

Prism Central Admin Center Guide

Prism pc.2024.2

April 29, 2025

Contents

About this Publication.....	5
Admin Center Overview.....	6
Application Switcher (Prism Central).....	7
Prism Central Licensing.....	8
Marketplace.....	9
Enabling Marketplace.....	12
Viewing Marketplace App Details.....	13
Searching or Filtering Marketplace Items.....	14
Nutanix Apps Deployment.....	15
Deploying Self-Service.....	16
Deploying Files.....	16
Deploying Foundation Central.....	17
Deploying Kubernetes Management.....	17
Deploying Objects.....	18
Deploying Database Service.....	20
Deploying Move.....	21
Preferred Partner App Deployment.....	22
Deploying OpenShift.....	22
Deploying Custom Apps from the Marketplace.....	25
Executing a Runbook from the Marketplace.....	28
Cloning a Marketplace Item.....	29
My Apps.....	30
Nutanix Apps Management.....	33
Registering Database Service or Move to My Apps.....	34
Other Apps Management.....	35
Tabs for Other Apps Management.....	37
Running a User-Level Action.....	48
Update VM Configurations of a Running App.....	48
Snapshot and Restore Overview.....	55
Running a System-Level Action.....	66
Updating Actions and Credentials of an App.....	70
Creating an Image on a Nutanix Platform.....	73
Cloning an App.....	75
Deleting an App.....	75
Accessing Web SSH Console.....	76
Project Management.....	78
Projects Summary View.....	80
Project Details View.....	82
Project Dashboard Tab.....	82
Project Environments Tab.....	87
Project Infrastructure Tab.....	88

Project Policies Tab.....	88
Project Tunnels Tab.....	88
Project Users & Groups Tab.....	89
Project Workloads Tab.....	89
Project Usage Tab.....	91
Creating a Project.....	91
Adding Users to a Project.....	95
Infrastructure in Projects.....	99
Adding Infrastructure in a Self-Service-Enabled Prism Central.....	100
Adding Infrastructure in a Self-Service-Disabled Prism Central.....	105
Modifying a Project.....	109
Deleting a Project.....	110
Environments in Projects.....	110
Configuring Environments in a Project.....	110
Configuring Check Log-in for your Environment.....	133
Support for Multiple Credential.....	134
Adding Credentials to the Environment.....	134
Quota Policy Overview.....	135
Managing Quota Limits for Projects.....	136
Creating a Snapshot Policy.....	138
Identity and Access Management.....	142
Security and User Management (Prism Central).....	142
User Management.....	142
Roles Management.....	143
LCM View.....	146
Admin Center Settings Options.....	147
Application Log Archive.....	148
Downloading Application Log Archive.....	148
Policy Engine Overview.....	149
Enabling Policy Engine.....	149
Enabling Policy Engine at a Dark Site.....	152
Viewing Policy Engine VM Details.....	153
Disabling Policy Checks.....	153
Configuring Rsyslog for Policy VM.....	153
Best Practices to Shut Down Policy Engine.....	160
Pulse Health Monitoring.....	160
Pulse Configuration.....	162
Mask Entity Names and IP Addresses.....	165
Pulse Health Monitoring Data Collection.....	166
Network Configuration for Pulse Health Monitoring.....	172
Prism Central Proxy for Pulse Data.....	172
Prism Central Configuration When a Cluster Uses Proxy Servers.....	173
Who Needs to Use the Allowlist Method.....	173
Configuring an HTTP Proxy.....	175
Allowing Prism Central and Its Managed Clusters (nCLI).....	177
Configuring Name Servers for Prism Central.....	178
Configuring NTP Servers for Prism Central.....	179
Configuring SNMP on Prism Central.....	180
Configuring Alert Emails in Prism Central.....	185
Configuring an SMTP Server for Prism Central.....	187

Configuring Syslog Monitoring.....	188
Syslog Modules.....	190
Setting the Prism Central Default Landing Page.....	194
Changing the Prism Central Language Settings.....	194
Internationalization (i18n) (Prism Central).....	196
Localization (L10n) (Prism Central).....	196
Configuring Prism Central UI Settings.....	197
Configuring the Prism Central Welcome Banner.....	199
Glossary.....	201
Copyright.....	202

ABOUT THIS PUBLICATION

This document provides information about how to configure and manage administrative tasks using Admin Center in Prism Central.

- For information on all the documents applicable for Prism Central, see [Prism Central Documentation Portfolio](#) in the *Prism Central Infrastructure Guide*.
- To access other Nutanix documents, see [Nutanix support portal](#).

ADMIN CENTER OVERVIEW

Admin Center facilitates you to configure and manage common administrative tasks that are applicable to the platform and various Nutanix apps.

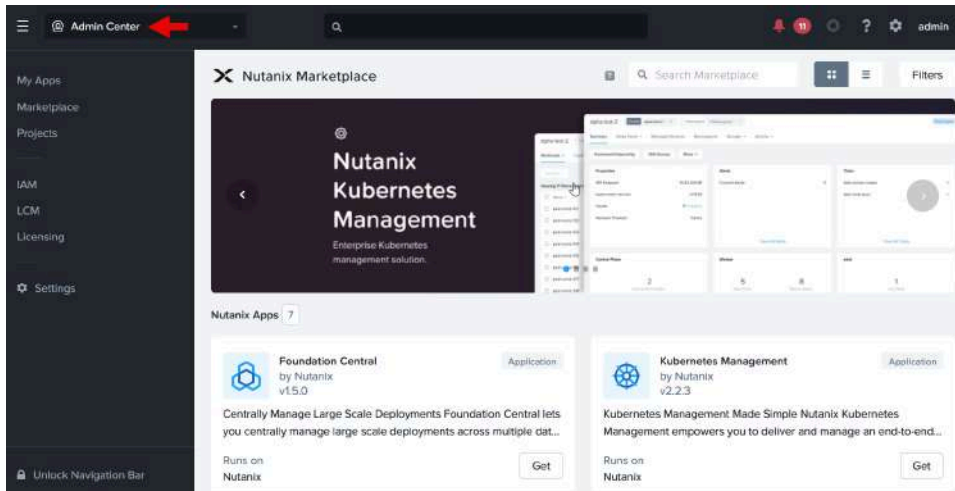


Figure 1: Admin Center

Admin Center Entities

You can use Admin Center to manage the following configurations:

- **Define Projects and Various App Management Features Within Projects**

Projects provide logical groupings of user roles for managing resource utilization within your organization. The project construct helps you define users and groups for marketplace app deployment, networks to use for app deployment, VM specifications and deployment options, quota policies, snapshot policies, and so on.

The project configuration options differ depending on whether you have enabled Self-Service in your Prism Central or not. In a Prism Central where Self-Service is disabled, you can configure projects with basic setups such as add users, add local Nutanix account, and manage usage and workloads. Application management features such as environment configuration, policy configuration, and cloud infrastructure setups are possible only when you deploy Self-Service in your Prism Central. For detailed information on projects, see [Project Management](#) on page 78.

- **Deploy Nutanix Apps and Other Supported Apps**

Nutanix Marketplace facilitates provisioning of Nutanix Apps such as Self-Service, Files, Objects, Foundation Central, Move, Database Service, and Kubernetes Management and preferred partner apps, such as OpenShift. Some of these apps used to be under the **Services** entities in pc.2022.9 and earlier versions.

The availability of apps in the Marketplace differs depending on whether you have deployed Self-Service in your Prism Central or not. When Self-Service is deployed, you can also access and deploy the preconfigured hybrid cloud apps, custom apps, and runbooks from the Other Apps section of the Marketplace.

For detailed information on Marketplace, see [Marketplace](#) on page 9.

- **Manage Deployed Nutanix Apps and Other Apps from a Common Workspace**

The My Apps entity of Admin Center provides a single workspace to manage all the apps that you deployed from the Marketplace. For detailed information on My Apps, see [My Apps](#) on page 30.

- **Manage Users, Roles, and Authentication Mechanisms**

Identity and Access Management (IAM) facilitates you to view details of local and directory-specified users, view system-defined roles and custom roles, and create custom roles. The [Security Management Using Prism Central](#) section in the *Security Guide* covers different aspects of users, roles, and authentication management, including SAML/IDP supported providers.

- **Perform inventory and Software Updates**

The life cycle manager (LCM) facilitates you to view information about the current inventory and lets you update the versions as needed. For detailed information on LCM, see the [Life Cycle Manager Guide](#).

- **Manage Licensing**

You can apply licenses to enable a variety of features. For detailed information, see the [Nutanix License Manager Guide](#).

- **Manage Global Prism Central Settings Applicable Across Prism Central**

The setting options in Admin Center are common across multiple Nutanix apps that you deploy and access using Prism Central. You can enable policy engine and configure various app management, networking, alerts, and appearance related settings. For more information on Admin Center settings, see [Admin Center Settings Options](#) on page 147.

Admin Center Navigation Bar

You get the Admin Center navigation bar when you select Admin Center in the Application Switcher. The Admin Center navigation bar has the following entities.

- My Apps
- Marketplace
- Projects
- IAM
- LCM
- Licensing
- Settings

You can click an entity in the navigation bar to view the summary or manage configurations associated with that entity. For example, you can click **Projects** to create a project, add users, add infrastructure, and configure various app management components.

You can also lock the navigation bar so that it remains open when you switch between different entities within the Admin Center.

Application Switcher (Prism Central)

Prism Central provides a centralized environment to procure multiple Nutanix Apps such as Self-Service, Files, Database Service, Objects, Kubernetes Management, and Foundation Central and manage these apps from a single workspace. The Application Switcher facilitates you to seamlessly switch between the Nutanix apps that you have enabled and have access to in Prism Central.

For more information on different Nutanix apps, see [Application Switcher Function](#) in the *Prism Central Infrastructure Guide*.

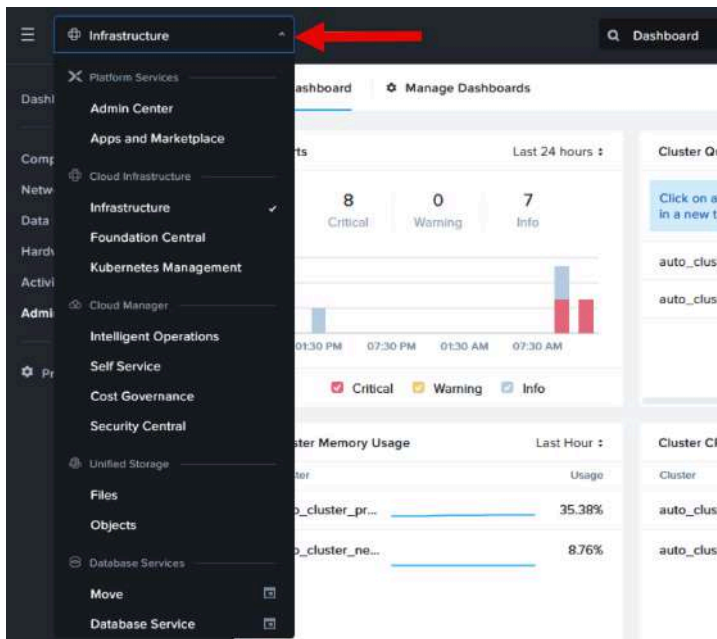


Figure 2: Application Switcher

Access to Apps in the Application Switcher

In a new deployment of Prism Central, the Application Switcher displays the Admin Center, Infrastructure, Apps and Marketplace, Cost Governance, and Security Central for users with the Nutanix admin role. You can view other Nutanix Apps in the Application Switcher after enabling them from the Marketplace. For more information, see [Marketplace](#) on page 9.

If you have upgraded to the current version of Prism Central, you can view all the Nutanix apps that you enabled in the previous version along with the Admin Center, Infrastructure, Apps and Marketplace, Cost Governance, and Security Central.

For users with the non-admin role, Application Switcher displays the apps based on the access policies. For example, if you are a user with system-defined consumer role, then you can view only Infrastructure, Apps and Marketplace, and any other Nutanix apps that you have access to.

Note: You cannot add any custom apps to the Application Switcher.

Prism Central Licensing

Nutanix provides licenses you can apply to enable a variety of features.

The Prism Element web console and Nutanix Support Portal provide the most current information on your licenses. For more information on licenses, see the [License Manager Guide](#).

MARKETPLACE

Nutanix Marketplace facilitates provisioning for authorized consumers to discover, procure, and deploy approved apps. The Marketplace has the following basic functions.

- It enables deployment and management of Nutanix Apps, such as Self-Service, Files, Objects, Foundation Central, Database Service, Move, and Kubernetes Management. As a Prism Central customer, you do not require any additional license to use the new Marketplace capabilities and access Nutanix apps.
- When Self-Service is deployed and licensed, Marketplace also provides users a workspace to instantly consume the app resources that are specifically tailored for their needs with a variety of provisioning options. Such apps can be deployed on-premise, in public clouds, or both, enabling IT services to implement a multi-cloud strategy. These apps are consumed in a repeatable way, saving time and effort used for routine app provisioning and management.

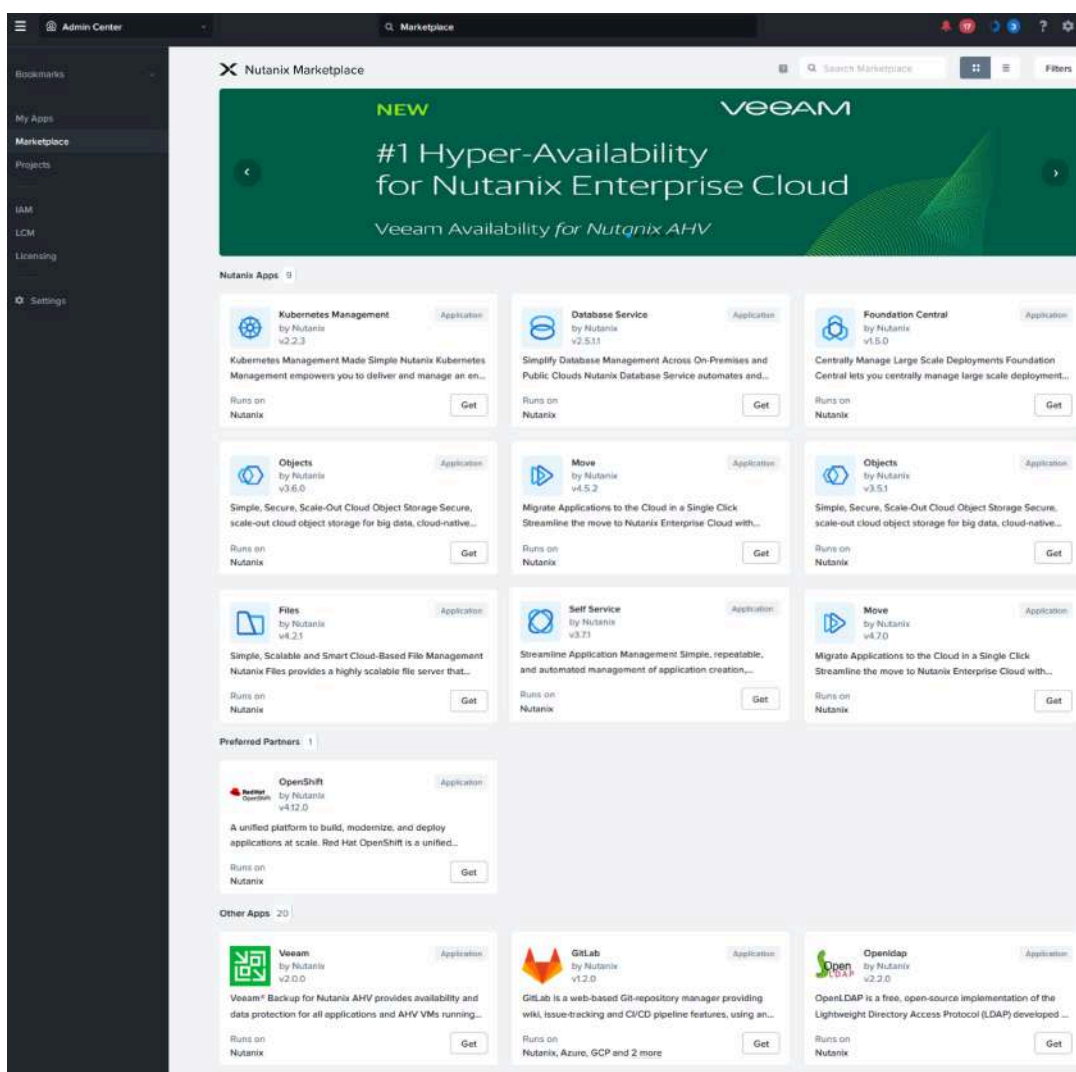


Figure 3: Marketplace

Marketplace App Categories

The following table lists the different categories of apps in the Marketplace and their availability depending on whether Self-Service is deployed in your Prism Central instance or not.

Table 1: Marketplace Items

Marketplace Item	Description	When Self-Service is Not Deployed	When Self-Service is Deployed
Nutanix Apps	<p>These apps are developed and certified by Nutanix. These apps include Self-Service, Files, Objects, Foundation Central, Move, Database Service, and Kubernetes Management. These apps run on the Nutanix platform. You might have to apply for necessary licenses to use these apps. For more information, see the Nutanix License Manager Guide.</p> <p>For information on how to deploy Nutanix apps, see Nutanix Apps Deployment on page 15.</p>	Yes	Yes
Preferred Partner Apps	<p>These apps are developed and certified by Nutanix in collaboration with the preferred partners. These apps are pre-seeded for deployment on the Nutanix platform. For example, OpenShift.</p> <p>For information on how to deploy Preferred Partner apps, see Preferred Partner App Deployment on page 22.</p>	Yes	Yes
Hybrid Cloud Apps (Other Apps)	<p>These are third-party or open-source apps that are developed and certified by Nutanix. This includes apps such as LAMP, Icinga, CouchDB, Docker Swarm, and so on.</p> <p>You can deploy these apps not just on Nutanix but also on VMware or public clouds.</p> <p>These apps are available for deployment only when Self-Service is deployed and licensed in your Prism Central instance.</p> <p>For information on how to deploy Hybrid Cloud apps, see the Self-Service Marketplace Items Guide.</p>	No	Yes

Marketplace Item	Description	When Self-Service is Not Deployed	When Self-Service is Deployed
Custom Apps (Other Apps)	<p>These apps are configured through app blueprints in Self-Service and are approved and published in the Marketplace by an administrator for the end-user consumption. For information on custom app configuration, see Self-Service Blueprints in the <i>Self-Service Administration and Operations Guide</i>.</p> <p>For information on approval and publishing of custom apps, see Marketplace Manager in the <i>Self-Service Administration and Operations Guide</i>.</p> <p>These apps are available for deployment only when Self-Service is deployed and licensed in your Prism Central instance.</p> <p>For information on how to deploy custom apps, see Deploying Custom Apps from the Marketplace on page 25.</p>	No	Yes
Runbooks (Other Apps)	<p>Runbooks are collection of tasks that are configured to run sequentially at different endpoints. Runbooks are configured in Self-Service and are approved and published in the Marketplace by an administrator for consumption.</p> <p>For more information, see the Runbooks section in the <i>Self-Service Administration and Operations Guide</i>.</p> <p>Runbooks are available in the Marketplace for execution only when Self-Service is deployed and licensed in your Prism Central instance.</p> <p>For information on how to execute a runbook, see Executing a Runbook from the Marketplace on page 28.</p>	No	Yes

Marketplace App Access and Availability

The first step to use Marketplace for app deployment is to enable the Marketplace in your Prism Central instance. For more information, see [Enabling Marketplace](#) on page 12.

While Nutanix apps and Preferred Partner apps are always available in the Marketplace for consumption, the custom apps, hybrid cloud apps, and runbooks are available only when:

- Self-Service is deployed and licensed in your Prism Central instance.
- The administrator has approved and published the apps or runbooks in the Marketplace for consumption. For more information, see the [Marketplace Manager](#) section in the *Self-Service Administration and Operations Guide*.

Marketplace Permissions

Important:

- Only a Prism Admin can view the Nutanix App and Preferred Partner App details or deploy the apps from the Marketplace.
- The IAM interface has been revamped to provide fine-grained RBAC so that custom roles can be created with granular operations or permissions. To access all marketplace blueprint items and their operations, the role must have access to the Marketplace Item entity and other dependent entities such as Marketplace Icon, Launch Blueprint, Import Blueprint, and so on. For information on how to create roles and add authorization policy, see the [Security Management using Prism Central \(PC\)](#) section of the *Security Guide*.

The following table lists the different Marketplace operations for hybrid cloud apps, custom apps, and runbooks and the permissions for different users to perform those operations:

Table 2: Permissions for Hybrid Cloud Apps, Custom Apps, and Runbooks

Operation	Prism Central Admin	Project Admin	Developer	Consumer	Operator
Deploy Apps	Yes	Yes	Yes	Yes	No
Clone Apps	Yes	Yes	Yes	No	No
View App List	Yes	Yes	Yes	Yes	No
Search or Filter Apps	Yes	Yes	Yes	Yes	No
View App Details	Yes	Yes	Yes	Yes	No
Execute Runbooks	Yes	Yes	Yes	Yes	No
View Runbooks	Yes	Yes	Yes	Yes	No
Filter Runbooks	Yes	Yes	Yes	Yes	No
View Runbook Details	Yes	Yes	Yes	Yes	No
Clone Runbooks	Yes	Yes	Yes	No	No

Enabling Marketplace

Enabling Marketplace is a prerequisite for app deployment from Nutanix Marketplace or app management in My Apps. Only a Prism Admin can enable Marketplace.

Before you begin

- Marketplace enablement requires an additional 2 GB of memory for a small Prism Central deployment and 4 GB of memory for a large Prism Central deployment. Ensure that your Prism Central has required resources to enable Marketplace in the same cluster, and the Prism Central is registered to the host Prism Element.

- A unique data service IP address should be configured in the Prism Element web console cluster that is running on Prism Central. For information on configuring data service IP address, see the [Modifying Cluster Details](#) in the *Prism Element Web Console Guide*.
- Nutanix recommends you to observe the impacts before you change the virtual IP address and Data Services IP address. For more information, see the [Virtual IP Address Impact](#) and [iSCSI Data Services IP Address Impact](#) sections in the *Prism Element Web Console Guide*. See the product-specific documents for detailed information on how to handle the changes for virtual IP address and Data Services IP address.

Note:

- Marketplace is supported only on AHV and ESXi hypervisors on a Nutanix cluster and is not supported on Hyper-V hypervisor.
- Marketplace is not supported on vSphere Essentials edition because vSphere Essentials edition does not support hot-pluggable virtual hardware.
- Marketplace is not supported on x-small Prism Central deployments.

About this task

Note: Once you enable Marketplace after an upgrade, you must perform the LCM inventory to see the correct version of Nutanix Apps on the My Apps page.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. Click **Enable Marketplace**.

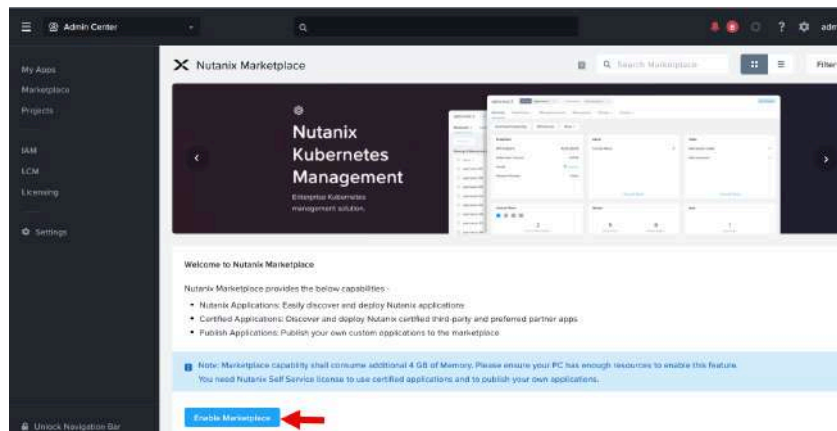


Figure 4: Enable Marketplace

Viewing Marketplace App Details

You can view the Marketplace app overview, version, and the actions included (available only for the apps in the Other Apps category) on the app details page.

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. Click **Get** for the application for which you want to view details.
5. On the app details page, view details such as the overview, version, and the actions included for the app.

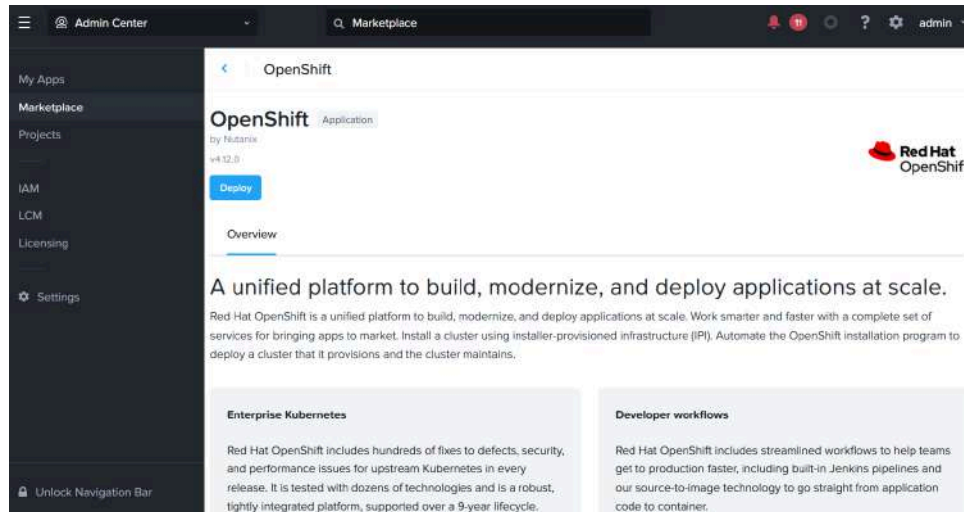


Figure 5: App Details

Searching or Filtering Marketplace Items

You can easily search a Marketplace item or filter to get a list of items based on a criteria.

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. To search an app or runbook, enter the name of the app or runbook that you want to search in the **Search marketplace** field.

5. To filter Marketplace items, click **Filters** and do the following.

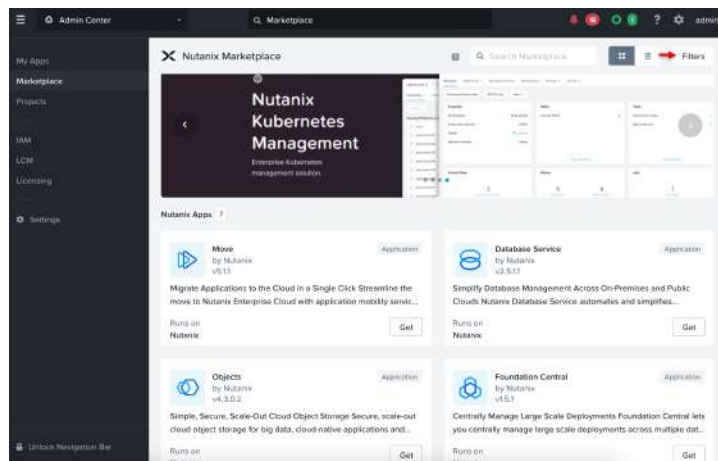


Figure 6: Marketplace Filters

a. Select the categories of items. Your options are:

- **Nutanix**
- **Preferred_Partners**
- **DevOps**
- **Networking**
- **Databases**
- **Containers**
- **Bi-Productivity**
- **Backup**

b. Select item types. You can select **Application**, **Runbook**, or both.

c. Select item sources. You can select **Global Store**, **Local**, or both.

Marketplace keeps filtering items as and when you select the filter options.

Nutanix Apps Deployment

As a Prism Admin, you can enable the following Nutanix apps from Admin Center.

- Self-Service. For information on the Self-Service deployment, see [Deploying Self-Service](#) on page 16
- Files. For information on the Files deployment, see [Deploying Files](#) on page 16
- Foundation Central. For information on the Foundation Central deployment, see [Deploying Foundation Central](#) on page 17
- Kubernetes Management. For information on the Kubernetes Management deployment, see [Deploying Kubernetes Management](#) on page 17
- Objects. For information on the Objects deployment, see [Deploying Objects](#) on page 18

- Database Service. For information on the Database Service deployment, see [Deploying Database Service](#) on page 20.
- Move. For information on the Move deployment, see [Deploying Move](#) on page 21.

Deploying Self-Service

About this task

Self-Service is integrated into Prism Central and does not require you to deploy any additional VMs. To start using Self-Service, you have to deploy the Self-Service app from the Marketplace.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.
- Ensure that you meet the prerequisites to enable Self-Service. For more information, see the [Prerequisites to Deploy Self-Service](#) section in the *Self-Service Administration and Operations Guide*.

Note:

If the Prism Element web console is not registered from a Prism Central and the app blueprints have a subnet, image, or VMs on the Prism Element web console, the Self-Service functionality is impacted.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Self-Service app.
5. On the app details page of Self-Service, click **Deploy**.

After deployment, the Self-Service app appears on the **My Apps** page. You can select **Self-Service** in the Application Switcher to access and use the app. See the [Self-Service Administration and Operations Guide](#) for information on configuring and using Self-Service.

Deploying Files

About this task

Nutanix Files provides file services to clients for file sharing across user work stations from a centralized and protected location.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.
- For deployment prerequisites, see the [Deploying Files on the Files Manager](#) section in the *Files Manager User Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Files app.
5. On the app details page, click **Deploy**.

After deployment, the Files app appears on the **My Apps** page. You can select **Files** in the Application Switcher to access and use the app. See the [Nutanix Files Manager Guide](#) for information on configuring and using Files in Prism Central.

Deploying 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.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.
- For deployment prerequisites, see the [Set Up Foundation Central](#) section in the *Foundation Central Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Foundation Central app.
5. On the app details page of Foundation Central, click **Deploy**.

After Foundation Central is enabled, the Foundation Central app appears on the **My Apps** page. You can select **Foundation Central** in the Application Switcher to access and use the app. See the [Foundation Central Documentation](#) for information on configuring and using Foundation Central.

Deploying Kubernetes Management

About this task

Nutanix Kubernetes Management 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.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.

- For more information about the deployment requirements, see the [Requirements](#) section in the *Nutanix Kubernetes Engine Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Kubernetes Management app.
5. On the app details page, click **Deploy**.

After deployment, the Kubernetes Management app appears on the **My Apps** page. You can select **Kubernetes Management** in the Application Switcher to access and use the app. See the [Nutanix Kubernetes Engine Guide](#) for information about configuring and using Kubernetes.

Deploying Objects

About this task

Nutanix Objects allows you to create and manage object stores in a cluster.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.
- For deployment prerequisites, see the [Deployment and Network Prerequisites](#) section in the *Objects User Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Objects app.
5. On the app details page, click **Deploy**.

6. Select **Objects** in the Application Switcher.
A welcome page appears with the prerequisite details to create an object store.

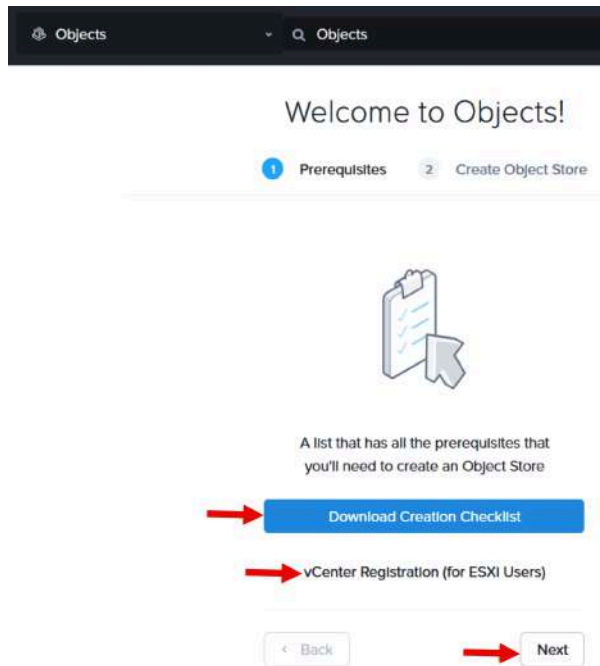


Figure 7: Objects Welcome Page - Prerequisites

7. Click **Download Creation Checklist** to download the list of prerequisites for deploying an object store.
8. (Only for ESXi clusters) Click **vCenter Registration** before deploying object store on an ESXi cluster.
To deploy object stores on the ESXi clusters, you need to provide the vCenter credentials and configure the IPAM for the ESXi networks. For more information, see the [Managing vCenter for Object Service](#) section in the *Objects User Guide*.

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

- Create Object Store
- View Object Store

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

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

For information on the built-in roles in Prism Central, see the [Built-in Role Management](#) section in the *Security Guide*.

9. Click **Next**.

10. Click **Create Object Store** to start creating the first object store.

For information on creating the object store, see the [Creating or Deploying an Object Store \(Prism Central\)](#) section in the *Objects User Guide*.

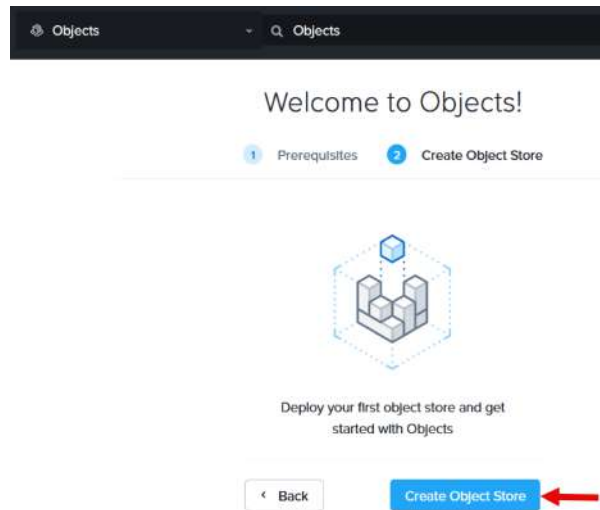


Figure 8: Objects Welcome Page - Create Object Store

After deployment, the Objects app appears on the **My Apps** page. You can also select **Objects** in the Application Switcher to access and use the app. See the [Nutanix Objects Guide](#) for information on using object stores.

Deploying Database Service

About this task

Nutanix Database Service (NDB) automates and simplifies database administration, bringing one-click simplicity and invisible operations to database provisioning and life-cycle management. You can deploy multiple instances of Database Service.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.
- For the network requirements for the Database Service, see the [NDB Network Requirements](#) section in the *Nutanix Database Service Administration Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Database Service app.
5. On the app details page, click **Deploy**.
6. Enter the **Application Name** and **Application Description**.

7. In the **EraService** section, keep the default values to deploy Database Service with a non-HA configuration.

For other deployment configurations, see the [NDB Control Plane Configuration and Scalability](#) section in the *Nutanix Database Service Administration Guide*.

8. In the **NETWORK ADAPTERS (NICs)** section, select an appropriate NIC from the dropdown menu.

9. Click **Deploy**.

After deployment, the Database Service app appears on the **My Apps** page. You can select **Database Service** in the Application Switcher to access and use the app. See the [Nutanix Database Service Administration Guide](#) for information on configuring and using Database Service. You can perform these steps again to deploy another instance of Database Service.

Deploying Move

About this task

Nutanix Move (Move) is a cross-hypervisor mobility solution to migrate virtual machines (VMs) with minimal downtime. You can deploy multiple instances of Move.

Note: Nutanix recommends to deploy Move using the Prism Element UI. For information on Move deployment using the Prism Element UI, see the [Move Deployment](#) section in the *Move User Guide*.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 12.
- Ensure that you have enabled DHCP in the network adapter.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Move app.

Note: The Move version that is available to you for deployment is the one that comes bundled with this Prism Central version. This bundled version may not necessarily be the latest Move version.

5. On the app details page, click **Deploy**.
The Deploy Move page appears.
6. Enter the **Application Name** and **Application Description**.

7. In the **MoveService** section, do the following:

- a. Update the value of **vCPUs** to 2.
- b. Update the value of **Cores per vCPU** to 2.
- c. Update the value of **Memory (GiB)** to 8.
- d. Click **+** in the **DISKS** section.
- e. For **Disk 2: SCSI**, enter the disk size as 50 in the **Size (GiB)** field.
- f. In the **NETWORK ADAPTERS (NICs)** section, select an appropriate NIC from the dropdown menu.

Note: You must select an appropriate NIC for a successful deployment.

8. Click **Deploy**.

After deployment, the Move app appears on the **My Apps** page. You can select **Move** in the Application Switcher to access and use the app. See the [Move User Guide](#) for information on configuring and using Move. You can perform these steps again to deploy another instance of Move.

Preferred Partner App Deployment

A Preferred Partner app is developed and certified by Nutanix in collaboration with the preferred partner. The app is pre-seeded for deployment on the Nutanix platform. For example, OpenShift.

Similar to Nutanix apps, the Preferred Partner app is available in the Marketplace for deployment immediately after you enable Marketplace in Prism Central. You do not have to enable Self-Service to access the Preferred Partner app in the Marketplace.

Ensure that you meet all the resource requirements specific to the app before you deploy the Preferred Partner app.

Deploying OpenShift

Red Hat OpenShift brings together tested and trusted services to reduce the friction of developing, modernizing, deploying, running, and managing applications. Built on Kubernetes, it delivers a consistent experience across public cloud, on-premise, hybrid cloud, or edge architecture.

About this task

Note: You can deploy OpenShift using the Internal Project only. For information on the Internal Project, see [Project Management](#) on page 78.

Before you begin

Ensure that you meet the following prerequisites:

- Your Prism Central has required resources to enable OpenShift. For more information, see [Red Hat Documented Cluster Resource Requirements](#).

- You have an Internet connection with or without proxy. If you use a firewall, you must configure access to the following URLs:

- *.docker.io
- *.production.cloudflare.docker.com
- *.nutanix.github.io
- download.nutanix.com

Allowlist these URLs in the firewalls to ensure that Prism Central and Prism Element clusters have access to these URLs. For URLs specific to OpenShift, see the [OpenShift Documentation](#).

- OpenShift deployment through Prism Central marketplace requires an AHV hypervisor. For Nutanix clusters with an ESXi hypervisor, use VMware Installer Provisioned Infrastructure (IPI) or User Provisioned Infrastructure (UPI) based installation.

Procedure

To deploy OpenShift, perform the following steps:

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. Under Preferred Partners, click **Get** for the OpenShift app.
5. Click **Deploy**.
6. Enter the name of the application in the **Application Name** field.
Following are the rules for naming convention.

- The name of the blueprint can start with an alphanumeric character or an underscore.
- The name must have at least one character.
- Use only space, underscore, and dash as special characters.
- Do not end the name with a dash.

7. Under Profile Variables, do the following.

- a. Select the versions from the **OpenShift minor release version** and **OpenShift patch release version** dropdown menus.
- b. Enter the base domain in the **OpenShift base domain** field.
Enter the base domain of your environment to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the base domain and cluster name parameter values that uses the `cluster_name.base_domain` format.
- c. Enter the Prism Central IP address or FQDN in the **Prism Central FQDN** field.
The value must be in the SAN properties of the Prism Central SSL certificate.
- d. Enter the Prism Element IP address or FQDN in the **Prism Element FQDN** field.
The value must point to the cluster virtual IP address where the provisioner VM is going to be deployed. Enter the value in the `your.prismcentral.domainname` format.
- e. Enter the Prism Element port in the **Prism Element port** field.
The default value for the Prism Element console port is 9440.
- f. Enter the IP address blocks for virtual machines in the **OpenShift Machine Network** field.
The block must not overlap with cluster and service networks. Enter the value in the `10.0.0.0/16` format.
- g. Enter the virtual IP (VIP) address that you configured for control plane API access in the **OpenShift API VIP** field.
The VIP address must be in the same subnet as the machine network. Enter the value in the `XXX.XXX.XXX.XXX` format.
- h. Enter the virtual IP (VIP) address that you configured for cluster ingress in the **OpenShift Ingress VIP** field.
The VIP address must be in the same subnet as the machine network. Enter the value in the `XXX.XXX.XXX.XXX` format.
- i. Enter the IP address block for Kubernetes pods in the **OpenShift Pod Network** field.
The default value is `10.128.0.0/14` with a host prefix of `/23`.
- j. Enter the subnet prefix length to assign to each individual node in the **OpenShift Pod addresses per host** field.
For example, if host prefix is set to 23 then each node is assigned a `/23` subnet out of the given CIDR. A host prefix value of 23 provides 510 pod IP addresses.
- k. Enter the IP address block for Kubernetes services in the **OpenShift Service Network** field.
The default value is `172.30.0.0/16`.
- l. Enter the number of CPUs per control plane host in the **Control Plane CPU** field (a minimum of 4 vCPUs and a maximum of 128 vCPUs).
- m. Enter the memory (MiB) per control plane host in the **Control Plane Memory** field (a minimum of 16384 MiB and a maximum of 1048576 MiB).
- n. Enter the number of compute machines (worker machines) to provision in the **Compute replicas** field.
- o. Enter the pull secret in the **OpenShift Pull Secret** field.
You can get a pull secret from the Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services, such as Quay.io. <https://console.redhat.com/openshift/install/pull-secret>.

- p. (Optional) Enter the details of the HTTP proxy server in the **(Optional) Internet proxy for HTTP** field.

The format for defining an HTTP proxy typically includes the proxy server address, port, and sometimes authentication credentials (when required). Enter the HTTP proxy server details in the following format.

- Without Authentication: The format must be `http://<ip>:<port>`, where `<ip>` is the domain name or IP address of the proxy server and `<port>` is the port number of the proxy server.
- With Authentication: The format must be `http://<username>:<pswd>@<ip>:<port>`, where `<username>` is the username for authenticating with the proxy and `<pswd>` is the corresponding password for the username.

- q. (Optional) Enter the details of the HTTPS proxy server in the **(Optional) Internet proxy for HTTPS** field.

The format for defining an HTTPS proxy typically includes the proxy server address, port, and sometimes authentication credentials (when required). Enter the HTTPS proxy server details in the following format.

- Without Authentication: The format must be `https://<ip>:<port>`, where `<ip>` is the domain name or IP address of the proxy server and `<port>` is the port number of the proxy server.
- With Authentication: The format must be `https://<username>:<pswd>@<ip>:<port>`, where `<username>` is the username for authenticating with the proxy and `<pswd>` is the corresponding password for the username.

- r. (Optional) Enter the destination domain names, IP addresses, or other network CIDRs to exclude from the proxy in the **(Optional) Bypass the proxy** field.

The details you enter in the **(Optional) Bypass the proxy** field must be a comma-separated list in the following format.

example.com,.example.com,192.168.1.1,192.168.0.0/24

- s. To diagnose and resolve issues when something goes wrong in the deployment process, select **True** in the **DEBUG** dropdown menu.

8. Under Provisioner, click **NETWORK ADAPTERS** and select a NIC from the **NIC** dropdown menu. The NIC dropdown menu options are based on the cluster configuration.

Note: Cluster details are optional and are inferred from the NIC you select.

9. Under Credentials, expand the credentials and enter the password in the **Password** field and private key in the **SSH Private Key** field.

Note: The user must be a local account with the User Admin role in Prism Central and Prism Element. Nutanix recommends you to create a service account in Prism Central and Prism Element.

10. Click **Deploy**.

Deploying Custom Apps from the Marketplace

The Other Apps category includes hybrid cloud apps and custom apps that are available for deployment only when Self-Service is deployed and licensed in your Prism Central instance.

About this task

Use this procedure to deploy custom apps from the Marketplace. The Marketplace shows only those custom apps in the Other Apps category that are approved in Self-Service and published in

the Marketplace by your Prism Admin. For information on approval and publishing of apps, see the [Marketplace Manager](#) section in the *Self-Service Administration and Operations Guide*.

For information on deployment of hybrid cloud apps (pre-seeded), see the [Self-Service Marketplace Items Guide](#).

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. Click **Get** for the app in the Other Apps category that you want to deploy.
5. Click **Deploy**.
The Deploy page is displayed.

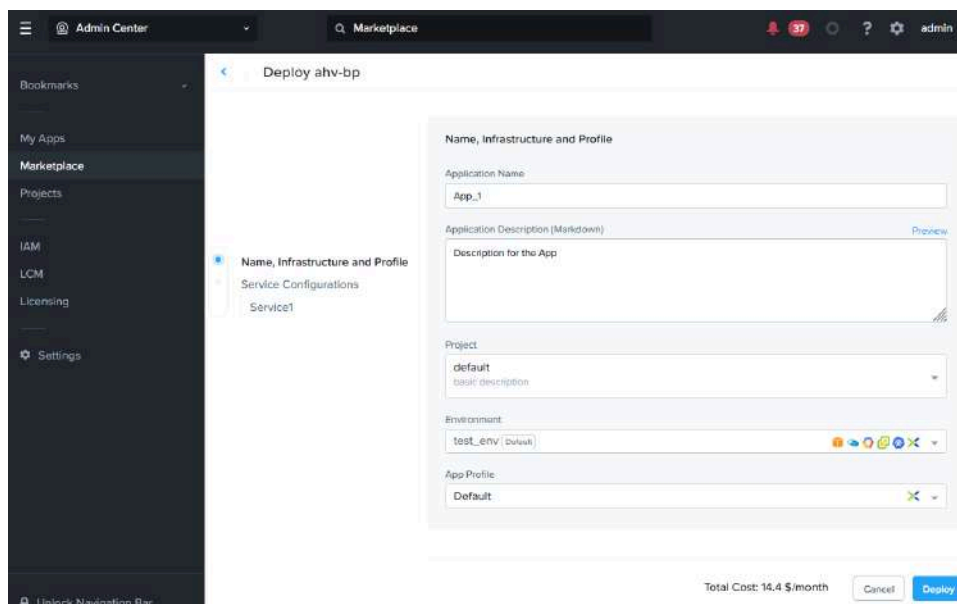


Figure 9: Deploy App

6. Enter a name for the application in the **Application Name** field.
Following are the rules for naming convention.
 - The name of the application can start with an alphanumeric character or an underscore.
 - The name must have at least one character.
 - Use only space, underscore, and dash as special characters.
 - Do not end the name with a dash.
7. Enter a description for the application in the **Application Description** field.
8. Select the project from the **Project** dropdown menu.

9. Select the environment from the **Environment** dropdown menu.

If you select an environment that is different from the account that you used for blueprint configuration, Self-Service updates all platform-dependent fields to match with the selected environment configuration.

For example, you created the application blueprint using an account with an environment (ENV1) so that the platform-dependent fields are similar to ENV1. While launching the application blueprint, if you select a different environment (ENV2), Self-Service updates all platform-dependent fields to match with the ENV2 configuration. For more information, see the [Environment Patching Behavior](#) and [Patching for Clusters and Subnets](#) sections in the *Self-Service Administration and Operations Guide*.

10. Select an application profile in the **App Profile** field.

Application profile provides different combinations of the service, package, and VM while configuring a blueprint.

11. In the section for the service configuration, verify the VM, disk, boot configuration, and network configuration. You can edit the fields based on your application requirements.



Figure 10: Application Launch - Service Configuration

12. If the blueprint is configured with a Nutanix account, do the following:

- a. Under Snapshot Configurations, select a snapshot policy from the **Snapshot Policy** dropdown menu.
- b. Based on the policy you select, select a rule from the **Select Local Rule** or **Select Remote Rule** dropdown menu.

The **Select Local Rule** or **Select Remote Rule** dropdown menu appears based on the **Snapshot Location** you defined in your blueprint. For more information, see the [Blueprint Configuration for Snapshots and Restore](#) section in the *Self-Service Administration and Operations Guide*. The values in the dropdown menu appear based on the snapshot policy you defined in the project and selected in the Snapshot Policy dropdown menu. For more information, see [Creating a Snapshot Policy](#) on page 138. The values also depend on the VM categories you configured in your blueprint.

The Snapshot Configuration section appears depending on the environment you select while launching the blueprint. If you select a specific environment, you must provide the snapshot policy and

snapshot rule to launch the blueprint. The Snapshot Configuration section does not appear in case you select the environment with all project accounts for the launch.

Note: Ensure that you have a valid NIC in the blueprint.

13. Click **Deploy**.

Executing a Runbook from the Marketplace

You can execute an approved and published runbook from the Marketplace.

Before you begin

Ensure that the runbook that you want to execute is approved by the administrator and published in the marketplace. For more information, see the [Approving and Publishing a Blueprint or Runbook](#) section in the *Self-Service Administration and Operations Guide*.

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. Click **Get** for the runbook in the Other Apps category that you want to execute.
5. Click **Execute**.
The **Execute Runbook** page appears.

Project

default

Default Endpoint (Optional) ⓘ

Linux

Linux IP

Endpoint Mapping (1)

Please select an endpoint belonging to the selected project for every endpoint used in the marketplace runbook.

Esript Endpoint need to be used, can be endpoint of any type

Linux

Variables (0)

Cancel Execute

Figure 11: Execute Runbook

6. Select the project for the runbook execution from the **Project** dropdown menu.

7. To change the default endpoint for the execution, select an endpoint from the **Default Endpoint** dropdown menu. This step is optional.

Note: If you published runbook without endpoints, then you must select the endpoint from the project in which you are executing the runbook.

8. To update the added variable in the runbook, click the respective variable field and edit the variable. This step is optional.

Note: You can update the variable only if the variable is marked as runtime editable while adding the variable in the runbook.

9. Click **Execute**.

Cloning a Marketplace Item

You can clone the apps in the Other Apps category from the marketplace.

Before you begin

The IAM interface has been revamped to provide fine-grained RBAC so that custom roles can be created with granular operations or permissions. The option to clone a marketplace item appears when the role assigned to you has permissions to clone marketplace items. For more information on roles and permissions, see the [Security Management using Prism Central \(PC\)](#) section of the *Security Guide*.

About this task

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. Click **Get** for the app or runbook in the Other Apps category that you want to clone.
5. Click **Clone**.
6. Enter the name for the clone.
7. Select the project that you want to assign to the cloned app or runbook from the **Project** dropdown menu.
8. Click **Clone**.

MY APPS

My Apps provides a single workspace to manage all the apps that you deployed from the Marketplace. These deployed apps can be Nutanix Apps such as Self-Service, Objects, and so on, or the apps deployed using preconfigured Marketplace blueprints or custom blueprints.

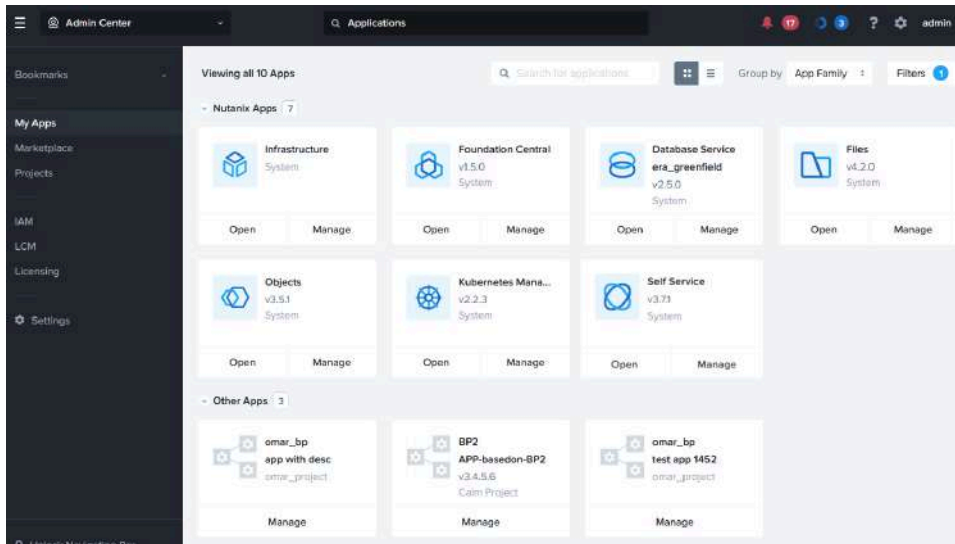


Figure 12: My Apps

Note: You can access and manage apps in My Apps only after your Prism Admin enables the Marketplace. For more information, see [Enabling Marketplace](#) on page 12.

The Nutanix Apps section of the My Apps page lists the Nutanix Apps that you deployed from the Marketplace, such as Self-Service, Objects, and so on. For information on managing Nutanix Apps, see [Nutanix Apps Management](#) on page 33.

The Other Apps section lists the hybrid cloud apps and custom apps that you deploy from the Marketplace or Self-Service. The Other Apps section appears when you deploy Self-Service in your Prism Central instance. For information on managing Other Apps, see [Other Apps Management](#) on page 35.

For more information on the app categories, see [Marketplace](#) on page 9.

For more information on deploying hybrid cloud apps and custom apps from Self-Service, see [Blueprint Management](#) in the *Self-Service Administration and Operations Guide*.

My Apps Views

The My Apps page allows you to view apps in a Grid view (default view) or a List view.

- Grid View

In the Grid view, you can view each app as a tile. This view displays the app name, associated project, version, and the options to open the app (only for Nutanix Apps) and manage apps.

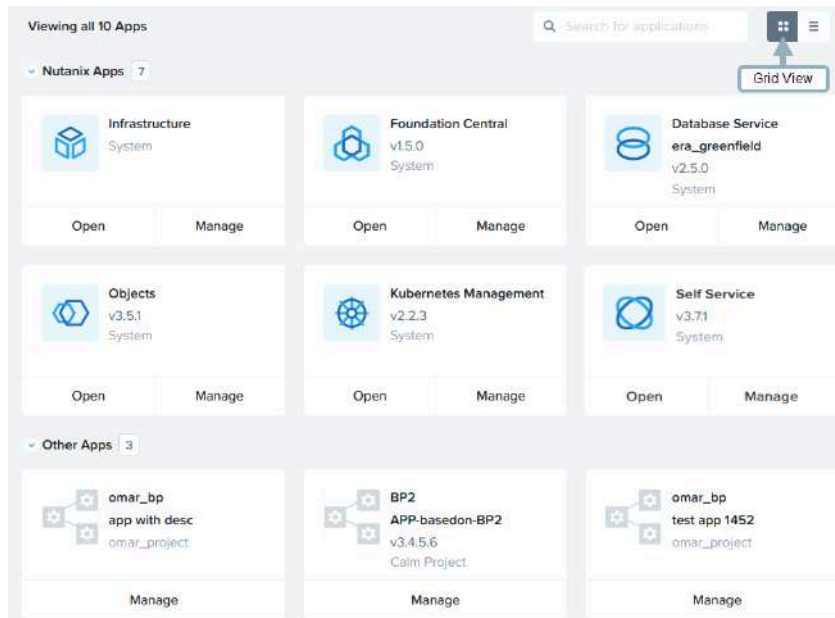


Figure 13: Grid View

- List View

The list view displays apps in multiple line items on separate tabs for Nutanix Apps and Other Apps when the **Group by** is selected as **App Family** (default view). You can also group apps by selecting **Project**, **Provider**, or **Categories**. For example, when you select **Project** in the **Group by** dropdown menu, the apps are grouped by their projects on separate tabs.

Note: You can group your apps by **Project** in the List view only when you have less than 10 projects.

The view displays the app name, version, source, associated project, state, owner, creation date, and last updated date. You can click the app name to see the app details.

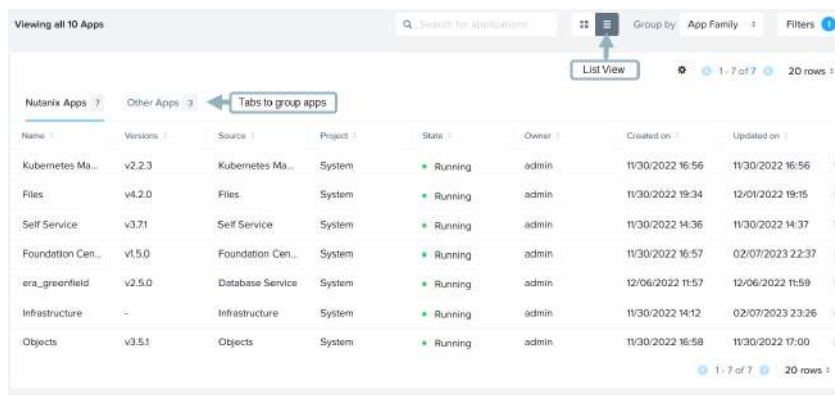


Figure 14: List View

You can type the name of the app in the **Search for applications** field to find the app on the My Apps page. You can also click **Filters** and filter your apps using the following criteria.

- **Source:** Type the source blueprint to filter associated apps.
- **States:** Select the app states from the dropdown menu to filter apps. See the App States section in this topic for more information on the states.
- **Owner:** Type the owner details to filter apps.
- **Project:** Select a project to which the app is associated.
- **VM Name:** Type the VM name of the app.
- **Account:** Select the associated account of the app.
- **Created:** Select a date range during which the apps were created.
- **Updated:** Select a date range during which the apps were updated.

App States

The My Apps page displays the state of the app based on the actions you perform on the **Manage** tab.

Table 3: App State

App State	Description
Provisioning	When you start an application.
Running	When the application is deployed and running after the provisioning state.
Stopping	When you have initiated an operation to stop the application.
Stopped	When the application is stopped.
Starting	When the application starts.
Restarting	When you have initiated an operation to restart the application after the application is stopped.
Deleting	When you have initiated an operation to delete the application.
Deleted	When the application is deleted.
Busy	When you have installed the NGT services on the VMs of an application.
Timed out	When the application is timed out.
Updating	When you are editing an application.
Error	When the application goes to error state due to any action you have performed in the Manage tab.
Failover-in-progress	When you have initiated a failover operation on Prism Central for the protected VMs of an application.

App State	Description
Failover-failed	<p>When the failover operation for the VMs has failed. The failure state mainly occurs in the following conditions.</p> <ul style="list-style-type: none"> • If there is any error from the Prism Central side. • If there is no NIC attached to the VM when you configure the recovery plan for the protected VM.

Note: The **Failover-in-progress** and **Failover-failed** states are only applicable for the applications that are running on the Nutanix platform.

App Access

The following table displays the different apps and the operations that you can perform depending on whether Self-Service is deployed in your Prism Central instance or not.

Table 4: App Access

App Type	Operations	When Self-Service is Not Deployed	When Self-Service is Deployed
Nutanix Apps	View apps that you deployed from the Marketplace.	Yes	Yes
Preferred Partner Apps	View apps that you deployed from the Marketplace.	Yes	Yes
Hybrid Cloud Apps (Other Apps)	View deployed apps or perform day 2 operations, such as creating snapshots, creating clones, or viewing metrics.	No	Yes
Custom Apps (Other Apps)	View deployed apps or perform day 2 operations, such as creating snapshots, creating clones, or viewing metrics.	No	Yes

Nutanix Apps Management

The App Details page of a Nutanix app provides the app summary and the audit trail of its deployment. The Overview tab on the App Details page displays the app UUID, name of the app with its version, and so on. The Audit tab displays the audit logs of the actions performed on the app.

The App Details page also provides the option to upgrade the app. For more information on upgrading Nutanix apps, see [Life Cycle Manager Guide](#).

The My Apps page also has the option to register existing instances of Database Service or Move so that you can manage them on the My Apps page. For more information, see [Registering Database Service or Move to My Apps](#) on page 34.

Note:

- When you upgrade Nutanix apps through LCM, you might notice the deployed app version on the My Apps page differs from the respective Marketplace Item version on the Marketplace Page. The Marketplace Item version refers to the version that is bundled with the Prism

Central version you have upgraded to. Whereas, the app version on the My Apps page shows the current version the app is running.

- When you upgrade from a previous version of Prism Central, you must enable Marketplace and perform the LCM inventory to see the correct versions of Nutanix Apps on the My Apps page.

Nutanix Apps Permissions

The following table lists the different operations for Nutanix Apps and permissions for different users to perform those operations in My Apps:

Table 5: My Apps Operations for Nutanix Apps

Operation	Prism Central Admin	Project Admin	Developer	Consumer	Operator
View or open Nutanix Apps	Yes	Yes	Yes	Yes	Yes
Search or Filter Nutanix Apps	Yes	Yes	Yes	No	No
Upgrade Nutanix Apps	Yes	No	No	No	No

Registering Database Service or Move to My Apps

About this task

You can register existing instances of Database Service or Move so that you can manage them through My Apps.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.

4. On the My Apps page, do the following:

- » To register an existing instance of Database Service, click **Register Database Service**.
- » To register an existing instance of Move, click **Register Move**.

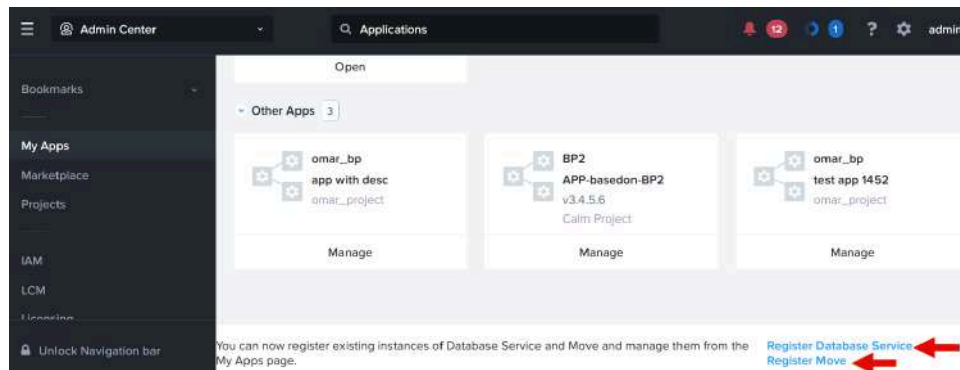


Figure 15: Register Database Service or Move

5. On the Register page, select the VM of the instance from the list.

You can also search the VMs on the basis of name, cluster name, or IP address.

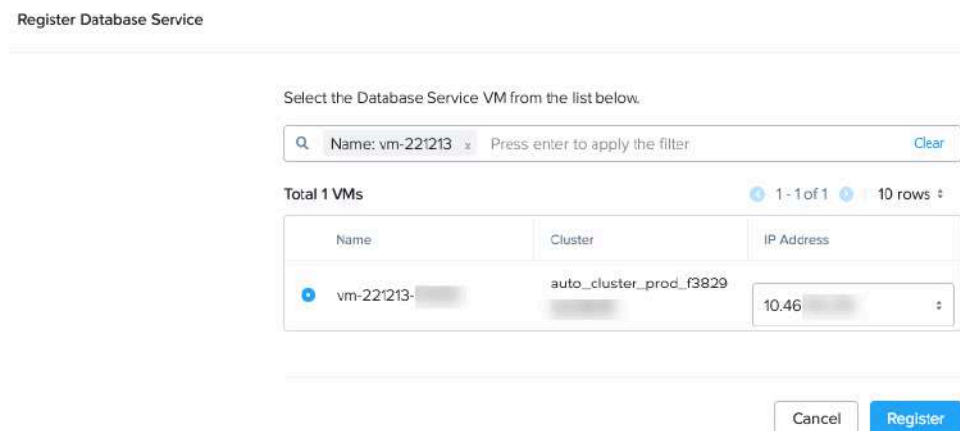


Figure 16: Database Service VMs

6. Click **Register**.

The registered instance appears on the My Apps page.

Other Apps Management

Applications that are categorized as Other Apps are deployed using preconfigured Marketplace blueprints or the custom blueprints that you configure in Self-Service. For more information about blueprint configuration, see [Self-Service Blueprints Overview](#) in the *Self-Service Administration and Operations Guide*.

Note: You can manage Other Apps only when you have deployed Self-Service in your Prism Central instance.

Click **Manage** for an app in the Other Apps section of the My Apps page whenever you want to view details or manage the configuration of the app. Each app has different tabs that you can use to manage the app. For information about these tabs, see [Tabs for Other Apps Management](#) on page 37.

Other Apps Configurations

You can perform the following operations on your apps in the Other Apps category.

- Run system-level and user-level actions on your app. For more information, see [Running a System-Level Action](#) on page 66 and [Running a User-Level Action](#) on page 48.
- Install and manage NGT Apps. For more information, see [Installing NGT Apps](#) on page 67 and [Managing NGT Apps](#) on page 69.
- Create an image. For more information, see [Creating an Image on a Nutanix Platform](#) on page 73.
- Create and restore snapshots. For more information, see [Snapshot and Restore Overview](#) on page 55.
- Update VM configurations. For more information, see [Update VM Configurations of a Running App](#) on page 48.
- Update actions and credentials. For more information, see [Updating Actions and Credentials of an App](#) on page 70.
- Clone an application. For more information, see [Cloning an App](#) on page 75.
- Delete an application. For more information, see [Deleting an App](#) on page 75.

My Apps Permissions

The following table lists the different app-specific operations for Other Apps and the permissions for different users to perform those operations.

Table 6: My Apps Permissions

Operation	Prism Central Admin	Project Admin	Developer	Consumer	Operator
Update VM Configuration	Yes	Yes	Yes	Yes	No
Day 2 Operations (Snapshot, Launch Console)	Yes	Yes	Yes	Yes	Yes
Add or Edit Actions or Tasks (Credentials, Packages, Post-Delete Actions)	Yes	Yes	Yes	Yes	No
Clone Apps	Yes	Yes	Yes	Yes	No
View App Details or Services	Yes	Yes	Yes	Yes	Yes
View Source Blueprint of Apps	Yes	Yes	Yes	Yes	No

Tabs for Other Apps Management

The App Details page of an app in the Other Apps category has different tabs to organize app-specific information or the actions that you perform on them.

Overview Tab

The **Overview** tab provides comprehensive details about the app such as the name, the state of the app, and so on. The tab consist of the following panels to categorize the app details for easy consumption.

- Cost Summary
- App Summary
- App Status
- VM Info

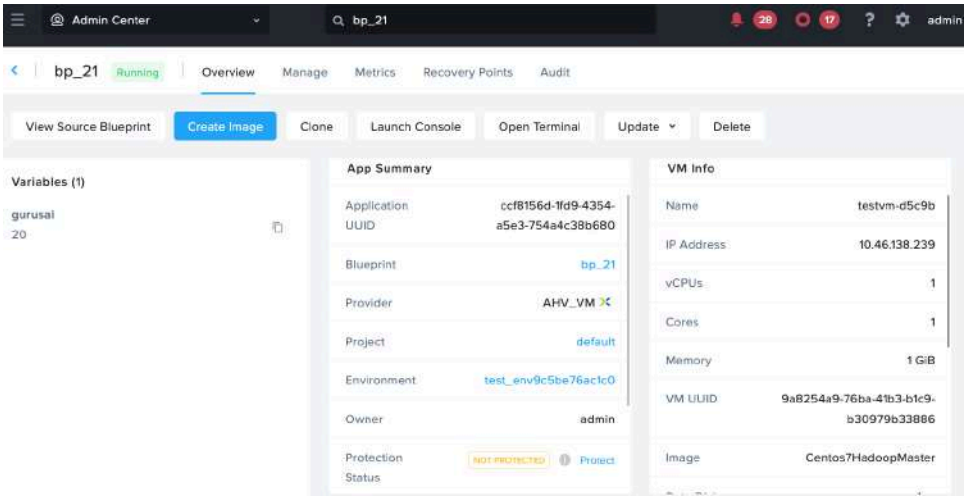


Figure 17: Overview Tab

Table 7: Overview Tab

Panel	Description
Cost Summary	Displays the total cost, current cost for each hour, and the cost incurred in a month for the resources that are running in the blueprint. The cost summary panel also displays a graphical representation of the incurred cost. <div>Note: The Cost Summary panel appears only for Nutanix and VMware providers.</div>

Panel	Description
App Summary	<p>Displays the following app details.</p> <ul style="list-style-type: none"> • Application UUID: Displays a unique identification code for the app. UUID is automatically generated after the app is created and in running state. • Blueprint: Displays the blueprint from which the app is created. • Cloud: Displays the cloud provider icon that hosts the app. You can click the blueprint name to view details of the blueprint. • Provider: Displays the provider associated with the app. The provider can be Nutanix, VMware, AWS, Azure, or GCP. • Project: Displays the project that you associated with the app. • Environment: The environment associated with the app. These environments are created within the project and associated with the app blueprints. • Owner: Displays the role of the user. • Created On: Displays the duration of existence of the app. • Last Updated On: Displays the duration of the latest app update.
App Status	<p>Displays the summary of virtual machines (VMs). The panel displays the number of VMs that are in the following state.</p> <ul style="list-style-type: none"> • On • Busy • Error • Off

Panel	Description
VM info	<p>Displays the following VM details of the app.</p> <ul style="list-style-type: none"> • Name: Displays the VM name. • IP Address: Displays the IP address of the VM. • Image: Displays the image from which the VM is created. • vCPUs: Displays the number of vCPU allocated to the VM. • Cores: Displays the number of cores allocated to the VM. • Memory: Displays the total memory allocated to the VM. • VM UUID: Display the unique identification code for the virtual machine. • Network Adapters: Displays the network adapters used in the VM. You can use the down arrow key to view the details of the network adapter. • VPC: Displays the associated VPC and the connection status. • Cluster UUID: Display the unique identification code for the associated cluster. • Cluster Name: Displays the name of the associated cluster. • Categories: Displays the categories added to the VMs. You can use the down arrow key to view the details of the network adapter. <div> <p>Note:</p> <p>For single-VM app, if the VM has network issues or is deleted from the provider side, the <i>Unreachable</i> state is displayed for the VM after the platform sync. This VM state indicates that any actions, such as stop, delete, restart, and so on that runs on the app will not be executed on this VM.</p> <p>In case the VM becomes available again, the Overview tab shows the state of the VM as <i>Restored</i>.</p> </div>

The **Overview** tab also has the following buttons to perform app-specific actions.

Table 8: Overview Tab Options

Button	Action
View Source Blueprint	Use this option to view the details of the blueprint associated with the app.
Create Image	Use this option to create images from an existing single-VM or multi-VM app running on a Nutanix platform. For more information, see Creating an Image on a Nutanix Platform on page 73.
Clone	Use this option to clone an app if you are using Nutanix, VMware, or AWS as the provider for the app. For more information, see Cloning an App on page 75.
Snapshot	Use this option to create a snapshot of a single-VM app running on a VMware, AWS, or Azure platform. For more information, see Creating Snapshots on a VMware Platform on page 60, Creating Snapshots on an AWS Platform on page 63, and Creating Snapshots on an Azure Platform on page 64.
Launch Console	Use this option to launch the console for the app.
Open Terminal	Use this option for web SSH or web RDP sessions in both single-VM and multi-VM apps for local and remote Prism Central. The sessions are supported for both Linux and Windows VM, including the VMs that are within the VPCs.
Update	Use this option to update actions and credentials of a single-VM app or update the VM configuration of a single-VM app running on VMware, AWS, or Azure. For more information, see Updating Actions and Credentials of an App on page 70 and Update VM Configurations of a Running App on page 48.
Delete	Use this option to delete an app. For more information, see Deleting an App on page 75.

Manage Tab

The **Manage** tab lists the system-level and user-level actions that you can perform on the app. When you click any of the listed actions, the editor displays the action dependencies.

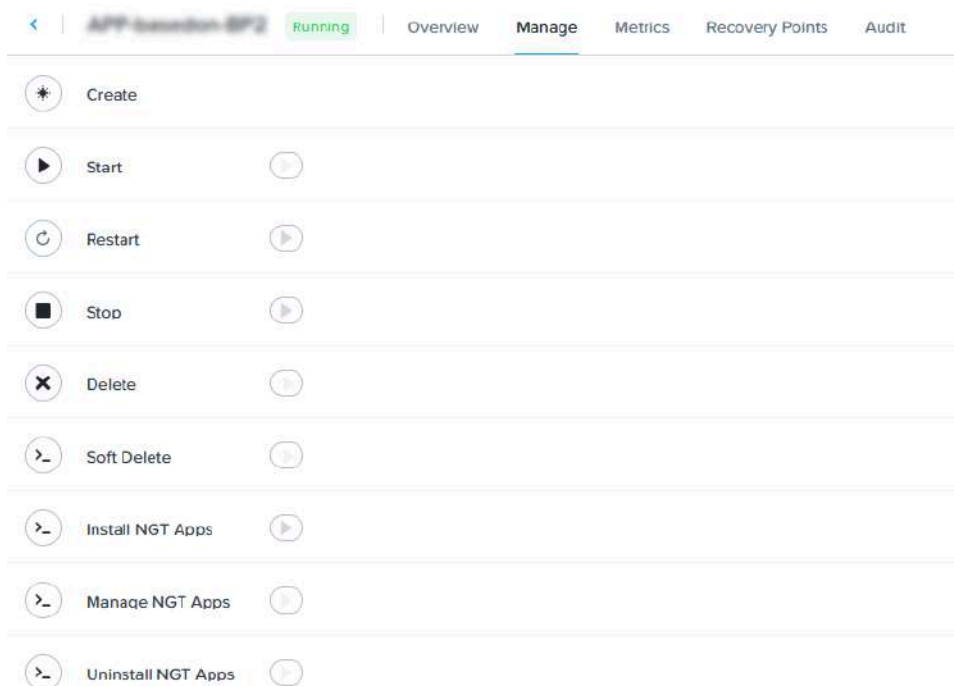


Figure 18: Manage Tab

You can perform the following system-level actions on an app.

Table 9: System-Level Actions

Start	Starts an app.
Restart	Restarts an app.
Stop	Stops an app.
Delete	Deletes an app including the underlying VMs on the provider side.
Soft Delete	Deletes the app from the Self-Service environment but does not delete the VMs on the provider side.

Install NGT Apps

Installs NGT service on your VM. For more information, see [Installing NGT Apps](#) on page 67.

Nutanix guest tools (NGT) is a software bundle that you can install on a guest virtual machine (Microsoft Windows or Linux) to enable the advanced functionalities provided by Nutanix. For more information on NGT, see the [Nutanix Guest Tools](#) section in the *Prism Central Infrastructure Guide*.

Note:

- NGT services applies only to single-VM apps running with Nutanix as the provider.
- For Kubernetes, the start, stop, and restart actions are disabled.

Manage NGT Apps

Manages NGT services for your app. You can enable or disable self-service restore (SSR) or Volume Shadow Copy Service (VSS) services.

The self-service restore (SSR) allows virtual machine administrators to do a self-service recovery from the Nutanix data protection snapshots with minimal administrator intervention. For more information, see [Self-Service Restore](#) section in the *Data Protection and Recovery with Prism Element Guide*.

Volume Shadow Copy Service (VSS), also known as Shadow Copy or Volume Snapshot Service, creates an app-consistent snapshot for a VM and is limited to consistency groups consisting of a single VM.

Uninstall NGT Apps

Uninstalls NGT services from the VM. For more information, see [Uninstalling NGT Apps](#) on page 69.

Note: You cannot perform the Create action once the blueprint is launched and app is created.

You can also perform the following user-level actions on the **Manage** tab.

Table 10: User-Level Actions

Action	Description
Update VM Configuration	You can update configurations of a single-VM or multi-VM app running on the Nutanix platform. For more information, see Update VM Configurations of a Running App on page 48.
Create and Restore Snapshots	You can create or restore snapshots of a single-VM or multi-VM app running on a Nutanix Platform. For more information, see Snapshot and Restore on Nutanix or VMware Platform on page 55.
Run Custom Actions	You can run custom actions that you configured in your blueprints in Self-Service (such as scale in, scale out, and so on). For more information about custom actions, see Action Overview in the <i>Self-Service Administration and Operations Guide</i> .

Metrics Tab

The **Metrics** tab allows you to view performance metrics of the VM. The **Metrics** tab displays a section on the left with a list of metrics.

Note:

- The Metrics tab applies only to single-VM apps running on the Nutanix platform.
 - The identified anomalies are based on VM behavioral machine-learning capabilities.
- 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 current value. To set the time interval (last 24 hours, last week, or last 21 days), select the duration from the dropdown menu on the right.

Note: The machine-learning algorithm uses 21 days of data to monitor performance. A graph does not appear for less than 21 days of data.
 - To create an alert for this VM based on either behavioral anomalies or status thresholds, click the **Set Alerts** link above the graph.

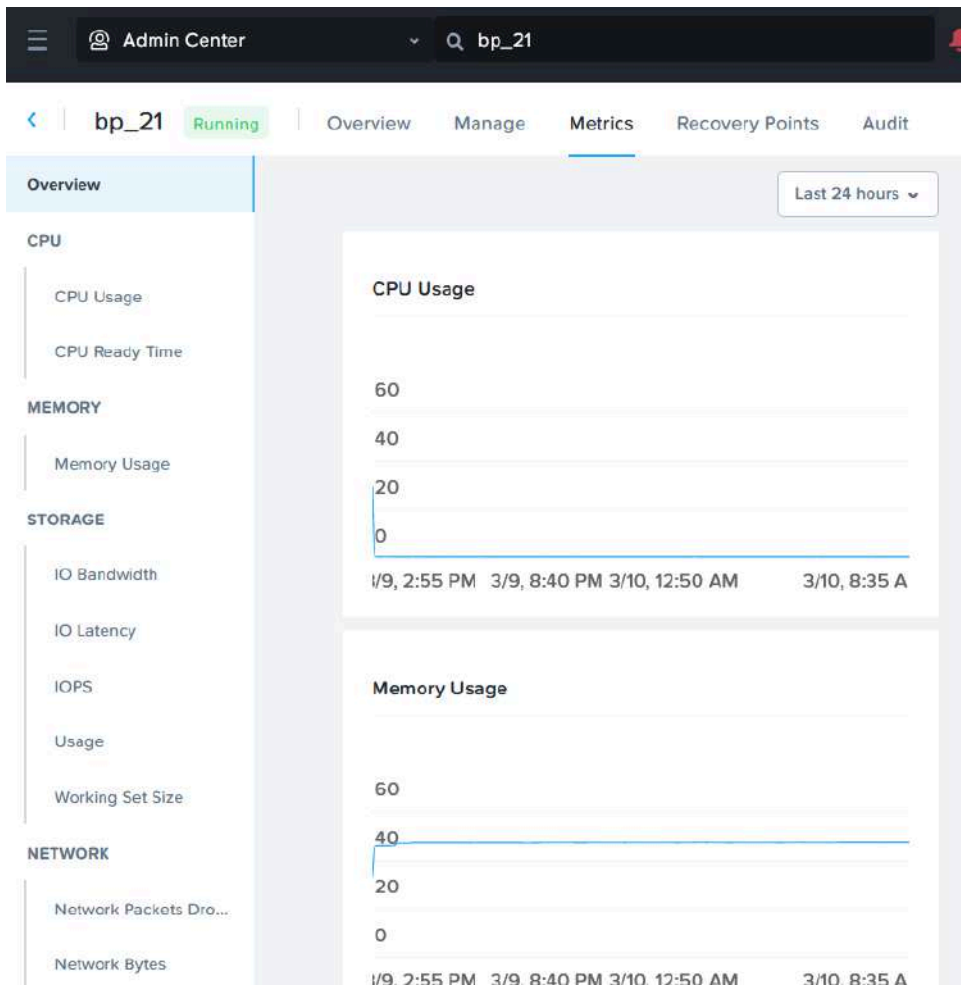


Figure 19: Metrics Tab

The following table describes the available metrics.

Table 11: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently the VM is using (0–100%).
CPU Ready Time	Displays the current, high, and low percentage of CPU ready time (0–100%).
Memory Usage	Displays the percentage of memory capacity currently the VM is using (0–100%).
I/O Bandwidth	Displays separate graphs for total, write (only), and read (only) I/O bandwidth used per second (Mbps or Kbps) for physical disk requests by the VM.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.

Metric	Description
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for 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).

Recovery Points Tab

The **Recovery Points** tab allows you to view the captured snapshots, restore apps from snapshots, and delete the snapshots for an app.

Note:

The Recovery Points tab applies only to apps running on a Nutanix or VMware platform.

To create a snapshot of a single-VM or multi-VM app that is running on a Nutanix or VMware platform, use the snapshot action on the **Manage** tab of the app. For more information, see [Snapshot and Restore on Nutanix or VMware Platform](#) on page 55.

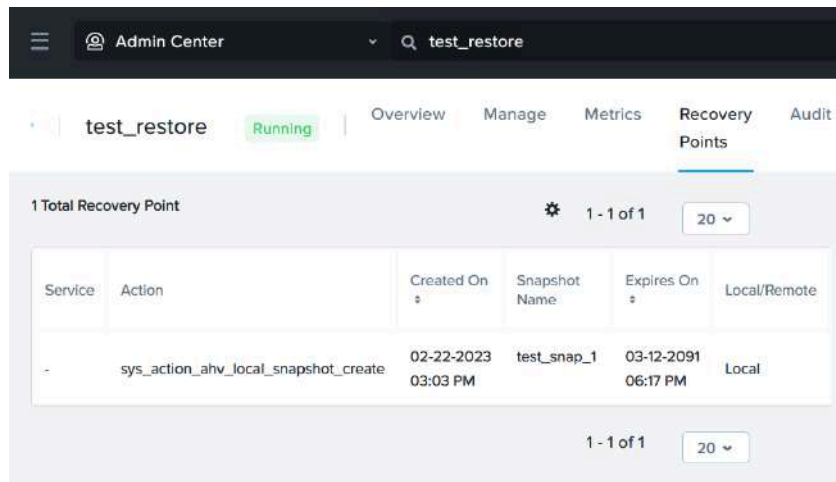


Figure 20: Recovery Points Tab

Table 12: Recovery Points Tab Fields

Fields	Description
Service Name	The name of the application service from which the snapshot is taken.
Cloud Provider	The name of the platform. This can be Nutanix or VMware.

Fields	Description
Snapshot Name	The name of the snapshot taken. This name is defined when you configure the snapshot/restore in the blueprint. By default, this name is generated using the VM name and the timestamp.
Action	The profile action that was used to create the snapshot.
Created On	The date and time of the snapshot creation.
Expires On	The expiration time of the snapshot.
Local/Remote	The location where the snapshot is managed. This can be both Local, Remote, or both for Nutanix and Local for VMware.
Recovery Point Type	Displays whether the snapshot type is app-consistent or crash-consistent.

Note: The VMware application snapshots that are taken prior to Self-Service version 3.7 do not show metadata such as snapshot parent information, memory, or quiescing in the Recovery Point table unless you trigger a new snapshot action.

Snapshots Tab

The **Snapshot** tab allows you to view the captured snapshots, restore apps from snapshots, and delete the snapshots for a single-VM app running on VMware or Azure. The tab displays the following fields for a snapshot.

Table 13: Snapshots Tab Fields

Fields	Description
ID	Displays the ID of the snapshot. Snapshot ID is unique and generated automatically when you take a snapshot.
Name	Displays the name of the snapshot.
Description	Displays the description of the snapshot.
Parent	Displays the parent blueprint app from which the snapshot is taken.
Creation Time	Displays the date and time when the snapshot is taken.

AMIs Tab

The **AMIs** tab allows you to view the captured snapshots, restore app from snapshots, and delete the snapshots for a single-VM app running on an AWS platform. The tab displays the following fields for a snapshot.

Table 14: AMI Tab Fields

Fields	Description
ID	Displays the ID of a snapshot. Snapshot ID is unique and generated automatically when you take a snapshot.
Name	Displays the name of the snapshot.

Fields	Description
Description	Displays the description of the snapshot.
Creation Time	Displays the date and time when the snapshot is taken.

Services Tab

The **Services** tab shows the services that are included in the app. You can select the service to view the configuration details, open terminal, or create image from the service inspector panel.

Note: Service tab is only applicable for multi-VM apps.

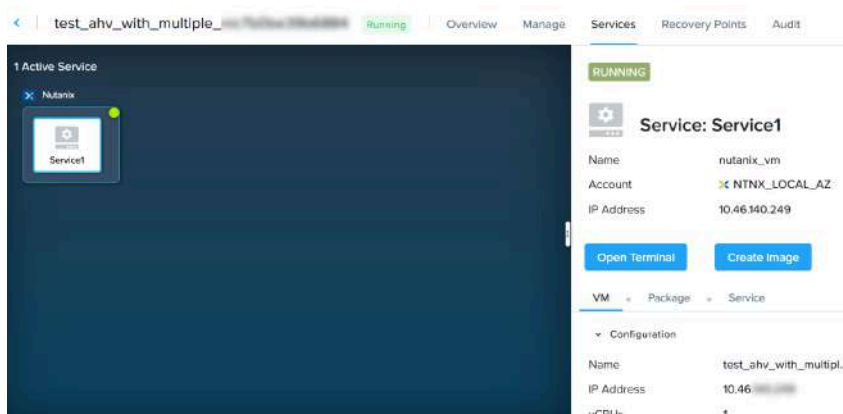


Figure 21: Services Tab

Audit Tab

The **Audit** tab lists the action or actions that are performed on an app. To view the detailed course of the action, click the action.

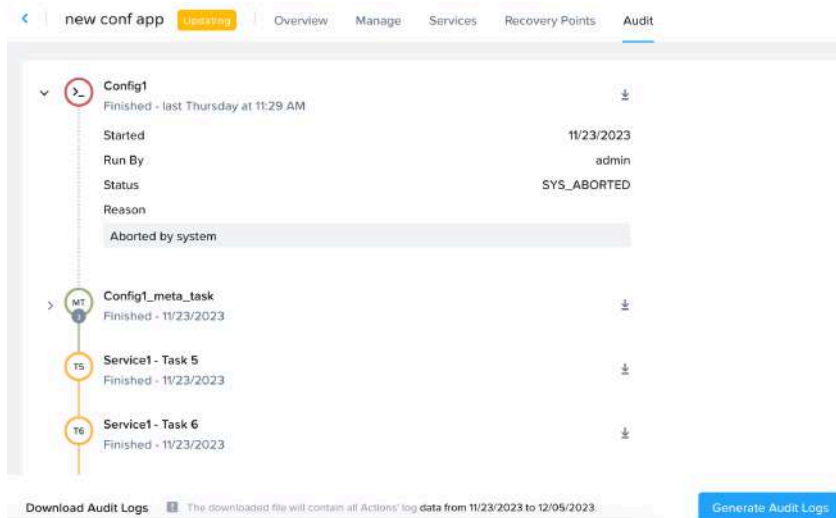


Figure 22: Audit Tab

On the **Audit** tab, you can either download the individual audit log of each action or bundle up the audit logs of all actions and download them as a compressed file. The compressed file contains the directory structure of all audit logs.

To download the log of an individual action, click the Download icon next to the action.

Downloading bundled audit logs is an asynchronous operation where you first generate the compressed file of audit logs of all application actions using the **Generate Audit Logs** button. You then download the compressed file using the **Download Audit Logs** button. The generated compressed file stays for 24 hours for you to download. You also have the option to regenerate the logs, if required.

Running a User-Level Action

You can configure user-level actions (custom actions or in-built user-level actions, such as Update VM Configuration) when you configure the blueprint of an app in the Other Apps category. Perform the following procedure to run a user-level action.

Before you begin

Ensure that you have configured a user-level action while configuring the blueprint of the app. For more information on user-level actions, see [Action Overview](#) in the *Self-Service Administration and Operations Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category that you want to run the user-level action on.
5. On the **Manage** tab, click the user-level action that you want to run.
The user-level action starts running for the app.

Update VM Configurations of a Running App

The update configuration feature allows you to update the virtual machine of a running app in the Other Apps category to a higher or lower configuration. Using this feature, you can modify VM specifications such as the vCPU, memory, disks, networking, or categories (tags) of a running production application with minimal downtime.

Note: Updating VM configuration of a running multi-VM app is supported on the Nutanix and VMware platforms.

The process to update VM configuration of a running app on Nutanix and VMware is different from other providers.

Update VM Configuration of an App on Nutanix or VMware

To update configurations of a single-VM or multi-VM app running on Nutanix or VMware, you need to perform the following steps:

- Add an update configuration to the app blueprint.

For more information, see [Update Configuration for VM](#) in the *Self-Service Administration and Operations Guide*.

- Run the corresponding action to update VM specifications.

You can update VM specifications from the **Manage** tab of the app. While launching the update, you can define the variables and verify the updates defined for the service by looking at the original value and updated value. You can also modify the values if the component is editable. You can also check the cost difference at the top of the page before applying the changes. For more information on updating VM specifications on Nutanix, see [Updating the VM Configuration of an App on a Nutanix Platform](#) on page 49. For more information on updating VM specifications on VMware, see [Updating the VM Configuration of an App on a VMware Platform](#) on page 51.

Update VM Configuration of an Application on Other Providers

The option to update VM configuration of a running single-VM app on AWS or Azure is available by default. The attributes that you can update depends on the provider account you selected for the app.

- For more information about updating VM configuration on a AWS platform, see [Updating the VM Configuration of an App on an AWS Platform](#) on page 53.
- For more information about updating VM configuration on a Azure platform, see [Update the VM Configuration of an App on an Azure Platform](#) on page 54.

Updating the VM Configuration of an App on a Nutanix Platform

You can run the update configuration to modify the VM specifications, such as the vCPU, memory, disks, networking, or categories of a single-VM or multi-VM app.

About this task

For an overview of VM configuration update, see [Update VM Configurations of a Running App](#) on page 48.

Note:

- If you update configuration of an app after cloning from a source app, the update fails if the source app has static IP address configured.
- When you update configuration of an app, the CD-ROM attached to mount NGT services is removed.

Before you begin

Ensure that your blueprint developer has added the update configuration before launching the app blueprint. For more information, see [Update Configuration for VM](#) in the *Self-Service Administration and Operations Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to update the VM configuration.

- On the **Manage** tab, click the action corresponding to the update configuration. The **Run Action** window appears.

The screenshot shows the 'Run Action: Config1' window. It has a sidebar on the right with a tree view containing 'VM Configuration', 'Disks', 'Categories', and 'NICs'. The main content area is divided into sections: 'VM Configuration' with a table for vCPUs, Cores per vCPU, and Memory (GiB); 'Disks' with 'Disk 1: SCSI | Centos7HadoopMaster' (Existing) and 'Disk 2: SCSI' (New) with a 'Value' field; 'Categories'; and 'Network Adapters'. At the bottom, there are buttons for 'API Equivalent', 'Cancel', and 'Run'.

	Original	Updated
vCPUs	1	2
Cores per vCPU	1	2
Memory (GiB)	1	2

Disks

Disk 1: SCSI | Centos7HadoopMaster Existing

Disk 2: SCSI New

Value

Size (GiB) 3

Categories

Network Adapters

API Equivalent Cancel Run

Figure 23: Update VM Configuration

- Under the **VM Configuration** section, enter the change factor value in the **Updated** field for the **vCPUs**, **Core per vCPU**, and **Memory (GiB)**.

The ability to edit a VM configuration attribute and the maximum or minimum value to which the attribute can be updated depend on the app blueprint configuration. You can update only those VM configuration attributes that your blueprint developer has enabled for editing. The **Updated** field does not allow you to enter a value that is beyond the minimum or maximum value configured for the attribute.

- Under the **Disks** section, edit the following.

- Enter the value in the **Updated** field to increase the size of the existing disk.

Note: You cannot decrease the size of an existing disk.

You can click the delete icon to remove the existing disk.

- Enter the value in the **Updated** field to increase or decrease the size of any new disk. The updated value must be within the maximum or minimum value your blueprint developer has configured in the app blueprint.

You can click the delete icon to remove any new disk if your blueprint developer has enabled it in the app blueprint.

8. Under the **Categories** section, delete any existing categories from the app if your blueprint developer has enabled it in the app blueprint configuration.
9. Under the **Network Adapters** section, delete any existing NICs from the app if your blueprint developer has enabled it in the app blueprint configuration.
10. To launch the update configuration, click **Run**.

Updating the VM Configuration of an App on a VMware Platform

You can run the update configuration to modify parameters, such as VM configurations, controllers, disks, tags, and network adapters of a single-VM or multi-VM app running on a VMware platform.

About this task

For an overview of VM configuration update, see [Update VM Configurations of a Running App](#) on page 48.

Note:

- If there is a mismatch of the NICs or Network setting count after updating an app VM and you try to clone the app, the cloned app fails.
- You cannot add or delete the app properties simultaneously.

Before you begin

Ensure that your blueprint developer has added the update configuration before launching the app blueprint. For more information, see [Update Configuration for VM](#) in the *Self-Service Administration and Operations Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to update the VM configuration.
5. On the **Manage** tab, click the action corresponding to the update configuration. The Update screen for the app VM appears.
6. In the **VM Location** field, specify the location of the folder in which the VM must reside when you update. 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** checkbox and specify a folder name in the **Folder/Directory Name** field.
Select the **Delete empty folder** checkbox to delete the subfolder created within the specified location, in case the folder does not contain any VM resources. This option helps you to keep a clean folder structure.

7. Select the **CPU Hot Add** checkbox if you want to increase the VCPU count of a running VM. CPU resources cannot be added to an ESXi VM when it is powered on. The **CPU Hot Add** option lets you add CPU resources to a running VM.

Note: Before you use this feature, ensure that:

- You have the latest version of VMware Tools installed.
- You have the guest operating system that supports the CPU hot add functionality.
- Your virtual machine is compatible with ESXi 4.x or later.

With **CPU Hot Add**, you can only increase the vCPU count. If you decrease the vCPU count or update the Cores Per Socket, the VM will require a restart.

8. Update the **vCPUs** and **Cores Per Socket** count.
9. Select the **Memory Hot Plug** checkbox if you want to increase the memory while the virtual machine is turned on.

Note: Before you use this feature, ensure that:

- You have the latest version of VMware Tools installed.
- You have the guest operating system that supports the memory hot plug functionality.
- Your virtual machine is compatible with ESXi 4.x or later.

With **Memory Hot Plug**, you can only increase the memory. If you decrease the memory, the VM will require a restart.

10. Update the memory in the **Memory** field.
11. Under the **Controllers** section, you can add or update the **SCSI** or **SATA** controllers.

Note: You cannot delete a controller if it is attached to a disk.

12. Under the **Disks** section, click the + icon to add vDisks and do the following:

- a. Select the device type from the **Device Type** dropdown menu.
You can either select **CD-ROM** or **DISK**.
- b. Select the adapter type from the **Adapter Type** dropdown menu.
You can select **IDE** for CD-ROM or **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** dropdown menu 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** dropdown menu, 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.

Note:

- You can also edit the disk size and disk mode. However, decreasing the disk size of a saved configuration is not allowed.
- You can delete a saved disk. However, you cannot add and delete the disks simultaneously.

13. Under the **Network Adapter** section, click the + icon to add an NIC and configure the **Adapter Type** and **Networks** fields.

Note: You can only update the **Networks** field of an existing NIC.

14. 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.

15. Click **Update** to run the update configuration.

Updating the VM Configuration of an App on an AWS Platform

You can run the update configuration to modify parameters, such as instance type, IAM role, security groups, tags, and storage of a single-VM app running on an AWS platform.

About this task

For an overview of VM configuration update, see [Update VM Configurations of a Running App](#) on page 48.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to update the VM configuration.
5. On the **Manage** tab, click **Update VM Configuration**.
The Update screen for the app VM appears.
6. Under the **VM Configuration** section, update the instance type from the **Instance Type** dropdown menu.
The **Region**, **Availability Zone**, **Machine Image**, **Key Pairs**, and **VPC** fields are automatically selected. You cannot update these fields.
7. To update the IAM role, select the role from the **IAM Role** dropdown menu.
An IAM role is an AWS Identity and Access Management entity with permissions to make AWS service requests.
8. To enable the security group rule, select the **Include Classic Security Group**.
9. From the **Security Groups** dropdown menu, select security groups.
10. To add tags to the app, add the key and value pair in the **Key** and **Value** fields respectively.
11. To update the storage of the app, do the following under the **Storage** section:
 - a. For the existing storage, update the memory in GB in the **Size (GiB)** field and volume type of the storage device from the **Volume Type** dropdown menu for the root storage.
 - b. Click the **+** icon to add a storage and specify the device, size, and volume type.
12. Click **Update** to run the update configuration.

Update the VM Configuration of an App on an Azure Platform

You can run the update configuration to modify parameters, such as VM configurations, controllers, disks, or network adapters of a single-VM app running on an Azure platform.

About this task

For an overview of VM configuration update, see [Update VM Configurations of a Running App](#) on page 48.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.

4. Click **Manage** for the app in the Other Apps category for which you want to update the VM configuration.
5. On the **Overview** tab, click **Update** and then click **Update VM Configuration**. The Update screen for the app VM appears.
6. Under the **VM Configuration** section, update the hardware profile from the **Hardware Profile** dropdown menu.

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.

The **Instance Name**, **Resource Group**, **Location**, and **Availability Option** fields are automatically selected. You cannot update these fields.
7. Under the **Storage Profile** section, do the following:
 - a. Select the **Storage Type** and **Disk Caching Type** and specify the **Size** and **Disk LUN** for the existing data disk.
 - b. Click the + icon to add a data disk. Select the **Storage Type** and **Disk Caching Type** and specify the **Size** and **Disk LUN** for the new data disk.
8. Under the **Network Profile** section, click the + icon to add NICs as per your requirement and do the following for each NIC:
 - a. Select a security group from the **Security Group** dropdown menu.
 - b. Select a virtual network from the **Virtual Network** dropdown menu.
 - 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.

You can also update the **Security Group**, **Subnet**, and public or private IP config **Allocation Method** of the existing NIC.
9. To add tags to the app, add the key and value pair in the **Key** and **Value** fields respectively.
10. Click **Update** to run the update configuration.

Snapshot and Restore Overview

A snapshot preserves the state and data of an app virtual machine at a specific point in time. You can create a snapshot of a virtual machine at a particular point in time and restore from the snapshot to recreate the app from that time.

On a Nutanix or VMware platform, you can add the snapshot/restore configuration and use the snapshot and restore feature on both single-VM and multi-VM apps. Adding the configuration to the blueprint generates separate profile actions for snapshot and restore. For more information, see [Blueprint Configuration for Snapshots and Restore](#) in the *Self-Service Administration and Operations Guide*.

On AWS and Azure platforms, you can use the built-in snapshot and restore feature only on a single-VM app.

Note: The snapshot or restore feature is used only for the apps in the Other Apps category.

Snapshot and Restore on Nutanix or VMware Platform

Snapshot and restore of an app VM that runs on a Nutanix or VMware platform involves the following configurations and actions:

- Policy Definition for Snapshots
- Blueprint Configuration for Snapshots
- Application Launch with Snapshot Policy
- Snapshot Creation
- Snapshot Restore

Policy Definition for Snapshots

As a project admin, you define snapshot policies in a project. Snapshot policies help you define rules for taking snapshots of app VM. The policy determines the overall intent of the snapshot creation process.

For applications on a Nutanix Platform, you can configure your snapshot policy to manage your snapshots on a local cluster, on a remote cluster, or both. You can also configure the duration of managing those snapshots.

- Local Snapshots: When you select local snapshots in the policy and use the policy for snapshots, the snapshots of the VMs reside on the same cluster as that of the VMs.




Figure 24: Local Snapshots

- Remote Snapshots: When you select remote snapshot in the policy and use the policy for snapshots, the snapshots of the VMs running on the primary cluster are stored on a remote cluster. The primary cluster and the remote cluster must be associated with the same Prism Central. When you restore the VMs, the snapshots are restored to the primary cluster to bring up the VMs.



Figure 25: Remote Snapshots

Remote snapshots are particularly useful when your Prism Central has a compute-intensive cluster managing workloads and a storage-intensive cluster managing your data, snapshots, and so on.

For applications on a VMware Platform, you can configure the policy for local snapshots by selecting clusters or hosts. You, however, cannot configure the duration of managing those snapshots.

For information on creating a snapshot policy, see [Creating a Snapshot Policy](#) on page 138.

Blueprint Configuration for Snapshots

You define snapshot and restore configuration for each service in a blueprint. In case your multi-VM blueprint has multiple replicas of the service, you can configure the action to take snapshot for the first replica or for the entire replica set. For blueprints on a Nutanix platform, you can also configure the service to create snapshots locally or on a remote cluster.

The snapshot-restore definition of a service generates the snapshot configuration and its corresponding restore configuration. You can use these configurations to modify your snapshot and restore setup. The snapshot-restore definition also generates app profile actions that you use to create or restore snapshots. You can add more tasks and actions as part of your snapshot and restore to define actions you might want to take on your services. For example, shutting down the app and the VM before taking the snapshot or restarting the VM or services before a restore.

Note: The snapshot and restore configurations in a service are integrated to each other and cannot be managed individually.

For more information, see [Blueprint Configuration for Snapshots and Restore](#) in the *Self-Service Administration and Operations Guide*.

Application Launch with Snapshot Policy

You associate a policy defined in a project when you launch the app. Depending on the snapshot configuration that you provide in the blueprint, you can select the policy and the cluster in which the snapshot must be stored.

When you define remote snapshot in the blueprint (on a Nutanix platform), you can view all the policies that allow you to take a remote snapshot. You can select a policy and the corresponding clusters before you launch the app.

For more information, see [Launching a Blueprint](#) in the *Self-Service Administration and Operations Guide*.

Snapshot Creation

Similar to other profile actions, the profile actions for snapshot and restore appear on the **Manage** tab of an app. The snapshots that you create are listed under the **Recovery Points** tab. When you create multiple snapshots as part of one action, they appear as a snapshot group. You can expand the group to view the snapshots, their corresponding services, and location.

For information on snapshot creation on a Nutanix platform, see [Creating Snapshots on a Nutanix Platform](#) on page 57.

For information on snapshot creation on a VMware platform, see [Creating Snapshots on a VMware Platform](#) on page 60.

Snapshot Restore

Restore follows the same configuration that the snapshot has. To restore, you specify the variables and select applicable recovery points depending on the VM.

For information on restoring VM details on a Nutanix platform, see [Restoring VM Details from Snapshots on a Nutanix Platform](#) on page 59.

For information on restoring VM details on a VMware platform, see [Restoring VM Details from Snapshots on a VMware Platform](#) on page 62.

Creating Snapshots on a Nutanix Platform

You can create app-consistent or crash-consistent snapshots on a Nutanix platform. App-consistent or crash-consistent snapshots are used to capture and recover all of the VM and app-level details. App-consistent snapshots also capture data stored in the memory and transactions in process.

About this task

Note: Only crash-consistent snapshots are supported for multi-VM apps on a Nutanix platform.

Before you begin

Ensure that you have installed NGT Apps to take app-consistent snapshots. For more information, see [Installing NGT Apps](#) on page 67.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to create snapshots.
5. On the **Manage** tab, click the snapshot action you created for the app VM.
The **Run Action: Snapshot** window appears.

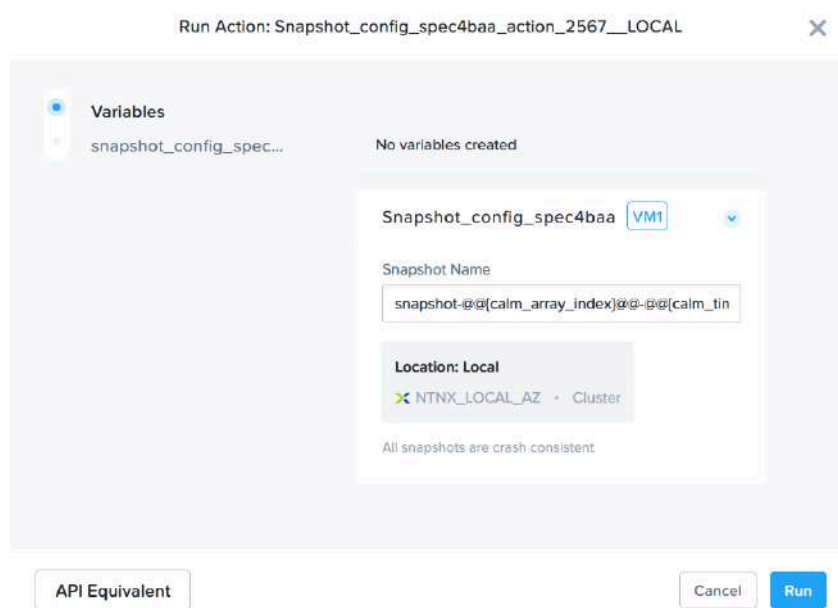


Figure 26: Snapshot Action

6. In the **Snapshot Name** field, enter a name for the snapshot.
You can use Self-Service macros to provide a unique name to the snapshot. For example, `snapshot-@@{calm_array_index}@@-@@{calm_time}@@`.
7. For single-VM apps, select **App consistent** for app-consistent snapshots or **Crash consistent** for crash-consistent snapshots.

Note:

- You can create app-consistent snapshots after you have installed NGT Apps with VSS service enabled. For more information about these snapshots, see [Conditions for](#)

[Application-consistent Snapshots](#) in the *Data Protection and Recovery with Prism Element Guide*.

- For multi-VM app, the default snapshot type is **Crash consistent**.

8. Click **Run**.

The saved snapshots are available under **Recovery Points** tab.

What to do next

You can recover the VM details for an app from the created snapshots. For more information, see [Restoring VM Details from Snapshots on a Nutanix Platform](#) on page 59.

Restoring VM Details from Snapshots on a Nutanix Platform

You can restore the VM details of an app after the host VM becomes unavailable.

Before you begin

Ensure that you have captured the snapshots for an app. For more details, see [Creating Snapshots on a Nutanix Platform](#) on page 57.

Note:

- A restored VM or a cloned VM does not have NGT service installed even if the snapshot or the source VM has NGT service installed.
- Restore operation for a VM fails if the snapshot is configured with static IP address and IP pool is not configured.
- When you perform a restore operation with a snapshot having static IP address configured, the restored VM comes up with a new IP address from the IP pool specified in IPAM. To ensure that the restored VM has the same static IP address as the old VM, remove the NIC that has this static IP address configured from the old VM, and attach the configuration to the new restored VM. If there is a failure during restore operation, perform an update operation on the VM to ensure that the VM is in valid state.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to restore the VM details from the snapshots.
5. On the **Manage** tab, click the restore action you created for the app.
The **Run Action: Restore** window appears.
6. Select a recovery point from the **Select Recovery Point** list.
The **Select Recovery Point** list shows all the snapshots taken for the app VM.
7. Click **Run**.
The app is restored from the snapshot in a new VM and the existing VM moves to power off state. If you selected the click the **Delete older VM after restore** checkbox while configuring the app blueprint, the existing VM is deleted after restoring the app VM.

Creating Snapshots on a VMware Platform

A snapshot preserves the state and data of a virtual machine at a specific point in time. You can create a snapshot of a virtual machine at any time and restore the snapshot to recreate the app from that time.

About this task

To create snapshots of an application VM on a VMware platform, you use the profile action generated in the application blueprint during snapshot-restore configuration. You can access the profile action on the **Manage** tab of the application.

Note: For VMware single-VM applications that you created prior to version 3.7, the Snapshot action is available as a Day-2 action on the **Manage** tab.

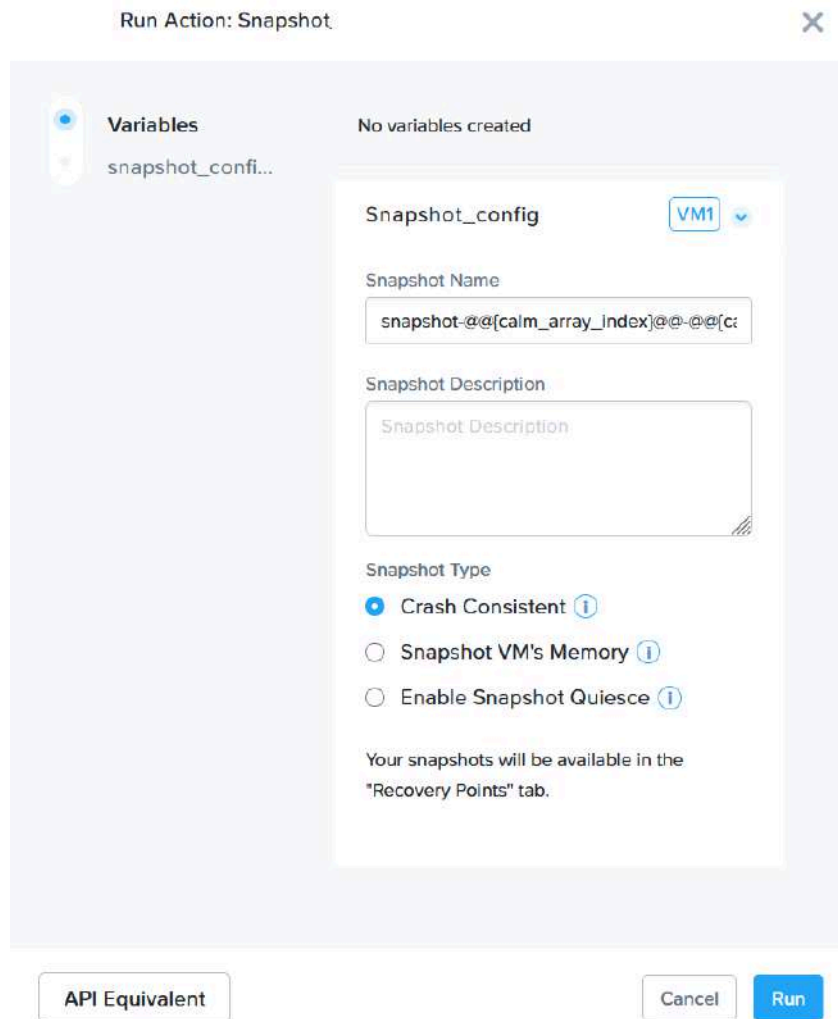
Before you begin

Ensure that the *VMware Tool* is installed and the VM is in powered on state to create the quiesced snapshots. For more information, see the [VMware Documentation](#).

Procedure

1. Log in to Prism Central as an administrator.
2. In the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **My Apps**.
4. Click **Manage** for the app in the Other Apps category for which you want to create snapshots.

5. On the **Manage** tab, click the snapshot action generated for the app VM.
The Run Action: Snapshot window appears.



The image shows a 'Run Action: Snapshot' dialog box. At the top, it says 'Run Action: Snapshot' with a close button. Below this, there's a 'Variables' section on the left with a dropdown menu showing 'snapshot_conf...' and a 'No variables created' message on the right. The main area is titled 'Snapshot_config' with a 'VM1' dropdown. It contains three input fields: 'Snapshot Name' with the text 'snapshot-@[calm_array_index]@[c...', 'Snapshot Description' with the placeholder 'Snapshot Description', and 'Snapshot Type' with three radio button options: 'Crash Consistent' (selected), 'Snapshot VM's Memory', and 'Enable Snapshot Quiesce'. Below these fields, a note states: 'Your snapshots will be available in the "Recovery Points" tab.' At the bottom, there are three buttons: 'API Equivalent', 'Cancel', and 'Run'.

Figure 27: Snapshot VMware

6. In the **Snapshot Name** field, enter a name for the snapshot. This step is optional.
7. In the **Snapshot Description** field, enter a brief description about the snapshot. This step is optional.

8. Do one of the following. This step is optional.

- » Select the **Crash Consistent** radio button to capture a snapshot that saves the data already written on the virtual disks, without capturing the data in memory or pending I/O operations in the snapshot.

Crash-consistent snapshots represent the VM state as it would be after an unexpected failure, similar to a system crash.

- » Select the **Snapshot VM's Memory** radio button to capture the memory state of the virtual machine and the power settings.

This option flushes the memory content of a VM to the disk as part of the snapshot, allowing the VM to go back to the exact state it was running in when the snapshot was taken. Memory snapshots take longer to create.

- » Select the **Enable Snapshot Quiesce** radio button to pause or alter the state of running processes on the virtual machine and take consistent and usable backup.

When you quiesce a virtual machine, VMware Tools quiesce the file system in the virtual machine. The quiesce operation pauses or alters the state of running processes on the virtual machine, especially processes that might modify information stored on the disk during a restore operation.

9. Click **Save**.

You can view the saved snapshots on the **Recovery Points** tab.

What to do next

You can recover the VM details for an app from the snapshots you created. For more information about recovering app-level information, see [Restoring VM Details from Snapshots on a VMware Platform](#) on page 62.

Restoring VM Details from Snapshots on a VMware Platform

You can restore the VM details of an app in case the host VM becomes unavailable.

About this task

Note: For VMware single-VM applications that are created prior to Self-Service version 3.7, the Restore action is available as a Day-2 action on the **Manage** tab.

Before you begin

Ensure that you have captured the snapshots for an app. For more information, see [Creating Snapshots on a VMware Platform](#) on page 60.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to restore the VM details from the snapshot you created.
5. On the **Manage** tab, click the restore action you created for the app.
6. Select a recovery point from the **Select Recovery Point** list.

The **Select Recovery Point** list shows all the snapshots taken for the app VM.

7. Click **Run**.

The app is restored from the snapshot.

Creating Snapshots on an AWS Platform

About this task

You can back up the data on your Amazon EBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. For more information, see [AWS Documentation](#).

Before you begin

Ensure that you have an AWS account with the required privileges to create a snapshot. For more information, see [Configuring AWS User Account with Minimum Privilege](#) and [AWS Policy Privileges](#) in the *Self-Service Administration and Operations Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the single-VM app in the Other Apps category for which you want to create snapshots.
5. On the **Overview** tab, click **Snapshot**.

Save Snapshot

An Amazon Machine Image (AMI) is created from the current VM.
Ensure that the AMI name is unique

AMI Name

ami-@[calm_array_index]@-@[calm_time]@

AMI Description

AMI Description

☐ No Reboot ⓘ

AMI will be available under "AMIs".

Cancel Save

Figure 28: Snapshot AWS

6. In the **AMI Name** field, enter a name for the snapshot.
7. In the **AMI Description** field, enter a brief description about the snapshot. This step is optional.
8. Select the **No Reboot** checkbox to avoid shutting down the Amazon EC2 instance before creating the image. This step is optional.

9. Click **Save**.

The **AMI** tab lists the snapshots you saved.

Restoring VM Details from Snapshots on an AWS Platform

You can restore the VM details of an app after the host VM becomes unavailable.

Before you begin

Ensure that you have captured the snapshots for the app VM.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the single-VM app in the Other Apps category for which you want to restore the VM details.
5. On the **AMIs** tab, click **Restore** next to a snapshot from which you want to restore the VM.
6. In the confirmation window, click **Confirm Restore**.
The restore action creates a new VM from the snapshot that has the same configuration as the source app with a different IP address.

Creating Snapshots on an Azure Platform

A snapshot of an app virtual machine on the Azure platform is a point-in-time copy of the operating system and the data disks that are associated with the VM. The snapshots you create can then be used to create a new VM with the same configurations as the source app VM.

Before you begin

Ensure that you have an Azure account with the required privileges to create a snapshot. For more information, see [Configuring Azure User Account with Minimum Privilege](#) in the *Self-Service Administration and Operations Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the single-VM app in the Other Apps category for which you want to create snapshots.

5. On the **Overview** tab, click **Snapshot**.

Save Snapshot

Ensure that the Snapshot name is unique

Snapshot Name

snap-@@(calm_time)@@

Snapshot Description

Snapshot Description

Storage Type

Standard HDD

Snapshot Pricing is based on the Snapshot Storage type

Created snapshots will be available in the "Snapshots" tab

Cancel Save

Figure 29: Snapshot Azure

6. In the **Snapshot Name** field, enter a name for the snapshot.
7. In the **Snapshot Description** field, enter a brief description about the snapshot. This step is optional.
8. From the **Snapshot Type** list, select the storage type to store your snapshot. Your options are:

- **Standard HDD**
- **Premium SSD**
- **Zone-redundant**

For more information on the storage type, see the [Azure Documentation](#).

9. Click **Save**.

You can track the progress of the snapshot creation process on the **Audit** tab. The snapshots are stored on the **Snapshots** tab.

Restoring VM Details from Snapshots on an Azure Platform

You can restore the VM details of an app after the host VM becomes unavailable. When you restore the VM details, a new VM is created using the snapshots of the disks.

About this task

The VM snapshot that you create on an Azure platform consists of the snapshot of operating system and data disks.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the single-VM app in the Other Apps category for which you want to restore the VM details.
5. On the **Snapshots** tab, click **Restore** next to the snapshot from which you want to restore the VM. The **Restore VM** dialog box appears.
6. Enter the restore name for the app VM.
7. Select the **Delete Previous VM** checkbox to delete the original VM from which the snapshot was created.
8. Click **Confirm Restore**.
The restore action creates a new VM from the snapshot that has the same configuration as the source app.

Deleting Snapshots

Perform the following procedure to delete the snapshots created for the VM under an app in the Other Apps category.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category for which you want to delete the snapshot.
5. Do one of the following.
 - » If your application is deployed on a Nutanix cluster, click the **Recovery Points** tab.
 - » If your application is deployed on a VMware platform, click the **Snapshots** tab.
 - » If your application is deployed on an AWS platform, click the **AMI** tab.
6. Click **Delete** next to the snapshot you want to delete.
7. Click **Confirm**.

Running a System-Level Action

System-level actions are predefined actions that you run on an app in the Other Apps category. Perform the following procedure to run a system-level action.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.

4. Click **Manage** for the app in the Other Apps category that you want to run the system generated action on.
5. Under the **Manage** tab, click one of the following type of action.

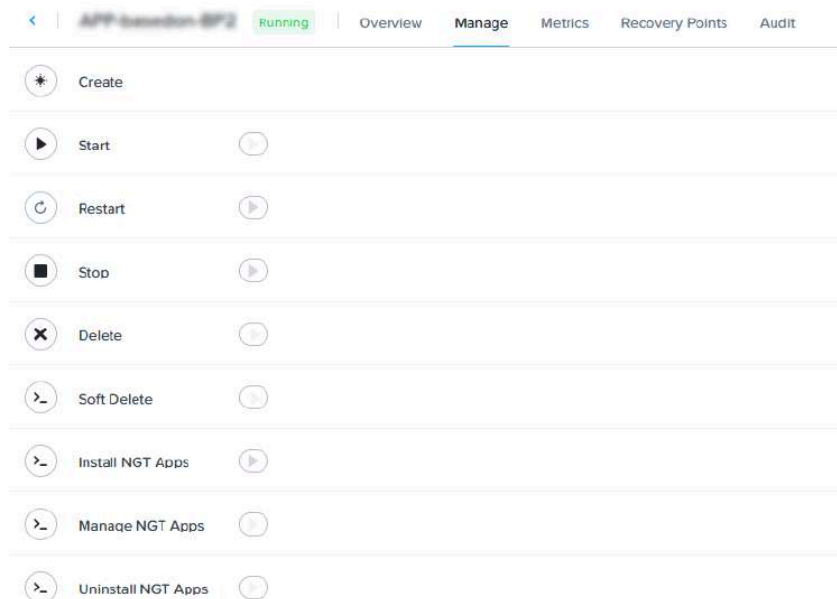


Figure 30: System-Level Action

- **Create:** Creates an app but cannot be performed once the blueprint is created.
- **Start:** Starts an app.
- **Restart:** Restarts an app.
- **Stop:** Stops an app.
- **Delete:** Deletes an app including the underlying VMs on the provider side.
- **Soft Delete:** Deletes the app from the Self-Service environment but does not delete the VMs on the provider side.
- **Install NGT Apps:** Installs NGT services for your app. To information on how to install NGT, see [Installing NGT Apps](#) on page 67.
- **Manage NGT Apps:** Manages NGT services for your app. You can enable or disable app-consistent and crash-consistent snapshots. For more information, see [Managing NGT Apps](#) on page 69.
- **Uninstall NGT Apps:** Uninstalls NGT services from the VM. For more information, see [Uninstalling NGT Apps](#) on page 69.

Installing NGT Apps

Nutanix Guest Tools (NGT) is a software bundle that you can install in a guest virtual machine (Microsoft Windows or Linux) to enable the advanced functionality provided by Nutanix. Perform the following procedure to install NGT services on your VM. NGT services are only applicable for AHV clusters.

About this task

Note: The Nutanix Guest Agent service is now upgraded to Python 3.6. For successful installation of NGT on Windows VMs, apply the Update for Universal C Runtime in Windows (Microsoft KB 2999226) to upgrade your Windows VMs to Python 3.6.

Before you begin

- Ensure that NGT requirements and limitations are met. For more information, see the [Prism Central Infrastructure Guide](#).
- Ensure that you have configured the cluster virtual IP address.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category that you want to install NGT on.
5. On the **Manage** tab, click the **Install NGT Apps** play button.

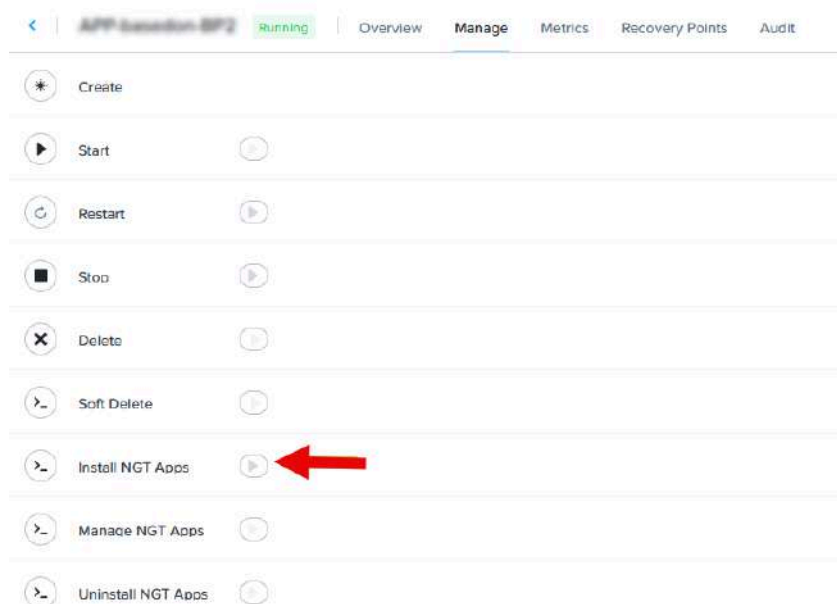


Figure 31: Install NGT

6. In the Install NGT Apps window, do the following:

- a. Select the **Enable Self Service Restore (SSR)** checkbox to restore desired files from the VM. This step is optional.
The self-service restore (SSR) allows virtual machine administrators to perform a self-service recovery from the Nutanix data protection snapshots with minimal administrator intervention.
- b. Select the **Enable Volume Snapshot Service (VSS)** checkbox to enable the Volume Snapshot Service (VSS). This step is optional.
VSS, also known as Shadow Copy, creates an app-consistent snapshot for a VM and is limited to consistency groups consisting of a single VM. Enabling VSS allows you to take app-consistent snapshots.
- c. To restart the VM after NGT installation, select **Restart as soon as the install is completed**.
You can select **Skip restart** to skip the restart of the VM after VM installation.
- d. Click **Enter Credentials**.
- e. Enter the guest operating username and password in the **Username** and **Password** fields.

7. Do one of the following.

- To mount the NGT on the VM and install it later, click **Skip and Mount**.
- If NGT is already mounted on a VM, click **Unmount** to unmount the NGT from the VM.
- To install NGT, click **Done**.

Managing NGT Apps

After you install NGT service on a VM, you can either enable or disable VSS and SSR services by using the **Manage NGT Apps** play button.

About this task

To know about VSS and SSR services, see the [Nutanix Guest Tools](#) section in the *Prism Element Web Console Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category that you want to manage NGT on.
5. On the **Manage** tab, click the **Manage NGT Apps** play button.
6. In the **Manage NGT Apps** window, click **Enable** or **Disable** to enable or disable self-service restore or volume snapshot service respectively.
7. Click **Confirm**.
The changes are saved and you can use the NGT services based on your selection.

Uninstalling NGT Apps

If you do not want to recover app details after the host VM becomes unavailable, uninstall the NGT app. Perform the following procedure to uninstall NGT services for your app.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category that you want to uninstall NGT from.
5. On the **Manage** tab, click the **Uninstall NGT Apps** play button.
A confirmation message appears to uninstall NGT.
6. Click **Uninstall**.

Updating Actions and Credentials of an App

You can add or update the credential, custom actions, post delete tasks, or package uninstall tasks from the **Overview** tab of a single-VM app.

About this task

Note: Dynamic variables are runtime editable by default, but you cannot mark variable as runtime editable if you add the variables while updating an app.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the single-VM app in the Other Apps category for which you want to update the credential or actions.
5. From the **Update** dropdown menu, select **Update Actions and Credentials**.
6. In the Credentials & Connection section, click **Edit** to open the Credentials & Connection page.
7. To add a credential, click **Add Credential** and do the following.
 - a. In the **Name** field, type a name for the credential.
 - b. Under the **Type** section, select the type of credential that you want to add.
 - » **Static:** Credentials store keys and passwords in the credential objects that are contained in the blueprints.
 - » **Dynamic:** Credentials fetch keys and passwords from an external credential store that you integrate as the credential provider.
 - c. In the **Username** field, type the user name.
For dynamic credentials, specify the @@(username)@@ that you defined while configuring the credential provider.

Note: A dynamic credential provider definition requires username and secret. The secret variable is defined by default when you configure your credential provider. However, you must configure a

runbook in the dynamic credential provider definition for the username variable before you use the variable in different entities.

- d. Select either **Password** or **SSH Private Key** as the secret type.
- e. Do one of the following to configure the secret type.
 - » If you selected **Static** as the credential type and **Password** as the secret type, then type the password in the **Password** field.
 - » If you selected **Static** as the credential type and **SSH Private Key** as the secret type, then enter or upload the key in the **SSH Private Key** field.
 - » If you selected **Dynamic** as the credential type and **Password** or **SSH Private Key** as the secret type, then select a credential provider in the **Provider** field. After you select the provider, verify or edit the attributes defined for the credential provider.

If the private key is password protected, click **+Add Passphrase** to provide the passphrase. For dynamic credentials, you must configure a runbook in the dynamic credential provider definition for the passphrase variable and then use the `@@{passphrase}@@` variable.

8. If you want this credential as your default credential, select the **Use as default** checkbox.
9. To delete an existing credential, click **Delete** next to the credential in the credential list.

Note: You can also update the user name or password of an existing credential. However, if you have logged on as an operator, you can only update the password.

10. On the **Connection** tab, select the credential from the **Credentials** list to update the credential and check the logon status after updating the app.

Note: You can update the credential to check the logon status only if you have enabled the **Check log-in upon create** option while configuring the blueprint.

11. Click **Done**.
12. To add a post delete task for the app, click **Edit** in the PostDelete section.

For more information, see [Adding a Pre-create, Post-create, or Post-delete Task](#) in the *Self-Service Administration and Operations Guide*. This step is optional.

13. To create a task to uninstall a package, click **Edit** next to the Package section and do the following.
- Click **+ Task**.
 - Enter the task name in the **Task Name** field.
 - To create the type of task, select the type from the **Type** dropdown menu.
The available options are:
 - **Execute**: To create the **Execute** task type, see [Creating an Execute Task](#) in the *Self-Service Administration and Operations Guide*.
 - **Set Variable**: To create the **Set Variable** task type, see [Creating an Set Variable Task](#) in the *Self-Service Administration and Operations Guide*.
 - **HTTP**: To create the **HTTP** task type, see [Creating an HTTP Task](#) in the *Self-Service Administration and Operations Guide*.
 - **Delay** : To create the **Delay** task type, see [Creating a Delay Task](#) in the *Self-Service Administration and Operations Guide*.
 - To add variables to the post delete task, click the **Package Uninstall Variables** tab.
 - In the **Variables** pane, click the **+** icon to add variable types
 - In the **Name** field, enter a name for the variable.
 - From the **Data Types** dropdown menu, select one of the base type variable or import a custom library variable type.
 - If you have selected a base type variable, configure all the variable parameters.
For more information, see [Creating Variable Types](#) in the *Self-Service Administration and Operations Guide*.
 - If you have imported a custom variable type, all the variable parameters are automatically filled.
 - Select the **Secret** checkbox if you want to hide the value of the variable.
 - To save the package uninstall task, click **Done**.
 - To establish a connection between tasks, click **Add Connector** and use the arrow to create connection between tasks.
 - To delete a task, click **Delete** next to the task.
You can delete a task only while adding a new task. If you are updating the existing task, you cannot delete the task.

14. To add another action to the app, click **+ Add Action** under the Actions section and do the following. This step is optional.
 - a. Click **+ Add Task**.
The task inspector panel is displayed.
 - b. In the task inspector panel, click **Task**.
 - c. Enter the task name in the **Task Name** field.
 - d. Select the type of tasks from the **Type** dropdown menu.
The available options are:
 - **Execute**: Use this task type to run eScripts on the VM. To create the **Execute** task type, see [Creating an Execute Task](#) in the *Self-Service Administration and Operations Guide*.
 - **Set Variable**: Use this task to change variables in a blueprint. To create the **Set Variable** task type, see [Creating an Set Variable Task](#) in the *Self-Service Administration and Operations Guide*.
 - **HTTP Task**: Use this task type to query REST calls from a URL. An HTTP task supports GET, PUT, POST, and DELETE methods. To create the **HTTP Task** type, see [Creating an HTTP Task](#) in the *Self-Service Administration and Operations Guide*.
 - **Delay** : Use this task type to set a time interval between two tasks or actions. To create the **Delay** task type, see [Creating a Delay Task](#) in the *Self-Service Administration and Operations Guide*.
 - e. To add another task, click **Add Task** in the task editor area.
 - f. To establish a connection between tasks, click **Add Connector** and use the arrow to create connection between tasks.
 - g. To delete a task, click **Delete** next to the task.
 - h. To add variables to the task, click the **Variables** tab.
 - i. In the **Variables** pane, click the **+** icon to add variable types in your blueprint.
 - j. In the **Name** field, enter a name for the variable.
 - k. From the **Data Types** dropdown menu, select one of the base type variable or import a custom library variable type.
 - l. If you have selected a base type variable, configure all the variable parameters.
For more information, see [Creating Variable Types](#) in the *Self-Service Administration and Operations Guide*.
 - m. If you have imported a custom variable type, all the variable parameters are auto filled.
 - n. Select the **Secret** checkbox if you want to hide the value of the variable.
 - o. To save the task, click **Done**.
15. To save the updated credentials and tasks for the app, click **Update**.

Creating an Image on a Nutanix Platform

An image is a template for creating new instance or VM. You can create images from an existing single-VM or multi-VM app running on a Nutanix platform. Perform the following procedure to create an image from an existing app.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category from which you want to create an image.
5. To create an image on a single-VM app, click **Create Image** on the **Overview** tab of the My Apps page.

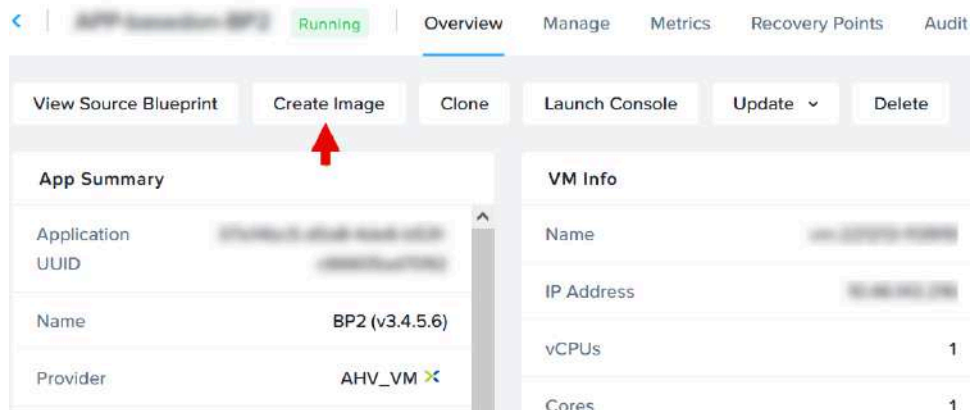


Figure 32: Create Image - Single VM Application

6. To create an image on a multi-VM app, select the service on the **Services** tab, and then click **Create Image** in the Inspector Panel.




Figure 33: Create Image - Single VM App

7. Select the checkbox next to the disks from which you want to create images.
8. Under the **Image Details** section, type a name and a description for the new image in the **Name** and **Description** fields respectively.
If you have selected multiple disk images, repeat the steps for all the **Image Details** sections.
9. Click **Save**.
The new image is created and available in the **Image** dropdown menu under the **VM Configuration** section when you create a blueprint. You can use the image to create a single-VM or multi-VM app.

Cloning an App

About this task

You can clone an app in the Other Apps category if you are using Nutanix, VMware, or AWS as the provider for the app. The cloned app has the same VM configuration as the source app from which it is cloned.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category that you want to clone.
5. On the **Overview** tab, click **Clone**.

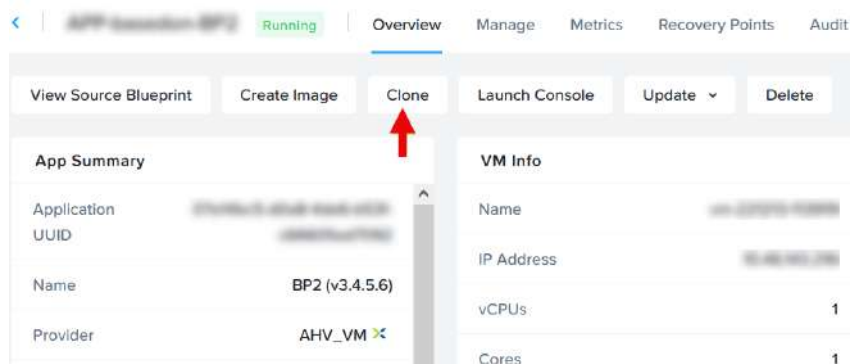


Figure 34: Clone

6. In the Clone window, do the following:
 - a. Enter a name for the clone in the **Cloned Application Name** field.
 - b. In the **Description** field, enter a brief description about the cloned app.
 - c. Click **Save**.

Note: In a Nutanix cluster, a restored VM or a cloned VM has NGT service installed if the snapshot or the source VM has NGT service installed.

What to do next

After you successfully cloned an app, you can view the link to the cloned app in the audit log of the source app on the **Audit** tab. You can click the link of the cloned app to view the cloned app details. To view the source app, click the **Clone From** field on the **Overview** tab.

Deleting an App

You can delete any redundant apps in the Other Apps category.

Procedure

1. Log in to Prism Central as an administrator.

2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category that you want to delete.
5. On the **Manage** tab, click the **Delete** play button.
6. In the Delete Application confirmation window, click **Delete**.

Accessing Web SSH Console

For VM management, web SSH or web RDP sessions are supported in both single-VM and multi-VM apps for local and remote Prism Central. The sessions are supported for both Linux and Windows VM, including the VMs that are within the VPCs.

About this task

Following protocols are supported for authentication:

- PROTOCOL_RDP: Standard RDP Security.
- PROTOCOL_SSL: TLS 1.0, 1.1, or 1.2.
- PROTOCOL_HYBRID: Credential Security Support Provider (CredSSP).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. Click **Manage** for the app in the Other Apps category.

5. Do one of the following:

- On a Single VM app, click **Open Terminal** at the top of the page.

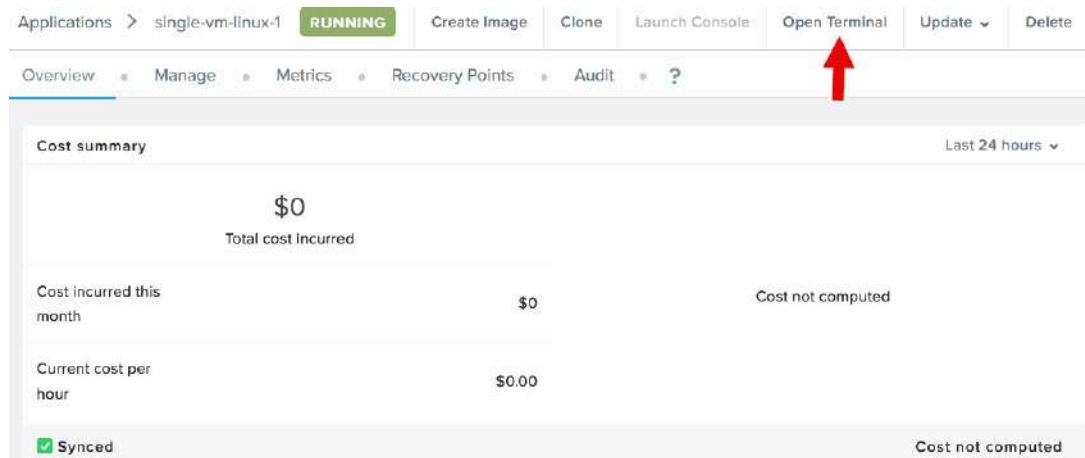


Figure 35: Single-VM Open Terminal

- On a multi-VM app, click the **Services** tab, click the service, and then click **Open Terminal** in the inspector panel.




Figure 36: Web SSH Console

For the **Open Terminal** option to be available, ensure that your image has the RDP server enabled.

PROJECT MANAGEMENT

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 Management](#) on page 143 and [Role-based Access Control in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

Projects provide a consistent experience when you access them from Prism Central or from Self-Service. However, the project features and configurations differ depending on whether you have deployed Self-Service in your Prism Central or not.

When Self-Service is not deployed, the project dashboard typically appears as follows:

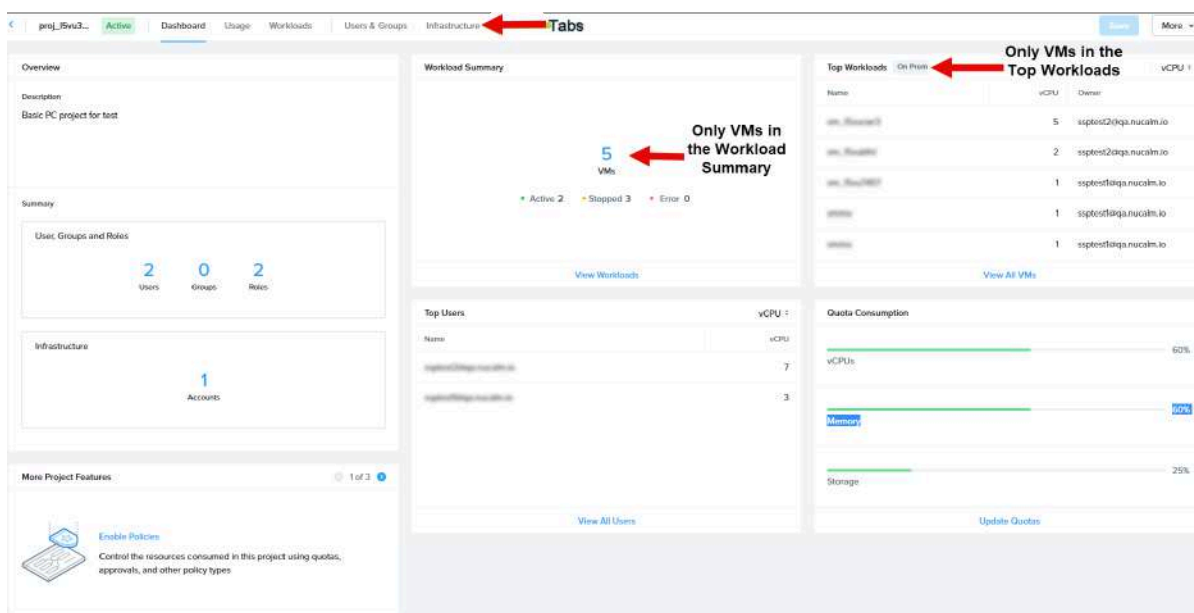


Figure 37: Projects Dashboard When Self-Service is Not Deployed

When Self-Service is deployed, the project dashboard typically appears as follows:

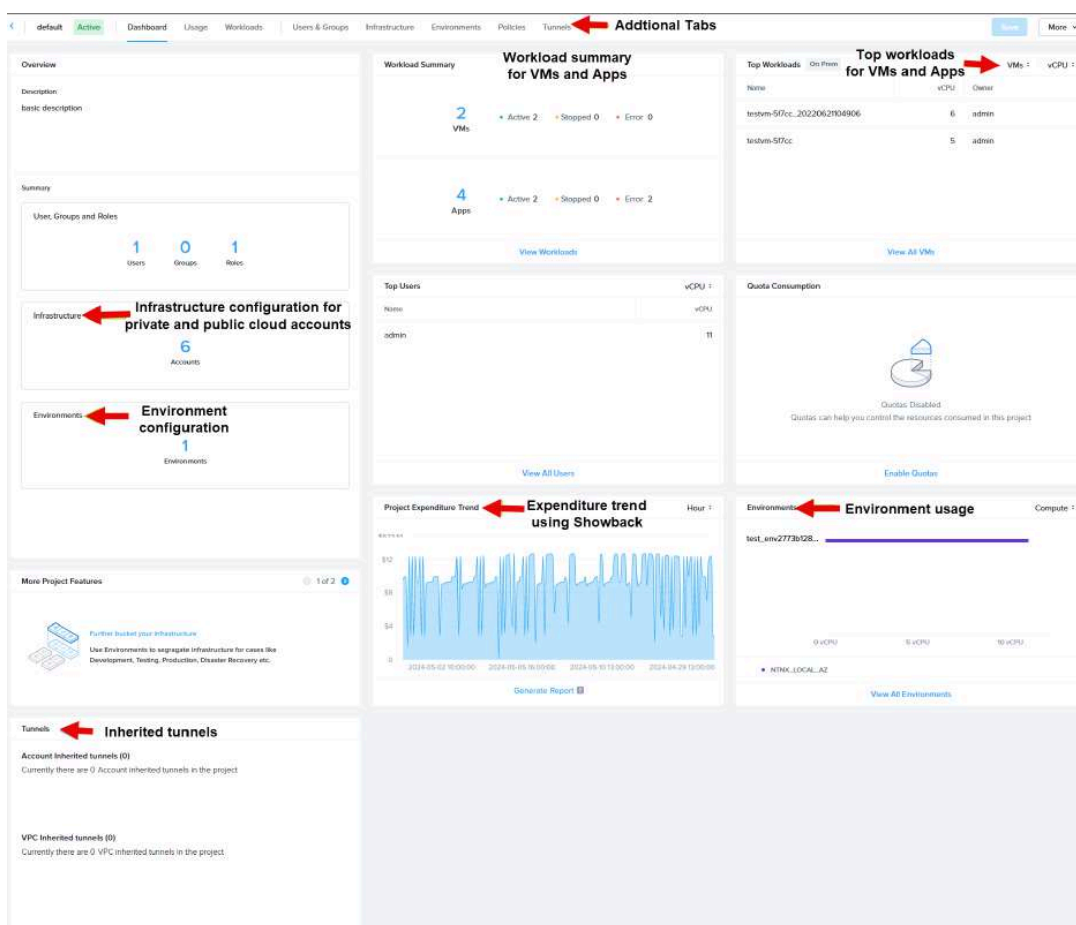


Figure 38: Projects Dashboard When Self-Service is Deployed

The following table lists the project features that are available with Prism Central when Self-Service is disabled and when Self-Service is enabled.

Table 15: Project Feature Availability in Prism Central

Project features	When Self-Service is disabled	When Self-Service 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
Allowing VPC Overlay subnets	Yes	Yes
Adding other private or public cloud accounts (ESXi, AWS, Azure or GCP)	No	Yes

Project features	When Self-Service is disabled	When Self-Service 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. For more information, see [VM Management](#) in the *Prism Central Infrastructure Guide*.

Note: With the refactored project, Self-Service 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 Self-Service-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.

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, select **Admin Center** in the Application Switcher, and click **Projects** in the navigation bar.

Name	Status	Infrastructure	Users	Apps	VMs	vCPU	Memory (GiB)	Storage (GiB)	Cost
test_vmware_pr object67f	Active	2 On Prem	4	1	0	1	1	40.05	\$12.97
test_vmware_pr object85d	Active	2 On Prem	4	1	0	1	1	40.05	\$12.97
test_vmware_pr object85d	Active	2 On Prem	4	1	0	0	0	0	\$2.98

Figure 39: Projects Summary

The following table describes the definitions that appear for a project on the Projects Summary page.

Table 16: Projects List Fields

Definition	Description
Name	Displays the project name.

Definition	Description
Status	Displays the status of the project.
Infrastructure	Displays the number of on-prem and cloud accounts added to the project.
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 select the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 17: Filter Pane Fields

Parameter	Description
Name	Filters based on the project name. Select a condition from the dropdown 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 Settings 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 also includes a **Create Project** to create projects. For more information, see [Creating a Project](#) on page 91.

The **Actions** menu appears when you select a project on the Projects Summary page. The menu allows you to delete the selected project.

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 82.

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 name on the Project Summary page. For information on the Project Summary page, see [Projects Summary View](#) on page 80.

A few tabs on the Project Details page appear only when you enable Self-Service in your Prism Central instance. The Project Details page has the following tabs:

- Dashboard
- Usage
- Workloads
- Users & Groups
- Infrastructure
- Environments (appears when you enable Self-Service in your Prism Central instance)
- Policies (appears when you enable Self-Service in your Prism Central instance)
- Tunnels (appears when you enable Self-Service in your Prism Central instance)

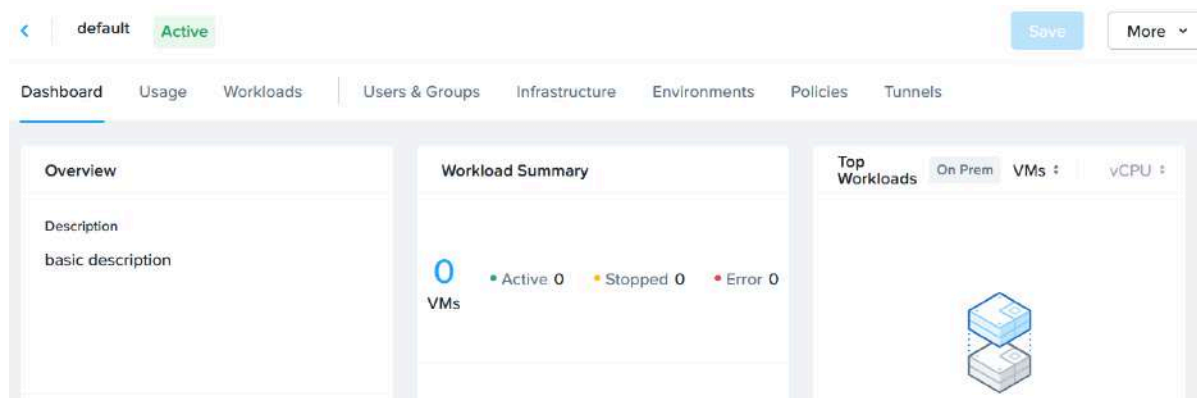


Figure 40: Project Details

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 Self-Service enablement (see [Deploying Self-Service](#) on page 16) and the overall project configuration and usage data.

When Self-Service is disabled, the project dashboard typically appears as follows:

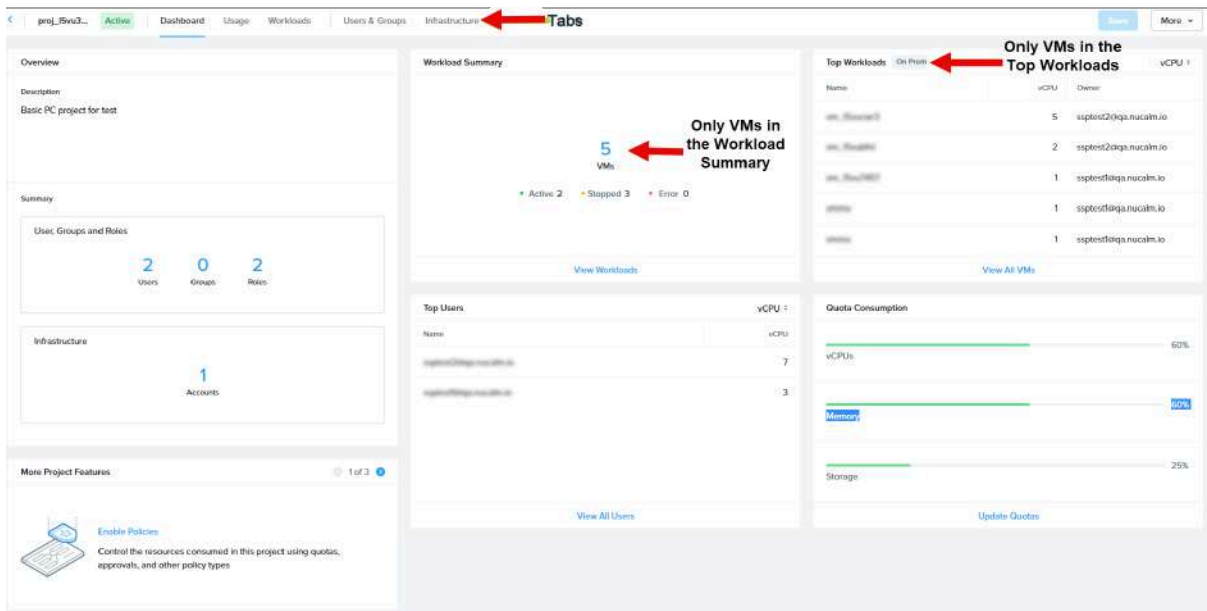


Figure 41: Project Dashboard when Self-Service is Disabled

When Self-Service is enabled, the project dashboard typically appears as follows:

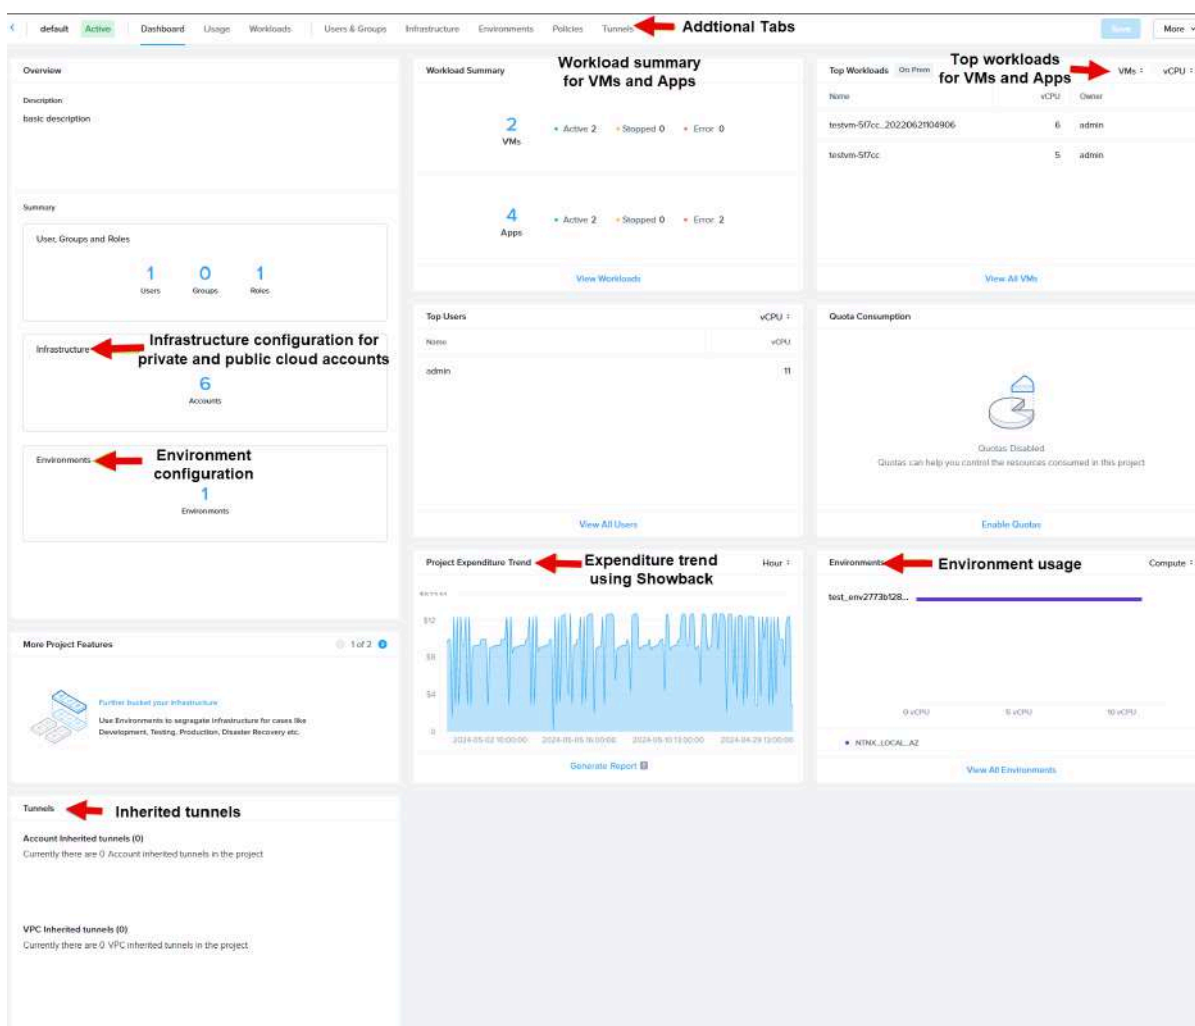


Figure 42: Project Dashboard when Self-Service is Enabled

The dashboard has the following tiles:

Table 18: Project Dashboard

Tiles	Description
Overview	<p>The Overview tile has the following sections.</p> <ul style="list-style-type: none">• 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. <p>For more information about the infrastructure (accounts), see Infrastructure in Projects on page 99.</p> <p>For more information about environments, see Environments in Projects on page 110.</p>
Workload Summary	<p>Provides information related to the overall workloads associated with the project.</p> <p>When Self-Service is disabled, the tile only 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 you hover your mouse, you can also view the number of local and remote VMs. The local VM count indicates the number of VMs associated with the local AHV clusters added to the project. The remote VM count indicates the number of VMs associated with the remote PC (remote clusters) added to the project.</p> <p>When Self-Service 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.</p> <p>You can click View Workloads to navigate to the Workloads tab and view the details of the VMs or applications.</p> <p>For more information, see Project Workloads Tab on page 89.</p>
Top Workloads	<p>Provides information related to the workloads that consume the maximum resources.</p> <p>When Self-Service is disabled, the tile displays the VMs that consume the maximum resources within the project. When Self-Service is enabled, the tile also displays the applications (on-prem AHV and ESXi) that consume the maximum resources within the project.</p> <p>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.</p> <p>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.</p> <p>For more information, see Project Workloads Tab on page 89.</p>

Tiles	Description
More Project Features	<p>Displays recommendations and options to enable additional features for your projects. You can click the options available on the tile to enable Self-Service and configure these additional features for your project.</p> <ul style="list-style-type: none"> • 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 in the <i>Self-Service Administration and Operations Guide</i>. • 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 Self-Service. You can configure your public clouds such as AWS, Azure, or GCP in Self-Service. For more information, see Provider Account Settings in Self-Service in the <i>Self-Service Administration and Operations Guide</i>. • Provision, scale, and manage multi-VM applications: You can design complex blueprints that span multiple environments and teams using Nutanix Self-Service. For more information, see Self-Service Blueprints Overview in the <i>Self-Service Administration and Operations Guide</i>. • 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 110.
Top Users	<p>Display the users that consume the maximum resources within the project.</p> <p>You can filter the users based on the vCPU, memory, or storage consumption using the resource switcher on the tile.</p> <p>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.</p> <p>For more information, see Project Users & Groups Tab on page 89.</p>

Tiles	Description
Quota Consumption	<p>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.</p> <p>When Self-Service 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.</p> <p>To get better quota consumption data with the upgraded quota capabilities, it is recommended to enable Self-Service and then the policy engine. You do not need to license Self-Service to enable policy engine.</p> <p>For more information about policy engine, see Policy Engine Overview in the <i>Self-Service Administration and Operations Guide</i>.</p>
Project Expenditure Trends	<p>Display a graph to report the expenditure trend within the project from an hour, day, or month viewpoint when Self-Service is enabled in your Prism Central instance.</p> <p>Expenditure trend is generated when you enable Showback in Self-Service. For more information, see Showback in the <i>Self-Service Administration and Operations Guide</i>.</p> <p>You can use the Generate Report option to generate a showback report. The showback report contains the cost data of the active applications within the project over the last 90 days. You can also use the Generate Report option to download the generated report.</p>
Environments	<p>Displays the environment usage within the project when you have configured environments in your project and have launched applications using those environments.</p> <p>You can sort the environment usage on the basis of compute, memory, or usage using the resource switcher on the tile.</p> <p>You can click View all Environments to navigate to the Environments tab and view or update your environment configuration.</p> <p>For more information on environment configuration, see Environments in Projects on page 110.</p>
Tunnels	<p>Displays the number of tunnels inherited from the account and the VPCs associated with the project.</p>

Project Environments Tab

The **Environments** tab appears when you enable Self-Service 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 110.

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 19: 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 Settings 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 110.

You can click the vertical ellipsis next to an environment to update or delete the environment.

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 Self-Service 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 Self-Service-Disabled Prism Central](#) on page 105.

When Self-Service 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 Self-Service. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 100.

For more information about account configuration in Self-Service, see [Provider Account Settings in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

Project Policies Tab

The **Policies** tab appears when you enable Self-Service 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 135 or [Managing Quota Limits for Projects](#) on page 136.

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 138.

Project Tunnels Tab

The **Tunnels** tab appears when you enable Self-Service 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](#) in the *Self-Service Administration and Operations Guide*.

Project Users & Groups Tab

The **Users & Groups** tab displays a list of users and groups assigned to the project.

Table 20: Users & Groups

Parameter	Description
Name	Displays the Active Directory name (typically in the form of <i>name@domain</i>) 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 95.

Project Workloads Tab

The **Workloads** tab displays the overall workloads associated with the project. The Workloads tab has the following tabs that you can access in the left pane.

Table 21: Tabs Within the Workloads Tab

Tabs	Description
VMs	<p>The VMs tab shows the total number of VMs and the number of local and remote VMs associated with the project.</p> <p>The local VM count indicates the number of VMs associated with the local AHV clusters that you add to the project. The remote VM count indicates the number of VMs associated with the remote PC (remote clusters) that you add to the project.</p> <p>The VMs tab displays the details of only those VMs that are associated with the local AHV clusters. You can view the following details:</p> <ul style="list-style-type: none">• 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.

Tabs	Description
Apps	<p>This tab appears only when Self-Service 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.</p> <ul style="list-style-type: none"> • 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 Self-Service.

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.

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 Self-Service in your Prism Central instance to configure application management features.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Click **+ Create Project** on the Project Summary page to create the project.

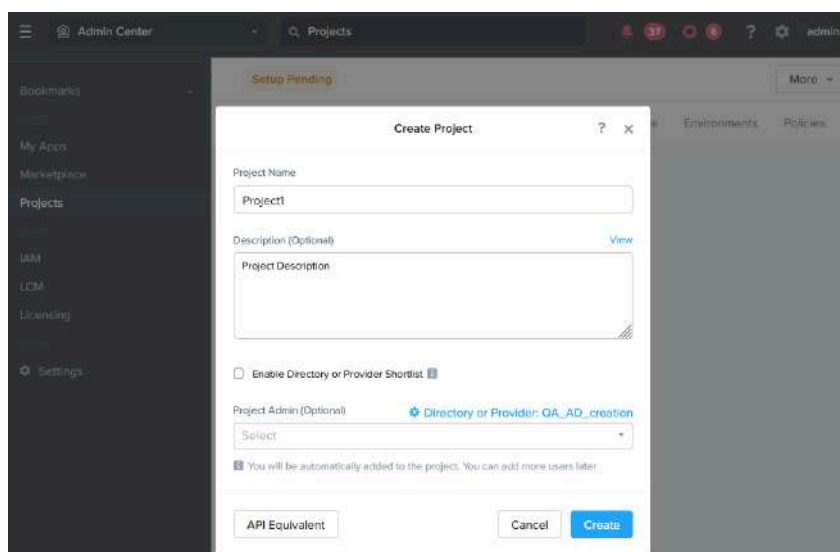


Figure 43: Create Project window

5. Type a name for the project in the **Project Name** field.
6. Type a description for the project in the **Description** field.
7. Select the **Enable Directory or Provider Shortlist** checkbox to shortlist Active Directories (ADs) or Identity Providers (IDPs) in the project and then select ADs or IDPs in the **Directory or Provider** dropdown menu.

After you shortlist ADs or IDPs in the project, project admins can add users or groups only from the shortlisted ADs or IDPs. By default, the option to shortlist ADs and IDPs remains disabled.

You can use this feature to ensure that the project do not show all ADs or IDPs during multi-tenancy.

8. Select an admin for the project from the **Project Admin** dropdown menu. This step is optional.
Type the first few letters of the project admin to get the list of matching results in the dropdown menu.

Note: The project automatically adds you as a Project Admin when you create it. You can add other users after you configure the basic setup of the project. For more information, see [Adding Users to a Project](#) on page 95.

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 Settings 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.

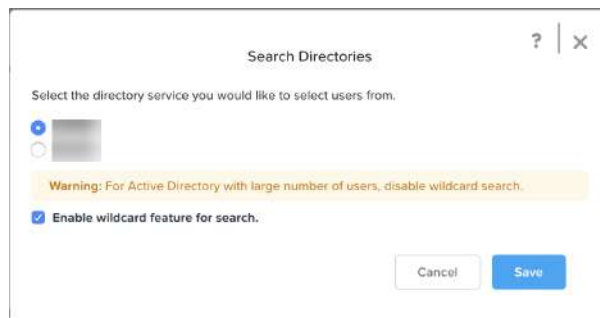


Figure 44: Search Directories Window

- Click **Save**.

Note: Local users are not supported in a project. You can only add an admin from your configured directory or identity provider.

9. Check the **Allow Collaboration** checkbox to allow project users to collaboratively manage VMs and applications within the project.

The **Allow Collaboration** checkbox appears when you add your first user to the project. By default, the **Allow Collaboration** checkbox is selected and enables a project user to view and manage VMs and applications of other users in the same project. If you clear the **Allow Collaboration** checkbox, project 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.

10. Click **Create**.
The **Dashboard** tab for the project appears.

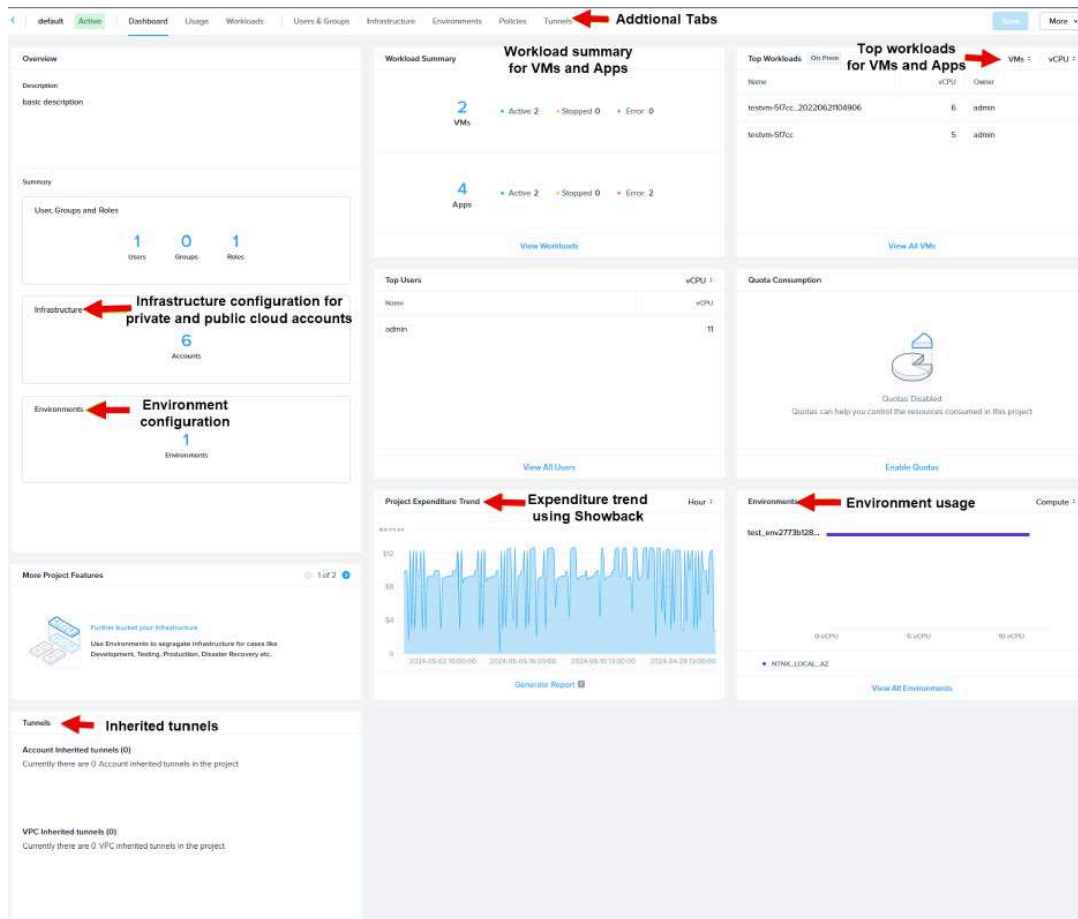


Figure 45: Project Dashboard

11. Add users or groups to your project on the **Users & Groups** tab. For more information, see [Adding Users to a Project](#) on page 95.
12. Add infrastructure to your project on the **Infrastructure** tab. For more information, see [Adding Infrastructure in a Self-Service-Disabled Prism Central](#) on page 105.
13. If you have enabled Self-Service 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 Self-Service-Enabled Prism Central](#) on page 100.
 - 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 110.
 - c. Use the **Policies** tab to define your quota and snapshot policies. For more information, see [Quota Policy Overview](#) on page 135 and [Creating a Snapshot Policy](#) on page 138.
14. Click **Save** on the Project Details page.

Adding Users to a Project

Use projects to assign roles to users or user groups. Based on the role assigned, users are allowed to configure and manage resources within the organization.

About this task

For information on the user roles, see [Roles Management](#) on page 143 and [Role-Based Access Control in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

Note: Local users are not supported in a project. You can only add users from your configured directory or identity provider.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click **+ Create Project** to create a new project and add users to the project. For information on creating a project, see [Creating a Project](#) on page 91.
 - » Click a project name in the list of existing projects to add users to that project.

5. Do the following to shortlist Active Directories (ADs) or Identity Providers (IDPs) in a project.

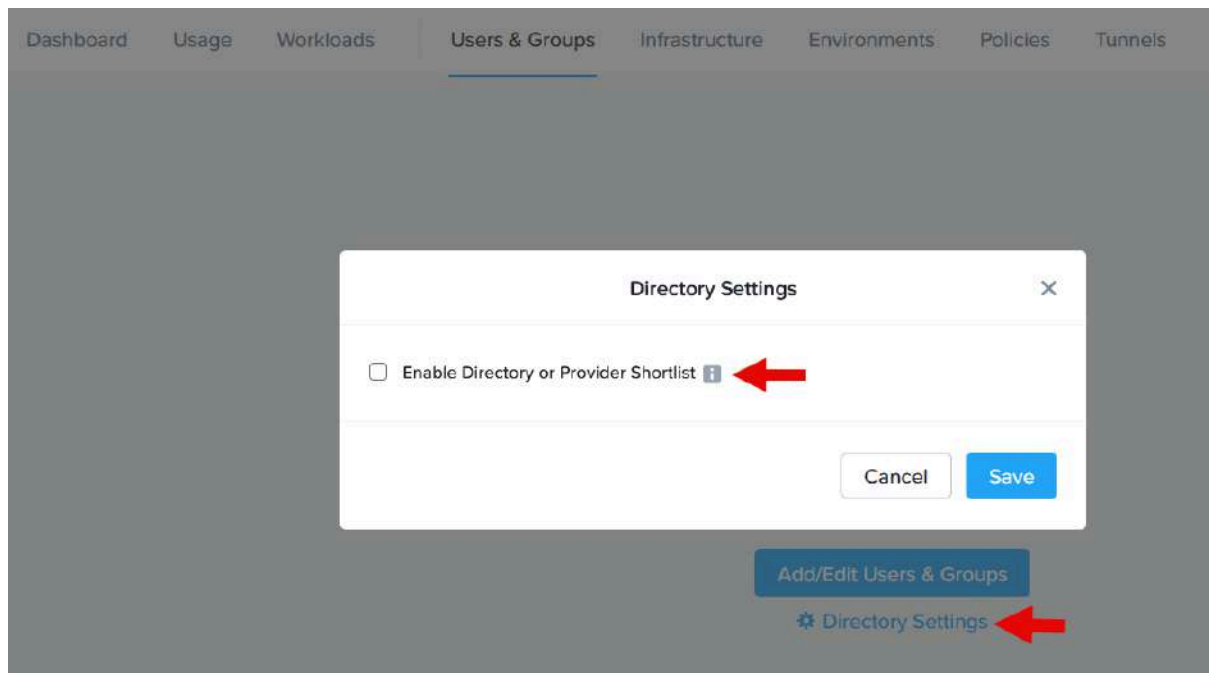


Figure 46: Directory Settings

- a. On the **Users & Groups** tab, click **Directory Settings**.
 - b. On the Directory Settings window, select the **Enable Directory or Provider Shortlist** checkbox. By default, the option to shortlist ADs and IDPs remains disabled.
 - c. From the **Directories or Providers** dropdown menu, select the ADs or IDPs that you want to shortlist in the Project.
You can select multiple ADs or IDPs to shortlist. Shortlisting ADs or IDPs in the project allows project admins to add only those users or user groups to the project that belong to the shortlisted ADs or IDPs. This feature is specifically helpful when you want to ensure that the project do not show all ADs or IDPs during multi-tenancy.
 - d. Click **Save**.
6. On the **Users & Groups** tab, click **Add/Edit Users & Groups**.
The **Add/Edit Users & Groups** window appears.
 7. Click the **Directory and Provider** settings icon.
The **Search Directories** window appears with options to select the AD or IDP from which you want to add the user or user group to the project.
You can add users and user groups from multiple ADs or IDPs configured in your Prism Central and shortlisted in the project. After you add users or user groups from one AD or IDP, you can select

another AD or IDP in the Search Directories window and add users or user groups belonging to the that AD or IDP.

Figure 47 shows the 'Add/Edit Users & Groups' dialog box. It features a '+ Add User' link, a search configuration tooltip, a table with 'Name' and 'Role' columns, and 'Cancel' and 'Save Users and Project' buttons.

Figure 47: Search Directories

8. Select the radio button for the directory service or identity provider that you want to use to add users or user groups.

Figure 48 shows the 'Search Directories' dialog box. It includes a selection area for directory services, a warning message about wildcard search for Active Directory, and 'Cancel' and 'Save' buttons.

Figure 48: Search Directories Window

9. To save the changes and close the **Search Directories** window, click **Save**.

10. 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 the **Role** columns.

b. In the **Name** column, enter the directory name or identify provider name of the user or a user group belonging to the AD or IDP you selected.

Note:

- Type the name of the AD user to search and add the user.
- For IDP user names, enter the complete IDP user name in the field and press Enter.
- The **Name** column accepts user group names. If you have created an IDP user group with a name, type the user group name in the **Name** column to add the group to the project.

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** checkbox appears when you add the first user to your project. By default, the **Allow Collaboration** checkbox is selected and enables a project user to view and manage VMs and applications of other users in the same project. If you clear the **Allow Collaboration** checkbox, project 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 **Save Users and Project**.

The **Users, Groups & Roles** section on the Overview tile displays the number of users you added to the project.

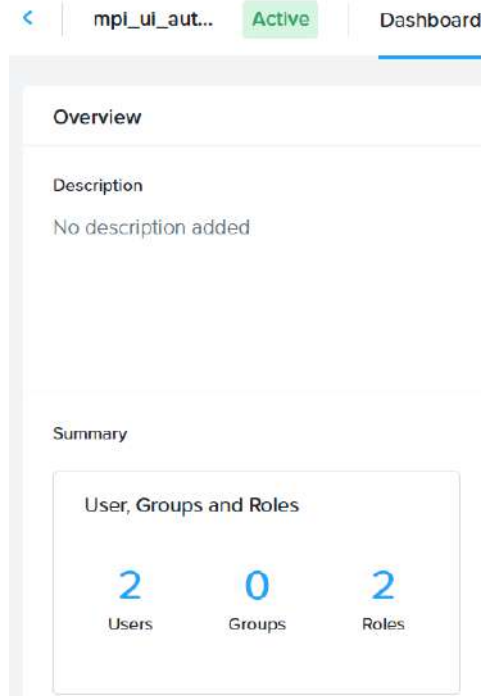


Figure 49: Users, Groups & Roles

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.

Infrastructure in Projects

Infrastructure allows you to add accounts to your project. When Self-Service is disabled, you can add only the local Nutanix account to your project. For more information, see [Adding Infrastructure in a Self-Service-Disabled Prism Central](#) on page 105). However, you can enable Self-Service and add multiple accounts of the same provider or different providers you configured in Self-Service to your projects. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 100).

Infrastructure supports accounts of the following providers:

- Nutanix
- VMware

- AWS
- Azure
- GCP
- Kubernetes

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.

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 accounts, you must enable policy engine from the Admin Center Settings. For more information, see [Enabling Policy Engine](#) on page 149 and [Managing Quota Limits for Projects](#) on page 136.

Adding Infrastructure in a Self-Service-Enabled Prism Central

When Self-Service is enabled in your Prism Central, you can add multiple provider accounts or credential provider accounts that you configured in Self-Service to your project. You can also define resource quota limits for quota checks for Nutanix and VMware accounts within the project if you enable policy engine in your Prism Central instance.

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 in Self-Service. For information on configuring a provider account, see [Provider Account Settings in Self-Service](#) in the *Self-Service Administration and Operations Guide*. For information on configuring a credential provider account, see [Credential Settings in Self-Service](#) in the *Self-Service Administration and Operations Guide*.
- For resource quota limit definition, ensure that you have enabled the policy engine in the Admin Center Settings. For information on enabling policy engine, see [Enabling Policy Engine](#) on page 149.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.

4. Do one of the following:
 - » Click **+Create Project** to create a new project and add providers to the project. For information on creating a project, see [Creating a Project](#) on page 91.
 - » Click a project name in the list of existing projects to add providers to that project.
5. Click the **Infrastructure** tab.
6. Click **Add Infrastructure** and select a provider account or a credential provider account from the list.

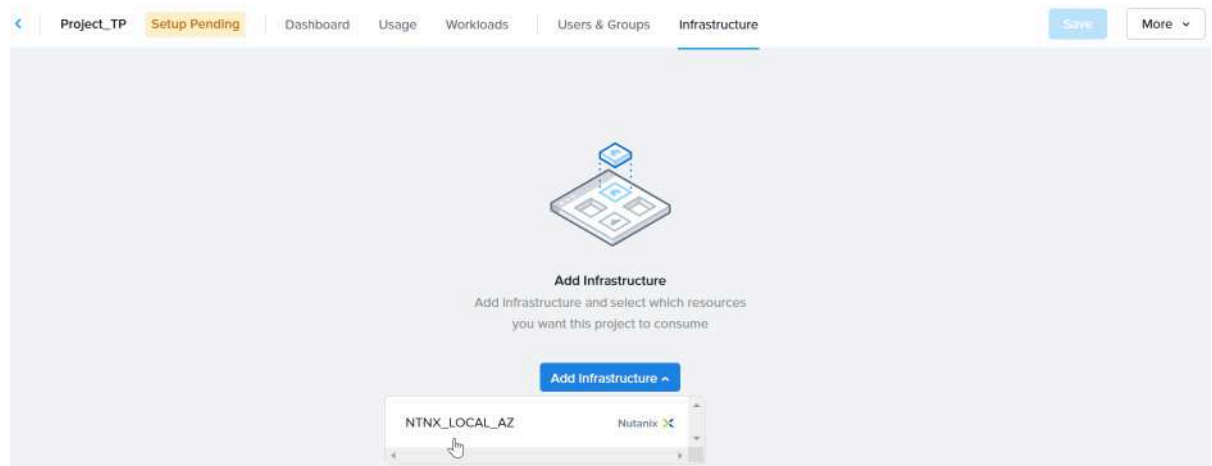


Figure 50: Add Infrastructure

7. If you selected a credential provider account, then go to step 11.

8. If you selected a Nutanix account, then do the following to configure resources in the project.
 - 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.

1 Select Clusters And VLANs 2 Select VPCs & Subnets 3 Confirm And Select Default

Select clusters to be added to this project

select_cluster_group_8000000000000000

Selected Cluster can only be removed once its whitelisted subnets are removed

Select VLANs for the above clusters (Optional)

2 VLANs selected

Name	VLAN ID
<input checked="" type="checkbox"/> vlan.800	800
<input checked="" type="checkbox"/> vlan1211	0

Cancel Select VPCs & Subnets >

Figure 51: 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 **Select VPCs & Subnets**.

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.

1 Select Clusters And VLANs 2 Select VPCs & Subnets 3 Confirm And Select Default

Select VPCs to view overlay subnets below (Optional)

vpc_name_2

Selected VPC can only be removed once its whitelisted subnets are removed

Select overlay subnets (Optional)

vpc_name_2
2 Overlay subnets selected

Name	IP Prefix
<input checked="" type="checkbox"/> vpc_subnet_1	10.10.10.0/24
<input checked="" type="checkbox"/> vpc_subnet_2	20.20.20.0/24

1 - 2 of 2 10

◀ Select Clusters and VLANs Cancel Confirm and Select Default ▶

Figure 52: 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](#) in the *Self-Service Administration and Operations Guide*.

- 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.

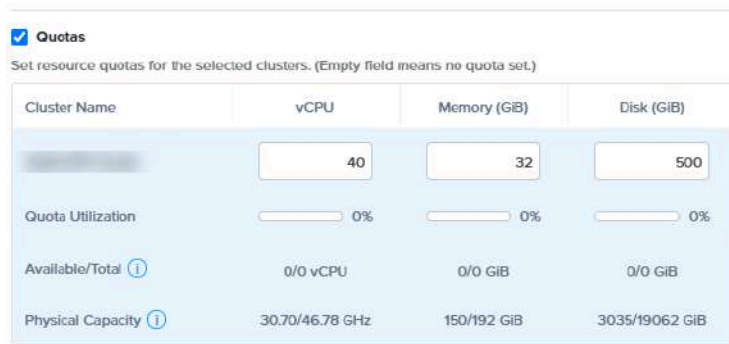
Note: You can select one or more AHV clusters. You can allow one or more subnets per cluster. When you configure a blueprint, only the allowed networks and clusters appear for the user to select during network configuration. If the network selection is a runtime attribute, only the allowed networks and clusters are available to update while launching a blueprint.

9. If you selected a Nutanix account, then do the following to configure services in the project.

A Nutanix service refers to a specific software component or functionality that Nutanix provides to enhance the capabilities, performance, and management of the Nutanix infrastructure and software stack. Currently, you can connect a Nutanix Database Service (NDB) account to your Nutanix account in Self-Service to leverage the database management capabilities of NDB with the application management workflows. For more information, see [Connecting NDB to an Account](#) in the *Self-Service Administration and Operations Guide*.

 - a. In the Nutanix Services section, click **Configure Services**.

- b. Under Add or Remote Accounts, select the NDB accounts that you want to allow in the project.
The list displays the name and the Server IP addresses of the NDB accounts that you connected to the selected Nutanix account.
 - c. Click **Next**.
 - d. View the list of selected accounts, and click **Confirm**.
The **NDB** section on the **Infrastructure** tab displays the list of selected NDB accounts with their Server IP addresses.
- 10.** If you selected a Nutanix or VMware account, you can define resource quota limits.
The options to define resource quota limits appear only when you enable policy engine in the Admin Center Settings of Prism Central. For information on enabling policy engine, see [Enabling Policy Engine](#) on page 149.
- a. Select the **Quotas**.
The **vCPU**, **Memory**, and **Disk** fields are enabled.



☒ **Quotas**

Set resource quotas for the selected clusters. (Empty field means no quota set.)

Cluster Name	vCPU	Memory (GiB)	Disk (GiB)
Cluster1	40	32	500
Quota Utilization	0%	0%	0%
Available/Total ⓘ	0/0 vCPU	0/0 GiB	0/0 GiB
Physical Capacity ⓘ	30.70/46.78 GHz	150/192 GiB	3035/19062 GiB

Figure 53: 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.

11. Click **Save**.

The Overview tile on the **Dashboard** tab displays the number of accounts you added to the project.

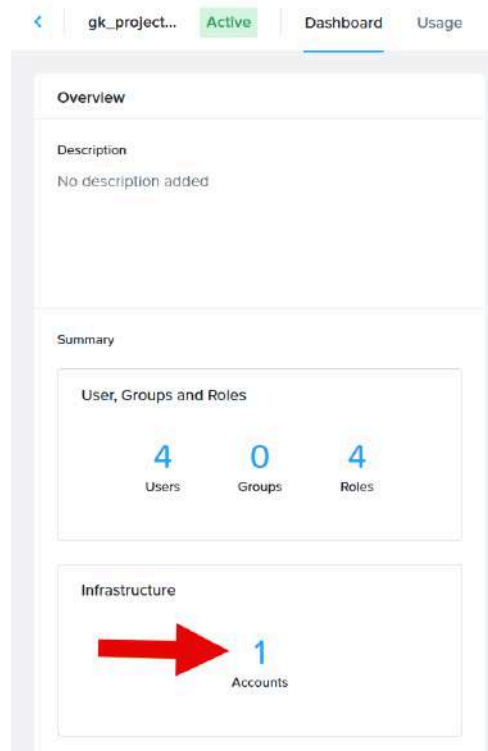


Figure 54: Accounts in a Project

Adding Infrastructure in a Self-Service-Disabled Prism Central

When Self-Service 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 within the project if you enable policy engine in your Prism Central instance.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. 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 on creating a project, see [Creating a Project](#) on page 91.
 - » Click a project name in the list of existing projects to add the Nutanix account to that project.
5. Click the **Infrastructure** tab.

6. Click **Add Infrastructure** and select the local Nutanix account.

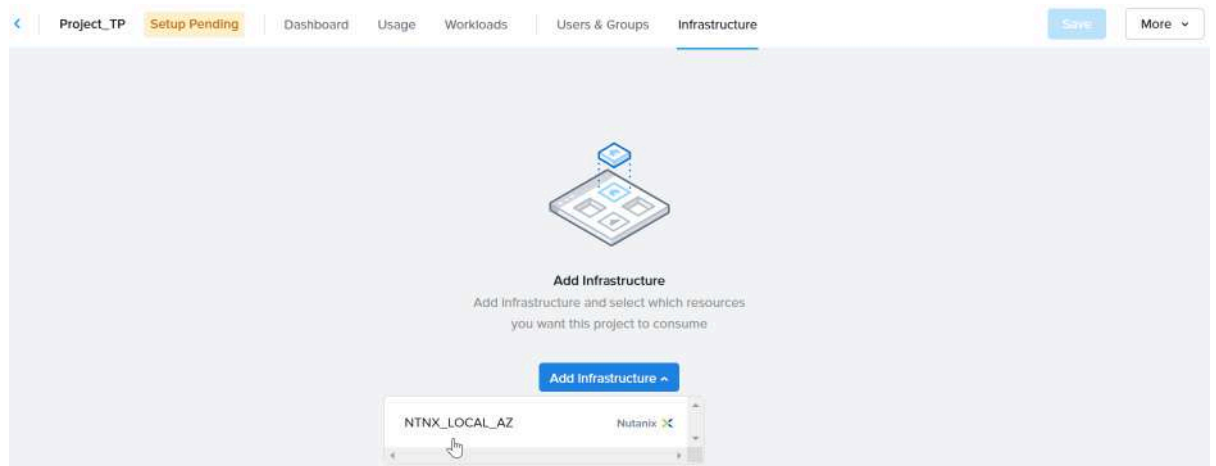


Figure 55: Add Infrastructure

7. 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** dropdown menu.

1 Select Clusters And VLANs 2 Select VPCs & Subnets 3 Confirm And Select Default

Select clusters to be added to this project

vdc1-cluster-prod-PROD000001

Selected Cluster can only be removed once its whitelisted subnets are removed

Select VLANs for the above clusters (Optional)

2 VLANs selected

Name	VLAN ID
<input checked="" type="checkbox"/> vlan.800	800
<input checked="" type="checkbox"/> vlan1211	0

Cancel Select VPCs & Subnets >

Figure 56: 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 **Select VPCs & Subnets**.
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.

1 Select Clusters And VLANs
2 Select VPCs & Subnets
3 Confirm And Select Default

Select VPCs to view overlay subnets below (Optional)

vpc_name_2

Selected VPC can only be removed once its whitelisted subnets are removed

Select overlay subnets (Optional)

vpc_name_2

2 Overlay subnets selected

1 - 2 of 2

10

Name	IP Prefix
<input checked="" type="checkbox"/> vpc_subnet_1	10.10.10.0/24
<input checked="" type="checkbox"/> vpc_subnet_2	20.20.20.0/24

Select Clusters and vLANs

Cancel

Confirm and Select Default

Figure 57: 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 Tunnel](#) in the *Self-Service Administration and Operations Guide*.

- f. Under the Select overlay subnets section, select the overlay subnets that you want to allow in the project. This step is optional.
 - g. Click **Confirm and Select Default**.
 - h. Select a default subnet.
Default subnets are used when you create a VM from Prism Central.
 - i. Click **Confirm**.
8. To define resource quota limits for the Nutanix account, under the Quotas section, enter quota values for **vCPU**, **Storage**, and **Memory**.
The options to define resource quota limits appear only when you enable policy engine in the Admin Center Settings of Prism Central. For information on enabling policy engine, see [Enabling Policy Engine](#) on page 149.

9. Click **Save**.

The Overview tile on the **Dashboard** tab displays the number of accounts you added to the project.

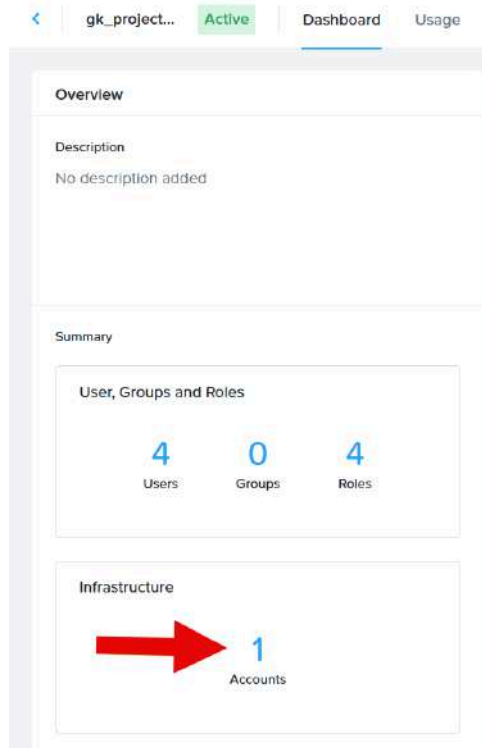


Figure 58: Accounts in a Project

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. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. 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.
5. Update the field values as required and then click **Save**.

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 window appears displaying the association of the project with different applications or blueprints.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. On the Project Summary page, select the checkbox adjacent to the project that you want to delete.
5. From the **Action** dropdown menu, select **Delete**.
Prism Central checks the association of the project with entities such as applications, blueprints, runbooks, endpoints, and so on. In case the project is associated with any entities, the details of the associated entities are displayed in a window. You cannot delete the project until you remove the project from these entities. In case the project is not associated with the entities, a confirmation window appears where you can confirm to delete the project.
6. In the confirmation window, click **Delete**.

Environments in Projects

If Self-Service 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 Self-Service 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. For more information, see [Configuring Environments in a Project](#) on page 110.

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. For more information, 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 in Self-Service or launch marketplace applications in the Admin Center. You can configure multiple environments in your project.

Before you begin

- Ensure that you have enabled Self-Service in your Prism Central instance. For more information, see [Deploying Self-Service](#) on page 16.
- 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 91 and [Infrastructure in Projects](#) on page 99.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. On the Projects Summary page, click the project in which you want to configure the environment.
5. Click the **Environments** tab.
6. Click **Create Environment**.
7. On the **General** tab, do the following:
 - a. Type 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** checkbox to use the environment as a default environment to launch applications for the project.

The **Set as default environment** checkbox is selected by default for the first environment you configure in your project.
 - d. Click **Next**.
8. On the **Infrastructure** tab, do the following:
 - a. Click the **Select Infrastructure** dropdown menu, and select an account.

The **Select Infrastructure** dropdown menu shows the accounts that you added to your project. For more information, see [Infrastructure in Projects](#) on page 99.

If you have already selected an account, then you can click **+ Add Infrastructure** to select and add more than one provider accounts to the environment.
 - b. Expand the **VM Configuration** section to configure the virtual machine details for the environment.
 - For information on how to configure a VM for Nutanix, see [Configuring Nutanix Environment](#) on page 112.
 - For information on how to configure a VM for AWS, see [Configuring AWS Environment](#) on page 121.
 - For information on how to configure a VM for VMware, see [Configuring VMware Environment](#) on page 116.
 - For information on how to configure a VM for GCP, see [Configuring GCP Environment](#) on page 130.
 - For information on how to configure a VM for Azure, see [Configuring Azure Environment](#) on page 124.
 - c. Click **Next**.
9. On the **Credentials** tab, add credentials for your environment. For more information, see [Adding Credentials to the Environment](#) on page 134.

10. 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.

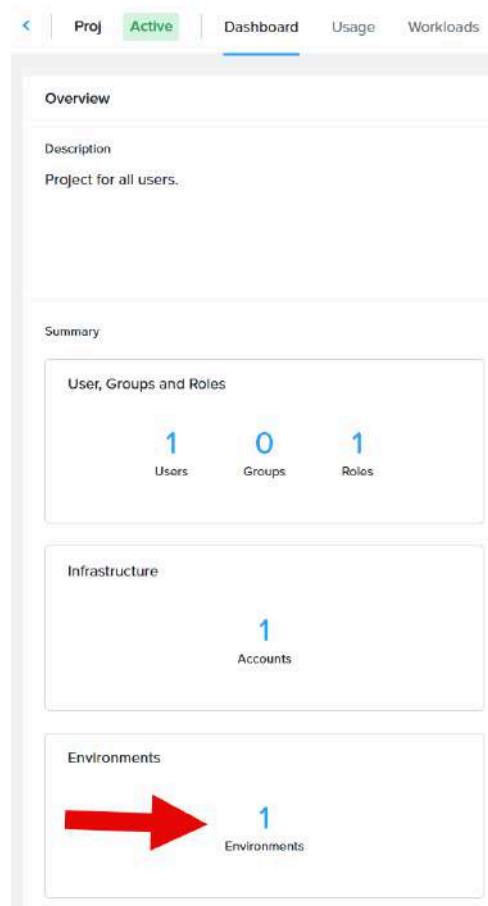


Figure 59: Environments on the Projects Dashboard

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 [Creating a Project](#) on page 91 and [Configuring Environments in a Project](#) on page 110.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, 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 **Configure Resources** 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.

The screenshot displays the 'NTNX_LOCAL_AZ' environment configuration page. At the top, there's a header with the environment name and a 'Configure Resources' button. Below this is the 'Resource Configurations' section, which includes a breadcrumb trail: '1 Cluster > 4 Total Subnets > Default Subnet: vlan.800 (auto_cluster_prod_4ff7f52803a9)'. The 'Clusters (1)' section shows a dropdown menu with 'auto_cluster_prod_4ff7f52803a9' selected. The 'VLANs (2)' section shows a table with 2 rows. The 'Overlay Subnets (2)' section shows a table with 2 rows. At the bottom, the 'VM Configuration' section is expanded, showing a message: 'Required for launching blueprints from marketplace' and 'Quickly launch blueprints in this environment with this configuration for VMs'. The status is 'Not configured'.

Name	VLAN ID	Cluster
vlan.800	800	auto_cluster_prod_4ff7f52803a9
vlan1211	0	auto_cluster_prod_4ff7f52803a9

Name	VPC Name	IP Prefix
vpc_subnet_1	vpc_name_1	10.10.10.0/24
vpc_subnet_2	vpc_name_1	20.20.20.0/24

VM Configuration Required for launching blueprints from marketplace
Quickly launch blueprints in this environment with this configuration for VMs
Not configured

Figure 60: VM Configuration

4. In the **Cluster** dropdown menu, select the cluster where you want to place the VM.

The **Cluster** dropdown menu 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** dropdown menu. 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 Self-Service macros to provide a unique name to the VM. For example, `vm-@@{calm_time}@@`. For more information on Self-Service 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**.

The screenshot shows the 'Guest Customization' configuration panel. At the top, the 'Guest Customization' checkbox is checked. Below it, the 'Type' section has two radio buttons: 'Cloud-init' and 'Sysprep', with 'Sysprep' selected. The 'Install Type' is a dropdown menu currently showing 'Select...'. The 'Script' section features a large dark text area with a small '1' in the top left corner and icons for 'Add' and 'Download' in the top right. Below the script area, the 'Join a Domain' checkbox is checked. Under 'Domain Setting', there is a blue informational message: 'Domain credentials and DNS details if already present in Sysprep XML will take precedence.' The 'Domain Name' field contains the text 'Windows Server Domain'. Another blue informational message states: 'Specify a user who has permission to add the computer to domain'. The 'Credential' field is a dropdown menu showing 'Credential'. The 'DNS IP' field contains the text 'DNS IP'. The 'DNS Search Path' field contains the text 'DNS Search Path'.

Figure 61: 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** checkbox 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** dropdown menu. 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.

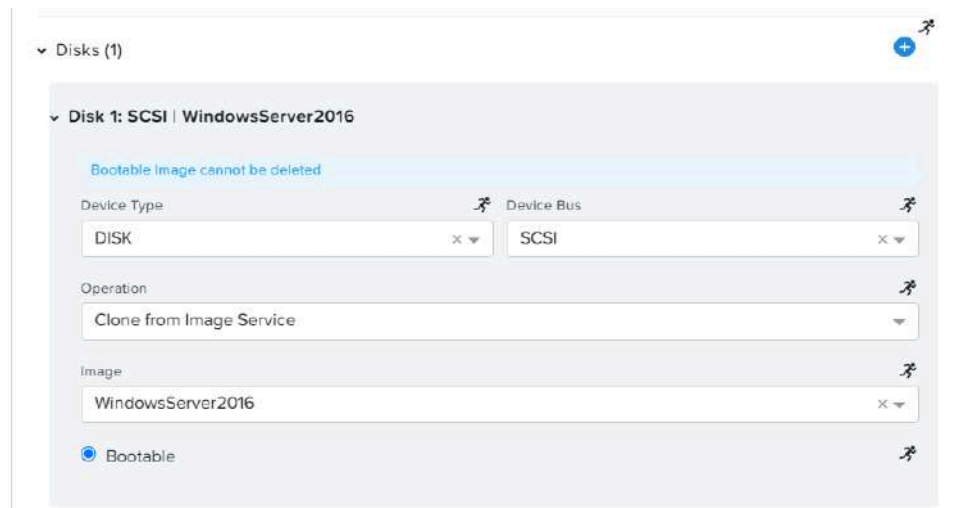


Figure 62: Disks

- a. Select the device for the image from the **Device Type** dropdown menu.
You can select either **CD-ROM** or **Disk**.
- b. Select the device bus from the **Device Bus** dropdown menu.
You can select **IDE** or **SATA** for CD-ROM and **SCSI**, **IDE**, **PCI**, or **SATA** for DISK.
- c. From the **Operations** dropdown menu, 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 the [Image Management](#) section in the *Prism Central Infrastructure Guide*.
- f. Select the **Bootable** checkbox 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** dropdown menu, select the GPU vendor.
 - b. From the **Device ID** dropdown menu, select the device ID of the GPU.
 - c. From the **Mode** dropdown menu, select the GPU mode.
11. Under the **Categories** section, select a category in the **Key: Value** dropdown menu.
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 63: NIC

The **NIC** dropdown menu 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** dropdown menu.

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 133.
14. Click **Next**.
15. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 134. 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 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 [Creating a Project](#) on page 91 and [Configuring Environments in a Project](#) on page 110.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, 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.

VM Configuration Required for launching blueprints from marketplace

Quickly launch blueprints in this environment with this configuration for VMs

Not configured

Windows Linux

VM Configuration

☒ **Compute DRS Mode**

Cluster

Template

☒ **Storage DRS Mode**

Datastore Cluster

Test

VM Location

VM Location

Existing path

☐ Create a folder/directory structure here

☒ Delete empty folder

@@[calm_array_index]@@ and @@(calm_time)@@ have been added to ensure a unique VM name. You can append or prepend a custom string to this value for identification.

Instance Name

vm-@[calm_array_index]@@-@[calm_time]@@

☐ **CPU Hot Add**

Sockets

2

Cores Per Socket

1

vCPU (Sockets * Cores Per Socket) = 2

☐ **Memory Hot Plug**

Memory (GiB)

2

Figure 64: VM Configuration

4. Select the **Compute DRS Mode** checkbox 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** dropdown menu.
 - » If you have not selected **Compute DRS Mode**, then select the host name of the VM from the **Host** dropdown menu.

5. Do one of the following:

- » Select **VM Templates** and then select a template from the **Template** dropdown menu.

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 **Content Library**, a content library in the **Content Library** dropdown menu, 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 Self-Service. 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** checkbox and a datastore cluster from the **Datastore Cluster** dropdown menu.

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** checkbox, and select a datastore from the **Datastore** dropdown menu.

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** and specify a folder name in the **Folder/Directory Name** field.

Note: Self-Service prefers 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, Self-Service considers the VM location you specified in the environment and creates the VM at the top-level folder.

Select the **Delete empty folder** checkbox 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** checkbox 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 **Cores Per Socket** count.
The number of sockets is calculated by dividing the total number of vCPUs by Cores Per Socket. For example, if the number of vCPUs is 4 and the Cores Per Socket is 2, then the number of sockets will be 2.
12. Select the **Memory Hot Plug** checkbox 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:
- Select the device type from the **Device Type** dropdown menu.
You can either select **CD-ROM** or **DISK**.
 - Select the adapter type from the **Adapter Type** dropdown menu.
You can select **IDE** for CD-ROM.
You can select **SCSI**, **IDE**, or **SATA** for DISK.
 - Enter the size of the disk in GiB.
 - In the **Location** field, select the disk location.
 - If you want to add a controller to the vDisk, select the type of controller in the **Controller** dropdown menu to attach to the disk.

Note: You can add either SCSI or SATA controllers. The available options depend on the adapter type.

- In the **Disk mode** dropdown menu, 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. If you want to customize the default OS properties of the VM, then click **Enable** under **VM Guest Customization** and select a customization from the **Predefined Guest Customization** dropdown menu.

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** dropdown menu.
 - Select **Hardware clock UTC** to enable hardware clock UTC.
 - Click the **+** icon to add network settings.
 - To automatically configure DHCP server, enable **Use DHCP** 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 133.
20. Click **Next**.
21. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 134. This step is optional.
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 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 [Creating a Project](#) on page 91 and [Configuring Environments in a Project](#) on page 110.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, 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.

VM Configuration Required for launching blueprints from marketplace

Quickly launch blueprints in this environment with this configuration for VMs

Not configured

Windows + Linux

▼ VM Configuration

@@(calm_array_index)@@ and @@(calm_time)@@ have been added to ensure a unique VM name. You can append or prepend a custom string to this value for identification.

Instance Name

vm-@@(calm_array_index)@@-@@(calm_time)@@

☒ Associate Public IP Address

Instance Type

a1.medium X ▼

Region

us-east-1 X ▼

Availability Zone

us-east-1a X ▼

Machine Image

Machine Image X ▼

IAM Role

aws-ec2-spot-fleet-role X ▼

Key Pairs

SID2 X ▼

VPC

VPC X ▼

☒ Include Classic Security Group

Security Groups

Security Groups X ▼

Subnet

Subnet X ▼

User Data

User Data X ▼

Figure 65: 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** checkbox to associate a public IP address with your AWS instance.

If you do not select the **Associate Public IP Address** checkbox, ensure that the AWS account and Self-Service are on the same network for the scripts to run.

6. Select an AWS instance type from the **Instance Type** dropdown menu.

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 dropdown menu displays the instances that are available in the AWS account. For more information, see *AWS documentation*.

7. Select the region from the **Region** dropdown menu and configure the following.

Note: The dropdown menu displays the regions which are selected while configuring the AWS setting.

- a. Select the availability zone from the **Availability Zone** dropdown menu.

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** dropdown menu.

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** dropdown menu.

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** dropdown menu.

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** dropdown menu.

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** checkbox to enable security group rules.
- Select security groups from the **Security Groups** dropdown menu.

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 66: 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** dropdown menu, 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. Select the **Delete on termination** checkbox to delete the storage when the instance is terminated. This step is optional.
- 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 133.
 12. Click **Next**.
 13. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 134. This step is optional.
 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 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.

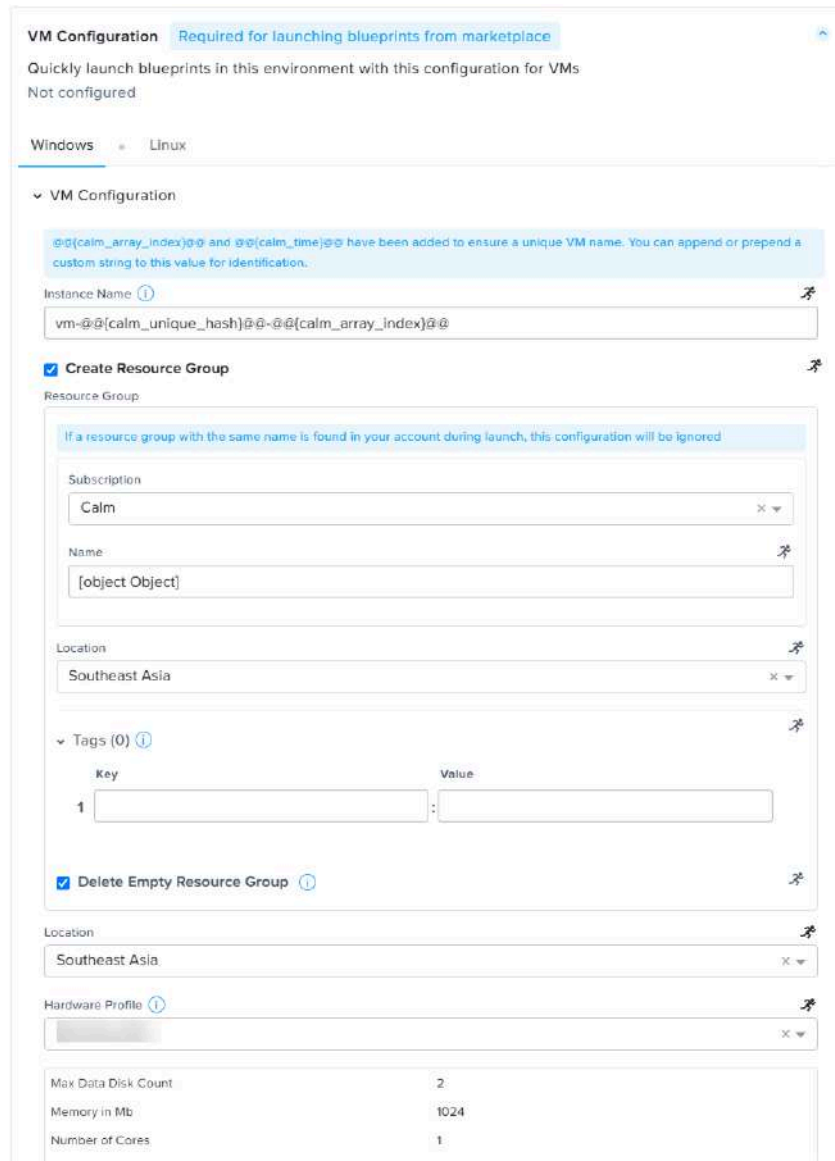
Before you begin

- 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 [Creating a Project](#) on page 91 and [Configuring Environments in a Project](#) on page 110.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, 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.



VM Configuration Required for launching blueprints from marketplace

Quickly launch blueprints in this environment with this configuration for VMs

Not configured

Windows * Linux

▼ VM Configuration

@{calm_array_index}@ and @{calm_time}@ have been added to ensure a unique VM name. You can append or prepend a custom string to this value for identification.

Instance Name ⓘ

vm-@{calm_unique_hash}@-@{calm_array_index}@

☒ Create Resource Group ⓘ

Resource Group

If a resource group with the same name is found in your account during launch, this configuration will be ignored

Subscription

Calm x ▼

Name ⓘ

[object Object]

Location ⓘ

Southeast Asia x ▼

▼ Tags (0) ⓘ

Key	Value
1	:

☒ Delete Empty Resource Group ⓘ

Location ⓘ

Southeast Asia x ▼

Hardware Profile ⓘ

Max Data Disk Count 2

Memory in Mb 1024

Number of Cores 1

Figure 67: 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** dropdown menu or select the **Create Resource Group** checkbox 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** dropdown menu displays the resource groups that are associated with the subscriptions you selected in your Azure account. In case you have not selected any subscriptions,

Self-Service considers all the subscriptions that are available in the Azure service principal to display the resource groups. Each resource group in the dropdown menu also displays the associated subscription.

6. If you selected a resource group from the **Resource Group** dropdown menu, then do the following:
 - a. Select the geographical location of the datacenter from the **Location** dropdown menu.
 - b. Select **Availability Sets** or **Availability Zones** from the **Availability Option** dropdown menu.

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** dropdown menu.

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** checkbox 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** dropdown menu.
 - 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** checkbox.
 - 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 `<UppercaseThumbprint>.prv` for private key. Both of these files are .pem formatted.

9. (For Windows) Select the **Provision Windows Guest Agent** checkbox.
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** checkbox.
11. Under the **Additional Unattended Content** section, click the **+** icon and do the following:
 - a. Select a setting from the **Setting Name** dropdown menu.
You can select **Auto Logon** or **First Logon Commands**.

Note: Guest customization is applicable only on images that allow 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** dropdown menu.
You can select **HTTP** or **HTTPS**.
 - b. If you selected **HTTPS**, then select the certificate URL from the **Certificate URL** dropdown menu.
13. Under the **Storage Profile** section, select the **Use Custom Image** checkbox 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** dropdown menu.
14. Under the **VM Image Details** section, select an image type in the **Source Image Type** dropdown menu.
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:

The screenshot shows the 'OS Disk Details' configuration section. It contains four rows, each with a label, a text input field, a help icon, and a clear button (X).

- Disk Name:** The input field contains the template `os-@@{calm_unique_hash}@@-@@{calm_array_index}@@-disk`.
- Storage Type:** The dropdown menu is set to 'Premium SSD'.
- Disk Caching Type:** The dropdown menu is set to 'Read Write'.
- Disk Create Options:** The dropdown menu is set to 'FromImage'.

Figure 68: OS Disk Details

- a. Select the storage type from the **Storage Type** dropdown menu.
You can select **Standard HDD**, **Standard SSD**, or **Premium SSD**.
 - b. Select a disk storage account from the **Disk Storage** dropdown menu.
This field is available only when the **Use Custom Image** is enabled.
 - c. Select disk caching type from the **Disk Caching Type** dropdown menu.
You can select **None**, **Read-only**, or **Read write**.
 - d. Select disk create option from the **Disk Create Option** dropdown menu.
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:
- a. Select a security group from the **Security Group** dropdown menu.
 - b. Select application security groups from the **ASGs** dropdown menu.
Unlike Network Security Groups that are defined at the network level, Application Security Groups (ASGs) are logical entities that are defined at the application level. ASGs help to manage the VM security by grouping the VMs according to the applications that run on them and allow application-centric use of Network Security Groups. You can select multiple ASGs in a NIC.
- Note:**

 - ASGs are supported when you clone your applications and are also supported in snapshot and restore.
 - ASGs are not supported in VM configuration updates.
- c. Select application security groups from the **ASGs** list.
Unlike Network Security Groups that are defined at the network level, Application Security Groups (ASGs) are logical entities that are defined at the application level. ASGs help to manage the VM

security by grouping the VMs according to the applications that run on them and allow application-centric use of Network Security Groups. You can select multiple ASGs in a NIC.

Note:

- ASGs are supported when you clone your applications and are also supported in snapshot and restore.
- ASGs are not supported in VM configuration updates.

- d. Select a virtual network from the **Virtual Network** dropdown menu.
- e. Under **Public IP Config**, enter a name and select an allocation method.
- f. 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. Enter tags in the **Tags** field. This step is optional.
18. Configure the connection in your environment. For more information, see [Configuring Check Log-in for your Environment](#) on page 133.
19. Click **Next**.
20. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 134. This step is optional.
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 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 [Creating a Project](#) on page 91 and [Configuring Environments in a Project](#) on page 110.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, 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.

VM Configuration

@@[calm_array_index]@@ and @@[calm_time]@@ have been added to ensure a unique VM name. You can append or prepend a custom string to this value for identification.

Instance Name

vm-@@[calm_array_index]@@-@@[calm_time]@@

Zone

us-central1-c

Machine Type

Figure 69: 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** dropdown menu.
Zone is a physical location where you can host the VM.
6. Select machine type from the **Machine Type** dropdown menu.
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** checkbox, and then select the persistent disk from the **Disk** dropdown menu.

Boot Disk

☐ Use existing disk

Storage Type

pd-balanced

Source Image

Size in GB

10

☒ Delete when instance is deleted

Figure 70: Disks

9. If you have not selected the **Use existing disk** checkbox, then do the following:
 - a. Select the type of storage from the **Storage Type** dropdown menu. 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** dropdown menu.

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** checkbox 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** checkbox and configure the following fields.
 - a. Select the network from the **Network** dropdown menu and the sub network from the **Subnetwork** dropdown menu.
 - b. Enter a name of the network in the **Access configuration Name** field and select the access configuration type from the **Access configuration type** dropdown menu.

These fields appear when you select the **Associate public IP Address** checkbox.
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** checkbox 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** dropdown menu.

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 133.
18. Click **Next**.
19. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 134. This step is optional.
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 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.

Before you begin

The network adapter must be specified in the VM configuration before you 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. Select the **Check log-in upon create** checkbox.
3. In the **Credential** list, select **Add New Credential** to add a new credential and do the following:
 - a. Type a name of the credential in the **Credential Name** field.
 - b. Type the 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, type 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.
If the private key is password protected, click **+Add Passphrase** to provide the passphrase. This step is optional.
 - e. If you want this credential as your default credential, select the **Use as default** checkbox.
 - f. Click **Done**.
4. Select address from the **Address** dropdown menu.
You can either select the public IP address or private IP address of a NIC.

5. Select the connection from the **Connection Type** dropdown menu.
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.
9. Click **Next** to add credentials to your environment or save the environment and the project.
For information on adding credentials, see [Adding Credentials to the Environment](#) on page 134.

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 Self-Service. Self-Service 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 [Creating a Project](#) on page 91 and [Configuring Environments in a Project](#) on page 110.

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** dropdown menu.
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.

Note:

Starting with version pc.2024.1, Prism Central discontinued supporting legacy quotas and enforcing any legacy quotas that you defined in your projects. To enforce quotas in projects, you must enable the policy engine. For more information, see [KB-15343](#).

When you upgrade your Prism Central instance to pc.2024.1 or later versions and enable the policy engine, your Prism Central instance migrates any legacy project quotas to policy-engine-based quotas. You cannot reverse the migration of legacy quotas.

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](#) on page 149.

- The resource quotas you allocate to the Nutanix and VMware provider accounts. For more information, see [Allocating Resource Quota to an Account](#) in the *Self-Service Administration and Operations Guide*.
- The resource quotas you allocate for a project at the project level. For more information, see [Managing Quota Limits for Projects](#) on page 136.
- The resource quotas you allocate to the Nutanix and VMware accounts within a project. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 100.

Project-Level Quota Allocation

You can define resource quota limits at the project level for the projects that you create in Prism Central. 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.

- 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.
- You can define resource quota limits for a project without defining quota limits for the associated Nutanix or VMware accounts.
- 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.
- 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 [Managing Quota Limits for Projects](#) on page 136.

For information on quota utilization report, see [Viewing Quota Utilization Report](#) in the *Self-Service Administration and Operations Guide*.

Managing Quota Limits for Projects

The **Policies** tab of a project provides you the options to define and manage quota limits for the project.

About this task

Quotas in projects work in the following ways:

- When Self-Service is disabled in your Prism Central instance, you can define and manage quotas only for Nutanix (AHV).
- When Self-Service is enabled in your Prism Central instance, you can define and manage quotas for Nutanix (AHV) and VMware accounts.

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 Self-Service-Enabled Prism Central](#) on page 100.
- View quota utilization report.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click **+Create Project** to create a new project and define quota limits. For information on creating a project, see [Creating a Project](#) on page 91.
 - » Click a project name in the list of existing projects to define quota limits for that project.
5. Configure your project with users, accounts, and environments. For more information, see [Adding Users to a Project](#) on page 95, [Infrastructure in Projects](#) on page 99, and [Environments in Projects](#) on page 110.
6. Click the **Policies** tab of the project.
7. Ensure that the **Quotas** tab is selected in the left pane.
8. To enable quota checks at the project level, enable the **Quotas** toggle button in the right pane.
9. 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.
10. To manage provider quota limits within the project, expand the provider account in the Provider Quotas section and do the following:

Note: The options to define and manage provider quota limits are available only when Self-Service is enabled in your Prism Central instance.

- 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.

11. To review quota utilization at the project level, click **Quota Utilization Report**.

Note: The Quota Utilization Report is available only when Self-Service is enabled in your Prism Central instance.

You can use the Quota Utilization Report to review the following details:

- Use the **Resource Utilization** tab to review the total utilization of vCPU, Memory, and Disk at the project level. You can also review the utilization of infrastructure resources by each cluster of the accounts associated with the project.
 - Use the **Quota Utilization** tab to review detailed utilization data at the level of each cluster of the associated accounts. You can review the quota allocated and quota used by each account. You can also review 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 Self-Service.
12. 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 or VMware platform. The policy determines the overall intent of the snapshot creation process. For applications VMs that run on Nutanix platform, you can create rules to manage your snapshots on a local cluster, on a remote cluster, or both and also configure snapshot expiration. For applications VMs that run on VMware platform, you can select allowed clusters and associated hosts for the local snapshots.

About this task

Perform the following steps to create your snapshot policy. For information on the snapshot configuration and creation, see [Snapshot and Restore Overview](#).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click **+Create Project** to create a new project and create snapshot policies within the project. For information on creating a project, see [Creating a Project](#) on page 91.
 - » Click a project name in the list of existing projects to create snapshot policies within the project.
5. Ensure that you have your project configured with users, accounts, and environments. For more information, see [Adding Users to a Project](#) on page 95, [Infrastructure in Projects](#) on page 99, and [Environments in Projects](#) on page 110.
6. On the **Policies** tab, click **Snapshot** in the left pane.

7. Click **+ Create Snapshot Policy**.

The screenshot shows the 'Create Snapshot Policy' form. The top section, 'General Information', contains a 'Policy Name' field with the value 'SnapshotPolicy1', a 'Policy Description (Optional)' text area with the text 'Policy for local snapshot.', and a checkbox labeled 'Set as default snapshot policy' which is currently unchecked. Below this is the 'Primary Site' section, which includes a 'Primary Environment' dropdown menu showing 'test_local_account_environment' and an 'Infrastructure' dropdown menu showing 'NTNX_LOCAL_AZ'. The bottom section, 'Local Snapshots', has an 'Enable' toggle switch that is turned on. A blue informational box states: 'All local clusters are added below by default. Enter the snapshot expiry duration associated with each rule.' Below this is a table for 'Snapshot Rules' with a '+ Add Rule' button. The table has two columns: 'Primary Cluster' and 'Snapshot Expiry (Days)'. One rule is listed with the cluster 'auto_cluster_prod_f3ae65b3a9d5' and an expiry of '15' days.

Primary Cluster	Snapshot Expiry (Days)
auto_cluster_prod_f3ae65b3a9d5	15

Figure 71: Create Snapshot Policy

8. In the **Policy Name** field, type a name for the snapshot policy.
9. In the **Policy Description** field, type a description for the snapshot policy.
10. Select the **Set as default snapshot policy** checkbox to make this your default snapshot policy.
11. Under **Primary Site**, select a primary environment.
You can select an environment you configured with your Nutanix or VMware account.

12. If you selected a Nutanix environment, then do the following to configure local snapshots in your policy.
- Select an account in the primary environment to which you want to associate the snapshot policy. Selecting an account in the **Infrastructure** dropdown menu enables **Local Snapshots**. You can view all the clusters that you allowed for the account in the **Primary Cluster** column under Snapshot Rules.
 - 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 **+ Add Rule** to include a deleted cluster to the policy.

13. If you selected a VMware environment, then do the following to configure local snapshots in your policy.

Note: The snapshot policy rules are created for VMware clusters and hosts. The Snapshot action is enabled only for the clusters or hosts that you select in the policy.

- Select an account in the primary environment to which you want to associate the snapshot policy. Selecting an account in the **Infrastructure** dropdown menu enables **Local Snapshots**. You can view all the clusters that you allowed for the account in the **Cluster** column under Snapshot Rules. The snapshot rule is applicable to the clusters that appear in the list.
- To remove an allowed cluster from the snapshot policy, click the Delete icon next to the cluster.
- To select a host for the snapshot policy, click **Add Rule**, and then select a host in the **Select Host** dropdown menu.
You can select multiple hosts for the snapshot policy.

14. If you selected a Nutanix environment, then do the following to enable remote snapshots in your policy.

The screenshot shows the 'Remote Snapshots' configuration window. At the top right, there is a toggle switch labeled 'Enable' which is currently turned on. Below this, a light blue banner contains the text: 'Remote snapshots can only be taken within the same Prism Central.' Under the 'Target Site' section, there are two dropdown menus. The first is labeled 'Target Environment' and has the value 'test_remote_account_environment'. The second is labeled 'Infrastructure' and has the value 'multiplc_account742b7335a375'. Below these is the 'Snapshot Rules' section, which includes a '+ Add Rule' button. A table below the button has four columns: 'Primary Clusters', 'Target Cluster', 'VM Categories', and 'Snapshot Expiry (Days)'. The first row of the table contains the following values: 'Select clusters', 'Select cluster', 'Select', and '15'. There is a red trash icon next to the '15'.

Figure 72: 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** dropdown menu, select the primary cluster where the snapshots of the VMs are taken.
 - d. From the **Target Cluster** dropdown menu, select the target cluster where you want to store the snapshots.
 - e. From the **VM Categories** dropdown menu, 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.
15. Click **Save Snapshot Policy**.

IDENTITY AND ACCESS MANAGEMENT

Identity and Access Management (IAM) is an authentication and authorization feature that uses attribute-based access control (ABAC). Starting from version pc.2024.1, you can find multiple improvements that focus on making the IAM feature better optimized and more autonomous, such as:

- Introduction of new system-defined roles to facilitate easier assignment of permissions (granular RBAC) and simplify the authorization process.
- Consolidation of role mapping and role assignment into the new Authorization Policies workflow.
- Simplified CAC authentication workflow.

Note: Prism Central welcome banner reverts to default after a series of disabling, upgrading and re-enabling IAM. For more information, see [KB-17644](#).

For information on IAM prerequisites, considerations, and enablement, see the [Identity and Access Management \(IAM\)](#) section in the *Security Guide*.

For IAM Software Support, see the [Prism Central Release Notes](#).

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 the *Security Guide*.

- Set the user authentication method to local, directory service, or both. For more information, see [Configuring Authentication](#).
- Create authorization policies to assign roles (built-in or custom) to identities (users or user groups) for a global or customized scope (of allowed actions). For more information, see [Authorization Policies](#).
- Add or edit local user accounts. For more information, see [Managing Local User Accounts](#).
 - Update your account. For more information, see [Updating My Account](#).
- Assign roles to users. For more information, see [Controlling User Access \(RBAC\)](#).

Note: For recovery plan RBAC, see [Nutanix Disaster Recovery Guide](#).

- Import SSL certificates. For more information, see [Importing an SSL Certificate](#).
- Control SSH access to Prism Central. For more information, see [Controlling Remote \(SSH\) Access](#).
- Create and apply security-related policies. For more information, see [Security Policies](#) in the *Flow Network Security Next-Gen Guide*.

User Management

Use the **Identities** tab on the Identify and Access Management page to view a list of local users, imported users, and user groups. The tab also provides you the option to add a local user. To access the **Identities** tab in Admin Center, click **IAM** in the navigation bar, and then click the **Identities** tab.

For information on creating and managing local user accounts, see the [Managing Local User Accounts](#) section in the *Security Guide*.

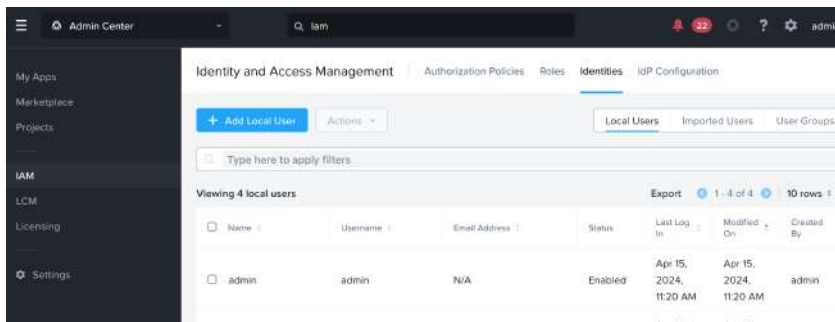


Figure 73: Identities Tab

The following table lists the actions that you can perform on the **Identities** tab and their associated options.

Table 22: Identities Tab - Options

Action	Option
Add a local user	To add a local user. For more information, see Managing Local User Accounts .
Edit a local user	To edit or disable a local user. For more information, see Editing or Disabling a Local User Account .
Order the user list alphabetically (A-Z or Z-A)	Click (toggle) the Name , Username , Type , Email Address , Last Log In , or Modified On column header.
View imported users	Click the Imported Users tab.
View Users Groups	Click the User Groups tab.
Filter by name or username	Type the name or username and press Enter. For User Groups, you can only filter by name.

Roles Management

The **Roles** tab on the Identify and Access Management page allows you to view summary information about built-in (default) and custom (created) roles and to access detailed information about each role. To access the list of roles in Admin Center, click **IAM** in the navigation bar, and then click the **Roles** tab.

For information on how to create, manage, and apply roles, see [Controlling User Access \(RBAC\)](#) in the *Security Guide*.

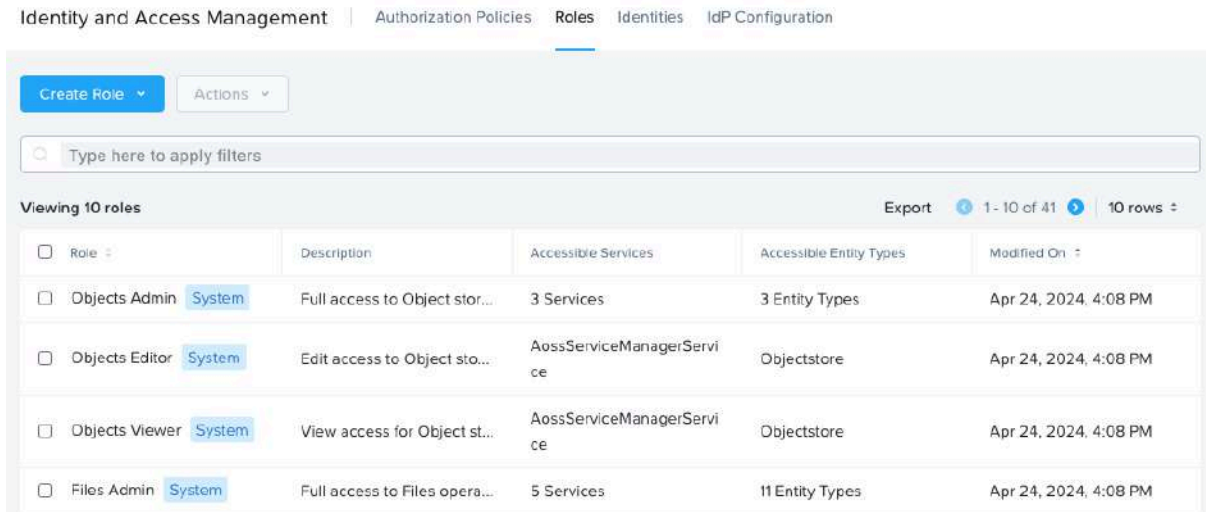


Figure 74: Roles Tab

By default, you can view the following information about the roles on the **Roles** tab.

Table 23: Roles List Fields

Parameter	Description
Role	Displays the role name.
Description	Describes the role (if a description was provided when the role was created or updated).
Accessible Services	The services that are accessible to the role. In case multiple services are accessible to the role, you can hover the mouse to view the list of services the role has access to.
Accessible Entity Types	The entity types that are accessible to the role. In case multiple entity types are accessible to the role, you can hover the mouse to view the list of entity types the role has access to.
Modified On	Displays the date and time when the role was modified the last time.

You can filter the list based on the role name.

You can use the **Create Role** button to create a custom role. For more information, see [Creating a Custom Role](#) in the *Security Guide*.

The **Actions** menu appears when you select a role. You can use the Actions menu to one or more roles are selected. It allows you to duplicate roles, to add authorization policy, to update a custom role, and to delete a custom role.

Note: You cannot delete or update the System roles.

To access details of a role, click the role name on the **Roles** tab. The Role page has the following action buttons.

- **Add Authorization Policy:** Use this button to add an authorization policy to the role. For information on authorization policies, see [Authorization Policies](#) in the *Security Guide*.

- **Duplicate:** Use this button to duplicate the role. It opens the role configuration screen preconfigured with the same permissions as this role. For more information, see [Creating a Custom Role](#) in the *Security Guide*.
- **Delete:** Use the button to delete the role. You can only delete a custom role.
- **Update:** Use this button to update the role permissions. You can only update a custom role. For more information, see [Modifying a Custom Role](#) in the *Security Guide*.

The role page also displays the entity types accessible to the role and their associated operations. You can expand the entity types to view more details about the operations.

LCM VIEW

The life cycle manager (LCM) tracks software and firmware versions of the various components in a cluster. It allows you to view information about the current inventory and update the versions as needed.

To access the LCM dashboard, select **Admin Center** from the Application Switcher and then click **LCM** in the navigation bar. You must have the Prism Admin role to access LCM.

For more information on life cycle manager, see the [Life Cycle Manager Guide](#).

ADMIN CENTER SETTINGS OPTIONS

You can use the Admin Center settings to configure various application management, networking, alerts, and appearance related settings. These administrative settings are common across multiple Nutanix applications that you deploy and access using Prism Central or affect the platform. You must log in as an administrator to configure these settings.

You have the following categories of settings in Admin Center.

- **App Management**

Nutanix Central is a unified cloud management console where you can access multiple Prism Centrals. Registration is a two-way API-key process. Nutanix Central uses role-based access control (RBAC) to manage access of the registered Prism Central domains.

For information on Nutanix Central and its registration process, see [Nutanix Central Admin Guide](#).

Prism Central automatically archives run logs of the deleted applications and custom actions that are older than three months to clean up resources. These archives can be downloaded and viewed by an administrator to get valuable insights within 7 days of archive creation. For more information, see [Application Log Archive](#) on page 148.

- **Governance Policies**

The policy engine enforces resource quota policy, orchestrates applications through tunnels on a virtual private network (VPC) on Nutanix accounts, enforces approval policies, and schedules app actions and runbook executions. For more information on policy engine, see [Policy Engine Overview](#) on page 149.

- **Setup**

Nutanix Cloud Manager Intelligent Operations (formerly Prism Pro/Ultimate) combines aspects of administration, reporting, and intelligent automation for IT Operations. Intelligent Operations empowers you to plan intelligently and optimize for capacity, proactively detect performance anomalies, and automate the IT operations. For information on Intelligent Operations enablement, see [Enabling Intelligent Operations](#) in the *Intelligent Operations Guide*.

The Pulse feature provides diagnostic system data about Prism Central to Nutanix to deliver proactive, context-aware support for Nutanix solutions. When enabled, Pulse sends a summary email of cluster configuration to a Nutanix server daily. With this information, Nutanix can apply advanced analytics to optimize your implementation and to address potential problems.

- **Network**

You can configure the following network settings:

- **HTTP Proxy:** A server to communicate with the Nutanix Service Center if Prism Central cannot send traffic to a Nutanix Service Center directly.
- **Name Server:** A server that hosts a network service to provide responses to queries against a directory service, such as a DNS server.
- **NTP Server:** A server to which Prism Central must connect to synchronize the system clock.
- **SNMP Server:** A server to facilitate the exchange of management information between network devices.

- **Security**

The Cluster Lockdown feature 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 the *Security Guide*. You can use the

SSL Certificate option to create a self-signed certificate. For more information, see [Installing an SSL Certificate](#) in the *Security Guide*.

- **Alerts and Notifications**

You can configure the alert messages that are sent from Prism Central, define rules, and the email content. As part of alert and notification settings, you can also configure the following servers:

- SMTP server: To allow Prism Central to send notifications.
- Syslog server: To allow 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.

- **Appearance**

You can configure settings related to accessibility and appearance, such as:

- Default landing page: The page that appears when you or other users first log in to Prism Central.
- Language: The language and locale for your Prism Central instance.
- UI Settings: The theme and timeout duration for admin and non-admin users.
- Banner page: The Welcome page appearance and content.

Application Log Archive

Prism Central automatically archives run logs of the deleted applications and custom actions that are older than three months. You can download the archives within 7 days from the time of archive creation.

For a running application, data is not archived for the system-generated Create actions.

You can get the following information for Start, Restart, Stop, Delete, and Soft Delete system-generated actions and user-created actions.

- Started by
- Run by
- Status

Prism Central archives all action details of a deleted application.

Only an administrator can view and download the application log archive. For more information, see [Downloading Application Log Archive](#) on page 148.

Downloading Application Log Archive

Prism Central periodically archives application logs to clear resources. You can download the archived application logs from the Admin Center settings.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under App Management, click **App Log Archive**.

The right pane displays the timestamp and the validity of the archive.

5. Click **Download** and navigate to the location on your computer where you want to save the tar.gz file.

Policy Engine Overview

The policy engine is a single-VM setup for the single or scale-out Prism Central. When you enable the policy engine for your Prism Central instance, a new VM is created and deployed for the policy engine. All you need is an available IP address that belongs to the same network as that of your Prism Central VM for the policy engine VM.

As an administrator, you can enable the policy engine to:

- Enforce resource quota policy for the infrastructure resources (compute, memory, and storage) on Nutanix and VMware. The quota policy enforcement allows better governance on resources across infrastructures at the provider and project levels. For more information, see [Quota Policy Overview](#) on page 135.
- Orchestrate apps through tunnels on a virtual private network (VPC) on Nutanix accounts. For more information, see [Tunnels for Orchestration within a VPC](#) in the *Self-Service Administration and Operations Guide*.
- Enforce approval policies to manage resources and control actions in your environment. For more information, see [Approval Policy Overview](#) in the *Self-Service Administration and Operations Guide*.
- Schedule app actions and runbook executions. For more information, see [Scheduler Overview](#) in the *Self-Service Administration and Operations Guide*.

Note:

- In versions earlier than pc.2024.1, the policy engine enablement was possible only when Self-Service was deployed and licensed in your Prism Central instance. Starting from Prism Central version pc.2024.1, you can enable the policy engine from the Admin Center irrespective of whether Self-Service is deployed or not.
- Nutanix recommends that you enable policy engine immediately after you upgrade your Prism Central to version pc.2024.1. Until the policy engine is enabled, Prism Central cannot perform any quota checks expected from the existing project quotas (legacy project quotas) or any validation on the resource consumption. The Quota section in your Projects pages also remains unavailable until you enable the policy engine.
- When you upgrade your Prism Central to version pc.2024.1 and enable the policy engine, Prism Central migrates any existing project quotas (legacy project quotas) to the new version of quotas designed as per the policy engine. You cannot reverse the migration of legacy quotas. For information on policy engine quotas, see [Quota Policy Overview](#) on page 135.

Enabling Policy Engine

The policy engine is a single-VM setup for the single or scale-out Prism Central. You can enable policy engine from the Admin Center Settings of your Prism Central instance.

Before you begin

- Ensure that you have an available IP address that belongs to the same network as that of your Prism Central VM. Prism Central uses the IP address to create and deploy the policy engine VM. If an HTTP proxy is configured, ensure that you add the policy engine IP address to the HTTP-proxy allowlist.
- Ensure that you enable Marketplace before enabling policy engine. For information on enabling marketplace, see [Enabling Marketplace](#) on page 12.

About this task

- When you enable policy engine for your Prism Central instance, a new VM with 4 vCPUs, 6 GiB memory, and 40 GiB disk space is created and deployed for the policy engine.
- When you upgrade Prism Central to the latest version and if the policy engine was enabled in the previous version, you must also upgrade policy engine to the latest version using LCM.

Note:

- In versions earlier than pc.2024.1, the policy engine enablement was possible only when Self-Service was deployed and licensed in your Prism Central instance. Starting from Prism Central version pc.2024.1, you can enable the policy engine from the Admin Center irrespective of whether Self-Service is deployed or not.
- When you upgrade your Prism Central to version pc.2024.1 and enable the policy engine, Prism Central migrates any existing project quotas (legacy project quotas) to the new version of quotas designed as per the policy engine. You cannot reverse the migration of legacy quotas. For information on policy engine quotas, see [Quota Policy Overview](#) on page 135.

To enable policy engine, follow these steps:

Procedure

1. Log in to your Prism Central instance as an administrator.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Settings**.
4. In the **Governance Policies** section, click **Policy Engine**.
5. If your Prism Central instance is on AHV, follow these steps:
 - a. In the **IP Address for Policy Engine** field, enter an IP address for the policy engine VM.

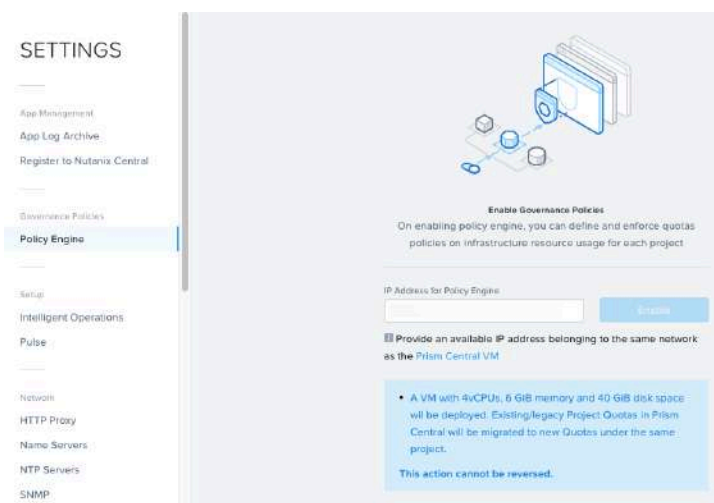


Figure 75: Policy Engine Page

- b. Click **Enable**.
- c. Click **Confirm** to enable the policy engine.

6. If your Prism Central instance is on ESXi and Self-Service is disabled, follow these steps to enter the details of your vCenter server where the Prism Central VM resides.

Prism Central requires the vCenter server details to establish communication with the server prior to creating the policy engine VM. Once you enter the server details, Prism Central validates the credentials and enables the **IP Address for Policy Engine** field so that you can enter the IP address of the policy engine VM.

Enable Governance Policies

On enabling policy engine, you can define and enforce quotas policies on infrastructure resource usage for each project.

• A VM with 4vCPU, 6 GiB memory and 40 GiB disk space will be deployed. Existing/legacy Project Quotas in Prism Central will be migrated to new Quotas under the same project.
This action cannot be reversed.

Please provide vCenter details where the Prism Central VM resides:

Server
Username
Password
Port
IP Address for Policy Engine

Provide an available IP address belonging to the same network as the Prism Central VM.

Enable

Figure 76: Policy Engine Page

- a. In the **Server** field, enter the IP address of the vCenter server.
 - b. In the **Username** field, enter the user name of the vCenter account.
If a domain is part of the username, use the <username>@<domain> username syntax.
 - c. In the **Password** field, enter the password of the account.
 - d. In the **Port** field, enter the port number as 443.
 - e. In the **IP Address for Policy Engine** field, enter an IP address for the policy engine VM.
 - f. Click **Enable**.
 - g. Click **Confirm** to enable the policy engine.
7. If your Prism Central instance is on ESXi and Self-Service is enabled, follow these steps:
 - a. Ensure that you added the VMware provider account for the vCenter that manages the host where the Prism Central VM resides to your Self-Service.
For more information, see [Configuring a VMware Account](#) in the *Self-Service Administration and Operations Guide*.
 - b. In the **IP Address for Policy Engine** field, enter an IP address for the policy engine VM.
 - c. Click **Enable**.
 - d. Click **Confirm** to enable the policy engine.

What to do next

- Allocate resource quotas to provider accounts. For more information, see [Allocating Resource Quotas to an Account](#) in the *Self-Service Administration and Operations Guide*.
- Allocate resource quotas to projects. For more information, see [Managing Quota Limits for Projects](#).
- Create VPC tunnels on Nutanix accounts. For more information, see [Creating VPC Tunnels](#) in the *Self-Service Administration and Operations Guide*.
- Create approval policies for runbook executions, app launch, or app day-2 operations. For more information, see [Creating an Approval Policy](#) in the *Self-Service Administration and Operations Guide*.
- Schedule app actions and runbook executions. For more information, see [Creating a Scheduler Job](#) in the *Self-Service Administration and Operations Guide*.

Enabling Policy Engine at a Dark Site

The policy engine is a single-VM setup for the single or scale-out Prism Central instance. You can enable the policy engine at a dark site.

Before you begin

If your Prism Central instance is on ESXi, add the VMware provider account for the vCenter that manages the host where the Prism Central VM resides to your Prism Central instance.

Procedure

1. Download the policy engine image from the [Support & Insights Portal](#).
When you download the image, ensure that the image version is compatible with your Prism Central version.
2. If your Prism Central is on AHV, upload the image on your Prism Central instance with the following name:
`<Self-Service version number>-CalmPolicyVM.qcow2`
3. If your Prism Central is on ESXi, follow these steps to upload and convert the image to a VM template:
 - a. Manually upload the image as a VM template on the vCenter host where the Prism Central VM resides with the following name:
`<Self-Service version number>-CalmPolicyVM.ova`
You must upload the OVA file as an OVF template in the same container where the Prism Central VM is located.
 - b. In the vCenter inventory, locate the generated VM to convert the VM into a template.
 - c. Right-click the VM (`<Self-Service version number>-CalmPolicyVM.ova`) and click **Template > Convert to Template**.
 - d. Verify that the VM template name is `<Self-Service version number>-CalmPolicyVM.ova`.
You must edit the VM template name in case the name is other than `<Self-Service version number>-CalmPolicyVM.ova`.
4. After uploading the image, enable the policy engine from the Admin Center Settings page. For information on enabling the policy engine, see [Enabling Policy Engine](#) on page 149.

Viewing Policy Engine VM Details

After enabling policy engine, you can review the policy engine VM configuration, network configuration, and cluster information on the Admin Center Settings page. For example, you can view the power status, protection status, or cluster name of the policy engine VM.

Before you begin

Ensure that you have enabled the policy engine for your Prism Central instance. For information on enabling the policy engine, see [Enabling Policy Engine](#) on page 149.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Governance Policies in the left pane, click **Policy Engine**.
5. Expand the **Policy Engine VM Details** section to view the policy engine VM details.

Disabling Policy Checks

Disable the policy checks for your Prism Central instance if the policy engine VM encounters any connectivity issues or the policy engine VM does not respond.

About this task

Note: Disabling policy checks disables approval policies, suspends quota policy enforcement, and pauses scheduled jobs.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Governance Policies in the left pane, click **Policy Engine**.
5. Select the **Skip Policy Checks** checkbox to disable the policy enforcement.
6. Click **Confirm**.

Configuring Rsyslog for Policy VM

Use the following procedure to configure rsyslog so that the Policy VM can send its logs to the rsyslog server.

About this task

Policy VM hardening happens during the last step of its image creation. Hardening is accomplished using the following files that the operating system vendor provides.

- centos-fedramp-os-hardening.yml
- centos-stig-mac1.yml
- centos-stig-mac2.yml

The following procedure provides you the steps to configure rsyslog for Policy VM. For the purpose of this procedure, 10.44.19.115 is considered as the IP address for the Policy VM.

Note: In case you are unable to connect to the server, ensure that there is no firewall rule blocking on port 514. You can check the connectivity using the following command.

```
[root@ntnx-10-44-19-115-calm-policy-vm ~]# nc -zv 10.46.138.90 514
Ncat: Version 7.50 ( https://nmap.org/ncat )
Ncat: Connected to 10.46.138.90:514.
Ncat: 0 bytes sent, 0 bytes received in 0.02 seconds.
[root@ntnx-10-44-19-115-calm-policy-vm ~]# nc -zvu 10.46.138.90 514
Ncat: Version 7.50 ( https://nmap.org/ncat )
Ncat: Connected to 10.46.138.90:514.
Ncat: UDP packet sent successfully
Ncat: 1 bytes sent, 0 bytes received in 2.01 seconds.
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
```

Before you begin

Ensure that the rsyslog server is running in your environment. For information on verifying the version or setting up the rsyslog server, see [Configuring an Rsyslog Server on CentOS Linux 7](#) on page 157.

About this task

Procedure

Perform the following steps to direct the Policy VM logs to the rsyslog server.

1. Log in to the Policy VM through Prism Central.

```
my.laptop ~ %ssh nutanix@10.44.19.115
Warning: Permanently added '10.44.19.115' (ED25519) to the list of known hosts.
nutanix@10.44.19.115's password:
Last login: Tue May 30 08:10:35 2023 from 10.138.224.88
Policy Engine Machine (CentOS Linux release 7.9.2009 (Core) 20220727)
[nutanix@ntnx-10-44-19-115-calm-policy-vm ~]$
[nutanix@ntnx-10-44-19-115-calm-policy-vm ~]$ sudo su -
Last login: Tue May 30 08:30:20 EDT 2023 on pts/0
Last failed login: Tue May 30 11:09:34 EDT 2023 from 10.46.8.7 on ssh:notty
There were 7 failed login attempts since the last successful login.
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
```

The logs are present at the following location. The configuration is done so that these log files can be sent to the rsyslog server.

```
[root@ntnx-10-44-19-115-calm-policy-vm ~]# ls -l /home/nutanix/data/log/*.log
-rw-r--r--. 1 333 333 649 May 26 01:55 /home/nutanix/data/log/agni.log
-rw-r--r--. 1 333 333 125994 May 30 11:25 /home/nutanix/data/log/
approval_v1_u8R9a_0.log
-rw-r--r--. 1 333 333 127581 May 30 11:25 /home/nutanix/data/log/
approval_v1_u8R9a_1.log
-rw-r--r--. 1 root root 233043 May 30 11:27 /home/nutanix/data/log/durga_0.log
-rw-r--r--. 1 root root 233043 May 30 11:27 /home/nutanix/data/log/durga_1.log
-rw-r--r--. 1 root root 2960 May 26 01:56 /home/nutanix/data/log/
elastic_search.log
-rw-r--r--. 1 root root 7450 May 26 01:55 /home/nutanix/data/log/
email_v1_3agzt_0.log
-rw-r--r--. 1 root root 7778 May 26 01:55 /home/nutanix/data/log/
email_v1_3agzt_1.log
-rw-r--r--. 1 333 333 9228768 May 30 11:27 /home/nutanix/data/log/ganga.log
-rw-r--r--. 1 root root 5420761 May 30 11:27 /home/nutanix/data/log/jove.log
```

```

-rw-r--r--. 1 333 333 1006 May 26 01:55 /home/nutanix/data/log/kali.log
-rw-r--r--. 1 333 333 1588866 May 30 11:26 /home/nutanix/data/log/leo.log
-rw-r--r--. 1 root root 232366 May 30 11:27 /home/nutanix/data/log/narad.log
-rw-r--r--. 1 333 333 94077 May 26 02:46 /home/nutanix/data/log/plum.log
-rw-r--r--. 1 333 333 4905 May 26 01:55 /home/nutanix/data/log/
policy_superevents.log
-rw-r--r--. 1 root root 1642881 May 30 11:25 /home/nutanix/data/log/quantum.log
-rw-r--r--. 1 root root 371 May 26 01:55 /home/nutanix/data/log/quota-
enforce_v1_ychqz_0.log
-rw-r--r--. 1 root root 586 May 26 01:55 /home/nutanix/data/log/quota-
enforce_v1_ychqz_1.log
-rw-r--r--. 1 root root 2132 May 26 01:55 /home/nutanix/data/log/redis.log
-rw-r--r--. 1 root root 5447 May 26 01:57 /home/nutanix/data/log/superevents.log
-rw-r--r--. 1 333 333 439 May 26 01:55 /home/nutanix/data/log/terraform-
plum_v1_EnkaH_0.log
-rw-r--r--. 1 333 333 1246 May 26 01:55 /home/nutanix/data/log/terraform-
plum_v1_EnkaH_1.log
-rw-r--r--. 1 root root 283753 May 30 11:27 /home/nutanix/data/log/vajra_0.log
-rw-r--r--. 1 root root 5189 May 26 01:57 /home/nutanix/data/log/zaffi.log
[root@ntnx-10-44-19-115-calm-policy-vm ~]#

```

2. Change SELinux from enforcing mode to permissive mode to ensure that the rsyslog functionality is not blocked.

```

root@ntnx-10-44-19-115-calm-policy-vm ~]# setenforce 0
[root@ntnx-10-44-19-115-calm-policy-vm ~]# sestatus
SELinux status:                enabled
SELinuxfs mount:              /sys/fs/selinux
SELinux root directory:       /etc/selinux
Loaded policy name:            targeted
Current mode:                  permissive
Mode from config file:         enforcing
Policy MLS status:             enabled
Policy deny_unknown status:    allowed
Max kernel policy version:     31

```

3. Uncomment the UDP protocol and add the rsyslog server IP address to the /etc/rsyslog.conf file.

```

[root@ntnx-10-44-19-115-calm-policy-vm ~]# cat /etc/rsyslog.conf
# rsyslog configuration file

# For more information see /usr/share/doc/rsyslog-*/rsyslog_conf.html
# If you experience problems, see http://www.rsyslog.com/doc/troubleshoot.html

#### MODULES ####

# The imjournal module below is now used as a message source instead of imuxsock.
$ModLoad imuxsock # provides support for local system logging (e.g. via logger
command)
$ModLoad imjournal # provides access to the systemd journal
#$ModLoad imklog # reads kernel messages (the same are read from journald)
#$ModLoad immark # provides --MARK-- message capability

# Provides UDP syslog reception
$ModLoad imudp
$UDPServerRun 514

# Provides TCP syslog reception
#$ModLoad imtcp
#$InputTCPServerRun 514

```

```
<- - - NOT SHOWING UNCHANGED DATA FOR BREVITY - - - >

# remote host is: name/ip:port, e.g. 192.168.0.1:514, port optional
#*. * @@remote-host:514
*. * @10.46.138.90:514
# ### end of the forwarding rule ###
$FileCreateMode 0640
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
```

4. Create the 50-default.conf file with the following content so that it is alphabetically ahead of the names of the log files and is parsed first.

```
[root@ntnx-10-44-19-115-calm-policy-vm ~]# cat /etc/rsyslog.d/50-default.conf
# Nutanix remote server rules
$ModLoad imfile
# Nutanix remote server rules end
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
```

Below are a few configuration examples for the log files. Similarly, you can add the remaining configurations of the log files that you want to send to the rsyslog server.

```
[root@ntnx-10-44-19-115-calm-policy-vm ~]# cat /etc/rsyslog.d/durga.conf
# Nutanix remote server rules
#$ModLoad imfile
$ActionForwardDefaultTemplate RSYSLOG_ForwardFormat
$InputFileName /home/nutanix/data/log/durga_*.log
$InputFileTag all-durga:
$InputFileFacility local0
$InputFileSeverity info
$InputRunFileMonitor
# Nutanix remote server rules end
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
[root@ntnx-10-44-19-115-calm-policy-vm ~]# cat /etc/rsyslog.d/jove.conf
# Nutanix remote server rules
#$ModLoad imfile
$ActionForwardDefaultTemplate RSYSLOG_ForwardFormat
$InputFileName /home/nutanix/data/log/jove.log
$InputFileTag jove:
$InputFileFacility local0
$InputFileSeverity info
$InputRunFileMonitor
# Nutanix remote server rules end
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
[root@ntnx-10-44-19-115-calm-policy-vm ~]# cat /etc/rsyslog.d/kali.conf
# Nutanix remote server rules
#$ModLoad imfile
$ActionForwardDefaultTemplate RSYSLOG_ForwardFormat
$InputFileName /home/nutanix/data/log/kali.log
$InputFileTag kali:
$InputFileFacility local0
$InputFileSeverity info
$InputRunFileMonitor
# Nutanix remote server rules end
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
[root@ntnx-10-44-19-115-calm-policy-vm ~]# cat /etc/rsyslog.d/narad.conf
# Nutanix remote server rules
#$ModLoad imfile
$ActionForwardDefaultTemplate RSYSLOG_ForwardFormat
$InputFileName /home/nutanix/data/log/narad.log
$InputFileTag narad:
$InputFileFacility local0
$InputFileSeverity info
$InputRunFileMonitor
```

```
# Nutanix remote server rules end
[root@ntnx-10-44-19-115-calm-policy-vm ~]#
```

5. Restart services for the changes to take effect.

```
[root@rsyslog-server ~]$systemctl restart systemd-journald
[root@rsyslog-server ~]$systemctl restart rsyslog
[root@rsyslog-server ~]$systemctl status rsyslog
```

6. Login to the rsyslog server and check in the custom configured directory to verify that the server is receiving the log files.

```
[root@rsyslog-server ~]$cd /home/nutanix/rsyslog-server/
[root@rsyslog-server rsyslog-server]$ls -l
total 0
drwx-----. 3 root root 22 May 30 15:32 10.44.19.115_10.44.19.115
drwx-----. 3 root root 22 May 30 15:16 localhost_10.46.138.90
drwx-----. 3 root root 22 May 30 15:12 localhost_127.0.0.1
drwx-----. 3 root root 22 May 30 15:32 ntnx-10-44-19-115-calm-policy-vm_10.44.19.115
[root@rsyslog-server ~]$
[root@rsyslog-server rsyslog-server]$
[root@rsyslog-server rsyslog-server]$ls -l ntnx-10-44-19-115-calm-policy-vm_10.44.19.115/20230530/
total 56
-rw-----. 1 root root 16420 May 30 15:34 all-durga.log
-rw-----. 1 root root 11074 May 30 15:34 jove.log
-rw-----. 1 root root 4096 May 30 15:34 narad.log
-rw-----. 1 root root 541 May 30 15:32 polkitd.log
-rw-----. 1 root root 358 May 30 15:32 rsyslogd.log
-rw-----. 1 root root 554 May 30 15:33 sshd.log
-rw-----. 1 root root 366 May 30 15:32 systemd.log
-rw-----. 1 root root 108 May 30 15:33 unix_chkpwd.log
[root@rsyslog-server rsyslog-server]$
```

Configuring an Rsyslog Server on CentOS Linux 7

Use the following procedure to configure the rsyslog server on your CentOS Linux 7 in case the server is not running in your environment.

About this task

CentOS Linux 7 typically comes with the rsyslog server installed. In case the rsyslog server is not installed or up-to-date, you can use YUM to install or update the server.

Procedure

1. Use the following command to verify the server version.

```
[root@rsyslog-server rsyslog-server]$rsyslogd -v
rsyslogd 8.24.0-52.el7_8.2, compiled with:
  PLATFORM:      x86_64-redhat-linux-gnu
  PLATFORM (lsb_release -d):
  FEATURE_REGEX: Yes
  GSSAPI Kerberos 5 support: Yes
  FEATURE_DEBUG (debug build, slow code): No
  32bit Atomic operations supported: Yes
  64bit Atomic operations supported: Yes
  memory allocator: system default
  Runtime Instrumentation (slow code): No
  uuid support:   Yes
  Number of Bits in RainerScript integers: 64
```

```
See http://www.rsyslog.com for more information.
[root@rsyslog-server rsyslog-server]$
```

2. Log in to the rsyslog server from your computer.

```
my.laptop ~ %ssh centos@10.46.138.90 -i ~/Downloads/my-secret-key.key
Warning: Permanently added '10.46.138.90' (ED25519) to the list of known hosts.
Last login: Tue May 30 13:15:40 2023 from 10.138.224.88
-bash: warning: setlocale: LC_CTYPE: cannot change locale (UTF-8): No such file or
directory
[centos@localhost ~]$
```

3. Change the prompt so that it is easier to recognise rsyslog server.

```
[centos@rsyslog-server ~]$sudo su -
Last login: Mon Feb 20 07:59:30 UTC 2023 on pts/0
[root@localhost ~]# OLDPS1=$PS1
[root@localhost ~]# PS1="[\u@rsyslog-server \W]\$"
[root@rsyslog-server ~]$
```

4. Create the directory where the log files must be collected on the server.

```
[root@rsyslog-server ~]$mkdir -p /home/nutanix/rsyslog-server
[root@rsyslog-server ~]$chown root:adm /home/nutanix/rsyslog-server
```

In the above command, a custom location is chosen instead of /var/log and the ownership of the directory is changed to root:adm

Ensure that this mount has sufficient storage space.

5. Change SELinux from enforcing mode to permissive mode to ensure that the rsyslog functionality is not blocked.

```
root@rsyslog-server ~]$setenforce 0
[root@rsyslog-server ~]$sestatus
SELinux status:                enabled
SELinuxfs mount:                /sys/fs/selinux
SELinux root directory:         /etc/selinux
Loaded policy name:              targeted
Current mode:                   permissive
Mode from config file:          enforcing
Policy MLS status:              enabled
Policy deny_unknown status:     allowed
Max kernel policy version:      31
```

This change reverts on reboot. For a permanent change, you need to add this change in a relevant config file.

6. Make the following changes in the /etc/rsyslog.conf file.

Uncomment the following:

```
$ModLoad imudp
$UDPServerRun 514
```

```
$ModLoad imtcp
$InputTCPServerRun 514
```

7. Add the server IP address at the end of the file. In this example, the server IP address is 10.46.138.90.

```
*.* @10.46.138.90:514
*.* @@10.46.138.90:514
```

The first line with a single @ indicates that UDP protocol is accepted and the second line indicates that TCP protocol is accepted.

8. Create a 50-remote-logs.conf file that points to the custom directory you configured.

The template directive in the file suggests you to create directories as per the IP address and hostname of the client machine. To make the segregation easier, each directory contains a folder for a particular date with the actual log files for the day in the folder. In addition, you can set up a Cron job to delete folders that are older than a certain duration and save storage space.

```
[root@rsyslog-server ~]$cat /etc/rsyslog.d/50-remote-logs.conf
# define template for remote login
# remote logs will be stored at /var/log/remotelogs directory
# each host will have specific directory based on the system %HOSTNAME%
# name of the log file is %PROGRAMNAME%.log such as sshd.log, su.log
# both %HOSTNAME% and %PROGRAMNAME% is the Rsyslog message properties
template (
    name="RemoteLogs"
    type="string"
    string="/home/nutanix/rsyslog-server/%hostname%_%fromhost-ip%/%$year%%$month%%$day%/%programname%.log"
)

# gather all log messages from all facilities
# at all severity levels to the RemoteLogs template
*. * -?RemoteLogs

# stop the process once the file is written
& stop
[root@rsyslog-server ~]$
```

9. Run a syntax check to ensure that there are no major errors in the above files.

```
[root@rsyslog-server ~]$rsyslogd -N1 -f /etc/rsyslog.conf
rsyslogd: version 8.24.0-52.el7_8.2, config validation run (level 1), master
config /etc/rsyslog.conf
rsyslogd: error during config processing: STOP is followed by unreachable
statements! [v8.24.0-52.el7_8.2 try http://www.rsyslog.com/e/2207 ]
[root@rsyslog-server ~]$
[root@rsyslog-server ~]$rsyslogd -N1 -f /etc/rsyslog.d/50-remote-logs.conf
rsyslogd: version 8.24.0-52.el7_8.2, config validation run (level 1), master
config /etc/rsyslog.d/50-remote-logs.conf
rsyslogd: End of config validation run. Bye.
[root@rsyslog-server ~]$
[root@rsyslog-server ~]$rsyslogd -N1
rsyslogd: version 8.24.0-52.el7_8.2, config validation run (level 1), master
config /etc/rsyslog.conf
rsyslogd: error during config processing: STOP is followed by unreachable
statements! [v8.24.0-52.el7_8.2 try http://www.rsyslog.com/e/2207 ]
[root@rsyslog-server ~]$
```

10. Restart the rsyslog process for these files to take effect.

```
[root@rsyslog-server ~]$systemctl restart rsyslog
[root@rsyslog-server ~]$systemctl status rsyslog
# rsyslog.service - System Logging Service
   Loaded: loaded (/usr/lib/systemd/system/rsyslog.service; enabled; vendor preset:
   enabled)
   Active: active (running) since Tue 2023-05-30 15:12:31 UTC; 6s ago
     Docs: man:rsyslogd(8)
           http://www.rsyslog.com/doc/
   Main PID: 208941 (rsyslogd)
    CGroup: /system.slice/rsyslog.service
            ##208941 /usr/sbin/rsyslogd -n
```

```
May 30 15:12:31 localhost systemd[1]: Starting System Logging Service...
May 30 15:12:31 localhost rsyslogd[208941]: [origin software="rsyslogd"
swVersion="8.24.0-52.el7_8.2" x-pid="208941" x-info="http://www.rsyslog.com"]
start
May 30 15:12:31 localhost rsyslogd[208941]: error during config processing:
STOP is followed by unreachable statements! [v8.24.0-52.el7_8.2 try http://
www.rsyslog.com/e/2207 ]
May 30 15:12:31 localhost systemd[1]: Started System Logging Service.
[root@rsyslog-server ~]$
```

In case the status shows errors, you can try restarting `systemd-journald` and then again restart `rsyslog` using the following commands.

```
[root@rsyslog-server ~]$systemctl restart systemd-journald
[root@rsyslog-server ~]$systemctl restart rsyslog
[root@rsyslog-server ~]$systemctl status rsyslog
```

Navigate to the `/home/nutanix/rsyslog-server/` directory and verify that localhost files are being created.

Best Practices to Shut Down Policy Engine

Consider the following points before you shut down your policy engine.

- Make sure that no VM related action such as VM create, update, delete, or so on is running in Prism Central.
- Make sure that no actions are running in your Self-Service (if enabled).

You can restart the policy engine VM from the Prism Central VM list page.

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 workflows. 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. For more information, see [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](#).

You can enable or disable Pulse in Prism Central or in a cluster registered to Prism Central. For more information, see [Pulse Configuration](#) on page 162.

Pulse Transport Methods

Nutanix recommends that you configure one of the following Pulse transport methods (in the order of preference):

1. Enable Pulse and use Prism Central as a proxy for the Pulse data transmitted by each node (for clusters registered with Prism Central). For more information, see [Prism Central Proxy for Pulse Data](#) on page 172.

Advantages: The configuration is automatic and no new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

2. Enable Pulse and configure an HTTP proxy server. For information on how to configure an HTTP proxy server, see [Configuring an HTTP Proxy](#) on page 175.

Advantage: No new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

3. Enable Pulse and configure your firewall. Enable Pulse by using each Controller VM IP address in each managed cluster. For information on enabling Pulse in a cluster, see [Enabling Pulse](#) in the *Prism Element Web Console Guide*. For information on configuring the firewall, see the *Pulse Access Requirements* section in the [Pulse Health Monitoring](#) topic of the *Prism Element 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 enables Nutanix Support to request granular diagnostic information from Pulse-enabled clusters. Pulse streams a collection of configuration data, metrics, alerts, events, and select logs to Nutanix Insights, providing a high-level representation of the cluster state. If the Pulse data stream is not detailed enough to diagnose a specific issue, Nutanix Support might need to collect more diagnostic data from the cluster. Remote Diagnostics allows Nutanix to remotely collect the following data only.

- Logs
 - Nutanix services logs
 - Custom gflags set for any Nutanix service
 - Activity traces for Nutanix services
 - Hypervisor logs
 - Hypervisor configuration
 - Cluster configuration
 - System statistics like memory usage
- Nutanix NCC health check reports
- Nutanix log summary report (Panacea)
- A curated set of read-only commands

Each time Remote Diagnostics triggers a collection, it adds an entry to the audit trail of the cluster. There are always two entries, the start (initiation) and finish (termination) of the diagnostics collection.

By default, every Pulse-enabled cluster has Remote Diagnostics enabled. If your security policy or other consideration does not allow cluster access to Nutanix Support for remote diagnostics collection, you can disable Remote Diagnostics without disabling Pulse. Nutanix Support continues to 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$ zcat /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.

Pulse Configuration

When you log in to Prism Central for the first time after an installation or upgrade, the system checks whether Pulse is enabled. If it is not enabled, a pop-up window appears recommending you enable Pulse.

- To enable Pulse in Prism Central and the clusters registered to Prism Central, click **Continue** in the message and follow the prompts.

- To continue without enabling Pulse in Prism Central and the clusters registered to Prism Central, select the **Disable Pulse (not recommended)** checkbox and click **Continue**. For subsequent enablement of Pulse, see [Enabling Pulse](#) on page 163.
- To enable Pulse in Prism Central but not in the clusters registered to Prism Central, select the **Do not enable Pulse for registered clusters. (not recommended)** checkbox and then click **Continue**.
- For Pulse configuration recommendations, see [Pulse Transport Methods](#).

Enabling Pulse

About this task

Perform the following steps to enable Pulse in Prism Central and the clusters registered to Prism Central.

Note:

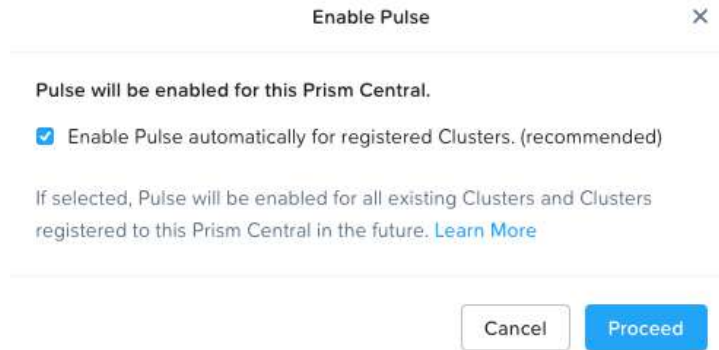
- To enable Pulse in an individual Prism Element cluster, see [Enabling Pulse](#) in the *Prism Element Web Console Guide*.
- Nutanix recommends that you enable Pulse to allow Nutanix Support to receive cluster data and deliver proactive and context-aware support.
- Nutanix does not collect any personally identifiable information (PII) through Pulse.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Setup, click **Pulse**.
The **Enable Insights powered by Pulse** window appears.
5. Click **Enable Pulse**.
The system displays the **Enable Pulse** window.

6. Click **Proceed**.

Note: To prevent enabling Pulse in the clusters registered to Prism Central, clear the **Enable Pulse automatically for registered clusters. (recommended)** checkbox. By default, this checkbox is selected.



Enable Pulse

Pulse will be enabled for this Prism Central.

☒ Enable Pulse automatically for registered Clusters. (recommended)

If selected, Pulse will be enabled for all existing Clusters and Clusters registered to this Prism Central in the future. [Learn More](#)

Cancel Proceed

Figure 77: Enable Pulse

Pulse is enabled in Prism Central and for the clusters registered to Prism Central, and the following message is displayed.

Pulse has been enabled for this Prism Central and enablement is in progress for registered Clusters.

Important:

- Enabling Pulse in Prism Central enables Pulse for all the existing clusters registered to Prism Central and those registered to Prism Central in the future. To prevent enabling Pulse in the clusters registered in the future, disable the **Enable Pulse automatically for all existing Clusters and Clusters registered to this Prism Central in the future. (recommended)** toggle in **Pulse Settings**. By default, this toggle is enabled. However, the existing registered clusters continue to have Pulse enabled in them.
- Nutanix recommends enabling Pulse to allow Nutanix Support to receive system data and deliver proactive and context-aware support.
- You can enable or disable the Pulse settings in an individual registered cluster through the Prism Element web console of the cluster. Changing the Pulse settings in a registered cluster through the Prism Element web console takes precedence over the Pulse setting changes done previously for that cluster through Prism Central.

For example, if a user enables Pulse through Prism Central on all the clusters registered to Prism Central and then disables Pulse on one of the clusters through the Prism Element web console of the cluster, Pulse remains disabled in the cluster, as the changes done through the Prism Element web console takes precedence.

- To reset the Pulse settings for all the registered clusters, you must disable the **Enable Pulse automatically for all existing Clusters and Clusters registered to this Prism Central in the future. (recommended)** toggle in the related Prism Central.

7. Verify the **Pulse Connection Status for Prism Central** section for connection information. You can view the following information.

- **Status:** Displays the transport mechanism status, which is **SUCCESS** when the HTTP client can access the Pulse server successfully or **FAILED** 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.

Disabling Pulse

About this task

Perform the following steps to disable Pulse from Prism Central.

Caution: Nutanix recommends that you enable Pulse to allow Nutanix Support to receive cluster data and deliver proactive and context-aware support. If Pulse is disabled, the clusters do not send alerts to Nutanix Support when problems occur.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Setup, click **Pulse**.
The **Enable Insights powered by Pulse** window appears.
5. Click **Disable Pulse**.
The system displays the **Disable Pulse** window.
6. Click **Disable**.

Note: Disabling Pulse from Prism Central does not disable Pulse from the clusters registered to Prism Central. To disable Pulse from the registered clusters, log in to the Prism Element Web console of each cluster and disable Pulse. For more information, see [Disabling Pulse](#) in the *Prism Element Web Console Guide*.

Pulse is disabled from Prism Central and the following message is displayed.

```
Pulse has been disabled for this Prism Central.
```

Mask Entity Names and IP Addresses

Nutanix processes data that Pulse sends in accordance with your agreement with Nutanix and, where applicable, by the Nutanix Privacy Statement. For more information, see [Nutanix Privacy Statement](#).

For the purpose of Pulse data, you can mask the entity names and IP addresses that are not masked by default. The entity names and IP addresses link to Nutanix entities, such as cluster names, and not to individuals.

To check masking settings on Prism Central, run these commands.

- Check if partial scrubbing is enabled.

```
nutanix@cvm$ curl -H "Content-Type: application/json" -X GET -H "X-Nutanix-Preauth-User:admin" http://localhost:9080/PrismGateway/services/rest/v1/pulse
```

If the output contains the string "identificationInfoScrubbingLevel":"PARTIAL" or "identificationInfoScrubbingLevel":"AUTO", partial scrubbing is enabled. By default, the scrubbing level is set to AUTO. If the output contains the string "identificationInfoScrubbingLevel":"ALL", it means partial scrubbing is not enabled, and the **Open In Prism** button on the Insights Portal remains deactivated.

- Mask the entity names and IP addresses and update the personally identifiable information (PII) scrub level for Prism Central.

```
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 check masking settings for all registered clusters, run these commands on Prism Central.

- Check if partial scrubbing is enabled for all the registered clusters.

```
nutanix@cvm$ curl -H "Content-Type: application/json" -X GET -H "X-Nutanix-Preauth-User:admin" http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
```

If the output contains the string "identificationInfoScrubbingLevel":"PARTIAL" or "identificationInfoScrubbingLevel":"AUTO", partial scrubbing is enabled. By default, the scrubbing level is set to AUTO. If the output contains the string "identificationInfoScrubbingLevel":"ALL", identificationInfoScrubbingLevel is still set to ALL, and the **Open In Prism** button on the Insights Portal remains deactivated.

- Mask the entity names and IP addresses and update the PII scrub level for all registered clusters.

```
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
```

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 24: Pulse Data Collection

Entity	Data Collected
Cluster	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none"> • 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

Note: In this context, hardware can include nodes, blocks, boards, disks, BMCs, fans, DIMMs, BIOS, CPUs, NICs, storage controllers, and power supplies.

Entity	Data Collected
Storage Pool	<ul style="list-style-type: none"> • Name (may be anonymized) • Capacity (logical used capacity and total capacity) • IOPS and latency
Container	<ul style="list-style-type: none"> • 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)	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • Performance stats • Usage
Hypervisor	<ul style="list-style-type: none"> • Hypervisor software and version • Uptime • Installed VMs • Memory usage • Attached datastore
Datastore	<ul style="list-style-type: none"> • Usage • Capacity • Name
Protection Domains	<ul style="list-style-type: none"> • Name (may be anonymized) • Count and names of VMs in each protection domain

Entity	Data Collected
Gflags	<ul style="list-style-type: none"> • Key and value • State (set) • Node ID • Service name • Time of modification
Feature	<ul style="list-style-type: none"> • Feature ID • Name • State (enabled or disabled) • Mode
License	License type (Starter, Ultimate, or Pro)
Alerts	<ul style="list-style-type: none"> • Alert ID • Type • Severity • Resolution status • Acknowledgement status • Impact type • Message • Creation time • Modification time
Tasks	<ul style="list-style-type: none"> • Task ID • Operation type • Status • Entities • Message Completion percentage • Creation time • Modification time
Logs	<ul style="list-style-type: none"> • 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 configured an HTTP proxy on Prism Element clusters registered to this Prism Central instance. If your Prism Element clusters are configured to use an HTTP proxy 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 by satisfying these requirements:

- You have configured an HTTP proxy on Prism Central but it does not require basic authentication (a user name and password). For information on setting up an HTTP proxy on Prism Central, see [Configuring an HTTP Proxy](#) on page 175.
- 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.

Prism Central Configuration 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 172.

To communicate with the Nutanix service center directly, you need to configure the HTTP proxy. For information on how to configure the HTTP proxy through the Prism Central web console, see [Configuring an HTTP Proxy](#) on page 175. 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 an allowed target entry indicates otherwise. For information on who needs to allow the IP addresses, see [Who Needs to Use the Allowlist Method](#) on page 173. To overcome a communication failure between a Prism Element and the registered Prism Central, you need to manually allow the target entries from the Prism Element web console or nCLI. These options enable you to add Prism Central and its managed/registered clusters to an allowlist, where any HTTP proxy settings are ignored. This configuration allows network traffic between them, bypassing any proxy servers configured in the cluster. The allowlist also enables you to register new clusters with Prism Central successfully where clusters are using an HTTP proxy.

For information on how to configure the allowlist IP address through the Prism Central web console, see [Configuring an HTTP Proxy](#) on page 175.

Alternatively, you can use the `ncli http-proxy add-to-whitelist` and `ncli http-proxy delete-from-whitelist` nCLI command options. For information on how to configure the allowlist IP address through the nCLI, see [Allowing Prism Central and Its Managed Clusters \(nCLI\)](#) on page 177.

- You can add or delete one allowlist entry at a time.
- Each allowlist entry cannot exceed 253 characters.
- A maximum of 1000 allowlist entries are supported.
- When deleting an entry from an allowlist, 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 allowlist.
- When applying an allowlist, 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 allow.

Who Needs to Use the Allowlist 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.

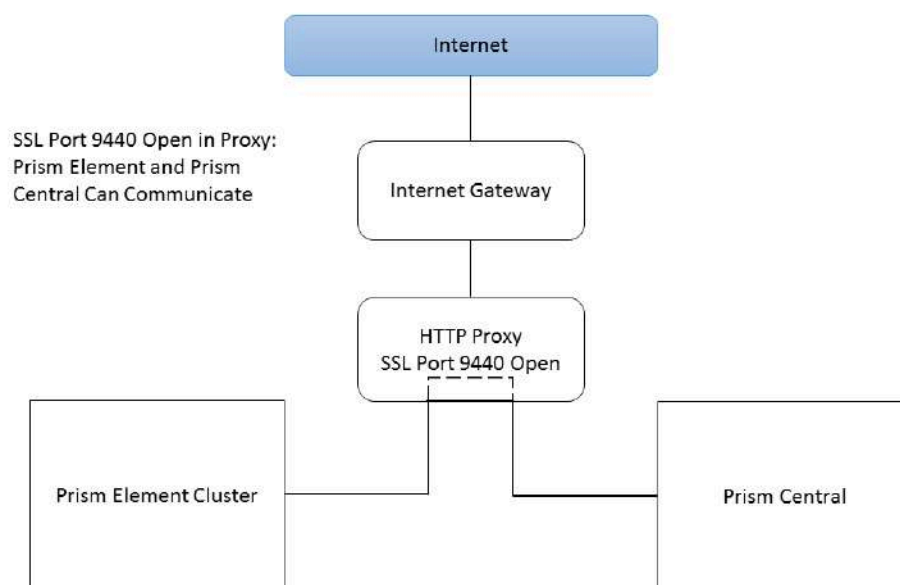


Figure 78: Prism Central, Proxy Server with SSL Port 9440 Open

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. For more information, see [Configuring an HTTP Proxy](#) on page 175 and [Allowing Prism Central and Its Managed Clusters \(nCLI\)](#) on page 177.

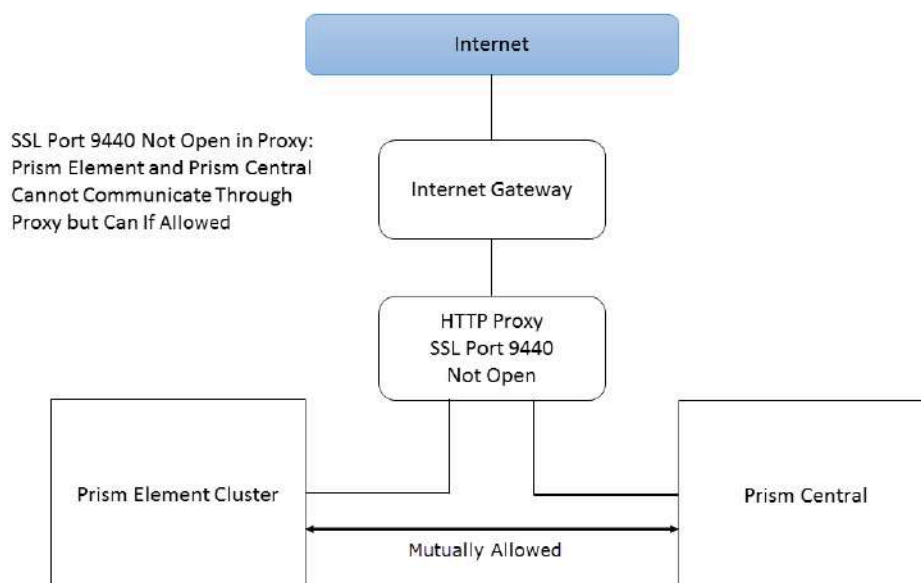


Figure 79: Prism Central, Proxy Server with SSL Port 9440 Closed and with Allowlisting

Configuring an HTTP Proxy

About this task

Note:

- 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 172.

About this task

If Prism Central cannot send traffic to a Nutanix Service Center directly, you need to configure an HTTP proxy. Use the following steps to configure an HTTP proxy.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Network, click **HTTP Proxy**.

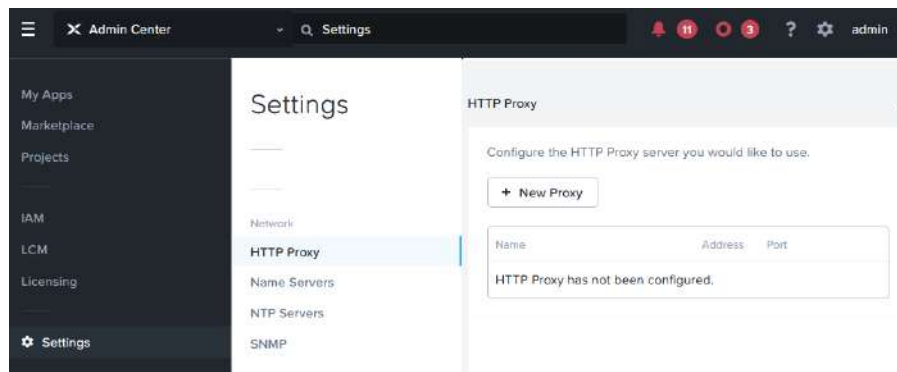


Figure 80: HTTP Proxy Window

5. To add an HTTP proxy, click **New Proxy** and configure the following fields:

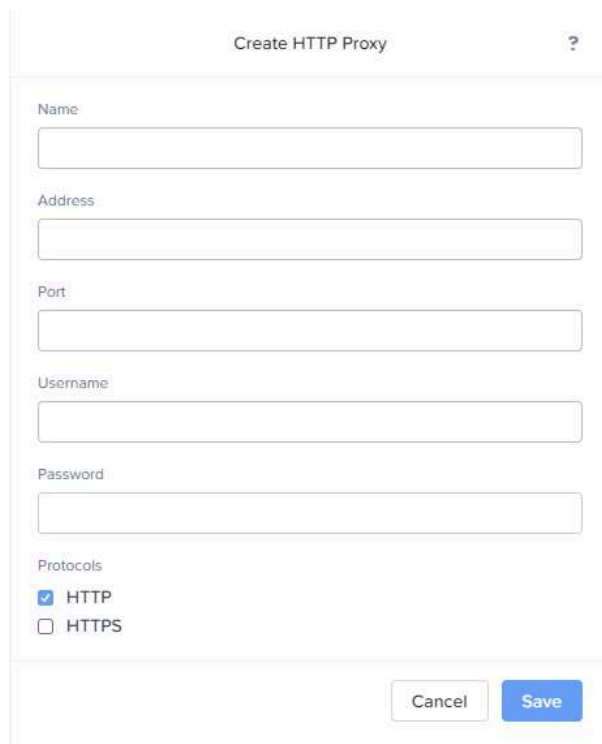
A screenshot of a 'Create HTTP Proxy' dialog box. The dialog has a title bar with the text 'Create HTTP Proxy' and a question mark icon. Inside, there are six input fields: 'Name', 'Address', 'Port', 'Username', and 'Password'. Below these fields is a 'Protocols' section with two checkboxes: 'HTTP' (checked) and 'HTTPS' (unchecked). At the bottom right, there are two buttons: 'Cancel' and 'Save'.

Figure 81: Create HTTP Proxy

Note: You can configure only one HTTP proxy at a time. If an HTTP proxy exists, you must delete the existing proxy 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 the protocol to use.
 - » 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. Click **Save**.

After you save the configuration, the dialog box displays the new HTTP proxy entry in the list.

6. To edit an HTTP proxy entry, click the pencil icon for that entry, update entries, and then click **Save**. Ensure that you add or update the allowlist entry in the correct format. For example, `10.0.0.0/255.255.255.0`

The **Update HTTP Proxy** window appears with the same fields as the **Create HTTP Proxy** window with an additional option to add allowlist entries.

To configure HTTP proxy allowlist entries, do the following:

- To add an allowlist 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 allowlist IP addresses directly rather than through the HTTP or HTTPS proxy.

- To edit an allowlist target, click the pencil icon for that target and update as needed.
- To delete an allowlist target, click the X icon for that target.

The screenshot shows the 'Whitelist' configuration window. At the top, there are checkboxes for 'HTTP' and 'HTTPS', both of which are checked. Below this is a table with two columns: 'TARGET' and 'ACTIONS'. The 'TARGET' column contains a text input field with the placeholder text 'Target'. The 'ACTIONS' column contains two icons: a pencil icon for editing and an 'X' icon for deleting. Below the table, there are 'Save' and 'Cancel' buttons.

Figure 82: Allowlist Targets

Note: Where proxy allowlist configuration is required, it is recommended that you add the cluster virtual IP and all external CVM IPs of each registered cluster to the proxy allowlist for Prism Central. The virtual IP of Prism Central and the IP of each Prism Central VM (PCVM) should be added to the proxy allowlist on the registered Prism Element clusters. For easier management, you can add the subnets instead of adding the individual IPs.

7. To delete an HTTP proxy entry, click the X icon for that entry, and then click **OK** in the confirmation window.

Allowing Prism Central and Its Managed Clusters (nCLI)

About this task

Using this example, you can bypass a proxy server used by a managed Prism Element cluster and allow network traffic between Prism Central and the cluster. This example helps if you attempted to register a cluster that implemented a proxy server, and the registration failed.

Procedure

1. Open an 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 allowlist, then ensure the Prism Central VM IP address was added to the allowlist.

```
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 scale-out Prism Central.

3. Open an SSH session to the Prism Central VM managing the cluster where you just modified the HTTP allowlist.
4. Add the cluster virtual IP address to the allowlist, then ensure the IP address was added to the allowlist.

```
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 allowlist. For example,

The following is a sample command to add a network to the http-proxy allowlist.

```
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 allowed 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 for 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.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.

4. Under Network, click **Name Servers**.

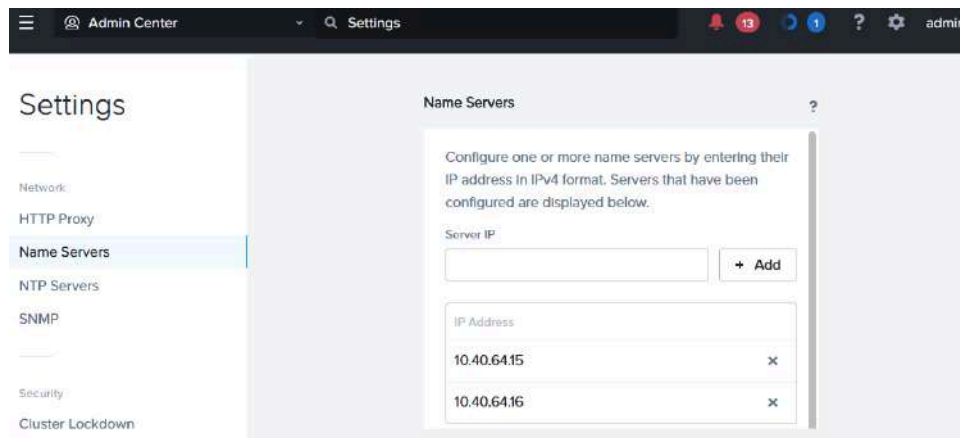


Figure 83: Name Servers

5. To add a name server, enter the server IP address in the **Server** field and then click **Add** to the right of that field.

The server is added to the **IP Address** list.

Note: Changes in the name server configuration might take up to 5 minutes to take effect. Functions that rely on DNS might not work properly during this time. You can configure a maximum of three name servers.

6. To delete a name server entry, click the X icon for that server in the **IP Address** list and then click **OK** in the confirmation window.

Configuring NTP Servers for 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.

Note: If Prism Central is running on Hyper-V, you must specify the IP address of the Active Directory Domain Controller server, not the hostname. Do not use DNS hostnames or external NTP servers.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.

4. Under Network, click **NTP Servers**.

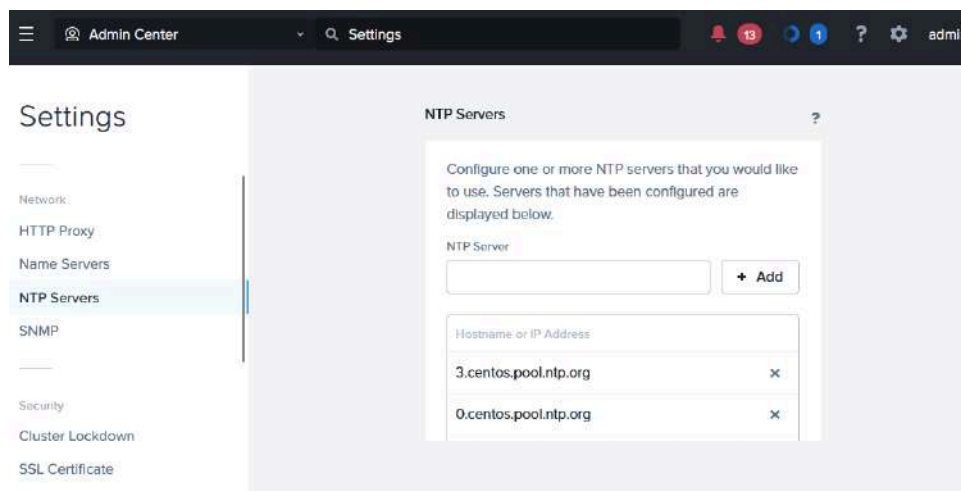


Figure 84: NTP Servers Window

5. To add an NTP server, enter the server IP address or fully-qualified host name in the **NTP Server** field and then click **Add**.
The name or address is added to the **Hostname or IP Address** list.
6. To delete an NTP server entry, click the delete icon for that server in the **Hostname or IP Address** list and then click **OK** in the confirmation window.

Configuring SNMP on 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.

Note:

- Prism Element (individual cluster) supports both the SNMP service (agent) and SNMP traps, but Prism Central supports only SNMP traps. The SNMP capability of Prism Central is limited to just sending traps. For information on configuring SNMP for an individual cluster and for details about the Nutanix MIB, see the [Prism Element Web Console Guide](#).
- The Net-SNMP package version 5.7.2 does not support 256-bit AES encryption.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.

4. Under Network, click **SNMP**.

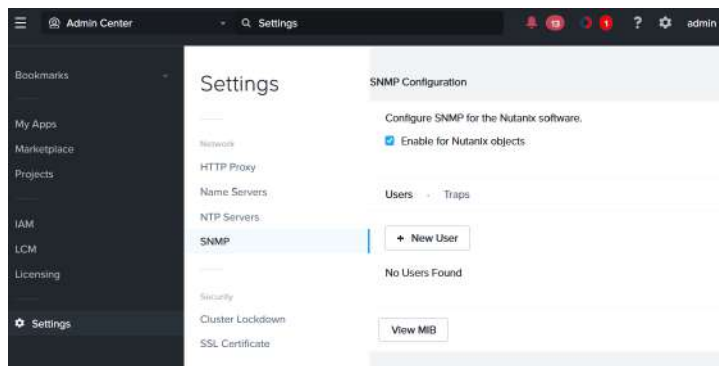


Figure 85: SNMP Configuration

5. To enable SNMP for Nutanix Objects, select the **Enable for Nutanix Objects** checkbox. For information on Nutanix Objects, see the [Nutanix Objects Guide](#).

6. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:

Configure SNMP for the Nutanix software.

☒ Enable for Nutanix objects

Users - Traps

Username

Priv Type Priv Key
AES

Auth Type Auth Key
SHA

Figure 86: SNMP Configuration: Users Tab

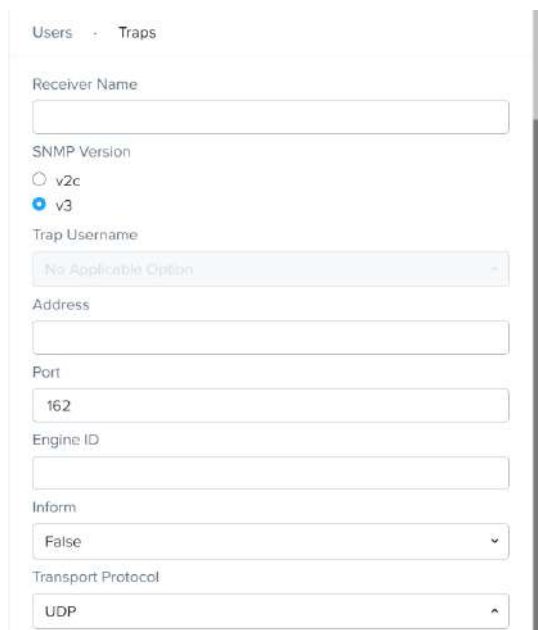
- a. **Username:** Enter a user name.
- b. **Priv Type:** Select **AES** (Advanced Encryption Standard) as the privacy encryption type from the dropdown menu.
- c. **Priv Key:** Enter a privacy key phrase (password) in the field.
 - The key phrase is AES encrypted when the user is created.
 - The password you choose must meet the following complexity requirements:
 - Contain a minimum 8 characters and a maximum 64 characters.
 - Contain one upper case letter (A-Z).
 - Contain one lower case letter (a-z).
 - Contain one digit (0-9).
 - Contain only the following special characters: ~ @ # \$ % * - = _ + { } | [] ? , . : &.
- d. **Auth Type:** Select **SHA** (Secure Hash Algorithm) as the authentication hash function type from the dropdown menu.
- e. **Auth Key:** Enter an authentication key phrase (password) in the field.
 - The key phrase is SHA-1 encrypted when the user is created.
 - The password you choose must meet the following complexity requirements:
 - Contain a minimum 8 characters and a maximum 64 characters.
 - Contain one upper case letter (A-Z).
 - Contain one lower case letter (a-z).
 - Contain one digit (0-9).
 - Contain only the following special characters: ~ @ # \$ % * - = _ + { } | [] ? , . : &.

- f. Click **Save**.

After you save the configuration, the page displays the new user entry in the list.

7. To view the Nutanix MIB (NUTANIX-MIB.txt), click **View MIB**. To download NUTANIX-MIB.txt, right-click and select the appropriate download action for your browser.
8. To add an SNMP trap receiver, click the **Traps** tab.

9. Click **New Trap Receiver** and then do the following in the indicated fields:



The screenshot shows the 'Traps' configuration tab. At the top, there are tabs for 'Users' and 'Traps'. The 'Traps' tab is active. The form contains the following fields:

- Receiver Name:** A text input field.
- SNMP Version:** Radio buttons for 'v2c' and 'v3'. 'v3' is selected.
- Trap Username:** A dropdown menu showing 'No Applicable Option'.
- Address:** A text input field.
- Port:** A text input field with the value '162'.
- Engine ID:** A text input field.
- Inform:** A dropdown menu showing 'False'.
- Transport Protocol:** A dropdown menu showing 'UDP'.

Figure 87: SNMP Configuration: Traps Tab

- Receiver Name:** Enter the receiver name.
- SNMP Version:** Select the SNMP version, either v3 or v2c. For SNMP v2c, Nutanix supports only SNMP TRAP and not SNMP GET.
- The following fields appear based on the SNMP Version you select:
 - Trap Username:** This field appears when you select v3 as the SNMP Version. Select a user from the list.
 - Community:** This field appears when you select v2c as the SNMP Version. The default value for v2c trap community is Public. You can enter any other value of your choice.

All users added previously (see step 7) appear in the list. You cannot add a trap receiver entry until you add at least one user.

- 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.
- 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](#).
- Engine ID:** Enter an engine identifier value, which must be a hexadecimal string between 5 and 32 characters long. This step is optional.

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 dropdown menu 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 list.

The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.

- i. Click **Save**.

After you save the configuration, the page displays the new trap entry 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.

- 10. 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 **Save**.

- 11. To delete an SNMP entry, click the appropriate tab (**Users** or **Traps**), click the X icon for that entry in the list, and then click **OK** in the confirmation box.

Configuring Alert Emails in Prism Central

You can configure the alert messages that are sent from Prism Central. You can configure alert settings, reporting rules, and message templates.

About this task

Alert emails sent from Prism Central are in addition to any alert emails you might have configured on individual clusters through the Prism Element web console. The system sends you email from both the entities in this case. Prism Central alert emailing is not enabled by default; you must explicitly enable it and specify the recipients. If you enable alerts through Prism Central and do not want to receive duplicate email notifications for the same alert, disable customer email notification for those alerts on the individual clusters through Prism Element web console, however keep the email notification support enabled for Nutanix customer support.

Note: Prism Central requires an SMTP server to send alert email messages. For information on how to configure an SMTP server, see [Configuring an SMTP Server for Prism Central](#).

Procedure

To configure alert emails in Prism Central, perform the following steps:

1. Log in to Prism Central as an administrator.

2. Select **Admin Center** from the [Application Switcher Function](#), click **Settings** in the navigation bar, and navigate to **Alerts and Notifications > Alert Email Configuration** or directly click the [Settings icon](#).

For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) information in *Prism Central Infrastructure Guide*.

You can also select the **Infrastructure** application from the [Application Switcher Function](#), and perform the following steps to access the **Alert Email Configuration** window:

- a. Navigate to **Activity > Alerts** from the **Navigation Bar**.
The system displays the **Alerts** page.
 - b. Click **Email Configuration**.
The system displays the **Alert Email Configuration** window.
3. Click the **Settings** tab, and perform the following actions:
 - a. Under **Email Preference**, set the following fields for alert emails:
 - Select the **Every Single Alert** checkbox to configure Prism Central to send an email whenever the event occurs.
 - Select the **Daily Digest** checkbox to configure Prism Central to send a cumulative (24 hour) list of alerts once a day. Set the time when you want the alert to be sent.
 - Select the **Skip the daily digest email if there are no alerts generated on a given day** checkbox to configure Prism Central to send an email only when there are any alerts. Prism Central does not send emails for the days when no alerts occur.
 - If you do not select any checkboxes, no alert emails are sent from Prism Central.
 - b. To send alert notifications to others, enter their email addresses as a comma-separated list in the **Email Recipients** field.
 - c. Click **Save**.

Note: The **Tunnel Connection** section displays mail transport status information.

If you want to configure rules or an email message template, perform the following steps before you click **Save**.

4. Click the **Custom Settings** tab to create a custom alert email setting.

5. Under **Create New Setting** set the following fields:

- a. Specify the conditions to generate the alert:
 - **Alert Severity**: Select one or more of the severities from the list (Critical, Warning, Info, or All).
 - **Impact Type**: Select one or more of the categories from the list (Availability, Capacity, Configuration, Performance, System Indicator, or All).
 - **Cluster**: Select one or more of the clusters from the list (*cluster_name* or 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. Select **Enable** checkbox to enable the new setting. If you have created custom rules (settings) in a version prior to Prism Central 2023.3, those rules are enabled by default after upgrade to Prism Central 2023.3.

Note: Custom settings override the **Email Preference** configured under **Settings** tab, and are honoured even if you do not specify **Email Preference**.

- d. Click **Save** to save the setting.
 - e. Repeat these steps to create more custom settings.
6. Click the **Email Content** tab to create a template for the email messages.

Enter the following field information in the **Email Content** tab:

- **Prepend Subject**: 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.
- **Append Body**: 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.

7. Click **Save**.

Configuring an SMTP Server for Prism Central

About this task

Simple Mail Transport Protocol (SMTP) is an Internet standard protocol for electronic mail transmission across Internet Protocol (IP) networks. Prism Central uses SMTP to send alert emails and to exchange emails. Use the following steps to configure an SMTP server entry.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.

4. Under Alerts and Notifications, click **SMTP Server**.

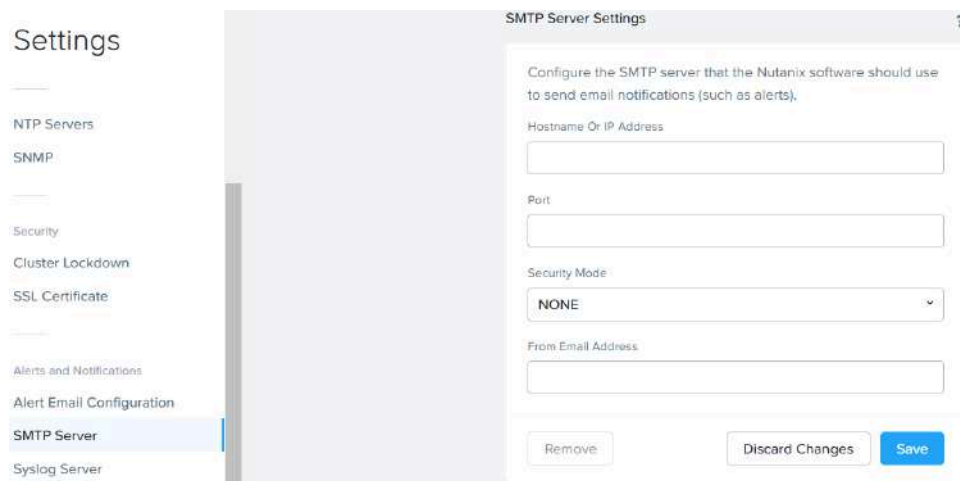


Figure 88: SMTP Server

5. 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 dropdown menu.
The options are **NONE** (unencrypted), **STARTTLS** (use TLS encryption), and **SSL** (use SSL encryption).
- d. **User:** Enter a user name.
The **User** field appears 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.
The **Password** field appears only when a secure option (**STARTTLS** or **SSL**) is selected.
- f. **From Email Address:** Enter an e-mail address that appears as the sender address. This step is optional.
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.

6. Click **Save**.

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. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Alerts and Notifications, click **Syslog Server**.
5. Click **Add Syslog Server**.
6. On the Syslog Server tab, do the following:

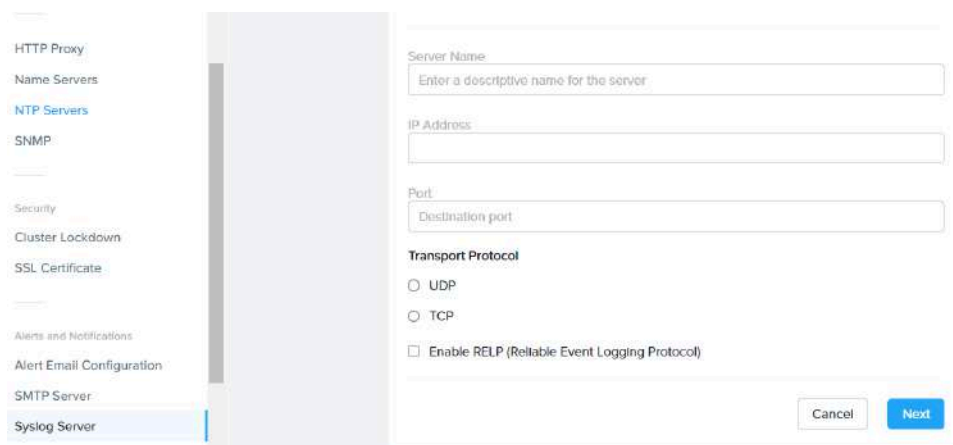


Figure 89: Syslog Server

- a. In **Server Name**, enter a descriptive name for the server.
- b. Enter the **IP Address** and **Port**.
- c. Select the **UDP** or **TCP** transport protocol.
- d. Select the **Enable RELP (Reliable Event Logging Protocol)** checkbox to enable RELP. This step is optional.

Reliable Event Logging Protocol (RELP) is a networking protocol for data logging in computer networks. It extends the functionality of the syslog protocol to provide reliable delivery of event messages.
- e. Click **Next**.

7. On the **Data Sources** tab, do the following:

a. Select one or more log modules from the following log types. For more information, see [Syslog Modules](#) on page 190.

- **API Audit**
- **Audit**
- **Security Policy Hit Logs** (policy hit log files logs)
- **Flow Service Logs** (Flow processes logs)
- **Calm**
- **Epsilon**

The screenshot shows the 'Syslog Servers' configuration window with the 'Data Sources' tab selected. The tab bar at the top has 'Syslog Server' (1) and 'Data Sources' (2). Below the tab bar, the section is titled 'Data Sources and Respective Severity Level'. It contains a table with two columns: 'Module Name' and 'Severity Level'. The table lists six modules: API Audit, Audit, Security Policy Hit Logs, Flow Service Logs, Calm, and Epsilon. Each module has a checkbox in the 'Module Name' column and a dropdown menu in the 'Severity Level' column. The 'Security Policy Hit Logs' dropdown is set to 'All Levels', while the others are set to 'Select Severity Level'. At the bottom of the window, there are three buttons: 'Back', 'Cancel', and 'Save'.

Module Name	Severity Level
<input type="checkbox"/> API Audit	Select Severity Level
<input type="checkbox"/> Audit	Select Severity Level
<input type="checkbox"/> Security Policy Hit Logs	All Levels
<input type="checkbox"/> Flow Service Logs	Select Severity Level
<input type="checkbox"/> Calm	Select Severity Level
<input type="checkbox"/> Epsilon	Select Severity Level

Figure 90: Data Sources

Caution: Users recommend that you only configure rsyslog modules listed in the [Syslog Modules](#) on page 190. Configuration of other modules using ncli commands will be over-written with any modifications done on the Prism Central web console.

b. Select the **Severity Level** for the data sources you selected. For more information, see [Syslog Modules](#) on page 190.

c. Click **Save**.

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 25: 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 <i>Syslog Module: AUDIT Fields</i> 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. <div> Note: Do not enable these unless instructed to do so by Nutanix support for troubleshooting purposes. </div> <p>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.</p>
	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.
Calm	0-7	Nucalm service logs are sent to the syslog server based on the log level you select. These service logs contain information about the module from the Nucalm container (such as styx, hercules, jove, helios, iris, and so on) based on the selected log-level priority. For example, if Warning (level-number: 4) is selected, all the logs that belong to the selected log level are stored on the syslog server.
Epsilon	0-7	Epsilon service logs are sent to the syslog server based on the log level you select. These service logs contain information about the module from the Epsilon container (such as jove, narad, indra, karan, arjun, durga, and so on) based on the selected log-level priority. For example, if Warning (level-number: 4) is selected, all the logs that belong to the selected log level are stored on the syslog server.

The following table provides information on the syslog severity levels.

Table 26: 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	Indication that an error might occur if 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 27: 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 <code>vm</code> 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. <code>VmMigrateAudit</code> and <code>VmCreateAudit</code> .
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. <i>Create</i> , <i>Delete</i> , and <i>Migrate</i> .
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 28: 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. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • Allow - traffic allowed by policy • Drop - traffic dropped by policy • Monitor - traffic allowed by policy in monitor mode

Field	Description
<Packet Stats>	Indicates the packet statistics.

Setting the Prism Central Default Landing Page

About this task

The default landing page is the page that appears when you first log in to Prism Central. You can set any of the deployed Nutanix apps, Infrastructure page, or the Admin Center page as your default landing page. As an administrator, you can also set the system default landing page that determines the default landing page for all other users. Use the following steps to set the default landing page.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **Default Landing Page**.

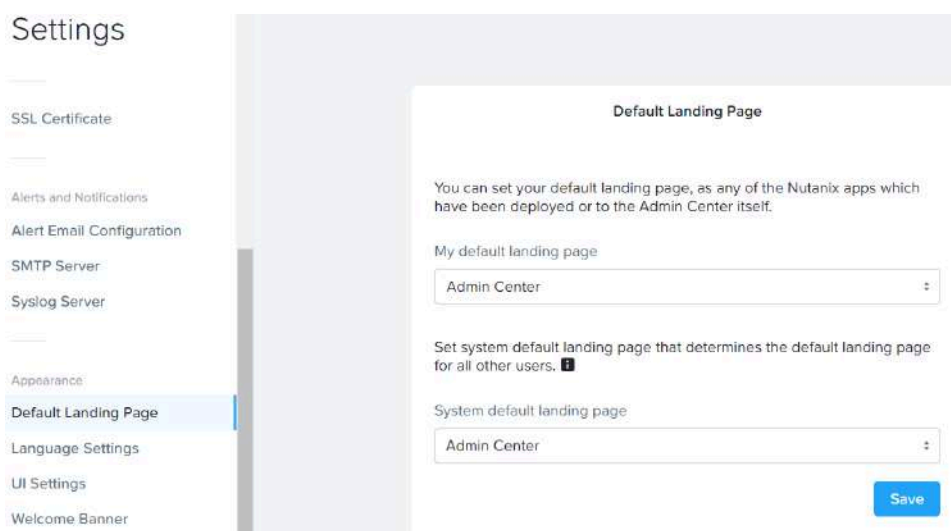


Figure 91: Default Landing Page

5. To set your default landing page, select an option from the **My default landing page** dropdown menu.
6. To set the system default landing page (for all other users), select an option from the **System default landing page** dropdown menu.

Note: Users can override the system default landing page and can set their own default landing page.

7. Click **Save**.

Changing the Prism Central Language Settings

You can change the language setting in Prism Central to Simplified Chinese, Japanese, or Korean (supported only in Self-Service).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **Language Settings**.
5. To change the language setting of the cluster, select a language from the **Language** dropdown menu. You have the following options:

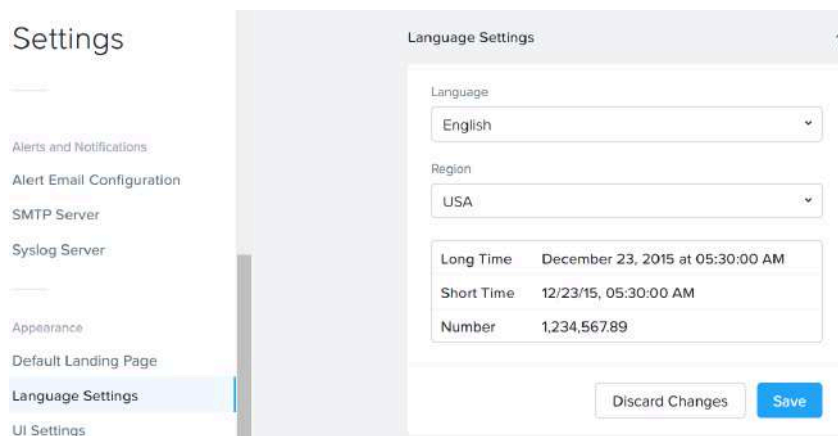


Figure 92: Language Settings

- » **Simplified Chinese**
- » **Japanese**
- » **Korean** (supported only in Self-Service)

By default, the **English** language is selected.

6. To change the locale settings (date, time, and calendar), select the appropriate region from the **Region** dropdown menu.

By default, the locale is set to the language setting that you have set in the **Language** dropdown 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.

7. Click **Save**.

The language and locale settings (date, time, and calendar) is changed according to the selection. For example, in the below image, once you click **Save** the language setting for the cluster is changed to Chinese and locale setting is changed to Russian. For information on the entities that are supported in Simplified Chinese, see [Internationalization \(i18n\) \(Prism Central\)](#) on page 196. Also, the user

interface is localized according to the selection. For information on localization, see [Localization \(L10n\) \(Prism Central\)](#) on page 196.

Figure 93: 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](#).

Internationalization (i18n) (Prism Central)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 29: Internationalization Support for Prism Central

Supported Entities	Unsupported Entities
User-defined dashboard name	Password fields
Custom widget name	Static dashboard name
First and last name under Update Profile	Static widget name
User name, first name, and last name under User Management	
Chart name	

Localization (L10n) (Prism Central)

Nutanix localizes the user interface in Simplified Chinese, Japanese, and Korean (supported only in Self-Service) 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, Japanese, or Korean (supported only in Self-Service). For more information, see [Changing the Prism Central Language Settings](#) on page 194.

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.

Configuring Prism Central UI Settings

About this task

The Prism Central login page includes background animation by default, and users are logged out automatically after being idle for 15 minutes. You can change one or both of these settings.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **UI Settings**.

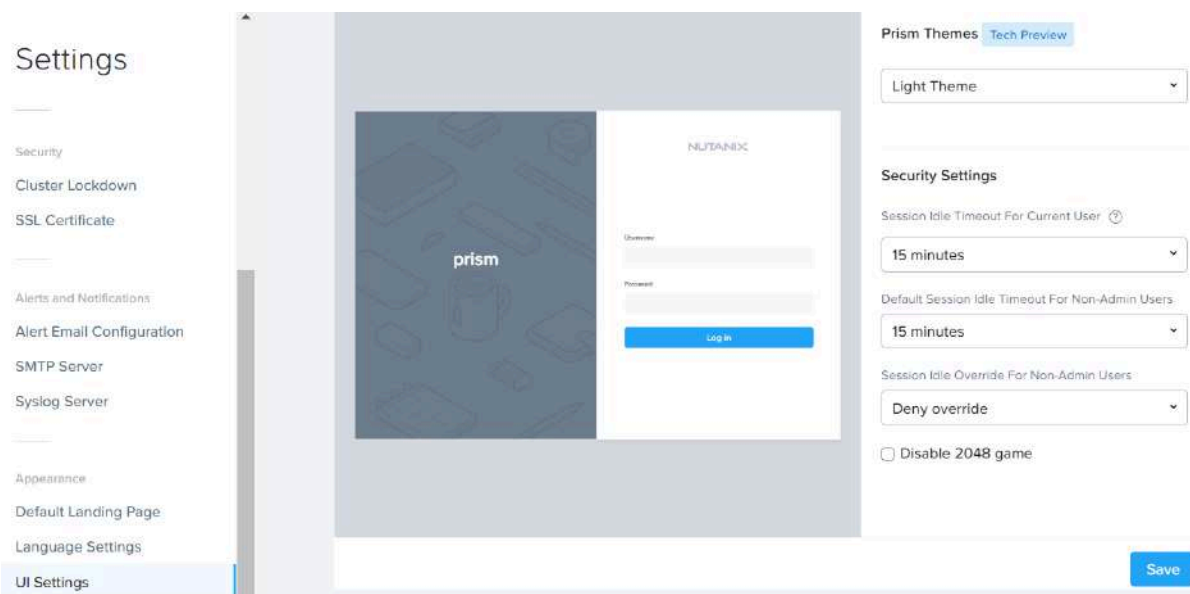


Figure 94: UI Settings

5. To set different variations of the Prism UI background themes, select any of the following options from the **Prism Themes** dropdown 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.

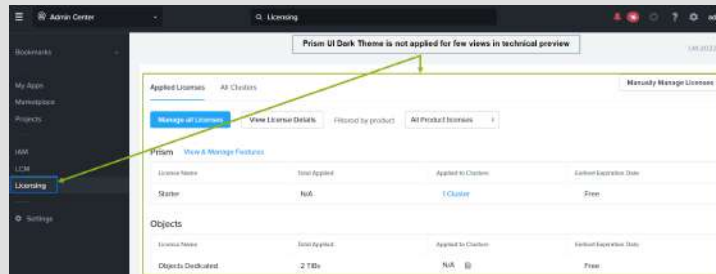


Figure 95: Example: Licensing View (Dark Theme Exception)

- After you change the background theme, save your selection and refresh any open Prism tabs for the changes to reflect.
 - The theme setting is not common between Prism Central and Prism Element web console. You have to configure background themes in Prism Central and Prism Element web console separately.
- Select **Light Theme** for the light background with high-contrast view. This is the default Prism background theme.
 - Select **Dark Theme** for the dark background with high-contrast view, and then click **Continue** to proceed.
 - Select **Auto (OS defined)** to apply background themes defined for the operating system, and then click **Continue** to proceed. For example, if you have defined **Dark** mode in the operating system setting, the Prism UI uses the setting and sets a dark background theme for Prism Central.

6. To customize the theme, background color, title text, or blurb text on the logon page, do the following:

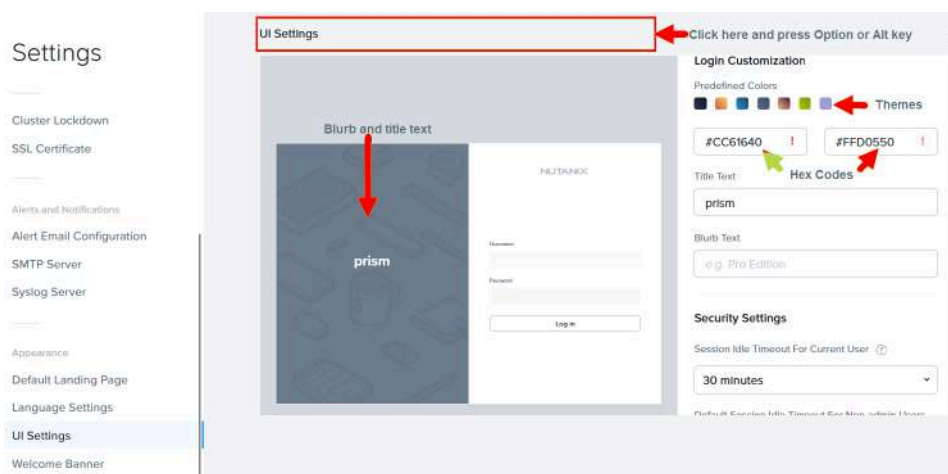


Figure 96: UI Settings - Theme, Title Text, and Blurb Text

- Simultaneously press the Alt key on the Windows system (Option key on the MAC system) and click on the top bar (as shown in the image) of the **UI Settings** page. The options to customize the theme, title text, and blurb text appear on the page.
 - 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.
7. To configure session timeout, do the following under **Security Settings**:
- Select the session idle timeout for the current user from the **Session Idle Timeout For Current User** dropdown menu.
 - Select the default session idle timeout for all non-administrative users from the **Default Session Idle Timeout For Non-Admin Users** dropdown menu.
 - Select the appropriate option from the **Session Idle Override For Non-Admin Users** dropdown menu to override the session timeout for non-administrative users.
- Note:** The Idle timeout interval for an administrator cannot be set for more than an hour.
8. Clear the **Disable 2048 game** checkbox to disable the 2048 game.
9. Click **Save**.

Configuring the Prism Central Welcome Banner

About this task

The welcome banner is the first screen that appears you attempt to log in to Prism Central. As an administrator, you can configure the content of the banner page by adding a custom message or any graphics. Use the following steps to configure the banner page.

Note: Prism Central welcome banner reverts to default after a series of disabling, upgrading and re-enabling IAM. For more information, see [KB-17644](#).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **Welcome Banner**.
5. Enter or 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.

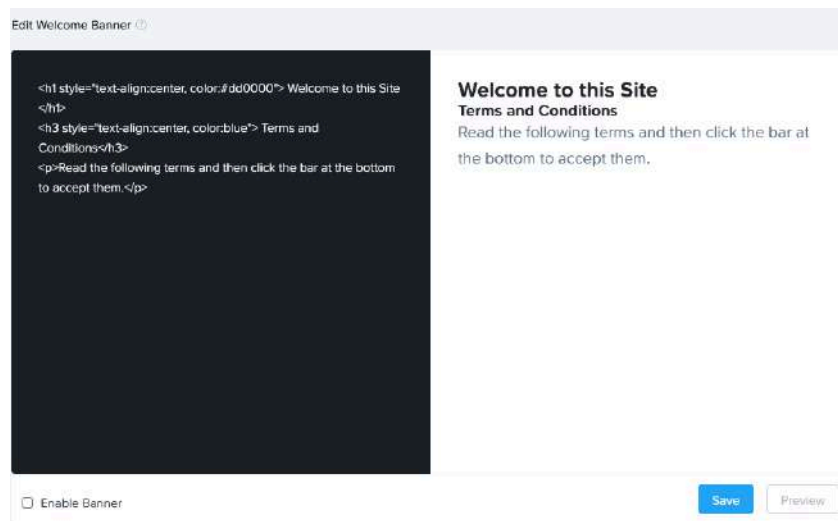


Figure 97: Welcome Banner

6. Click **Preview** to display the banner in the pane on the right.
7. In case the banner is not correct, update the HTML code until the preview pane displays the desired message.
8. To enable the banner, select the **Enable Banner** checkbox and then click **Save**.

A live banner page includes an "Accept terms and conditions" bar at the bottom. Clicking the bar sends the user to the login page.

To disable the banner, clear the **Enable Banner** checkbox.

GLOSSARY

For terms used in this guide, see [Nutanix Glossary](#).

COPYRIGHT

Copyright 2025 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.