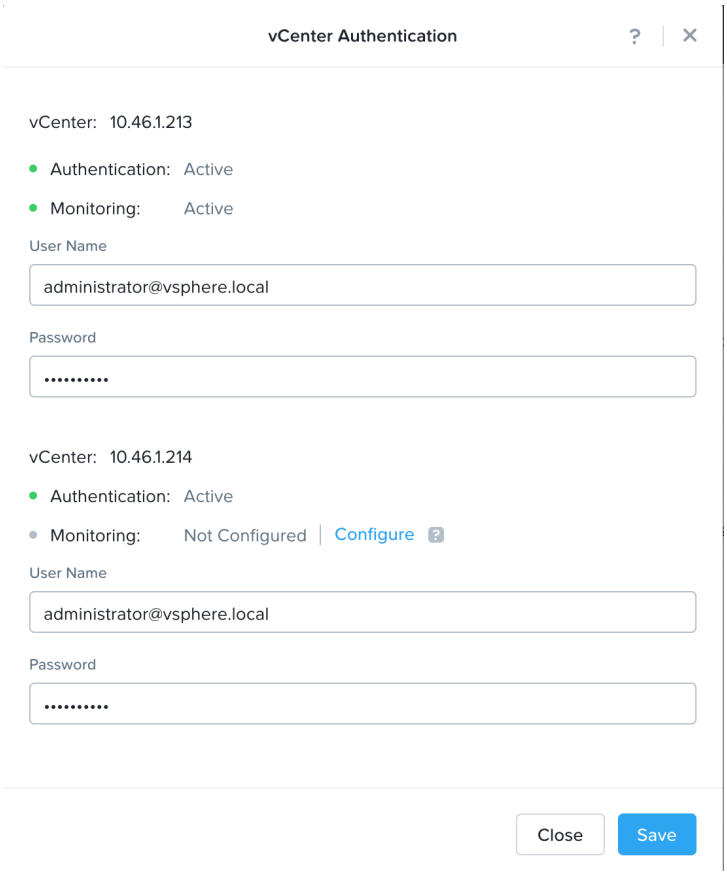

Figure 447: vCenter Authentication Window

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 832).
   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.

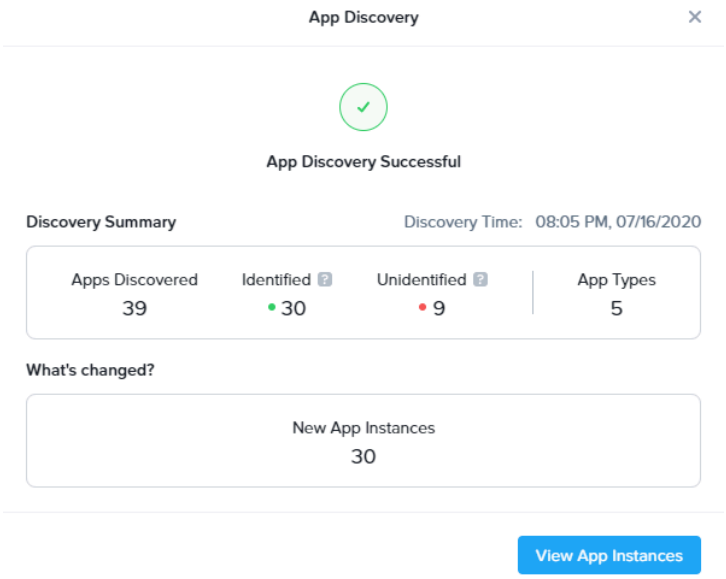

Figure 448: App Discovery Window

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 832).

Publishing Applications

You can publish the application discovery findings to a database.

About this task

Applications are discovered based on discovery policies that map applications and ports. You may validate, fix, and sanitize the list of discovered applications, and then publish the final list of discovered application. 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 832).
   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 832).

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 832).

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.

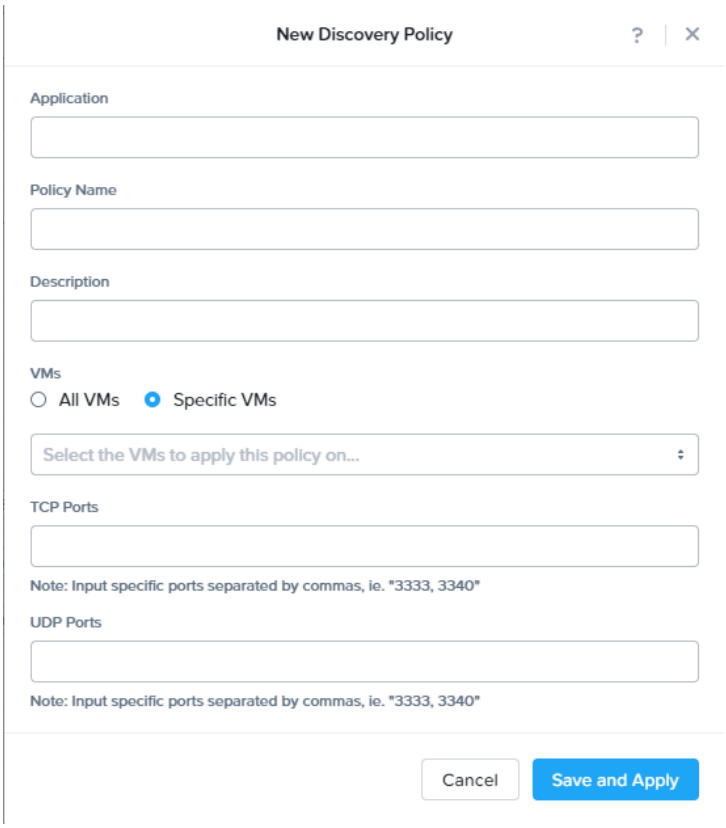

Figure 449: New Discovery Policy Window

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 832).

   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 852), change the entry in the appropriate field(s) and then click the **Save and Apply** button.

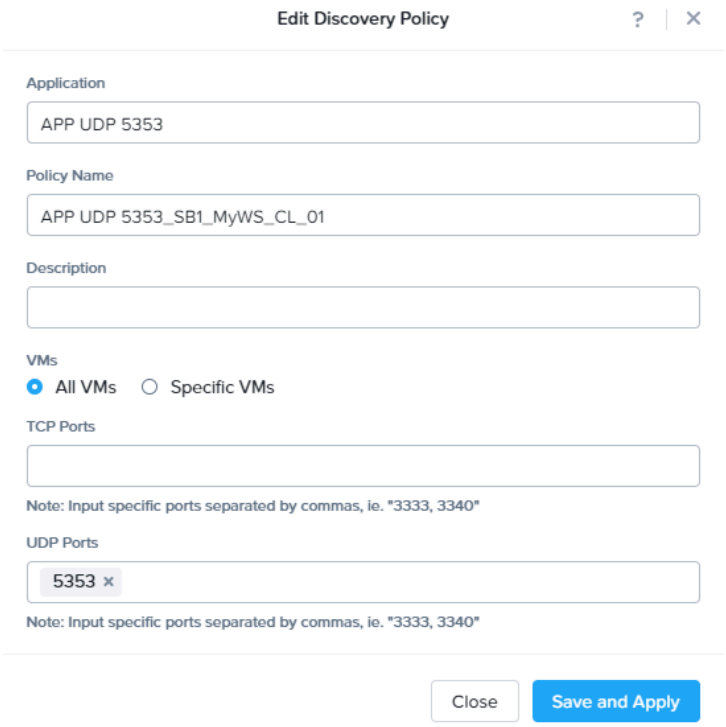

Figure 450: Edit Discovery Policy Window

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 997).
2. Set up a Xi Beam (SaaS) account.
3. Click the collapse (hamburger) menu button in the main menu (see Main Menu (Prism Central) 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 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:

- VMware vCenter Server: When you configure a vCenter Server for monitoring, the **Monitoring Configurations** page lists the vCenter Server. You can also view the vCenter Server listed on the **Non-Nutanix** page (**Compute & Storage** > **VMs** > **Non-Nutanix**). The vCenter details are also available on the **Hardware** > **Clusters** > **Non-Nutanix** page.

  See Application Monitoring Summary View on page 856 and Application Instance Details View (vCenter) on page 866.
- Microsoft SQL Server: When you configure a Microsoft SQL Server for monitoring, only the **Monitoring Configurations** page lists the SQL Server.

  See Application Monitoring Summary View on page 856 and Application Instance Details View (SQL Server) on page 858.

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 89 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 868 .

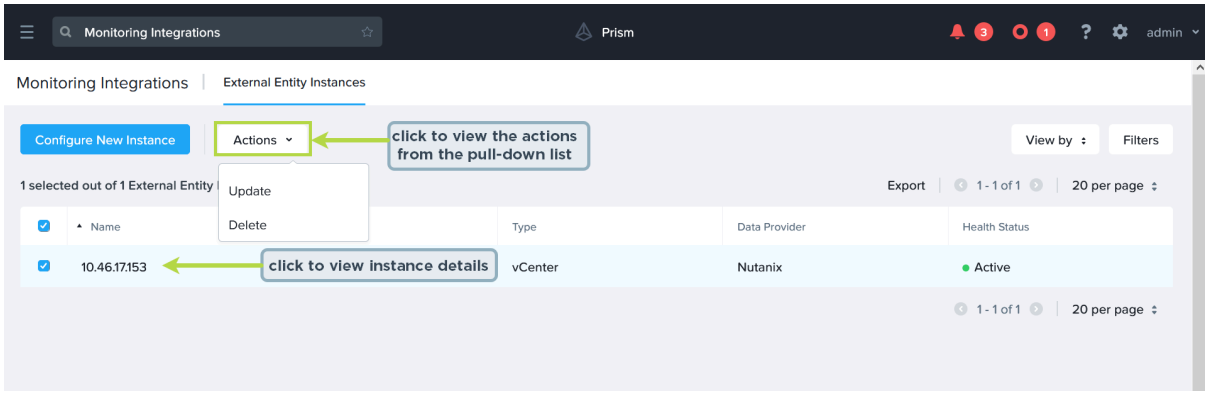

Figure 451: Monitoring Integrations Dashboard

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 756: 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 858 or Application Instance Details View (vCenter) on page 866).	(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 868).

- An **Actions** menu appears when an instance is selected:
  - Select **Update** from the pull-down list to update the instance configuration. See Modifying Application Monitoring on page 872 .
  - 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 856). 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.

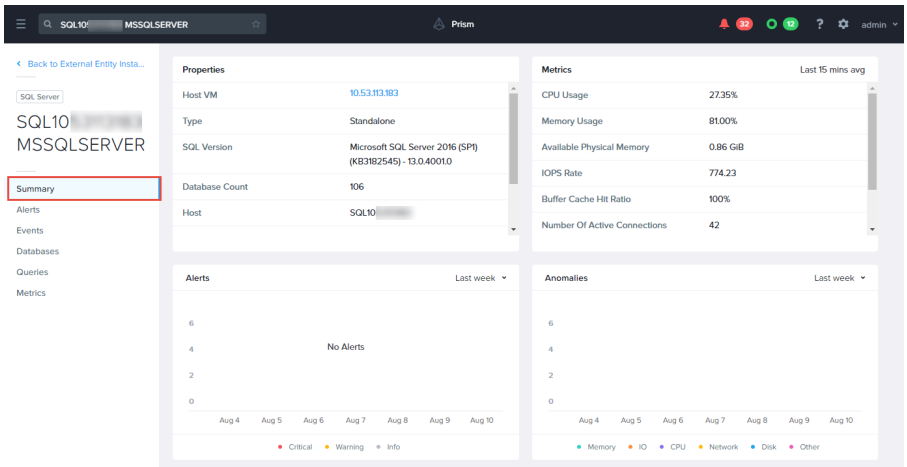

Figure 452: Application Instance Summary Tab

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 757: 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 105 .	(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 associated 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]
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 (Prism Central) on page 269 .

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 (Prism Central) on page 279 .

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.

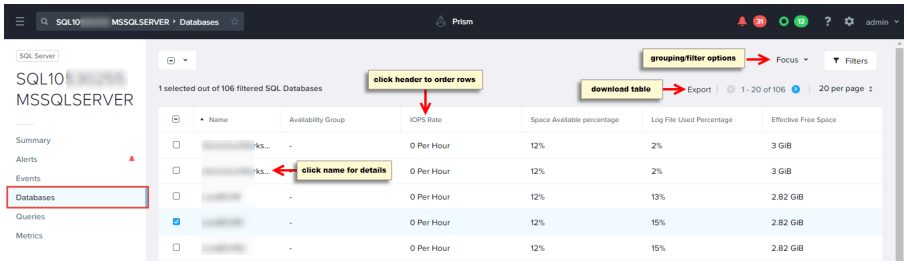

Figure 453: Application Instance Databases Tab

- 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 89).
- 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 758: 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

Parameter	Description	Values
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.

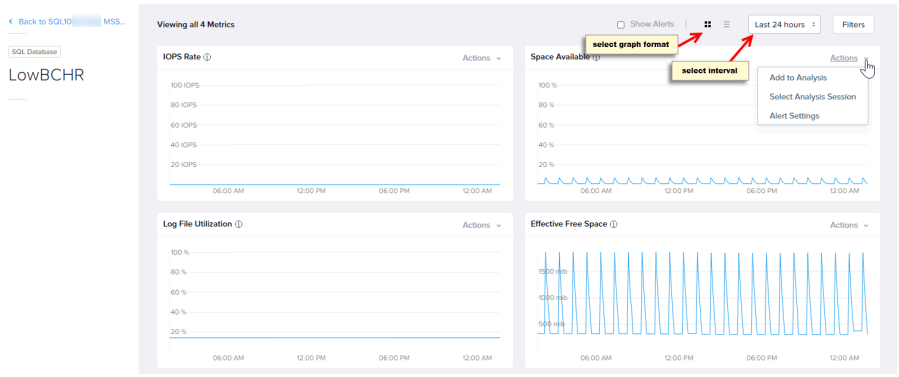

Figure 454: Database Instance Page

- 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 (Prism Central) on page 799 .
  - 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.

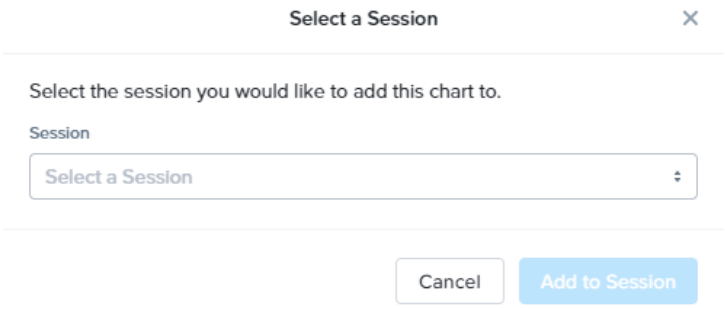

Figure 455: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric. See Creating Custom Alert Policies on page 288 .
- 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.

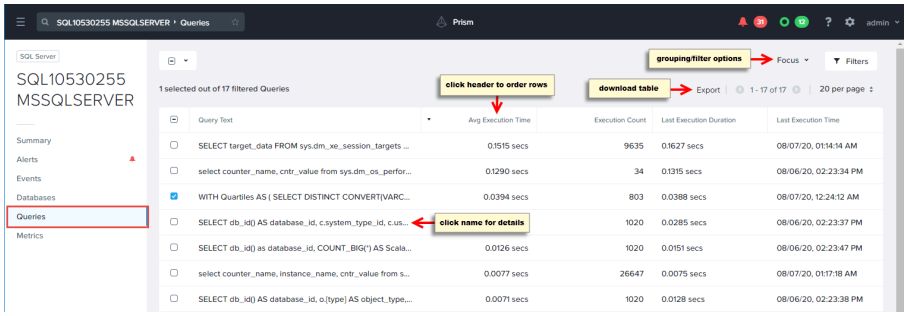

Figure 456: Application Instance Queries Tab

- 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 89 .
- 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 759: 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.

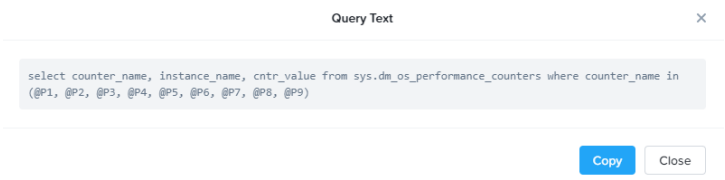

Figure 457: Query Statement Window

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

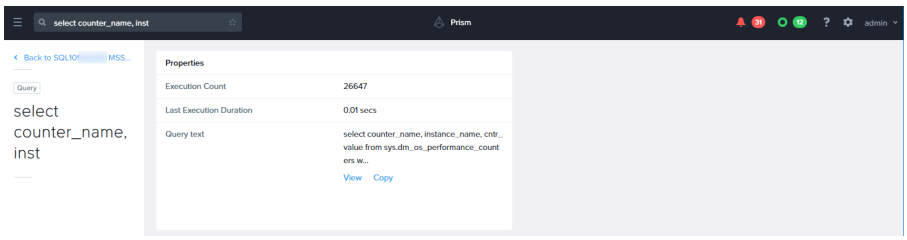

Figure 458: Query Details Page

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.

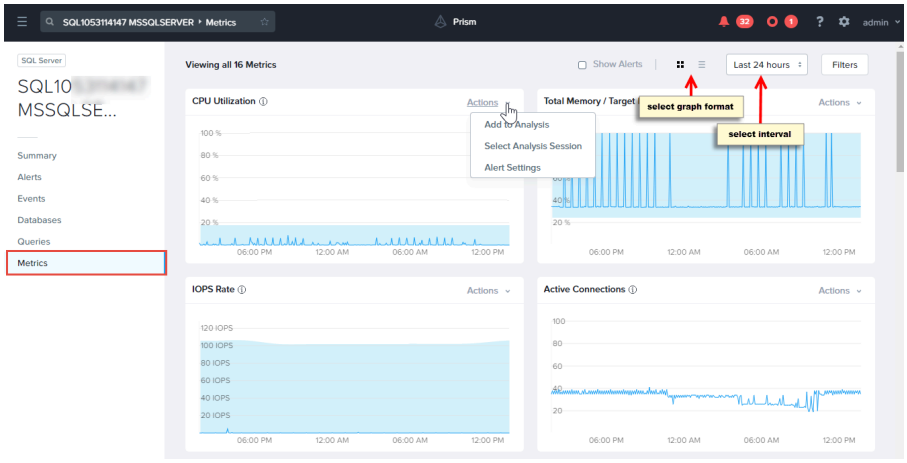

Figure 459: Application Instance Metrics Tab

- 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 (Prism Central) on page 799 .
  - 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.

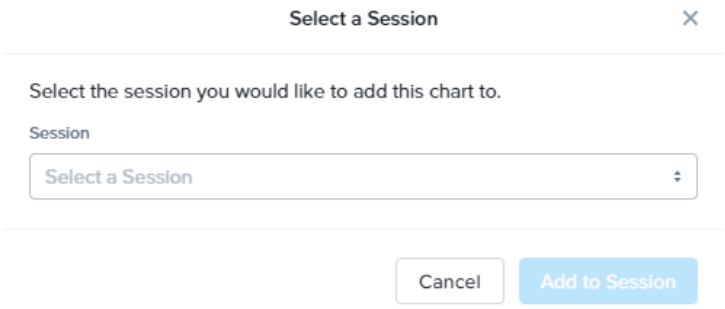

Figure 460: Session Window

  - Select **Alert Settings** to create a custom alert policy for the metric. See Creating Custom Alert Policies on page 288 .
- 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 760: 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.
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.

Metric	Description
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 856). 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.

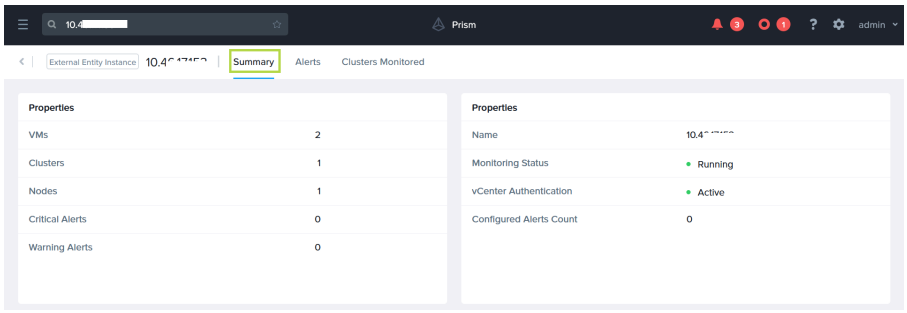

Figure 461: Application Instance Summary Tab

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 761: 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 (Prism Central) on page 269 .

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 94 . 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.

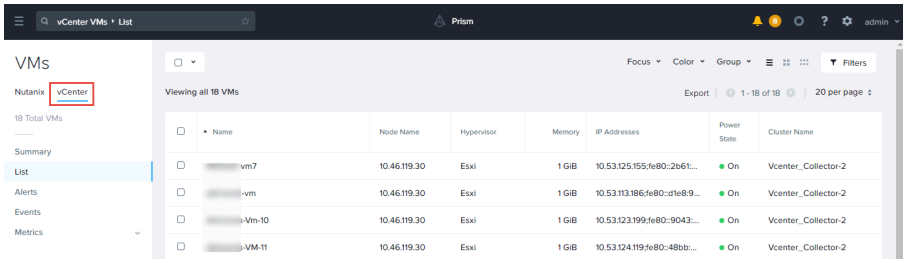

Figure 462: VMs dashboard (vCenter tab)

- 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 202 . 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.

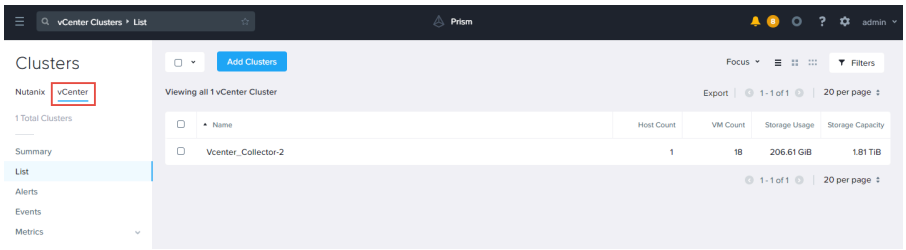

Figure 463: Clusters dashboard (vCenter tab)

- 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 221 .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.

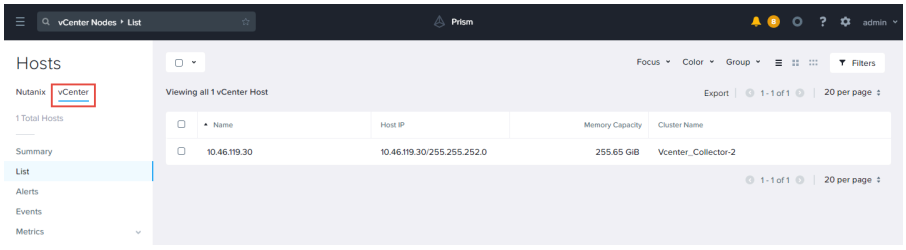

Figure 464: Hosts dashboard (vCenter tab)

Configuring Application Monitoring

Monitoring an application instance requires configuring an instance profile.

Before you begin

Complete the prerequisites for configuring application Monitoring (see Prerequisites for Configuring Application Monitoring on page 871).

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 856).

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. You need to log in as a user with super admin role to perform this operation.

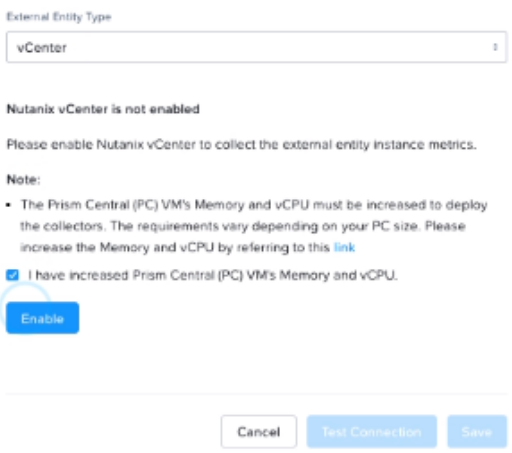

Figure 465: Configuration Page (enable)

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

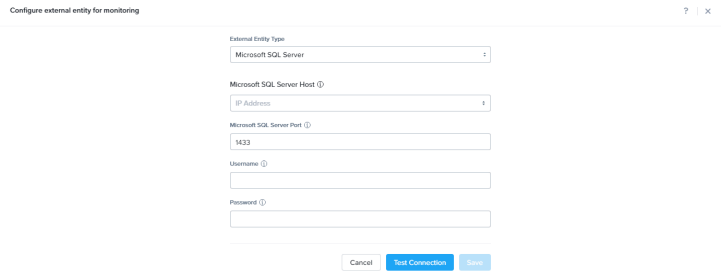

Figure 466: Configuration Page (SQL server example)

   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 only] **Microsoft SQL Server Port**: Enter the target port number.

      The port number defaults to the standard port number for SQL Server (1433). For SQL Server named instance, enter the port number of the SQL server instance you want to monitor.

   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.

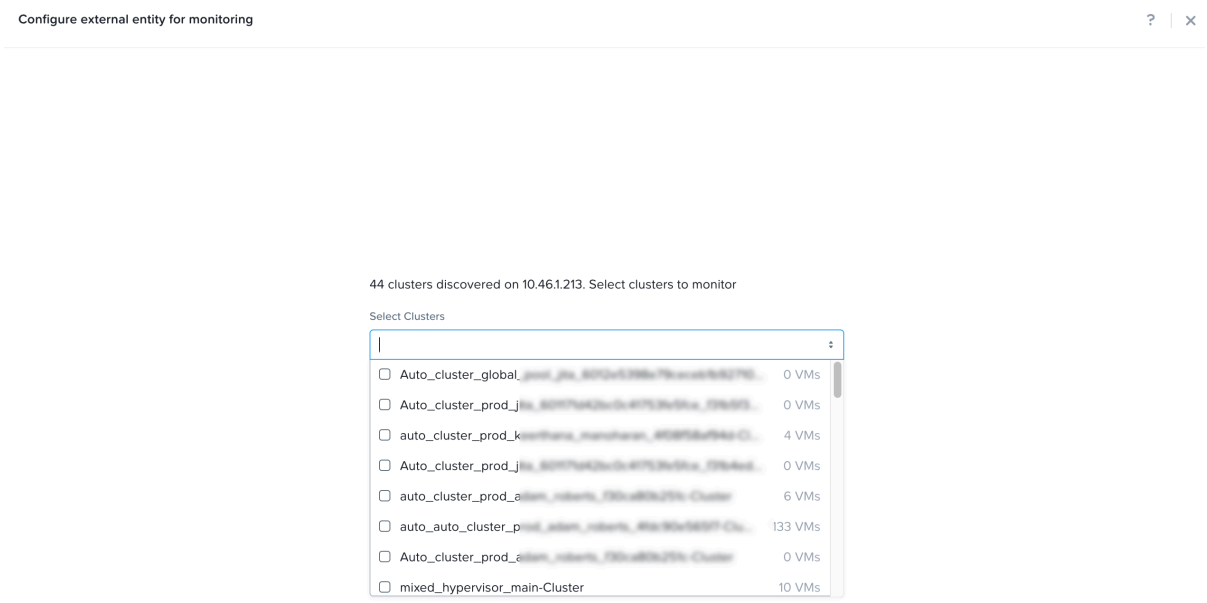

Figure 467: Select Clusters

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 866).

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 841).

Prerequisites for Configuring Application Monitoring

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.

- TLS version 1.2.
- SQL Server-specific requirements:
  - Sufficient access permissions to create and view monitoring configurations for vCenter Servers and Microsoft SQL Servers.
  - Microsoft 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.
    - To view the Microsoft SQL Servers you have configured for monitoring, you need the following user permissions configured:
      - **Roles** > **Monitoring Configurations**: Select **View Only**. You can provide **Update Access** or **Full Access** (Create, Update, Delete, View instance configuration and View MSSQL instance) permissions if required. **Monitoring Configurations** > **View Only** permissions are minimum required permissions to **View** MS SQL Servers that are already configured for monitoring.
      - **Roles** > **SQL Server**: Select **View Only**.

        Note: You cannot provide **SQL Server** > **View Only** permission without providing **Monitoring Configurations** > **View Only** or greater permissions for **Monitoring Configurations**.

      Therefore, ensure that you have a minimum of **Monitoring Configurations** > **View Only** and **SQL Server** > **View Only** permissions to view the monitored SQL Servers on the **Monitoring Configurations** page.

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 856).

   The **Configure external entity for monitoring** page for that instance appears (see Configuring Application Monitoring on page 868).

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 899). 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. For more information, see Managing a VM through Prism Central (AHV) on page 552.
- **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. For more information, see Managing a VM through Prism Central (AHV) on page 552.

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 89 for instructions on how to view and organize that information in a variety of ways.

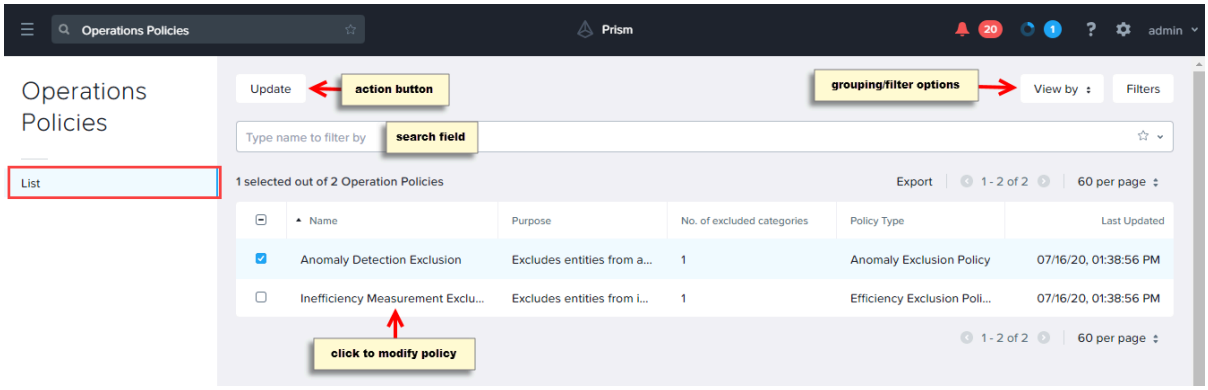

Figure 468: Operations Policies Dashboard

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 762: Operations Policy Fields

Parameter	Description	Values
Name	Displays the operations policy name. Click the name to display the policy details for **Anomaly Detection Exclusion** or **Inefficiency Measurement Exclusion** operations policies.	Anomaly Detection Exclusion | Inefficiency Measurement Exclusion
Purpose	Describes the purpose of the operation policy.	(description, string)
No. of Excluded Categories	Displays the number of categories that are excluded for anomaly detection or inefficient measurements checks. For information on how to create a category, see Category Management on page 795.	(integer)
Policy Type	Describes the type of operations policy.	Anomaly Detection Exclusion | Inefficiency Measurement Exclusion
Last Updated	Displays the date and time when the operations 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 874).

Modifying an Operations Policy

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

Before you begin

Ensure that the VM categories are created, and the target VMs are added to the VM categories.

For information about how to create a category, see Creating a Category on page 795.

For information about how to assign a category to the VMs, see Assigning a Category on page 797.

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 873).

2. In the policy page, click the **Add** link.

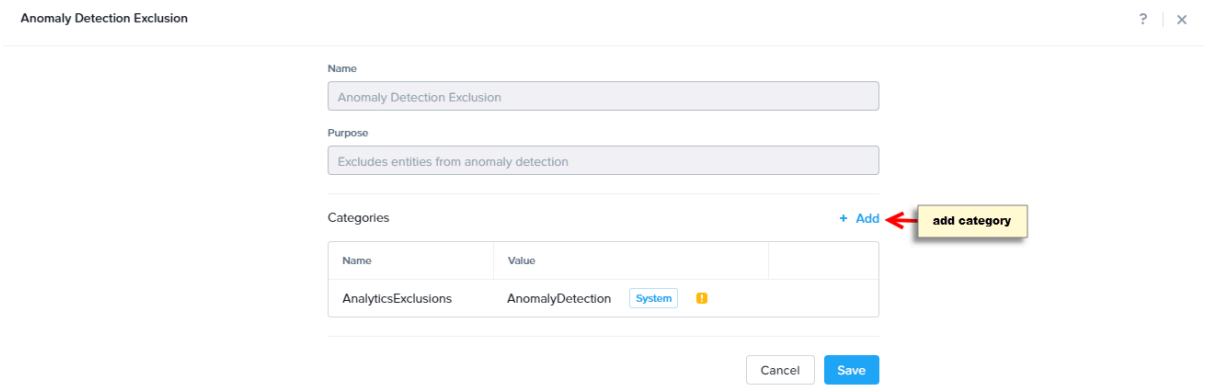

Figure 469: Anomaly Detection Exclusion Page (add category)

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.

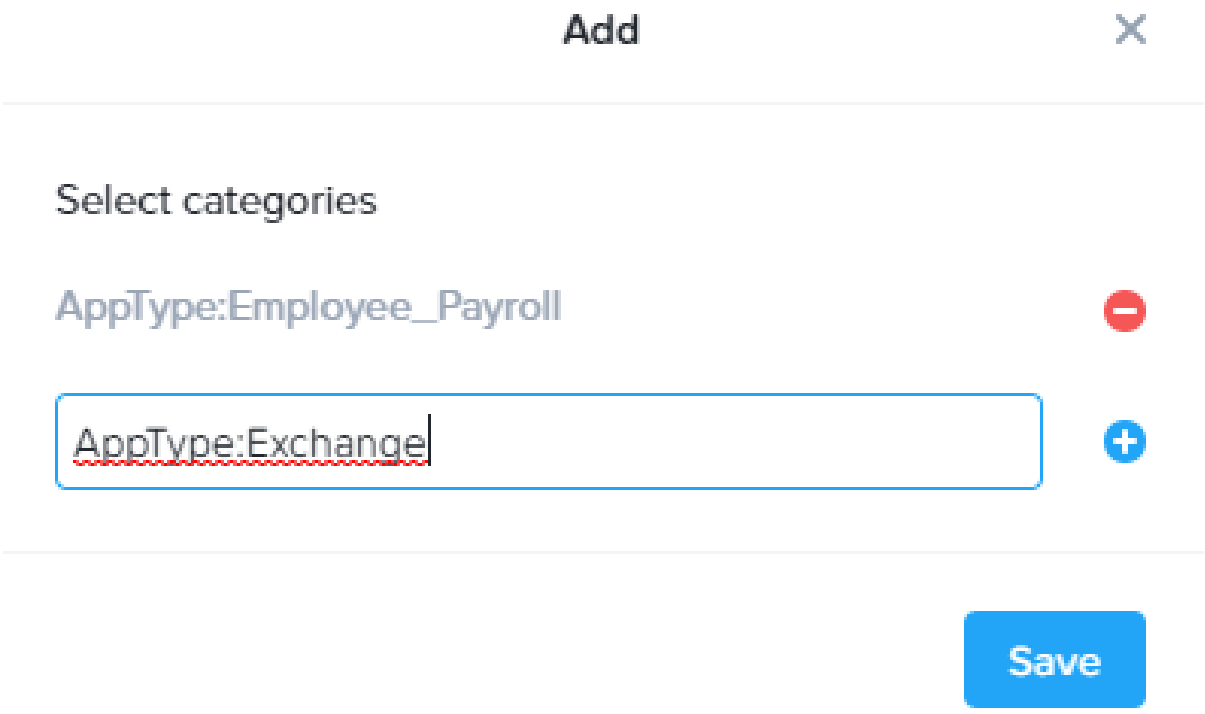

Figure 470: Add Window

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.

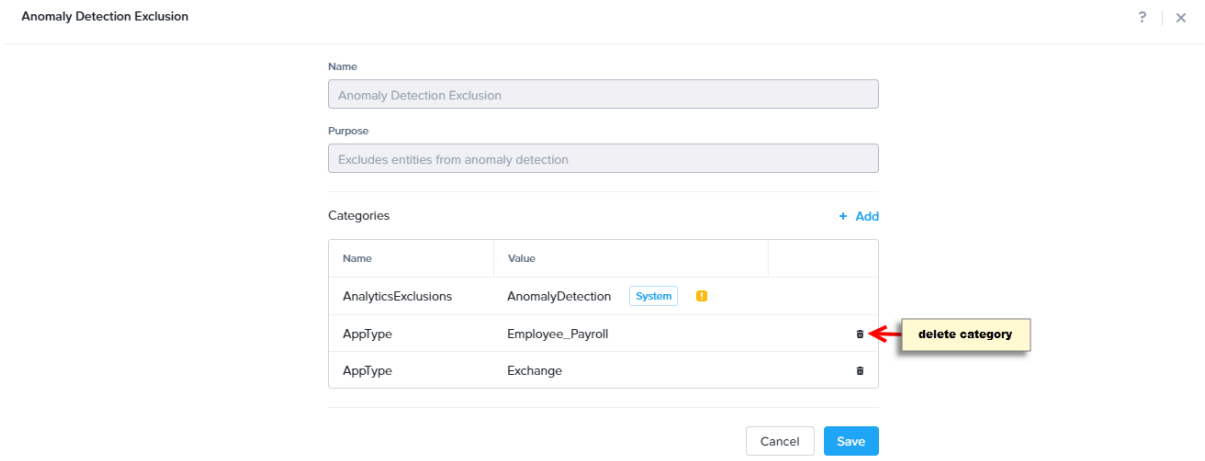

Figure 471: Anomaly Detection Exclusion Page (delete category)

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.

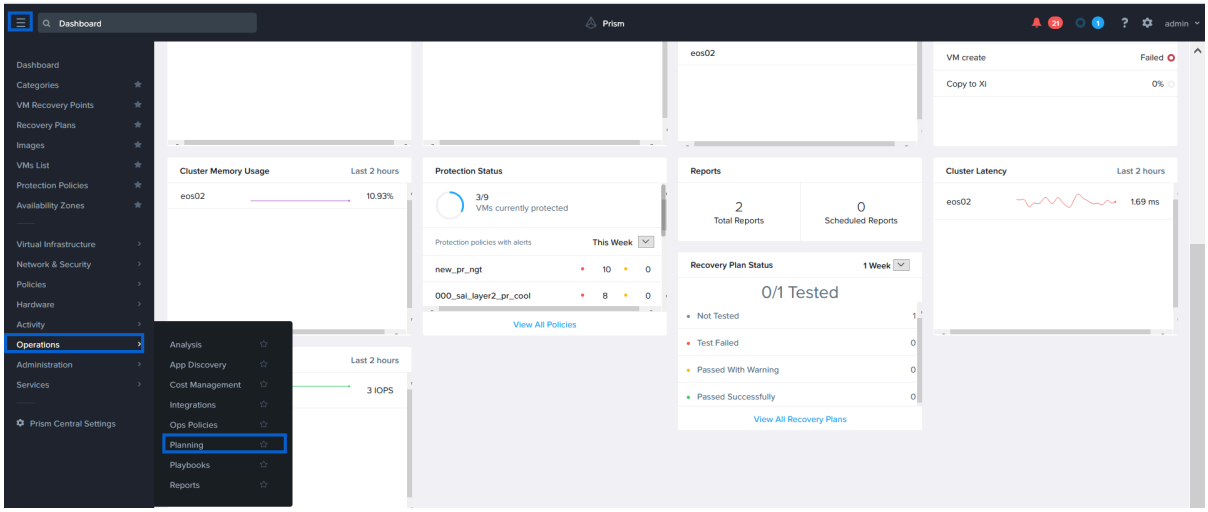

Figure 472: Planning

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 877).

- 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 878).
- 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 898.

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 899).

Capacity Runway View

Click the **Capacity Runway** tab in the planning dashboard (see Resource Planning on page 876) 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.
- When you determine the current usage (Total Used Space) for a cluster that has Rebuild Capacity Reservation enabled, Nutanix recommends that you ignore the Effective Capacity threshold displayed in the Storage Runway graph of the Prism Central Capacity planning report. Instead, you must rely on the Total Capacity and Resilient Capacity numbers on the Storage Summary widget in the Prism Element web console. This is because the Storage Runway report in Prism Central for a cluster does not take into account the Rebuild Capacity Reservation of the cluster when calculating the Effective Capacity threshold in the report.
- 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 856 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 209).

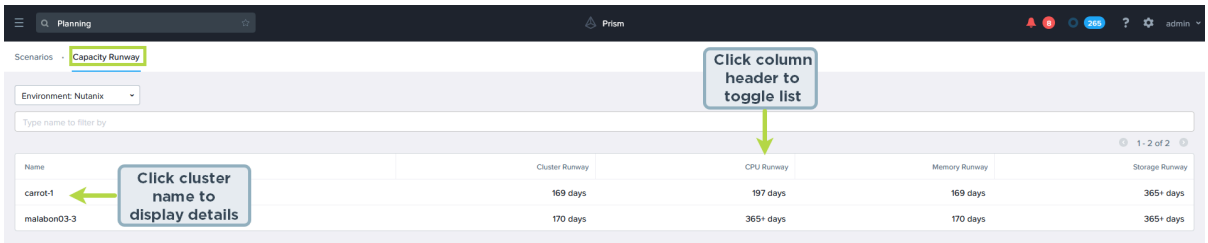

Figure 473: Capacity Runway View (Nutanix)

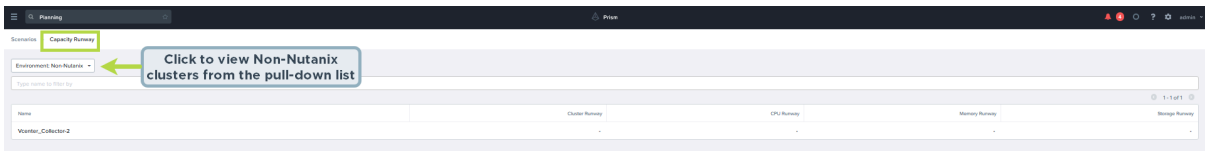

Figure 474: Capacity Runway View (Non-Nutanix)

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

Table 763: 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 209).	(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 876) 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 879).
- 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.

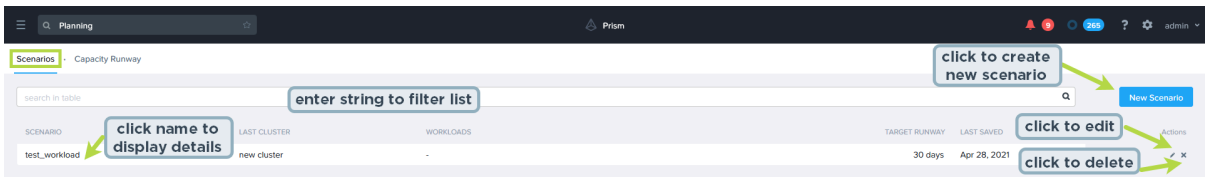

Figure 475: Scenarios View

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

Table 764: 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 879).	(scenario name)
Last Cluster	Displays the name of the cluster to which the scenario applies. (The column title is "Last Cluster" because you can change the target cluster when modifying a scenario.)	(cluster name)
Workloads	Displays the workloads included in the scenario. This field lists workloads added to the scenario. In the figure above, the "tom" scenario lists "sql" in this field indicating an SQL workload was added to the scenario. A blank field indicates either there is no workload (for a scenario based on a new cluster) or only the existing (current) workload is applied (for a scenario based on an existing cluster).	(blank),
Target Runway	Displays the target runway duration for this scenario, that is how many days am I planning for. This is specified when the scenario is created or modified ("Target" parameter).	xxx [days]
Last Saved	Displays the time and date for the last (most recent) time the scenario was saved.	(time and date)
Actions	• To edit the scenario, click the pencil icon (see Modifying a Scenario on page 895). • To share the scenario, click the share icon. • 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 878) and click the **New Scenario** button.
   The Scenario configuration page appears.

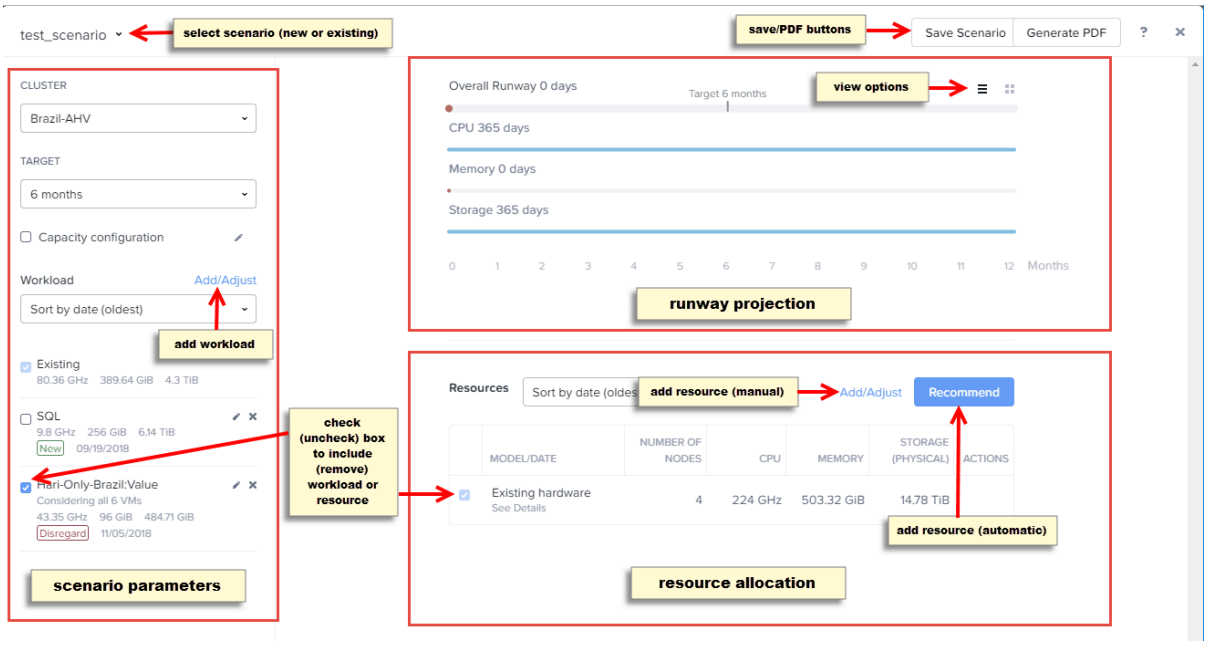

Figure 476: Scenario Configuration Page (example)

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 878), 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 Workloads**: Specify the ratio of virtual CPU to physical CPU. (The default is 1.)
      6. **Memory Overcommit Ratio for New VM Workloads**: 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.

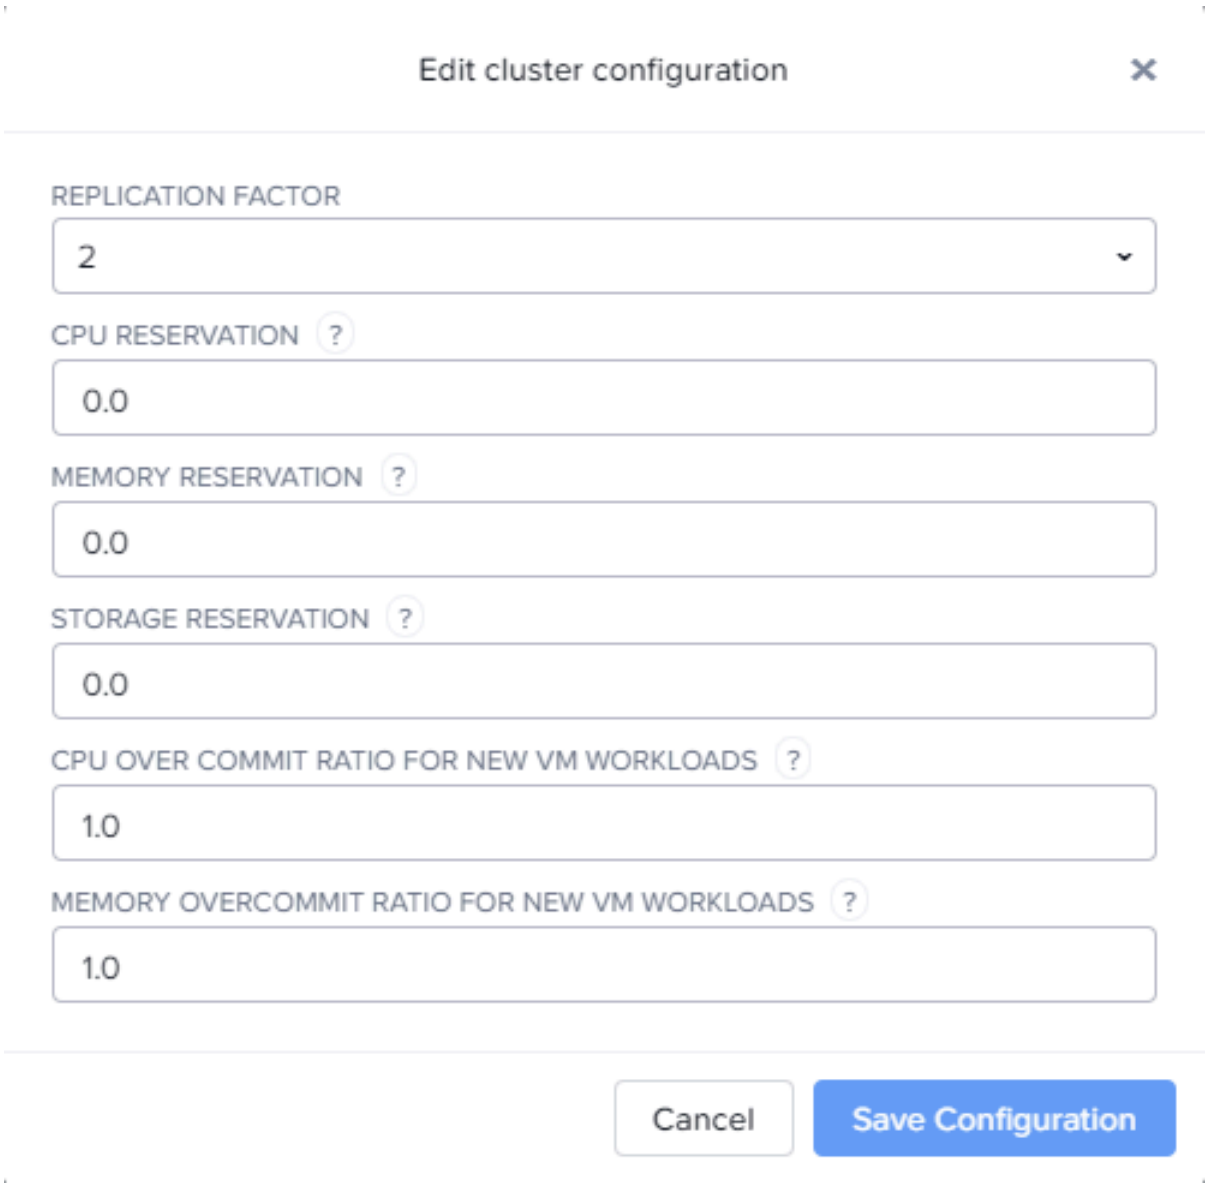

Figure 477: Edit Cluster Configuration Window

> **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.

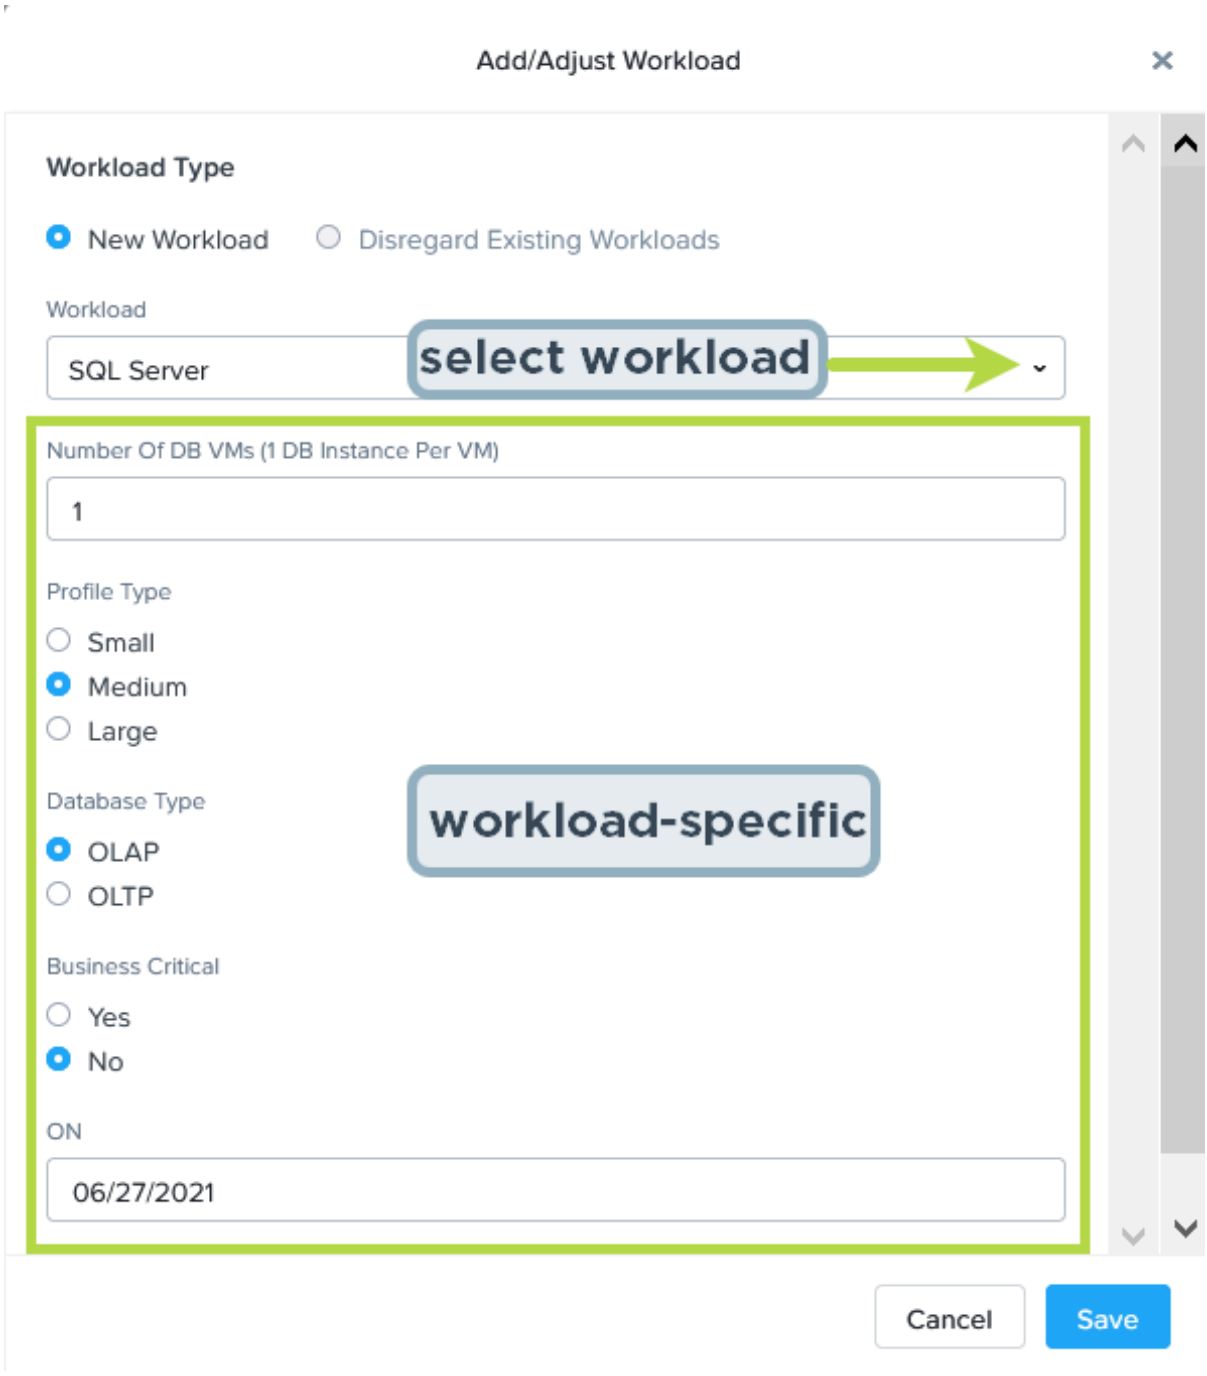

Figure 478: New Workload Fields (SQL Server example)

Table 765: 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)
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

Workload	Field	Description	Values
	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 from the following options: • **Task worker (default)**: The Task Worker profile is designed to run fewer applications than the other workloads. • **Knowledge worker**: The Knowledge Worker profile is designed to run the standard office and email applications with browsers. • **Power user**: The Power user profile is the most intensive of the standard workloads. • **Developer**: The Developer profile is for super heavy user demands. See the table Table 766: Resources for VDI User Types on page 887 for resources for each of the profiles.	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
	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)

Workload	Field	Description	Values
	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]
	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)

Workload	Field	Description	Values
	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)

Table 766: Resources for VDI User Types

User Type	vCPU	vCPU: pCore Ratio	RAM	User Data Capacity	Gold Image Capacity
Task worker (default)	2	12	2 GiB	5 GiB	20 GiB
Knowledge worker	2	10	4 GiB	10 GiB	20 GiB
Power user	2	6	8 GiB	25 GiB	20 GiB
Developer	2	2	16 GiB	50 GiB	20 GiB

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 892 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 795).

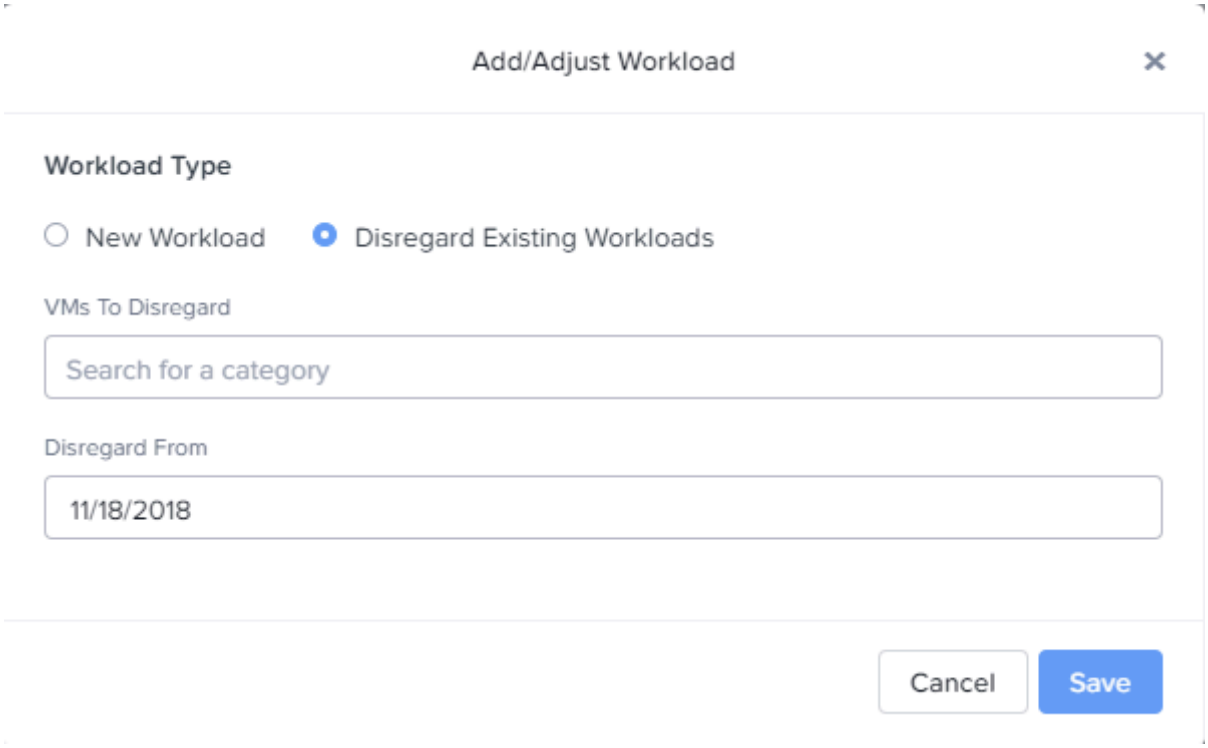

Figure 479: Disregard Existing Workloads Fields

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.

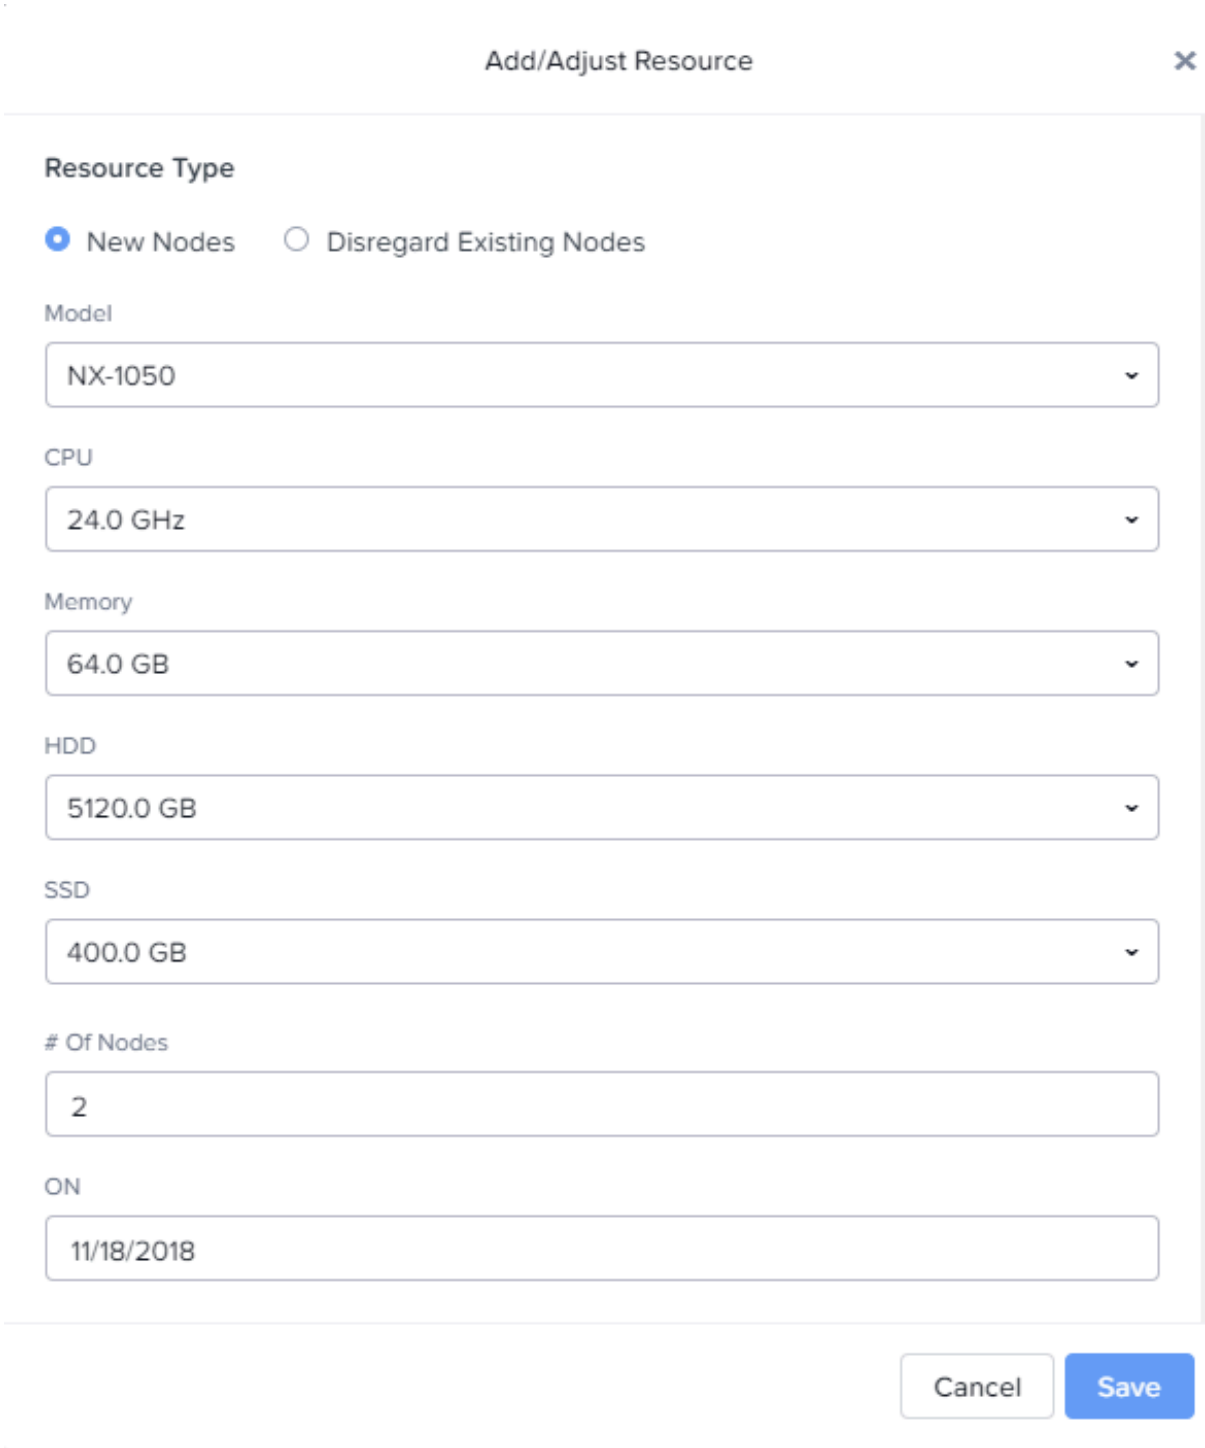

Figure 480: New Nodes Fields

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.

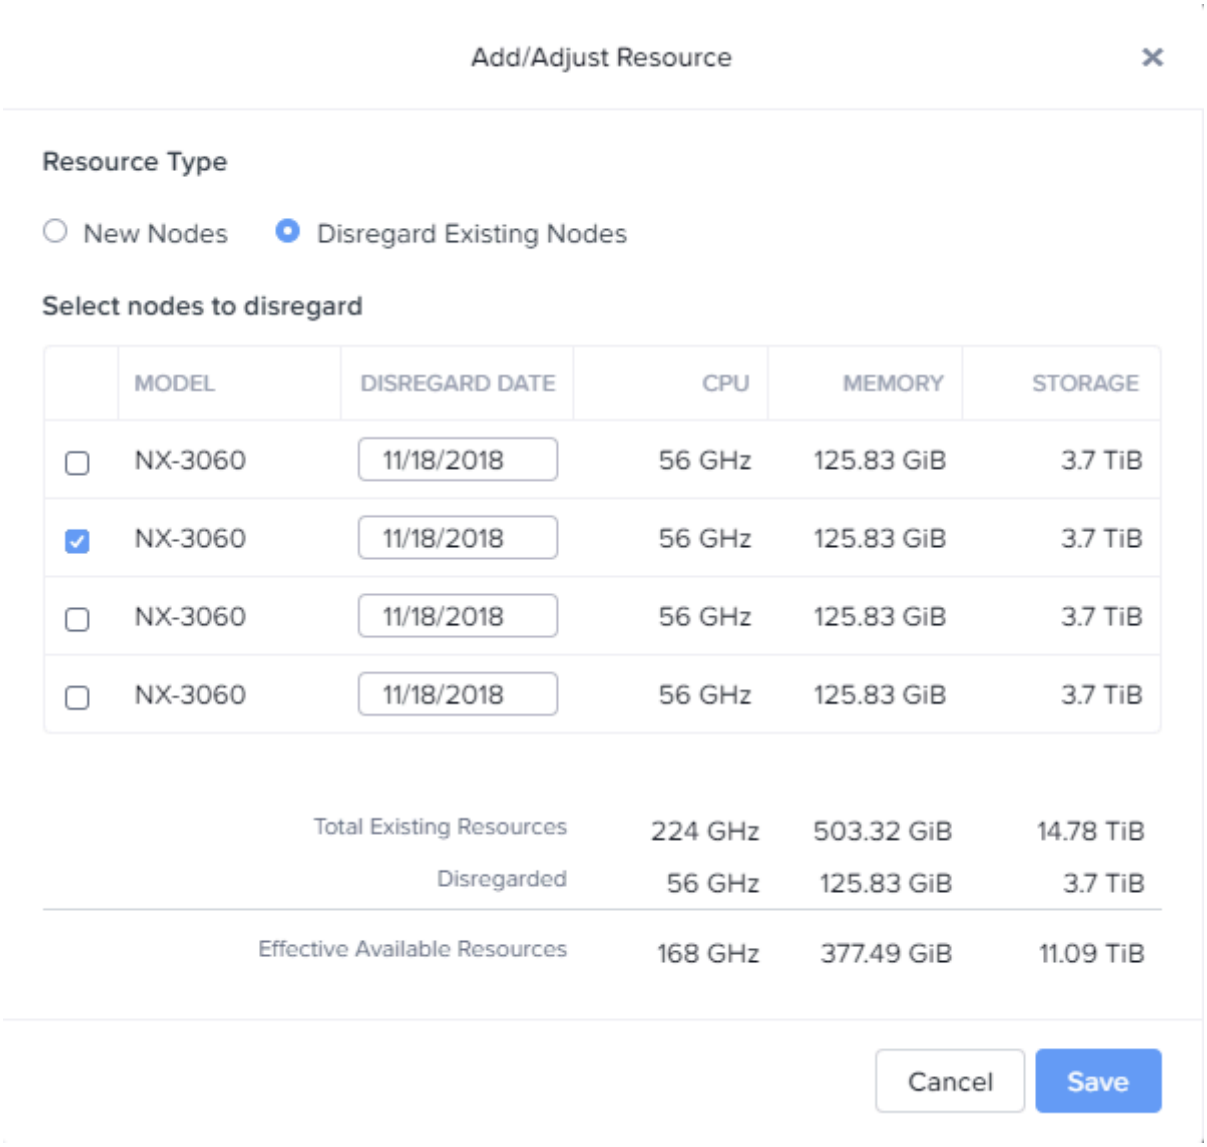

Figure 481: Disregard Existing Nodes Fields

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 878).

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.

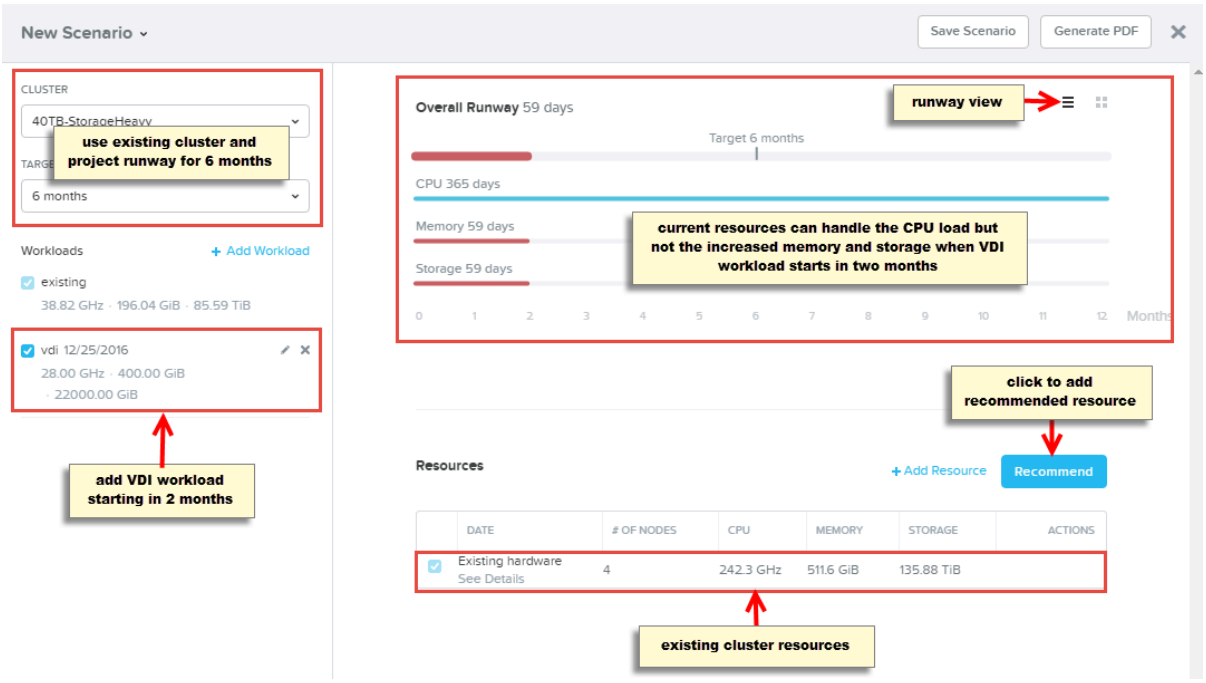

Figure 482: Scenario Screen (with added workload)

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.

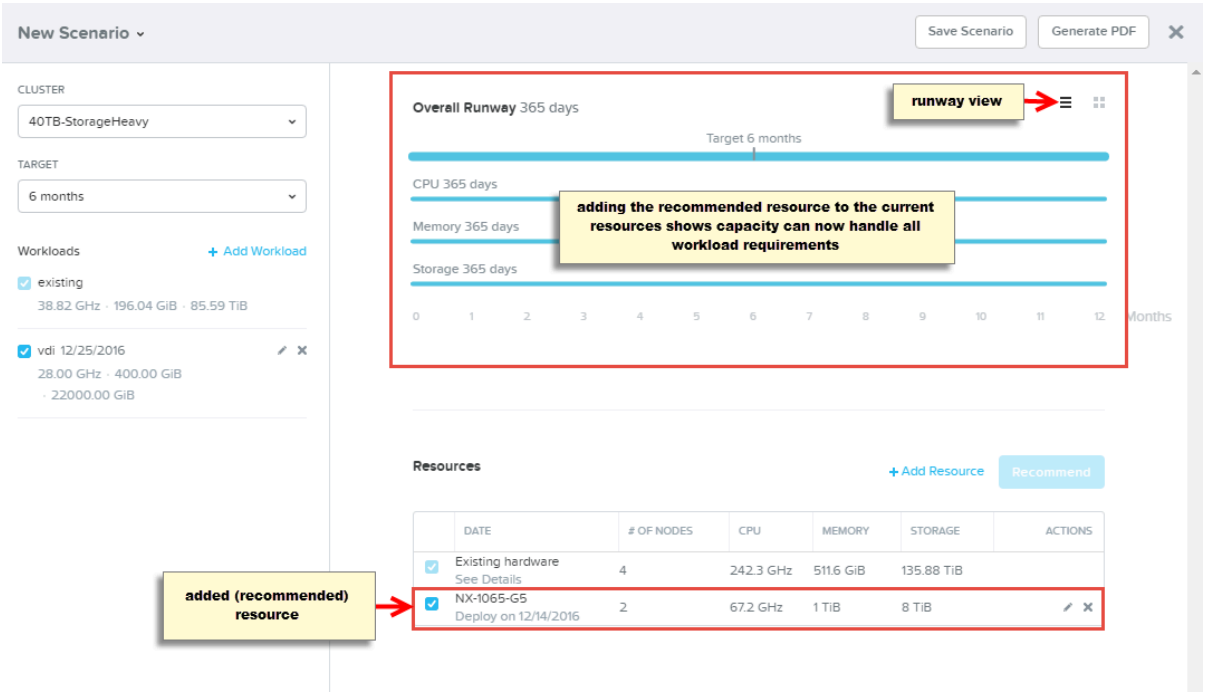

Figure 483: Scenario Screen (with added workload and recommended resource)

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.

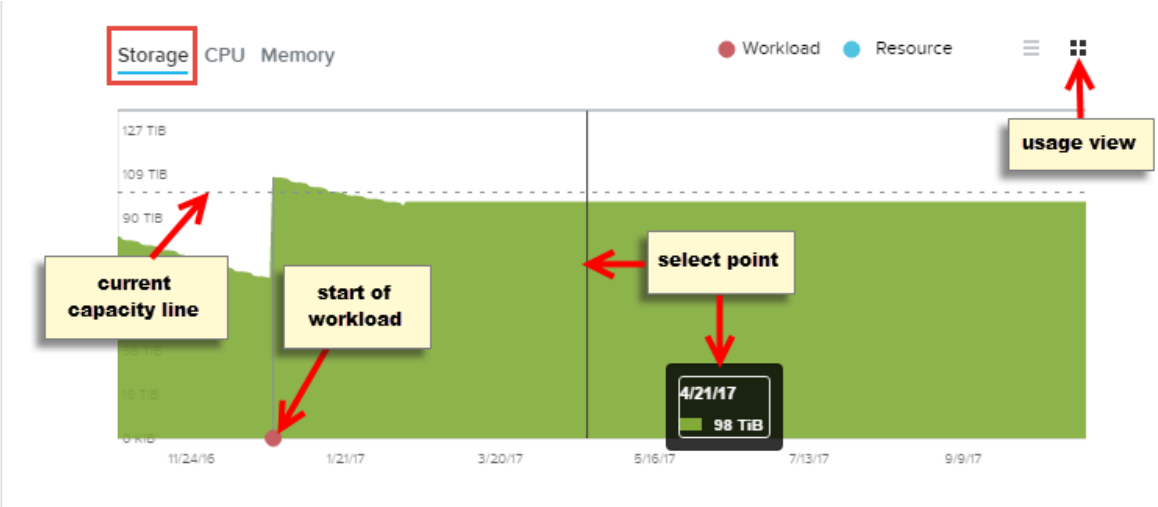

Figure 484: Storage Usage View

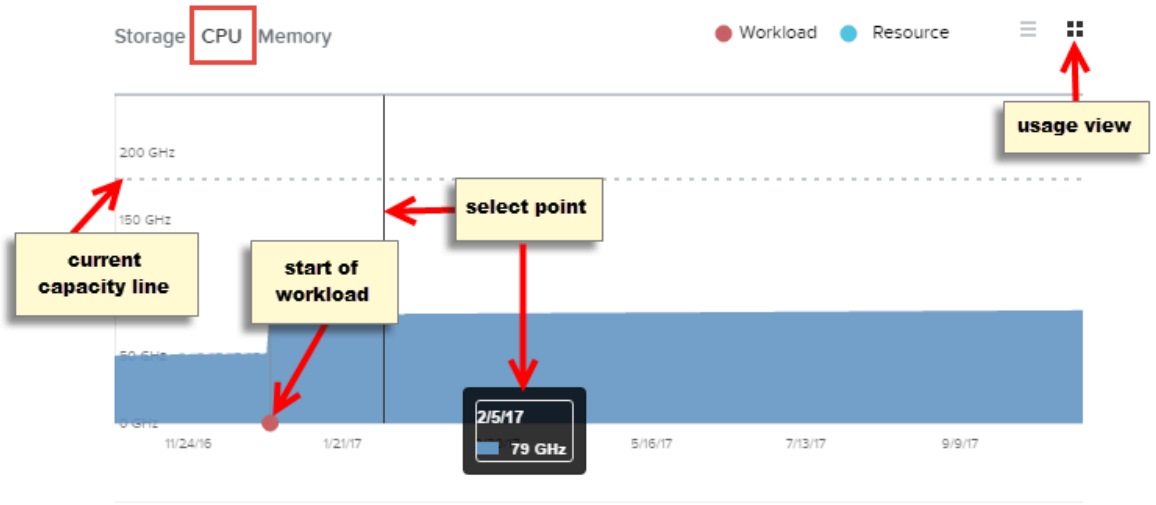

Figure 485: CPU Usage View

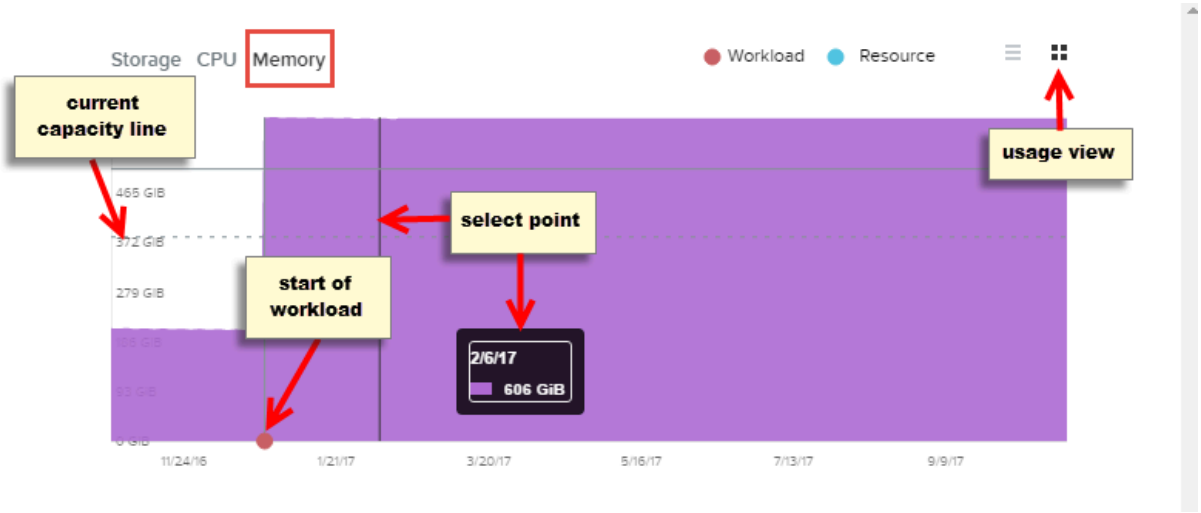

Figure 486: Memory Usage View

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 878), 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 879 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.

Sharing a Scenario

About this task

Any user with edit permission can share a scenario with other users and provide view or edit access to them.

Note: Enable CMSP to use this feature.

Procedure

1. Click the hamburger icon.
2. Go to **Operations** and click **Planning**.

   The **Planning** page with **Scenarios** and **Capacity Runway** tab appears.

3. In the **Scenarios** tab, select the scenario that you want to share and click share icon under **Actions**.

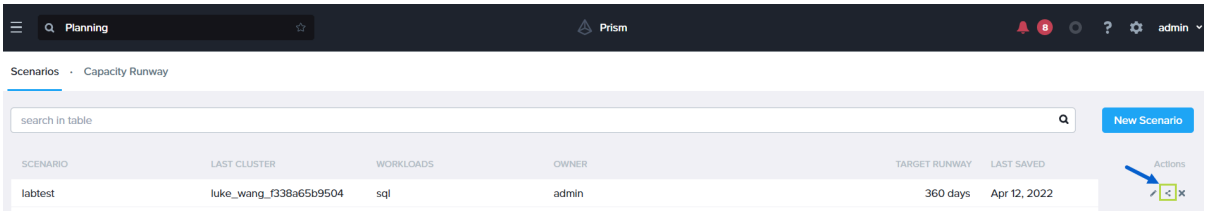

Figure 487: Scenarios Tab

The **Sharing Scenario** window appears.

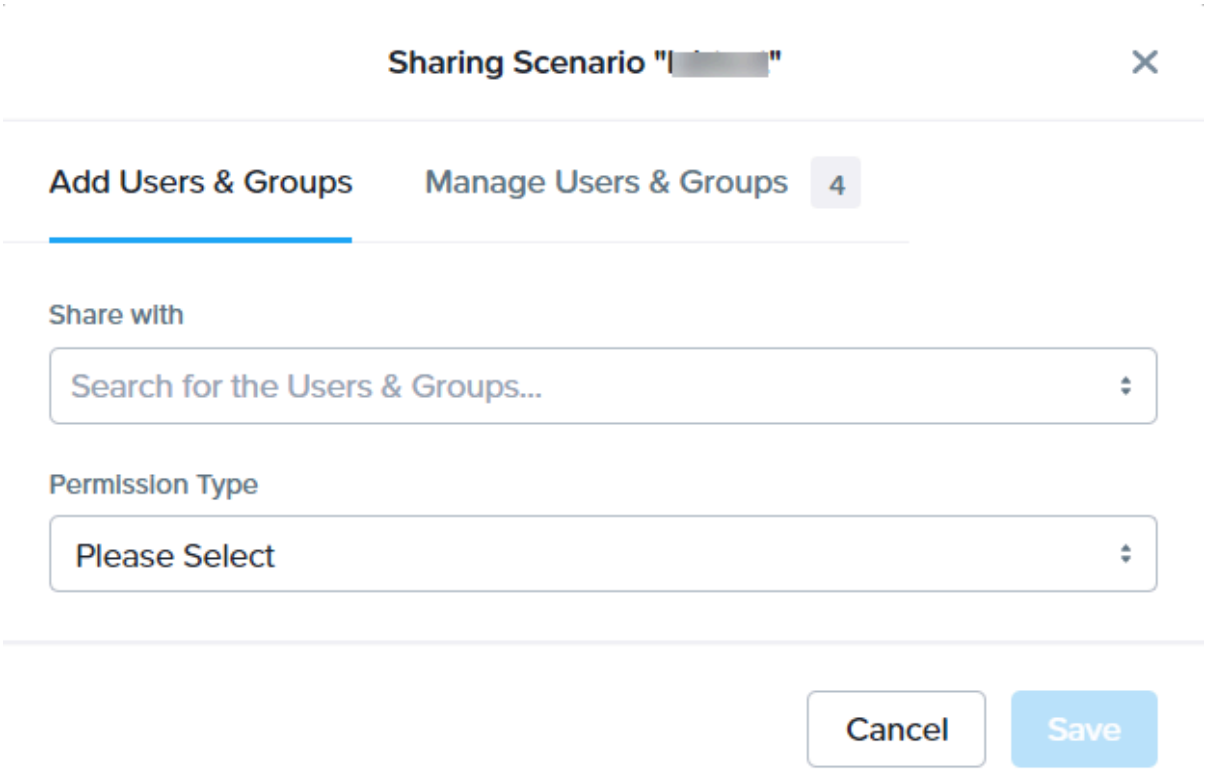

Figure 488: Sharing Scenario

4. In the **Add Users & Groups** tab, do the following:

   a. **Share with**. Add the users or groups to share the scenario. You can type first few characters of a user or group, as you type a pull-down list appears. Click the username or group name to add. You can add more than one user or a group.

   b. **Permission Type**. You can provide View only or Edit permissions to the user or group. User with Edit permission can perform these tasks: View, Edit or Share the scenario.

   The **Manage Users & Groups** tab allows you to change the permission type for users or groups with whom you have shared the scenario earlier. You can also remove the users or groups.

5. Click **Save**.

   A message that the scenario is shared successfully appears.

   After sharing the scenario, wait for few minutes for changes to take effect.

   Users will not be able to view any information about a cluster with permissions related only to scenarios and not a cluster.

Updating Capacity Configurations

About this task

The Prism Central planning tools (see Resource Planning on page 876) 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.

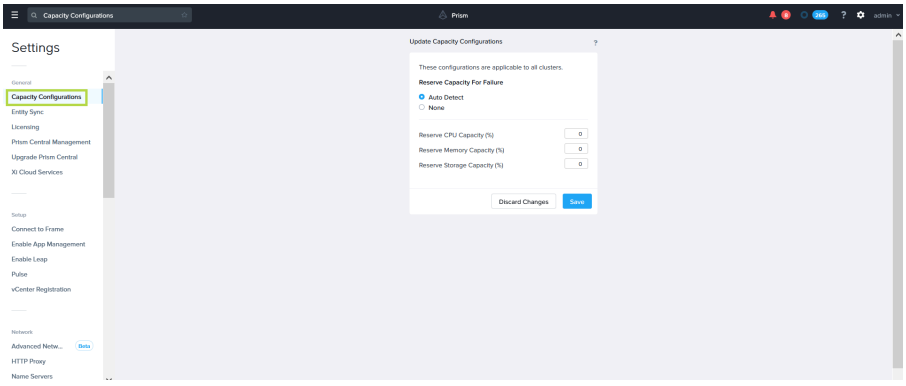

Figure 489: Update Capacity Configurations Window

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. Setting 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" means 10% 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 767: 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.

Type	Description
Over-provisioned VM	An "over-provisioned" VM is the opposite of a constrained VM, meaning it is a VM that is over-sized and wasting resources which are not needed. A VM is considered over-provisioned when it exhibits one or more of the following baseline values, based on the past 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 30 days. • *Zombie* VM: A VM is considered zombie and marked inactive if it has very few read/write IOs and minimal network traffic every day for the past 21 days. To prevent host resource wastage, delete the dead and zombie VMs.

Note: If a VM is already classified as inactive, it is not marked as over-provisioned. Inactive category takes precedence over over-provisioned category.

You can view the information about inefficient VMs from the **VM Efficiency** and **Impacted Cluster** widgets on the home or a custom dashboard (see Main Dashboard on page 70 and Adding Dashboard Widgets on page 80) and from the **Capacity** tab of a cluster details page (see Cluster Details View on page 209).

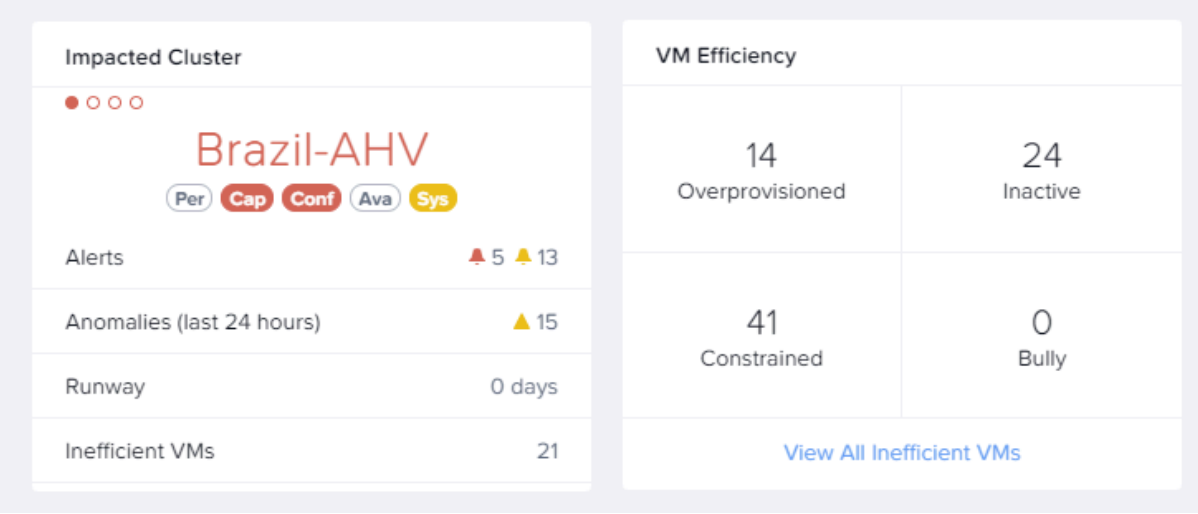

Figure 490: Impacted Cluster and VM Efficiency Widgets

You can also use a search filter from the VMs Summary View on page 94 to identify inefficient VMs.

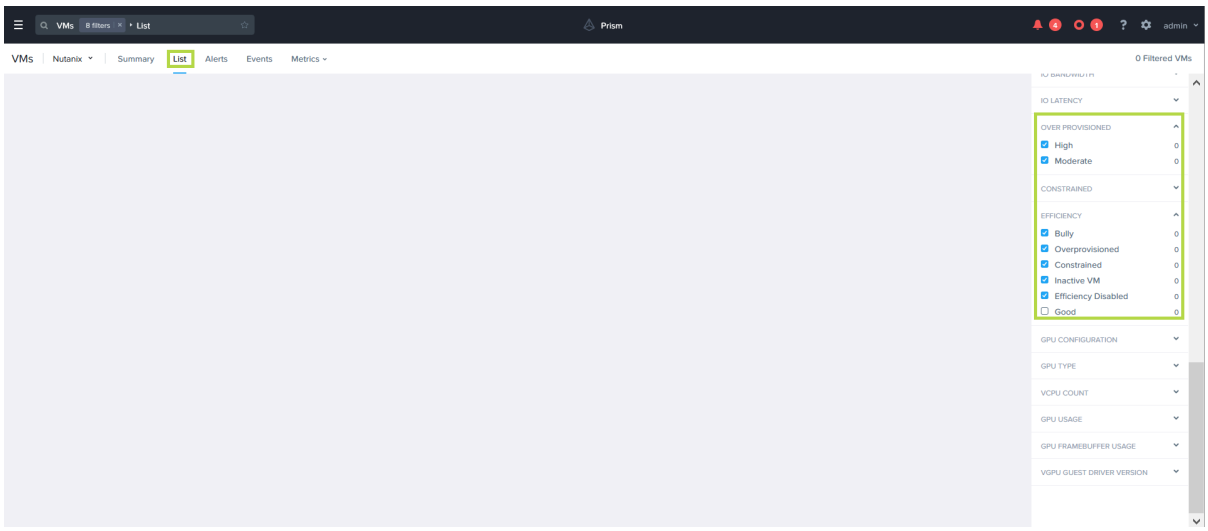

Figure 491: Filtering VMs

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 280). Anomalies appears as outliers in the behavioral anomaly event details screen and the VM details Metrics tab.

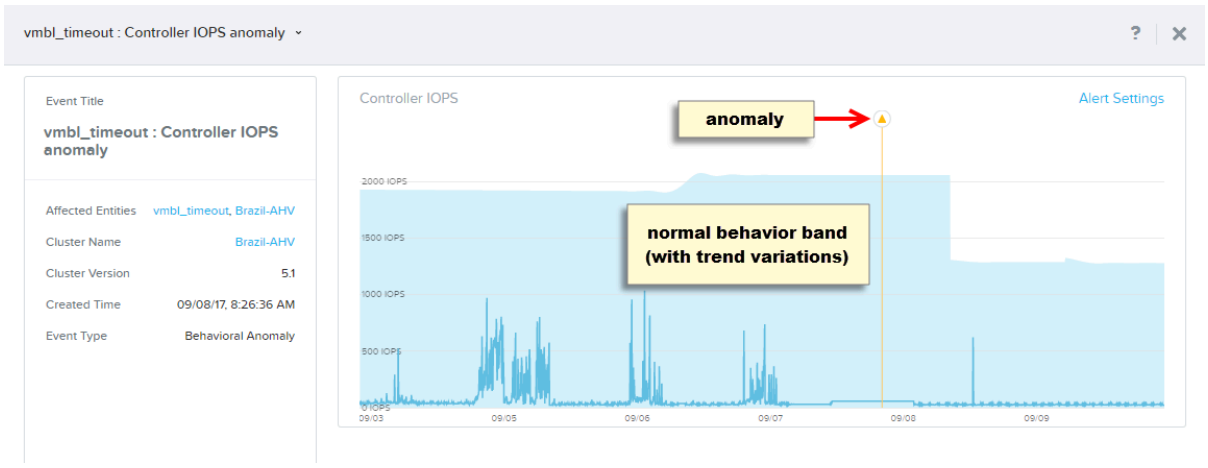

Figure 492: Anomaly Event Details Screen

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

Cluster Efficiency Reports

One of the default reports (see Reports View on page 246) 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 922).

You can view the instances of playbook execution using the Plays on page 903 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, clone, export, update, delete, enable, disable and run playbooks.
- Clone, read, update, and delete action template in the action gallery.
- View executed plays, and stop or resume a playbook that is in running state.

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.

Limitation

When Prism Central Disaster Recovery is enabled in your setup, playbooks with application passwords, authentication tokens, or secrets fail on a recovered PC. You must manually re-enter the credentials in playbooks.

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 903 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 925).

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.

To access the **Actions** menu:

1. Log in to Prism Central.
2. Click the hamburger icon from the main menu and go to **Operations** > **Playbooks**.
3. Click **Create Playbook**.

   The **Create Playbook** window opens.
4. Select a Trigger of your choice.
5. Click **Add Action**.

   Prism Central displays the Actions page. 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. For more information, see Prerequisites for VM Actions in Playbooks on page 910.

You can integrate the following applications for Playbook Actions:

- Slack - see **Slack** under *Communication Actions* in the table.
- MS Teams - see **MS Teams** under *Communication Actions* in the table.
- Ansible Tower - see **Ansible** under *Execution Actions* in the table.
- ServiceNow - see **Send Alert to ServiceNow** under *Alert Actions* in the table.
- PagerDuty - see **Send Alert to PagerDuty** under *Alert Actions* in the table.

Table 768: Playbook Actions

Action	Action Description	Configurable attributes
VM Actions		
Add to Category	Add an entity like a VM, Host or Cluster to one or more Categories.	1. Select an **Entity Type** such as VM, Host or Cluster. 2. Select the target entity in the **Target <entity>** field. 3. Select the **Category** or categories you want to add the entity to.
Power Off VM	Power off a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. 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** dropdown menu.
Remove from Category	Remove an entity like a VM, Host or Cluster from one or more Categories.	1. Select an **Entity Type** such as VM, Host or Cluster. 2. Select the target entity in the **Target <entity>** field. 3. Select the **Category** or categories you want to remove the entity from.
VM Add CPU	Add CPU for a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **vCPUs to Add** and **Maximum Number of vCPUs**. For more information, see Prerequisites for VM Actions in Playbooks on page 910.

Action	Action Description	Configurable attributes
VM Add Disk	Add vDisk to a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Select the **Bus Type** and enter the **Disk Size** in GiB.
VM Add Memory	Add memory for a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **Add Absolute Memory** and **Absolute Maximum**.
VM Expand Disk	Expand vDisk for a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **Expand Disk By** in GiB. 3. Enter the **Maximum Disk Size** in GiB. 4. Choose from: • **Specific Disk** - Select the **Bus Type** and enter the **Index** value for the selected disk type. • **Most Used Disk**
VM Recovery Point	Captures the state of a VM and saves it.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **Time To Live** in number of days.
VM Reduce CPU	Reduce CPU for a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **vCPUs to Remove** and **Minimum Number of vCPUs**.
VM Reduce Memory	Reduce memory for a VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the value for **Reduce Memory by** and **Minimum Limit**.
Communication Actions		

Action	Action Description	Configurable attributes
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. • This action requires the SMTP server to be configured. For more information, see Configuring an SMTP Server in the *Prism Web Console Guide*.
Microsoft Teams	Send a message to Microsoft Teams	Enter the **Webhook URL** that you configured in Microsoft Teams and the **Message**. For more information, see Integrating MS Teams Action in Playbooks on page 913.
Slack	Send a Slack message.	Enter the **Slack Token**, **Channel Name**, and **Message**. For more information, see Integrating Slack Action in Playbooks on page 912.
Alert Actions		
Acknowledge Alert	Mark an Alert as acknowledged.	Select the target Alert from the **Target Alert** dropdown menu.
Resolve Alert	Mark an Alert as resolved.	Select the alert from the **Target Alert** dropdown menu.
Send Alert to PagerDuty	Send an Alert to PagerDuty.	Enter the **Routing Key**. The Routing Key is an integration key that you find in your PagerDuty instance. For more information, see Configuring a Playbook with the PagerDuty Integration on page 939.
Send Alert 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.

Action	Action Description	Configurable attributes
	• Mid Server	Enter the **Mid Server Address**, **Mid Server Port**, **Username**, and **Password**.
Send Alert to Webhook	Send a JSON Alert Object to a REST API endpoint.	Enter the **Alert**, **Method**, and **URL** for the outgoing webhook. Optionally, select the **Username**, **Password**, and **Request Headers**.
Report Actions		
Generate Forecast Report	Generate a just in time forecast report.	Select the cluster for which you want to generate the report and enter the days for the runway period.
Generate Report	Generate a custom report.	1. Select the Report Type as **New Report**. You can also select any of previously created reports to be generated. 2. Enter the **Report Instance Name**. 3. Optionally, enter a **Description**. 4. 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		
Ansible	Run an Ansible **Workflow** or **Job** template.	Select the **Template Type**, and enter the Ansible **Tower/ AWX Hostname** and OAuth authentication **Token** to log in to the Tower or AWX host. For more information, see Integrating Ansible Action in Playbooks on page 915.

Action	Action Description	Configurable attributes
IP Address PowerShell	Remotely run a PowerShell script for a given VM IP address.	1. Enter the **VM IP Address**, **Username**, **Password**, and the **Path to Script**. 2. 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.
Puppet	Run a Puppet Task or Plan.	1. Enter **Puppet Enterprise HostName**. Fully qualified domain name of a host where the Puppet server is running. 2. The default **Puppet Enterprise Port** is 8143. The port number in which Puppet accepts connections. You can change the default port number. 3. Enter **Authorization Token**. The authorization token provided by the Puppet server. 4. Select **Command Type** Task or Plan. 5. Enter **Task** or **Plan name** depending on the Command Type. 6. Enter **Scope**. This is the JSON payload that determines set of nodes on which this task runs. Optionally, add other parameters shown in the screen.
REST API	Make a call to a REST API endpoint.	1. Select the **Method**. 2. Enter the **URL**, **Username**, **Password**, **Request Body**, and **Request Headers**.

Action	Action Description	Configurable attributes
VM Powershell	Remotely execute PowerShell script for a given VM.	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **Username**, **Password**, and the **Path to Script**. 3. Enable or disable **HTTPS**.
VM SSH	Remotely execute an ssh command for a given VM	1. Select the VM from the **Target VM** dropdown menu. 2. Enter the **Username**, **Password**, and the **Command to Run**.
Utility Actions		
Lookup Cluster Details	Captures cluster details	Select the cluster from the **Target Cluster** dropdown menu.
Lookup Host Details	Captures host details	Select the host from the **Target Host** dropdown menu.
Lookup VM Details	Captures VM details	Select the VM from the **Target VM** dropdown menu.

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

Branch Actions		
Branch	Switch execution based on a matching criteria.	For more information, see Using Branch Action (Conditional Execution) on page 930.

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 769: VM Action Prerequisites

Action	Hypervisor	Prerequisites	Notes
VM Add CPU (Hot Add)	AHV	AOS 5.6 and above	The VM has a guest operating system that supports CPU hot add functionality.
	ESXi	None	The VM has a guest operating system that supports CPU hot add functionality.
VM Reduce CPU	AHV	VM must be powered off	
	ESXi	VM must be powered off	
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 Action in Playbooks

You can create a Playbook to send a Slack message. The Playbook sends the message from Prism Central to the defined channels in a Slack instance.

Before you begin

You must integrate Slack to automate receiving messages in Slack. To do this, you must create a new app in Slack, set the necessary app permissions, install the app in your workspace, and retrieve the Slack authorization token. For more information, see *Basic app setup* in Slack documentation.

About this task

Perform the following steps to create a Playbook to send a Slack message.

Procedure

1. Log in to Prism Central.
2. Click the hamburger icon from the main menu and go to **Operations** > **Playbooks**.
3. Click **Create Playbook**.

   > **Note:** If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

   The **Select a Trigger** page appears.
4. Select a trigger of your choice.
5. Click **Add Actions**.
   A list of all the available actions on the deck appears.
6. Click **Select** within the **Slack** (option) card.
7. Enter the **Slack Token**.

   The Slack Token is an OAuth access token that you can find in your Slack instance. For more information, see Slack documentation.
8. Enter the **Channel Name**.

   Type the channel name that should receive the slack message that you describe in Step 9.
9. Type a **Message**.
10. (Optional) 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.
11. 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**.

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.

Before you begin

You must configure a webhook for your Microsoft Teams channel to allow Prism Central to share content in the Microsoft Teams channel. For more information, see *Create Incoming Webhooks* in Microsoft documentation.

About this task

Perform the following steps to create a Playbook to send a message to a Microsoft Teams channel.

Procedure

1. Log in to Prism Central.
2. Click the hamburger icon from the main menu and go to **Operations** > **Playbooks**.
3. Click **Create Playbook**.

   > **Note:** If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

   The **Select a Trigger** page appears.
4. Select a trigger of your choice.
5. Click **Add Actions**.
   A list of all the available actions on the deck appears.
6. Click **Select** within the **Microsoft Teams** (option) card.
7. Enter the **Webhook URL**.

   The Webhook URL is a unique URL found in your MS Teams instance. You can use it to send information to Microsoft Teams. For more information, see *Create Incoming Webhooks* in Microsoft documentation.
8. Type a **Message**.
9. (Optional) 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.
10. 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**.

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.

- For 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.

- For information about how to create an application in Ansible Tower, see *Applications* information in Ansible documentation.
- For information about how to create application token in Ansible Tower, see *Users - Tokens* information in Ansible documentation.

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/

Integrating Ansible Action in Playbooks

The Ansible action uses a token-based authentication to integrate with Ansible Tower.

Before you begin

You must create a job template, a workflow template, an application, and an application token in Ansible Tower to integrate Prism Central with Ansible Tower. For more information, see Ansible documentation.

About this task

Perform the following steps to create a Playbook to trigger job or workflow templates in your Ansible Tower or AWX instance.

Procedure

1. Log in to Prism Central.
2. Click the hamburger icon from the main menu and go to **Operations** > **Playbooks**.
3. Click **Create Playbook**.

   Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

   The **Select a Trigger** page appears.
4. Select a trigger of your choice.
5. Click **Add Actions**.
   A list of all the available actions on the deck appears.
6. Click **Select** within the **Ansible** (option) card.
7. Select **Template Type**.

8. Enter the appropriate details into the following fields:

   - **Workflow Template Name** or **Job Template Name** (as per the selected template type): Enter the name for the workflow template.
   - **Tower/AWX Hostname**: Enter the name of the Ansible tower or the AWX host.
   - **Token**: Enter the OAuth authentication token created in the Ansible Tower.

     This token is required to log in to the Tower or AWX host. For information, see *Users - Tokens* information in Ansible documentation.
   - (Optional) **Request Body JSON (Optional)**: Enter any other parameters that you want to use in JSON format.

9. (Optional) 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.

10. (Optional) Turn on the toggle for **Wait for job completion?** if you want Prism Central to wait for the current action to complete before beginning the next action.

11. 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**.

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. For information about how to enable the plugin for your production instance, see ServiceNow documentation.

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

For information about how to enable X-Play support on the ServiceNow developer instance, see *Activate a plugin on a personal developer instance* in ServiceNow documentation.

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 923.
   - **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 925.
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.

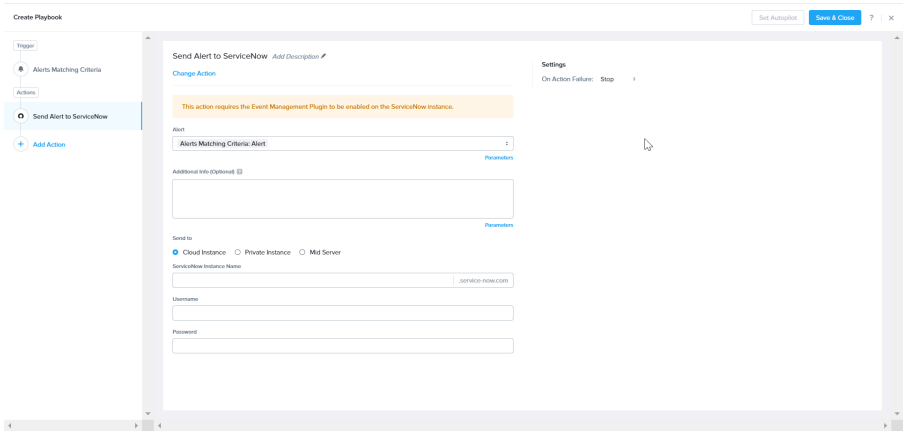

Figure 493: ServiceNow Action

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 919 and Activating the Create Incident Rule on page 919.

You can also create playbooks for actions that you want ServiceNow to initiate on Prism Central for the alerts sent to Service now.

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 Nutanix XML download link.

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.

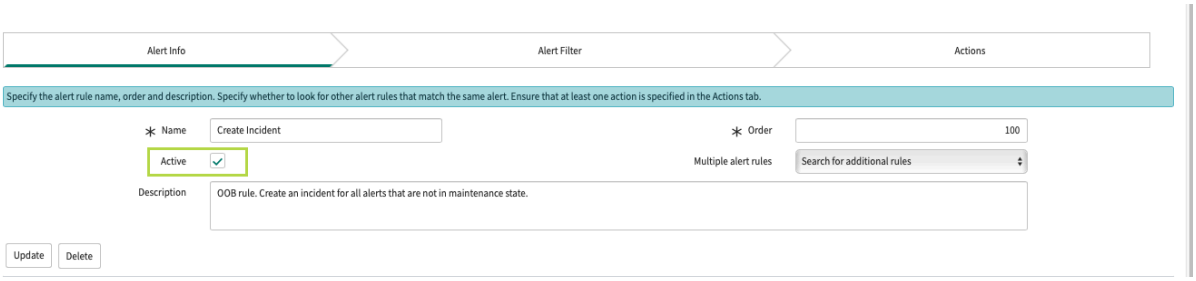

Figure 494: Activate the Create Incident Rule

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. Log in to Prism Central.
2. Click the hamburger icon from the main menu and go to **Operations** > **Playbooks**.

3. Click **Create Playbook**, then click **Select** in the **Webhook** (option) card.

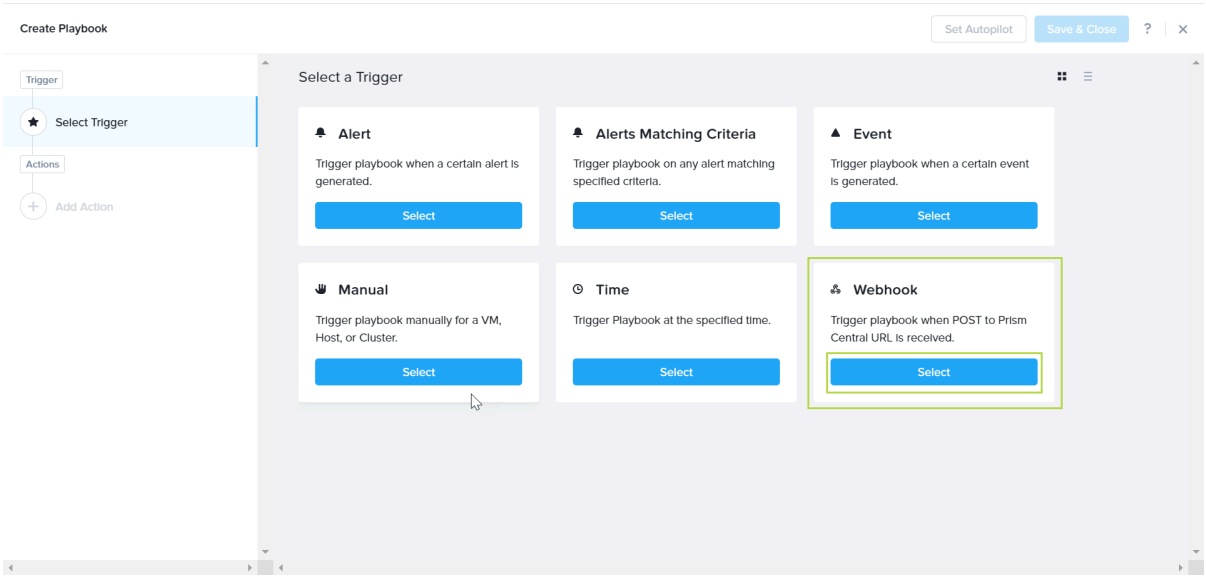

Figure 495: Select a Trigger

The **Webhook** page displays a message informing you to return to this playbook and review the instructions configured for this trigger.

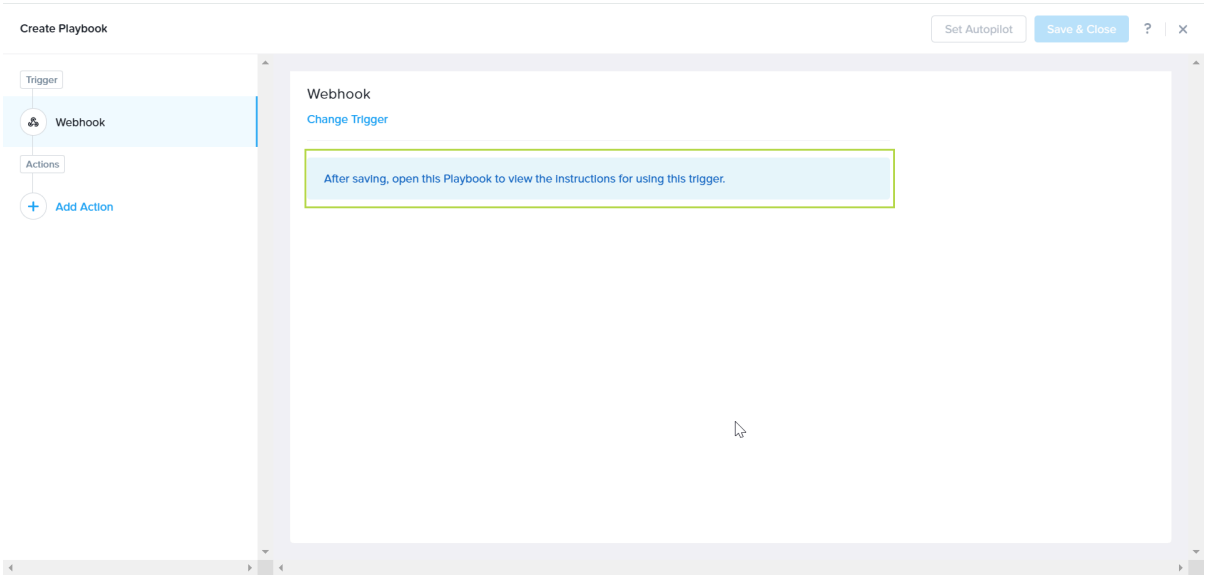

Figure 496: Webhook page

4. 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.

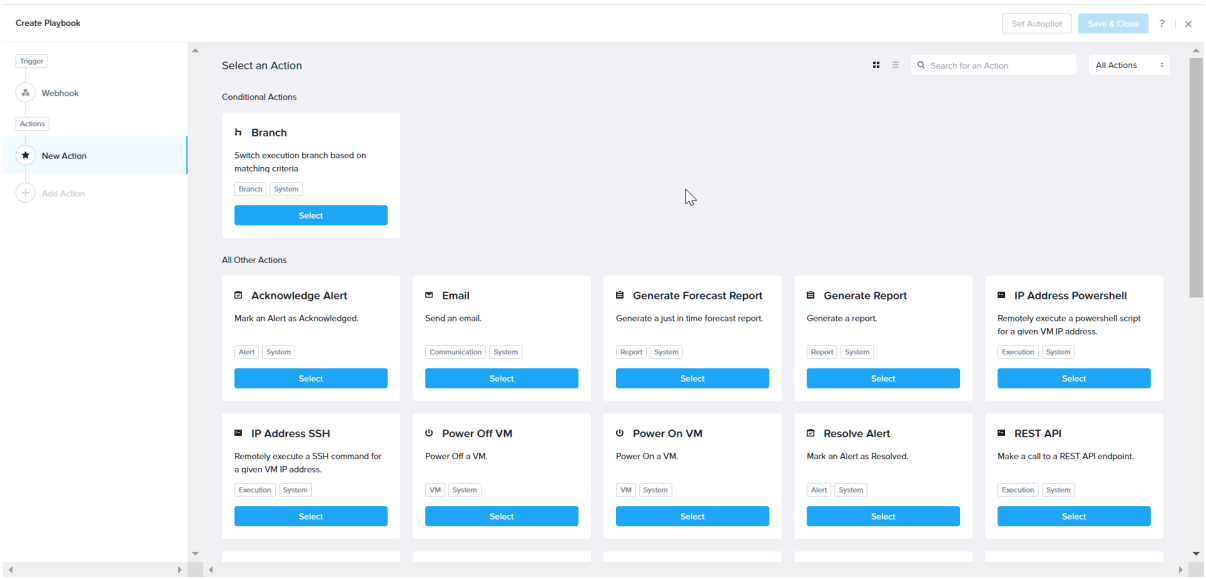

Figure 497: Select an Action page

   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.

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**.

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 925.

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.
   See Alerts Summary View (Prism Central) on page 269 for details on alert messages.
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.

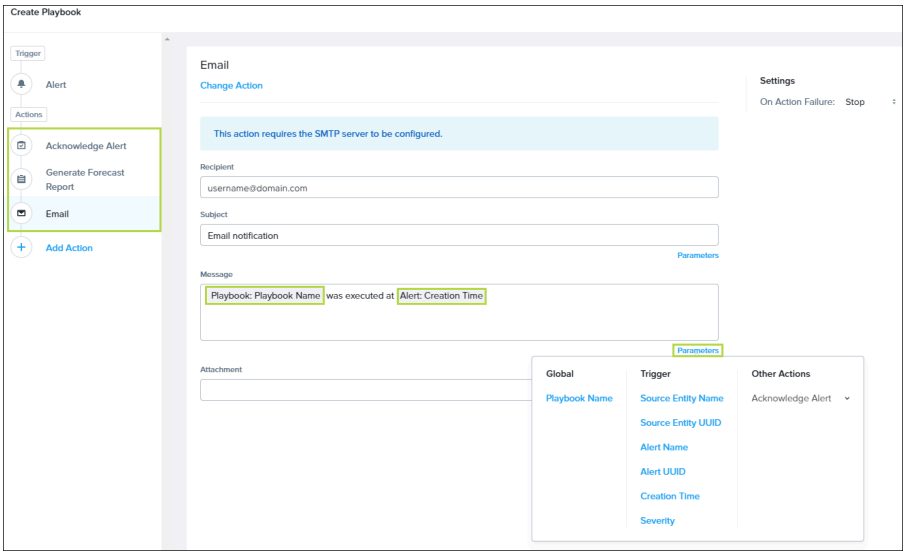

Figure 498: Parameters

   To configure an action, set the action attributes based on your requirement. Refer to the Playbook Actions on page 903 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 925.

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 met.

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 903 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.
   See Events Summary View (Prism Central) on page 279 for details on event messages.

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.

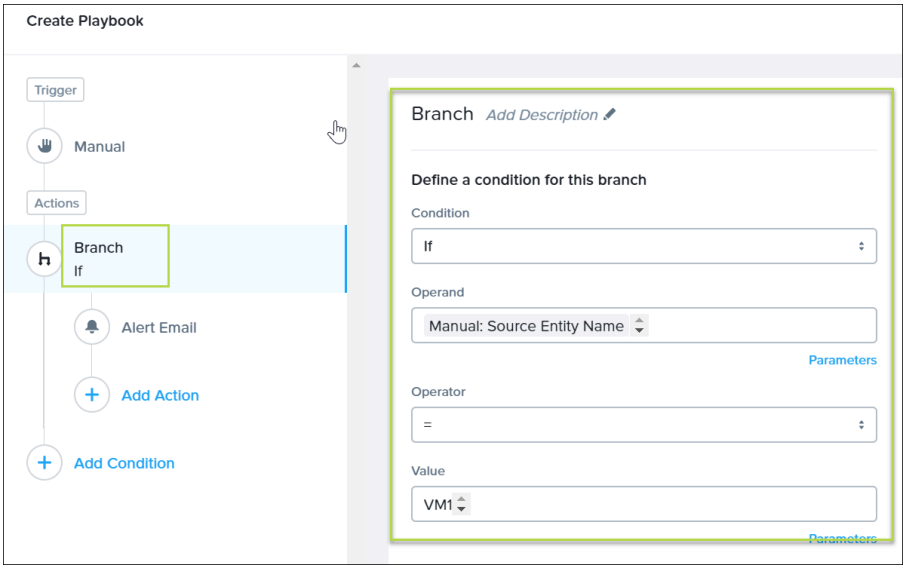

Figure 499: Branch Action

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 903 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 934 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.

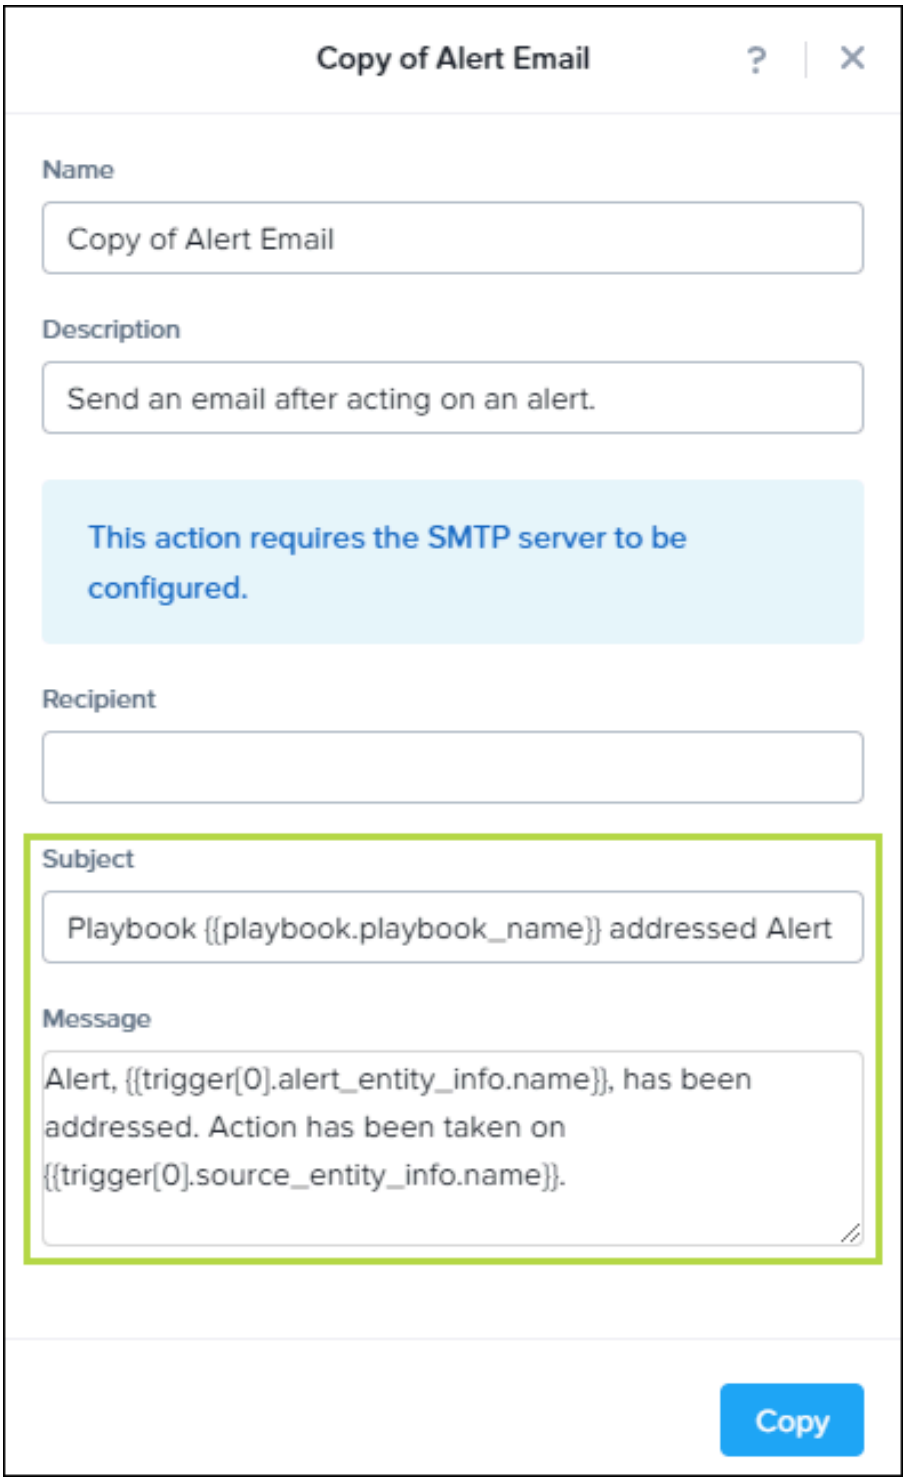

Figure 500: Configuring parameters manually during Clone operation

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 934 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 output parameters available for Playbook triggers.

Table 770: Triggers and Parameters

Trigger	Output Parameter	API Parameter
Alert	Alert (object)	{{trigger[0].alert_entity_info}}
	Source Entity (object)	{{trigger[0].source_entity_info}}
	Severity (string)	{{trigger[0].severity}}
	Creation Time (long)	{{trigger[0].creation_time}}
	Description (string)	{{trigger[0].description}}
Alerts Matching Criteria	Alert (object)	{{trigger[0].alert_entity_info}}
	Source Entity (object)	{{trigger[0].source_entity_info}}
	Severity (string)	{{trigger[0].severity}}
	Source Cluster (object)	{{trigger[0].cluster_entity_info}}
	Impact Type (string)	{{trigger[0].impact_type}}
	Creation Time (long)	{{trigger[0].creation_time}}
	Description (string)	{{trigger[0].description}}
Event	Event (object)	{{trigger[0].audit_entity_info}}
	Source Entity (object)	{{trigger[0].source_entity_info}}
	Creation Time (long)	{{trigger[0].creation_time}}
	Username (string)	{{trigger[0].user_name}}
	Client IP (string)	{{trigger[0].client_ip}}
	Source Cluster (object)	{{trigger[0].cluster_name}}
Manual	Source Entity (object)	{{trigger[0].source_entity_info}}
	Playbook UUID (string)	{{trigger[0].action_rule_uuid}}
Time	Trigger Time (string)	{{trigger[0].timestamp}}

Trigger	Output Parameter	API Parameter
	Playbook (object)	{{trigger[0].rule_entity}}
Webhook	Webhook Id (string)	{{trigger[0].webhook_id}}

Refer to the following table for the complete list of output parameters available for Playbook actions.

Table 771: Actions and Parameters

Action	Output Parameter	API Parameter
Branch Actions		
Branch	Selected Branch (string)	{{action[index].selected_branch}}
VM Actions		
Add to Category	Task (object)	{{action[index].task_entity}}
Power Off VM	Task (object)	{{action[index].task_entity}}
Power On VM	Task (object)	{{action[index].task_entity}}
Remove from Category	Task (object)	{{action[index].task_entity}}
VM Add CPU	Task (object)	{{action[index].task_entity}}
	Resultant vCPU count (long)	{{action[index].result_vcpu_count}}
VM Add Disk	Added Disk (string)	{{action[index].disk_name}}
	Task (object)	{{action[index].task_entity}}
VM Add Memory	Task (object)	{{action[index].task_entity}}
	Resultant Memory (size)	{{action[index].total_memory}}
VM Expand Disk	Expanded Disk (object)	{{action[index].disk}}
	Resultant Disk Size (size)	{{action[index].resultant_disk_size}}
VM Recovery Point	Recovery Point (object)	{{action[index].vm_recovery_point_entity}}
	Task (object)	{{action[index].task_entity}}
VM Reduce CPU	Task (object)	{{action[index].task_entity}}
	Resultant vCPU count (long)	{{action[index].result_vcpu_count}}
VM Reduce Memory	Task (object)	{{action[index].task_entity}}
	Resultant Memory (size)	{{action[index].total_memory}}
Communication Actions		
Email	None	None
Microsoft Teams	None	None
Slack	Channel Id (string)	{{action[index].channel_id}}
	Message Timestamp (string)	{{action[index].timestamp}}

Action	Output Parameter	API Parameter
Alert Actions		
Acknowledge Alert	Task (object)	{{action[index].task_entity}}
Resolve Alert	Task (object)	{{action[index].task_entity}}
Send Alert to PagerDuty	None	None
Send Alert to ServiceNow	None	None
Send Alert to Webhook	Response Status Code (long)	{{action[index].status_code}}
	Response Headers (string)	{{action[index].response_headers}}
	Response Body (string)	{{action[index].response_body}}
Report Actions		
Generate Forecast Report	Report PDF (file)	{{action[index].report}}
Generate Report	Report (file)	{{action[index].report}}
	Task (object)	{{action[index].task_entity}}
Wait Actions		
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}}
Execution Actions		
Ansible	Job ID (string)	{{action[index].job_id}}
IP Address Powershell	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
IP Address SSH	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
Puppet	Job Name (string)	{{action[index].job_name}}
REST API	Response Status Code (long)	{{action[index].status_code}}
	Response Headers (string)	{{action[index].response_headers}}
	Response Body (string)	{{action[index].response_body}}
VM Powershell	Standard Output (string)	{{action[index].stdout}}

Action	Output Parameter	API Parameter
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
VM SSH	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
Utility Actions		
Lookup Cluster Details	Cluster Name (string)	{{action[index].cluster_name}}
	Cluster UUID (string)	{{action[index].cluster_uuid}}
	Cluster VIP (string)	{{action[index].cluster_vip}}
	CVM IP List (string)	{{action[index].cvm_ip_list}}
	Storage Capacity (size)	{{action[index].storage_capacity}}
	CPU Count (long)	{{action[index].cpu_counts}}
	Memory Size (size)	{{action[index].memory}}
	Categories (object)	{{action[index].categories}}
Lookup Host Details	Host Name (string)	{{action[index].host_name}}
	Host UUID (string)	{{action[index].host_uuid}}
	Host IP (string)	{{action[index].host_ip}}
	VM Count (long)	{{action[index].vm_count}}
	CPU Count (long)	{{action[index].cpu_count}}
	Memory Size (size)	{{action[index].memory}}
	Disk Capacity (size)	{{action[index].disk_capacity}}
	Hosting Cluster VIP (string)	{{action[index].hosting_cluster_vip}}
	Hosting CVM IP (string)	{{action[index].hosting_cvm_ip}}
	Hosting Maintenance Mode Status (string)	{{action[index].host_maintenance_mode_status}
	Categories (object)	{{action[index].categories}}
Lookup VM Details	VM Name (string)	{{action[index].vm_name}}
	VM UUID (string)	{{action[index].vm_uuid}}
	VM IP Address (string)	{{action[index].vm_ip}}
	Host IP (string)	{{action[index].host_ip}}
	VM Hypervisor Type (string)	{{action[index].hypervisor_type}}
	Cores per vCPU (long)	{{action[index].vm_cores_per_vcpu}}
	vCPU Count (long)	{{action[index].vcpus}}

Action	Output Parameter	API Parameter
	Memory Size (size)	{{action[index].memory_size}}
	Disk Capacity (size)	{{action[index].disk_capacity}}
	Power State (string)	{{action[index].power_state}}
	Hosting Cluster VIP (string)	{{action[index].hosting_cluster_vip}}
	Hosting CVM IP (string)	{{action[index].hosting_cvm_ip}}
	VM Owner Name (string)	{{action[index].owner_name}}
	Categories (object)	{{action[index].categories}}
	Number of Cores (long)	{{action[index].total_cores}}
String Parser	Parsed String (string)	{{action[index].parsed_data}}

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.

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.

Before you begin

You must set up the PagerDuty integration key. This key is used as the routing key to integrate PagerDuty with Playbooks in Prism Central. For information about how to generate an integration key, see *Create a Generic Events API Integration* in PagerDuty documentation.

About this task

Perform the following steps to create a Playbook to send Prism alerts to PagerDuty.

Procedure

1. Log in to Prism Central.
2. Click the hamburger icon and go to **Operations** > **Playbooks**.
3. Click **Create Playbook**.
   The **Select a Trigger** page appears.
4. Select **Alerts Matching Criteria** as a trigger.

   The **Alerts Matching Criteria** allows you to choose all or specific alert policies.
5. 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.

6. 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

7. 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.

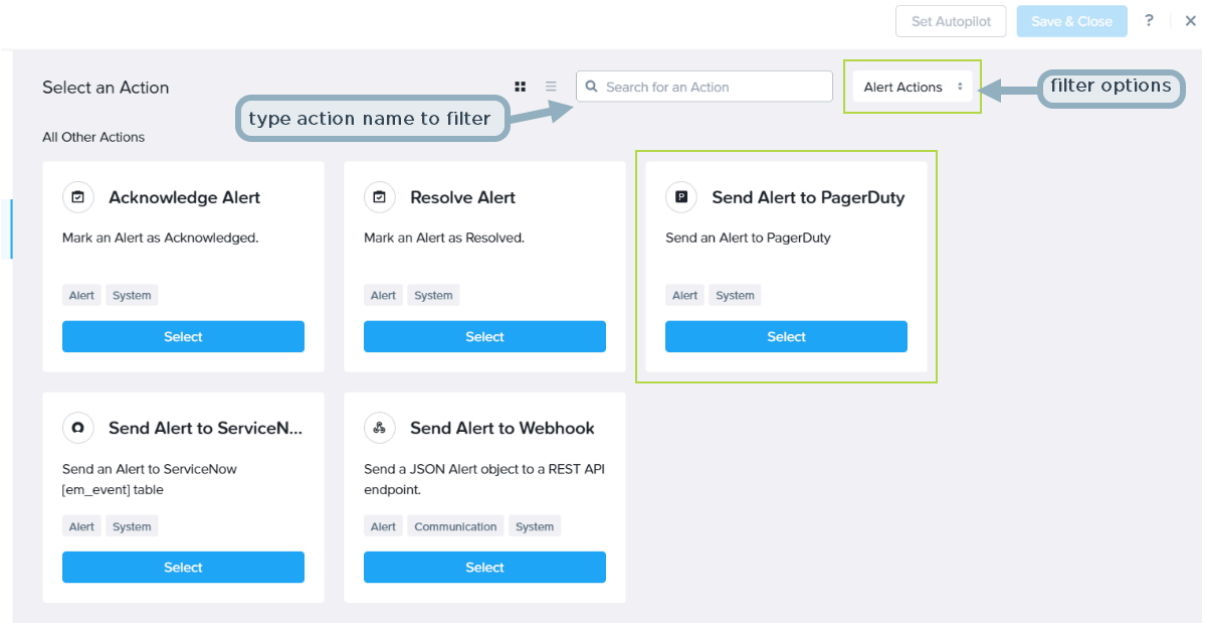

Figure 501: Send Alert to PagerDuty

8. 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.

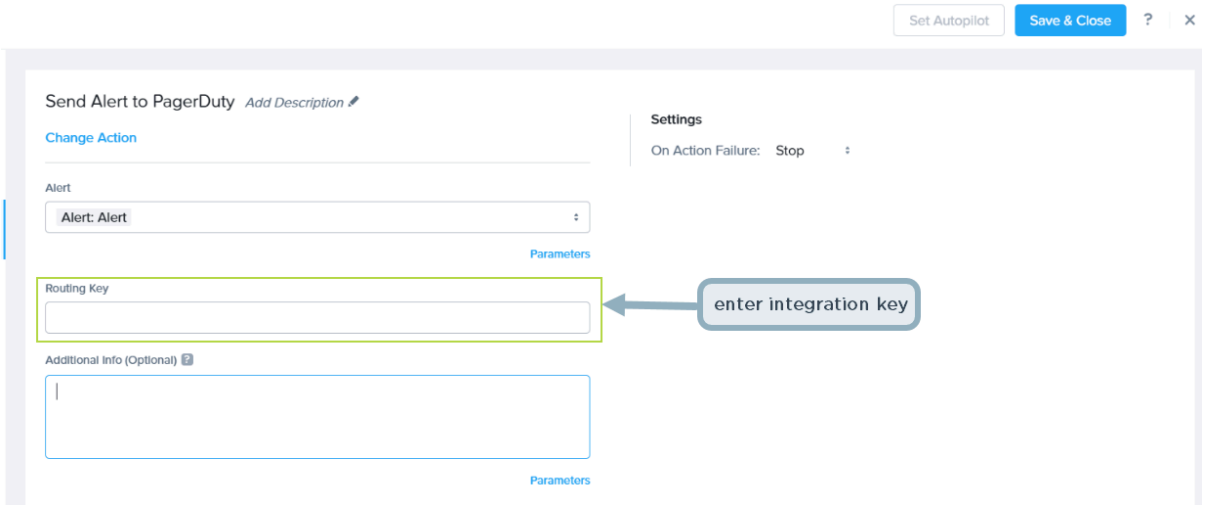

Figure 502: Routing Key

9. 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.

10. 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:**
>
> - All the built-in roles except the **Prism Viewer** role can manage the reports (create, update, share, and delete). Users with **Prism Viewer** privileges can only view the reports.
> - 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 943.
> - When cluster-based licensing is in use, data from an unlicensed cluster (where cluster-based license is not applied) is filtered out from the report. For more information, refer to the License Manager Guide.

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 share 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 246), 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 959 .
3. Configure a schedule for the report. For more information about scheduling, see Scheduling a Report on page 986.
4. Configure the settings specific for this report. For more information about configuring report settings, see Configuring Report Settings on page 988 .
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 944 .

   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).

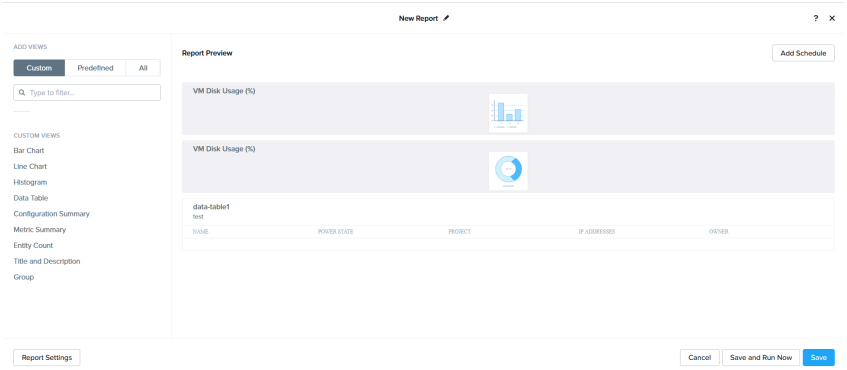

Figure 503: Create Report

Managing a Report

After you create a report, you can generate the report, edit the report, view instances of the report, clone the report, delete the report, or share the report.

- To generate a report, see Generating a Report on page 944.
- To edit a report, see Editing a Custom Report on page 945.
- To clone a report, see Cloning a Report on page 945.
- To delete a report, see Deleting a Custom Report on page 946.
- To share a report, see Sharing a Report on page 946.

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 246), 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 988 .

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 246), 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 943 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 979.

     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 246), 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 943 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 246)), 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.

Sharing a Report

About this task

Any user with edit permission can share a report with other users and provide view or edit access to them.

> **Note:** Enable CMSP to use this feature.

Procedure

1. In the **Reports** dashboard, select a report, go to **Actions** and click **Share**.
   The **Sharing Report** window appears.

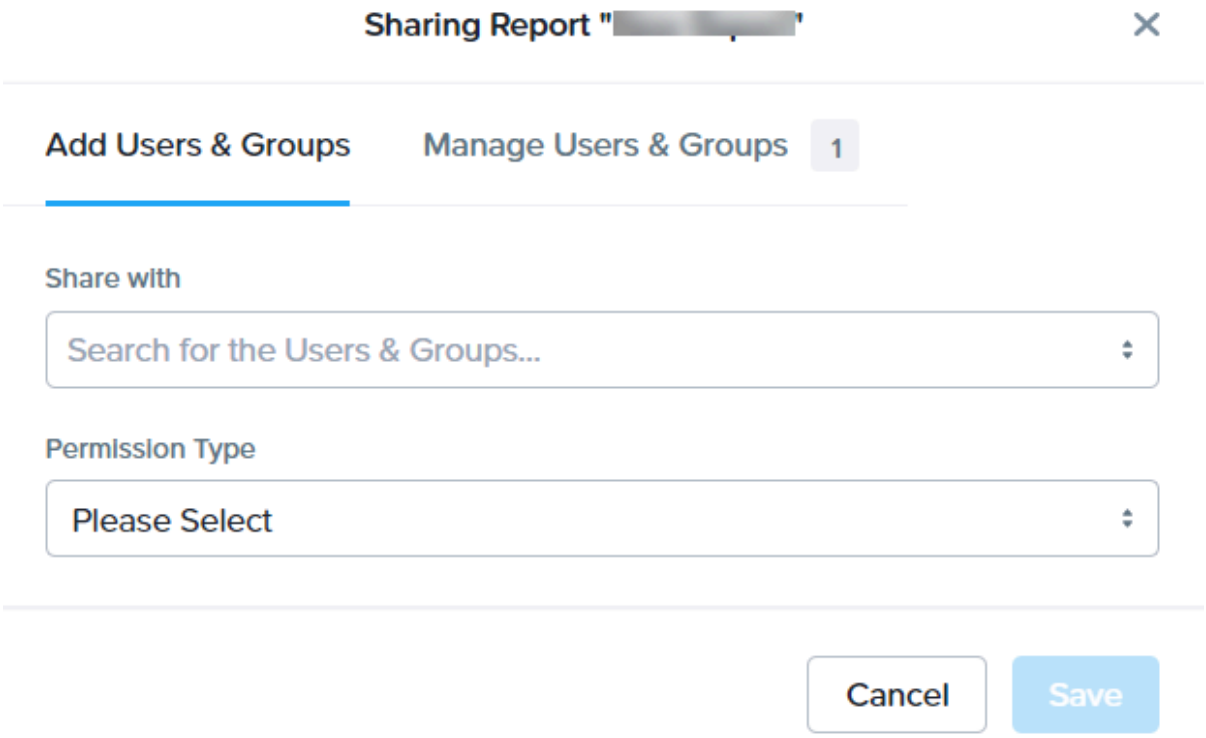

Figure 504: Sharing Report

2. In the **Add Users & Groups** tab, do the following:

   a. **Share with**. Add the users or groups to share the report. You can type first few characters of a user or group, as you type a pull-down list appears. Click the username or group name to add. You can add more than one user or a group.

   b. **Permission Type**. You can provide View only or Edit permissions to the user or group. User with Edit permission can perform these tasks: View, Edit, Run or Share the report.

   The **Manage Users & Groups** tab allows you to change the permission type for users or groups with whom you have shared the report earlier. You can also remove the users or groups.

3. Click **Save**.
   A message that the report is shared successfully appears.

   After sharing the report, wait for few minutes for the changes to take effect.

   Users will not be able to view any information about entities in the reports without entity permissions.

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 246), 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.

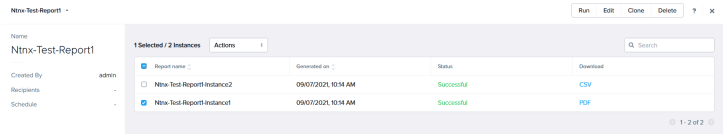

Figure 505: Viewing Instance Information (Nutanix)

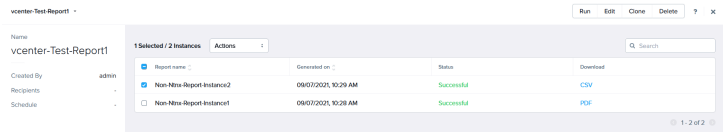

Figure 506: Viewing Instance Information (vCenter)

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.

- 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 507: 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.

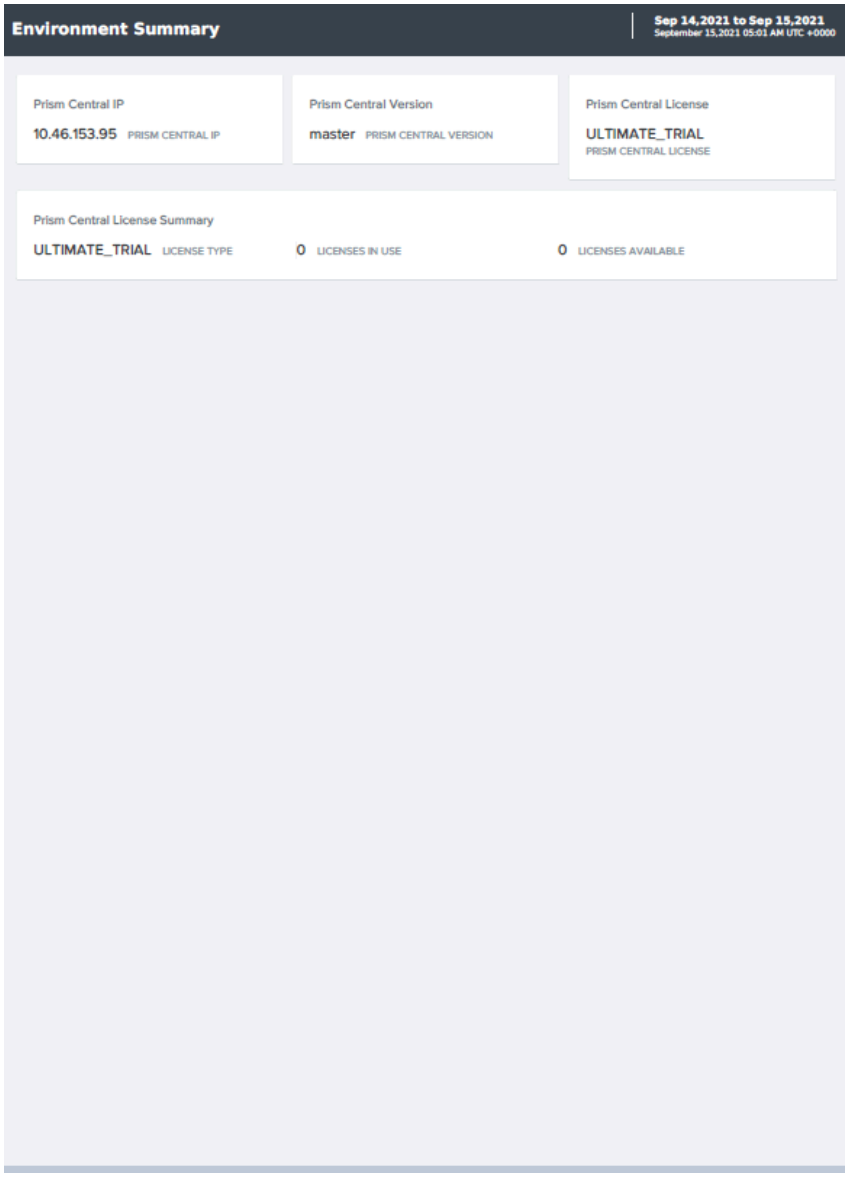

Figure 508: Detailed Metrics Information (Nutanix: Image 1)

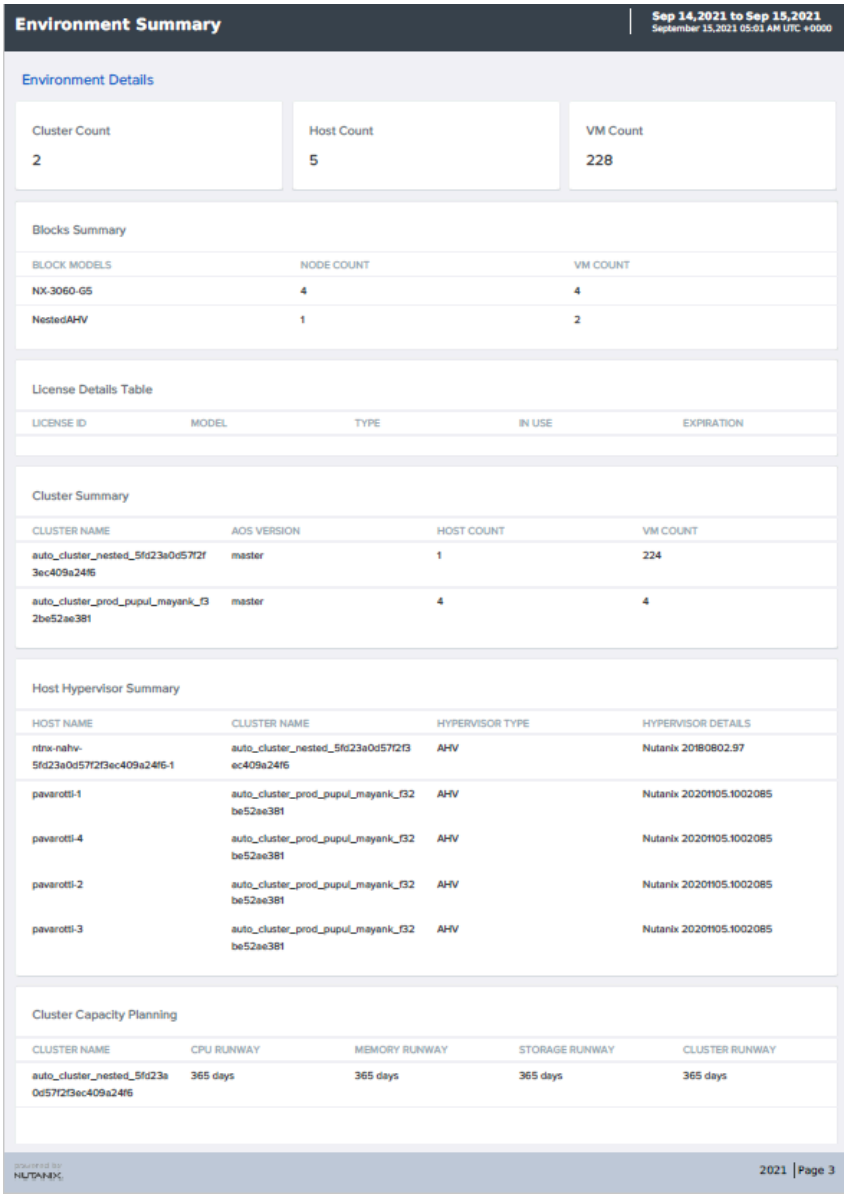

Environment Summary

Sep 14,2021 to Sep 15,2021
September 15,2021 05:01 AM UTC +0000

Environment Details

Cluster Count	Host Count	VM Count
2	5	228

Blocks Summary

BLOCK MODELS	NODE COUNT	VM COUNT
NX-3060-G5	4	4
NestedAHV	1	2

License Details Table

LICENSE ID	MODEL	TYPE	IN USE	EXPIRATION

Cluster Summary

CLUSTER NAME	AOS VERSION	HOST COUNT	VM COUNT
auto_cluster_nested_5fd23a0d57f2f3ec409a24f6	master	1	224
auto_cluster_prod_pupul_mayank_f32be52ae381	master	4	4

Host Hypervisor Summary

HOST NAME	CLUSTER NAME	HYPERVISOR TYPE	HYPERVISOR DETAILS
ntnx-nahv-5fd23a0d57f2f3ec409a24f6-1	auto_cluster_nested_5fd23a0d57f2f3ec409a24f6	AHV	Nutanix 20180802.97
pavarotti-1	auto_cluster_prod_pupul_mayank_f32be52ae381	AHV	Nutanix 20201105.1002085
pavarotti-4	auto_cluster_prod_pupul_mayank_f32be52ae381	AHV	Nutanix 20201105.1002085
pavarotti-2	auto_cluster_prod_pupul_mayank_f32be52ae381	AHV	Nutanix 20201105.1002085
pavarotti-3	auto_cluster_prod_pupul_mayank_f32be52ae381	AHV	Nutanix 20201105.1002085

Cluster Capacity Planning

CLUSTER NAME	CPU RUNWAY	MEMORY RUNWAY	STORAGE RUNWAY	CLUSTER RUNWAY
auto_cluster_nested_5fd23a0d57f2f3ec409a24f6	365 days	365 days	365 days	365 days

NUTANIX

2021 | Page 3

Figure 509: Detailed Metrics Information (Nutanix: Image 2)

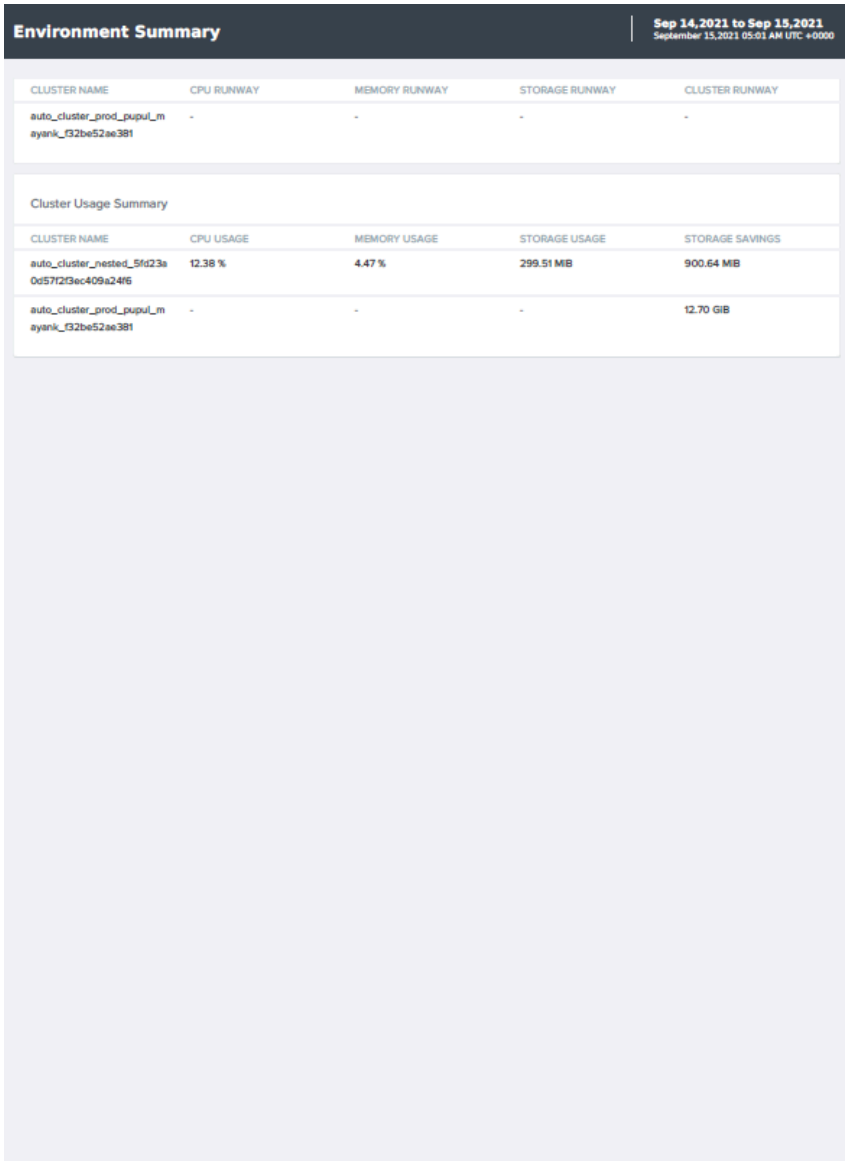

Figure 510: Detailed Metrics Information (Nutanix: Image 3)

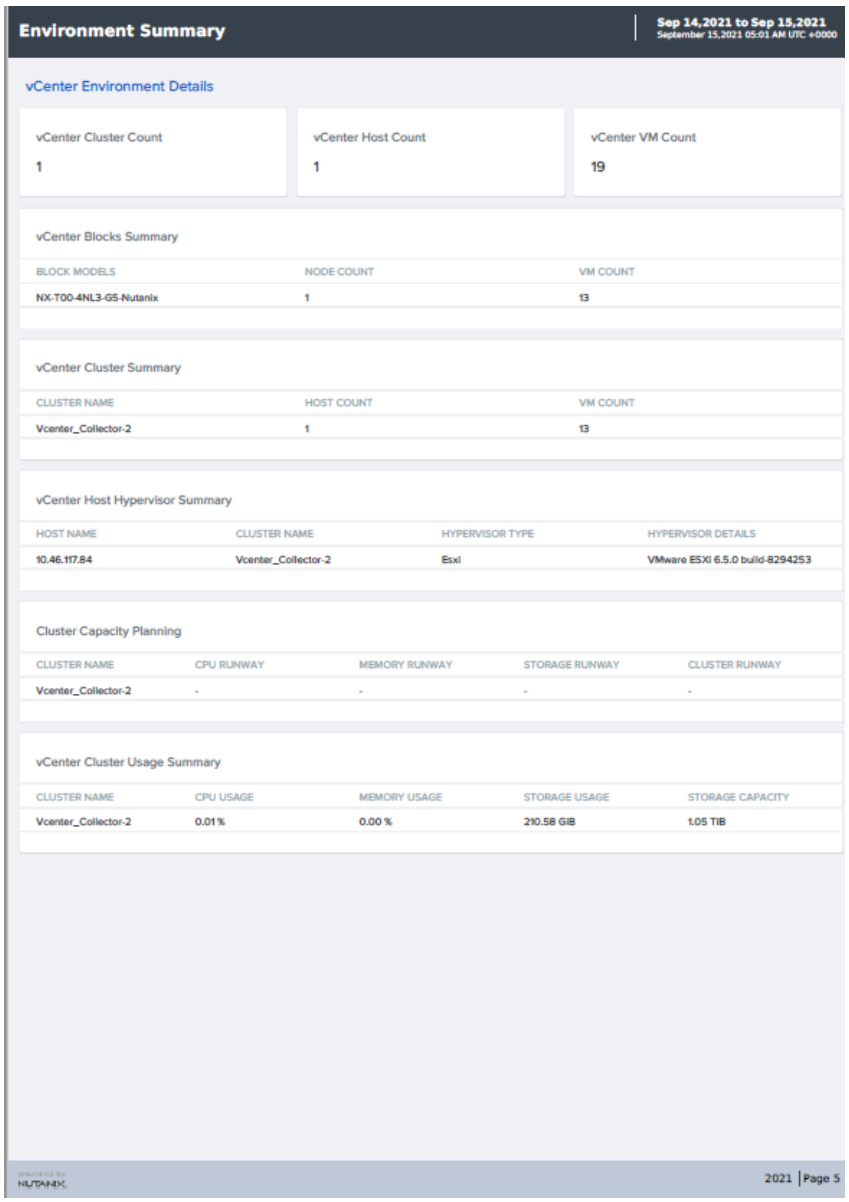

Environment Summary

Sep 14,2021 to Sep 15,2021
September 15,2021 05:01 AM UTC +0000

vCenter Environment Details

vCenter Cluster Count	vCenter Host Count	vCenter VM Count
1	1	19

vCenter Blocks Summary

BLOCK MODELS	NODE COUNT	VM COUNT
NX-T00-4NL3-G5-Nutanix	1	13

vCenter Cluster Summary

CLUSTER NAME	HOST COUNT	VM COUNT
Vcenter_Collector-2	1	13

vCenter Host Hypervisor Summary

HOST NAME	CLUSTER NAME	HYPERVISOR TYPE	HYPERVISOR DETAILS
10.46.117.84	Vcenter_Collector-2	Esxi	VMware ESXi 6.5.0 build-8294253

Cluster Capacity Planning

CLUSTER NAME	CPU RUNWAY	MEMORY RUNWAY	STORAGE RUNWAY	CLUSTER RUNWAY
Vcenter_Collector-2	-	-	-	-

vCenter Cluster Usage Summary

CLUSTER NAME	CPU USAGE	MEMORY USAGE	STORAGE USAGE	STORAGE CAPACITY
Vcenter_Collector-2	0.01 %	0.00 %	210.58 GiB	1.05 TiB

powered by NUTANIX

2021 | Page 5

Figure 511: Detailed Metrics Information (vCenter: Image 1)

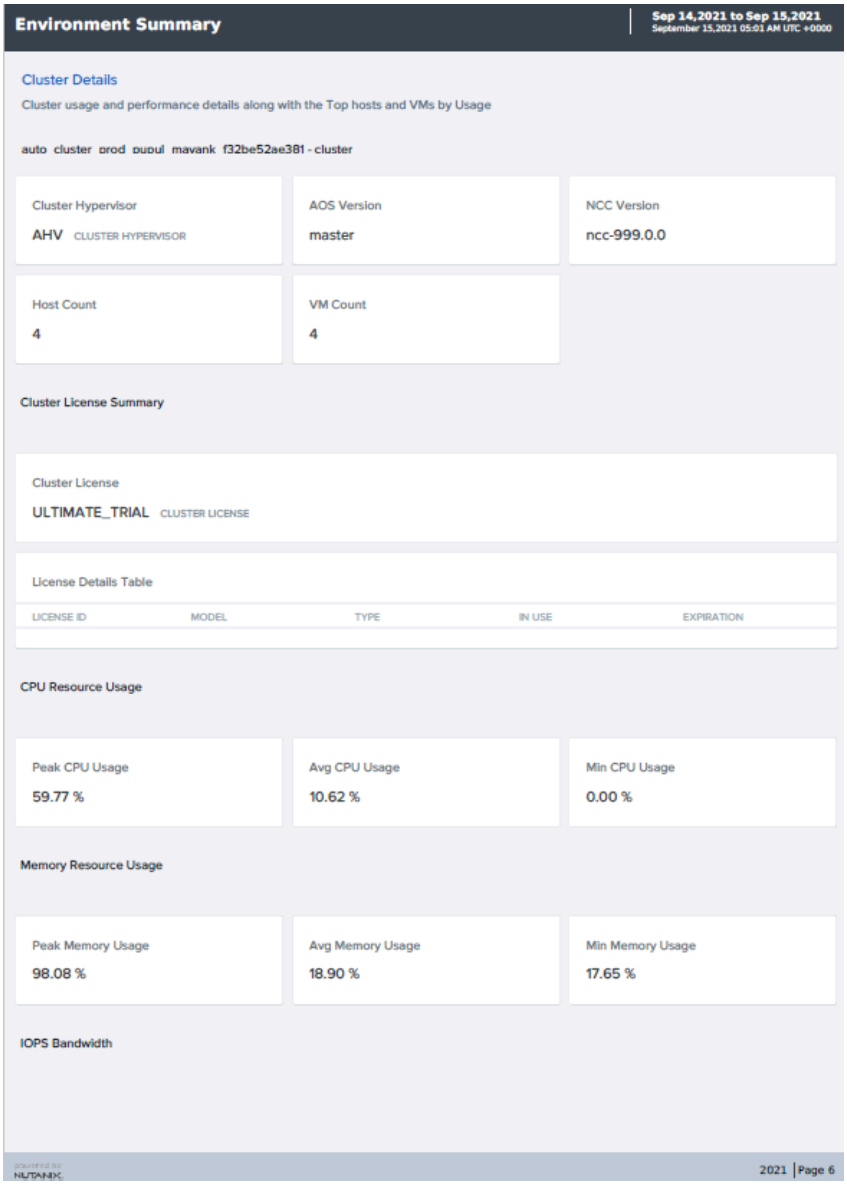

Figure 512: Detailed Metrics Information (vCenter: Image 2)

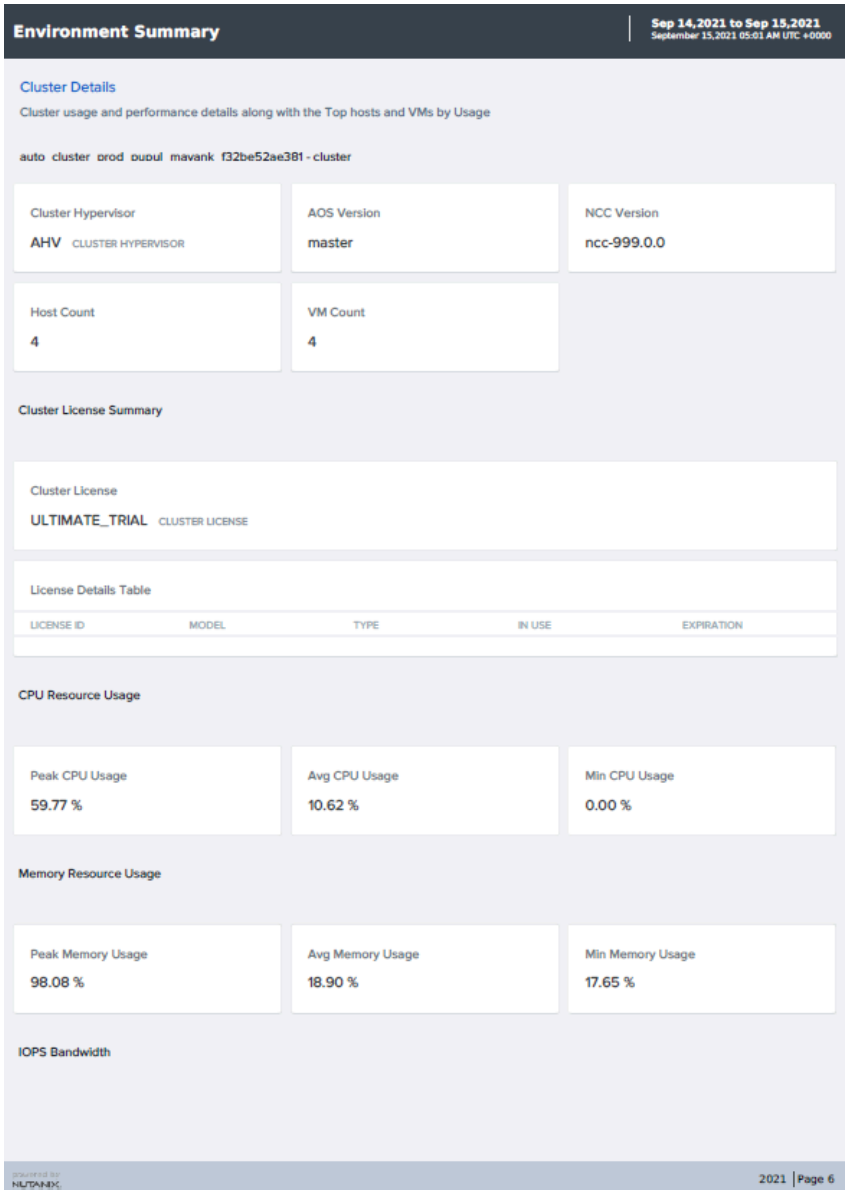

Figure 513: Detailed Metrics Information (vCenter: Image 3)

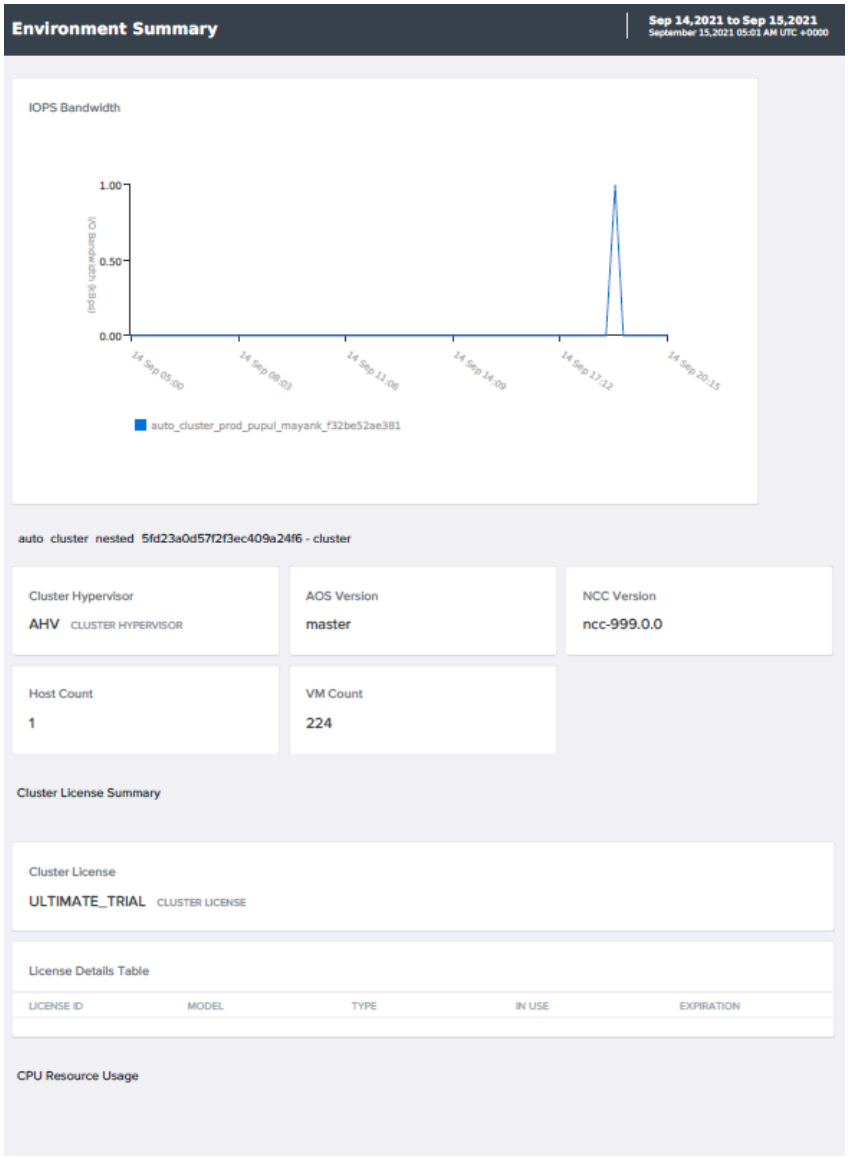

Figure 514: Detailed Metrics Information (vCenter: Image 4)

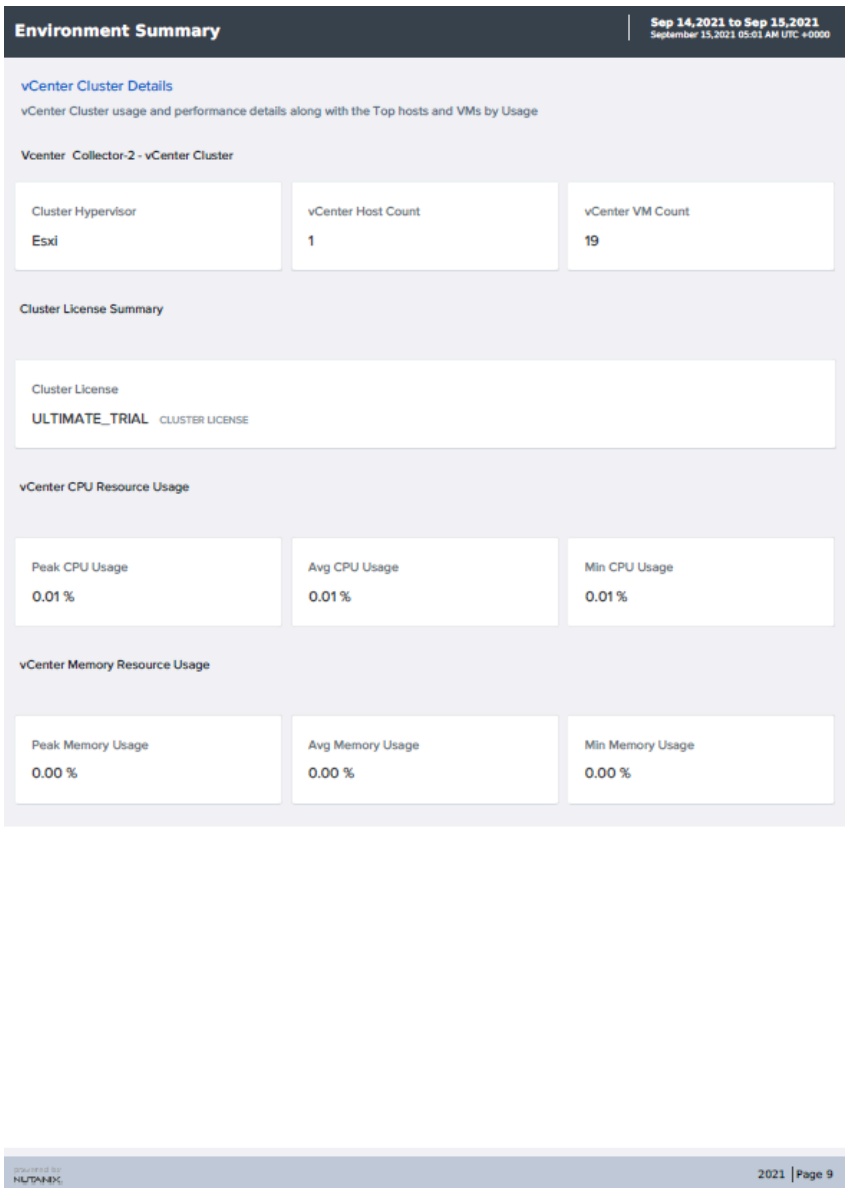

Figure 515: Detailed Metrics Information (vCenter: Image 5)

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.

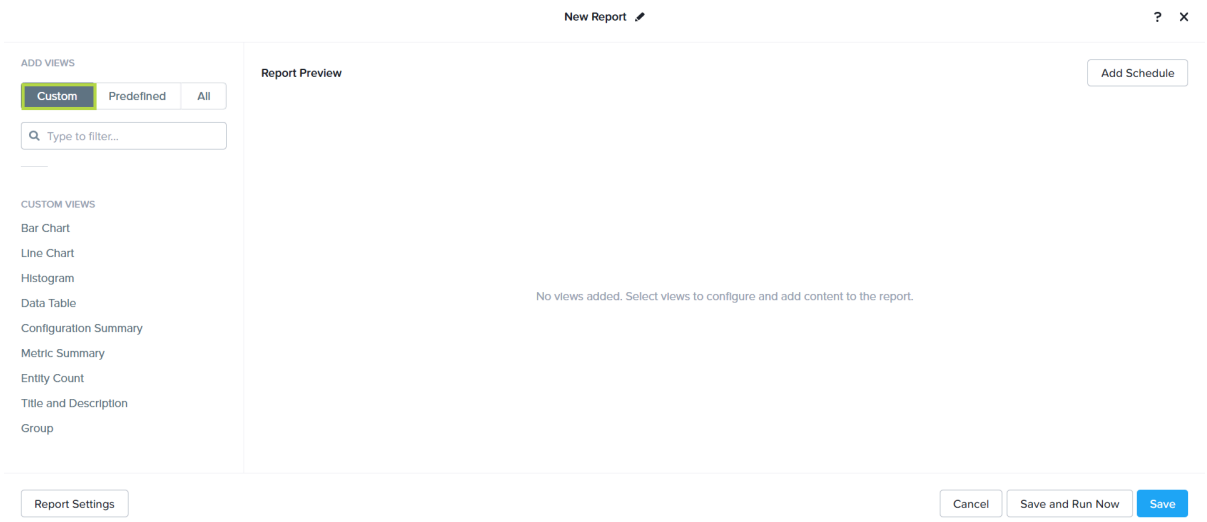

Figure 516: Custom View

Figure 517: Predefined View

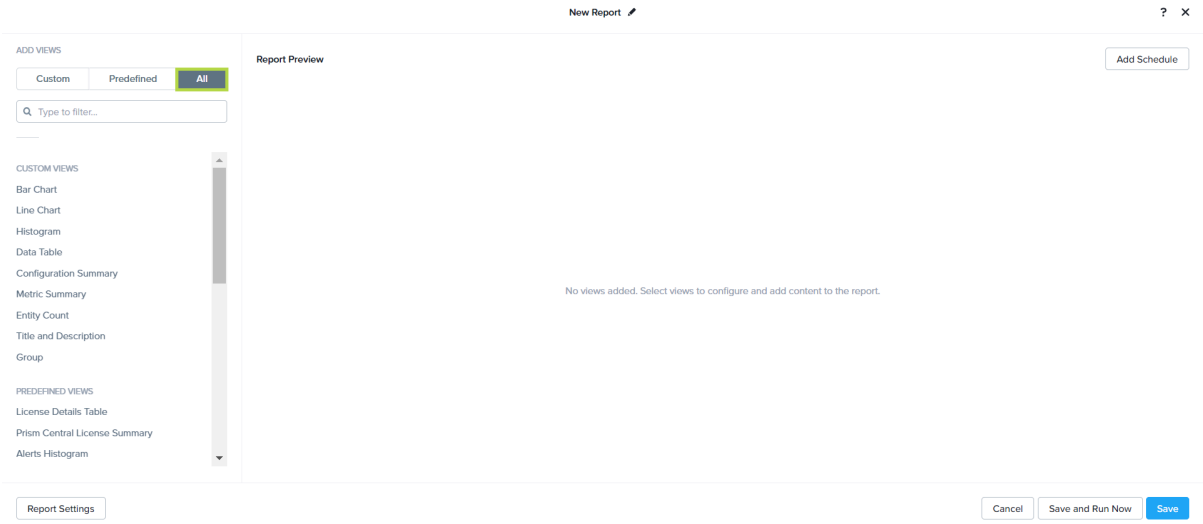

Figure 518: All View (Custom and Predefined Combined)

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 960 and Adding a Bar Chart or a Line Chart or a Histogram View to a Report (vCenter) on page 963.
- To add a data table view, see Adding a Data Table View to a Report (Nutanix) on page 965 and Adding a Data Table View to a Report (vCenter) on page 968.
- To add a configuration summary view, see Adding Configuration Summary View on page 970.
- To add a metric summary view, see Adding Metric Summary View (Nutanix) on page 972 and Adding Metric Summary View (vCenter) on page 974.
- To add an entity count view, see Adding Entity Count View (Nutanix) on page 976 and Adding Entity Count View (vCenter) on page 977.
- To add a title and description view, see Adding Title and Description View on page 979.
- To add a group view, see Adding a Group View on page 979.
- To add a predefined view, see Adding a Predefined View on page 983.

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.

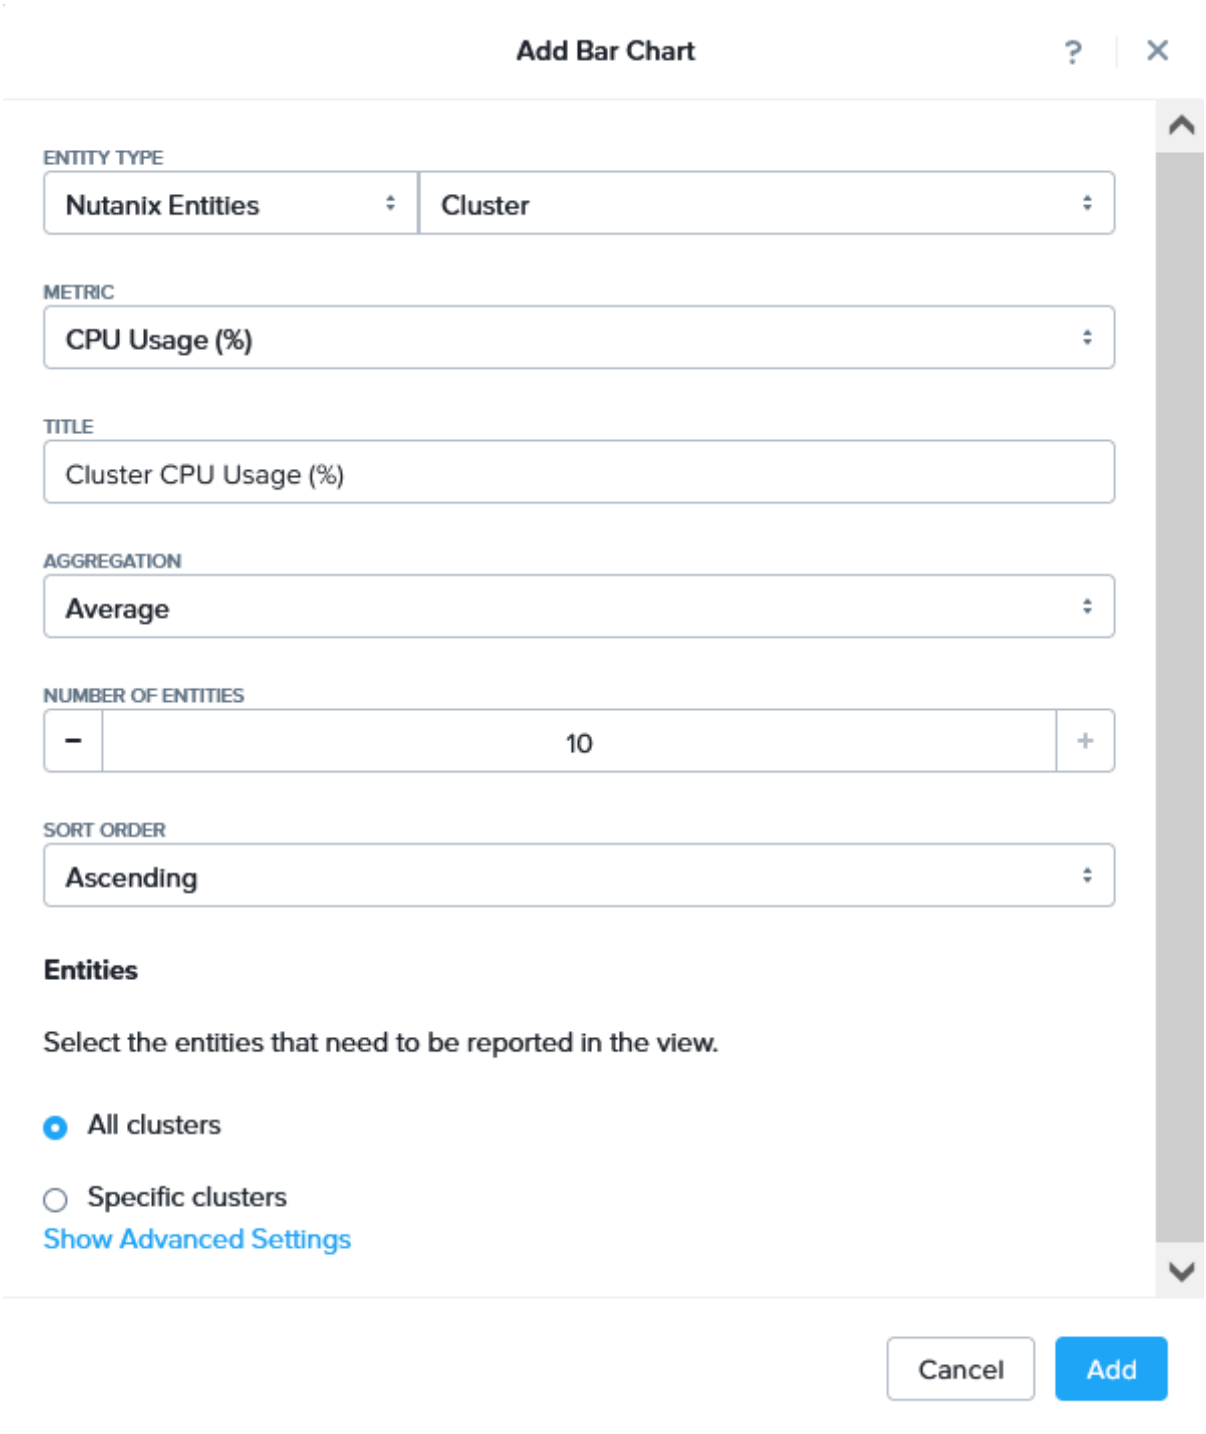

Figure 519: Bar Chart View

   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) on page 963

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 Bar Chart or a Line Chart or a Histogram View to a Report (vCenter)

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.

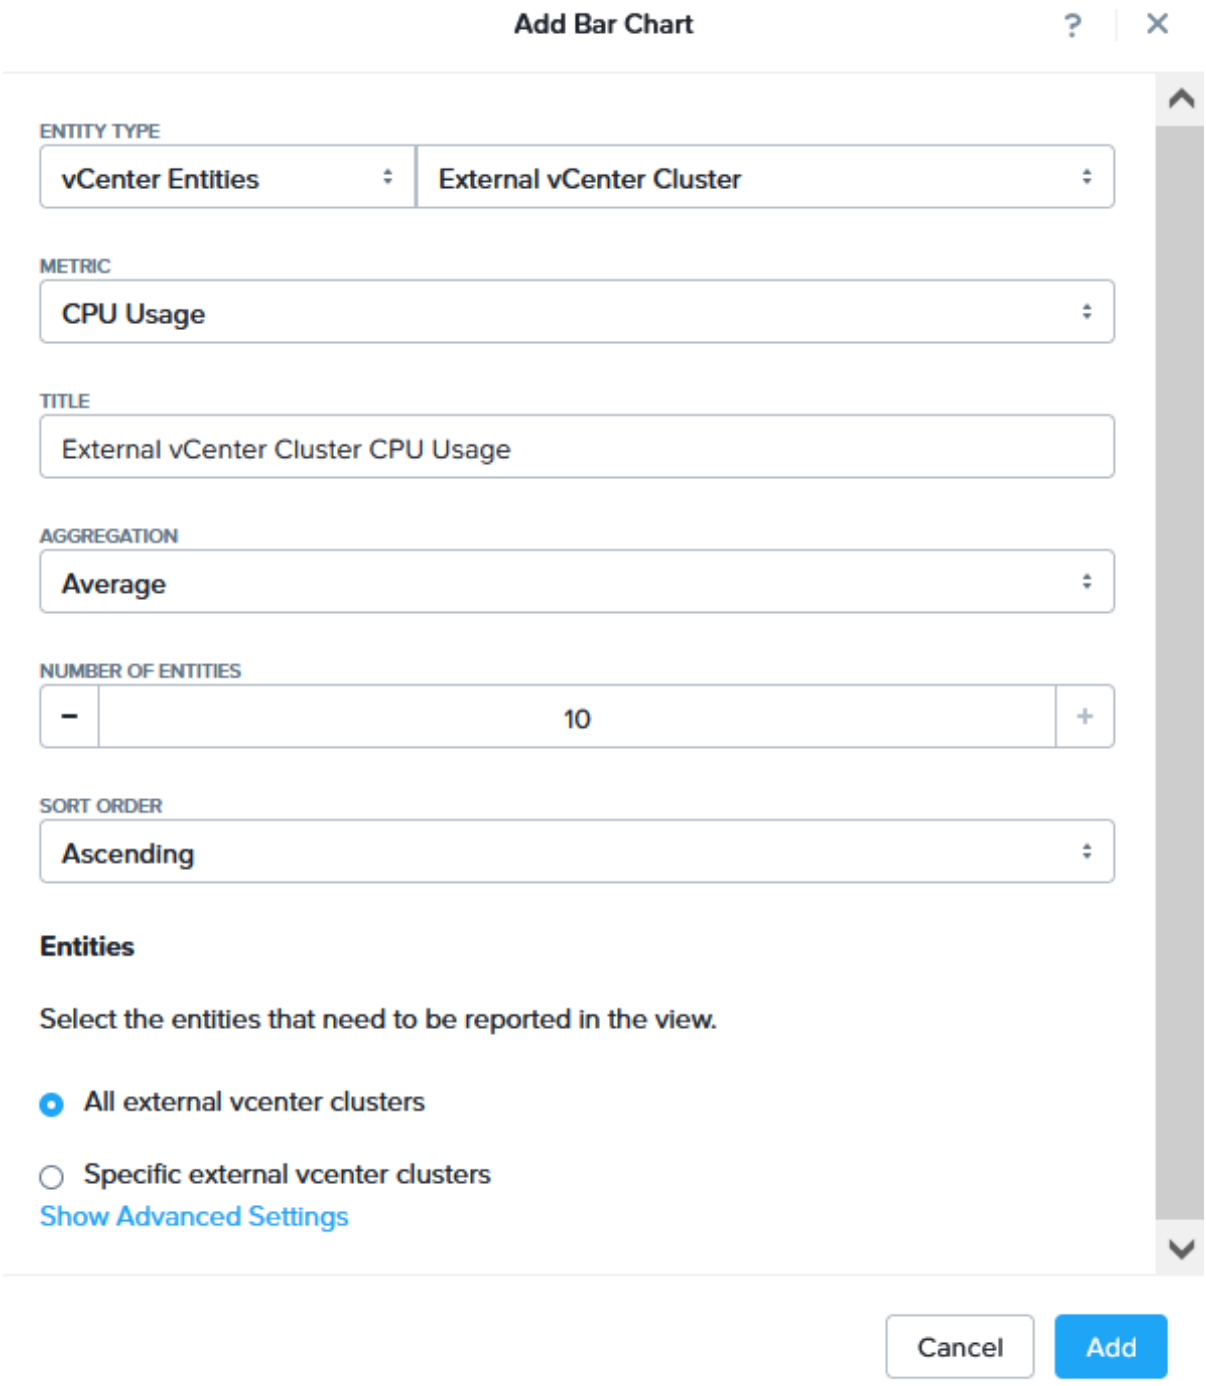

Figure 520: Bar Chart View

   a. Select **vCenter Entities** and the entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** drop-down menu for which you want to configure the view.

b. Select a specific metric of the entity type.

The metric selection depends on what you have selected in the **Entity Type** drop-down menu. For example, if you select **External vCenter VM** as an entity type, you can select one of the following metric.

Ready, **CPU Usage**, **Highest latency**, **Average read requests per second**, **Average read requests per second**, **Read rate**, **Read latency**, **Write latency**, **Controller Write I/O Bandwidth**, **Controller Read I/O Latency**, **Controller Write I/O Latency**, **Storage Logical Usage**, **Controller Read IOPS**, **Controller Write IOPS**.

c. Enter the title of the bar chart, line chart, or histogram view in the **Title** field.

d. Select the aggregation from the **Aggregation** menu.

For the metrics that has time-series data, you can define the aggregation field. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric and select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the report.

You can define the aggregation value of **Sum**, **Max**, **Min**, **Average**, **Last**, or **Count**.

> **Note:** Aggregation is not supported for Line Chart view.

e. Enter the number of entities that you want to include in the report from the **Number of Entities** text box.

> **Note:** This field does not appear for the Histogram view.

f. Select the **Limit entities in the chart** option to limit the number of entities that you want to include in the view and how you represent the entities (**Ascending** or **Descending**).

g. (Only for Histogram view) Enter number of buckets in the **Number of Objects** field.

You can define the frequency distribution for all the entities in the entity type for the selected metric by using this field. For example, for a selected metric if the minimum value is 10 and maximum value is 100 and you have entered number of buckets as 5 then you will have 18 buckets configured (100-10/5). Depending on the value of the metric, different metrics will fall in the appropriate buckets.

h. Select the entities that you want to add.

- **All Entities**: Select this option if you want to include all the entities that you have selected in the view
- **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** drop-down menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

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.

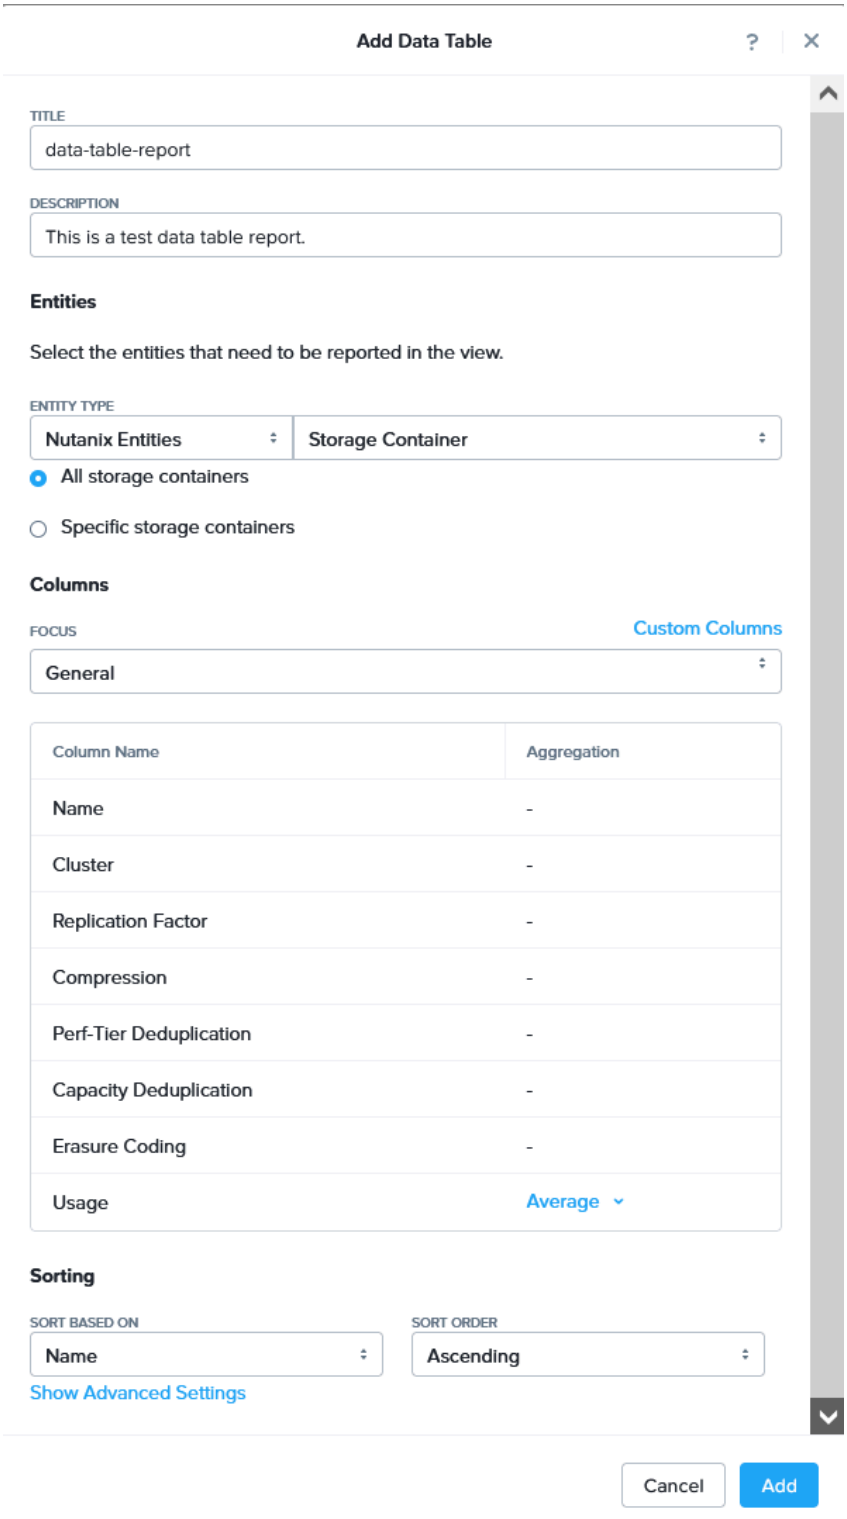

Figure 521: Data Table View

category

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.

   If you select *Clusters*, *VMs*, or *Hosts* as entities along with the specific entities option, and then select *Categories* in the first drop-down list for defining the rules, you can select *Equal to* or *Not Equal to* in the second drop-down list, and then enter a combination of category name and category value in the format *category name:category value* in the last text box to filter the relevant data. For example, you must enter *AppType:Hadoop* to filter the VMs that are assigned *Apptype* category and *Hadoop* as its value (see the image below). You can see the category name and category value correlation in the Categories view (see

Categories Summary View on page 164). You can leave the last text box blank to filter all the entities that do not have a user-defined category assigned to them.

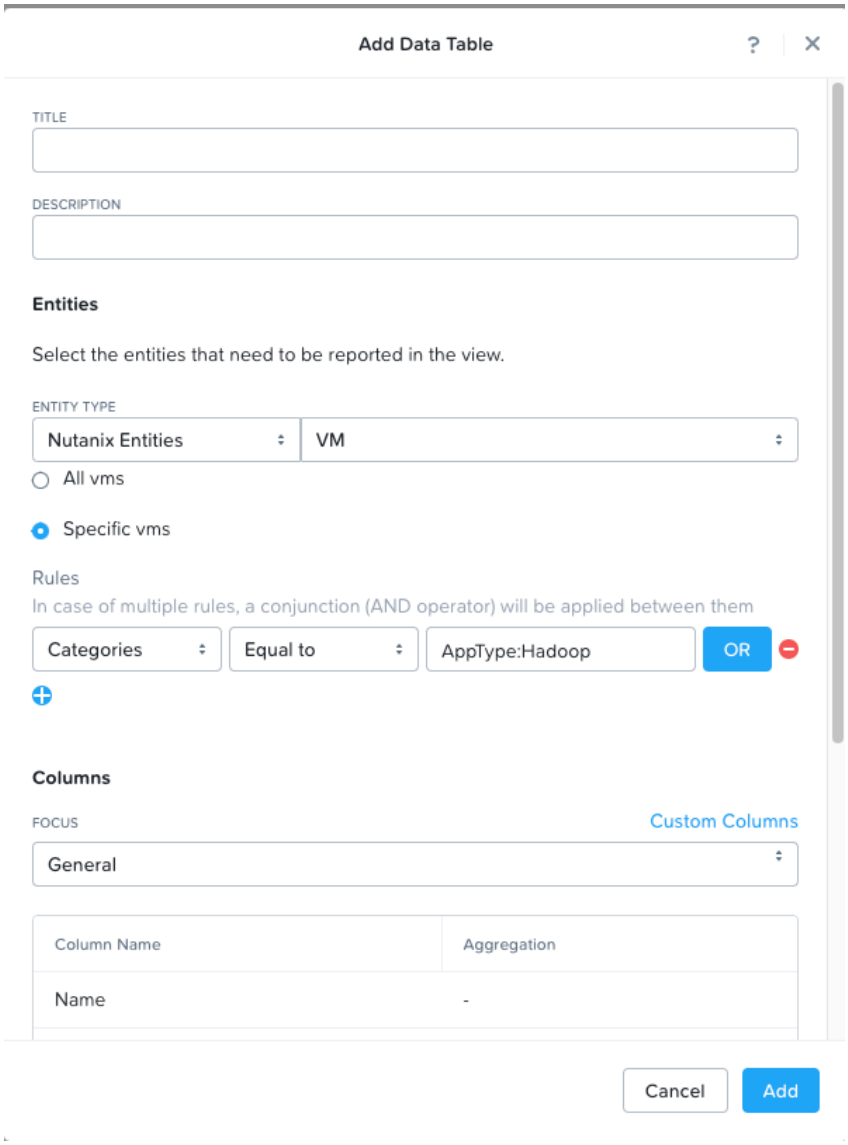

Figure 522: Using Category with Data Table Rules

3. Click **Add**.
   The configured view gets added to the report.

Adding a Data Table View to a Report (vCenter)

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.

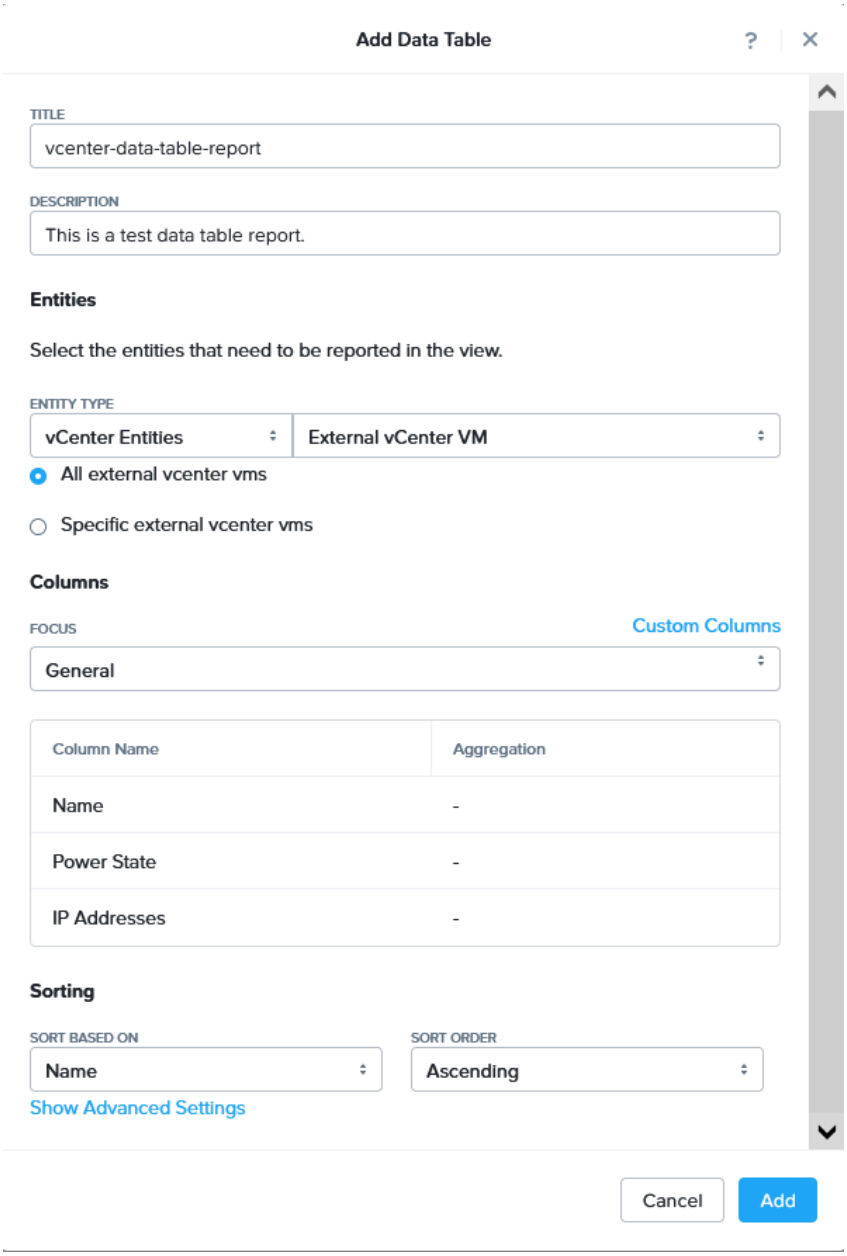

Figure 523: Data Table View

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 **vCenter Entities** and the entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** drop-down menu 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 external vcenter 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 External vCenter VM and desire to limit the view for the 50 powered on VMs, you can select **Specific external vcenter cms** 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.

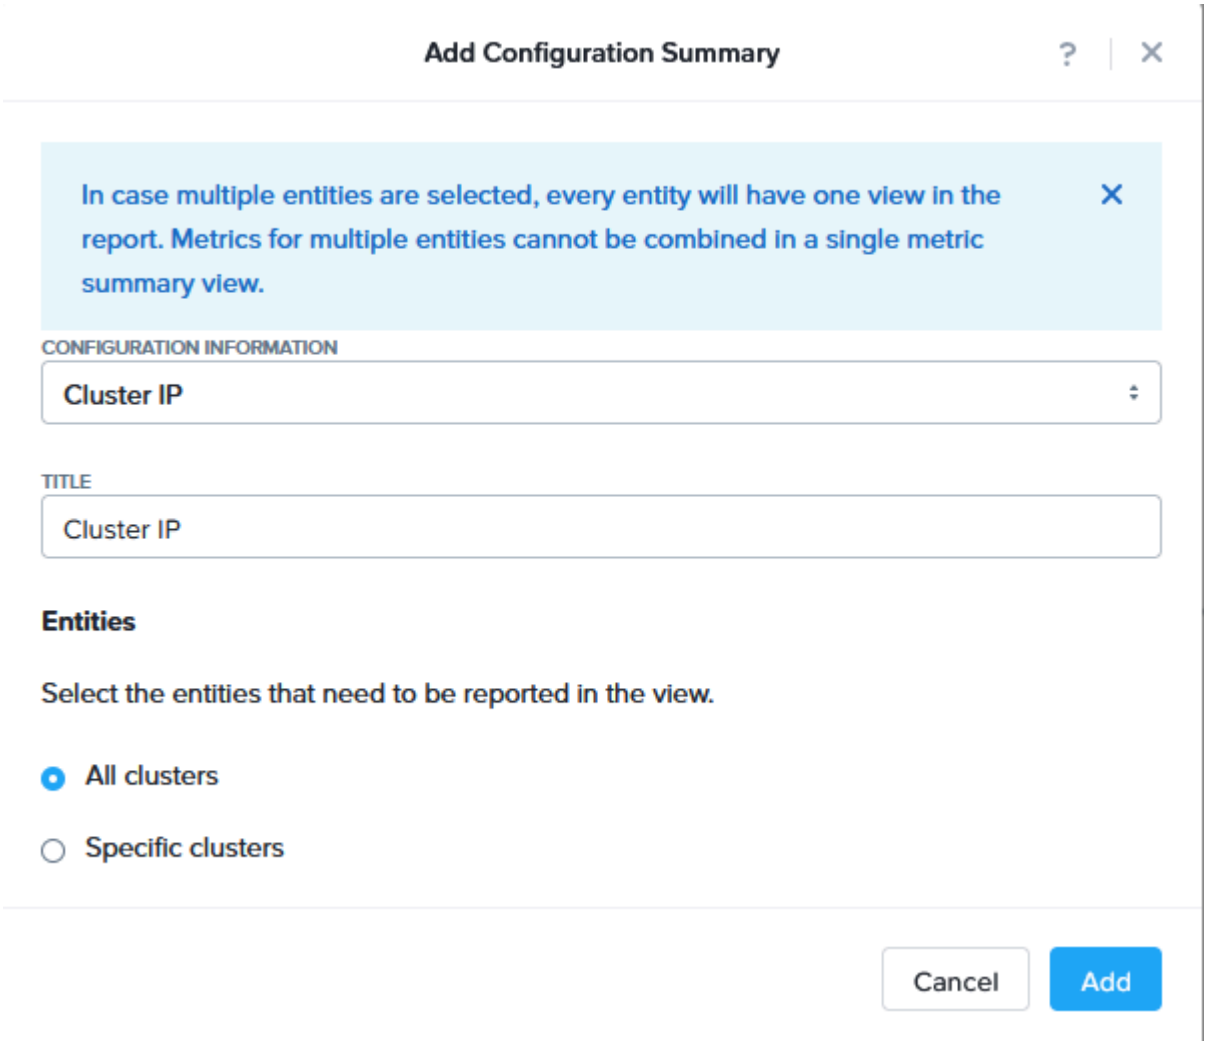

Figure 524: Configuration Summary View

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.

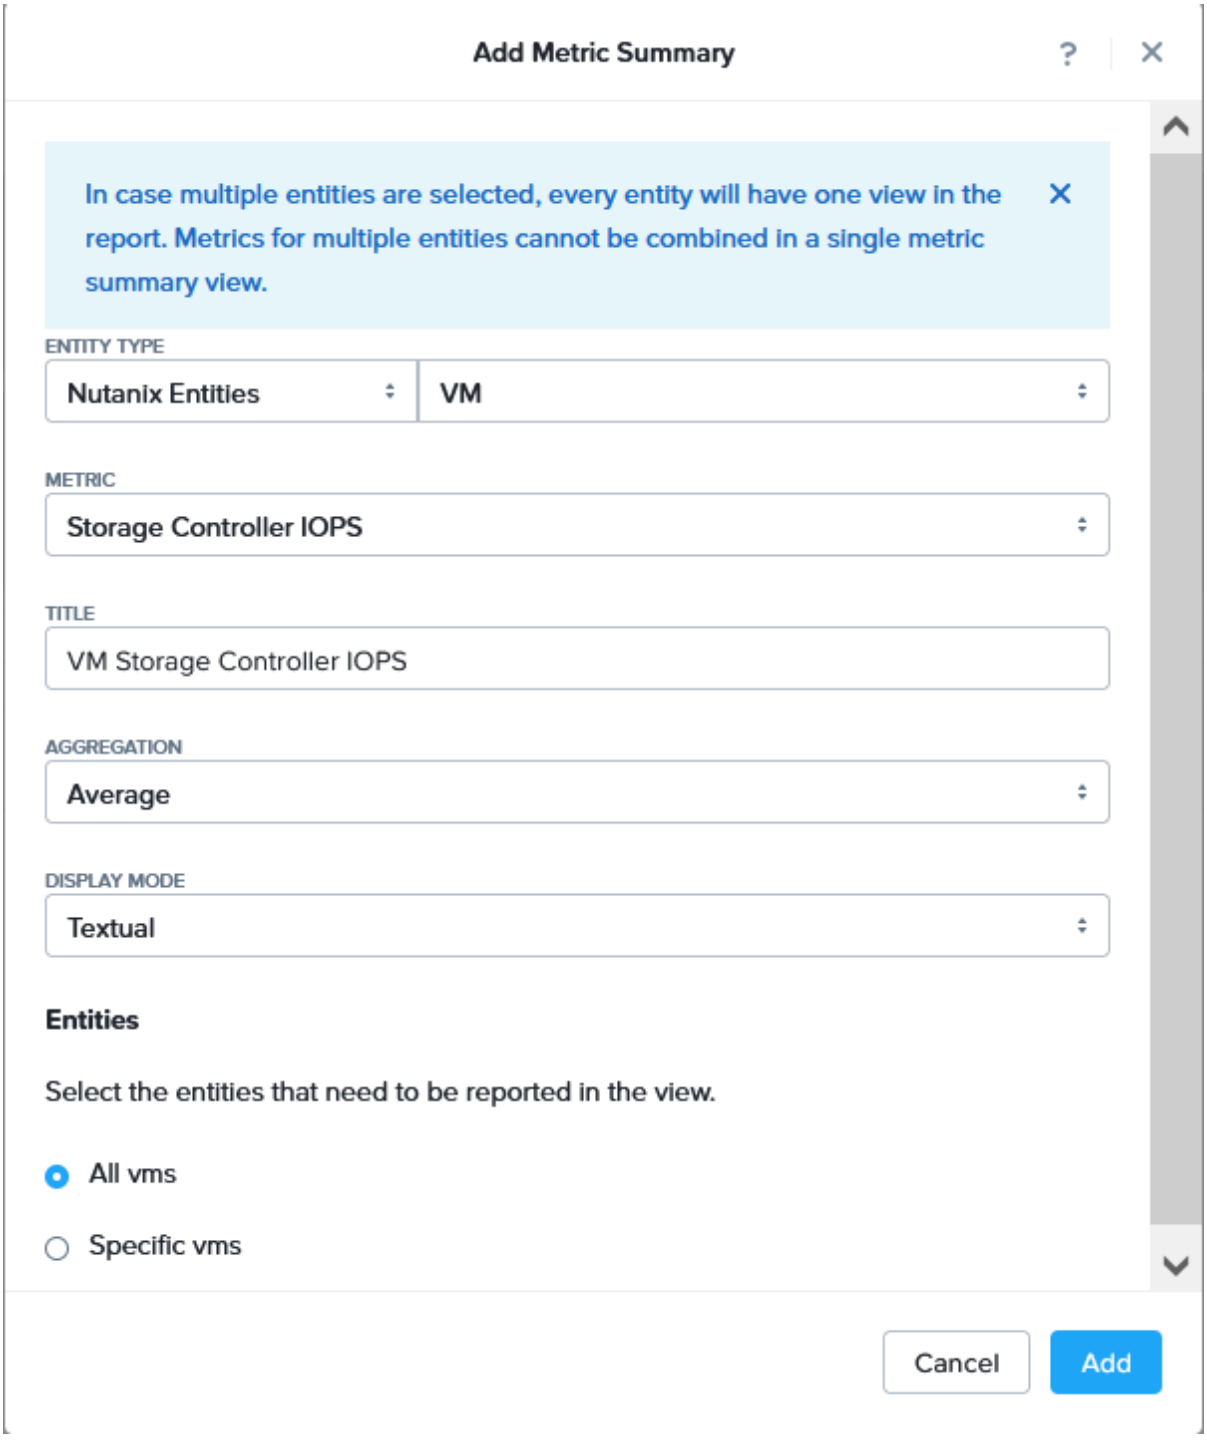

Figure 525: Metric Summary View

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 Metric Summary View (vCenter)

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.

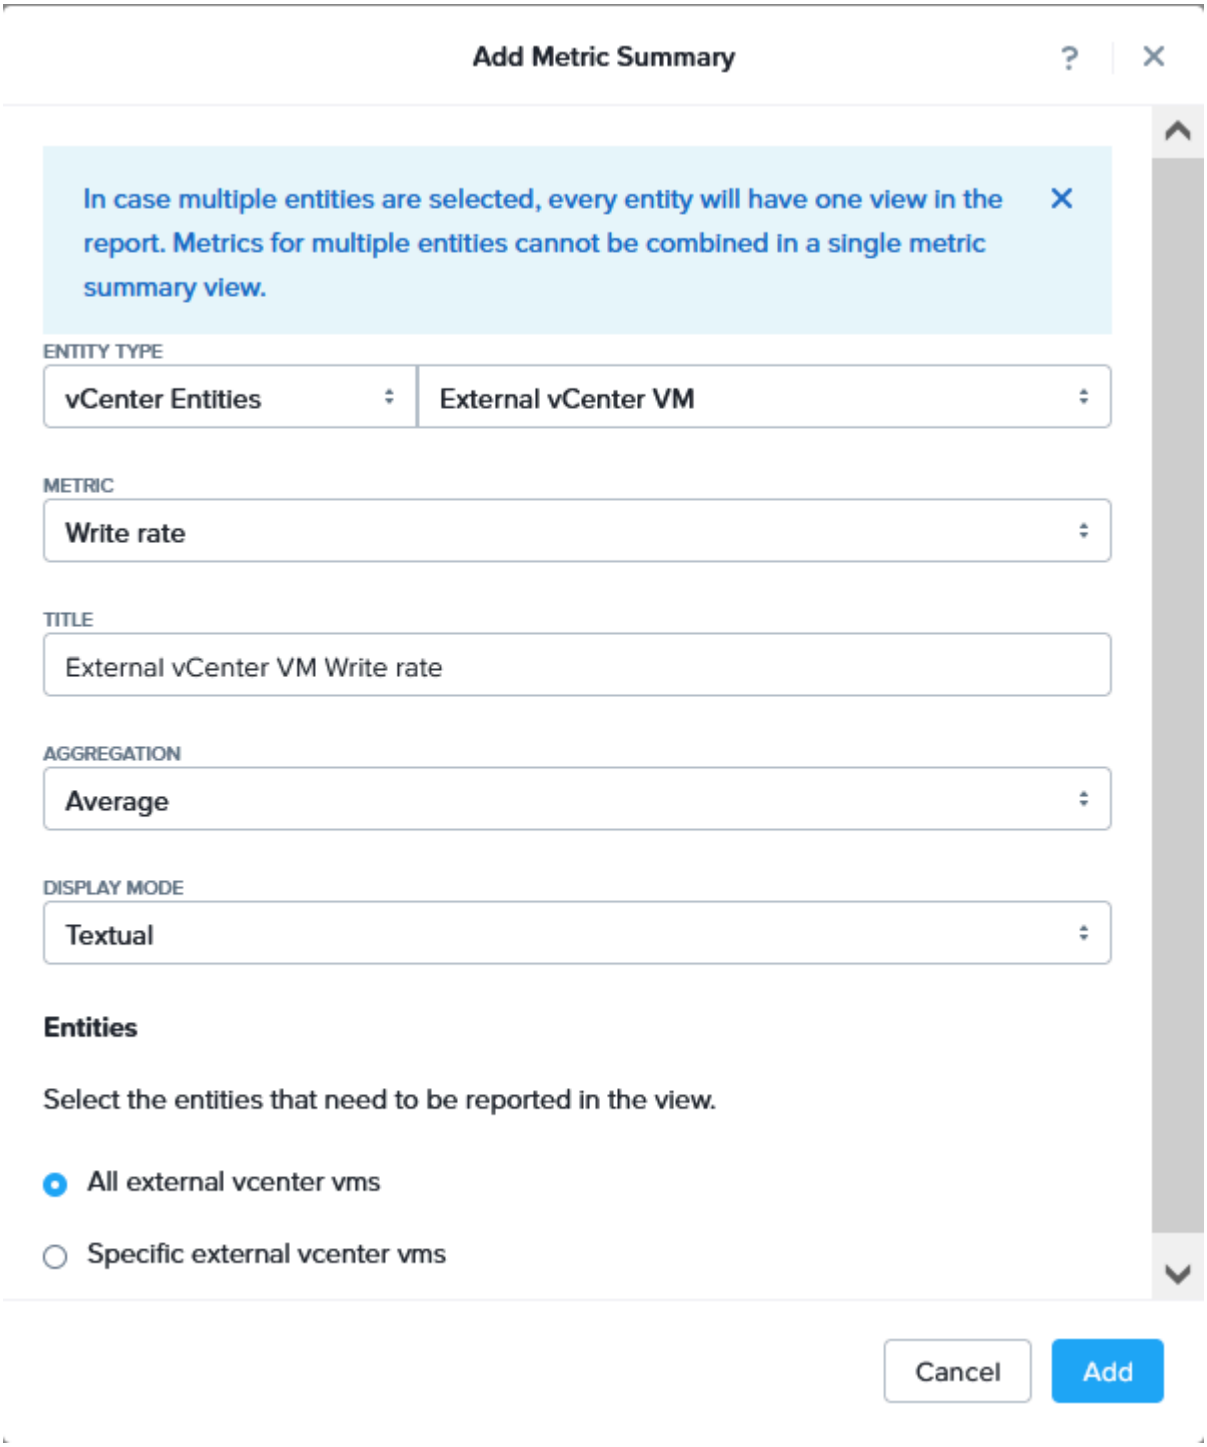

Figure 526: Metric Summary View

a. Select **vCenter Entities** and the entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** drop-down menu for which you want to configure the view.

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 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.

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.

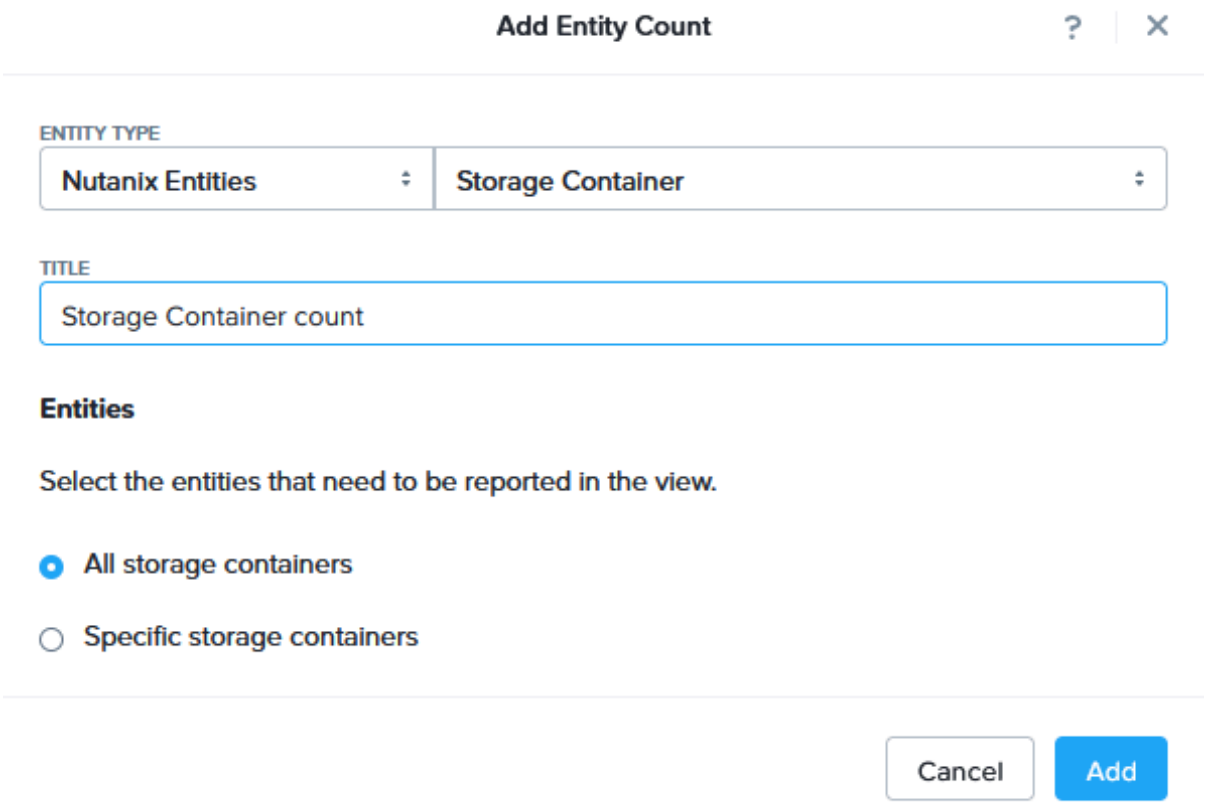

Figure 527: Entity Count View

   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 Entity Count View (vCenter)

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.

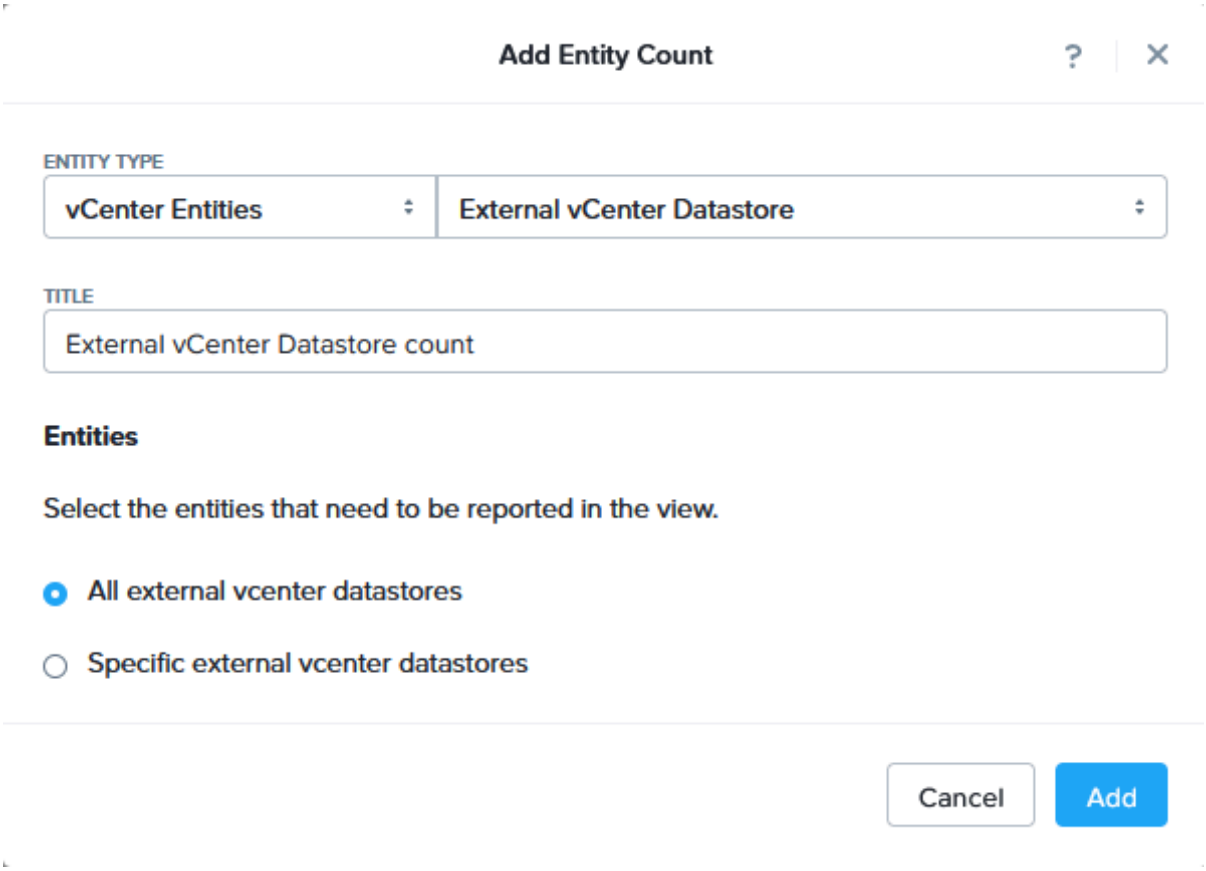

Figure 528: Entity Count View

   a. Select **vCenter Entities** and entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** drop-down menu for which you want to configure the view.
   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 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.

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.

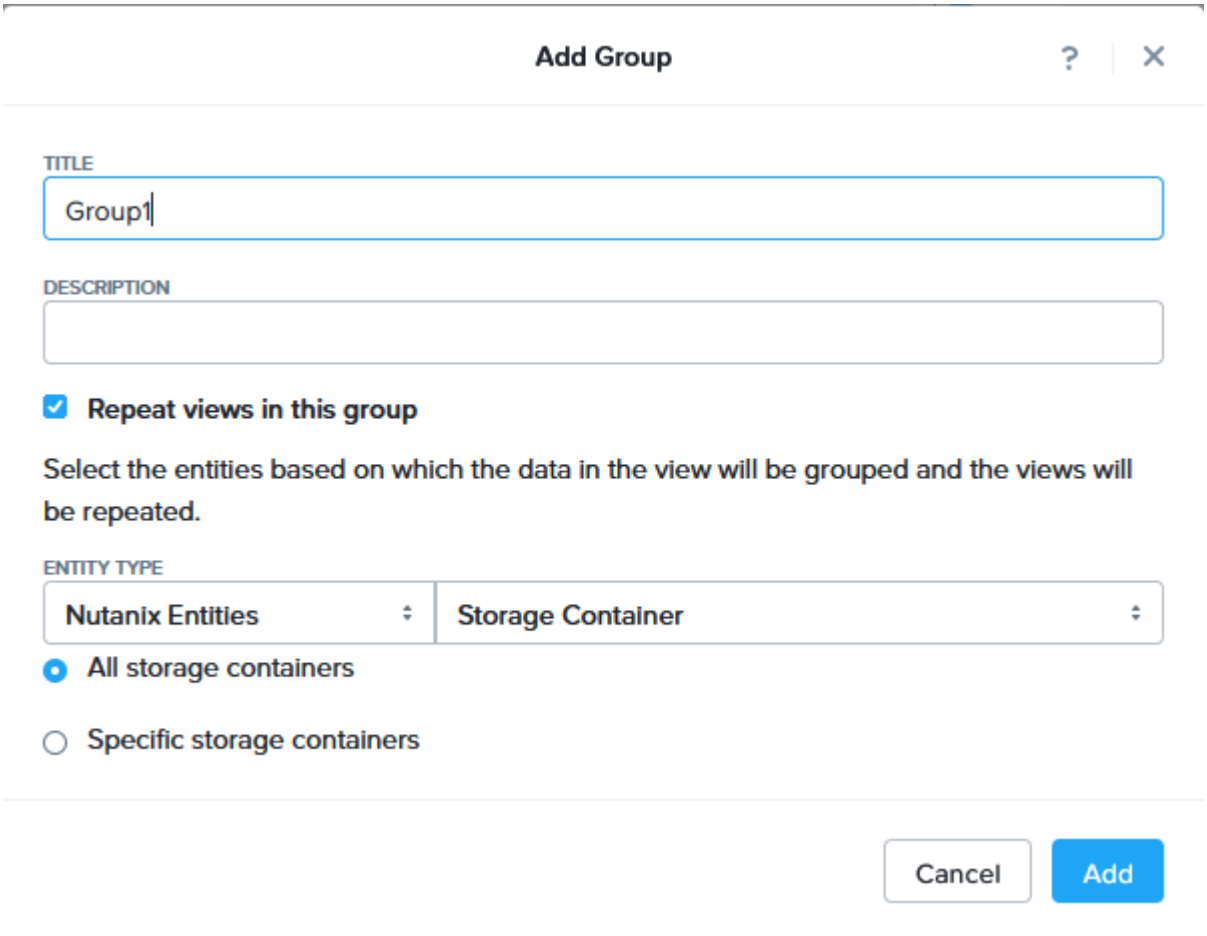

Figure 529: Group View (Nutanix)

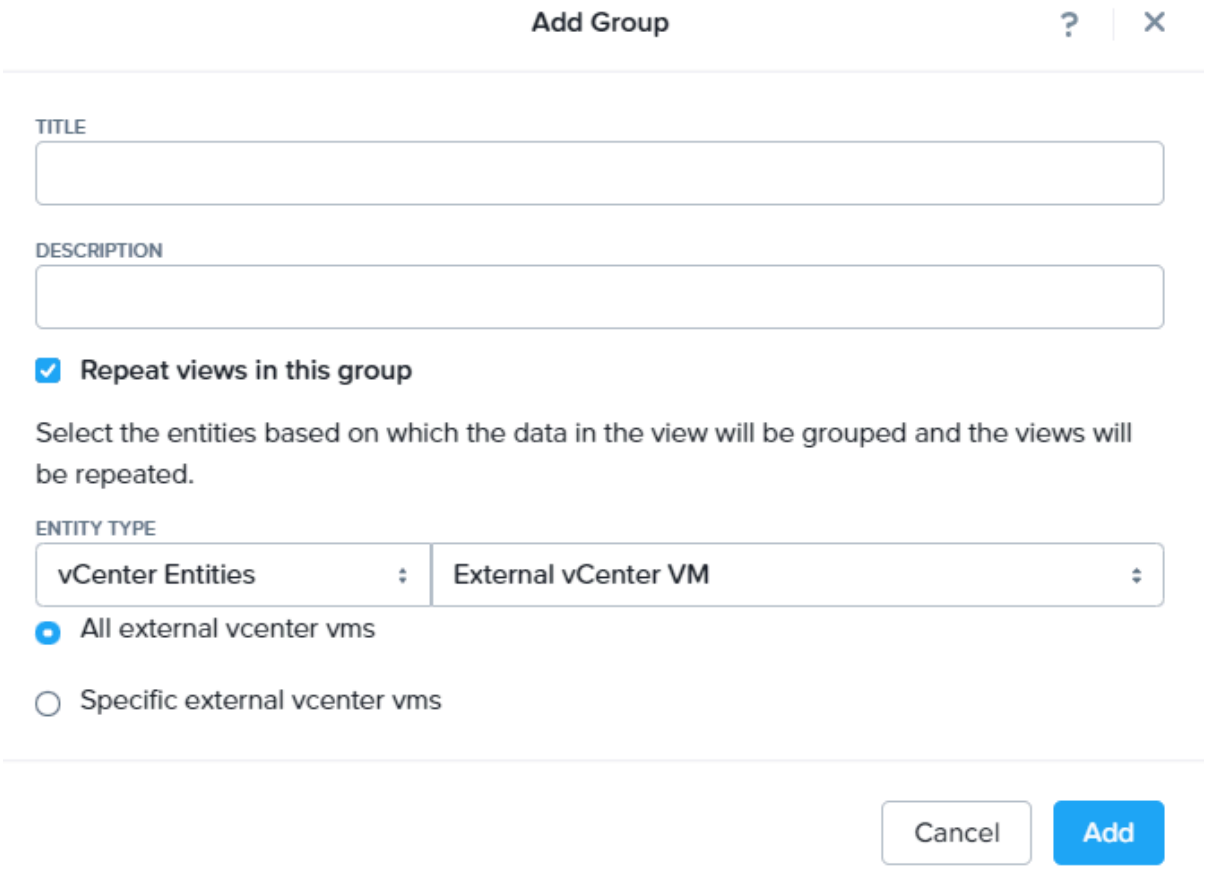

Figure 530: Group View (vCenter)

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 982

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 979

   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**.

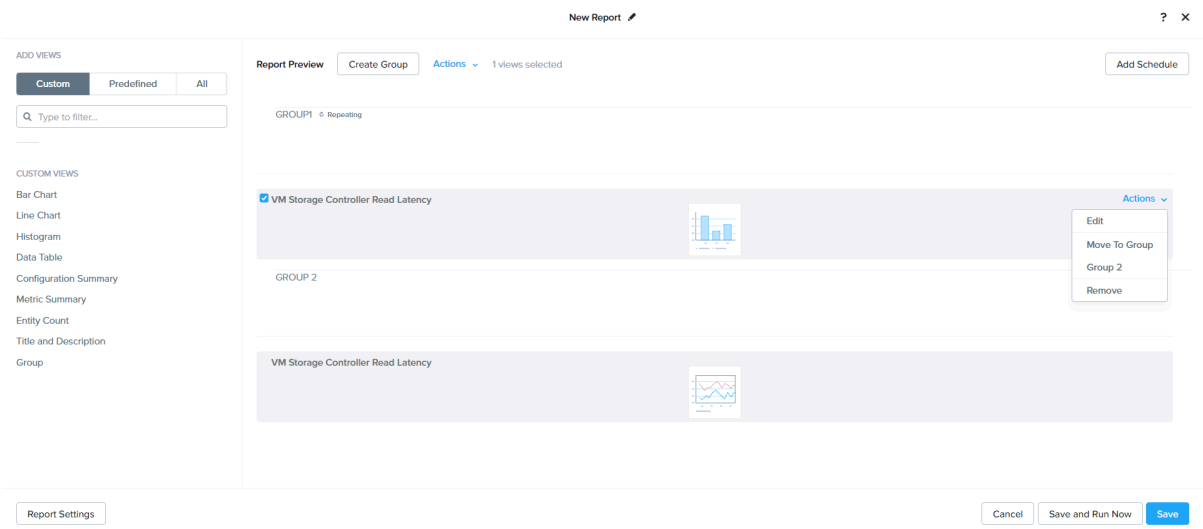

Figure 531: Adding Views to Group view

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 979.
   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 can 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.

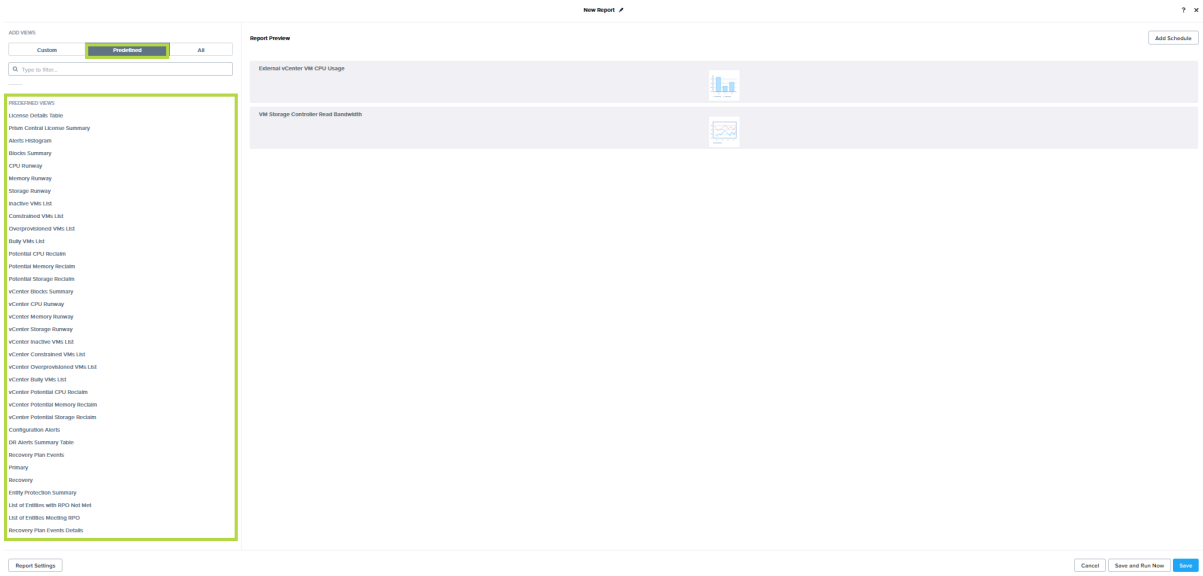

Figure 532: Pre-Defined Views

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.

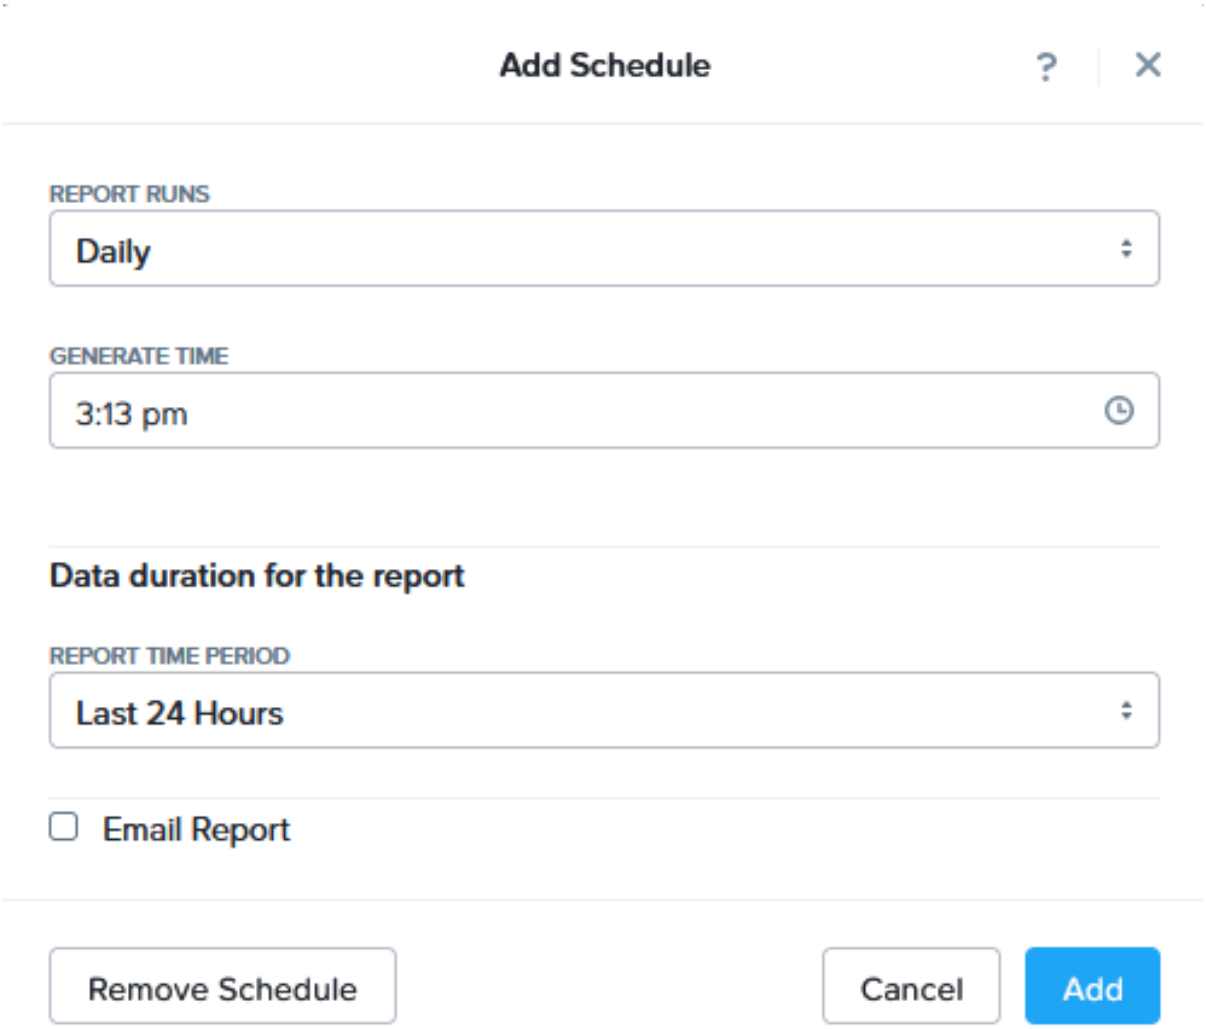

Figure 533: Report Scheduling

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 246), 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 943.

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 246).
   - » 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. For colors, you can specify any of the 16 color names

supported for HTML: aqua, black, blue, fuchsia, gray (or grey), green, lime, maroon, navy, olive, purple, red, silver, teal, white, and yellow.

> **Note:** The copyright information will not be displayed in the footer section, therefore the ability to customize copyright information is not allowed.

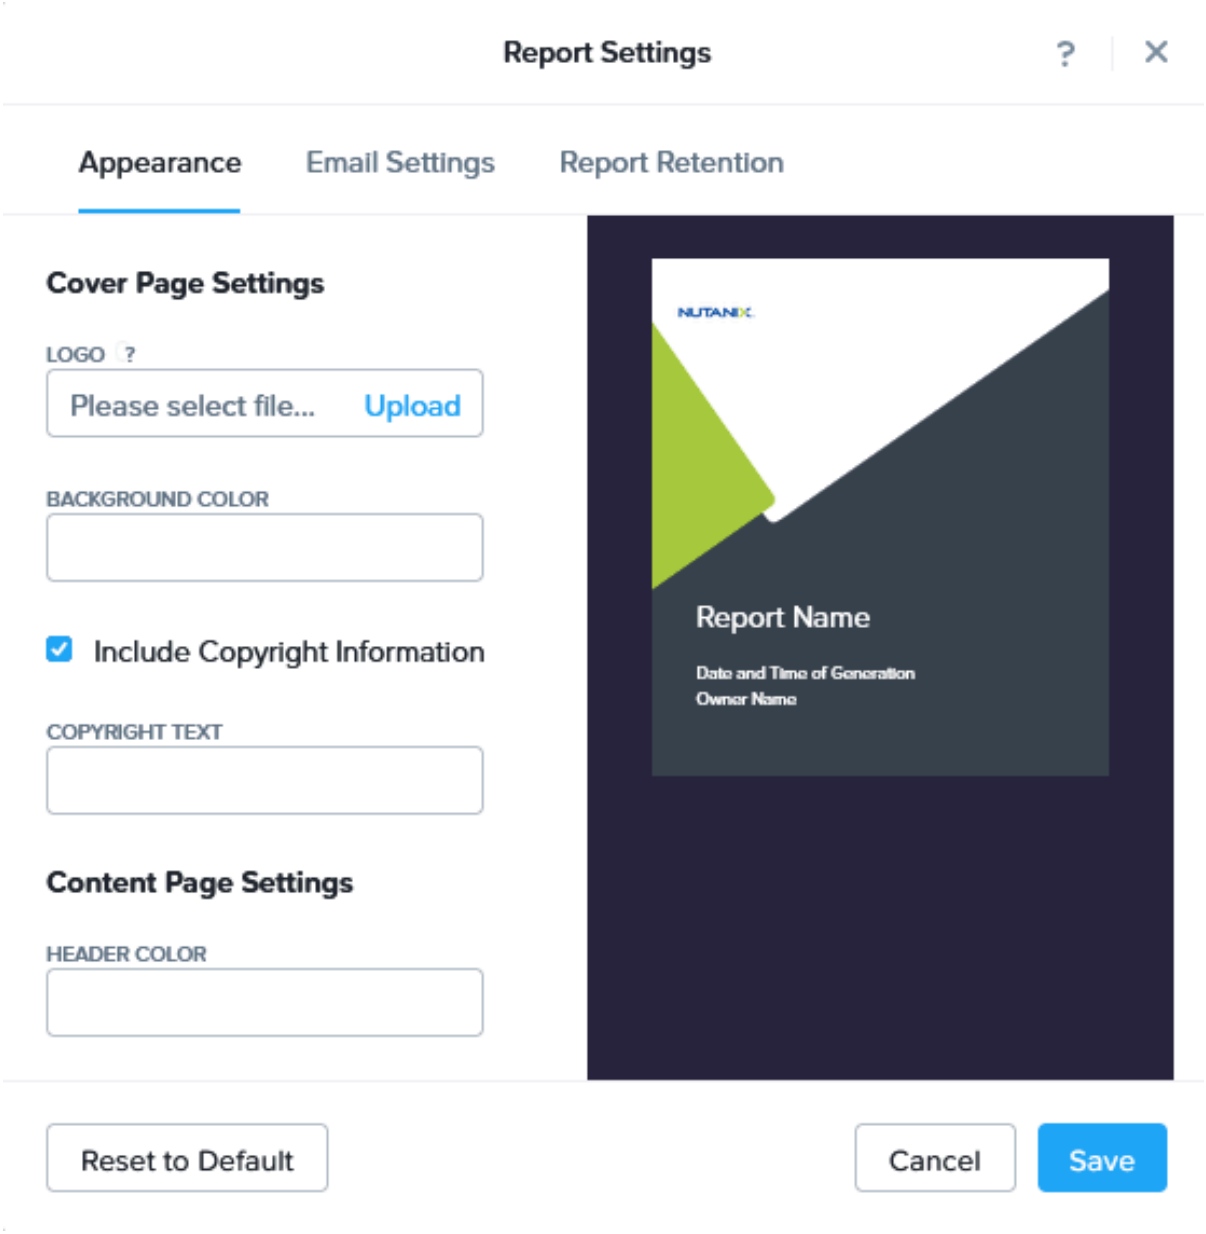

Figure 534: Report Settings Configuration

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 772: Roles and Permissions

Access/ Permissio	Consumer	Developer	Operator	Prism Admin	Prism Viewer	Project Admin	Self-Service Admin	Super Admin
Create Common Report Config				#		#	#	#
Create Report Config				#		#	#	#
Create Report Instance				#		#	#	#
Delete Common Report Config				#		#	#	#
Delete Report Config				#		#	#	#
Delete Report Instance				#		#	#	#
Notify Report Instance	#	#	#	#	#	#	#	#
Update Common Report Config				#		#	#	#
Update Report Config				#		#	#	#
View Common Report Config	#	#	#	#	#	#	#	#
View Report Config	#	#	#	#		#	#	#
View Report Instance	#	#	#	#	#	#	#	#

SERVICES ENABLEMENT

You can enable selected services through Prism Central.

- Enabling Calm on page 992
- Enabling Files on page 992
- Enabling Foundation Central on page 993
- Enabling Karbon on page 993
- Enabling Objects on page 993

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 (Prism Central) 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 (Prism Central) 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 (Prism Central) 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 (Prism Central) 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 Documentation 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 (Prism Central) 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 (Prism Central) 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.

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 995 and Configuring Pulse on page 997).
- 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 (Prism Central) on page 1011).
- If you need help, you can create a support ticket directly from Prism Central (see Creating a Support Case on page 1006).

Pulse Health Monitoring

The Pulse feature provides diagnostic system data about Prism Central or a Prism Element cluster to Nutanix Support to help deliver proactive, context-aware support for Nutanix solutions. When you enable Pulse, Pulse periodically sends data to Nutanix Insights and the Nutanix Support team for troubleshooting. Pulse is the underlying technology that securely transmits system-level diagnostic data to the Insights platform, enabling predictive health and context-aware support automation work#ows. Nutanix Insights is an integrated service that utilizes this data to augment product support, reducing customer case volume and expediting issue resolution time.

Pulse collects data automatically without affecting the system performance and shares only the basic system-level information required for monitoring the health and status of Prism Central or a Prism Element cluster. This information includes the following items.

- System alerts
- System tasks
- System logs
- System configuration
- Performance metrics
- Current Nutanix software version
- Nutanix processes and Controller VM information
- Hypervisor details such as type and version
- Cumulative data about the monitored clusters

All the information in the above categories is specific to the internal processes of Nutanix and does not contain information regarding customer workloads or applications.

Nutanix processes data that Pulse sends consistent with your agreement with Nutanix and where applicable, by the Nutanix Privacy Statement. Some of the Nutanix products require Pulse enablement for functionality and features. See the Nutanix Privacy Statement and applicable product documentation for more details.

Pulse frequently collects important data, like system-level statistics and configuration information, to automatically detect issues and help simplify troubleshooting. With this information, Nutanix Support can apply advanced analytics to optimize your implementation and address potential problems. Nutanix can use pulse data to trend the adoption of versions, features, and configurations.

Pulse sends messages through HTTPS (port 443) to insights.nutanix.com:443 using TLS 1.2. The HTTPS request uses certificate authentication to validate that Pulse has established communication with the Nutanix Remote Diagnostics service. The TLS 1.2 protocol uses public key cryptography and server authentication to provide

confidentiality, message integrity, and authentication for traffic passed over the Internet. For the complete list of required ports, see Ports and Protocols.

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 1005), 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 509). 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, log on to a Controller VM through SSH and run the following command.

```
nutanix@cvm$ zkcat /appliance/logical/nusights/collectors/kCommand/override_config
```

> **Note:** This command prints the Remote Diagnostics status only if the Remote Diagnostics status is set explicitly. The command does not print anything if the status is the default status.

- To disable Remote Diagnostics, log on to a Controller VM through SSH and run the following command.

  ```
  nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --
  enable=false --reason="text"
  ```

  Replace *text* with a string describing the reason for disabling Remote Diagnostics. The --reason parameter is optional.

- To enable Remote Diagnostics, log on to a Controller VM through SSH and run the following command.

  ```
  nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --
  enable=true --reason="text"
  ```

  Replace *text* with a string describing the reason for disabling Remote Diagnostics. The --reason parameter is optional.

  > **Note:** To enable remote diagnostics, clusters must run NCC 3.7.0.1 or later versions.

Configuring Pulse

About this task

- See Pulse Health Monitoring on page 995 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 (Prism Central) on page 516). 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 (Prism Central) on page 16) and select **Pulse**.

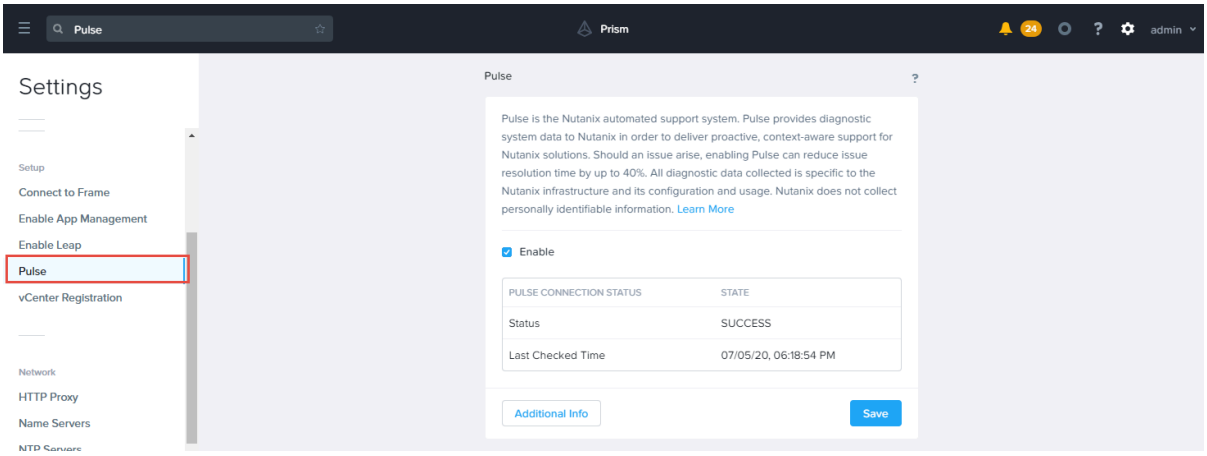

Figure 535: Pulse Window

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.

   - 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.

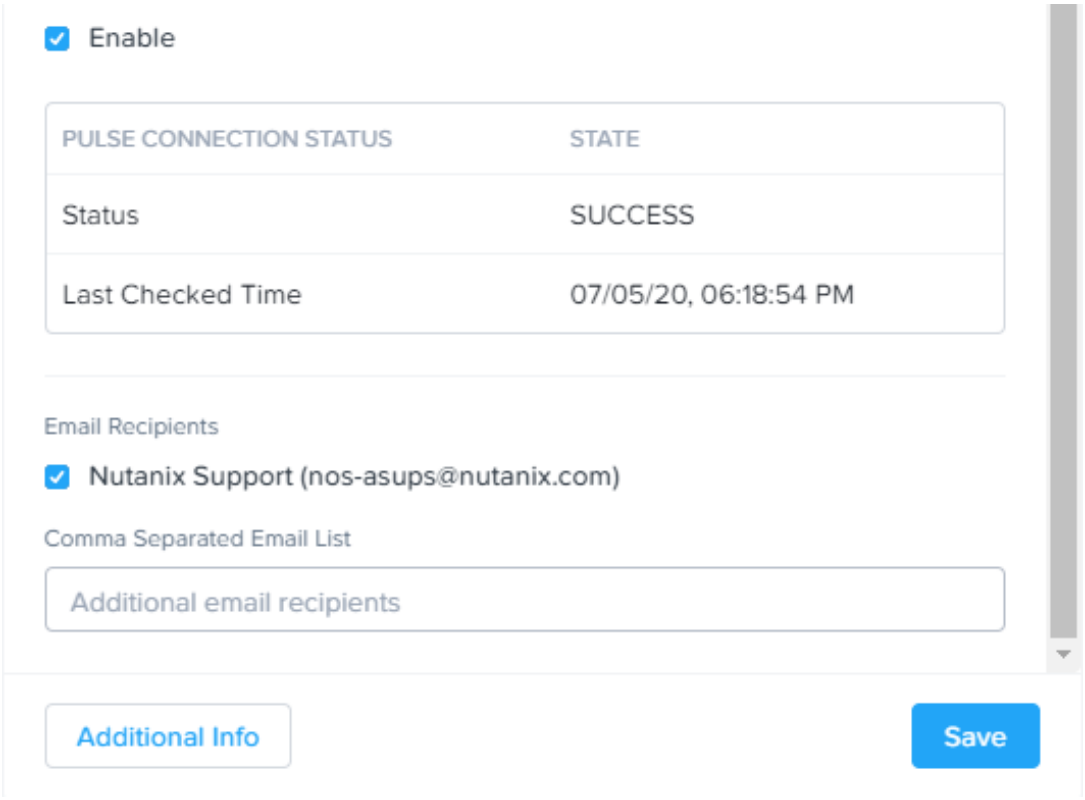

Figure 536: Pulse Additional Info

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.

Pulse Health Monitoring Data Collection

The following categories serve as a general overview of the types of information that Pulse gathers from the clusters. Note that some of this information may be anonymized depending on your settings. This list is not exhaustive; for more details about the Pulse information your clusters send to Nutanix, contact Nutanix Support.

Table 773: Pulse Data Collection

Entity	Data Collected
Cluster	• Cluster name (may be anonymized) • Uptime • AOS version • Cluster ID • Block serial number • HW model • Cluster IOPS • Cluster latency • Cluster memory

Entity	Data Collected
Hardware **Note:** In this context, hardware can include nodes, blocks, boards, disks, BMCs, fans, DIMMs, BIOS, CPUs, NICs, storage controllers, and power supplies.	• Model number • Serial number • Part number • Block number • Node UUID • Type • Size • Version • Name (may be anonymized) • Manufacturer • Status • Memory (size) • Hypervisor type • Hypervisor version • Firmware version • Disk type • Disk model • Disk capacity • Node temperature • Network interface model • SATADOM firmware • PSU status • Node location • IPMI version • Fan RPM • Component location • DIMM bank connection • Clock speed • DIMM temperature • BIOS release date • BIOS ROM size • CPU signature • CPU core count • CPU cores enabled • CPU thread count • CPU temperature

Entity	Data Collected
Storage Pool	• Name (may be anonymized) • Capacity (logical used capacity and total capacity) • IOPS and latency
Container	• Container name (may be anonymized) • Capacity (logical used and total) • IOPS and latency • Replication factor • Compression ratio • Deduplication ratio • Inline or post-process compression • Inline deduplication • Post-process deduplication • Space available • Space used • Erasure coding and savings
Controller VM (CVM)	• Details of logs, attributes, and configurations of services on each CVM • CVM memory • vCPU usage • Uptime • Network statistics • IP addresses (may be anonymized)

Entity	Data Collected
VM	• Name (may be anonymized) • VM state • vCPU • Memory • Disk space available • Disk space used • Number of vDisks • Name of the container that contains the VM (may be anonymized) • VM operating system • IOPS • Latency • VM protection status • Management VM (yes or no) • I/O pattern (read, read/write, random, sequential) • IP address (may be anonymized)
Disk Status	• Performance stats • Usage
Hypervisor	• Hypervisor software and version • Uptime • Installed VMs • Memory usage • Attached datastore
Datastore	• Usage • Capacity • Name
Protection Domains	• Name (may be anonymized) • Count and names of VMs in each protection domain

Entity	Data Collected
Gflags	• Key and value • State (set) • Node ID • Service name • Time of modification
Feature	• Feature ID • Name • State (enabled or disabled) • Mode
License	License type (Starter, Ultimate, or Pro)
Alerts	• Alert ID • Type • Severity • Resolution status • Acknowledgement status • Impact type • Message • Creation time • Modification time
Tasks	• Task ID • Operation type • Status • Entities • Message Completion percentage • Creation time • Modification time
Logs	• Component • Timestamp • Source file name • Line number • Message

Entity	Data Collected
Nutanix Services	Service-specific metrics

Network Configuration for Pulse Health Monitoring

When Pulse Health Monitoring is enabled, all telemetry data streams to the Insights service hosted at insights.nutanix.com over port 443. Therefore, telemetry traffic must be allowed to reach this destination.

The source of Pulse data varies depending on the deployment:

- If Prism Central is deployed, Pulse routes all the information from every node in a managed cluster through this Prism Central. You can opt for setting up a Prism Central Proxy and allowing traffic from Proxy Server IPs to insights.nutanix.com over port 443.
- If Prism Central is not deployed, Pulse routes this information from each CVM for every node in the cluster. You can opt for setting up a Prism Element Proxy and allowing traffic from Proxy Server IPs to insights.nutanix.com over port 443.

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 each registered Prism Element cluster.
- Prism Central and each Prism Element cluster node are running NCC 3.5.2 or later.
- The Prism Element clusters do not have direct internet connectivity, and you have not con#gured an HTTP proxy on Prism Element clusters registered to this Prism Central instance. If your Prism Element clusters are con#gured to use an HTTP proxy or direct internet connectivity, 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 509 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 (Prism Central) on page 1011). However, you can also create a support case directly from Prism Central. To create a support case or view information about your open cases, do the following:

> **Note:** A support portal connection is required before you can create a case (see Licensing on page 31). In addition, at least one cluster must be registered with Prism Central. Also, this feature is available only to customers who get support directly from Nutanix (such as NX and SX model customers), not customers who get initial support from third parties.

Procedure

1. Click the question mark icon in the main menu and select **Create Support Case** from the pull-down list.

   The **Create new support case** page appears. One of the following appears on this page:

   - If there are no open cases currently, a blank page with fields for creating a case appears.
   - If you have one or more open support cases, summary information about those cases appears (see Viewing Case Status on page 1010). 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.

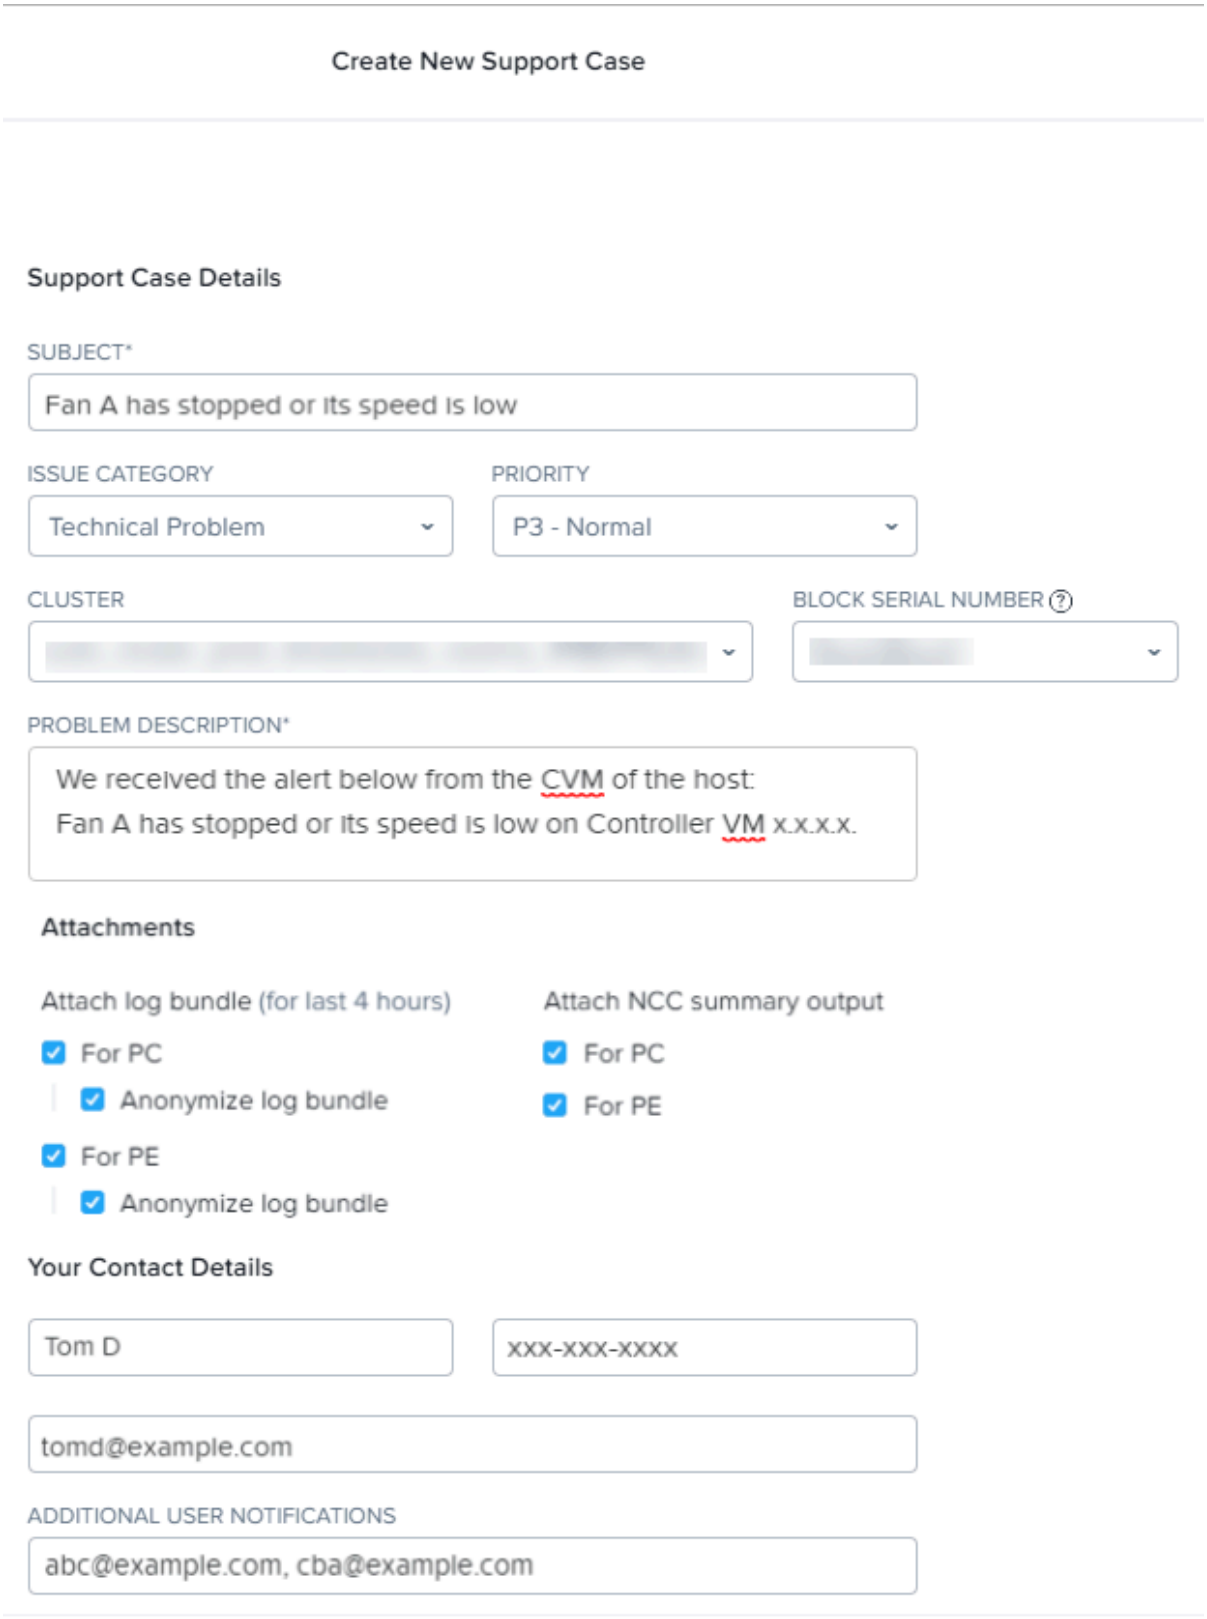

Figure 537: Create New Support Case Page

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 244). 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 1010).

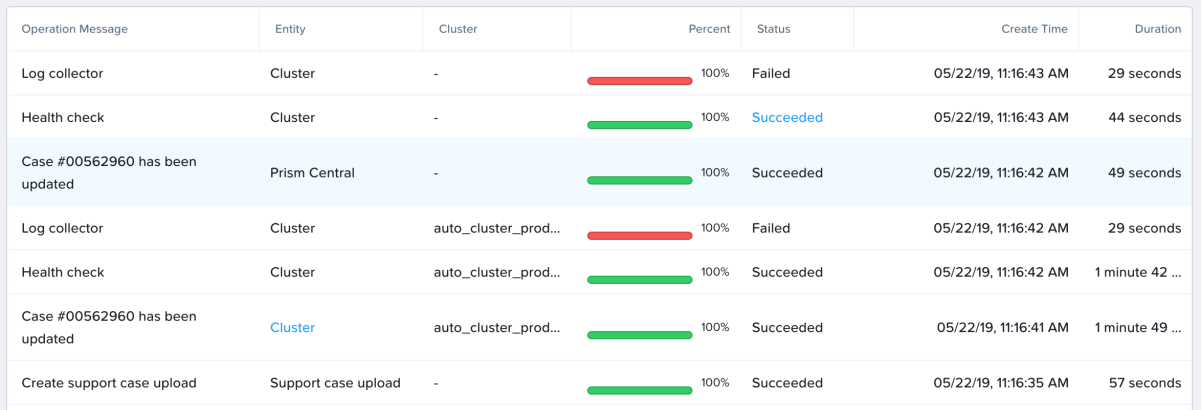

Operation Message	Entity	Cluster	Percent	Status	Create Time	Duration
Log collector	Cluster	-	100%	Failed	05/22/19, 11:16:43 AM	29 seconds
Health check	Cluster	-	100%	Succeeded	05/22/19, 11:16:43 AM	44 seconds
Case #00562960 has been updated	Prism Central	-	100%	Succeeded	05/22/19, 11:16:42 AM	49 seconds
Log collector	Cluster	auto_cluster_prod...	100%	Failed	05/22/19, 11:16:42 AM	29 seconds
Health check	Cluster	auto_cluster_prod...	100%	Succeeded	05/22/19, 11:16:42 AM	1 minute 42 ...
Case #00562960 has been updated	Cluster	auto_cluster_prod...	100%	Succeeded	05/22/19, 11:16:41 AM	1 minute 49 ...
Create support case upload	Support case upload	-	100%	Succeeded	05/22/19, 11:16:35 AM	57 seconds

Figure 538: Tasks Dashboard for Prism Central (create case tasks)

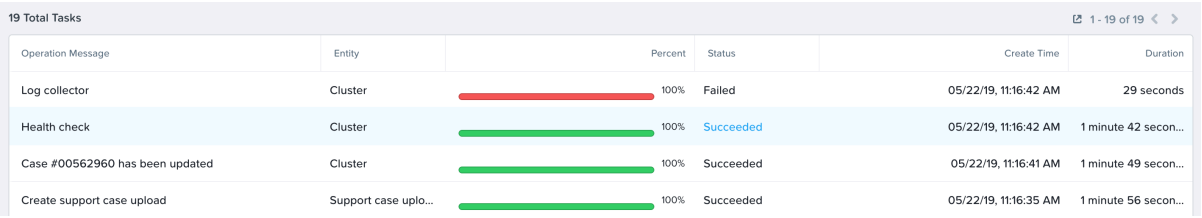

19 Total Tasks — 1 - 19 of 19

Operation Message	Entity	Percent	Status	Create Time	Duration
Log collector	Cluster	100%	Failed	05/22/19, 11:16:42 AM	29 seconds
Health check	Cluster	100%	Succeeded	05/22/19, 11:16:42 AM	1 minute 42 secon...
Case #00562960 has been updated	Cluster	100%	Succeeded	05/22/19, 11:16:41 AM	1 minute 49 secon...
Create support case upload	Support case uplo...	100%	Succeeded	05/22/19, 11:16:35 AM	1 minute 56 secon...

Figure 539: 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.

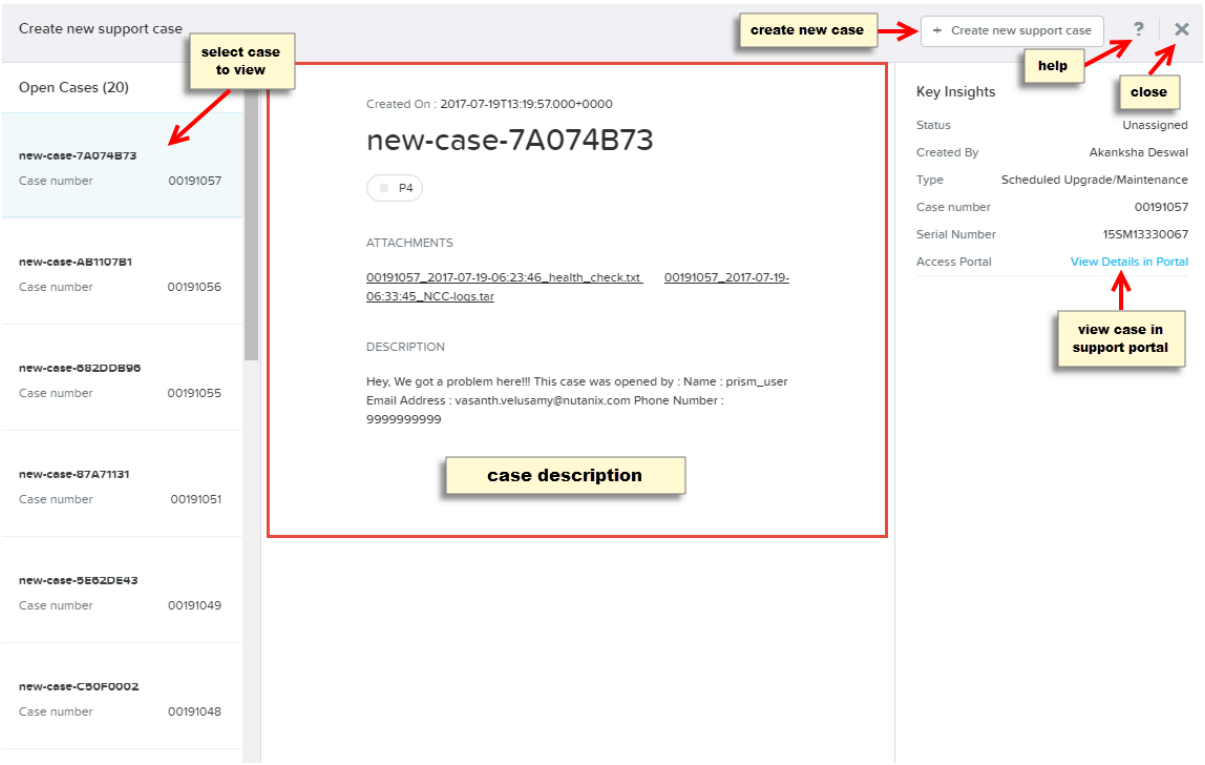

Figure 540: Create New Support Case Page (open cases)

Accessing the Nutanix Support Portal (Prism Central)

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 774: 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.

Category	Option	Description
	.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.
	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.

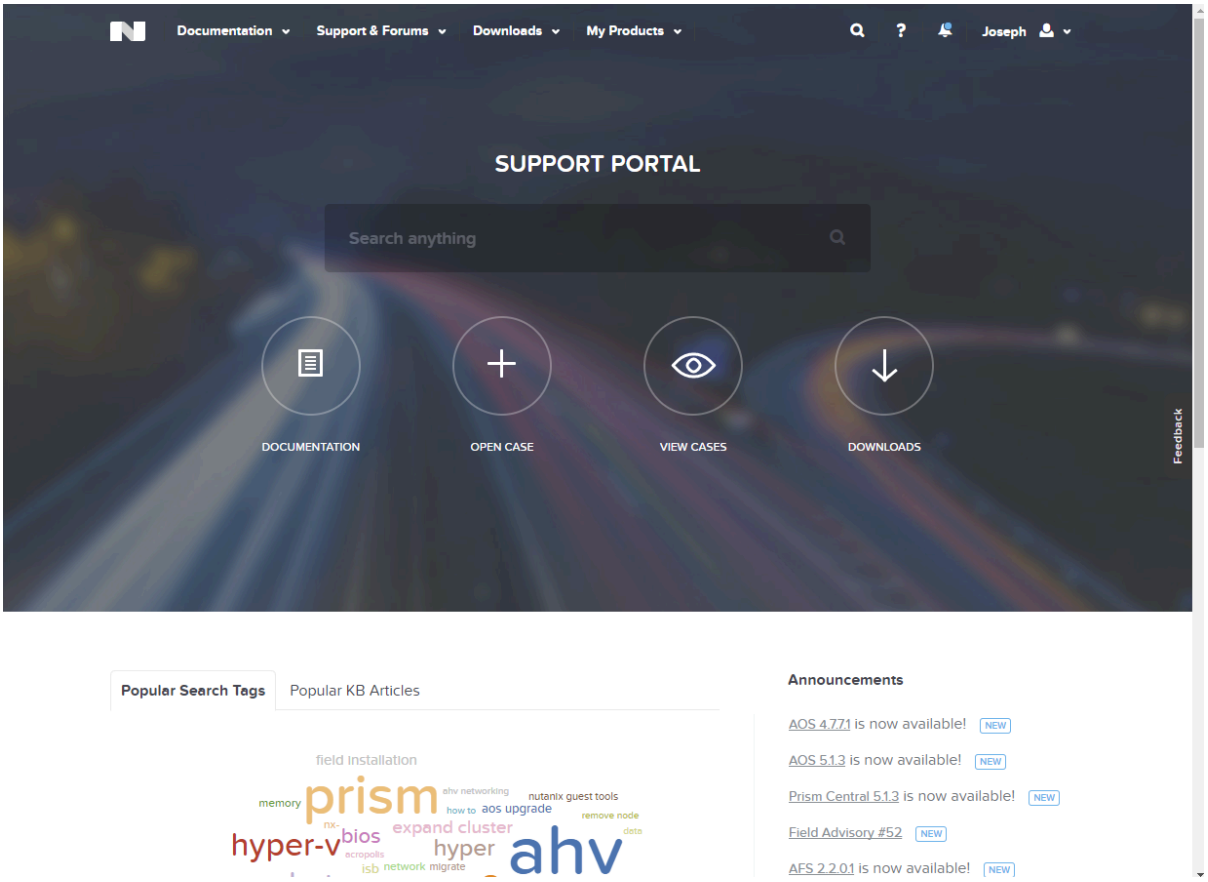

Figure 541: Nutanix Support Portal

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. Only an admin user can access the REST API Explorer.

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.

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 (Prism Central) on page 1016).
- Nutanix customer support portal (see Accessing the Nutanix Support Portal (Prism Central) on page 1011).
- Nutanix community forum (see Accessing the Nutanix Next Community on page 1018).
- Glossary of terms (see Glossary on page 1018).

Accessing Online Help (Prism Central)

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.

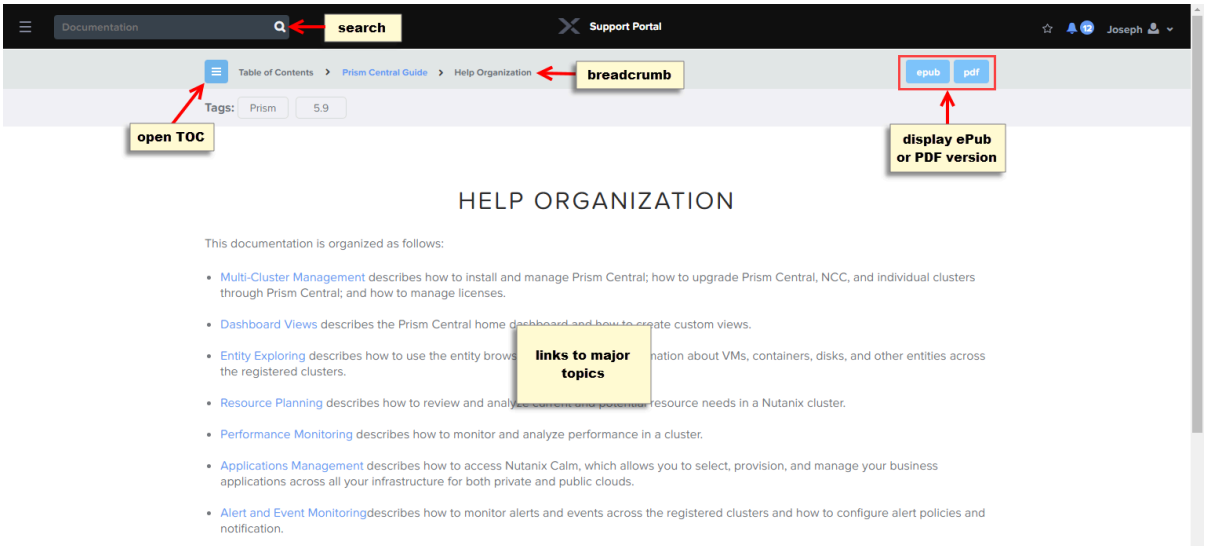

Figure 542: Help Organization Page

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.

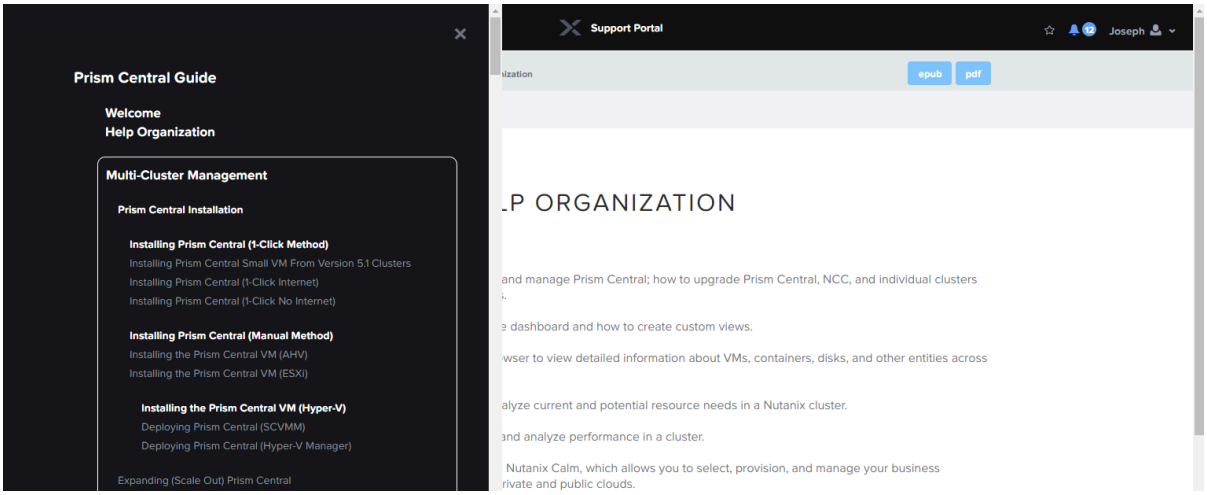

Figure 543: Table of Contents Pane

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.

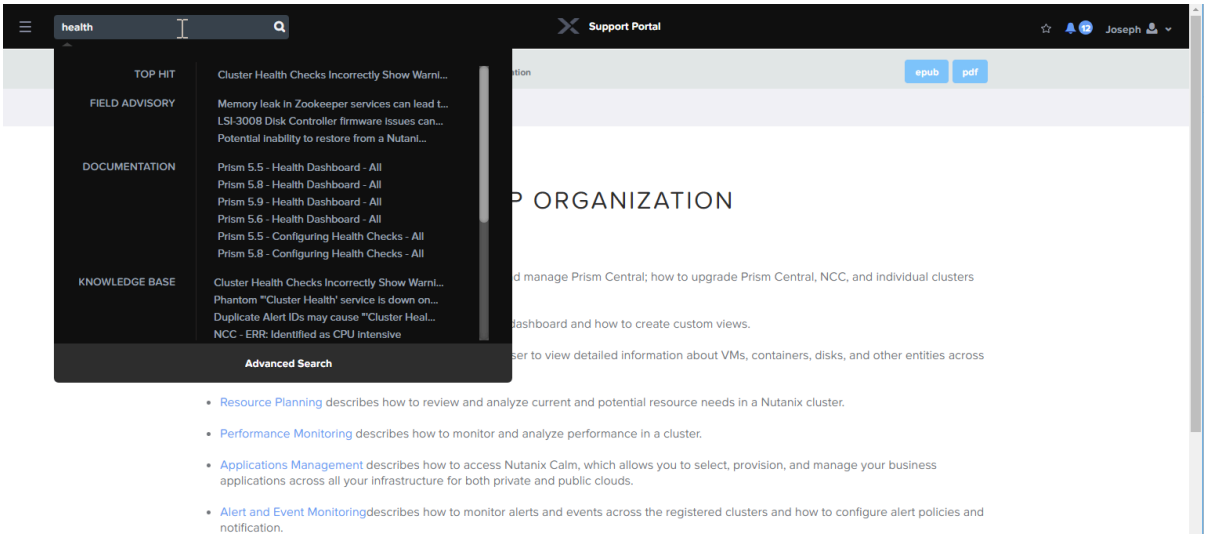

Figure 544: Search Results Example

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 (see Main Menu (Prism Central) 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.

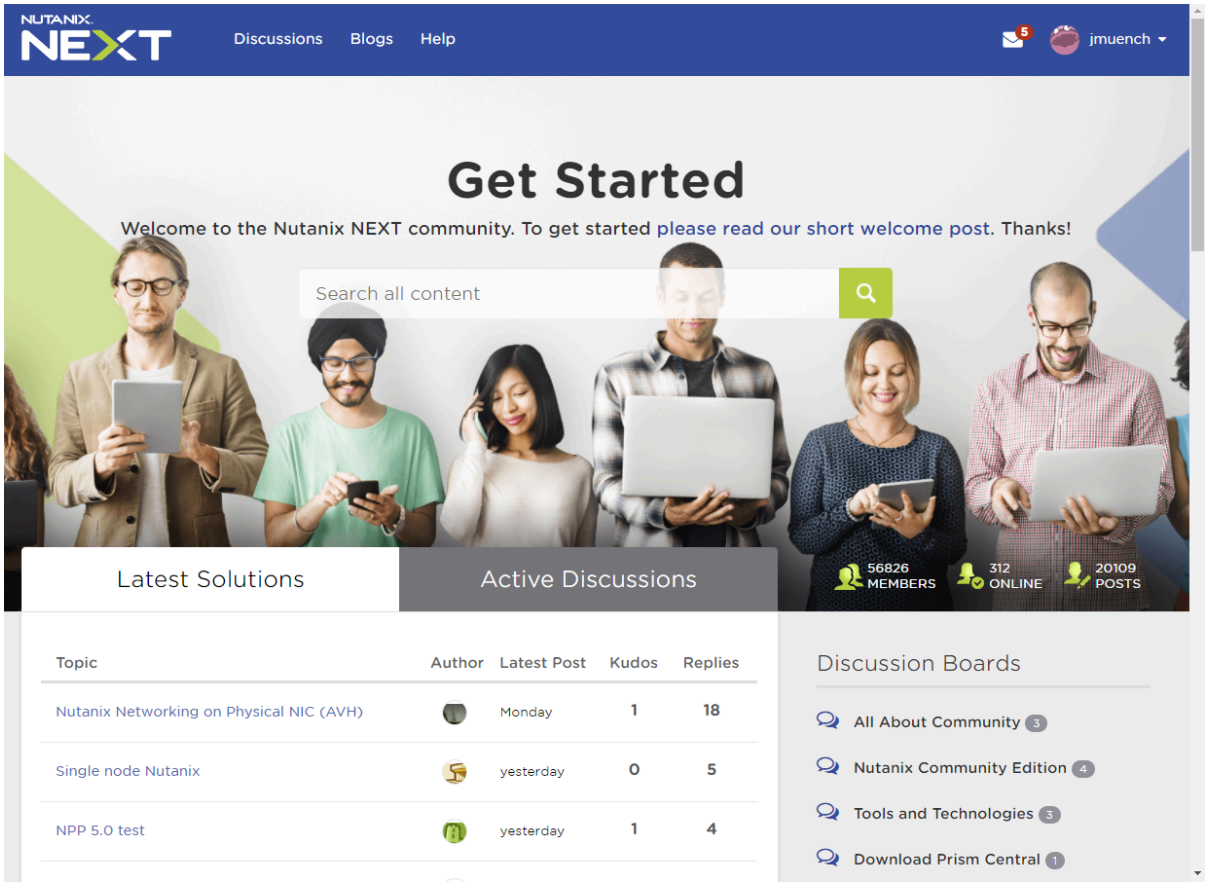

Figure 545: Next Community Screen

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 2024 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.