

Security Guide

AOS Security 6.8

May 17, 2024

Contents

Audience & Purpose.....	4
Nutanix Security Infrastructure.....	5
SCMA Implementation.....	5
RHEL 8 STIG Implementation in Nutanix Controller VM.....	5
Generating STIG Compliance Report.....	6
Security Updates.....	6
Nutanix Security Landscape.....	6
Security Management Using Prism Central (PC).....	8
Identity and Access Management (IAM).....	8
Configuring Authentication.....	10
User Management.....	20
SSL Certificate Management.....	32
Importing an SSL Certificate.....	32
Generating a Self-signed SSL Certificate with Subject Alternative Name.....	36
Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority (CA).....	39
Controlling Remote (SSH) Access.....	43
Security Policies using Flow.....	44
Data-in-Transit Encryption.....	44
Enabling Data-in-Transit Encryption.....	44
Securing AHV VMs with Virtual Trusted Platform Module.....	46
Considerations for Enabling vTPM in AHV VMs.....	47
Creating a VM with vTPM.....	47
Enabling vTPM on an Existing VM.....	48
Removing vTPM from an Existing VM.....	49
Security Dashboard.....	50
Upgrading Security Dashboard.....	51
Security Widget (PC Main Dashboard).....	51
Security Dashboard Wizard.....	52
Using the Security Dashboard.....	53
Managing Security Dashboard.....	56
Security Management Using Prism Element (PE).....	58
Configuring Authentication.....	58
Assigning Role Permissions.....	67
Authentication Best Practices.....	70
Emergency Local Account Usage.....	70
Modifying Default Passwords.....	70
Controlling Cluster Access.....	71
Setup Admin Session Timeout.....	72
Password Retry Lockout.....	73
Internationalization (i18n).....	73
User Management.....	74
Creating a User Account.....	74

Deleting a User Account (Local).....	82
SSL Certificate Management.....	83
Importing an SSL Certificate.....	83
Generating a Self-signed SSL Certificate with Subject Alternative Name.....	87
Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority (CA).....	90
Exporting an SSL Certificate for Third-party Backup Applications.....	94
Controlling Cluster Access.....	95
Data-at-Rest Encryption.....	96
Data-at-Rest Encryption (SEDs).....	98
Data-at-Rest Encryption (Software Only).....	110
Switching from SED-EKM to Software-LKM.....	124
Configuring Dual Encryption.....	124
Backing up Keys.....	125
Importing Keys.....	126
Securing Traffic Through Network Segmentation.....	127
Traffic Types In a Segmented Network.....	128
Segmented and Unsegmented Networks.....	129
Implementation Considerations.....	134
Configuring the Network on an AHV Host.....	136
Network Segmentation for Traffic Types (Backplane, Management, and RDMA).....	137
Service-Specific Traffic Isolation.....	168
Configuring Backplane IP Pool.....	176
Enabling Backplane Network Segmentation on a Mixed Hypervisor Cluster.....	177
Updating Backplane Portgroup.....	178
IP Address Customization for each CVM and Host.....	179
Enabling Physical Backplane Segmentation on Hyper-V Using CLI.....	180
Network Segmentation during Cluster Expansion.....	181
Network Segmentation-Related Changes During an AOS Upgrade.....	181
Firewall Requirements.....	182
Log management.....	182
Log Forwarding.....	182
Documenting the Log Fingerprint.....	182

Security Management Using Nutanix Command Line Interface

(nCLI).....	183
Hardening Instructions (nCLI).....	183
Hardening AHV.....	183
Hardening Controller VM.....	187
Hardening PCVM.....	191
Common Criteria.....	192
Certificate Revocation Checking (nCLI).....	193
Enabling Certificate Revocation Checking using Online Certificate Status Protocol (nCLI).....	193
Enabling Certificate Revocation Checking using Certificate Revocation Lists (nCLI).....	193
Eliminate Default Passwords during Cluster Creation.....	194
Eliminating Default Passwords during Cluster Creation (CVM-only).....	194

Accessing a List of Open Source Software Running on a Cluster..... 196

Copyright..... 197

AUDIENCE & PURPOSE

This Security Guide is intended for security-minded people responsible for architecting, managing, and supporting infrastructures, especially those who want to address security without adding more human resources or additional processes to their datacenters.

This guide offers an overview of the security development life cycle (SecDL) and host of security features supported by Nutanix. It also demonstrates how Nutanix complies with security regulations to streamline infrastructure security management. In addition to this, this guide addresses the technical requirements that are site specific or compliance-standards (that should be adhered), which are not enabled by default.

Note:

Hardening of the guest OS or any applications running on top of the Nutanix infrastructure is beyond the scope of this guide. We recommend that you refer to the documentation of the products that you have deployed in your Nutanix environment.

NUTANIX SECURITY INFRASTRUCTURE

Nutanix takes a holistic approach to security with a secure platform, extensive automation, and a robust partner ecosystem. The Nutanix security development life cycle (SecDL) integrates security into every step of product development, rather than applying it as an afterthought. The SecDL is a foundational part of product design. The strong pervasive culture and processes built around security harden the Enterprise Cloud Platform and eliminate zero-day vulnerabilities. Efficient one-click operations and self-healing security models easily enable automation to maintain security in an always-on hyperconverged solution.

Since traditional manual configuration and checks cannot keep up with the ever-growing list of security requirements, Nutanix conforms to RHEL 7 Security Technical Implementation Guides (STIGs) that use machine-readable code to automate compliance against rigorous common standards. With Nutanix Security Configuration Management Automation (SCMA), you can quickly and continually assess and remediate your platform to ensure that it meets or exceeds all regulatory requirements.

Nutanix has standardized the security profile of the Controller VM to a security compliance baseline that meets or exceeds the standard high-governance requirements.

The most commonly used references in United States to guide vendors to build products according to the set of technical requirements are as follows.

- The National Institute of Standards and Technology Special Publications Security and Privacy Controls for Federal Information Systems and Organizations (NIST 800.53)
- The US Department of Defense Information Systems Agency (DISA) Security Technical Implementation Guides (STIG)

SCMA Implementation

The Nutanix platform and all products leverage the Security Configuration Management Automation (SCMA) framework to ensure that services are constantly inspected for variance to the security policy.

Nutanix has implemented security configuration management automation (SCMA) to check multiple security entities for both Nutanix storage and AHV. Nutanix automatically reports log inconsistencies and reverts them to the baseline.

With SCMA, you can schedule the STIG to run hourly, daily, weekly, or monthly. STIG has the lowest system priority within the virtual storage controller, ensuring that security checks do not interfere with platform performance.

Note: Only the SCMA schedule can be modified. The AIDE schedule is run on a fixed weekly schedule. To change the SCMA schedule for AHV or the Controller VM, see [Hardening Instructions \(nCLI\)](#) on page 183.

RHEL 8 STIG Implementation in Nutanix Controller VM

Nutanix leverages SaltStack and SCMA to self-heal any deviation from the security baseline configuration of the operating system and hypervisor to remain in compliance. If any component is found as non-compliant, then the component is set back to the supported security settings without any intervention. To achieve this objective, Nutanix has implemented the Controller VM to support STIG compliance with the RHEL 8 STIG as published by DISA.

The STIG rules are capable of securing the boot loader, packages, file system, booting and service control, file ownership, authentication, kernel, and logging.

Example: STIG rules for Authentication

Prohibit direct `root` login, lock system accounts other than `root`, enforce several password maintenance details, cautiously configure SSH, enable screen-locking, configure user shell defaults, and display warning banners.

For more information on Nutanix STIGs, see [Security Dashboard STIG Guidance Reference](#) and [Security Dashboard](#) on page 50.

Generating STIG Compliance Report

You can leverage OpenSCAP report to ensure that the Controller VM is STIG compliant.

About this task

To generate the STIG compliance report, do the following:

Procedure

1. Log on to any Controller VM in the cluster with SSH.
2. Create an output directory for the STIG report.

```
nutanix@cvm$ mkdir /tmp/stig_report
```

A directory named "stig_report" is created under the "tmp" directory.

3. Generate the STIG compliance report.

```
nutanix@cvm$ sudo python3 -B /home/nutanix/config/security/stig_scripts/  
stig_scan_driver.py /home/nutanix/config/security/stig_scripts/  
U_RHEL_7_V3R8_STIG_SCAP_1-2_Benchmark.xml /tmp/stig_report/
```

The STIG compliance report is generated in the "/tmp/stig_report" directory.

Security Updates

Nutanix provides continuous fixes and updates to address threats and vulnerabilities. *Nutanix Security Advisories* provide detailed information on the available security fixes and updates, including the vulnerability description and affected product/version.

To see the list of security advisories or search for a specific advisory, log on to the [Support Portal](#) and select **Documentation**, and then **Security Advisories**.

Nutanix Security Landscape

This topic provides highlights on Nutanix security landscape and its highlights. The following table helps to identify the security features offered out-of-the-box in Nutanix infrastructure.

Topic	Highlights
Authentication and Authorization	<ul style="list-style-type: none">• Support for Authentication types and directories• Role Permissions• Password Complexity Support with standard Pluggable Authentication Module (PAM) library
Network segmentation	VLAN-based, data driven segmentation
Security Policy Management	Implement security policies using Microsegmentation .

Topic	Highlights
Data security and integrity	<ul style="list-style-type: none"> • Cluster access control • SSL certificate management • Data-at-rest (DAR) encryption
Hardening Instructions	<ul style="list-style-type: none"> • Hardening AHV • Hardening Controller VM • TCP Wrapper Integration
Log monitoring and analysis	<ul style="list-style-type: none"> • Cluster-wide log shipping • Documenting the log fingerprint • See Pulse Health Monitoring in the <i>Prism Element Web Console Guide</i>
Flow Networking	See Flow Networking Guide
UEFI	See UEFI Support for VMs in the <i>AHV Administration Guide</i>
Secure Boot	See Secure Boot Support for VMs in the <i>AHV Administration Guide</i>
Windows Credential Guard support	See Windows Defender Credential Guard Support in AHV in the <i>AHV Administration Guide</i>
RBAC	See Controlling User Access (RBAC)

SECURITY MANAGEMENT USING PRISM CENTRAL (PC)

Prism Central provides several mechanisms and features to enforce security of your multi-cluster environment.

Identity and Access Management (IAM)

IAM Introduction

Administered from Prism Central, Identity and Access Management (IAM) is an authentication and authorization feature that uses fine-grained RBAC that allows users to create custom roles from granular operations/permissions. Using IAM, you can configure the scope of the roles on subset of entities, like VMs, clusters, and categories.

IAM is enabled by default. When Microservices Infrastructure is enabled (whether by default or manually) on Prism Central, IAM is automatically enabled. See [Microservices Infrastructure](#).

Using Prism Central IAM, you can configure authentication for local users, directory services and a wide selection of identity providers, including SAML-based identity providers, see [Configuring Authentication](#) on page 10.

For authorization, IAM supports granular role-based access control (RBAC) that you can configure to provide customized access permissions for users based on their assigned roles, see [Controlling User Access \(RBAC\)](#) on page 25

IAM Features

Highly Scalable Architecture

Based on the Kubernetes open source platform, IAM uses independent pods for authentication (AuthN), authorization (AuthZ), and IAM data storage and replication.

- IAM uses a rolling upgrade method to help ensure zero downtime.

Secure by Design

- Mutual TLS authentication (mTLS) secures IAM component communication.
- The Microservices Infrastructure on Prism Central provisions certificates for mTLS.

More SAML Identity Providers (IDP)

IAM supports multiple IDPs, but Nutanix has tested the following IDPs for SAML IDP authentication in Prism Central.

- Microsoft Active Directory Federation Services (ADFS)
- Microsoft Entra ID (formerly Microsoft Azure ADFS)

Note: IAM supports SAML-based authentication with Microsoft Entra ID. LDAP-based authentication with Microsoft Entra ID is not supported.

- Okta
- PingOne
- Shibboleth
- Keycloak

Nutanix supports SAML 2.0 compliant IDPs.

Users can log on from the Prism Central web console only. IDP-initiated authentication work flows are not supported. That is, logging on or signing on from an IDP web page or site is not supported.

IAM Prerequisites

For specific minimum software support and requirements for IAM, see the [Prism Central release notes](#).

For microservices infrastructure requirements, see [Microservices Infrastructure](#).

Prism Central

See [Microservices Infrastructure](#).

Prism Element Clusters

See [Microservices Infrastructure](#).

IAM Considerations

First Log on after Upgrading to Prism Central Version

After you upgrade Prism Central to a minimum version of pc.2022.9, when you log in to Prism Central for the first time using IAM, click **Log in with your Nutanix Local Account** and log in using the default **admin** credentials.

After this first time login, for subsequent login you can use your Active Directory (AD)

Existing Authentication and Authorization Migrated When IAM is Enabled

- IAM migrates existing authentication and authorization configurations, including Common Access Card client authentication configurations.

If Security Assertion Markup Language (SAML) IDP is configured

Signed single logout (SAML SLO) is not supported. This limitation results in a error on the SAML ADFS page when you logout of Prism Central

Upgrading Prism Central After Enabling CMSP

After you upgrade Prism Central, if Microservices Infrastructure (and IAM) was previously enabled, both the services are enabled by default if the cluster meets all the requirements as provided in [Microservices Infrastructure](#). You must contact Nutanix Support for any custom requirement.

Note: After upgrade to pc.2022.9 if the Security Assertion Markup Language (SAML) IDP is configured, you need to download the Prism Central metadata and re-configure the SAML IDP to recognize Prism Central as the service provider. See [Updating ADFS When Using SAML Authentication](#) on page 17 to create the required rules for ADFS.

User Session Lifetime

- Each session has a maximum lifetime of 8 hours
- Session idle time is 15 minutes. After 15 minutes, a user or client is logged out and must re-authenticate.

Client Authentication and Common Access Card (CAC) Support

- IAM supports deployments where CAC authentication and client authentication are enabled on Prism Central. If you want to enable a client to authenticate by using certificates, you must also enable CAC authentication.

Note: IAM does not support client authentication if CAC authentication is not enabled.

- Ensure that port 9441 is open in your firewall if you are using CAC client authentication.

Hypervisor Support

- IAM is enabled only on an on-premise Prism Central (PC) deployment hosted on an AOS cluster running AHV or ESXi. Clusters running other hypervisors are not supported.

Deleting Users from Prism Central

- Once local or directory users are configured, they cannot be deleted from Prism Central. However, you can disable any user account, and re-enable as needed.

Updating Authorization Policy Scope

- When you change the scope in an authorization policy to include or exclude individual entities, it takes approximately five minutes for the filtered items to update and reflect on the entity page. During this period, the entities may display incorrectly.

Migration of Built-in Roles in Prism Central Version pc.2024.1

- Starting with Prism Central version pc.2024.1, some built-in roles such as "User Admin" and "Viewer" are migrated with different role names but retain similar permissions. If these roles were associated with any users, the role mapping configuration will be migrated as new authorization policies created automatically upon upgrade to Prism Central version pc.2024.1.

Limited Role Details View

- The role details view does not display the associated authorization policies and users.

Configuring Authentication

Caution: Prism Central does not support the SSLv2 and SSLv3 ciphers. Therefore, you must disable the SSLv2 and SSLv3 options in a browser before accessing Prism Central. This disabling avoids an SSL Fallback and access denial situations. However, you must enable TLS protocol in the browser.

Prism Central IAM supports user authentication with these authentication options:

- SAML authentication. Users can authenticate through a supported identity provider when SAML support is enabled for Prism Central. The Security Assertion Markup Language (SAML) is an open standard for exchanging authentication and authorization data between two parties: an identity provider (IDP) and Prism Central as the service provider.

With IAM, in addition to ADFS, other IDPs are available. For more information, see [#unique_27/unique_27_Connect_42_saml-idp-support](#) and [Updating ADFS When Using SAML Authentication](#) on page 17.

- Local user authentication. Users can authenticate if they have a local Prism Central account. For more information, see [Managing Local User Accounts](#).
- Active Directory authentication. Users can authenticate using their Active Directory (or OpenLDAP) credentials when Active Directory support is enabled for Prism Central.

The Prism Central login page is updated depending on your IAM configuration. For example, if you have configured local user account and Active Directory authentication, this default page appears for directory (AD) users as follows. To log in as a local user, click the **Log In with your Nutanix Local Account** link.

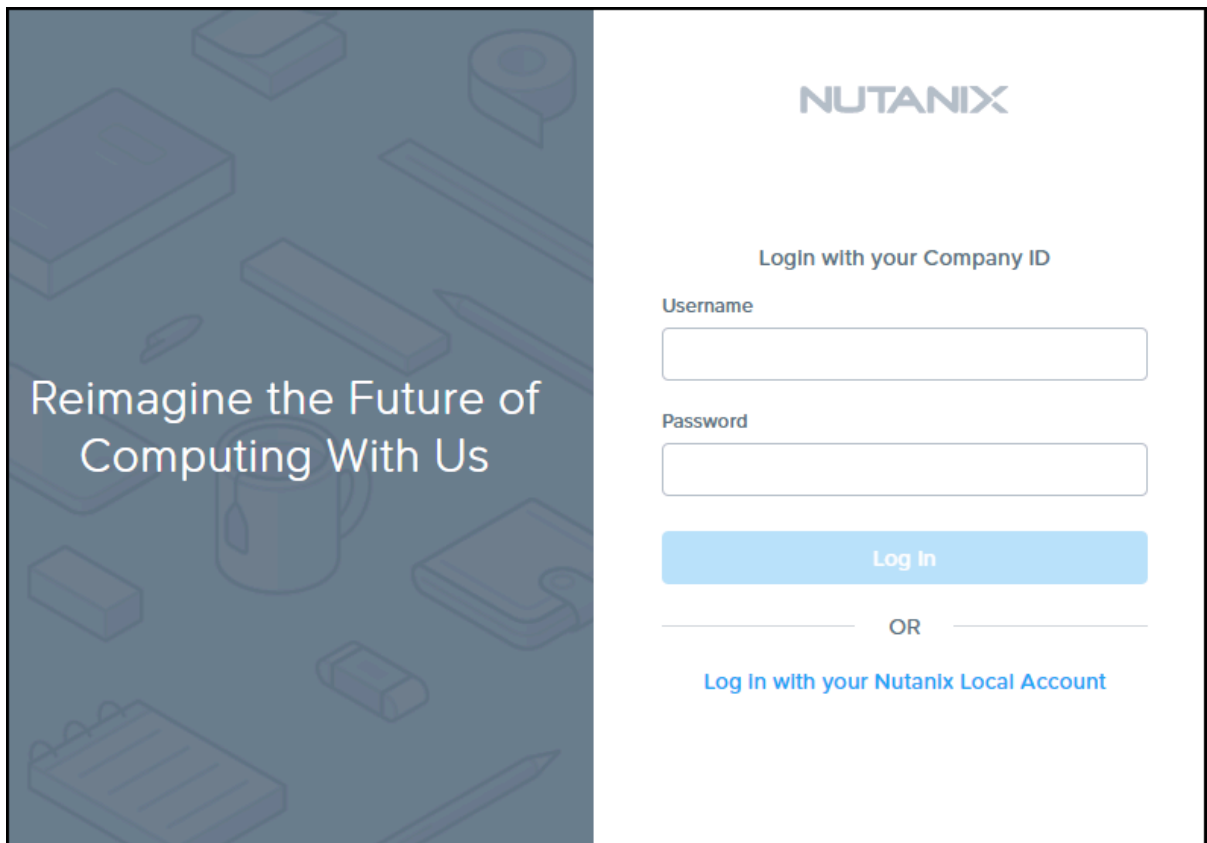


Figure 1: Sample Default Prism Central IAM Logon Page, Active Directory And Local User Authentication

In another example, if you have configured SAML authentication instances named Shibboleth and AD2, Prism Central displays this page.

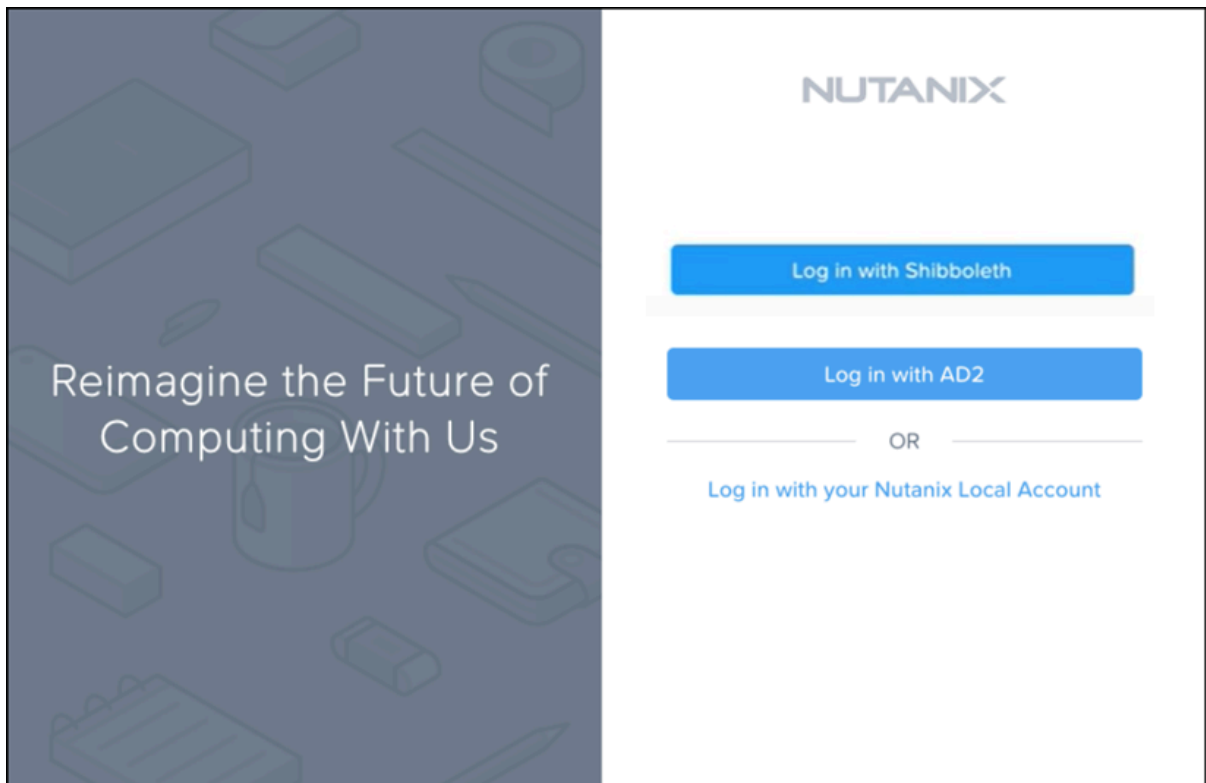


Figure 2: Sample Prism Central IAM Logon Page, Active Directory , Identity Provider, And Local User Authentication

Note: After upgrade to pc.2022.9 if the Security Assertion Markup Language (SAML) IDP is configured, you need to download the Prism Central metadata and re-configure the SAML IDP to recognize Prism Central as the service provider. See [Updating ADFS When Using SAML Authentication](#) on page 17 to create the required rules for ADFS.

Adding An Authentication Directory (Active Directory/OpenLDAP)

Before you begin

Caution: Prism Central does not allow the use of the (not secure) SSLv2 and SSLv3 ciphers. To eliminate the possibility of an SSL Fallback situation and denied access to Prism Central, disable (uncheck) SSLv2 and SSLv3 in any browser used for access. However, TLS must be enabled (checked).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to the **IdP Configuration** tab.
4. Click **Add Identity Provider > Active Directory/OpenLDAP**.

The **Configure Directory** window appears.

5. In the **Configure Directory** window, a set of fields is displayed. Do the following in the indicated fields:

- **Active Directory:** Active Directory (AD) is a directory service implemented by Microsoft for Windows domain networks.

Note:

- Users with the "User must change password at next logon" attribute enabled will not be able to authenticate to Prism Central. Ensure users with this attribute first login to a domain workstation and change their password prior to accessing Prism Central. Also, if SSL is enabled on the Active Directory server, make sure that Nutanix has access to that port (open in firewall).
- Use of the "Protected Users" group is currently unsupported for Prism authentication. For more details on the "Protected Users" group, see "Guidance about how to configure protected accounts" on Microsoft documentation website.
- An Active Directory user name or group name containing spaces is not supported for Prism Central authentication.
- The Microsoft AD is LDAP v2.
- The Microsoft AD servers supported are Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019.

- **OpenLDAP:** OpenLDAP is a free, open source directory service, which uses the Lightweight Directory Access Protocol (LDAP), developed by the OpenLDAP project.

Note: Prism Central uses a service account to query OpenLDAP directories for user information and does not currently support certificate-based authentication with the OpenLDAP directory.

- a. **Name:** Enter a directory name.

This is a name you choose to identify this entry; it need not be the name of an actual directory.

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

Enter the domain name in DNS format, for example, nutanix.com.

- c. **Directory URL:** Enter the URL address to the directory.

The URL format is as follows for an LDAP entry: `ldap://host:ldap_port_num`. The host value is either the IP address or fully qualified domain name. (In some environments, a simple domain name is sufficient.) The default LDAP port number is 389. Nutanix also supports LDAPS (port 636)

and LDAP/S Global Catalog (ports 3268 and 3269). The following are example configurations appropriate for each port option:

Note: LDAPS support does not require custom certificates or certificate trust import.

- Port 389 (LDAP). Use this port number (in the following URL form) when the configuration is single domain, single forest, and not using SSL.

```
ldap://ad_server.mycompany.com:389
```

- Port 636 (LDAPS). Use this port number (in the following URL form) when the configuration is single domain, single forest, and using SSL. This requires all Active Directory Domain Controllers have properly installed SSL certificates.

```
ldaps://ad_server.mycompany.com:636
```

- Port 3268 (LDAP - Global Catalog). Use this port number when the configuration is multiple domain, single forest, and not using SSL.
- Port 3269 (LDAPS - Global Catalog). Use this port number when the configuration is multiple domain, single forest, and using SSL.

Note:

- When constructing your LDAP/S URL to use a Global Catalog server, ensure that the Domain Control IP address or name being used is a global catalog server within the domain being configured. If not, queries over 3268/3269 may fail.
- Cross-forest trust between multiple AD forests is not supported.

For the complete list of required ports, see [Port Reference](#).

- d. [OpenLDAP only] Configure the following additional fields:

Note:

The value for the following variables depend on your OpenLDAP configuration.

- **User Object Class:** Enter the value that uniquely identifies the object class of a user.
- **User Search Base:** Enter the base domain name in which the users are configured.
- **Username Attribute:** Enter the attribute to uniquely identify a user.
- **Group Object Class:** Enter the value that uniquely identifies the object class of a group.
- **Group Search Base:** Enter the base domain name in which the groups are configured.
- **Group Member Attribute:** Enter the attribute that identifies users in a group.
- **Group Member Attribute Value:** Enter the attribute that identifies the users provided as value for **Group Member Attribute**.

Here are some of the possible options for the fields:

- User Object Class: user | person | inetOrgPerson | organizationalPerson | posixAccount
- User Search Base: ou=<organizational unit>, dc=<domain>
- Username Attribute: uid

- Group Object Class: posixGroup | groupOfNames
 - Group Search Base: ou=<organizational unit>, dc=<domain>
 - Group Member Attribute: member | memberUid
 - Group Member Attribute Value: uid
- e. **Search Type.** How to search your directory when authenticating. Choose **Non Recursive** if you experience slow directory logon performance. For this option, ensure that users listed in Role Mapping are listed flatly in the group (that is, not nested). Otherwise, choose the default **Recursive** option.
- f. **Service Account Username:** Depending upon the **Directory type** you select in step 2.a, the service account user name format as follows:
- For *Active Directory*, enter the service account user name in the `user_name@domain.com` format.
 - For *OpenLDAP*, enter the service account user name in the following Distinguished Name (DN) format:

`cn=username, dc=company, dc=com`

A service account is created to run only a particular service or application with the credentials specified for the account. According to the requirement of the service or application, the administrator can limit access to the service account.

A service account is under the Managed Service Accounts in the Active Directory and openLDAP server. An application or service uses the service account to interact with the operating system. Enter your Active Directory and openLDAP service account credentials in this (username) and the following (password) field.

Note: Be sure to update the service account credentials here whenever the service account password changes or when a different service account is used.

- g. **Service Account Password:** Enter the service account password.

- h. When all the fields are correct, click the **Save** button (lower right).

This saves the configuration and redisplay the Authentication Configuration dialog box. The configured directory now appears in the **Directory List** tab.

- i. Repeat this step for each authentication directory you want to add.

Note:

- No permissions are granted to the directory users by default. To grant permissions to the directory users, you must specify roles for the users in that directory (see [Configuring an Authorization Policy](#) on page 30).
- Service account for both Active directory and openLDAP must have full read permission on the directory service.

Adding a SAML-based Identity Provider

Before you begin

- An identity provider (typically a server or other computer) is the system that provides authentication through a SAML request. There are various implementations that can provide authentication services in line with the SAML standard.

- You can specify other tested standard-compliant IDPs in addition to ADFS. See the [Prism Central release notes](#) topic *Identity and Access Management Software Support* for specific support requirements.
- You must configure the identity provider to return the `NameID` attribute in SAML response. Prism Central uses the `NameID` attribute for role mapping.
- IAM only supports redirect binding for IDP sign-on URLs. You must configure and include redirect binding in IDP metadata.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to the **IdP Configuration** tab.
4. Click **Add Identity Provider > SAML Identity Provider**.
5. In the **Configure Identity Provider** window, a set of fields is displayed. Do the following in the indicated fields:
 - a. **Configuration name:** Enter a name for the identity provider. This name appears in the logon authentication screen.
 - b. **Username Attribute:** Enter the username attribute of the identity provider configuration.
 - c. **Email Attribute:** Enter the email attribute of the identity provider configuration.
 - d. **Group Attribute Name (Optional):** Optionally, enter the group attribute name such as **groups**. Ensure that this name matches the group attribute name provided in the IDP configuration.
 - e. **Group Attribute Delimiter (Optional):** Optionally, enter a delimiter that needs to be used when multiple groups are selected for the Group attribute.
 - f. **Import Metadata:** Click this option to upload a metadata file that contains the identity provider information.

Identity providers typically provide an XML file on their website that includes metadata about that identity provider, which you can download from that site and then upload to Prism Central. Click **Import Metadata** to open a search window on your local system and then select the target XML file that you downloaded previously. Click the **Save** button to save the configuration.

This step completes configuring an identity provider in Prism Central, but you must also configure the callback URL for Prism Central on the identity provider. To configure the callback URL, click the **Download Metadata** link just below the Identity Providers table to download an XML file that describes Prism Central and then upload this metadata file to the identity provider.

If your identity provider asks you to enter the Prism Central callback URL manually, you can find it in the `AssertionConsumerService` tag with attribute `Location` in the downloaded XML metadata file. The following sample XML metadata file shows `AssertionConsumerService` with attribute `Location`.

```
<md:AssertionConsumerService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
  Location="https://10.46.153.235:9440/api/iam/authn/callback"
  index="1" />
```

6. To edit a directory entry, select the directory from the **Identity Provider** list and click **Actions > Edit**.

7. To delete a directory entry, select the directory from the **Identity Provider** list and click **Actions > Delete**.

Updating ADFS When Using SAML Authentication

With Nutanix IAM, to maintain compatibility with new and existing IDP/SAML authentication configurations, update your Active Directory Federated Services (ADFS) configuration - specifically the Prism Central Relying Party Trust settings. For these configurations, you are using SAML as the open standard for exchanging authentication and authorization data between ADFS as the identity provider (IDP) and Prism Central as the service provider. See the Microsoft Active Directory Federation Services documentation for details.

About this task

In your ADFS Server configuration, update the Prism Central Relying Party Trust settings by creating claim rules to send the selected LDAP attribute as the SAML NameID in email address format. For example, map the User Principal Name to NameID in the SAML assertion claims.

As an example, this topic uses UPN as the LDAP attribute to map. You could also map the email address attribute to NameID. See the *Microsoft Active Directory Federation Services* documentation for details about creating a claims aware Relying Party Trust and claims rules.

Procedure

1. In the Relying Party Trust for Prism Central, configure a claims issuance policy with two rules.
 - a. One rule based on the **Send LDAP Attributes as Claims** template.
 - b. One rule based on the **Transform an Incoming Claim** template
2. For the rule using the **Send LDAP Attributes as Claims** template, select the **LDAP Attribute** as **User-Principal-Name** and set **Outgoing Claim Type** to **UPN**.
For User group configuration using the **Send LDAP Attributes as Claims** template, select the **LDAP Attribute** as **Token-Groups - Unqualified-Names** and set **Outgoing Claim Type** to **Group**.
3. For the rule using the **Transform an Incoming Claim** template:
 - a. Set **Incoming claim type** to **UPN**.
 - b. Set the **Outgoing claim type** to **Name ID**.
 - c. Set the **Outgoing name ID format** to **Email**.
 - d. Select **Pass through all claim values**.

Enabling and Configuring Client Authentication/CAC

Before you begin

- To enable a client to authenticate by using certificates, you must also enable CAC authentication.
- Ensure that port 9441 is open in your firewall if you are using CAC client authentication. After enabling CAC client authentication, your CAC logon redirects the browser to use port 9441.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.

3. Select **IAM** and go to the **IdP Configuration** tab.

4. Click **Add Identity Provider > Common Access Cards**.

The **Configure Client Chain Certificate** window appears.

5. In the **Configure Client Chain Certificate** window, do the following steps.

- a. Click the **Upload** button, browse to and select a client chain certificate to upload, and then click the **Open** button to upload the certificate.

Note: Uploaded certificate files must be PEM encoded. The web console restarts after the upload step.

- b. To enable client authentication, click **Enable Client Authentication**.

- c. To modify client authentication, do one of the following:

- Click **Enable Client Authentication** to disable client authentication.
- Click **Remove** to delete the current certificate. (This deletion also disables client authentication.)
- To enable OCSP or CRL-based certificate revocation checking, see [Certificate Revocation Checking](#).

Client authentication allows you to securely access the Prism by exchanging a digital certificate. Prism validates if the certificate is signed by the trusted signing certificate of your organization.

Client authentication ensures that the Nutanix cluster gets a valid certificate from the user. Normally, a one-way authentication process occurs where the server provides a certificate so the user can verify the authenticity of the server. When client authentication is enabled, this process becomes a two-way authentication where the server also verifies the authenticity of the user. A user must provide a valid certificate when accessing the console either by installing the certificate on the local machine, or by providing it through a smart card reader.

Note: The CA must be the same for both the client chain certificate and the certificate on the local machine or smart card.

6. Select the active directory for CAC authentication.

- a. From the **Select Active Directory** drop-down list, select the authentication directory that contains the CAC users that you want to authenticate.

This list includes the directories that are configured on the **Identity Provider** tab.

The Common Access Card (CAC) is a smart card about the size of a credit card, which some organizations use to access their systems. After you insert the CAC into the CAC reader connected to your system, the software in the reader prompts you to enter a PIN. After you enter a valid PIN, the software extracts your personal certificate that represents you and forwards the certificate to the server using the HTTP protocol.

Nutanix Prism verifies the certificate as follows:

- Validates that the certificate has been signed by the trusted signing certificate of your organization.
- Extracts the Electronic Data Interchange Personal Identifier (EDIP) from the certificate and uses the EDIP to check the validity of an account within the Active Directory. The security context from the EDIP is used for your PRISM session.
- Prism Central supports both certificate authentication and basic authentication in order to handle both Prism Central login using a certificate and allowing REST API to use basic authentication. It is physically not possible for REST API to use CAC certificates. With this behavior, if the certificate is

present during Prism Central login, the certificate authentication is used. However, if the certificate is not present, basic authentication is enforced and used.

If you map a Prism Central role to a CAC user and not to an Active Directory group or organizational unit to which the user belongs, specify the EDIPI (User Principal Name, or UPN) of that user in the role mapping. A user who presents a CAC with a valid certificate is mapped to a role and taken directly to the web console home page. The web console login page is not displayed.

Note: If you have logged on to Prism Central by using CAC authentication, to successfully log out of Prism Central, close the browser after you click **Log Out**.

Restoring Identity and Access Management Configuration Settings

Prism Central regularly backs up the Identity and Access Management (IAM) database, typically every 15 minutes. This procedure describes how to restore a specific time-stamped IAM backup instance.

About this task

You can restore authentication and authorization configuration settings available from the IAM database. For example, use this procedure to restore your authentication and authorization configuration to a previous state. You can choose an available time-stamped backup instance when you run the shell script in this procedure, and your authentication and authorization configuration is restored to the settings in the point-in-time backup.

Note:

- The script in this procedure is specifically designed for recreating and redeploying IAM components in Prism Central.
- The script is intended for performing complete restoration operations and not for inspecting backups.
- Before you begin the restoration operations, ensure that you have verified the timestamp of backups. Backups are generated every 15 minutes; select the most current backup to avoid IAM data loss.

Caution: Interrupting the script can lead to the unavailability of Prism Central UI due to the redeployment processes. Ensure that the script runs without interruption.

Tip: Contact Nutanix Support for assistance with script execution questions or issues.

Procedure

1. Log in to the Prism Central VM through an SSH session as the `nutanix` user.
2. Run the backup shell script `restore_iamv2.sh`

```
nutanix@pcvm$ sh /home/nutanix/cluster/bin/restore_iamv2.sh
```

The script displays a numbered list of available backups, including the backup file time-stamp.

```
Enter the Backup No. from the backup list (default is 1):
```

3. Select a backup by number to start the restore process.

The script displays a series of messages indicating restore progress, similar to:

```
You Selected the Backup No 1
Stopping the IAM serviceso
Waiting to stop all the IAM services and to start the restore process
Restore Process Started
```

```
Restore Process Completed
...
Restarting the IAM services
IAM Services Restarted Successfully
```

After the script runs successfully, the command shell prompt returns and your IAM configuration is restored.

4. To validate that your settings have been restored, do the following:
 - a. Log in to Prism Central as an administrator.
 - b. Select **Admin Center** in the Application Switcher.
 - c. Select **IAM** and verify the settings.

User Management

Managing Local User Accounts

About this task

The Prism Central admin user is created automatically, but you can add more (locally defined) users as needed. To add or update a user account, do the following:

Note:

- To add user accounts through Active Directory, see [Configuring Authentication](#). If you enable the Prism Self Service feature, an Active Directory is assigned as part of that process.
- No permissions are granted to the local users by default. To grant permissions to the directory users, you must specify roles for the users in that directory (see [Configuring an Authorization Policy](#) on page 30).
- Changing the Prism Central admin user password does not impact registration (re-registering clusters is not required).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to **Identities**.

4. To add a user account, click the **+ Add Local User** button and do the following in the displayed fields:
- First Name:** Enter a first name.
 - Last Name:** Enter a last name.
 - Username:** Enter a user name.
 - Email:** Enter a valid user email address.
 - Password:** Enter a password (maximum of 255 characters).

Note: A second field to verify the password is not included, so be sure to enter the password correctly in this field.

Note: Ensure that you meet the following password requirements:

- Contain at least eight characters
- Contain at least one lowercase character
- Contain at least one uppercase character
- Contain at least one digit
- Contain at least one special character
- Does not contain more than three consecutive characters

- Language:** Select the language setting for the user.

English is selected by default. You have an option to select **Simplified Chinese**, **Japanese** or **Korean**. If you select either of these, the cluster locale is updated for the new user. For example, if you select **Simplified Chinese**, the user interface is displayed in Simplified Chinese when the new user logs in.

- When all the fields are correct, click the **Save** button (lower right).

This saves the configuration and redisplay the dialog box with the new user appearing in the list.

Editing or Disabling a Local User Account

About this task

To update or disable a local user account, do the following:

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to **Identities**.
4. To modify a user account, select the user from the list of **Local Users** and click **Actions > Edit**.

5. To disable login access for a user account, select the user from the list of **Local Users** and click **Actions > Disable**.

Note: You cannot delete local users, however, you can restore access for a disabled user account by using the **Enable** option. This action will automatically restore all previously assigned permissions.

6. To change the password for a user account, select the user from the list of **Local Users** and click **Actions > Change Password**.
A window prompt appears to verify the action; click the **Save** button. The user account is removed and the user no longer appears in the list.

Updating My Account

About this task

To update your account credentials (that is, credentials for the user you are currently logged in as), do the following:

Procedure

1. To update your password, select **Change Password** from the user icon pull-down list of the main menu. The **Change Password** dialog box appears. Do the following in the indicated fields:
 - a. **Current Password:** Enter the current password.
 - b. **New Password:** Enter a new password.
 - c. **Confirm Password:** Re-enter the new password.
 - d. When the fields are correct, click the **Save** button (lower right). This saves the new password and closes the window.

Note: Password complexity requirements might appear above the fields; if they do, your new password must comply with these rules.

Note: By default, there is no password expiration day set for a local user account.

2. To update other details of your account, select **Update Profile** from the user icon pull-down list. The **Update Profile** dialog box appears. Do the following in the indicated fields for any parameters you want to change:
 - a. **First Name:** Enter a different first name.
 - b. **Last Name:** Enter a different last name.
 - c. **Email Address:** Enter a different valid user email address.
 - d. **Language:** Select a different language for your account from the pull-down list.
 - e. When all the fields are correct, click the **Save** button (lower right). This saves the changes and closes the window.

Local Account Passwords - Centralized Management

Password Manager Overview

The Password Management feature enables you to centrally manage all the local account passwords using Prism Central. The centralized management of passwords ensures enhanced account security by providing a direct view of the status of passwords (default or secure) and the ability to change the passwords of both individual accounts and the accounts that are grouped based on the cluster, controller VM, or Prism Central scope.

You can also utilize different viewing options for the local account passwords using the **Filters** option. Using filters, you can refine the view based on the following categories:

Status of the account password

- Default
- Secure
- No password

Cluster Name

Account Type

- admin
- Nutanix

Tip:

- You can change the account passwords in bulk.
- You can view the password change history using the [Audit Dashboard](#) or [Tasks](#) for traceability.

Password Manager Requirements

Supported Software Versions:

- Prism Central version pc.2023.4 or above
- AOS version 6.7.1 or above

Password Manager Limitations

- The account passwords for a maximum of 10 clusters can be changed simultaneously.
- Bulk password change for multiple user accounts on different clusters is not supported.

Changing Local Account Passwords

About this task

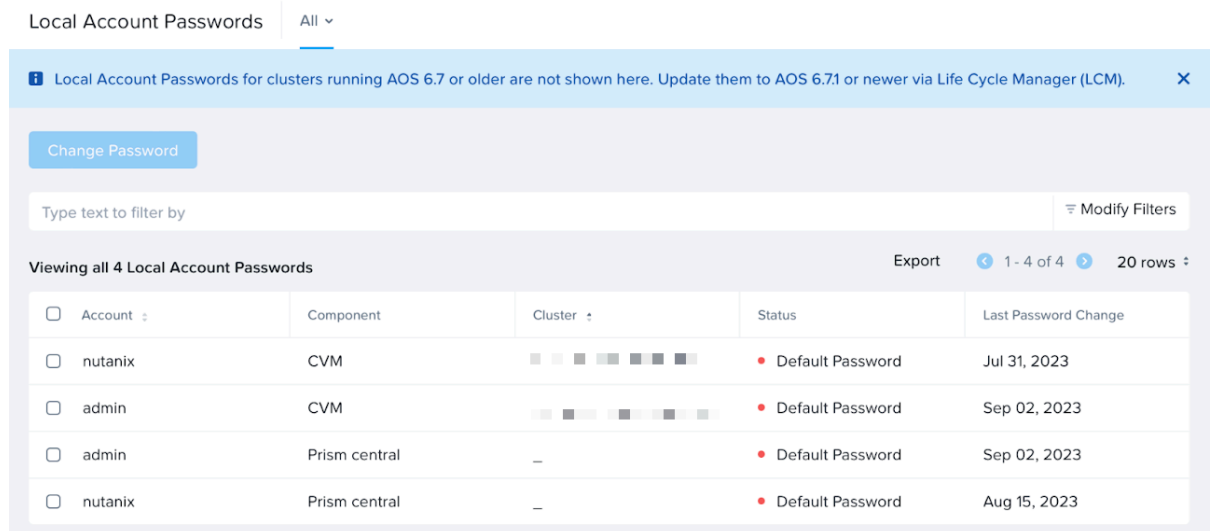
To change the password of one or multiple local accounts, do the following.

Procedure

1. Log on to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and go to **Network & Security > Local Account Passwords** from the **Navigation Bar**. Alternatively, go to the [Security Dashboard](#) and click **View all Local Account Passwords**.

The **Local Account Passwords** page opens. This page provides the list of accounts and the password status for all the accounts configured on the registered clusters.



The screenshot shows the 'Local Account Passwords' window. At the top, there's a header 'Local Account Passwords' with a dropdown menu set to 'All'. Below this is a blue banner with a message: 'Local Account Passwords for clusters running AOS 6.7 or older are not shown here. Update them to AOS 6.71 or newer via Life Cycle Manager (LCM).' A 'Change Password' button is visible. Below the banner is a search bar 'Type text to filter by' and a 'Modify Filters' button. The main content area shows 'Viewing all 4 Local Account Passwords' and an 'Export' button. A table lists the accounts with columns: Account, Component, Cluster, Status, and Last Password Change. The table contains 4 rows of data.

Account	Component	Cluster	Status	Last Password Change
<input type="checkbox"/> nutanix	CVM		Default Password	Jul 31, 2023
<input type="checkbox"/> admin	CVM		Default Password	Sep 02, 2023
<input type="checkbox"/> admin	Prism central	—	Default Password	Sep 02, 2023
<input type="checkbox"/> nutanix	Prism central	—	Default Password	Aug 15, 2023

Figure 3: Local Account Passwords Window

3. To change the password of an individual account, select the account from the **Account** column and click **Change Password**.
4. Optionally, you can select multiple accounts and change the passwords of the accounts in bulk. To view and select the required accounts, click **Modify Filters**, select the accounts, and then click **Change Password**.
5. In the **Change Password** window, enter the following and click **Change Password**.
 - Existing account password
 - New password
 - Confirm new password

The passwords are updated for the selected accounts.

Admin Account Password Retry Lockout

For enhanced security, Prism Central and Prism Element locks out the default 'admin' account for a period of 15 minutes after five unsuccessful login attempts. Once the account is locked out, the following message is displayed at the login screen.

Account locked due to too many failed attempts

You can attempt entering the password after the 15 minutes lockout period, or contact Nutanix Support in case you have forgotten your password.

Note: You cannot modify the default 15 minutes lock out period.

Controlling User Access (RBAC)

Prism Central supports role-based access control (RBAC) that you can configure to provide customized access permissions for users based on their assigned roles. The roles dashboard allows you to view information about all defined roles.

- Prism Central includes a set of predefined roles (see [Built-in Role Management](#) on page 25).
- You can also define additional custom roles (see [Custom Role Management](#) on page 26).
- After user authentication is configured, the users or user groups are not assigned the permissions by default. The required permissions must be explicitly assigned to users by creating an authorization policy (see [Configuring an Authorization Policy](#) on page 30).
- You can use Granular RBAC to assign fine-grained VM operation permissions to users based on your specific requirements.

Note: Defining custom roles and assigning roles are supported on AHV only.

Built-in Role Management

The following built-in roles are defined by default. You can see a more detailed list of permissions for any of the built-in roles through the details view for that role (see [Displaying Role Permissions](#) on page 29).

Role	Privileges
Prism Admin	Day-to-day admin of a Nutanix deployment. Manages the infrastructure and platform, but cannot entitle other users to be admins.
Prism Viewer	View-only admin of a Nutanix deployment. Has access to all infrastructure and platform features, but cannot make any changes.
Operator	Owner of team applications at runtime. Works on existing application deployments, exercises blueprint actions.
Developer	Application developer within a team. Authors blueprints, tests deployments, and publishes applications for other project members.
Self-Service Admin	Cloud admin for a Nutanix tenant. Manages virtual infrastructure, oversees self service, and can delegate end user management.
Super Admin	Highest-level admin with full infrastructure and tenant access. Manages a Nutanix deployment and can set up, configure, and make use of every feature in the platform.
Project Admin	Team lead to whom cloud administration gets delegated in the context of a project. Manages end users within the project and has full access to their entities.
Consumer	Lifecycle manager for team applications. Launches blueprints and controls their lifecycle and actions.
Network Infra Admin	Network infrastructure admin of a Nutanix deployment. Manages the infrastructure and underlay networking.
VPC Admin	VPC admin of a Nutanix deployment. Manages VPCs and related entities. Agnostic of the physical network infrastructure.
Flow Policy Author	Full Access to flow operations, except categories provisioning
Flow Admin	Full access to Flow operations, with categories provisioning

Role	Privileges
Flow Viewer	View access for Flow operations.
Monitoring Admin	Full access to perform all Monitoring operations
Monitoring Viewer	View access to all API in Monitoring
Storage Admin	Storage admin of a Nutanix deployment. This user can view and perform actions on Storage entities.
Action Service User	Basic Playbook access for all users
Category Viewer	View access for category object
Category Admin	Full access for category object
CSI System	Full access for Kubernetes cluster infrastructure resources for CSI
Kubernetes Data Services System	Full access for Kubernetes cluster infrastructure resources for Kubernetes Data Services
Kubernetes Infrastructure Provision	Access for Kubernetes cluster infrastructure VMs resources
Storage Viewer	View access for Storage entities.
Objects Admin	Full access to Object store operations.
Objects Editor	Edit access to Object store operations.
Objects Viewer	View access for Object store operations.
Files Admin	Full access to Files operations.
Files Viewer	View access for Files operations.
File Server Security Admin	All File Server security related permissions.
File Server Share Admin	Full access to File Server Share operations.
Disaster Recovery Admin	Full access to Disaster recovery operations.
Disaster Recovery Viewer	View access for Disaster recovery operations.
Cluster Admin	Full access to Cluster operations.
Cluster Viewer	View access for Cluster operations.
Virtual Machine Viewer	View access for Virtual Machines.
Virtual Machine Operator	Gives access for day-to-day activities on Virtual Machines.
Virtual Machine Admin	Full access to Virtual Machines.
Storage Admin	Full access to all storage entities' operations.
Storage Viewer	View-only access to all storage entities' operations.

Custom Role Management

If the built-in roles are not sufficient for your needs, you can create one or more custom roles (AHV only).

Creating a Custom Role

About this task

To create a custom role, do the following:

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to **Roles**
The **Roles** page appears.
4. In the **Roles** page, click **+ Create Role** and select one of the following options:
 - » **From existing role** - In the **Choose Role** window, type the first few characters of the existing built-in or custom role to search and select the role that you want to use as a template. Once you select the existing role, the list of operations associated with the role is displayed. You can configure the operations as shown in following steps. Click **Next**, and then click **Duplicate**.

Note: Any system operations in the existing role does not get added to the new role that you are create.

- » **New role**

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

Note: You must enter a unique name for the custom role; built-in role names are not allowed (including some Nutanix-internal role names).

6. **Description** (optional): Enter a description of the role.

7. Select the operations to grant access.

- All operations are listed by default. Similar operations are grouped in an expandable list. Click **Filter by** to view a subset of the operations list based on the **Entity Type** or **Operation** by entering a string in the search field.
 - Some operations are pre-selected if you are creating a role from an existing role.
- a. Click the "blue plus" icon to add individual operations or a group of operations. You can click



to view and select the related operations.

Note: To ensure that the role is granted sufficient permissions within the authorization policy, it is recommended to select all the related operations.

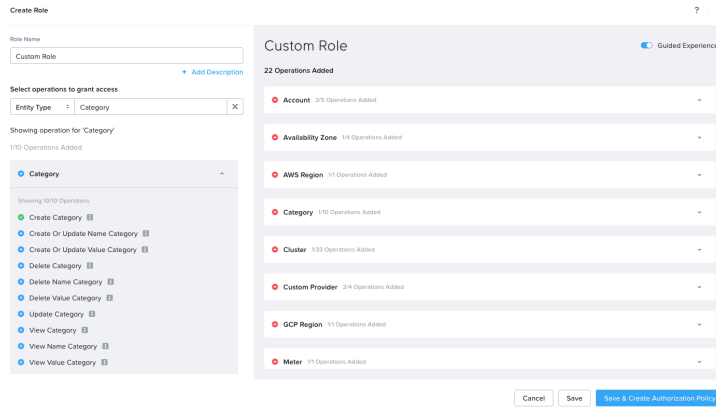


Figure 4: Create Custom Role

8. Choose one of the following actions:

- » Click **Save** to create the role. The page closes and the new role appears in the **Roles** view list.
- » Click **Save & Create Authorization Policy** to grant access to users based on the new role created in this procedure (see [Configuring an Authorization Policy](#) on page 30).

Modifying a Custom Role

About this task

Perform the following procedure to modify or delete a custom role.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to **Roles**.
The **Roles** page appears.
4. Do one of the following:
 - » To modify a custom role, select the role from the roles list and click **Update** from the **Actions** pull-down list. The **Update Roles** page for that role appears. Update the field values as desired and then click **Save**. See [Creating a Custom Role](#) on page 26 for field descriptions.
 - » To delete a custom role, select the role from the roles list and click **Delete** from the **Action** pull-down list. A confirmation message is displayed. Click **Delete** to delete and remove the role from the list.

Displaying Role Permissions

About this task

Do the following to display the privileges associated with a role.

Procedure

1. Log in to Prism Central as an administrator.

2. Select **Admin Center** in the Application Switcher.

3. Select **IAM** and go to **Roles**.

The page displays system defined and custom roles.

For example, if you click the **Consumer** role, the details page for that role appears, and you can view all the privileges associated with the **Consumer** role.

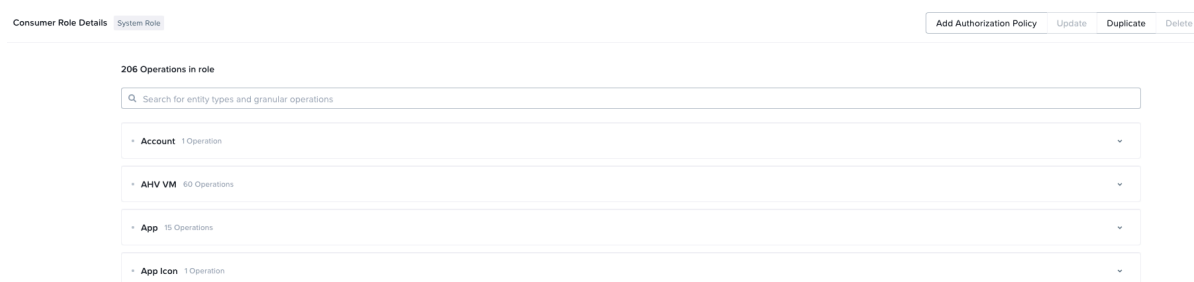


Figure 5: Role Summary Tab

4. Click **Add Authorization Policy** to add a new authorization policy to the role.

5. Click **Duplicate** tab to create a new role based on the selected role.

6. Click **Update** update the role permissions (applicable to custom roles only).

7. Click **Delete** to delete the selected role (applicable to custom roles only).

Authorization Policies

Authorization policy is an IAM mechanism to define access to Nutanix resources. You create an authorization policy to assign a role (built-in or custom) to an identity (user or user group) for a global or customized scope (of allowed actions). Authorization Policy ensures that the identity accessing a resource for a specific operation has the appropriate permissions.

After user authentication is configured (see [Configuring Authentication](#)), the users or user groups are not assigned the permissions by default. The required permissions must be explicitly assigned to users by creating an authorization policy.

Note: Default authorization policies are available for System Admin, Prism Admin and Prism Viewer roles. You can not modify or delete these policies.

To create an authorization policy, see [Configuring an Authorization Policy](#) on page 30.

Configuring an Authorization Policy

About this task

This procedure allows you to assign users to the built-in or custom roles and define the scope of access by creating an authorization policy. To configure an authorization policy, do the following:

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **IAM** and go to **Authorization Policies**.

4. To create an authorization policy, click the **+ Create New Authorization Policy** button.
The **Create New Authorization Policy** window appears.
5. In the **Choose Role** step, enter a role name by typing in the **Select the role you'd like to add to this policy** field, and click **Next**.
You can enter any built-in or custom roles.
The selected role and role details are displayed.
6. In the **Define Scope** step, select one of the following:
 - » **Full Access: all entity type & instances**
Selecting **Full Access** gives all added users access to all entity types in the associated role.
Optionally, you can select the **Automatically grant access to new entity types that are added to this role in the future** option. This allows assigned users to automatically gain access to any new entity types that are added to this role in the future.
 - » **Configure Access: select entity types & instances**
Selecting **Configure Access** provides you with the option to configure the entity types and instances for the added users in the associated role.
In **Entity Type** selection, select the entities on which this authorization policy will be applied. The list of available entities depends on the role selected in Step 5.
In **Filter** selection, refine your chosen entity type by selecting clusters, categories, or other available criteria to further specify access as necessary. The list of available filter options depends on the entity selected.
Repeat the entity type and filter selection for any combination of entity/filter that you want to define.
Optionally, you can select the **Allow users access to entities created by them** option. This option allows the assigned users to access any new entity instances created by them. For instance, if a user has permission to create a new virtual machine, they will also be able to access it.
7. Click **Next**.
8. In the **Assign Users** step, do the following:
 - From the pull-down menu, select **Local User** to add a local user or group to the policy. Search a user or group by typing the first few letters in the text field.
 - From the pull-down menu, select the available directory to add a directory user or group. Search a user or group by typing the first few letters in the text field.
9. Click **Save**.
The authorization policy configurations are saved, and the authorization policy is listed in the **Authorization Policies** window.

Editing, Duplicating, and Deleting an Authorization Policy

About this task

To edit, duplicate, or delete an authorization policy, do the following.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.

3. Select **IAM** and go to **Authorization Policies**.

4. To edit an authorization policy, select the policy from the list of available policies, and go to **Actions > Edit**. Edit the required configuration and click the **Save** button to update the changes. See [Configuring an Authorization Policy](#) on page 30 for information on authorization policy configuration.

Note: You can not edit system defined authorization policies.

5. To duplicate an authorization policy, select the policy from the list of available policies, and go to **Actions > Duplicate**. Edit the required configuration in the copy of the selected authorization policy and click the **Save** button to update the changes. See [Configuring an Authorization Policy](#) on page 30 for information on authorization policy configuration.

6. To delete an authorization policy, select the policy from the list of available policies, and go to **Actions > Delete**. Click **Delete** again to confirm.

Note:

- You can not delete system defined authorization policies.
- Any users associated with the authorization policy lose their assigned access after deleting the authorization policy.

SSL Certificate Management

Prism Central supports SSL certificate-based authentication for console access. To enable secure communication with a cluster, Prism Central includes a default self-signed SSL certificate. You can replace the default self-signed SSL certificate with your own self-signed SSL certificate or a certificate authority (CA) signed SSL certificate.

For production purposes, Nutanix recommends that you replace the default self-signed certificate with a CA signed SSL certificate.

Note:

- You can import only a cluster-wide SSL certificate in Prism Central. The SSL certificate cannot be customized for an individual controller VM (CVM).
- Nutanix recommends that you check for the validity of the certificate periodically and replace the certificate if it is invalid.

Importing an SSL Certificate

Nutanix simplifies the SSL certificate import process into Prism Central.

Before you begin

Depending upon your requirements, you need to either generate a self-signed SSL certificate or generate a Certificate Signing Request (CSR) for submission to a certificate Authority (CA) to get a CA signed certificate.

For more information, see

- [Generating a Self-signed SSL Certificate with Subject Alternative Name](#) on page 36
- [Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority \(CA\)](#) on page 39

About this task

The following procedure explains how to import a self-signed SSL certificate or a CA signed SSL certificate into Prism Central.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **Settings** and go to **SSL Certificate**.
4. To replace an SSL certificate, click **Replace Certificate**.
5. Do one of the following:
 - » To regenerate Nutanix default self-signed certificate, select **Regenerate Self Signed Certificate** and then click **Apply**.

A dialog box appears to verify the action; click **OK**. A new RSA 2048 bit self-signed certificate is generated and applied for Prism Central.
 - » To import self-signed SSL certificate or CA signed certificate, select **Import Key and Certificate** and then click **Next**.

6. To import self-signed SSL certificate or CA signed certificate files, do the following:

For self-signed certificate:

- **Private Key Type:** Select the appropriate private key type for the self-signed certificate from the dropdown list.
- **Private Key:** Click **Choose file** and select the private key.
- **Public Certificate:** Click **Choose file** and select the self-signed certificate corresponding to the private key.
- **CA Certificate/Chain:** Click **Choose file** select the self-signed certificate corresponding to the private key.

SSL Certificate ?

Note: Any change made to a certificate will restart the prism session and the user will be logged out automatically.

Guidelines for RSA: Please use a SHA-256, SHA-384, or SHA-512 signature algorithm with RSA certificates for security and performance.

Private Key Type

RSA 2048 bit

Private Key ⓘ

Choose file my_key_name.key

Public Certificate

Choose file my crt_name.crt

CA Certificate/Chain

Choose file my crt_name.crt

Cancel Import Files

Figure 6: Importing self-signed certificate

The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	my crt_name.crt

Certificate Components	File type
CA Certificate/Chain	my_crt_name.crt

For CA signed certificate:

- **Private Key Type:** Select the appropriate private key type for the CA signed certificate from the dropdown list.
- **Private Key:** Click **Choose file** and select the private key.
- **Public Certificate:** Click **Choose file** and select the CA signed public portion of the certificate corresponding to the private key.
- **CA Certificate/Chain:** Click **Choose file** and select the certificate or chain of the signing authority for the public certificate.

Note: To create a chain file from the list of CA certificates, see [Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority \(CA\)](#) on page 39.

SSL Certificate ?

Note: Any change made to a certificate will restart the prism session and the user will be logged out automatically.

Guidelines for RSA: Please use a SHA-256, SHA-384, or SHA-512 signature algorithm with RSA certificates for security and performance.

Private Key Type

RSA 2048 bit

Private Key ⓘ

Choose file my_key_name.key

Public Certificate

Choose file ca_signed_public_cert.cer

CA Certificate/Chain

Choose file ca_public_cert.crt

Cancel Import Files

Figure 7: Importing CA signed certificate

The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	ca_signed_public_cert.cer
CA Certificate/Chain	ca_public_cert.crt or ca_chain_certs.crt

7. To begin SSL certificate import, click **Import Files**.

Results

After generating or importing the new certificate, the interface gateway restarts. If the certificate and credentials are valid, the interface gateway uses the new certificate immediately, which means that your browser session (and all other open browser sessions) is invalid until you reload the page and accept the new certificate. If anything is wrong with the certificate (such as a corrupted file or wrong certificate type), the new certificate is discarded, and the system reverts to the original default certificate provided by Nutanix.

Note: The system holds only one custom SSL certificate. If a new certificate is uploaded, it replaces the existing certificate. The previous certificate is discarded.

Generating a Self-signed SSL Certificate with Subject Alternative Name

This task explains how to generate a self-signed SSL certificate using OpenSSL commands.

About this task

To comply with the security standards of NIST SP800-131a and the requirements of the RFC 6460 for NSA Suite B, the certificate import process validates that the correct signature algorithm is used for a given key and certificate pair. Use proper set of key types, sizes and curves, and signature algorithms. For more information, see [Recommended Key Configurations](#) on page 87.

Nutanix recommends including a DNS name for all controller VMs (CVMs) in the self-signed SSL certificate using the SAN extension. This scheme helps avoid SSL certificate errors when you access a CVM by direct DNS instead of the shared cluster IP address.

Procedure

1. Log in to any of the controller VMs (CVMs) with SSH using the management IP address of the CVM:

```
$ ssh admin@cvm_ip_address
```

2. To generate a private key with a bit length of your choice, run one of the following commands:

- RSA private key with a bit length of 2048:

```
admin@cvm$ openssl genrsa -out my_key_name.key 2048
```

- RSA private key with a bit length of 4096:

```
admin@cvm$ openssl genrsa -out my_key_name.key 4096
```

- ECDSA private key using the prime256v1 curve:

```
admin@cvm$ openssl ecparam -name prime256v1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp384r1 curve:

```
admin@cvm$ openssl ecparam -name secp384r1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp521r1 curve:

```
admin@cvm$ openssl ecparam -name secp521r1 -genkey -out my_key_name.pem
```

Important: While you are generating the private key, ensure that the private key is not password protected.

3. To generate the CSR of your choice, run one of the following commands:

- For RSA 2048 and RSA 4096 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -signature_algorithm -out my_csr_name.csr
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for RSA 2048 private key using SHA-256 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -sha256 -out my_csr_name.csr
```

- For ECDSA 256, ECDSA 384, ECDSA 521 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -signature_algorithm -out my_csr_name.csr
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for ECDSA 384 private key using SHA-384 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -sha384 -out my_csr_name.csr
```

4. Enter the information in the command output to incorporate into your certificate request:

```
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []:US
State or Province Name (full name) []:CA
Locality Name (eg, city) []:San Jose
Organization Name (eg, company) []: Nutanix Inc
```

```
Organizational Unit Name (eg, section) []: IT
Common Name (eg, fully qualified host name) []: ntx.com
Email Address []: myname@domain.com
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []: Enter your password
```

5. Create a configuration file in your home directory with your preferred text editor named `san.cnf` that contains the following text:

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
[v3_req]
basicConstraints = CA:FALSE
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
subjectAltName = @alt_names

[alt_names]
DNS.0 = example1.domain.com
DNS.1 = example2.domain.com
DNS.2 = example3.domain.com
DNS.3 = *.domain.com
IP.0 = x.x.x.x

[alt_names]
```

Specify your DNS and IP addresses. If you have a range of hosts, use wildcards (*) to match any subdomain of the domain name.

6. Generate a self-signed certificate:

```
admin@cvm$ openssl x509 -req -days number_of_days -in my_csr_name.csr -signkey
my_key_name.key -out my crt_name.crt -signature_algorithm -extensions v3_req -
extfile san.cnf
```

`number_of_days`: Specify the number of days until a newly generated certificate expires.

`signature_algorithm`: Specify sha256 or sha384 or sha512.

Example:

```
admin@cvm$ openssl x509 -req -days 1460 -in my_csr_name.csr -signkey my_key_name.key
-out my crt_name.crt -sha256 -extensions v3_req -extfile san.cnf
```

7. Copy `my_key_name.key` and `my crt_name.crt` from the CVM to your local machine:

```
admin@cvm$ scp my_key_name.key my crt_name.crt username@local-machine:/
local_file_path/
```

8. Log out of the CVM.

What to do next

After you successfully create a self-signed certificate with a private key, follow the procedure described in [Importing an SSL Certificate](#) on page 32 to replace the default certificate with your self-signed SSL certificate. The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Table 1: SSL Certificate Import Files

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	my_crt_name.crt
CA Certificate/Chain	my_crt_name.crt

Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority (CA)

This task explains how to generate a certificate signing request using OpenSSL commands.

About this task

To comply with the security standards of NIST SP800-131a and the requirements of the RFC 6460 for NSA Suite B, the certificate import process validates the correct signature algorithm is used for a given key and certificate pair. Use proper set of key types, sizes and curves, and signature algorithms. For more information, see [Recommended Key Configurations](#) on page 87.

Nutanix recommends including a DNS name for all CVMs in the CSR using the Subject Alternative Name (SAN) extension. This avoids SSL certificate errors when you access a CVM by direct DNS instead of the shared cluster IP address.

Procedure

1. Log in to any of the controller VMs (CVMs) with SSH using the management IP address of the CVM:

```
$ ssh admin@cvm_ip_address
```

2. Create a configuration file in your home directory with your preferred text editor named `ssl.cnf` that contains the following text:

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
countryName = Country Name (2 letter code)
stateOrProvinceName = State or Province Name (full name)
localityName = Locality Name (eg, city)
organizationName = Organization Name (eg, company)
organizationalUnitName = Organizational Unit Name (eg, BU)
commonName = Common Name (e.g. server FQDN or YOUR name)
emailAddress = Email Address

[v3_req]
subjectAltName = @alt_names

[alt_names]
DNS.0 = example1.domain.com
DNS.1 = example2.domain.com
DNS.2 = example3.domain.com
DNS.3 = *.domain.com
IP.0 = x.x.x.x
```

```
[alt_names]
```

Specify your DNS and IP addresses. If you have a range of hosts, use wildcards (*) to match any subdomain of the domain name.

3. To generate a private key with a bit length of your choice, run one of the following commands:

- RSA private key with a bit length of 2048:

```
admin@cvm$ openssl genrsa -out my_key_name.key 2048
```

- RSA private key with a bit length of 4096:

```
admin@cvm$ openssl genrsa -out my_key_name.key 4096
```

- ECDSA private key using the prime256v1 curve:

```
admin@cvm$ openssl ecparam -name prime256v1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp384r1 curve:

```
admin@cvm$ openssl ecparam -name secp384r1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp521r1 curve:

```
admin@cvm$ openssl ecparam -name secp521r1 -genkey -out my_key_name.pem
```

Important: While you are generating the private key, ensure that the private key is not password protected.

4. To generate the CSR of your choice, run one of the following commands:

- For RSA 2048 and RSA 4096 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -signature_algorithm -out my_csr_name.csr -config ssl.cnf
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for RSA 2048 private key using SHA-256 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -sha256 -out my_csr_name.csr -config ssl.cnf
```

- For ECDSA 256, ECDSA 384, ECDSA 521 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -signature_algorithm -out my_csr_name.csr -config ssl.cnf
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for ECDSA 384 private key using SHA-384 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -sha384 -out my_csr_name.csr -config ssl.cnf
```

5. Enter the information in the command output to incorporate into your certificate request:

```
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []: US
State or Province Name (full name) []: CA
Locality Name (eg, city) []: San Jose
Organization Name (eg, company) []: Nutanix Inc
```



```
Organizational Unit Name (eg, BU) []:IT
Common Name (e.g. server FQDN or YOUR name) []:ntx.com
Email Address []:myname@domain.com
```

6. Copy `my_key_name.key` and `my_crt_name.crt` from the CVM to your local machine:

```
nutanix@cvm$ scp my_key_name.key my_csr_name.csr username@local-machine:/
local_file_path/
```

7. Log out of the CVM.

8. Send your CSR file to the CA of your choice.

After receiving your CSR, the CA sends the following files:

- CA signed public certificate
- CA's public certificate
- Root CA public certificate (if the CA is intermediate)

The public certificate is validated by the issuing CA, and if the issuing CA is intermediate, the issuing CA certificate is validated by the root CA. A chain of certificates is validated to establish the trust.

9. Download all the certificate files received from CA to the local file directory.

10. (Optional) If the CA chain certificate provided by the certificate authority is not in a single file, run the following command to concatenate the list of CA certificates into a chain file:

```
$ cat intermediateCAcert.crt rootCAcert.crt > ca_chain_certs.crt
```

Note:

- The chain must start with the certificate of the signer and ends with the root CA certificate.
- Ensure that the chain file only has the root and intermediate certificates. If your chain file has public or private certificates, it will fail to import in Prism Central.

What to do next

Follow [Importing an SSL Certificate](#) on page 32 section to replace the default certificate with a CA signed certificate. The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Table 2: SSL Certificate Import Files

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	ca_signed_public_cert.cer
CA Certificate/Chain	ca_public_cert.crt or ca_chain_certs.crt

Verifying the Certificate Generation Request

Run the following commands to verify the certificate generation request.

- Verify that the generated certificate chain is OK:

```
admin@cvm$ openssl verify -CAfile ca_chain_certs.crt myPublicCert.cer
```

Example output:

```
myPublicCert.cer: OK
```

- Verify the private key and signature algorithm details:

```
admin@cvm$ openssl x509 -in my_cert_name.crt -text -noout | grep -i 'rsa\|ecdsa\|Public'
```

Example output:

```
Signature Algorithm: ecdsa-with-SHA512
      Subject Public Key Info:
        Public Key Algorithm: id-ecPublicKey
          Public-Key: (521 bit)
        Signature Algorithm: ecdsa-with-SHA512
```

- Verify that the CA certificate chain uses SHA 256 as a signature algorithm:

```
admin@cvm$ openssl crl2pkcs7 -nocrl -certfile ca_chain_certs.crt | openssl pkcs7 -print_certs -noout -text | grep -Ew '(Subject|Issuer|Signature Algorithm):' | grep -C1 Issuer
```

Troubleshooting the Certificate Generation Request

The following troubleshooting tips can help you resolve common issues that can occur when generating certificates.

Chain certificate format

If your chain certificate file has public or private certificates, it will fail to import in Prism Central. Ensure that the chain certificate file only has the root and intermediate certificates.

For example, if a public certificate is present in a chain file, you can remove it by opening your chain file in your preferred text editor. Ensure that there are no extra white spaces at the bottom of the file.

DER-encoded certificate issue

If the certificate is DER encoded, it fails to import in Prism Central. You can resolve the issue by converting it to PEM-encoded ASCII format.

- Ensure that the certificate is DER encoded:

```
admin@cvm$ openssl x509 -in cert.crt -inform der -text -noout
```

- If the certificate is DER encoded, run the following command to convert the certificate from DER to PEM-encoded ASCII format:

```
admin@cvm$ openssl x509 -in certDER.crt -inform der -outform pem -out cert.crt
```

Certificate format

Ensure that all the certificates do not have any extra data (or custom attributes) before the beginning (`(-----BEGIN CERTIFICATE-----)`) or after the end (`(-----END CERTIFICATE-----)`) of the block.

Controlling Remote (SSH) Access

About this task

Nutanix supports key-based SSH access to Prism Central. Enabling key-based SSH access ensures that password authentication is disabled and only the keys you have provided can be used to access the Prism Central (only for nutanix/admin users). Thus making the Prism Central more secure.

Nutanix supports the following key-based SSH encryption algorithms:

- AES128-CTR
- AES192-CTR
- AES256-CTR

You can create a key pair (or multiple key pairs) and add the public keys to enable key-based SSH access. However, when site security requirements do not allow such access, you can remove all public keys to prevent SSH access.

Note: When you add a public key, SSH access is enabled for both the "nutanix" and "admin" user accounts on the controller VM and the AHV host.

Caution: You should use the SSH-based command-line interface only when directed by Nutanix support or as specified in Nutanix documentation.

Tip: In addition to configuring the SSH Security Level, you can also consider cluster lockdown to disable password-based SSH authentication by adding SSH keys, see [Controlling Cluster Access](#) on page 71.

To control key-based SSH access to Prism Central, do the following:

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Select **Settings** and go to **Cluster Lockdown**.

The **Cluster Lockdown** dialog box appears. Enabled public keys (if any) are listed in this window.

4. To disable (or enable) remote login access, uncheck (check) the **Enable Remote Login with Password** box.

Remote login access is enabled by default.

5. To add a new public key, click the **New Public Key** button and then do the following in the displayed fields:

- a. **Name:** Enter a key name.
- b. **Key:** Enter (paste) the key value into the field.
- c. Click the **Save** button (lower right) to save the key and return to the main **Cluster Lockdown** window.

There are no public keys available by default, but you can add any number of public keys.

6. To delete a public key, click the **X** on the right of that key line.

Note: Deleting all the public keys and disabling remote login access locks down the cluster from SSH access.

Security Policies using Flow

Nutanix Flow includes a policy-driven security framework that inspects traffic within the data center. For more information, see the [Flow Microsegmentation Guide](#).

Data-in-Transit Encryption

Data-in-Transit Encryption allows you to encrypt service level traffic between the cluster nodes. Data-in-Transit Encryption, along with [Data-at-Rest Encryption](#) on page 96, protects the entire life cycle of data and is an essential countermeasure for unauthorized access of critical data.

To enable Data-in-Transit Encryption, see [Enabling Data-in-Transit Encryption](#) on page 44.

Note:

- Data-in-Transit Encryption can have an impact on I/O latency and CPU performance.
- Intra-cluster traffic encryption is supported only for the Stargate service.
- RDMA traffic encryption is not supported.
- When a Controller VM goes down, the traffic from guest VM to remote Controller VM is not encrypted.
- Traffic between guest VMs connected to Volume Groups is not encrypted when the target disk is on a remote Controller VM.

Enabling Data-in-Transit Encryption

About this task

Data-in-Transit Encryption allows you to encrypt service level traffic between the cluster nodes. To enable Data-in-Transit Encryption, do the following.

Before you begin

1. Ensure that the cluster is running AOS version 6.1.1 and Prism Central version pc.2022.4.
2. Ensure that you allow port 2009, which is used for Data-in-Transit Encryption.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Infrastructure** in the Application Switcher.

3. Go to **Hardware > Clusters > Actions** and select **Enable Data-In-Transit Encryption**.

The following confirmation dialog box is displayed.

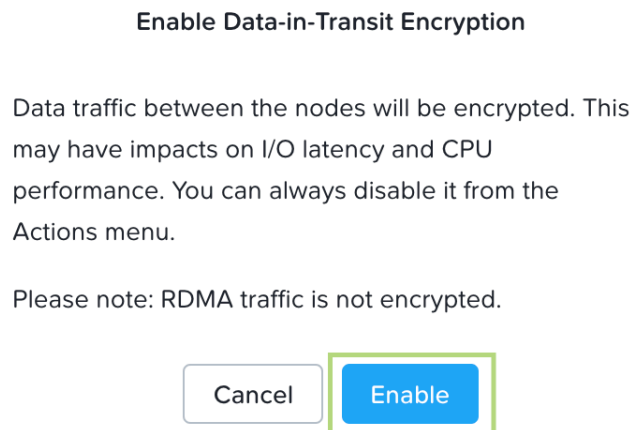


Figure 8: Enable Data-in-Transit Encryption

4. Click **Enable** to confirm.

What to do next

You can disable Data-in-Transit Encryption after you have enabled it. To disable Data-in-Transit Encryption, see [Disabling Data-in-Transit Encryption](#) on page 45.

Disabling Data-in-Transit Encryption

About this task

You can disable Data-in-Transit Encryption after you have enabled it. To disable Data-in-Transit Encryption, do the following.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Infrastructure** in the Application Switcher.

3. Go to **Hardware > Clusters > Actions** and select **Disable Data-In-Transit Encryption**.

The following confirmation dialog box is displayed.

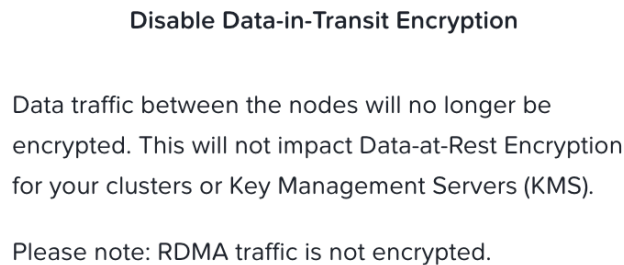


Figure 9: Disable Data-in-Transit Encryption

4. Click **Disable** to confirm.

Securing AHV VMs with Virtual Trusted Platform Module

Overview

A Trusted Platform Module (TPM) is used to manage cryptographic keys for security services like encryption and hardware (and software) integrity protection. AHV virtual trusted platform Module (vTPM) is software-based emulation of the TPM 2.0 specification that works as a virtual device.

Note:

- You can enable vTPM using Prism Central UI or aCLI. To enable vTPM using aCLI, see [Securing AHV VMs with Virtual Trusted Platform Module \(aCLI\)](#) in the *AHV Administration Guide*.
- AHV vTPM does NOT require OR use a hardware TPM.

You can use the AHV vTPM feature to secure virtual machines running on AHV.

vTPM Use Cases

AHV vTPM provides virtualization-based security support for the following primary use cases.

- Support for storing cryptographic keys and certificates for Microsoft Windows BitLocker
- TPM protection for storing VBS encryption keys for Windows Defender Credential Guard

See *Microsoft documentation* for details on Microsoft Windows Defender Credential Guard and Microsoft Windows BitLocker.

Tip: Windows 11 installation requires both TPM 2.0 and secure boot enabled for the guest VM. For more information, see Microsoft website for Windows 11 specs, features, and computer requirements.

For information on how to create or update a guest VM with secure boot enabled, see [Creating a VM through Prism Central \(AHV\)](#) and [Managing a VM through Prism Central \(AHV\)](#) sections in *Prism Central Infrastructure Guide*.

Considerations for Enabling vTPM in AHV VMs

Consider the following requirements and limitations when configuring vTPM using Prism Central.

Requirements

Supported Software Versions:

- Prism Central version pc.2022.9 or above
- AHV version 20220304.242 or above
- AOS version 6.5.1 or above

VM Requirements:

- You must enable UEFI on the VM on which you want to enable vTPM, see [UEFI Support for VM](#).
- You must enable Secure Boot (applicable if using Microsoft Windows BitLocker). To enable Secure Boot, see [Creating/Updating a VM with Secure Boot Enabled](#).

Limitations

- All [Secure Boot limitations](#) apply to vTPM VM.
- [Disaster Recovery limitations](#) apply when protecting vTPM-enabled VMs.

Creating a VM with vTPM

About this task

You can create a virtual machine with the vTPM configuration enabled using the following procedure.

Procedure

1. Log on to Prism Central as an administrator.
2. Select **Infrastructure** in the Application Switcher.
3. Go to **Compute & Storage > VMs** and click **List** tab of the VMs dashboard (see [VM Summary View](#)).
4. Click the **Create VM** button.
The **Create VM** wizard appears. Follow the instructions in the [Creating a VM](#) topic for details on the **Create VM** wizard.

5. At the **Shield VM Security Settings**, click the **Attach vTPM** check box to enable vTPM for the selected VM.

Note: **Shield VM Security Settings** is available only if you have enabled **UEFI BIOS Mode** for **Boot Configuration**.




Figure 10: Attach vTPM

6. Click **Next** at the subsequent VM setting tabs and then click **Save**.

What to do next

You can now start the VM to verify if the vTPM configuration is applied on the VM.

Enabling vTPM on an Existing VM

About this task

You can update the settings of an existing virtual machine to enable vTPM using the following procedure.

Procedure

1. Log on to Prism Central as an administrator.
2. Select **Infrastructure** in the Application Switcher.
3. Go to **Compute & Storage > VMs** to view the VMs dashboard (see [VM Summary View](#)).
4. You can choose to manage or update an existing VM configuration using any of the following methods, see [Managing a VM \(AHV\)](#) for details.
 - Select the target VM in the **List** tab of the VMs dashboard (see [VMs Summary View](#)) and choose the required action from the **Actions** menu.
 - Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
 - Go to the details page of a selected VM (see [VM Details View](#)) and select the desired action.
5. VM must be powered off before you can update the **Shield VM Security Settings** configuration. To power off the VM, select **More** and then **Power Off**.
6. To modify the VM configuration, select **Update**.

7. Go to **Shield VM Security Settings** and click the **Attach vTPM** check box to enable vTPM for the selected VM.

Note:

- **Shield VM Security Settings** is available only if you have enabled **UEFI BIOS Mode for Boot Configuration**
- Shield VM settings cannot be selected when the VM is running. Shut down the VM to update these settings.




Figure 11: Attach vTPM

8. Click **Next** at the subsequent VM setting tabs and then click **Save**.

What to do next

You can now start the VM to verify if the vTPM configuration is applied on the VM.

Removing vTPM from an Existing VM

About this task

You can update the settings of an existing virtual machine to remove the previously enabled vTPM setting using the following procedure.

Procedure

1. Log on to Prism Central as an administrator.
2. Select **Infrastructure** in the Application Switcher.
3. Go to **Compute & Storage > VMs** to view the VMs dashboard (see [VM Summary View](#)).
4. You can choose to manage or update an existing VM configuration using any of the following methods, see [Managing a VM \(AHV\)](#) for details.
 - Select the target VM in the **List** tab of the VMs dashboard (see [VMs Summary View](#)) and choose the required action from the **Actions** menu.
 - Right-click on the target VM in the **List** tab of the VMs dashboard and select the required action from the drop-down list.
 - Go to the details page of a selected VM (see [VM Details View](#)) and select the desired action.
5. VM must be shut down before you can update the **Shield VM Security Settings** configuration. To shut down the VM, select **More** and then **Power Off**.
6. To modify the VM configuration, select **Update**.

- Go to **Shield VM Security Settings** and click (to uncheck) the **Attach vTPM** check box to disable vTPM for the selected VM.

Warning: Disabling vTPM may severely affect VM functionality or result in data loss. For example, if Microsoft Windows BitLocker key is stored in vTPM, then you will require the recovery key to unlock the encrypted disk.

Note: Shield VM settings cannot be selected when the VM is running. Shut down the VM to update these settings.

- Click **Next** at the subsequent VM setting tabs and then click **Save**.

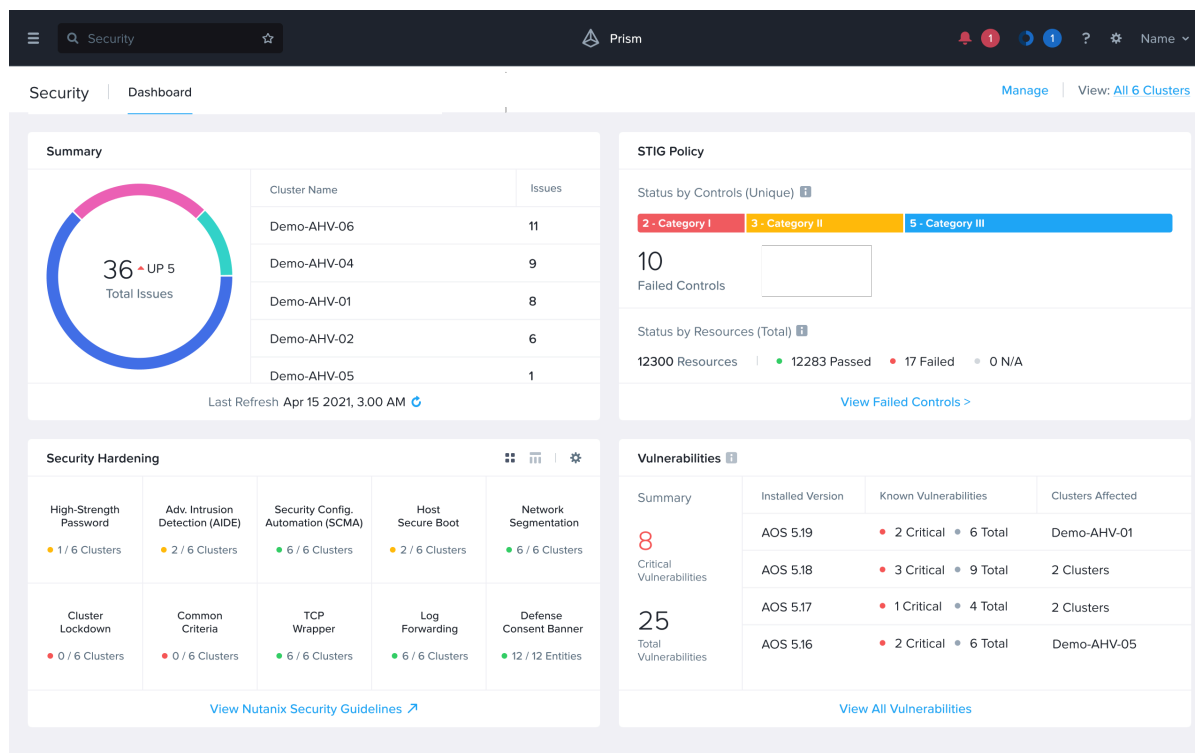
Security Dashboard

This topic provides an overview of the Security Dashboard.

Overview

The Security Dashboard provides dynamic summary of the security posture across all registered clusters. The Security Dashboard allows you to view the most critical security parameters like cluster-based issue summary, STIG policy compliance, security hardening, identified vulnerabilities, and the local account password status. The security dashboard is divided into multiple widgets to represent different security focus areas.

You can customize the security dashboard based on your preference. See [Managing Security Dashboard](#) on page 56 to customize the security dashboard. A sample view of the Security Dashboard with the default widgets is displayed in the figure.



Requirements

Ensure that meet the following requirements to use Security Dashboard.

- Enable Microservices infrastructure on Prism Central, see [Enabling Microservices Infrastructure Manually](#) in the *Prism Central Infrastructure Guide*.
- Ensure that your cluster is running any software edition of AOS version 6.6 or later and Prism Central version 2022.9 or later. Ensure that you are also running a compatible version of the AHV or ESXi hypervisor in your cluster. For software compatibility details, see the [Compatibility and Interoperability Matrix](#) at the Nutanix Support portal.
- Security Dashboard is available only for the users with the **Prism Admin** role. You must login with an account that has Prism Admin role assignment to view and manage the Security Dashboard.

Note: The Security Dashboard feature is not supported for Prism Central instance deployed on Nutanix Cloud Clusters (NC2).

Upgrading Security Dashboard

About this task

Life Cycle Manager (LCM) allows you to upgrade security dashboard by upgrading the "PC Core Services" module. To upgrade security dashboard manually using LCM, do the following:

Procedure

1. Log on to Prism Central.
2. Open the drop-down menu on the upper left and select **Administration > LCM**.
3. Perform inventory using the procedure shown in [LCM Inventory](#).
4. From the list of softwares available for upgrade, select **PC Core Services** and click **Upgrade**.

Note:

- Manual upgrade of security dashboard (PC Core Services module) is applicable to the following software versions:
 - Prism Central version pc.2022.9 and later
 - Life Cycle Manager version 2.6 and later
- Not every upgrade of the PC Core Services module includes updates to the security dashboard. Some upgrades may only include internal service upgrades.

Security Widget (PC Main Dashboard)

This topic provides the information on the **Security** information tile or widget available in the [Prism Central Main Dashboard](#).

The Security widget displays the total number of security issues in your clusters. The issue count is also available based on categories like security hardening, STIG issues, and vulnerabilities. Click **View All Issues** to go to the [Security Dashboard](#) that contains detailed information on the security posture of all registered clusters.

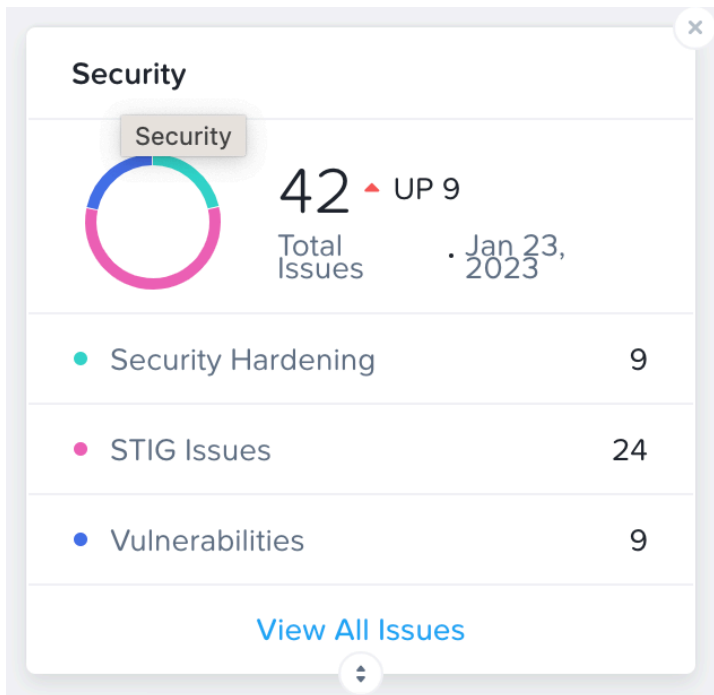


Figure 12: Security Dashboard Widget

Security Dashboard Wizard

The Security Dashboard wizard allows you to take a guided tour and navigate through the various tasks and workflows in the **Security Dashboard**.

The Security Wizard is automatically presented as a dialog-box when you access the dashboard feature for the first time. Click **Start Tour** to begin the dashboard walkthrough.

Optionally, click **Skip for Now** to go the dashboard directly. To access this tour later, click the help menu icon ("?"), expand **New in Prism Central**, and select **Security Dashboard**.



Welcome to Security Dashboard!

Let's walk you through the dashboard features, navigation and upkeep.

Skip for Now

Start Tour

Figure 13: Security Dashboard Wizard

Using the Security Dashboard

Security Dashboard is a customizable and dynamic console to monitor the security posture of your Nutanix infrastructure in a single view.

Accessing the Security Dashboard

You can view the security dashboard in Prism Central using any of the following methods.

- Go to **Prism Central Main Dashboard** > **Security widget** and click **View All Issues**.
- Click the hamburger icon, go to **Network & Security** > **Security Dashboard**.

Note: You must login with an account that has Prism Admin role assignment to view the Security Dashboard.

Tip: The Security Dashboard Wizard is automatically presented as a pop-up when you access the dashboard feature for the first time. Click **Start Tour** to begin the dashboard walkthrough. Optionally, click **Skip for Now** to go the dashboard directly.

Navigating the Security Dashboard

The security dashboard displays the following information tiles (widgets) by default.

- [Summary](#) on page 54
- [STIG Policy](#) on page 54
- [Security Hardening](#) on page 55
- [Vulnerabilities](#) on page 55
- [Local Account Passwords](#) on page 56

Tip: You can change the view of the security dashboard based on the following options.

- All clusters
- Individual cluster
- Selection of clusters

The **View** selection menu on the top-right corner of the dashboard allows you to switch between the different views. The values displayed in the widgets are dynamically updated based your selection.

Summary

The **Summary** widget allows you to view your open security issues or focus on the clusters that have the most number of security issues. You can click the Summary pie graph to view the following information.

- Total number of issues in the clusters
- Number of issues in the clusters categorized based on the following issue categories:
 - Security Hardening
 - STIG Issues
 - Vulnerabilities
 - Local Account Passwords

Tip: The Security Dashboard refreshes once daily. Updating the STIG check based vulnerabilities information requires a manual refresh. You can initiate a manual refresh by clicking the refresh icon at the bottom of the widget. Manual refresh process takes some approximately 20 minutes to 2 hours to complete and depends on the number of clusters in your environment.

STIG Policy

STIG helps detect deviation from the security baseline configuration of the operating system and hypervisor to remain in compliance. Nutanix has implemented the Controller VM and Prism Central VM to

support STIG compliance with the RHEL 8 STIG as published by DISA. For more information, see [Security Dashboard STIG Guidance Reference](#).

The **STIG Policy** widget provides an accurate snapshot of policy violations or deviations (resulting in failure) from the baseline STIG policy. This widget displays the STIG policy compliance status according to controls and resources.

Note: The STIG checks are run only on the node running primary Prism Element.

Status by Controls (Unique)	Number of unique STIG controls that are not met.
Status by Resources (Total)	Total number of individual resources that have failed a set of STIG controls.

Click **View Failed Controls** to view the list of STIG controls that are not met.

STIG Policy - Failed Controls

The STIG Policy -Failed Controls dialog box allows you to view STIG controls that do not meet the required settings.

Security Hardening

The **Security Hardening** widget displays the status of security hardening controls applied on your clusters. This widget also allows you to configure multiple security hardening controls from the widget directly.

You can configure the following Security Hardening configurations from the security hardening widget.

Hardening Parameter	Description
High-Strength Password	Enable high strength password for clusters. Note: Ensure that you update the Controller VM and AHV passwords after enabling this setting, see Resetting Password (CVM) and Changing Admin User Password (AHV) .
Advanced Intrusion Detection Environment (AIDE)	Enable AIDE.
Security Configuration Management Automation (SCMA)	Enable SCMA frequency for AHV hosts and Controller VM. Note: Higher frequency SCMA scan can decrease performance on the respective cluster.
Cluster Lockdown	Enable cluster lockdown mode.
Defense Consent Banner	Enable defense consent banner for AHV hosts and Controller VM.

Refer to the [Nutanix Security Guide](#) to enable other hardening settings using Prism or nCLI.

Vulnerabilities

The **Vulnerabilities** widget displays a list of vulnerabilities (or CVEs) associated with your clusters based on the Acropolis Operating System versions.

Click **View All Vulnerabilities** to view the list of all vulnerabilities and the recommended upgrade path for mitigating the identified vulnerabilities.

Note: Not all the listed CVEs might be resolved when upgraded to a new release.

Local Account Passwords

The **Local Account Passwords** widget displays a list of local accounts and the associated password state (default or secure).

You can view the local accounts passwords **By Component** or **By Cluster**.

Click **View all Local Account Passwords** to change (or bulk change) the account passwords using the [Local Account Passwords](#) feature.

Managing Security Dashboard

About this task

To manage the security dashboard, do the following.

Procedure

1. Access the Security Dashboard, see [Accessing the Security Dashboard](#).
2. Click the **Manage Dashboard** hyperlink available at the top-right corner of the **Security Dashboard**.
3. Click **Refresh Dashboard** to refresh the security dashboard manually.

Note: Security Dashboard refreshes once daily. This process takes some time to complete.

4. View **Security Dashboard Version**.

Ensure that your Security Dashboard version is updated in Life Cycle Manager (LCM) to ensure security details are correct for all your clusters.

Manage Dashboard

Dashboard Refresh

Security Dashboard refreshes once daily. You can manually refresh below if you know changes have occurred. Please note, this process can take some time to complete.

Refresh Dashboard

Last Refresh Oct 20 2022, 5:30 AM

Security Dashboard Version List

1 - 1 of 120 rows

Clusters	Dashboard Version	Release Date	Status
cluster_	security_aos.2022.9	1 Sep 2022	Up-to-date

Please check for updates in [Life Cycle Manager \(LCM\)](#) if any of your clusters are outdated.

Close

Figure 14: Manage Security Dashboard

SECURITY MANAGEMENT USING PRISM ELEMENT (PE)

Nutanix provides several mechanisms to maintain security in a cluster using Prism Element.

Configuring Authentication

About this task

Nutanix supports user authentication. To configure authentication types and directories and to enable client authentication or to enable client authentication only, do the following:

Caution: The web console (and nCLI) does not allow the use of the not secure SSLv2 and SSLv3 ciphers. There is a possibility of an SSL Fallback situation in some browsers which denies access to the web console. To eliminate this, disable (uncheck) SSLv2 and SSLv3 in any browser used for access. However, TLS must be enabled (checked).

Procedure

1. Click the gear icon in the main menu and then select **Authentication** in the **Settings** page. The **Authentication Configuration** window appears.

Note: The following steps combine three distinct procedures, enabling authentication (step 2), configuring one or more directories for LDAP/S authentication (steps 3-5), and enabling client authentication (step 6). Perform the steps for the procedures you need. For example, perform step 6 only if you intend to enforce client authentication.

2. To enable server authentication, click the **Authentication Types** tab and then check the box for either **Local** or **Directory Service** (or both). After selecting the authentication types, click the **Save** button.
The **Local** setting uses the local authentication provided by Nutanix (see [User Management](#) on page 74). This method is employed when a user enters just a login name without specifying a domain (for example, `user1` instead of `user1@nutanix.com`). The **Directory Service** setting validates `user@domain` entries and validates against the directory specified in the **Directory List** tab. Therefore, you need to configure an authentication directory if you select **Directory Service** in this field.

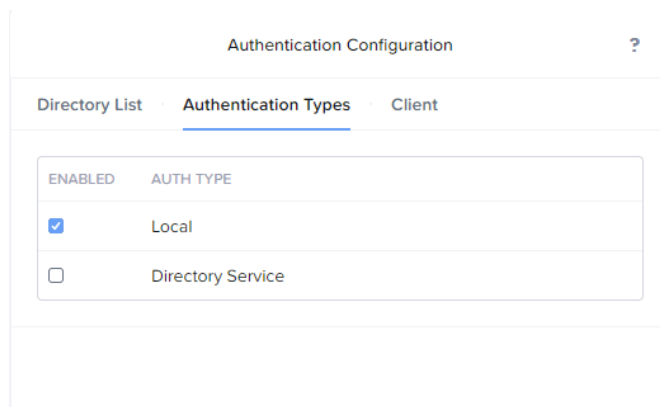


Figure 15: Authentication Types Tab

Note: The Nutanix admin user can log on to the management interfaces, including the web console, even if the **Local** authentication type is disabled.

3. To add an authentication directory, click the **Directory List** tab and then click the **New Directory** option.

A set of fields is displayed. Do the following in the indicated fields:

- a. **Directory Type:** Select one of the following from the pull-down list.

- **Active Directory:** Active Directory (AD) is a directory service implemented by Microsoft for Windows domain networks.

Note:

- Users with the "User must change password at next logon" attribute enabled will not be able to authenticate to the web console (or nCLI). Ensure users with this attribute first login to a domain workstation and change their password prior to accessing the web console. Also, if SSL is enabled on the Active Directory server, make sure that Nutanix has access to that port (open in firewall).
- An Active Directory user name or group name containing spaces is not supported for Prism Element authentication.
- Active Directory domain created by using non-ASCII text may not be supported. For more information about usage of ASCII or non-ASCII text in Active Directory configuration, see the [Internationalization \(i18n\)](#) on page 73 section.
- Use of the "Protected Users" group is currently unsupported for Prism authentication. For more details on the "Protected Users" group, see "Guidance about how to configure protected accounts" on Microsoft documentation website.
- The Microsoft AD is LDAP v2 and LDAP v3 compliant.
- The Microsoft AD servers supported are Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019.

- **OpenLDAP:** OpenLDAP is a free, open source directory service, which uses the Lightweight Directory Access Protocol (LDAP), developed by the OpenLDAP project. Nutanix currently supports the OpenLDAP 2.4 release running on CentOS distributions only.

- b. **Name:** Enter a directory name.

This is a name you choose to identify this entry; it need not be the name of an actual directory.

- c. **Domain:** Enter the domain name.

Enter the domain name in DNS format, for example, nutanix.com.

- d. **Directory URL:** Enter the URL address to the directory.

The URL format is as follows for an LDAP entry: `ldap://host:ldap_port_num`. The host value is either the IP address or fully qualified domain name. (In some environments, a simple domain name is sufficient.) The default LDAP port number is 389. Nutanix also supports LDAPS (port 636)

and LDAP/S Global Catalog (ports 3268 and 3269). The following are example configurations appropriate for each port option:

Note: LDAPS support does not require custom certificates or certificate trust import.

- Port 389 (LDAP). Use this port number (in the following URL form) when the configuration is single domain, single forest, and not using SSL.

```
ldap://ad_server.mycompany.com:389
```

- Port 636 (LDAPS). Use this port number (in the following URL form) when the configuration is single domain, single forest, and using SSL. This requires all Active Directory Domain Controllers have properly installed SSL certificates.

```
ldaps://ad_server.mycompany.com:636
```

Note: The LDAP server SSL certificate must include a Subject Alternative Name (SAN) that matches the URL provided during the LDAPS setup.

- Port 3268 (LDAP - GC). Use this port number when the configuration is multiple domain, single forest, and not using SSL.
- Port 3269 (LDAPS - GC). Use this port number when the configuration is multiple domain, single forest, and using SSL.

Note: When constructing your LDAP/S URL to use a Global Catalog server, ensure that the Domain Control IP address or name being used is a global catalog server within the domain being configured. If not, queries over 3268/3269 may fail.

Note: When querying the global catalog, the users **sAMAccountName** field must be unique across the AD forest. If the **sAMAccountName** field is not unique across the subdomains, authentication may fail intermittently or consistently.

Note: For the complete list of required ports, see [Port Reference](#).

- e. (OpenLDAP only) Configure the following additional fields:
 1. **User Object Class:** Enter the value that uniquely identifies the object class of a user.
 2. **User Search Base:** Enter the base domain name in which the users are configured.
 3. **Username Attribute:** Enter the attribute to uniquely identify a user.
 4. **Group Object Class:** Enter the value that uniquely identifies the object class of a group.
 5. **Group Search Base:** Enter the base domain name in which the groups are configured.
 6. **Group Member Attribute:** Enter the attribute that identifies users in a group.
 7. **Group Member Attribute Value:** Enter the attribute that identifies the users provided as value for **Group Member Attribute**.
- f. **Search Type.** How to search your directory when authenticating. Choose **Non Recursive** if you experience slow directory logon performance. For this option, ensure that users listed in Role

Mapping are listed flatly in the group (that is, not nested). Otherwise, choose the default **Recursive** option.

- g. **Service Account Username:** Enter the service account user name in the `user_name@domain.com` format that you want the web console to use to log in to the Active Directory.

A service account is created to run only a particular service or application with the credentials specified for the account. According to the requirement of the service or application, the administrator can limit access to the service account.

A service account is under the Managed Service Accounts in the Active Directory server. An application or service uses the service account to interact with the operating system. Enter your Active Directory service account credentials in this (username) and the following (password) field.

Note: Be sure to update the service account credentials here whenever the service account password changes or when a different service account is used.

- h. **Service Account Password:** Enter the service account password.
- i. When all the fields are correct, click the **Save** button (lower right). This saves the configuration and redisplay the Authentication Configuration dialog box. The configured directory now appears in the **Directory List** tab.
- j. Repeat this step for each authentication directory you want to add.

Note:

- The Controller VMs need access to the Active Directory server, so open the standard Active Directory ports to each Controller VM in the cluster (and the virtual IP if one is configured).
- No permissions are granted to the directory users by default. To grant permissions to the directory users, you must specify roles for the users in that directory (see [Assigning Role Permissions](#) on page 67).
- Service account for both Active directory and openLDAP must have full read permission on the directory service. Additionally, for successful Prism Element authentication, the users must also have search or read privileges.

Authentication Configuration ?

Directory List · Authentication Types · Client

Directory Type
Active Directory

Name
Name

Domain
e.g., eng.company.com

Directory URL
e.g., ldap://10.1.4.111:389

Search Type ?
Non Recursive (Default)


Service Account
We need to query your company's active directory for details of users. Ideally a service account with no time limit should be used.

Service Account Username
Username

Service Account Password
Password

< Back Save

Figure 16: Directory List Tab

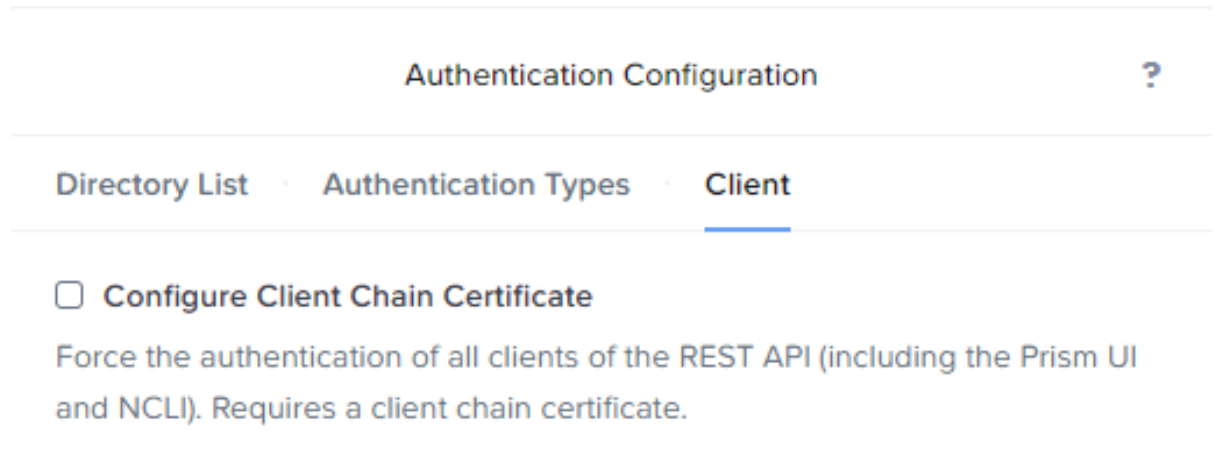
4. To edit a directory entry, click the **Directory List** tab and then click the pencil icon  for that entry. After clicking the pencil icon, the Directory List fields reappear (see step 3). Enter the new information in the appropriate fields and then click the **Save** button.
5. To delete a directory entry, click the **Directory List** tab and then click the X icon for that entry. After clicking the X icon, a window prompt appears to verify the delete action; click the **OK** button. The entry is removed from the list.

6. To enable client authentication, do the following:

- a. Click the **Client** tab.
- b. Select the **Configure Client Chain Certificate** check box.

Client Chain Certificate is a list of certificates that includes all intermediate CA and root-CA certificates.

Note: To authenticate on the PE with Client Chain Certificate the 'Subject name' field must be present. The subject name should match the userPrincipalName (UPN) in the AD. The UPN is a username with domain address. For example *user1@nutanix.com*.



The screenshot shows the 'Authentication Configuration' page with the 'Client' tab selected. The 'Configure Client Chain Certificate' checkbox is unchecked. Below the checkbox, a description states: 'Force the authentication of all clients of the REST API (including the Prism UI and NCLI). Requires a client chain certificate.'

Figure 17: Client Tab (1)

- c. Click the **Choose File** button, browse to and select a client chain certificate to upload, and then click the **Open** button to upload the certificate.

Note: Uploaded certificate files must be PEM encoded. The web console restarts after the upload step.

Authentication Configuration ?

Directory List · Authentication Types · **Client**

☒ **Configure Client Chain Certificate**

Force the authentication of all clients of the REST API (including the Prism UI and NCLI). Requires a client chain certificate.

Client Chain Certificate

No file chosen

Figure 18: Client Tab (2)

- d. To enable client authentication, click **Enable Client Authentication**.
- e. To modify client authentication, do one of the following:

Note: The web console restarts when you change these settings.

- Click **Enable Client Authentication** to disable client authentication.
- Click **Remove** to delete the current certificate. (This also disables client authentication.)
- To enable OCSP or CRL based certificate revocation checking, see [Certificate Revocation Checking \(nCLI\)](#) on page 193.

☒ **Configure Client Chain Certificate**

Force the authentication of all clients of the REST API (including the Prism UI and NCLI). Requires a client chain certificate.

CLIENT CHAIN CERTIFICATE

CHOOSE FILE

UNIVERSAL_CA_CHAIN.CER

REMOVE

☒ Enable Client Authentication

Figure 19: Authentication Window: Client Tab (3)

Client authentication allows you to securely access the Prism by exchanging a digital certificate. Prism will validate that the certificate is signed by your organization's trusted signing certificate.

Client authentication ensures that the Nutanix cluster gets a valid certificate from the user. Normally, a one-way authentication process occurs where the server provides a certificate so the user can verify the authenticity of the server (see [Importing an SSL Certificate](#) on page 83). When client authentication is enabled, this becomes a two-way authentication where the server also verifies the authenticity of the user. A user must provide a valid certificate when accessing the console either by installing the certificate on their local machine or by providing it through a smart card reader. Providing a valid certificate enables user login from a client machine with the relevant user certificate without utilizing user name and password. If the user is required to login from a client machine which does not have the certificate installed, then authentication using user name and password is still available.

Note: The CA must be the same for both the client chain certificate and the certificate on the local machine or smart card.

7. To specify a service account that the web console can use to log in to Active Directory and authenticate Common Access Card (CAC) users, select the **Configure Service Account** check box, and then do the following in the indicated fields:

The screenshot shows a configuration interface for Common Access Card (CAC) authentication. At the top, there is a toggle switch labeled "Enable Client Authentication" which is currently turned on. Below this, there is a checked checkbox labeled "Configure Service Account" with an "Edit" link next to it. A descriptive text states: "Configure to enable swipe access authentication instead of two-step token authentication." This is followed by a form with three fields: "DIRECTORY" with a dropdown menu showing "AD1", "SERVICE USERNAME" with a text input field containing "cac_auth_acct@mycompany.com", and "SERVICE PASSWORD" with a password input field showing masked characters. At the bottom, there is a "Note" stating: "Enabling CAC Authentication will also enable Client Authentication." and another toggle switch labeled "Enable CAC Authentication" which is currently turned off.

Figure 20: Common Access Card Authentication

- a. **Directory:** Select the authentication directory that contains the CAC users that you want to authenticate.
This list includes the directories that are configured on the **Directory List** tab.
- b. **Service Username:** Enter the user name in the `user_name@domain.com` format that you want the web console to use to log in to the Active Directory.
- c. **Service Password:** Enter the password for the service user name.

d. Click **Enable CAC Authentication**.

Note: Enabling CAC disables all other directory service and local user logons, only the local admin user logon is permitted in this case.

Note: For federal customers only.

Note: The web console restarts after you change this setting.

The Common Access Card (CAC) is a smart card about the size of a credit card, which some organizations use to access their systems. After you insert the CAC into the CAC reader connected to your system, the software in the reader prompts you to enter a PIN. After you enter a valid PIN, the software extracts your personal certificate that represents you and forwards the certificate to the server using the HTTP protocol.

Nutanix Prism verifies the certificate as follows:

- Validates that the certificate has been signed by your organization's trusted signing certificate.
- Extracts the Electronic Data Interchange Personal Identifier (EDIPI) from the certificate and uses the EDIPI to check the validity of an account within the Active Directory. The security context from the EDIPI is used for your PRISM session.
- Prism Element supports both certificate authentication and basic authentication in order to handle both Prism Element login using a certificate and allowing REST API to use basic authentication. It is physically not possible for REST API to use CAC certificates. With this behavior, if the certificate is present during Prism Element login, the certificate authentication is used. However, if the certificate is not present, basic authentication is enforced and used.

Note: Nutanix Prism does not support OpenLDAP as directory service for CAC.

If you map a Prism role to a CAC user and not to an Active Directory group or organizational unit to which the user belongs, specify the EDIPI (User Principal Name, or UPN) of that user in the role mapping. A user who presents a CAC with a valid certificate is mapped to a role and taken directly to the web console home page. The web console login page is not displayed.

Note: If you have logged on to Prism by using CAC authentication, to successfully log out of Prism, close the browser after you click **Log Out**.

8. Click the **Close** button to close the **Authentication Configuration** dialog box.

Assigning Role Permissions

About this task

When user authentication is enabled for a directory service (see [Configuring Authentication](#) on page 58), the directory users do not have any permissions by default. To grant permissions to the directory users, you must specify roles for the users (with associated permissions) to organizational units (OUs), groups, or individuals within a directory.

If you are using Active Directory, you must also assign roles to entities or users, especially before upgrading from a previous AOS version.

To assign roles, do the following:

Procedure

1. Log in to Prism Element web console as an administrator.

2. In the web console, click the gear icon in the main menu and then select **Role Mapping** in the **Settings** page.

The **Role Mapping** window appears.

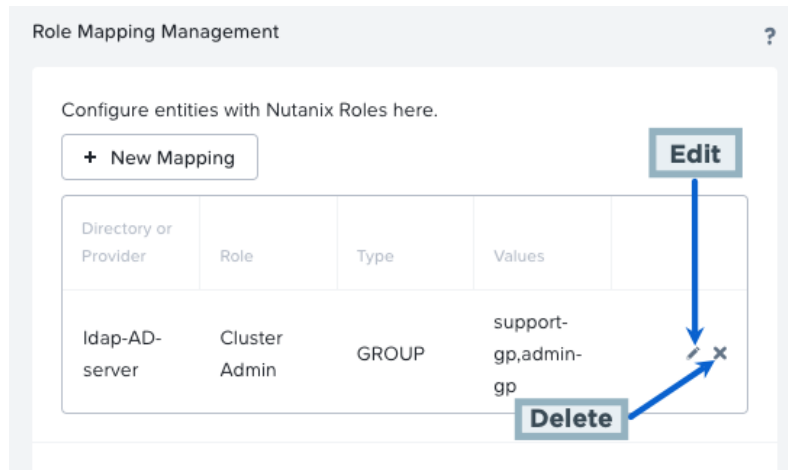


Figure 21: Role Mapping Window

3. To create a role mapping, click the **New Mapping** button.

The **Create Role Mapping** window appears. Do the following in the indicated fields:

- a. **Directory:** Select the target directory from the pull-down list.

Only directories previously defined when configuring authentication appear in this list. If the desired directory does not appear, add that directory to the directory list (see [Configuring Authentication](#) on page 58) and then return to this procedure.

- b. **LDAP Type:** Select the desired LDAP entity type from the pull-down list.

The entity types are **GROUP**, **USER**, and **OU**.

- c. **Role:** Select the user role from the pull-down list.

There are three roles from which to choose:

- **Viewer:** This role allows a user to view information only. It does not provide permission to perform any administrative tasks.
- **Cluster Admin:** This role allows a user to view information and perform any administrative task (but not create or modify user accounts).
- **User Admin:** This role allows the user to view information, perform any administrative task, and create or modify user accounts.

Note: After updating to AOS 6.0 or later, users and user groups with Active Directory (AD) Backup Admin role assigned cannot log in to Prism using their AD credentials or access v3 APIs. For more information, see [KB-14105](#).

- d. **Values:** Enter the case-sensitive entity names (in a comma separated list with no spaces) that should be assigned this role.

The values are the actual names of the organizational units (meaning it applies to all users in those OUs), groups (all users in those groups), or users (each named user) assigned this role. For example, entering value "admin-gp,support-gp" when the LDAP type is **GROUP** and the role is

Cluster Admin means all users in the admin-gp and support-gp groups should be assigned the cluster administrator role.

Note:

- Do not include a domain in the value, for example enter just admin-gp, not admin-gp@nutanix.com. However, when users log into the web console, they need to include the domain in their user name.
- The AD user UPN must be in the user@domain_name format.
- When an admin defines user role mapping using an AD with forest setup, the admin can map to the user with the same name from any domain in the forest setup. To avoid this case, set up the user-role mapping with AD that has a specific domain setup.

e. When all the fields are correct, click **Save**.

This saves the configuration and redisplay the **Role Mapping** window. The new role map now appears in the list.

Note: All users in an authorized service directory have full administrator permissions when role mapping is not defined for that directory. However, after creating a role map, any users in that directory that are not explicitly granted permissions through the role mapping are denied access (no permissions).

f. Repeat this step for each role map you want to add.

You can create a role map for each authorized directory. You can also create multiple maps that apply to a single directory. When there are multiple maps for a directory, the most specific rule for a user applies. For example, adding a **GROUP** map set to **Cluster Admin** and a **USER** map set to **Viewer** for select users in that group means all users in the group have administrator permission except those specified users who have viewing permission only.

Create Role Mapping ?

Enter the attributes for this role mapping.

Directory or Provider
ldap-AD-server


Type
group

Role
Cluster Admin

Values
admin-gp,support-gp

< Back Save

Figure 22: Create Role Mapping Window

4. To edit a role map entry, click the pencil icon  for that entry.
After clicking the pencil icon, the **Edit Role Mapping** window appears, which contains the same fields as the **Create Role Mapping** window (see step 2). Enter the new information in the appropriate fields and then click the **Save** button.
5. To delete a role map entry, click the "X" icon for that entry.
After clicking the X icon, a window prompt appears to verify the delete action; click the **OK** button. The entry is removed from the list.
6. Click the **Close** button to close the **Role Mapping** window.

Authentication Best Practices

The authentication best practices listed here are guidance to secure the Nutanix platform by using the most common authentication security measures.

Emergency Local Account Usage

You must use the admin account as a local emergency account. The admin account ensures that both the Prism Web Console and the Controller VM are available when the external services such as Active Directory is unavailable.

Note: Local emergency account usage does not support any external access mechanisms, specifically for the external application authentication or external Rest API authentication.

For all the external authentication, you must configure the cluster to use an external IAM service such as Active Directory. You must create service accounts on the IAM and the accounts must have access grants to the cluster through Prism web console user account management configuration for authentication.

Modifying Default Passwords

You must change the default Controller VM password for **nutanix** user account by adhering to the password complexity requirements.

Procedure

1. SSH to the Controller VM.
2. Change the "nutanix" user account password.

```
nutanix@cvm$ sudo passwd nutanix
```

3. Respond to the prompts and provide the current and new root password.

```
Changing password for nutanix.  
New password:  
Retype new password:  
passwd: all authentication tokens updated successfully.
```

Note:

- Changing the user account password on one of the Controller VMs is applied to all Controller VMs in the cluster.
- Ensure that you preserve the modified nutanix user password, since the local authentication (PAM) module requires the previous password of the nutanix user to successfully start the password reset process.
- For the root account, both the console and SSH direct login is disabled.

- It is recommended to use the admin user as the administrative emergency account.

Controlling Cluster Access

About this task

Nutanix supports the Cluster lockdown feature. This feature enables key-based SSH access to the Controller VM and AHV on the Host (only for nutanix/admin users).

Enabling cluster lockdown mode ensures that password authentication is disabled and only the keys you have provided can be used to access the cluster resources. Thus making the cluster more secure.

You can create a key pair (or multiple key pairs) and add the public keys to enable key-based SSH access. However, when site security requirements do not allow such access, you can remove all public keys to prevent SSH access.

Note: When you add a public key, SSH access is enabled for both the "nutanix" and "admin" user accounts on the controller VM and the AHV host.

Caution: You should use the SSH-based command-line interface only when directed by Nutanix support or as specified in Nutanix documentation.

Tip: For additional security, you can configure SSH Security Level for SSH access to the Controller VM. See [Hardening Controller VM](#) on page 187 for details.

To control key-based SSH access to the cluster, do the following:

Note: Use this procedure to lock down access to the Controller VM and hypervisor host. In addition, it is possible to lock down access to the hypervisor.

Procedure

1. Click the gear icon in the main menu and then select **Cluster Lockdown** in the **Settings** page. The **Cluster Lockdown** dialog box appears. Enabled public keys (if any) are listed in this window.

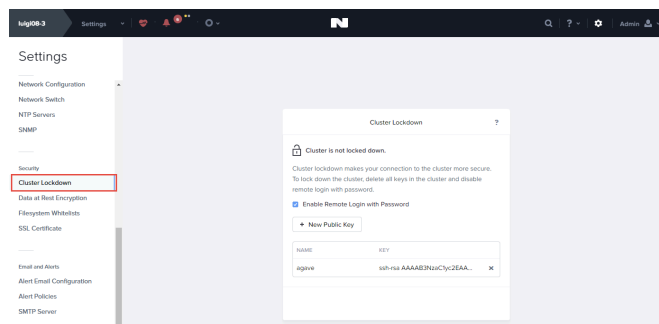


Figure 23: Cluster Lockdown Window

2. To disable (or enable) remote login access, uncheck (check) the **Enable Remote Login with Password** box.
Remote login access is enabled by default.

3. To add a new public key, click the **New Public Key** button and then do the following in the displayed fields:

- a. **Name:** Enter a key name.
- b. **Key:** Enter (paste) the key value into the field.

Note: Prism supports the following key types.

- RSA
- ECDSA

- a. Click the **Save** button (lower right) to save the key and return to the main **Cluster Lockdown** window.

There are no public keys available by default, but you can add any number of public keys.

4. To delete a public key, click the **X** on the right of that key line.

Note: Deleting all the public keys and disabling remote login access locks down the cluster from SSH access.

Setup Admin Session Timeout

By default, the users are logged out automatically after being idle for 15 minutes. You can change the session timeout for users and configure to override the session timeout by following the steps shown below.

Procedure

1. Click the gear icon in the main menu and then select **UI Settings** in the **Settings** page.
2. Select the session timeout for the current user from the **Session Timeout For Current User** drop-down list.

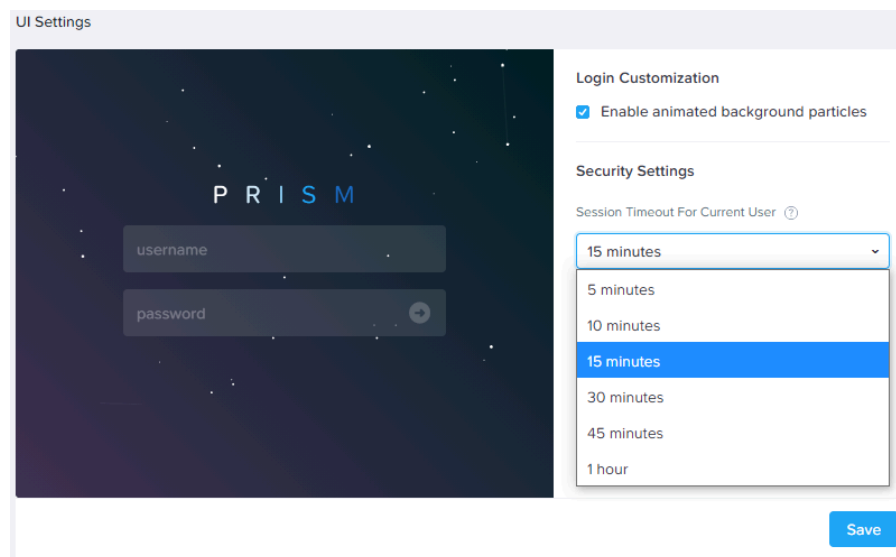


Figure 24: Session Timeout Settings

3. Select the appropriate option from the **Session Timeout Override** drop-down list to override the session timeout.

Password Retry Lockout

For enhanced security, Prism Element locks out the admin account for a period of 15 minutes after a default number of unsuccessful login attempts. Once the account is locked out, the following message is displayed at the logon screen.

Account locked due to too many failed attempts

You can attempt entering the password after the 15 minutes lockout period, or contact Nutanix Support in case you have forgotten your password.

Note: You cannot modify the default 15 minutes lock out period.

Internationalization (i18n)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 3: Internationalization Support

Supported Entities	Unsupported Entities
Cluster name	Acropolis file server
Storage Container name	Share path
Storage pool	Internationalized domain names
VM name	E-mail IDs
Snapshot name	Hostnames
Volume group name	Integers
Protection domain name	Password fields
Remote site name	Any Hardware related names (for example, vSwitch, iSCSI initiator, vLAN name)
User management	
Chart name	

Caution: The creation of none of the above entities are supported on Hyper-V because of the DR limitations.

Entities Support (ASCII or non-ASCII) for the Active Directory Server

- In the New Directory Configuration, **Name** field is supported in non-ASCII.
- In the New Directory Configuration, **Domain** field is not supported in non-ASCII.
- In Role mapping, **Values** field is supported in non-ASCII.
- User names and group names are supported in non-ASCII.

User Management

Nutanix user accounts can be created or updated as needed using the Prism web console.

- The web console allows you to add (see [Creating a User Account](#) on page 74), edit (see [Updating a User Account](#) on page 76), or delete (see [Deleting a User Account \(Local\)](#) on page 82) local user accounts at any time.
- You can reset the local user account password using nCLI if you are locked out and cannot login to the Prism Element or Prism Central web console (see [Resetting Local User Account Password](#) on page 81).
- You can also configure user accounts through Active Directory and LDAP (see [Configuring Authentication](#) on page 58). Active Directory domain created by using non-ASCII text may not be supported.

Note: In addition to the Nutanix user account, there are IPMI, Controller VM, and hypervisor host users. Passwords for these accounts cannot be changed through the web console.

Creating a User Account

About this task

The admin user is created automatically when you get a Nutanix system, but you can add more users as needed. Note that you cannot delete the admin user. To create a user, do the following:

Note: You can also configure user accounts through Active Directory (AD) and LDAP (see [Configuring Authentication](#) on page 58).

Procedure

1. Click the gear icon in the main menu and then select **Local User Management** in the **Settings** page. The **User Management** dialog box appears.

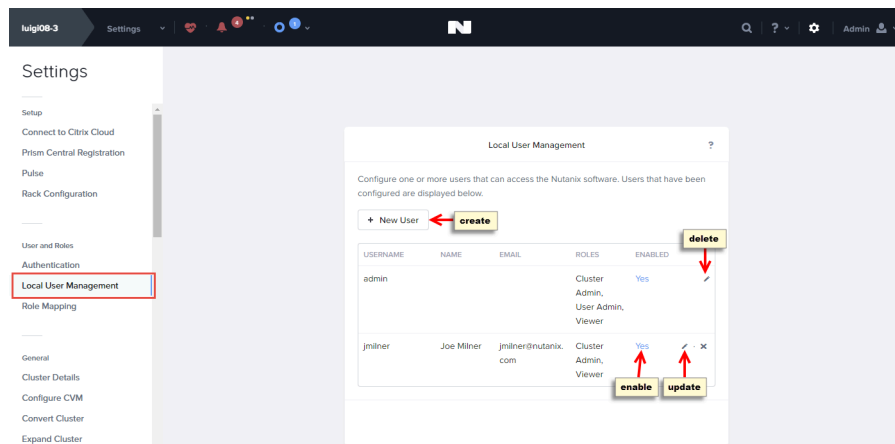


Figure 25: User Management Window

2. To add a user, click the **New User** button and do the following in the displayed fields:

- a. **Username:** Enter a user name.
- b. **First Name:** Enter a first name.
- c. **Last Name:** Enter a last name.
- d. **Email:** Enter a valid user email address.

Note: AOS uses the email address for client authentication and logging when the local user performs user and cluster tasks in the web console.

- e. **Password:** Enter a password (maximum of 255 characters).

A second field to verify that the password is not included, so be sure to enter the password correctly in this field.

- f. **Language:** Select the language setting for the user.

By default **English** is selected. You can select **Simplified Chinese** or **Japanese**. Depending on the language that you select here, the cluster locale is be updated for the new user. For example, if you select **Simplified Chinese**, the next time that the new user logs on to the web console, the user interface is displayed in Simplified Chinese.

- g. **Roles:** Assign a role to this user.

- Select the **User Admin** box to allow the user to view information, perform any administrative task, and create or modify user accounts. (Checking this box automatically selects the **Cluster Admin** box to indicate that this user has full permissions. However, a user administrator has full permissions regardless of whether the cluster administrator box is checked.)
- Select the **Cluster Admin** box to allow the user to view information and perform any administrative task (but not create or modify user accounts).
- Select the **Backup Admin** box to allow the user to perform backup-related administrative tasks. This role does not have permission to perform cluster or user tasks.

Note: Backup admin user is designed for Nutanix Mine integrations as of AOS version 5.19 and has minimal functionality in cluster management. This role has restricted access to the Nutanix Mine cluster.

- **Health, Analysis, and Tasks** features are available in read-only mode.
- The **File server** and **Data Protection** options in the web console are not available for this user.
- The following features are available for Backup Admin users with limited functionality.
 - **Home** - The user cannot a register a cluster with Prism Central. The registration widget is disabled. Other read-only data is displayed and available.
 - **Alerts** - Alerts and events are displayed. However, the user cannot resolve or acknowledge any alert or event. The user cannot configure **Alert Policy** or **Email configuration**.
 - **Hardware** - The user cannot expand the cluster or remove hosts from the cluster. Read-only data is displayed and available.
 - **Network** - Networking data or configuration is displayed but configuration options are not available.

- **Settings** - The user can only upload a new image using the **Settings** page.
- **VM** - The user cannot configure options like **Create VM** and **Network Configuration** in the **VM** page. The following options are available for the user in the **VM** page:

Unpowered On

Unpowered Off

- Leaving all the boxes unchecked allows the user to view information, but it does not provide permission to perform cluster or user tasks.
- h. When all the fields are correct, click **Save**.
This saves the configuration and the web console redisplay the dialog box with the new user-administrative appearing in the list.

Figure 26: Create User Window

Updating a User Account

About this task

Update credentials and change the role for an existing user by using this procedure.

Note: To update your account credentials (that is, the user you are currently logged on as), see [Updating My Account](#) on page 78. Changing the password for a different user is not supported; you must log in as that user to change the password.

Procedure

1. Click the gear icon in the main menu and then select **Local User Management** in the **Settings** page. The **User Management** dialog box appears.

2. Enable or disable the login access for a user by clicking the toggle text **Yes** (enabled) or **No** (disabled) in the **Enabled** column.

A **Yes** value in the **Enabled** column means that the login is enabled; a **No** value in the **Enabled** column means it is disabled.

Note: A user account is enabled (login access activated) by default.

3. To edit the user credentials, click the pencil icon for that user and update one or more of the values in the displayed fields:

- a. **Username:** The username is fixed when the account is created and cannot be changed.
- b. **First Name:** Enter a different first name.
- c. **Last Name:** Enter a different last name.
- d. **Email:** Enter a different valid email address.

Note: AOS Prism uses the email address for client authentication and logging when the local user performs user and cluster tasks in the web console.

- e. **Roles:** Change the role assigned to this user.

- Select the **User Admin** box to allow the user to view information, perform any administrative task, and create or modify user accounts. (Checking this box automatically selects the **Cluster Admin** box to indicate that this user has full permissions. However, a user administrator has full permissions regardless of whether the cluster administrator box is checked.)
- Select the **Cluster Admin** box to allow the user to view information and perform any administrative task (but not create or modify user accounts).
- Select the **Backup Admin** box to allow the user to perform backup-related administrative tasks. This role does not have permission to perform cluster or user administrative tasks.

- Leaving all the boxes unchecked allows the user to view information, but it does not provide permission to perform cluster or user-administrative administrative tasks.

f. **Reset Password:** Change the password of this user.

Enter the new password for **Password** and **Confirm Password** fields. Click the info icon to view the password complexity requirements.

Note: By default, there is no password expiration day set for a local user account.

g. When all the fields are correct, click **Save**.

This saves the configuration and redisplay the dialog box with the new user appearing in the list.

Figure 27: Update User Window

Updating My Account

About this task

To update your account credentials (that is, credentials for the user you are currently logged in as), do the following:

Procedure

1. To update your password, select **Change Password** from the user icon  pull-down list in the web console.

The **Change Password** dialog box appears. Do the following in the indicated fields:

- a. **Current Password:** Enter the current password.
- b. **New Password:** Enter a new password.
- c. **Confirm Password:** Re-enter the new password.
- d. When the fields are correct, click the **Save** button (lower right). This saves the new password and closes the window.

Note: You can change the password for the "admin" account only once per day. Please contact Nutanix support if you need to update the password multiple times in one day

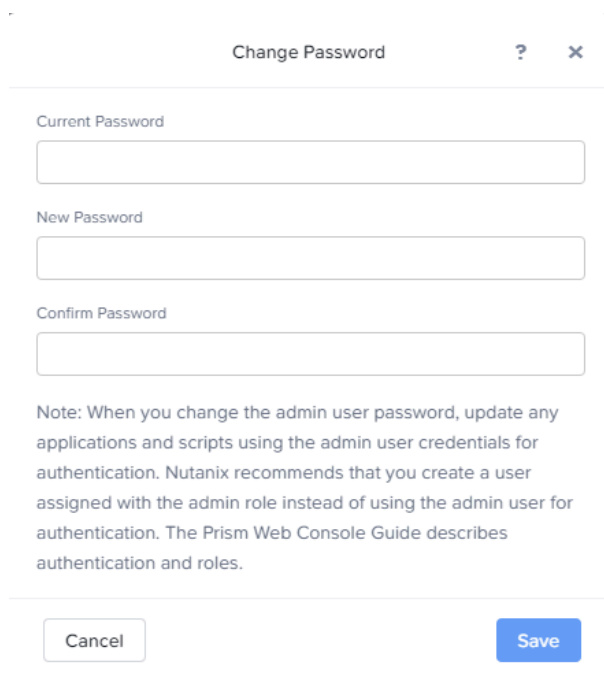
The image shows a 'Change Password' dialog box with a title bar containing a question mark and a close button. Inside the dialog, there are three text input fields labeled 'Current Password', 'New Password', and 'Confirm Password'. Below these fields is a note: 'Note: When you change the admin user password, update any applications and scripts using the admin user credentials for authentication. Nutanix recommends that you create a user assigned with the admin role instead of using the admin user for authentication. The Prism Web Console Guide describes authentication and roles.' At the bottom of the dialog are two buttons: 'Cancel' and 'Save'.

Figure 28: Change Password Window

2. To update other details of your account, select **Update Profile** from the user icon pull-down list. The **Update Profile** dialog box appears. Update (as desired) one or more of the following fields:
 - a. **First Name**: Enter a different first name.
 - b. **Last Name**: Enter a different last name.
 - c. **Email**: Enter a different valid user email address.
 - d. **Language**: Select a language for your account.
 - e. **API Key**: Enter the key value to use a new API key.
 - f. **Public Key**: Click the **Choose File** button to upload a new public key file.
 - g. When all the fields are correct, click the **Save** button (lower right). This saves the changes and closes the window.

Update Profile ? X

Profile settings for **admin**.

General

First Name

Last Name

Email Address

Language

Portal Connection

API Key

Public Key No file chosen Optional

Your keys can be managed from the [API Keys](#) page on the Support Portal. Your connection will be secure without the optional public key, and the public key option is provided in the event that your default public key expires.

Figure 29: Update Profile Window

Resetting Local User Account Password

This procedure describes how to reset a local user's password on the Prism Element or the Prism Central web consoles.

About this task

To reset the password, do the following:

Note:

Only a user with admin privileges can reset a password for other users.

Procedure

1. Access the CVM via SSH.
2. Log in with the admin credentials.
3. Use the `ncli user reset-password` command and specify the username and password of the user whose password is to be reset:

```
nutanix@cvm$ ncli user reset-password user-name=xxxxx password=yyyyy
```

- Replace `user-name=xxxxx` with the name of the user whose password is to be reset.
- Replace `password=yyyyy` with the new password.

What to do next

You can relaunch the Prism Element or the Prism Central web console and verify the new password setting.

Exporting an SSL Certificate for Third-party Backup Applications

Nutanix allows you to export an SSL certificate for Prism Element on a Nutanix cluster and use it with third-party backup applications.

Procedure

1. Log on to a Controller VM in the cluster using SSH.
2. Run the following command to obtain the virtual IP address of the cluster:

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

The current cluster configuration is displayed.

```
Cluster Id       : 0001abl2-abcd-efgh-0123-012345678m89::123456
Cluster Uuid     : 0001abl2-abcd-efgh-0123-012345678m89
Cluster Name     : three
Cluster Version  : 6.0
Cluster Full Version : e17.3-release-fraser-6.0-
a0b1c2345d6789ie123456fg789h1212i34jk5lm6
External IP address : 10.10.10.10
Node Count       : 3
Block Count      : 1
. . . . .
```

Note: The external IP address in the output is the virtual IP address of the cluster.

- 3. Run the following command to enter into the Python prompt:**

```
nutanix@cvm$ python
```

The Python prompt appears.

- 4.** Run the following command to import the SSL library.

```
$ import ssl
```

- 5.** From the Python console, run the following command to print the SSL certificate.

```
$ print ssl.get_server_certificate(('virtual_IP_address',9440),
ssl_version=ssl.PROTOCOL_TLSv1_2)
```

Example: Refer to the following example where `virtual_IP_address` value is replaced by 10.10.10.10.

```
$ print ssl.get_server_certificate(('10.10.10.10', 9440),
ssl_version=ssl.PROTOCOL_TLSv1_2)
```

The SSL certificate is displayed on the console.

[illegible]

Deleting a User Account (Local)

About this task

To delete an existing user, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Local User Management** in the **Settings** page. The **User Management** dialog box appears.

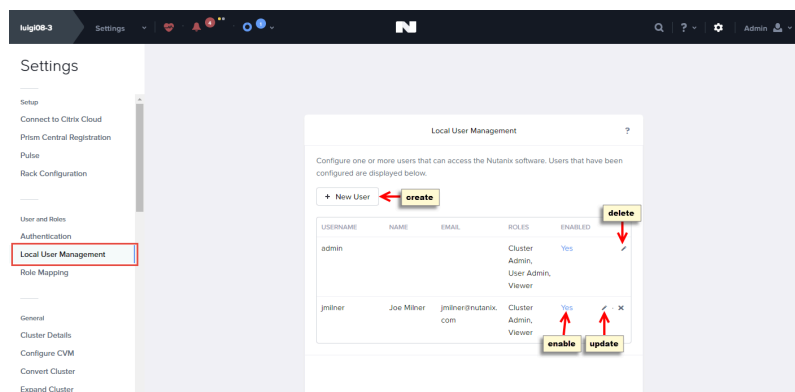


Figure 30: User Management Window

2. Click the **X** icon for that user. Note that you cannot delete the admin user. A window prompt appears to verify the action; click the **OK** button. The user account is removed and the user no longer appears in the list.

SSL Certificate Management

Prism web console supports SSL certificate-based authentication for console access. To enable secure communication with a cluster, Prism web console includes a default self-signed SSL certificate. You can replace the default self-signed SSL certificate with your own self-signed SSL certificate or a certificate authority (CA) signed SSL certificate.

For production purposes, Nutanix recommends that you replace the default self-signed certificate with a CA signed SSL certificate.

Note:

- You can import only a cluster-wide SSL certificate in Prism web console. The SSL certificate cannot be customized for an individual controller VM (CVM).
- Nutanix recommends that you check for the validity of the certificate periodically and replace the certificate if it is invalid.

Importing an SSL Certificate

Nutanix simplifies the SSL certificate import process into Prism web console.

Before you begin

Depending upon your requirements, you need to either generate a self-signed SSL certificate or generate a Certificate Signing Request (CSR) for submission to a certificate Authority (CA) to get a CA signed certificate.

For more information, see

- [Generating a Self-signed SSL Certificate with Subject Alternative Name](#) on page 87
- [Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority \(CA\)](#) on page 90

About this task

The following procedure explains how to import a self-signed SSL certificate or a CA signed SSL certificate into Prism web console.

Procedure

1. Click the gear icon in the main menu and in the **Settings** page, select **SSL Certificate**. The **SSL Certificate** dialog box appears.

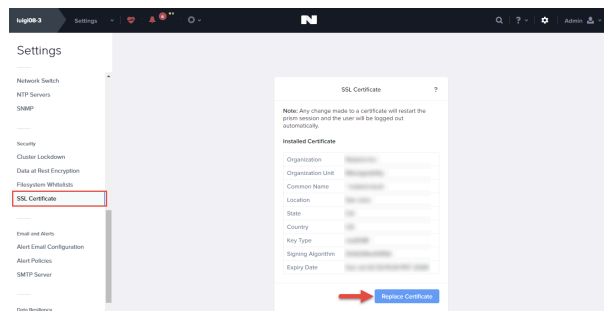


Figure 31: SSL Certificate Window

2. To replace an SSL certificate, click **Replace Certificate**.
3. Do one of the following:
 - » To regenerate Nutanix default self-signed certificate, select **Regenerate Self Signed Certificate** and then click **Apply**.
A dialog box appears to verify the action; click **OK**. A new RSA 2048 bit self-signed certificate is generated and applied for Prism web console.
 - » To import self-signed SSL certificate or CA signed certificate, select **Import Key and Certificate** and then click **Next**.

4. To import self-signed SSL certificate or CA signed certificate files, do the following:

For self-signed certificate:

- **Private Key Type:** Select the appropriate private key type for the self-signed certificate from the dropdown list.
- **Private Key:** Click **Choose file** and select the private key.
- **Public Certificate:** Click **Choose file** and select the self-signed certificate corresponding to the private key.
- **CA Certificate/Chain:** Click **Choose file** select the self-signed certificate corresponding to the private key.

SSL Certificate ?

Note: Any change made to a certificate will restart the prism session and the user will be logged out automatically.

Guidelines for RSA: Please use a SHA-256, SHA-384, or SHA-512 signature algorithm with RSA certificates for security and performance.

Private Key Type

RSA 2048 bit

Private Key ⓘ

Choose file my_key_name.key

Public Certificate

Choose file my crt_name.crt

CA Certificate/Chain

Choose file my crt_name.crt

Cancel Import Files

Figure 32: Importing self-signed certificate

The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	my crt_name.crt

Certificate Components	File type
CA Certificate/Chain	my_crt_name.crt

For CA signed certificate:

- **Private Key Type:** Select the appropriate private key type for the CA signed certificate from the dropdown list.
- **Private Key:** Click **Choose file** and select the private key.
- **Public Certificate:** Click **Choose file** and select the CA signed public portion of the certificate corresponding to the private key.
- **CA Certificate/Chain:** Click **Choose file** and select the certificate or chain of the signing authority for the public certificate.

Note: To create a chain file from the list of CA certificates, see [Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority \(CA\)](#) on page 90.

SSL Certificate ?

Note: Any change made to a certificate will restart the prism session and the user will be logged out automatically.

Guidelines for RSA: Please use a SHA-256, SHA-384, or SHA-512 signature algorithm with RSA certificates for security and performance.

Private Key Type

RSA 2048 bit

Private Key ⓘ

Choose file my_key_name.key

Public Certificate

Choose file ca_signed_public_cert.cer

CA Certificate/Chain

Choose file ca_public_cert.crt

Cancel Import Files

Figure 33: Importing CA signed certificate

The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	ca_signed_public_cert.cer
CA Certificate/Chain	ca_public_cert.crt or ca_chain_certs.crt

5. To begin SSL certificate import, click **Import Files**.

Results

After generating or importing the new certificate, the interface gateway restarts. If the certificate and credentials are valid, the interface gateway uses the new certificate immediately, which means that your browser session (and all other open browser sessions) is invalid until you reload the page and accept the new certificate. If anything is wrong with the certificate (such as a corrupted file or wrong certificate type), the new certificate is discarded, and the system reverts to the original default certificate provided by Nutanix.

Note: The system holds only one custom SSL certificate. If a new certificate is uploaded, it replaces the existing certificate. The previous certificate is discarded.

Recommended Key Configurations

This table provides the Nutanix recommended set of key types, sizes/curves, and signature algorithms.

Key Type	Size/Curve	Signature Algorithm
RSA	4096	SHA-256, SHA-384 or SHA512
RSA	2048	SHA-256, SHA-384 or SHA512
EC DSA 256	prime256v1	ecdsa-with-sha256
EC DSA 384	secp384r1	ecdsa-with-sha384
EC DSA 521	secp521r1	ecdsa-with-sha512

Note:

- Client and CAC authentication only supports RSA 2048 bit certificate.
- RSA 4096 bit certificates might not work with certain AOS and Prism Central releases, see the release notes for your AOS and Prism Central versions for details. Specifying an RSA 4096 bit certificate might cause multiple cluster services to restart frequently. To work around the issue, see [KB 12775](#).
- Certificate import fails if you attempt to upload SHA-1 certificate (including root CA).

Generating a Self-signed SSL Certificate with Subject Alternative Name

This task explains how to generate a self-signed SSL certificate using OpenSSL commands.

About this task

To comply with the security standards of NIST SP800-131a and the requirements of the RFC 6460 for NSA Suite B, the certificate import process validates that the correct signature algorithm is used for a given key and certificate pair. Use proper set of key types, sizes and curves, and signature algorithms. For more information, see [Recommended Key Configurations](#) on page 87.

Nutanix recommends including a DNS name for all controller VMs (CVMs) in the self-signed SSL certificate using the SAN extension. This scheme helps avoid SSL certificate errors when you access a CVM by direct DNS instead of the shared cluster IP address.

Procedure

1. Log in to any of the controller VMs (CVMs) with SSH using the management IP address of the CVM:

```
$ ssh admin@cvm_ip_address
```

2. To generate a private key with a bit length of your choice, run one of the following commands:

- RSA private key with a bit length of 2048:

```
admin@cvm$ openssl genrsa -out my_key_name.key 2048
```

- RSA private key with a bit length of 4096:

```
admin@cvm$ openssl genrsa -out my_key_name.key 4096
```

- ECDSA private key using the prime256v1 curve:

```
admin@cvm$ openssl ecparam -name prime256v1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp384r1 curve:

```
admin@cvm$ openssl ecparam -name secp384r1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp521r1 curve:

```
admin@cvm$ openssl ecparam -name secp521r1 -genkey -out my_key_name.pem
```

Important: While you are generating the private key, ensure that the private key is not password protected.

3. To generate the CSR of your choice, run one of the following commands:

- For RSA 2048 and RSA 4096 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -signature_algorithm -out my_csr_name.csr
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for RSA 2048 private key using SHA-256 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -sha256 -out my_csr_name.csr
```

- For ECDSA 256, ECDSA 384, ECDSA 521 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -signature_algorithm -out my_csr_name.csr
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for ECDSA 384 private key using SHA-384 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -sha384 -out my_csr_name.csr
```

4. Enter the information in the command output to incorporate into your certificate request:

```
You are about to be asked to enter information that will be incorporated
```


into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.

```
-----
Country Name (2 letter code) []:US
State or Province Name (full name) []:CA
Locality Name (eg, city) []:San Jose
Organization Name (eg, company) []: Nutanix Inc
Organizational Unit Name (eg, section) []: IT
Common Name (eg, fully qualified host name) []:ntx.com
Email Address []:myname@domain.com
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []: Enter your password
```

5. Create a configuration file in your home directory with your preferred text editor named `san.cnf` that contains the following text:

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
[v3_req]
basicConstraints = CA:FALSE
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
subjectAltName = @alt_names

[alt_names]
DNS.0 = example1.domain.com
DNS.1 = example2.domain.com
DNS.2 = example3.domain.com
DNS.3 = *.domain.com
IP.0 = x.x.x.x

[alt_names]
```

Specify your DNS and IP addresses. If you have a range of hosts, use wildcards (*) to match any subdomain of the domain name.

6. Generate a self-signed certificate:

```
admin@cvm$ openssl x509 -req -days number_of_days -in my_csr_name.csr -signkey
my_key_name.key -out my crt_name.crt -signature_algorithm -extensions v3_req -
extfile san.cnf
```

`number_of_days`: Specify the number of days until a newly generated certificate expires.

`signature_algorithm`: Specify sha256 or sha384 or sha512.

Example:

```
admin@cvm$ openssl x509 -req -days 1460 -in my_csr_name.csr -signkey my_key_name.key
-out my crt_name.crt -sha256 -extensions v3_req -extfile san.cnf
```

7. Copy `my_key_name.key` and `my crt_name.crt` from the CVM to your local machine:

```
admin@cvm$ scp my_key_name.key my crt_name.crt username@local-machine:/
local_file_path/
```

8. Log out of the CVM.

What to do next

After you successfully create a self-signed certificate with a private key, follow the procedure described in [Importing an SSL Certificate](#) on page 83 to replace the default certificate with your self-signed SSL certificate. The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Table 4: SSL Certificate Import Files

Certificate Components	File type
Private Key	my_key_name.key
Public Certificate	my_crt_name.crt
CA Certificate/Chain	my_crt_name.crt

Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority (CA)

This task explains how to generate a certificate signing request using OpenSSL commands.

About this task

To comply with the security standards of NIST SP800-131a and the requirements of the RFC 6460 for NSA Suite B, the certificate import process validates the correct signature algorithm is used for a given key and certificate pair. Use proper set of key types, sizes and curves, and signature algorithms. For more information, see [Recommended Key Configurations](#) on page 87.

Nutanix recommends including a DNS name for all CVMs in the CSR using the Subject Alternative Name (SAN) extension. This avoids SSL certificate errors when you access a CVM by direct DNS instead of the shared cluster IP address.

Procedure

1. Log in to any of the controller VMs (CVMs) with SSH using the management IP address of the CVM:

```
$ ssh admin@cvm_ip_address
```

2. Create a configuration file in your home directory with your preferred text editor named `ssl.cnf` that contains the following text:

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
countryName = Country Name (2 letter code)
stateOrProvinceName = State or Province Name (full name)
localityName = Locality Name (eg, city)
organizationName = Organization Name (eg, company)
organizationalUnitName = Organizational Unit Name (eg, BU)
commonName = Common Name (e.g. server FQDN or YOUR name)
emailAddress = Email Address

[v3_req]
subjectAltName = @alt_names

[alt_names]
DNS.0 = example1.domain.com
```

```
DNS.1 = example2.domain.com
DNS.2 = example3.domain.com
DNS.3 = *.domain.com
IP.0 = x.x.x.x
```

[alt_names]

Specify your DNS and IP addresses. If you have a range of hosts, use wildcards (*) to match any subdomain of the domain name.

3. To generate a private key with a bit length of your choice, run one of the following commands:

- RSA private key with a bit length of 2048:

```
admin@cvm$ openssl genrsa -out my_key_name.key 2048
```

- RSA private key with a bit length of 4096:

```
admin@cvm$ openssl genrsa -out my_key_name.key 4096
```

- ECDSA private key using the prime256v1 curve:

```
admin@cvm$ openssl ecparam -name prime256v1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp384r1 curve:

```
admin@cvm$ openssl ecparam -name secp384r1 -genkey -out my_key_name.pem
```

- ECDSA private key using the secp521r1 curve:

```
admin@cvm$ openssl ecparam -name secp521r1 -genkey -out my_key_name.pem
```

Important: While you are generating the private key, ensure that the private key is not password protected.

4. To generate the CSR of your choice, run one of the following commands:

- For RSA 2048 and RSA 4096 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -signature_algorithm -out my_csr_name.csr -config ssl.cnf
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for RSA 2048 private key using SHA-256 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.key -sha256 -out my_csr_name.csr -config ssl.cnf
```

- For ECDSA 256, ECDSA 384, ECDSA 521 private keys:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -signature_algorithm -out my_csr_name.csr -config ssl.cnf
```

signature_algorithm: Specify sha256 or sha384 or sha512.

For example, to generate a CSR for ECDSA 384 private key using SHA-384 signature algorithm:

```
admin@cvm$ openssl req -new -nodes -key my_key_name.pem -sha384 -out my_csr_name.csr -config ssl.cnf
```

5. Enter the information in the command output to incorporate into your certificate request:

You are about to be asked to enter information that will be incorporated into your certificate request.

```

What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []: US
State or Province Name (full name) []: CA
Locality Name (eg, city) []: San Jose
Organization Name (eg, company) []: Nutanix Inc
Organizational Unit Name (eg, BU) []: IT
Common Name (e.g. server FQDN or YOUR name) []: ntx.com
Email Address []: myname@domain.com

```

6. Copy `my_key_name.key` and `my_cert_name.crt` from the CVM to your local machine:

```

nutanix@cvm$ scp my_key_name.key my_cert_name.crt username@local-machine:/
local_file_path/

```

7. Log out of the CVM.
8. Send your CSR file to the CA of your choice.

After receiving your CSR, the CA sends the following files:

- CA signed public certificate
- CA's public certificate
- Root CA public certificate (if the CA is intermediate)

The public certificate is validated by the issuing CA, and if the issuing CA is intermediate, the issuing CA certificate is validated by the root CA. A chain of certificates is validated to establish the trust.

9. Download all the certificate files received from CA to the local file directory.
10. (Optional) If the CA chain certificate provided by the certificate authority is not in a single file, run the following command to concatenate the list of CA certificates into a chain file:

```

$ cat intermediateCAcert.crt rootCAcert.crt > ca_chain_certs.crt

```

Note:

- The chain must start with the certificate of the signer and ends with the root CA certificate.
- Ensure that the chain file only has the root and intermediate certificates. If your chain file has public or private certificates, it will fail to import in Prism web console.

What to do next

Follow [Importing an SSL Certificate](#) on page 83 section to replace the default certificate with a CA signed certificate. The following table lists certificate components and its corresponding file type to choose when SSL certificate window prompts:

Table 5: SSL Certificate Import Files

Certificate Components	File type
Private Key	<code>my_key_name.key</code>
Public Certificate	<code>ca_signed_public_cert.cer</code>

Certificate Components	File type
CA Certificate/Chain	ca_public_cert.crt or ca_chain_certs.crt

Verifying the Certificate Generation Request

Run the following commands to verify the certificate generation request.

- Verify that the generated certificate chain is OK:

```
admin@cvm$ openssl verify -CAfile ca_chain_certs.crt myPublicCert.cer
```

Example output:

```
myPublicCert.cer: OK
```

- Verify the private key and signature algorithm details:

```
admin@cvm$ openssl x509 -in my_cert_name.crt -text -noout | grep -i 'rsa\|ecdsa\|Public'
```

Example output:

```
Signature Algorithm: ecdsa-with-SHA512
    Subject Public Key Info:
        Public Key Algorithm: id-ecPublicKey
        Public-Key: (521 bit)
        Signature Algorithm: ecdsa-with-SHA512
```

- Verify that the CA certificate chain uses SHA 256 as a signature algorithm:

```
admin@cvm$ openssl crl2pkcs7 -nocrl -certfile ca_chain_certs.crt | openssl pkcs7 -print_certs -noout -text | grep -Ew '(Subject|Issuer|Signature Algorithm):' | grep -C1 Issuer
```

Troubleshooting the Certificate Generation Request

The following troubleshooting tips can help you resolve common issues that can occur when generating certificates.

Chain certificate format

If your chain certificate file has public or private certificates, it will fail to import in Prism web console. Ensure that the chain certificate file only has the root and intermediate certificates.

For example, if a public certificate is present in a chain file, you can remove it by opening your chain file in your preferred text editor. Ensure that there are no extra white spaces at the bottom of the file.

DER-encoded certificate issue

If the certificate is DER encoded, it fails to import in Prism web console. You can resolve the issue by converting it to PEM-encoded ASCII format.

- Ensure that the certificate is DER encoded:

```
admin@cvm$ openssl x509 -in cert.crt -inform der -text -noout
```

- If the certificate is DER encoded, run the following command to convert the certificate from DER to PEM-encoded ASCII format:

```
admin@cvm$ openssl x509 -in certDER.crt -inform der -outform pem -out cert.crt
```

Certificate format

Ensure that all the certificates do not have any extra data (or custom attributes) before the beginning (-----BEGIN CERTIFICATE-----) or after the end (-----END CERTIFICATE-----) of the block.

Exporting an SSL Certificate for Third-party Backup Applications

Nutanix allows you to export an SSL certificate for Prism Element on a Nutanix cluster and use it with third-party backup applications.

Procedure

1. Log on to a Controller VM in the cluster using SSH.
2. Run the following command to obtain the virtual IP address of the cluster:

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

The current cluster configuration is displayed.

```
Cluster Id       : 0001abl2-abcd-efgh-0123-012345678m89::123456
Cluster Uuid     : 0001abl2-abcd-efgh-0123-012345678m89
Cluster Name     : three
Cluster Version  : 6.0
Cluster Full Version : e17.3-release-fraser-6.0-
a0b1c2345d6789ie123456fg789h1212i34jk5lm6
External IP address : 10.10.10.10
Node Count       : 3
Block Count      : 1
. . . . .
```

Note: The external IP address in the output is the virtual IP address of the cluster.

3. Run the following command to enter into the Python prompt:

```
nutanix@cvm$ python
```

The Python prompt appears.

4. Run the following command to import the SSL library.

```
$ import ssl
```

5. From the Python console, run the following command to print the SSL certificate.

```
$ print ssl.get_server_certificate(('virtual_IP_address',9440),
ssl_version=ssl.PROTOCOL_TLSv1_2)
```

Example: Refer to the following example where `virtual_IP_address` value is replaced by 10.10.10.10.

```
$ print ssl.get_server_certificate(('10.10.10.10', 9440),
ssl_version=ssl.PROTOCOL_TLSv1_2)
```

The SSL certificate is displayed on the console.

```
-----BEGIN CERTIFICATE-----
0123456789ABCDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz01
23456789ABCDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123
456789ABCDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz012345
6789ABCDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz01234567
89ABCDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789
ABCDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789AB
CDEFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCD
EFGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDE
FGHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEF
GHIJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGH
IJKLMNopQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJ
```

```
KLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
MNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
OPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
QRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
STUVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
UVWXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
WXYZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
YZabcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxyz0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ
-----END CERTIFICATE-----
```

Controlling Cluster Access

About this task

Nutanix supports the Cluster lockdown feature. This feature enables key-based SSH access to the Controller VM and AHV on the Host (only for nutanix/admin users).

Enabling cluster lockdown mode ensures that password authentication is disabled and only the keys you have provided can be used to access the cluster resources. Thus making the cluster more secure.

You can create a key pair (or multiple key pairs) and add the public keys to enable key-based SSH access. However, when site security requirements do not allow such access, you can remove all public keys to prevent SSH access.

Note: When you add a public key, SSH access is enabled for both the "nutanix" and "admin" user accounts on the controller VM and the AHV host.

Caution: You should use the SSH-based command-line interface only when directed by Nutanix support or as specified in Nutanix documentation.

Tip: For additional security, you can configure SSH Security Level for SSH access to the Controller VM. See [Hardening Controller VM](#) on page 187 for details.

To control key-based SSH access to the cluster, do the following:

Note: Use this procedure to lock down access to the Controller VM and hypervisor host. In addition, it is possible to lock down access to the hypervisor.

Procedure

1. Click the gear icon in the main menu and then select **Cluster Lockdown** in the **Settings** page. The **Cluster Lockdown** dialog box appears. Enabled public keys (if any) are listed in this window.

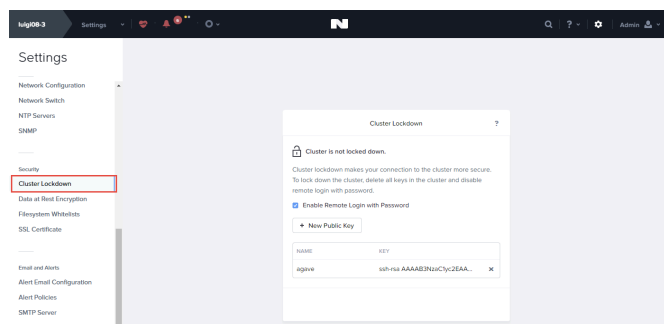


Figure 34: Cluster Lockdown Window

2. To disable (or enable) remote login access, uncheck (check) the **Enable Remote Login with Password** box.
Remote login access is enabled by default.
3. To add a new public key, click the **New Public Key** button and then do the following in the displayed fields:
 - a. **Name:** Enter a key name.
 - b. **Key:** Enter (paste) the key value into the field.

Note: Prism supports the following key types.

- RSA
- ECDSA

- a. Click the **Save** button (lower right) to save the key and return to the main **Cluster Lockdown** window.

There are no public keys available by default, but you can add any number of public keys.

4. To delete a public key, click the **X** on the right of that key line.

Note: Deleting all the public keys and disabling remote login access locks down the cluster from SSH access.

Data-at-Rest Encryption

Nutanix provides an option to secure data while it is at rest using either self-encrypted drives or software-only encryption and key-based access management (cluster's native or external KMS for software-only encryption).

Encryption Methods

Nutanix provides you with the following options to secure your data.

- **Self Encrypting Drives (SED) Encryption** - You can use a combination of SEDs and an external KMS to secure your data while it is at rest.

- **Software-only Encryption** - Nutanix AOS uses the AES-256 encryption standard to encrypt your data. Once enabled, software-only data-at-rest encryption cannot be disabled, thus protecting against accidental data leaks due to human errors. Software-only encryption supports both Nutanix Native Key Manager (local and remote) and External KMS to secure your keys.

Note the following points regarding data-at-rest encryption.

- Encryption is supported for AHV, ESXi, and Hyper-V.
 - For ESXi and Hyper-V, software-only encryption can be implemented at a cluster level or container level.
 - For AHV, encryption can be implemented at the cluster level, VM level, or VG level. For more information about VM or VG level encryption, see [Storage Policy Based Encryption](#).
- Nutanix recommends using cluster-level encryption. With the cluster-level encryption, the administrative overhead of selecting different containers for the data storage gets eliminated.
- Encryption cannot be disabled once it is enabled at a cluster level or container level.
- Encryption can be implemented on an existing cluster with data that exists. If encryption is enabled on an existing cluster (AHV, ESXi, or Hyper-V), the unencrypted data is transformed into an encrypted format in a low priority background task that is designed not to interfere with other workload running in the cluster.
- Data-at-rest encryption choices can be implemented at the entity level using storage policies in Prism Central.

Deployments can continue to have data-at-rest encryption capabilities scoped for the entire cluster. Storage policy provides the additional option to control the encryption scope decisions at the entity (VM or VG) level. For more information, see [Storage Policies Based Encryption](#) in the *Prism Central Infrastructure Guide*

- Data can be encrypted using either self-encrypted drives (SEDs) or software-only encryption. You can change the encryption method from SEDs to software-only. You can perform the following configurations.
 - For ESXi and Hyper-V clusters, you can switch from SEDs and External Key Management (EKM) combination to software-only encryption and EKM combination. First, you must disable the encryption in the cluster where you want to change the encryption method. Then, select the cluster and enable encryption to transform the unencrypted data into an encrypted format in the background.
 - For AHV, background encryption is supported.
- Once the task to encrypt a cluster begins, you cannot cancel the operation. Even if you stop and restart the cluster, the system resumes the operation.
- In the case of mixed clusters with ESXi and AHV nodes, where the AHV nodes are used for storage only, the encryption policies consider the cluster as an ESXi cluster. So, the cluster-level and container-level encryption are available.
- You can use a combination of SED and non-SED drives in a cluster. After you encrypt a cluster using the software-only encryption, all the drives are considered as unencrypted drives. In case you switch from the SED encryption to the software-only encryption, you can add SED or non-SED drives to the cluster.
- Data is not encrypted when it is replicated to another cluster. You must enable the encryption for each cluster. Data is encrypted as a part of the write operation and decrypted as a part of the read operation. During the replication process, the system reads, decrypts, and then sends the data over to the other

cluster. You can use a third-party network solution if there a requirement to encrypt the data during the transmission to another cluster.

- Software-only encryption does not impact most of data efficiency features such as deduplication, compression, zero block suppression, and so on. The software encryption is the last data transformation performed. For example, during the write operation, compression is performed first, followed by encryption.

Note: Software-only encryption requires additional space if erasure coding (EC) is enabled on the cluster (or container). The additional space requirement is temporary, and EC space savings are restored once the encryption process is complete.

Enabling encryption on EC-enabled clusters involves completely decoding EC-encoded data and re-encoding the data for EC. This process of decoding and re-encoding requires additional space on the cluster. Encryption is initiated only if the aggregate of EC-based space saving and current space usage is less than 85 percent of the cluster capacity. Additionally, the cluster continuously checks if the aggregate space usage stays below 85 percent for the encryption to progress.

Key Management

Nutanix supports a Native Key Management Server, also called Local Key Manager (LKM), thus avoiding the dependency on an External Key Manager (EKM). Cluster localised Key Management Service support requires a minimum of 3-node in a cluster and is supported only for software-only encryption. So, 1-node and 2-node clusters can use either the Native KMS (remote) option or an EKM. .

The following types of keys are used for encryption.

- Data Encryption Key (DEK) - A symmetric key, such as AES-256, that is used to encrypt the data.
- Key Encryption Key (KEK) - This key is used to encrypt or decrypt the DEK.

Note the following points regarding the key management.

- Nutanix does not support the use of the Local Key Manager with a third party External Key Manager.
- Dual encryption (both SED and software-only encryption) requires an EKM. For more information, see [Configuring Dual Encryption](#) on page 124.
- You can switch from an EKM to LKM, and inversely. For more information, see [Switching between Native Key Manager and External Key Manager](#) on page 121.
- Rekey of keys stored in the Native KMS is supported for the Leader Keys. For more information, see [Changing Key Encryption Keys \(SEDs\)](#) on page 109 and [Changing Key Encryption Keys \(Software Only\)](#) on page 122.
- You must back up the keys stored in the Native KMS. For more information, see [Backing up Keys](#) on page 125.
- You must backup the encryption keys whenever you create a new container or remove an existing container. Nutanix Cluster Check (NCC) checks the status of the backup and sends an alert if you do not take a backup at the time of creating or removing a container.

Data-at-Rest Encryption (SEDs)

For customers who require enhanced data security, Nutanix provides a data-at-rest security option using Self Encrypting Drives (SEDs) included in the Ultimate license.

Note: If you are running the AOS Pro License on G6 platforms and above, you can use SED encryption by installing an add-on license.

Following features are supported:

- Data is encrypted on all drives at all times.
- Data is inaccessible in the event of drive or node theft.
- Data on a drive can be securely destroyed.
- A key authorization method allows password rotation at arbitrary times.
- Protection can be enabled or disabled at any time.
- No performance penalty is incurred despite encrypting all data.
- Re-key of the leader encryption key (MEK) at arbitrary times is supported.

Note: If an SED cluster is present, then while executing the data-at-rest encryption, you will get an option to either select data-at-rest encryption using SEDs or data-at-rest encryption using AOS.

Encryption Type

Encrypt the cluster using:

☒ Drive-based encryption

Use the cluster's self-encrypting drives (SEDs) to encrypt all cluster data.

☐ Software-based encryption

Figure 35: SED and AOS Options

Note: This solution provides enhanced security for data on a drive, but it does not secure data in transit.

Data Encryption Model

To accomplish these goals, Nutanix implements a data security configuration that uses SEDs with keys maintained through a separate key management device. Nutanix uses open standards (TCG and KMIP protocols) and FIPS validated SED drives for interoperability and strong security.



Figure 36: Cluster Protection Overview

This configuration involves the following workflow:

1. The security implementation begins by installing SEDs for all data drives in a cluster.

The drives are FIPS 140-2 validated and use FIPS 140-2 validated cryptographic modules.

Creating a new cluster that includes SEDs only is straightforward, but an existing cluster can be converted to support data-at-rest encryption by replacing the existing drives with SEDs (after migrating all the VMs/vDisks off of the cluster while the drives are being replaced).

Note: Contact Nutanix customer support for assistance before attempting to convert an existing cluster. A non-protected cluster can contain both SED and standard drives, but Nutanix does not support a mixed cluster when protection is enabled. All the disks in a protected cluster must be SED drives.

2. Data on the drives is always encrypted but read or write access to that data is open. By default, the access to data on the drives is protected by the in-built manufacturer key. However, when data protection for the cluster is enabled, the Controller VM must provide the proper key to access data on a SED. The Controller VM communicates with the SEDs through a Trusted Computing Group (TCG) Security Subsystem Class (SSC) Enterprise protocol.

A symmetric data encryption key (DEK) such as AES 256 is applied to all data being written to or read from the disk. The key is known only to the drive controller and never leaves the physical subsystem, so there is no way to access the data directly from the drive.

Another key, known as a key encryption key (KEK), is used to encrypt/decrypt the DEK and authenticate to the drive. (Some vendors call this the authentication key or PIN.)

Each drive has a separate KEK that is generated through the FIPS compliant random number generator present in the drive controller. The KEK is 32 bytes long to resist brute force attacks. The KEKs are sent to the key management server for secure storage and later retrieval; they are not stored locally on the node (even though they are generated locally).

In addition to the above, the leader encryption key (MEK) is used to encrypt the KEKs.

Each node maintains a set of certificates and keys in order to establish a secure connection with the external key management server.

3. Keys are stored in a key management server that is outside the cluster, and the Controller VM communicates with the key management server using the Key Management Interoperability Protocol (KMIP) to upload and retrieve drive keys.

Only one key management server device is required, but it is recommended that multiple devices are employed so the key management server is not a potential single point of failure. Configure the key manager server devices to work in clustered mode so they can be added to the cluster configuration as a single entity that is resilient to a single failure.

4. When a node experiences a full power off and power on (and cluster protection is enabled), the controller VM retrieves the drive keys from the key management server and uses them to unlock the drives.

If the Controller VM cannot get the correct keys from the key management server, it cannot access data on the drives.

If a drive is re-seated, it becomes locked.

If a drive is stolen, the data is inaccessible without the KEK (which cannot be obtained from the drive). If a node is stolen, the key management server can revoke the node certificates to ensure they cannot be used to access data on any of the drives.

Preparing for Data-at-Rest Encryption (External KMS for SEDs and Software Only)

About this task

Caution: DO NOT HOST A KEY MANAGEMENT SERVER VM ON THE ENCRYPTED CLUSTER THAT IS USING IT!

Doing so could result in complete data loss if there is a problem with the VM while it is hosted in that cluster.

If you are using an external KMS for encryption using AOS, preparation steps outside the web console are required. The information in this section is applicable if you choose to use an external KMS for configuring encryption.

You must install the license of the external key manager for all nodes in the cluster. See [Compatibility and Interoperability Matrix](#) for a complete list of the supported key management servers. For instructions on how to configure a key management server, refer to the documentation from the appropriate vendor.

The system accesses the EKM under the following conditions:

- Starting a cluster
- Regenerating a key (key regeneration occurs automatically every year by default)
- Adding or removing a node (only when Self Encrypting Drives is used for encryption)
- Switching between Native to EKM or EKM to Native
- Starting, and restarting a service (only if Software-based encryption is used)
- Upgrading AOS (only if Software-based encryption is used)
- NCC heartbeat check if EKM is alive

Procedure

1. Configure a key management server.

The key management server devices must be configured into the network so the cluster has access to those devices. For redundant protection, it is recommended that you employ at least two key management server devices, either in active-active cluster mode or stand-alone.

Note: The key management server must support KMIP version 1.0 or later.

» SafeNet

Ensure that **Security > High Security > Key Security > Disable Creation and Use of Global Keys** is checked.

» Vormetric

Set the appliance to compatibility mode. Suite B mode causes the SSL handshake to fail.

2. Generate a certificate signing request (CSR) for each node in the cluster.

- The Common Name field of the CSR is populated automatically with `unique_node_identifier.nutanix.com` to identify the node associated with the certificate.

Tip: After generating the certificate from Prism, (if required) you can update the custom common name (CN) setting by running the following command using nCLI.

```
ncli data-at-rest-encryption-certificate update-csr-information domain-name=abcd.test.com
```

In the above command example, replace "abcd.test.com" with the actual domain name.

- A UID field is populated with a value of `Nutanix`. This can be useful when configuring a Nutanix group for access control within a key management server, since it is based on fields within the client certificates.

Note: Some vendors when doing client certificate authentication expect the client username to be a field in the CSR. While the CN and UID are pre-generated, many of the user populated fields can be used instead if desired. If a node-unique field such as CN is chosen, users must be created on a per node basis for access control. If a cluster-unique field is chosen, customers must create a user for each cluster.

3. Send the CSRs to a certificate authority (CA) and get them signed.

» Safenet

The SafeNet KeySecure key management server includes a local CA option to generate signed certificates, or you can use other third-party vendors to create the signed certificates.

To enable FIPS compliance, add user nutanix to the CA that signed the CSR. Under **Security > High Security > FIPS Compliance** click **Set FIPS Compliant**.

Note: Some CAs strip the UID field when returning a signed certificate.

To comply with FIPS, Nutanix does not support the creation of global keys.

In the SafeNet KeySecure management console, go to **Device > Key Server > Key Server > KMIP Properties > Authentication Settings**.

Then do the following:

- Set the **Username Field in Client Certificate** option to **UID (User ID)**.
- Set the **Client Certificate Authentication** option to **Used for SSL session and username**.

If you do not perform these settings, the KMS creates global keys and fails to encrypt the clusters or containers using the software only method.

4. Upload the signed SSL certificates (one for each node) and the certificate for the CA to the cluster. These certificates are used to authenticate with the key management server.

5. Generate keys (KEKs) for the SED drives and upload those keys to the key management server.

Configuring Data-at-Rest Encryption (SEDs)

Nutanix offers an option to use self-encrypting drives (SEDs) to store data in a cluster. When SEDs are used, there are several configuration steps that must be performed to support data-at-rest encryption in the cluster.

Before you begin

A separate key management server is required to store the keys outside of the cluster. Each key management server device must be configured and addressable through the network. It is recommended that multiple key manager server devices be configured to work in clustered mode so they can be added to the cluster configuration as a single entity (see step 5) that is resilient to a single failure.

About this task

To configure cluster encryption, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page. The **Data at Rest Encryption** dialog box appears. Initially, encryption is not configured, and a message to that effect appears.

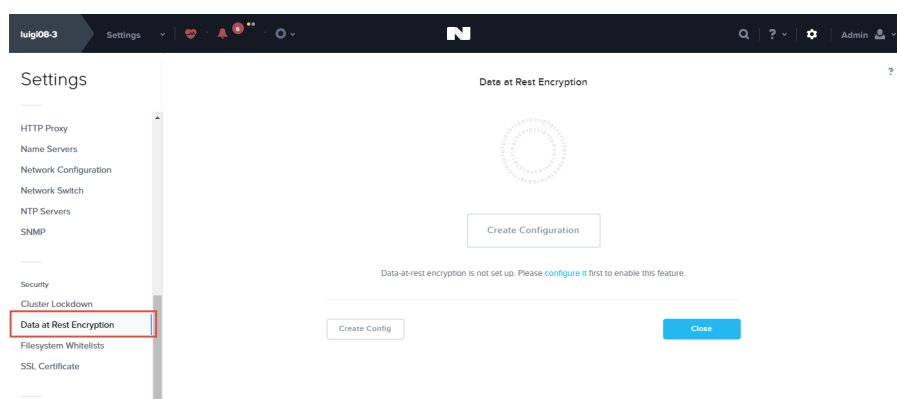


Figure 37: Data at Rest Encryption Screen (initial)

2. Click the **Create Configuration** button.
Clicking the **Continue Configuration** button, **configure it** link, or **Edit Config** button does the same thing, which is display the Data-at-Rest Encryption configuration page.
3. Select the Encryption Type as **Drive-based Encryption**. This option is displayed only when SEDs are detected.

4. In the **Certificate Signing Request Information** section, do the following:

Data-at-Rest Encryption

Certificate Signing Request Information

Enter the following information to generate the Certificate Signing Requests for the cluster

EMAIL ginger@nutanix.com	ORGANIZATION ntnx	ORGANIZATIONAL UNIT OU
COUNTRY CODE US	CITY la	STATE CA

Save CSR Info Download CSRs

Figure 38: Certificate Signing Request Section

- a. Enter appropriate credentials for your organization in the **Email**, **Organization**, **Organizational Unit**, **Country Code**, **City**, and **State** fields and then click the **Save CSR Info** button. The entered information is saved and is used when creating a certificate signing request (CSR). To specify more than one **Organization Unit** name, enter a comma separated list.

Note: You can update this information until an SSL certificate for a node is uploaded to the cluster, at which point the information cannot be changed (the fields become read only) without first deleting the uploaded certificates.

- b. Click the **Download CSRs** button, and then in the new screen click the **Download CSRs for all nodes** to download a file with CSRs for all the nodes or click a **Download** link to download a file with the CSR for that node.

Data-at-Rest Encryption

Certificate Signing Requests

Download the CSRs for the Nodes

Download CSRs for all nodes

Node Address	Action
10.4.36.142	Download
10.4.36.143	Download
10.4.36.144	Download

Back

Figure 39: Download CSRs Screen

- c. Send the files with the CSRs to the desired certificate authority. The certificate authority creates the signed certificates and returns them to you. Store the returned SSL certificates and the CA certificate where you can retrieve them for step 6.
- The certificates must be X.509 format. (DER, PKCS, and PFX formats are not supported.)
 - The certificate and the private key should be in separate files.

5. In the **Key Management Server** section, do the following:

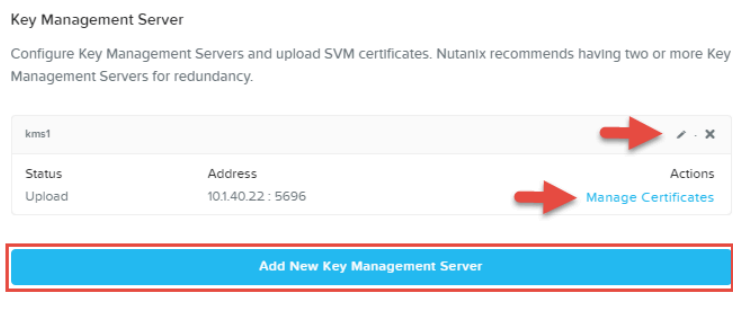


Figure 40: Key Management Server Section

- Click the **Add New Key Management Server** button.
- In the **Add a New Key Management Server** screen, enter a name, IP address, and port number for the key management server in the appropriate fields.

The port is where the key management server is configured to listen for the KMIP protocol. The default port number is 5696. For the complete list of required ports, see [Port Reference](#).

 - » If you have configured multiple key management servers in cluster mode, click the **Add Address** button to provide the addresses for each key management server device in the cluster.
 - » If you have stand-alone key management servers, click the **Save** button. Repeat this step (**Add New Key Management Server** button) for each key management server device to add.

Note: If your key management servers are configured into a leader/follower (active/passive) relationship and the architecture is such that the follower cannot accept write requests, do not add the follower into this configuration. The system sends requests (read or write) to any configured

key management server, so both read and write access is needed for key management servers added here.

Note: To prevent potential configuration problems, always use the **Add Address** button for key management servers configured into cluster mode. Only a stand-alone key management server should be added as a new server.

Data-at-Rest Encryption ? X

Add a New Key Management Server

Enter a name and at least one address for the Key Management Server.

Add Address

NAME

ADDRESS PORT 5696

Back Save

Figure 41: Add Key Management Server Screen

- c. To edit any settings, click the pencil icon for that entry in the key management server list to redisplay the add page and then click the **Save** button after making the change. To delete an entry, click the X icon.
6. In the **Add a New Certificate Authority** section, enter a name for the CA, click the **Upload CA Certificate** button, and select the certificate for the CA used to sign your node certificates (see step 4c). Repeat this step for all CAs that were used in the signing process.

Data-at-Rest Encryption ? X

Add a New Certificate Authority

Upload the Certificate Authority (CA) and enter a name for the CA certificate.

Upload CA Certificate

CERTIFICATE AUTHORITY NAME

Back Save

Figure 42: Certificate Authority Section

7. Go to the **Key Management Server** section (see step 5) and do the following:

- Click the **Manage Certificates** button for a key management server.
- In the **Manage Signed Certificates** screen, upload the node certificates either by clicking the **Upload Files** button to upload all the certificates in one step or by clicking the **Upload** link (not shown in the figure) for each node individually.
- Test that the certificates are correct either by clicking the **Test all nodes** button to test the certificates for all nodes in one step or by clicking the **Test CS** (or **Re-Test CS**) link for each node individually. A status of **Verified** indicates the test was successful for that node.

Note: Before removing a drive or node from an SED cluster, ensure that the testing is successful and the status is **Verified**. Otherwise, the drive or node will be locked.

- Repeat this step for each key management server.

Note: Before removing a drive or node from an SED cluster, ensure that the testing is successful and the status is **Verified**. Otherwise, the drive or node will be locked.

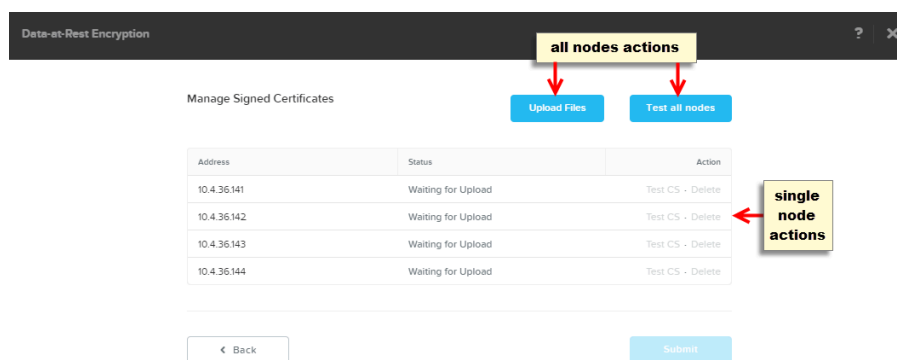


Figure 43: Upload Signed Certificates Screen

- When the configuration is complete, click the **Protect** button on the opening page to enable encryption protection for the cluster.

A clear key icon appears on the page.

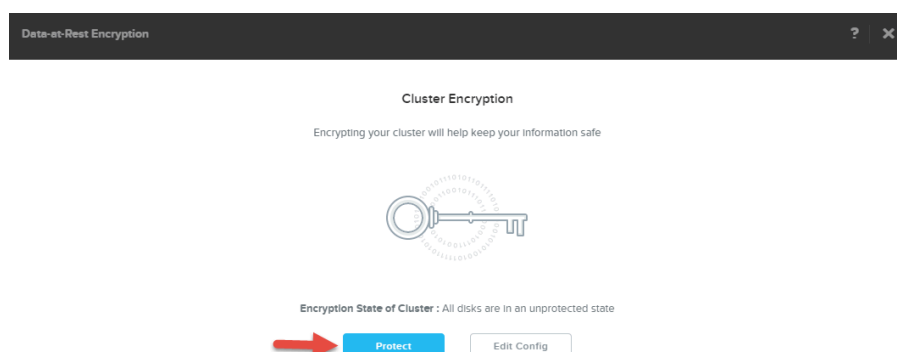


Figure 44: Data-at-Rest Encryption Screen (unprotected)

The key turns gold when cluster encryption is enabled.

Note: If changes are made to the configuration after protection has been enabled, such as adding a new key management server, you must rekey the disks for the modification to take full effect (see [Changing Key Encryption Keys \(SEDs\)](#) on page 109).

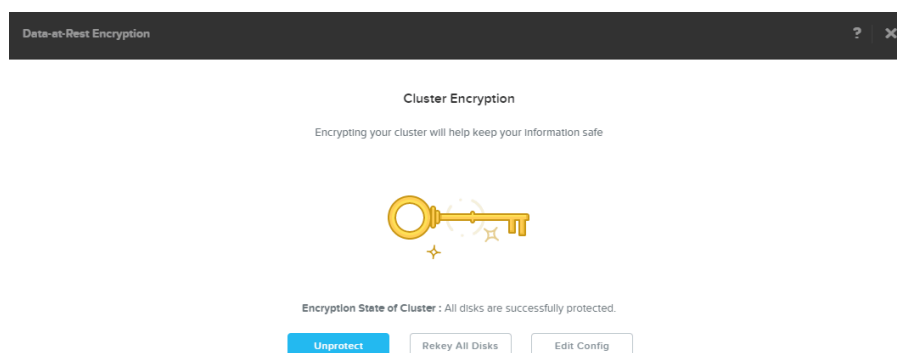


Figure 45: Data-at-Rest Encryption Screen (protected)

Enabling/Disabling Encryption (SEDs)

Data on a self encrypting drive (SED) is always encrypted, but enabling/disabling data-at-rest encryption for the cluster determines whether a separate (and secured) key is required to access that data.

About this task

To enable or disable data-at-rest encryption after it has been configured for the cluster (see [Configuring Data-at-Rest Encryption \(SEDs\)](#) on page 102), do the following:

Note: The key management server must be accessible to disable encryption.

Procedure

- Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page.

2. In the **Cluster Encryption** page, do one of the following:

- » If cluster encryption is enabled currently, click the **Unprotect** button to disable it.
- » If cluster encryption is disabled currently, click the **Protect** button to enable it.

Enabling cluster encryption enforces the use of secured keys to access data on the SEDs in the cluster; disabling cluster encryption means the data can be accessed without providing a key.

Changing Key Encryption Keys (SEDs)

The key encryption key (KEK) can be changed at any time. This can be useful as a periodic password rotation security precaution or when a key management server or node becomes compromised. If the key management server is compromised, only the KEK needs to be changed, because the KEK is independent of the drive encryption key (DEK). There is no need to re-encrypt any data, just to re-encrypt the DEK.

About this task

To change the KEKs for a cluster, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page.
2. In the **Cluster Encryption** page, select **Manage Keys** and click the **Rekey All Disks** button under **Hardware Encryption**.

Rekeying a cluster under heavy workloads may result in higher-than-normal IO latency, and some data may become temporarily unavailable. To continue with the rekey operation, click **Confirm Rekey**.

This step resets the KEKs for all the self encrypting disks in the cluster.

Note:

- The **Rekey All Disks** button appears only when cluster protection is active.
- If the cluster is already protected and a new key management server is added, you must press the **Rekey All Disks** button to use this new key management server for storing secrets.

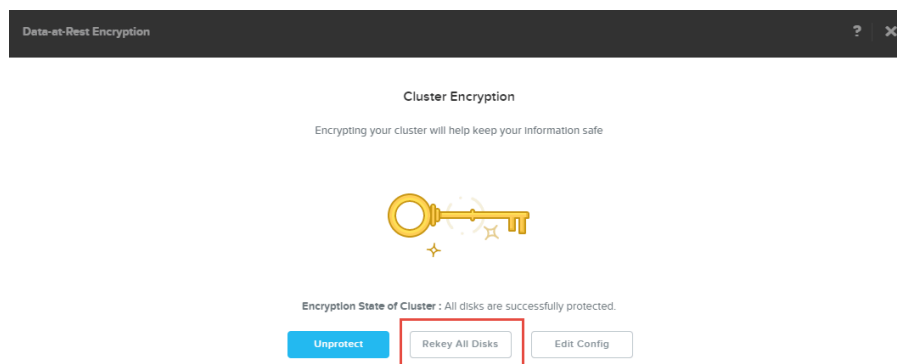


Figure 46: Cluster Encryption Screen

Destroying Data (SEDs)

Data on a self encrypting drive (SED) is always encrypted, and the data encryption key (DEK) used to read the encrypted data is known only to the drive controller. All data on the drive can effectively be destroyed

(that is, become permanently unreadable) by having the controller change the DEK. This is known as a crypto-erase.

About this task

To crypto-erase a SED, do the following:

Procedure

1. In the web console, go to the Hardware dashboard and select the **Diagram** tab.
2. Select the target disk in the diagram (upper section of screen) and then click the **Remove Disk** button (at the bottom right of the following diagram).

As part of the disk removal process, the DEK for that disk is automatically cycled on the drive controller. The previous DEK is lost and all new disk reads are indecipherable. The key encryption key (KEK) is unchanged, and the new DEK is protected using the current KEK.

Note:

- When a node is removed, all SEDs in that node are crypto-erased automatically as part of the node removal process.
- When you run the cluster destroy command to decommission your entire cluster, the command automatically performs a crypto erase on the SED as the final step.

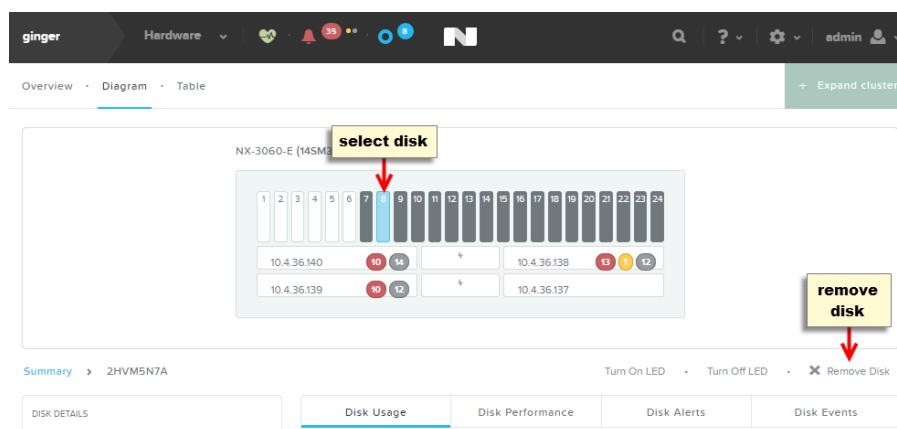


Figure 47: Removing a Disk

Data-at-Rest Encryption (Software Only)

For customers who require enhanced data security, Nutanix provides a software-only encryption option for data-at-rest security (SEDs not required) included in the Ultimate license.

Note: On G6 platforms running the AOS Pro license, you can use software encryption by installing an add-on license.

Software encryption using a local key manager (LKM) supports the following features:

- For AHV, the data can be encrypted on a cluster level. This is applicable to an empty cluster or a cluster with existing data.

- For ESXi and Hyper-V, the data can be encrypted on a cluster or container level. The cluster or container can be empty or contain existing data. Consider the following points for container level encryption.
 - Once you enable container level encryption, you can not change the encryption type to cluster level encryption later.
 - After the encryption is enabled, the administrator needs to enable encryption for every new container.
- Data is encrypted at all times.
- Data is inaccessible in the event of drive or node theft.
- Data on a drive can be securely destroyed.
- Re-key of the leader encryption key at arbitrary times is supported.
- Cluster's native KMS is supported.

Note: In case of mixed hypervisors, only the following combinations are supported.

- ESXi and AHV
- Hyper-V and AHV

Note: This solution provides enhanced security for data on a drive, but it does not secure data in transit.

Data Encryption Model

To accomplish the above mentioned goals, Nutanix implements a data security configuration that uses AOS functionality along with the cluster's native or an external key management server. Nutanix uses open standards (KMIP protocols) for interoperability and strong security.




Figure 48: Cluster Protection Overview

This configuration involves the following workflow:

- For software encryption, data protection must be enabled for the cluster before any data is encrypted. Also, the Controller VM must provide the proper key to access the data.

- A symmetric data encryption key (DEK) such as AES 256 is applied to all data being written to or read from the disk. The key is known only to AOS, so there is no way to access the data directly from the drive.
- In case of an external KMS:
Each node maintains a set of certificates and keys in order to establish a secure connection with the key management server.
Only one key management server device is required, but it is recommended that multiple devices are employed so the key management server is not a potential single point of failure. Configure the key manager server devices to work in clustered mode so they can be added to the cluster configuration as a single entity that is resilient to a single failure.

Configuring Data-at-Rest Encryption (Software Only)

Nutanix offers a software-only option to perform data-at-rest encryption in a cluster or container.

Before you begin

- Nutanix provides the option to choose the KMS type as the Native KMS (local), Native KMS (remote), or External KMS.
- Cluster Localised Key Management Service (Native KMS (local)) requires a minimum of 3-node cluster. 1-node and 2-node clusters are not supported.
- Software encryption using Native KMS is supported for remote office/branch office (ROBO) deployments using the Native KMS (remote) KMS type.
- For external KMS, a separate key management server is required to store the keys outside of the cluster. Each key management server device must be configured and addressable through the network. It is recommended that multiple key manager server devices be configured to work in clustered mode so they can be added to the cluster configuration as a single entity that is resilient to a single failure.

Caution: DO NOT HOST A KEY MANAGEMENT SERVER VM ON THE ENCRYPTED CLUSTER THAT IS USING IT!!

Doing so could result in complete data loss if there is a problem with the VM while it is hosted in that cluster.

Note: You must install the license of the external key manager for all nodes in the cluster. See [Compatibility and Interoperability Matrix](#) for a complete list of the supported key management servers. For instructions on how to configure a key management server, refer to the documentation from the appropriate vendor.

- This feature requires an Ultimate license, or as an Add-On to the PRO license (for the latest generation of products). Ensure that you have procured the add-on license key to use the data-at-rest encryption using AOS, contact Sales team to procure the license.

Caution: For security, you can't disable software-only data-at-rest encryption once it is enabled.

About this task

To configure cluster or container encryption, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page. The **Data at Rest Encryption** dialog box appears. Initially, encryption is not configured, and a message to that effect appears.

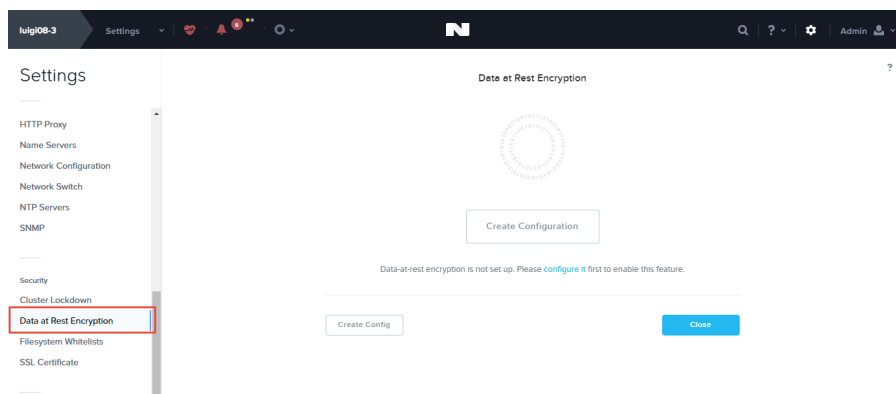


Figure 49: Data at Rest Encryption Screen (initial)

2. Click the **Create Configuration** button.
Clicking the **Continue Configuration** button, **configure it** link, or **Edit Config** button does the same thing, which is display the Data-at-Rest Encryption configuration page

3. Select the Encryption Type as **Encrypt the entire cluster** or **Encrypt storage containers**. Then click **Save Encryption Type**.

Caution: You can enable encryption for the entire cluster or just the container. However, if you enable encryption on a container; and there are any encryption key issue like loss of encryption key, you can encounter the following:

- The entire cluster data is affected, not just the encrypted container.
- All the user VMs of the cluster will not be able to access the data.

The hardware option is displayed only when SEDs are detected. Else, software based encryption type will be used by default.

Data-at-Rest Encryption

Encryption Type

☒ Encrypt the entire cluster
All data in the cluster will be encrypted. Once the cluster has been encrypted, you will no longer be able to encrypt individual storage containers.

☐ Encrypt storage containers
To encrypt data you'll need to create new encrypted storage containers. Only data in encrypted storage containers will be encrypted.

Save Encryption Type

Figure 50: Select encryption type

Note: For ESXi and Hyper-V, the data can be encrypted on a cluster or container level. The cluster or container can be empty or contain existing data. Consider the following points for container level encryption.

- Once you enable container level encryption, you can not change the encryption type to cluster level encryption later.
- After the encryption is enabled, the administrator needs to enable encryption for every new container.

To enable encryption for every new storage container, do the following:

- a. In the web console, select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Container** tabs.
- b. To enable encryption, select the target storage container and then click the **Update** link. The **Update Storage Container** window appears.
- c. In the **Advanced Settings** area, select the **Enable** check box to enable encryption for the storage container you selected.

Update Storage Container

STORAGE POOL

default-storage-pool

MAX CAPACITY

15.24 TiB

NFS DATASTORE

☒ Mount on all ESXi hosts

☐ Unmount on all ESXi hosts

☐ Mount/Unmount on the following ESXi hosts

Advanced Settings

ENCRYPTION

☒ Enable

Figure 51: Update storage container

- d. Click **Save** to complete.

4. Select the Key Management Service.

To keep the keys safe with the native KMS, select **Native KMS (local)** or **Native KMS (remote)** and click **Save KMS type**. If you select this option, skip to step 9 to complete the configuration.

Note:

- Cluster Localised Key Management Service (**Native KMS (local)**) requires a minimum of 3-node cluster. 1-node and 2-node clusters are not supported.
- For enhanced security of ROBO environments (typically, 1 or 2 node clusters), select the **Native KMS (remote)** for software based encryption of ROBO clusters managed by Prism Central.

Note: This option is available only if the cluster is registered to Prism Central.

For external KMS type, select the **External KMS** option and click **Save KMS type**. Continue to step 5 for further configuration.

Select Key Management Server (KMS)

The KMS manages the encryption keys used to encrypt data.

☐ Native KMS (local)

Cluster Localised Key Management Service is not possible with 2 or less nodes

☒ Native KMS (remote)

Keep your keys secured with a registered remote PC. This can be used for software encryption with any cluster size.

☐ An external KMS

Configure Key Management Servers and upload SVM certificates. You will manually download and upload certificates to validate the KMS. Nutanix recommends having two or more Key Management Servers for redundancy.

Figure 52: Select KMS Type

Note: You can switch between the KMS types at a later stage if the specific KMS prerequisites are met, see [Switching between Native Key Manager and External Key Manager](#) on page 121.

5. In the **Certificate Signing Request Information** section, do the following:

Data-at-Rest Encryption

Certificate Signing Request Information

Enter the following information to generate the Certificate Signing Requests for the cluster

EMAIL ginger@nutanix.com	ORGANIZATION ntnx	ORGANIZATIONAL UNIT OU
COUNTRY CODE US	CITY la	STATE CA

Save CSR Info Download CSRs

Figure 53: Certificate Signing Request Section

- a. Enter appropriate credentials for your organization in the **Email**, **Organization**, **Organizational Unit**, **Country Code**, **City**, and **State** fields and then click the **Save CSR Info** button.
- The entered information is saved and is used when creating a certificate signing request (CSR). To specify more than one **Organization Unit** name, enter a comma separated list.

Note: You can update this information until an SSL certificate for a node is uploaded to the cluster, at which point the information cannot be changed (the fields become read only) without first deleting the uploaded certificates.

- b. Click the **Download CSRs** button, and then in the new screen click the **Download CSRs for all nodes** to download a file with CSRs for all the nodes or click a **Download** link to download a file with the CSR for that node.

Data-at-Rest Encryption

Certificate Signing Requests

Download the CSRs for the Nodes

Download CSRs for all nodes

Node Address	Action
10.4.36.142	Download
10.4.36.143	Download
10.4.36.144	Download

Back

Figure 54: Download CSRs Screen

- c. Send the files with the CSRs to the desired certificate authority.
- The certificate authority creates the signed certificates and returns them to you. Store the returned SSL certificates and the CA certificate where you can retrieve them in step 5.
- The certificates must be X.509 format. (DER, PKCS, and PFX formats are not supported.)
 - The certificate and the private key should be in separate files.

6. In the **Key Management Server** section, do the following:

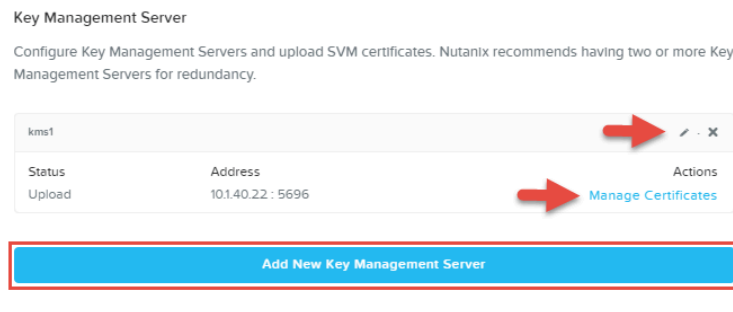


Figure 55: Key Management Server Section

- a. Click the **Add New Key Management Server** button.
- b. In the **Add a New Key Management Server** screen, enter a name, IP address, and port number for the key management server in the appropriate fields.

The port is where the key management server is configured to listen for the KMIP protocol. The default port number is 5696. For the complete list of required ports, see [Port Reference](#).

 - » If you have configured multiple key management servers in cluster mode, click the **Add Address** button to provide the addresses for each key management server device in the cluster.
 - » If you have stand-alone key management servers, click the **Save** button. Repeat this step (**Add New Key Management Server** button) for each key management server device to add.

Note: If your key management servers are configured into a master/slave (active/passive) relationship and the architecture is such that the follower cannot accept write requests, do not add the follower into this configuration. The system sends requests (read or write) to

any configured key management server, so both read and write access is needed for key management servers added here.

Note: To prevent potential configuration problems, always use the **Add Address** button for key management servers configured into cluster mode. Only a stand-alone key management server should be added as a new server.

Data-at-Rest Encryption

Add a New Key Management Server

Enter a name and at least one address for the Key Management Server.

NAME

ADDRESS

PORT
5696

< Back

Save

Figure 56: Add Key Management Server Screen

- c. To edit any settings, click the pencil icon for that entry in the key management server list to redisplay the add page and then click the **Save** button after making the change. To delete an entry, click the X icon.
7. In the **Add a New Certificate Authority** section, enter a name for the CA, click the **Upload CA Certificate** button, and select the certificate for the CA used to sign your node certificates (see step 3c). Repeat this step for all CAs that were used in the signing process.

Data-at-Rest Encryption

Add a New Certificate Authority

Upload the Certificate Authority (CA) and enter a name for the CA certificate.

CERTIFICATE AUTHORITY NAME

< Back

Save

Figure 57: Certificate Authority Section

8. Go to the **Key Management Server** section (see step 4) and do the following:
 - a. Click the **Manage Certificates** button for a key management server.
 - b. In the **Manage Signed Certificates** screen, upload the node certificates either by clicking the **Upload Files** button to upload all the certificates in one step or by clicking the **Upload** link (not shown in the figure) for each node individually.
 - c. Test that the certificates are correct either by clicking the **Test all nodes** button to test the certificates for all nodes in one step or by clicking the **Test CS** (or **Re-Test CS**) link for each node individually. A status of **Verified** indicates the test was successful for that node.
 - d. Repeat this step for each key management server.

Note: Before removing a drive or node from an SED cluster, ensure that the testing is successful and the status is **Verified**. Otherwise, the drive or node will be locked.

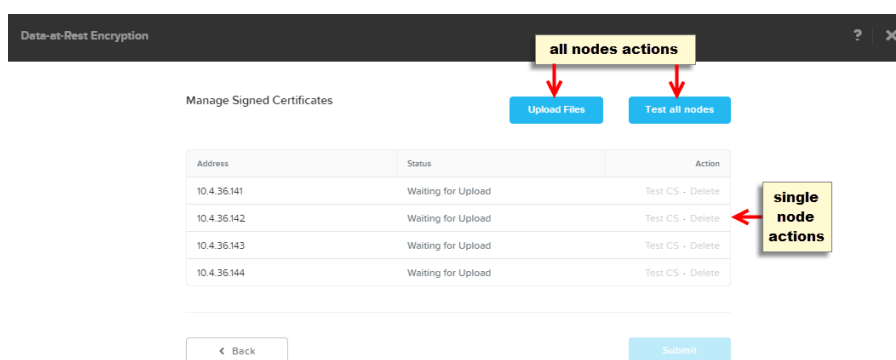


Figure 58: Upload Signed Certificates Screen

9. When the configuration is complete, click the **Enable Encryption** button. Enable Encryption window is displayed.



Figure 59: Data-at-Rest Encryption Screen (unprotected)

Caution: To help ensure that your data is secure, you cannot disable software-only data-at-rest encryption once it is enabled. Nutanix recommends regularly backing up your data, encryption keys, and key management server.

10. Enter **ENCRYPT**.

11. Click **Encrypt** button.

The data-at-rest encryption is enabled. To view the status of the encrypted cluster or container, go to **Data at Rest Encryption** in the **Settings** menu.

When you enable encryption, a low priority background task runs to encrypt all the unencrypted data. This task is designed to take advantage of any available CPU space to encrypt the unencrypted data within a reasonable time. If the system is occupied with other workloads, the background task consumes less CPU space. Depending on the amount of data in the cluster, the background task can take 24 to 36 hours to complete.

Note: If changes are made to the configuration after protection has been enabled, such as adding a new key management server, you must do the rekey operation for the modification to take full effect. In case of EKM, rekey to change the KEKs stored in the EKM. In case of LKM, rekey to change the leader key used by native key manager, see [Changing Key Encryption Keys \(Software Only\)](#) on page 122) for details.

Note: Once the task to encrypt a cluster begins, you cannot cancel the operation. Even if you stop and restart the cluster, the system resumes the operation.

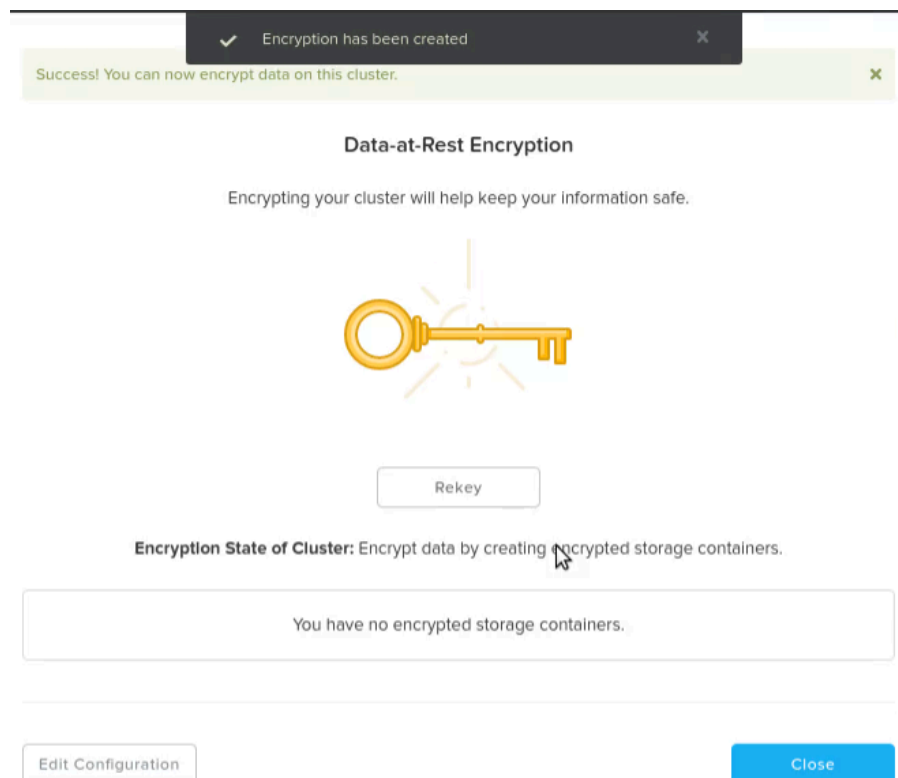


Figure 60: Data-at-Rest Encryption Screen (protected)

Switching between Native Key Manager and External Key Manager

After Software Encryption has been established, Nutanix supports the ability to switch the KMS type from the External Key Manager to the Native Key Manager or from the Native Key Manager to an External Key Manager, without any down time.

Note:

- The Native KMS requires a minimum of 3-node cluster.

- For external KMS, a separate key management server is required to store the keys outside of the cluster. Each key management server device must be configured and addressable through the network. It is recommended that multiple key manager server devices be configured to work in clustered mode so they can be added to the cluster configuration as a single entity that is resilient to a single failure.
- It is recommended that you backup and save the encryption keys with identifiable names before and after changing the KMS type. For backing up keys, see [Backing up Keys](#) on page 125.

To change the KMS type, change the KMS selection by editing the encryption configuration. For details, see step 4 in [Configuring Data-at-Rest Encryption \(Software Only\)](#) on page 112 section.

Select Key Management Server (KMS)

The KMS manages the encryption keys used to encrypt data.

☐ Native KMS (local)

Cluster Localised Key Management Service is not possible with 2 or less nodes

☒ Native KMS (remote)

Keep your keys secured with a registered remote PC. This can be used for software encryption with any cluster size.

☐ An external KMS

Configure Key Management Servers and upload SVM certificates. You will manually download and upload certificates to validate the KMS. Nutanix recommends having two or more Key Management Servers for redundancy.

Figure 61: Select KMS type

Note: This operation completes in a few minutes, depending on the number of encrypted objects and network speed.

Changing Key Encryption Keys (Software Only)

The key encryption key (KEK) can be changed at any time. This can be useful as a periodic password rotation security precaution or when a key management server or node becomes compromised. If the key management server is compromised, only the KEK needs to be changed, because the KEK is independent of the drive encryption key (DEK). There is no need to re-encrypt any data, just to re-encrypt the DEK.

About this task

To change the KEKs for a cluster, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page.

2. In the **Cluster Encryption** page, select **Manage Keys** and click the **Rekey** button under **Software Encryption**.

Note: The **Rekey** button appears only when cluster protection is active.

Note: If the cluster is already protected and a new key management server is added, you must press the **Rekey** button to use this new key management server for storing secrets.



Figure 62: Cluster Encryption Screen

Note: The system automatically regenerates the leader key yearly.

Destroying Data (Software Only)

Data on the AOS cluster is always encrypted, and the data encryption key (DEK) used to read the encrypted data is known only to the AOS. All data on the drive can effectively be destroyed (that is, become permanently unreadable) by deleting the container or cluster. This is known as a crypto-erase.

About this task

Note: To help ensure that your data is secure, you cannot disable software-only data-at-rest encryption once it is enabled. Nutanix recommends regularly backing up your data, encryption keys, and key management server.

To crypto-erase the container or cluster, do the following:

Procedure

1. Delete the storage container or destroy the cluster.

- For information on how to delete a storage container, see [Modifying a Storage Container](#) in the *Prism Element Web Console Guide*.
- For information on how to destroy a cluster, see [Destroying a Cluster](#) in the *Acropolis Advanced Administration Guide*.

Note:

When you delete a storage container, the Curator scans and deletes the DEK and KEK keys automatically.

When you destroy a cluster, then:

- the Native Key Manager (local) destroys the master key shares and the encrypted DEKs/KEKs.
- the Native Key Manager (remote) retains the root key on the PC if the cluster is still registered to a PC when it is destroyed. You must unregister a cluster from the PC and then destroy the cluster to delete the root key.
- the External Key Manager deletes the encrypted DEKs. However, the KEKs remain on the EKM. You must use an external key manager UI to delete the KEKs.

2. Delete the key backup files, if any.

Switching from SED-EKM to Software-LKM

This section describes the steps to switch from SED and External KMS combination to software-only and LKM combination.

About this task

To switch from SED-EKM to Software-LKM, do the following.

Procedure

1. Perform the steps for the software-only encryption with External KMS. For more information, see [Configuring Data-at-Rest Encryption \(Software Only\)](#) on page 112.

After the background task completes, all the data gets encrypted by the software. The time taken to complete the task depends on the amount of data and foreground I/O operations in the cluster.

2. Disable the SED encryption. Ensure that all the disks are unprotected.

For more information, see [Enabling/Disabling Encryption \(SEDs\)](#) on page 108.

3. Switch the key management server from the External KMS to Local Key Manager. For more information, see [Switching between Native Key Manager and External Key Manager](#) on page 121.

Configuring Dual Encryption

About this task

Dual Encryption protects the data on the clusters using both SED and software-only encryption. An external key manager is used to store the keys for dual encryption, the Native KMS is not supported.

To configure dual encryption, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page.
2. In the Cluster Encryption page, check to enable both **Drive-based** and **Software-based** encryption.
3. Click **Save Encryption Type**.

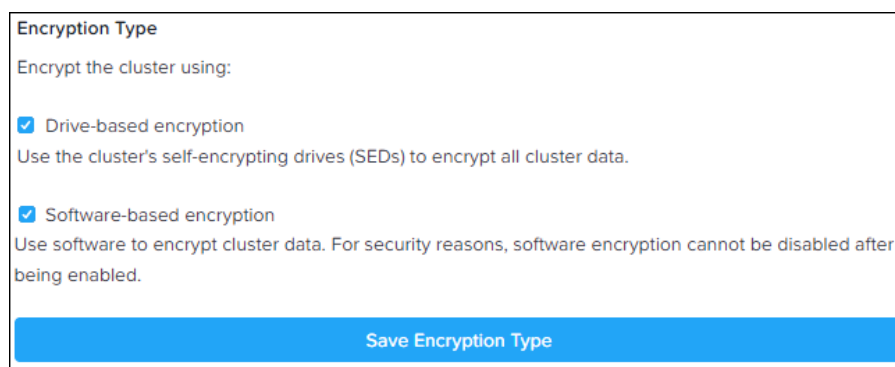


Figure 63: Dual Encryption

4. Continue with the rest of the encryption configuration, see:
 - [Configuring Data-at-Rest Encryption \(Software Only\)](#) on page 112
 - [Configuring Data-at-Rest Encryption \(SEDs\)](#) on page 102

Backing up Keys

About this task

You can take a backup of encryption keys:

- when you enable Software-only Encryption for the first time
- after you regenerate the keys

Backing up encryption keys is critical in the very unlikely situation in which keys get corrupted.

You can download key backup file for a cluster on a PE or all clusters on a PC. To download key backup file for all clusters, see [Taking a Consolidated Backup of Keys \(Prism Central\)](#).

To download the key backup file for a cluster, do the following:

Procedure

1. Log on to the Prism Element web console.
2. Click the gear icon in the main menu and then select **Data at Rest Encryption** in the **Settings** page.
3. In the Cluster Encryption page, select **Manage Keys**.
4. Enter and confirm the password.

5. Click the **Download Key Backup** button.

The backup file is saved in the default download location on your local machine.

Note: Ensure you move the backup key file to a safe location.

Tip: If vTPM is enabled, the downloaded key backup includes vTPM keys, in addition to the encryption keys.

Taking a Consolidated Backup of Keys (Prism Central)

If you are using the Native KMS option with software encryption for your clusters, you can take a consolidated backup of all the keys from Prism Central.

About this task

To take a consolidated backup of keys for software encryption-enabled clusters (Native KMS-only), do the following:

Procedure

1. Log on to the Prism Central web console.
2. Click the hamburger icon, then select **Clusters > List** view.
3. Select a cluster, go to **Actions**, then select **Manage & Backup Keys**.
4. Download the backup keys:
 - a. In **Password**, enter your password.
 - b. In **Confirm Password**, reenter your password.
 - c. To change the encryption key, select the **Rekey Encryption Key (KEK)** box .
 - d. To download the backup key, click **Backup Key**.

Note: Ensure that you move the backup key file to a safe location.

Tip: If vTPM is enabled, the downloaded key backup includes vTPM keys, in addition to the encryption keys.

Importing Keys

You can import the encryption keys from backup. You must note the specific commands in this topic if you backed up your keys to an external key manager (EKM)

About this task

Note: Nutanix recommends that you contact Nutanix Support for this operation. Extended cluster downtime might result if you perform this task incorrectly.

Procedure

1. Log on to any Controller VM in the cluster with SSH.

2. Retrieve the encryption keys stored on the cluster and verify that all the keys you want to retrieve are listed.

In this example, the password is Nutanix.123. *date* is the timestamp portion of the backup file name.

```
mantle_recovery_util --backup_file_path=/home/nutanix/encryption_key_backup_date \
--password=Nutanix.123 --list_key_ids=true
```

3. Import the keys into the cluster.

```
mantle_recovery_util --backup_file_path=/home/nutanix/key_backup \
--password=Nutanix.123 --interactive_mode
```

4. If you are using an external key manager such as IBM Security Key Lifecycle Manager, Gemalto Safenet, or Vormetric Data Security Manager, use the `--store_kek_remotely` option to import the keys into the cluster.

In this example, *date* is the timestamp portion of the backup file name.

```
mantle_recovery_util --backup_file_path path/encryption_key_backup_date \
--password key_password --store_kek_remotely
```

Tip: The imported key backup includes vTPM keys (if vTPM is enabled during backup), in addition to the encryption keys.

Securing Traffic Through Network Segmentation

Network segmentation enhances security, resilience, and cluster performance by isolating a subset of traffic to its own network.

You can achieve traffic isolation in one or more of the following ways:

Isolating Backplane Traffic by using VLANs (Logical Segmentation)

You can separate management traffic from storage replication (or backplane) traffic by creating a separate network segment (LAN) for storage replication. For more information about the types of traffic seen on the management plane and the backplane, see [Traffic Types In a Segmented Network](#) on page 128.

To enable the CVMs in a cluster to communicate over these separated networks, the CVMs are multihomed. Multihoming is facilitated by the addition of a virtual network interface card (vNIC) to the Controller VM and placing the new interface on the backplane network. Additionally, the hypervisor is assigned an interface on the backplane network.

The traffic associated with the CVM interfaces and host interfaces on the backplane network can be secured further by placing those interfaces on a separate VLAN.

In this type of segmentation, both network segments continue to use the same external bridge and therefore use the same set of physical uplinks. For physical separation, see [Physically Isolating the Backplane Traffic on an AHV Cluster](#) on page 158.

Isolating backplane traffic from management traffic requires minimal configuration through the Prism web console. No manual host (hypervisor) configuration steps are required.

For information about isolating backplane traffic, see [Isolating the Backplane Traffic Logically on an Existing Cluster \(VLAN-Based Segmentation Only\)](#) on page 138.

Isolating Backplane Traffic Physically (Physical Segmentation)

You can physically isolate the backplane traffic (intra cluster traffic) from the management traffic (Prism, SSH, SNMP) in to a separate vNIC on the CVM and using a dedicated virtual network that has its own physical NICs. This type of segmentation therefore offers true physical separation of the backplane traffic from the management traffic.

You can use Prism to configure the vNIC on the CVM and configure the backplane traffic to communicate over the dedicated virtual network. However, you must first manually configure the virtual network on the hosts and associate it with the physical NICs that it requires for true traffic isolation.

For more information about physically isolating backplane traffic, see [Physically Isolating the Backplane Traffic on an AHV Cluster](#) on page 158.

Isolating service-specific traffic

You can also secure traffic associated with a service (for example, Nutanix Volumes) by confining its traffic to a separate vNIC on the CVM and using a dedicated virtual network that has its own physical NICs. This type of segmentation therefore offers true physical separation for service-specific traffic.

You can use Prism to create the vNIC on the CVM and configure the service to communicate over the dedicated virtual network. However, you must first manually configure the virtual network on the hosts and associate it with the physical NICs that it requires for true traffic isolation. You need one virtual network for each service you want to isolate. For a list of the services whose traffic you can isolate in the current release, see [Cluster Services That Support Traffic Isolation](#) on page 136.

For information about isolating service-specific traffic, see [Isolating Service-Specific Traffic](#) on page 169.

Isolating Stargate-to-Stargate traffic over RDMA

Some Nutanix platforms support remote direct memory access (RDMA) for Stargate-to-Stargate service communication. You can create a separate virtual network for RDMA-enabled network interface cards. If a node has RDMA-enabled NICs, Foundation passes the NICs through to the CVMs during imaging. The CVMs use only the first of the two RDMA-enabled NICs for Stargate-to-Stargate communications. The virtual NIC on the CVM is named `rdma0`. Foundation does not configure the RDMA LAN. After creating a cluster, you need to enable RDMA by creating an RDMA LAN from the Prism web console. For more information about RDMA support, see [Remote Direct Memory Access](#) in the *NX Series Hardware Administration Guide*.

For information about isolating backplane traffic on an RDMA cluster, see [Isolating the Backplane Traffic on an Existing RDMA Cluster](#) on page 143.

Traffic Types In a Segmented Network

The traffic entering and leaving a Nutanix cluster can be broadly classified into the following types:

Backplane traffic

Backplane traffic is intra-cluster traffic that is necessary for the cluster to function, and it comprises traffic between CVMs and traffic between CVMs and hosts for functions such as storage RF replication, host management, high availability, and so on. This traffic uses `eth2` on the CVM. In AHV, VM live migration traffic is also backplane, and uses the AHV backplane interface, VLAN, and virtual switch when configured. For nodes that have RDMA-enabled NICs, the CVMs use a separate RDMA LAN for Stargate-to-Stargate communications.

Management traffic

Management traffic is administrative traffic, or traffic associated with Prism and SSH connections, remote logging, SNMP, and so on. The current implementation simplifies the definition of management traffic to be any traffic that is not on the backplane network, and therefore also includes communications between user VMs and CVMs. This traffic uses `eth0` on the CVM.

Traffic on the management plane can be further isolated per service or feature. An example of this type of traffic is the traffic that the cluster receives from external iSCSI initiators (Nutanix Volumes iSCSI traffic). For a list of services supported in the current release, see [Cluster Services That Support Traffic Isolation](#) on page 136.

Segmented and Unsegmented Networks

In the default unsegmented network in a Nutanix cluster (ESXi and AHV), the Controller VM has two virtual network interfaces—eth0 and eth1.

Interface eth0 is connected to the default external virtual switch, which is in turn connected to the external network through a bond or NIC team that contains the host physical uplinks.

Interface eth1 is connected to an internal network that enables the CVM to communicate with the hypervisor.

In the below unsegmented network (see figure *Unsegmented Network - ESXi Cluster*, and *Unsegmented Network - AHV Cluster*) all external CVM traffic, whether backplane or management traffic, uses interface eth0. These interfaces are on the default VLAN on the default virtual switch.

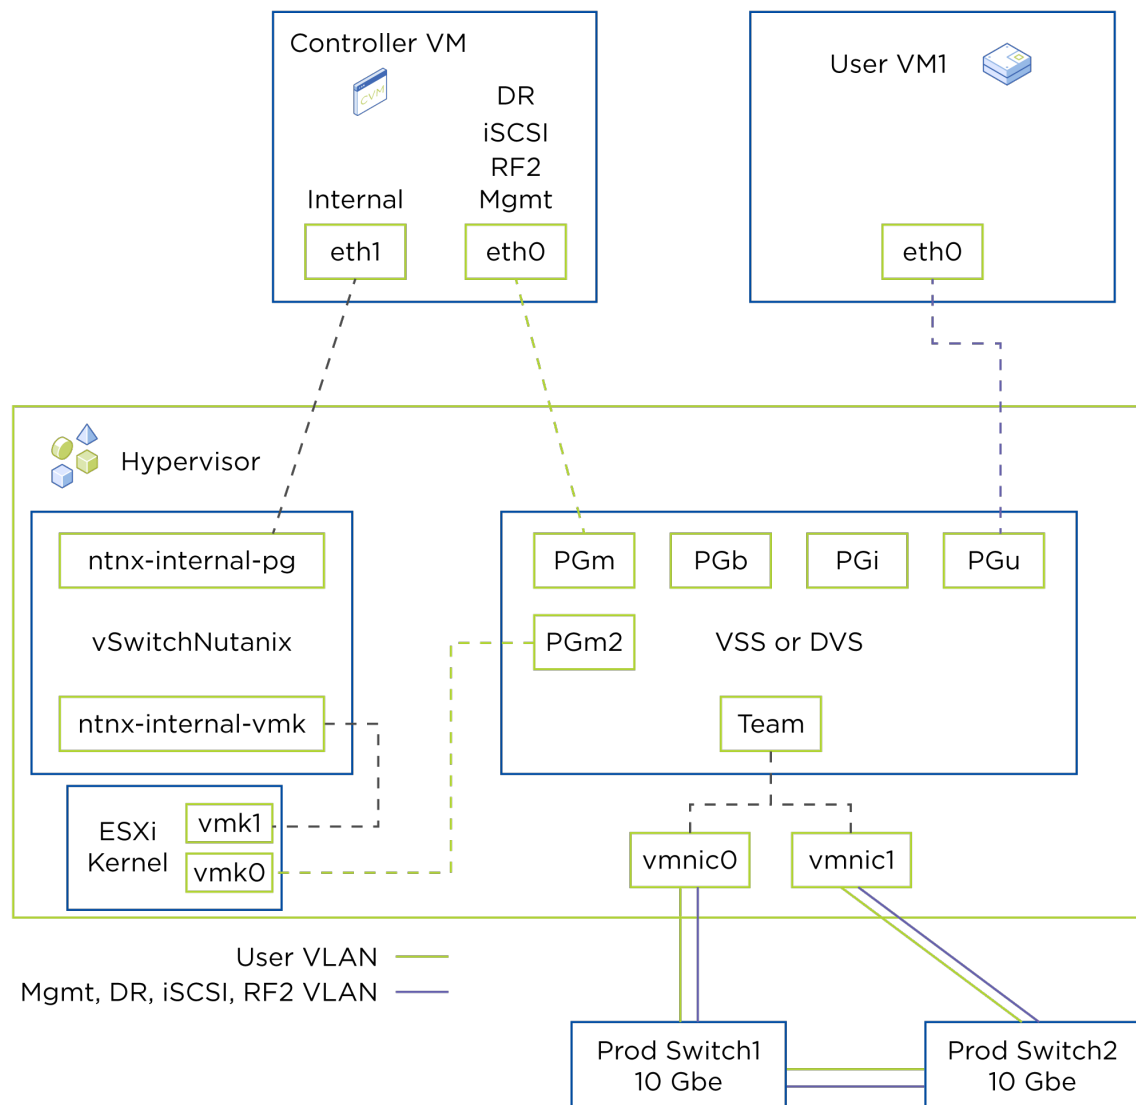


Figure 64: Unsegmented Network- ESXi Cluster

This figure shows an unsegmented network AHV cluster.

In AHV, VM live migration traffic is also backplane, and uses the AHV backplane interface, VLAN, and virtual switch when configured.

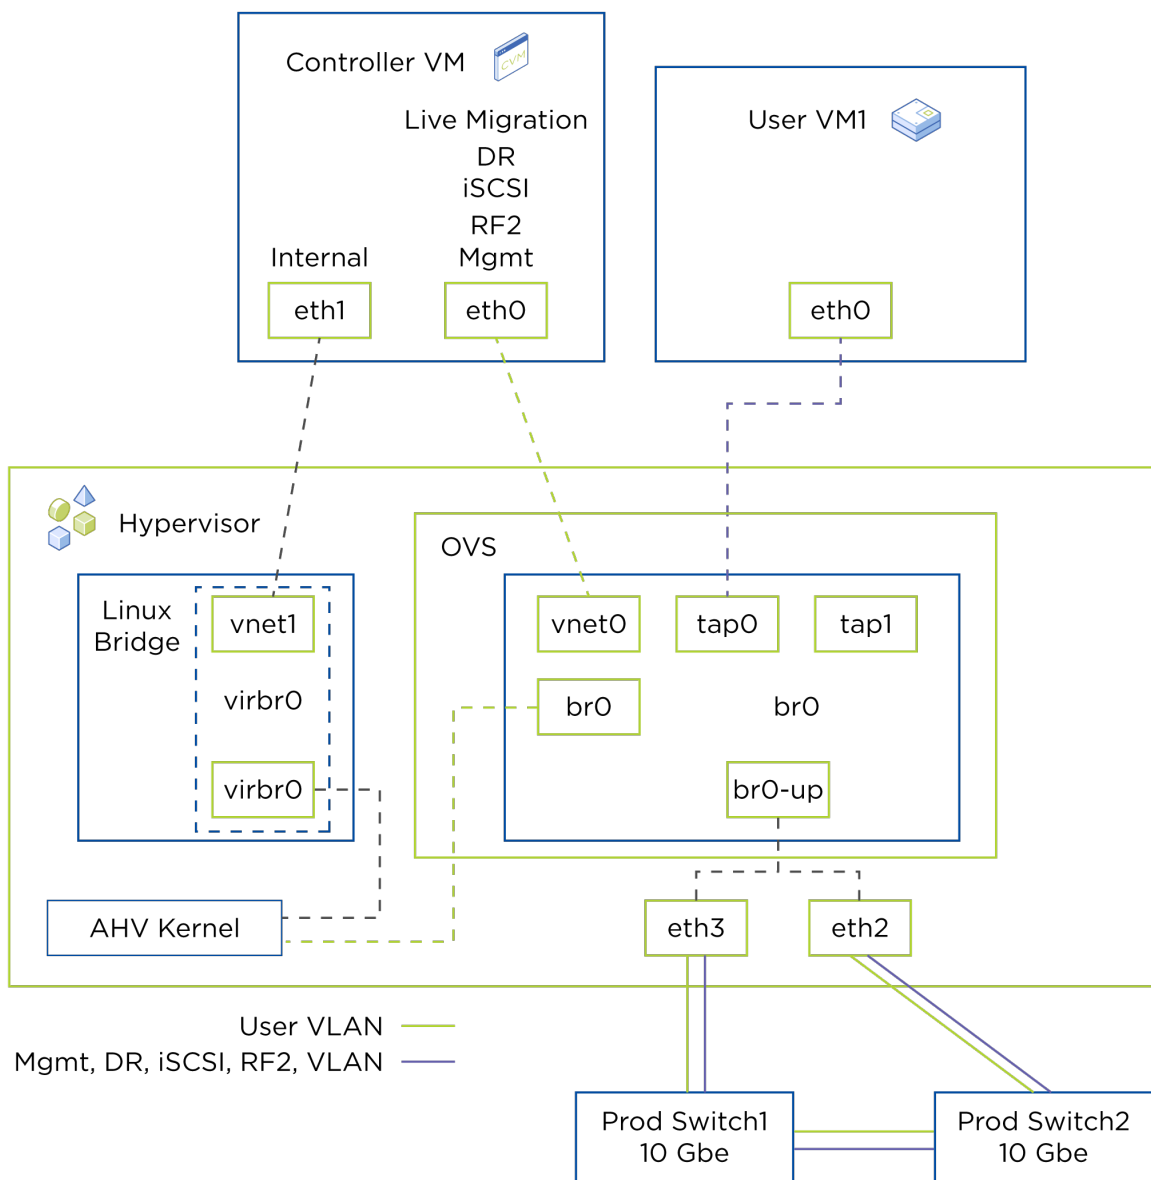


Figure 65: Unsegmented Network- AHV Cluster

If you further isolate service-specific traffic, additional vNICs are created on the CVM. Each service requiring isolation is assigned a dedicated virtual NIC on the CVM. The NICs are named **ntnx0**, **ntnx1**, and so on. Each service-specific NIC is placed on a configurable existing or new virtual network (vSwitch or bridge) and a VLAN and IP subnet are specified.

Network with Segmentation

In a segmented network, management traffic uses CVM interface **eth0** and additional services can be isolated to different VLANs or virtual switches. In backplane segmentation, the backplane traffic uses interface **eth2**. The backplane network uses either the default VLAN or, optionally, a separate VLAN that you specify when segmenting the network. In ESXi, you must select a port group for the new vmkernel

interface. In AHV this internal interface is created automatically in the selected virtual switch. For physical separation of the backplane network, create this new port group on a separate virtual switch in ESXi, or select the desired virtual switch in the AHV GUI.

If you want to isolate service-specific traffic such as Volumes or Disaster Recovery as well as backplane traffic, then additional vNICs are needed on the CVM, but no new vmkernel adapters or internal interfaces are required. AOS creates additional vNICs on the CVM. Each service that requires isolation is assigned a dedicated vNIC on the CVM. The NICs are named ntnx0, ntnx1, and so on. Each service-specific NIC is placed on a configurable existing or new virtual network (vSwitch or bridge) and a VLAN and IP subnet are specified.

You can choose to perform backplane segmentation alone, with no other forms of segmentation. You can also choose to use one or more types of service specific segmentation with or without backplane segmentation. In all of these cases, you can choose to segment any service to either the existing, or a new virtual switch for further physical traffic isolation. The combination selected is driven by the security and networking requirements of the deployment. In most cases, the default configuration with no segmentation of any kind is recommended due to simplicity and ease of deployment.

The following figure shows an implementation scenario where the backplane and service specific segmentation is configured with two vSwitches on ESXi hypervisors.

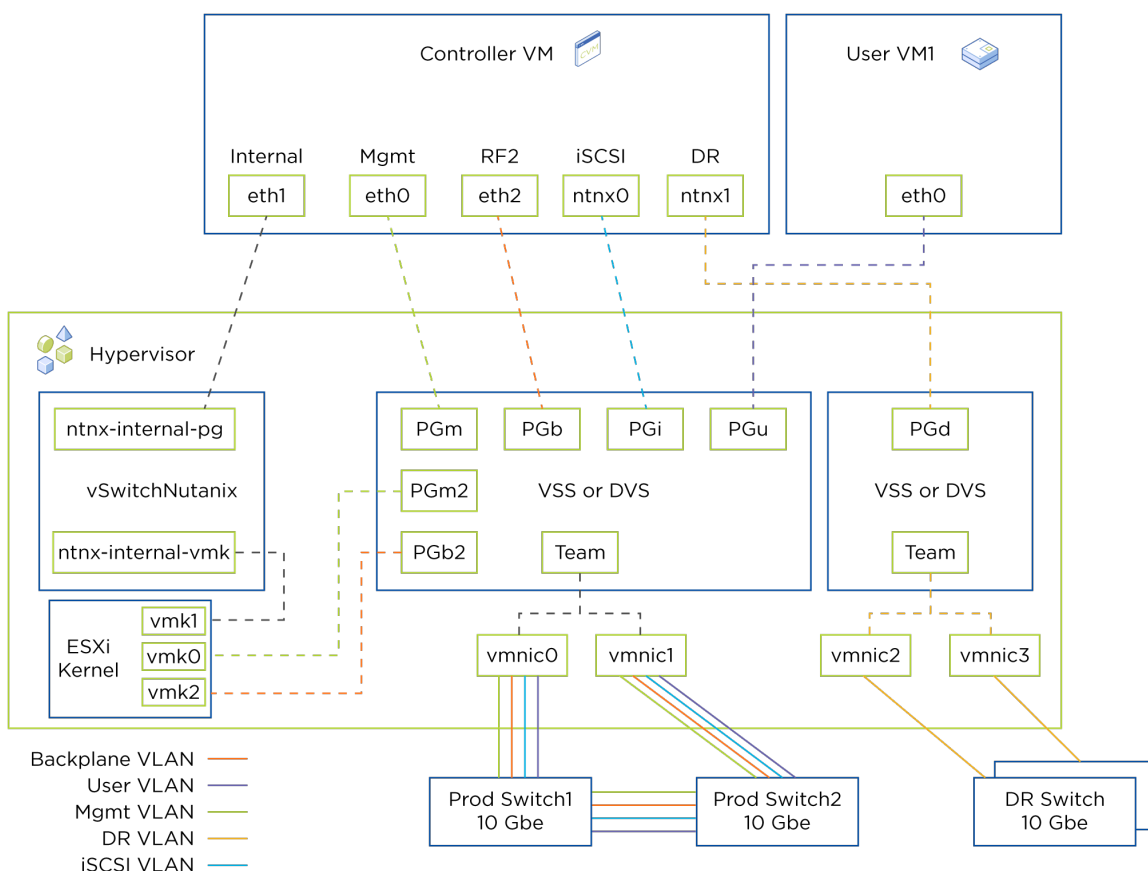


Figure 66: Backplane and Service Specific Segmentation Configured with two vSwitches on an ESXi Cluster

Here are the CVM to ESXi hypervisor connection details:

- The eth0 vNIC on the CVM and vmk0 on the host are carrying management traffic and connected to the hypervisor through the existing PGm (portgroup) on vSwitch0.

- The eth2 vNIC on the CVM and vmk2 on the host are carrying backplane traffic and connected to the hypervisor through a new user created PGb on the existing vSwitch.
- The ntnx0 vNIC on the CVM is carrying iSCSI traffic and connected to the hypervisor through PGi on the vSwitch1. No new vmkernel adapter is required.
- The ntnx1 vNIC on the CVM is carrying DR traffic and connected to the hypervisor through PGd on the vSwitch2. Here as well, there is no new vmkernel adapter required.

The following figure shows an implementation scenario where the backplane and service specific segmentation is configured with two vSwitches on an AHV hypervisors.

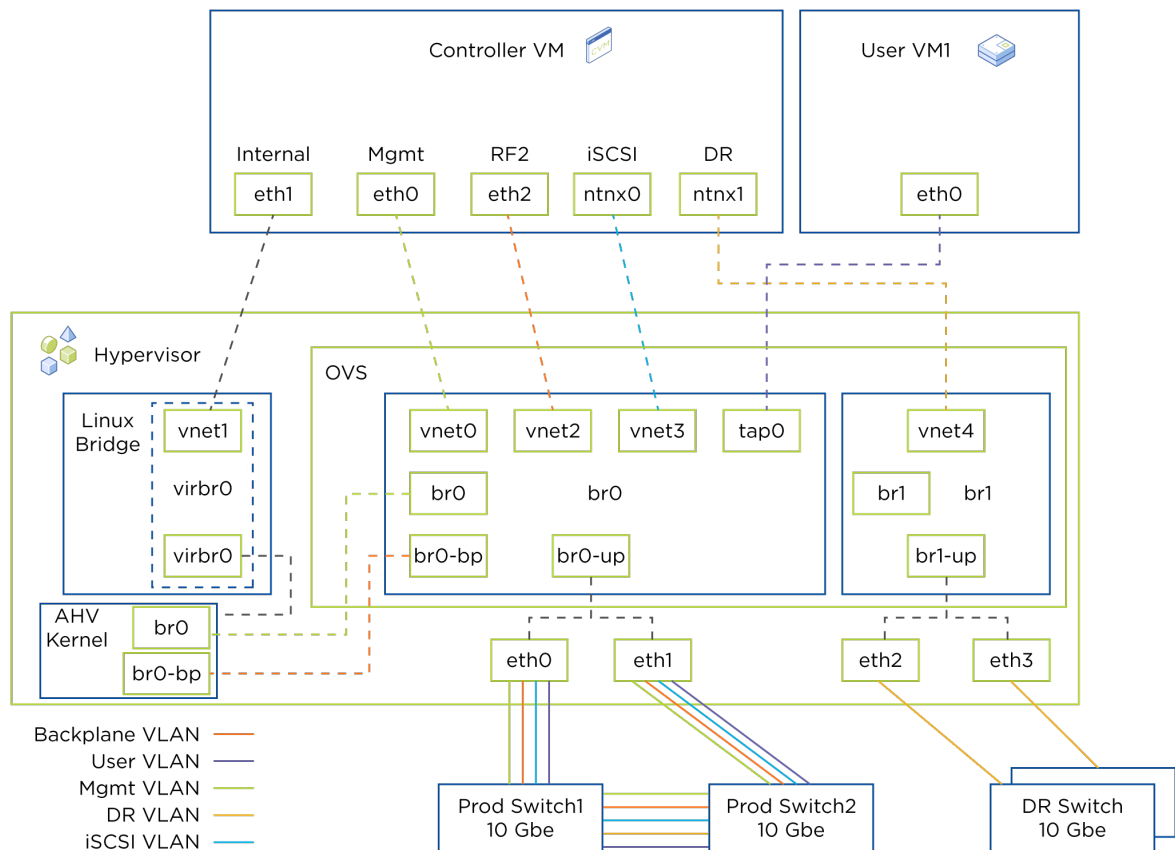


Figure 67: Backplane and Service Specific Segmentation Configured with two vSwitches on an AHV Cluster

Here are the CVM to AHV hypervisor connection details:

- The eth0 vNIC on the CVM is carrying management traffic and connected to the hypervisor through the existing vnet0.
- Other vNICs such as eth2, ntnx0, and ntnx1 are connected to the hypervisor through the auto created interfaces on either the existing or new vSwitch.

Note: In the above figure the interface name 'br0-bp' is read as 'br0-backplane'.

The following table describes the vNIC, port group (PG), VM kernel (vmk), virtual network (vnet) and virtual switch connections for CVM and hypervisor in different implementation scenarios. The tables capture information for ESXi and AHV hypervisors:

Table 6:

Implementation Scenarios	vNICs on CVM	Connected to ESXi Hypervisor	Connected to AHV Hypervisor
Backplane Segmentation with 1 vSwitch	eth0: DR, iSCSI, and Management traffic	vmk0 via existing PGm on vSwitch	Existing vnet0
	eth2: Backplane traffic	New vmk2 via PGb on vSwitch0 CVM vNIC via PGb on vSwitch0	Auto created interfaces on bridge br0
Backplane Segmentation with 2 vSwitches	eth0: Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0
	eth2: Backplane traffic	New vmk2 via PGb on new vSwitch CVM vNIC via PGb on new vSwitch	Auto created interfaces on new virtual switch
Service Specific Segmentation for Volumes with 1 vSwitch	eth0: DR, Backplane, and Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0
	ntnx0: iSCSI (Volumes) traffic	CVM vNIC via PGi on vSwitch0	Auto created interface on existing br0
Service Specific Segmentation for Volumes with 2 vSwitches	eth0: DR, Backplane, and Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0
	ntnx0: iSCSI (Volumes) traffic	CVM vNIC via PGi on new vSwitch	Auto created interface on new virtual switch
Service Specific Segmentation for DR with 1 vSwitch	eth0: iSCSI, Backplane, and Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0
	ntnx1: DR traffic	CVM vNIC via PGd on vSwitch0	Auto created interface on existing br0
Service Specific Segmentation for DR with 2 vSwitches	eth0: iSCSI, Backplane, and Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0

Implementation Scenarios	vNICs on CVM	Connected to ESXi Hypervisor	Connected to AHV Hypervisor
Backplane and Service Specific Segmentation with 1 vSwitch	ntnx1: DR traffic	CVM vNIC via PGd on new vSwitch	Auto created interface on new virtual switch
	eth0: Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0
	eth2: Backplane traffic	New vmk2 via PGb on vSwitch0 CVM vNIC via PGb on vSwitch0	Auto created interfaces on br0
	ntnx0: iSCSI traffic	CVM vNIC via PGi on vSwitch0	Auto created interface on br0
Backplane and Service Specific Segmentation with 2 vSwitches	ntnx1: DR traffic	CVM vNIC via PGd on vSwitch0	Auto created interface on br0
	eth0: Management traffic	vmk0 via existing PGm on vSwitch0	Existing vnet0
	eth2: Backplane traffic	New vmk2 via PGb on new vSwitch CVM vNIC via PGb on new vSwitch	Auto created interfaces on new virtual switch
	ntnx0: iSCSI traffic	CVM vNIC via PGi on vSwitch1 No new user defined vmkernel adapter is required.	Auto created interface on new virtual switch
	ntnx1: DR traffic	CVM vNIC via PGd on vSwitch2. No new user defined vmkernel adapter is required.	Auto created interface in new virtual switch

Implementation Considerations

Supported Environment

Network segmentation is supported in the following environment:

- The hypervisor must be one of the following:
 - For network segmentation by traffic type (separating backplane traffic from management traffic):
 - AHV
 - ESXi
 - Hyper-V
 - For service-specific traffic isolation:
 - AHV
 - ESXi
- For logical network segmentation, AOS version must be 5.5 or later. For physical segmentation and service-specific traffic isolation, the AOS version must be 5.11 or later.
- RDMA requirements:
 - Network segmentation is supported with RDMA for AHV and ESXi hypervisors only.
 - For more information about RDMA, see [Remote Direct Memory Access](#) in the *NX Series Hardware Administration Guide*.

Prerequisites

Disable proxy ARP within the Nutanix VLAN before configuring network segmentation.

For Nutanix Volumes

Stargate does not monitor the health of a segmented network. If physical network segmentation is configured, network failures or connectivity issues are not tolerated. To overcome this issue, configure redundancy in the network. That is, use two or more uplinks in a fault tolerant configuration, connected to two separate physical switches.

For Disaster Recovery

- Ensure that the VLAN and subnet that you plan to use for the network segment are routable.
- Make sure that you have a pool of IP addresses to specify when configuring segmentation. For each cluster, you need $n+1$ IP addresses, where n is the number of nodes in the cluster. The additional IP address is for the virtual IP address requirement.
- Enable network segmentation for disaster recovery at both sites (local and remote) before configuring remote sites at those sites.

Limitations

For Nutanix Volumes

- If network segmentation is enabled for Volumes, volume group attachments are not recovered during VM recovery.
- Nutanix service VMs such as Objects worker nodes continue to communicate with the CVM eth0 interface when using Volumes for iSCSI traffic. Other external clients such as Files use the new service-specific CVM interface.

Cluster Services That Support Traffic Isolation

You can isolate traffic associated with the following services to its own virtual network:

- Management (The default network that cannot be moved from CVM eth0)
- Backplane
- RDMA
- Service Specific Disaster Recovery
- Service Specific Volumes

Configurations in Which Network Segmentation Is Not Supported

Network segmentation is not supported in the following configurations:

- Clusters on which the CVMs have a manually created eth2 interface.
- Clusters on which the eth2 interface on one or more CVMs have been assigned an IP address manually. During an upgrade to an AOS release that supports network segmentation, an eth2 interface is created on each CVM in the cluster. Even though the cluster does not use these interfaces until you configure network segmentation, you must not manually configure these interfaces in any way.
- Clusters on which CVM interfaces are connected to port groups backed by NSX NVDS switches.

Caution:

Nutanix has deprecated support for manual multi-homed CVM network interfaces from AOS version 5.15 and later. Such a manual configuration can lead to unexpected issues on these releases. If you have configured an eth2 interface on the CVM manually, refer to the [KB-9479](#) and Nutanix Field Advisory #78 for details on how to remove the eth2 interface.

Troubleshooting Tips

This section provides information to assist troubleshooting of Network Segmentation deployments.

The Failed to restart one or more services after Backplane was enabled error may occur while enabling network segmentation. In such cases, the network segmentation task gets completed, however, restarting one or more services fails to complete on time.

To ensure the necessary services starts on time, login to a CVM over SSH and run the following command:

```
nutanix@CVM:~$ cluster start
```

Configuring the Network on an AHV Host

These steps describe how to configure host networking for physical and service-specific network segmentation on an AHV host. These steps are prerequisites for physical and service-specific network segmentation and you must perform these steps before you perform physical or service-specific traffic isolation. If you are configuring networking on an ESXi host, perform the equivalent steps by referring to the ESXi documentation. On ESXi, you create vSwitches and port groups to achieve the same results.

About this task

For information about the procedures to create, update and delete a virtual switch in Prism Element Web Console, see [Configuring a Virtual Network for Guest VMs](#) in the *Prism Element Web Console Guide*.

Note: The term *unconfigured node* in this procedure refers to a node that is not part of a cluster and is being prepared for cluster expansion.

To configure host networking for physical and service-specific network segmentation, do the following:

Note: If you are segmenting traffic on nodes that are already part of a cluster, perform the first step. If you are segmenting traffic on an unconfigured node that is not part of a cluster, perform the second step directly.

Procedure

1. If you are segmenting traffic on nodes that are already part of a cluster, do the following:
 - a. From the default virtual switch vs0, remove the uplinks that you want to add to the virtual switch you created by updating the default virtual switch.

For information about updating the default virtual switch vs0 to remove the uplinks, see [Creating or Updating a Virtual Switch](#) in the *Prism Element Web Console Guide*.

- b. Create a virtual switch for the backplane traffic or service whose traffic you want to isolate.
Add the uplinks to the new virtual switch.

For information about creating a new virtual switch, see [Creating or Updating a Virtual Switch](#) in the *Prism Web Console Guide*.

2. If you are segmenting traffic on an unconfigured node (new host) that is not part of a cluster, do the following:

- a. Create a bridge for the backplane traffic or service whose traffic you want to isolate by logging on to the new AHV host.

```
ovs-vsctl add-br br1
```

- b. From the default bridge br0, log on to the host CVM and keep only eth0 and eth1 in br0.

```
manage_ovs --bridge_name br0 --interfaces eth0,eth1 --bond_name br0-up --bond_mode  
active-backup update_uplinks
```

- c. Log on to the host CVM and then add eth2 and eth3 to the uplink bond of br1.

```
manage_ovs --bridge_name br1 --interfaces eth2,eth3 --bond_name br1-up --bond_mode  
active-backup update_uplinks
```

Note: If this step is not done correctly, a network loop can be created that causes a network outage. Ensure that no other uplink interfaces exist on this bridge before adding the new interfaces, and always add interfaces into a bond.

What to do next

Prism can configure a VLAN only on AHV hosts. Therefore, if the hypervisor is ESXi, in addition to configuring the VLAN on the physical switch, make sure to configure the VLAN on the port group.

If you are performing physical network segmentation, see [Physically Isolating the Backplane Traffic on an Existing Cluster](#) on page 158.

If you are performing service-specific traffic isolation, see [Service-Specific Traffic Isolation](#) on page 168.

Network Segmentation for Traffic Types (Backplane, Management, and RDMA)

You can segment the network on a Nutanix cluster in the following ways:

- You can segment the network on an existing cluster by using the Prism web console.
- You can segment the network when creating a cluster by using Nutanix Foundation 3.11.2 or higher versions.

The following topics describe network segmentation procedures for existing clusters and changes during AOS upgrade and cluster expansion. For more information about segmenting the network when creating a cluster, see the [Field Installation Guide](#).

Isolating the Backplane Traffic Logically on an Existing Cluster (VLAN-Based Segmentation Only)

You can segment the network on an existing cluster by using the Prism web console. You must configure a separate VLAN for the backplane network to achieve logical segmentation. The network segmentation process creates a separate network for backplane communications on the existing default virtual switch. The process then places the eth2 interfaces (that the process creates on the CVMs during upgrade) and the host interfaces on the newly created network. This method allows you to achieve logical segmentation of traffic over the selected VLAN. From the specified subnet, assign IP addresses to each new interface. You, therefore, need two IP addresses per node. When you specify the VLAN ID, AHV places the newly created interfaces on the specified VLAN.

Before you begin

If your cluster has RDMA-enabled NICs, follow the procedure in [Isolating the Backplane Traffic on an Existing RDMA Cluster](#) on page 143.

- For ESXi clusters, it is mandatory to create and manage port groups that networking uses for CVM and backplane networking. Therefore, ensure that you create port groups on the default virtual switch vs0 for the ESXi hosts and CVMs.

Since backplane traffic segmentation is logical, it is based on the VLAN that is tagged for the port groups. Therefore, while creating the port groups ensure that you tag the new port groups created for the ESXi hosts and CVMs with the appropriate VLAN ID. Consult your networking team to acquire the necessary VLANs for use with Nutanix nodes.

- For new backplane networks, you must specify a non-routable subnet. The interfaces on the backplane network are automatically assigned IP addresses from this subnet, so reserve the entire subnet for the backplane network segmentation. See the [Configuring Backplane IP Pool](#) on page 176 topic to create an IP pool for backplane interfaces.

About this task

You need separate VLANs for Management network and Backplane network. For example, configure VLAN 100 as Management network VLAN and VLAN 200 as Backplane network VLAN on the Ethernet links that connect the Nutanix nodes to the physical switch.

Note: Nutanix does not control these VLAN IDs. Consult your networking team to acquire VLANs for the Management and Backplane networks.

To segment the network on an existing ESXi and Hyper-V clusters for a backplane LAN, do the following:

To segment the network on an existing AHV cluster for a backplane LAN, follow the procedure described in the [Physically Isolating the Backplane Traffic on an AHV Cluster](#) on page 158 topic.

Note:

In this method, for AHV nodes, logical segmentation (VLAN-based segmentation) is done on the default bridge. The process creates the host backplane interface (VMkernel) on the **Backplane Network** port group on ESXi or **br0-backplane** (interface) on br0 bridge in case of AHV. The eth2 interface on the CVM is on **CVM Backplane Network** by default.

You don't need to manually create the VMkernel adapter (vmk) for backplane segmentation since the workflow takes care of creating it on the port group selected as the host port group or default **Backplane Network** port group if none was selected.

Procedure

1. Log on to the Prism web console, click the gear icon in the top-right corner, and then click **Network Configuration** in the **Settings** page.

The **Network Configuration** dialog box appears.

2. In the **Network Configuration > Internal Interfaces > Backplane LAN** row, click **Configure**.

The **Create Interface** dialog box appears.

3. In the **Create Interface** dialog box, provide the necessary information.

- 1. In the **Subnet IP** field, specify a non-routable subnet.

Ensure that the subnet has sufficient IP addresses. The segmentation process requires two IP addresses per node. Reconfiguring the backplane to increase the size of the subnet involves cluster downtime, so you might also want to make sure that the subnet can accommodate new nodes in the future.

- 2. In the **Netmask** field, specify the netmask.

- 3. If you want to assign the interfaces on the network to a VLAN, specify the VLAN ID in the **VLAN ID** field.

Nutanix recommends that you use a VLAN. If you do not specify a VLAN ID, Prism uses the default VLAN on the virtual switch.

- 4. In the **Host Port Group** list, select the port group you created for the host.

- 5. In the **CVM Port Group** list, select the port group you created for the CVM.

Note: For VLAN-based segmentation, Nutanix recommends you leave the Host Port Group and CVM Port Group fields blank. Prism selects the default port group from vSwitch0 when you do not provide the Host Port Group and CVM Port Group details.

4. Click **Verify and Save**.

The network segmentation process creates the backplane network if the network settings that you specified pass validation.

RDMA over Converged Ethernet (RoCE)

This section provides the information about RDMA, RDMA port pass-through mechanism, and how to configure the RDMA network segmentation using Zero-Touch RoCE (ZTR) or Priority-Based Flow Control (PFC) mechanism.

RDMA Overview

Remote Direct Memory Access (RDMA) enables you to directly transfer data between multiple hosts without involving CPU, OS, or system cache. RDMA reduces the communication latency and increases the bandwidth output for the data transfer. It directly uses the network adapters for data transfer and never creates a data copy between network layers.

When RDMA is enabled, the CPU resources available for other applications running in the cluster enhance the AOS data acceleration mechanism.

RDMA Port Pass-through Mechanism

The following table describes the supported RDMA port pass-through mechanisms:

Table 7: RDMA Port Pass-through Mechanism

AOS release	Hypervisor	RDMA Port Pass-through Mechanism
Prior to AOS 6.6	<ul style="list-style-type: none"> • AHV • ESXi 	<p>RDMA port pass-through can be done during Foundation setup only. In this case, the CVM reserves the entire NIC for RDMA port pass-through. The host cannot use the empty NIC port for other operations nor can it change the selected port. If the NIC or port fails, or any connectivity issue occurs, the I/O traffic falls back to the TCP/IP connection mode.</p>
AOS 6.6 and above	AHV	<p>RDMA port pass-through can be done either during Foundation setup or after the cluster creation is completed.</p> <p>The following behavior is applicable for RDMA port pass-through:</p> <ul style="list-style-type: none"> • If RDMA port pass-through is done during Foundation setup, the CVM reserves the entire NIC for RDMA port pass-through. In this case, the host cannot use the empty NIC port for other operations nor can it change the selected port. • If RDMA port pass-through is not done during Foundation setup, the system provides you an option to reserve the port for RDMA. <div> <p>Note: Ensure the port that you reserve for RDMA is not allocated to the uplink vswitch of the host. If the port that you want to reserve for RDMA is already allocated for uplink vswitch, the system displays it as disabled for selection and you need to manually remove it from the vswitch.</p> </div>

AOS release	Hypervisor	RDMA Port Pass-through Mechanism
	ESXi	RDMA port pass-through can be done during Foundation setup only. In this case, the system doesn't provide you an option to change the port for RDMA after cluster creation.

For more information on how to enable RDMA port pass-through to the CVM during Foundation, see [Configuring Foundation VM by Using the Foundation GUI](#).

Cluster Expansion Case

When a new node is added to increase the capacity of a cluster, the RDMA ports cannot be changed after the cluster is created. In this case, the system checks if RDMA is enabled in the cluster. If enabled, the system provisions the RDMA for the new node, and selects the first port available in the NIC. To change the RDMA port, you need to disable the RDMA for the whole cluster and reconfigure it with the new port selection.

ZTR Specifications

The ZTR is a deployment mechanism for RDMA setup in which the Mellanox NIC firmware handles the entire configurations (card optimizations) without any user intervention or dependency on the customized switch profiles or switch compatibility. ZTR reduces RDMA deployment duration and does not require the support of PFC and End-to-end Congestion Notification (ECN) settings.

ZTR functionality is supported with both ESXi and AHV hypervisors. For information about how to enable ZTR for RDMA network segmentation, see [Isolating the Backplane Traffic on an Existing RDMA Cluster](#) on page 143.

Nutanix recommends you use ZTR only if NVIDIA Mellanox Connect X-5 Ethernet Adapters (Cx5 NICs) or NVIDIA Mellanox Connect X-6 Ethernet Adapters (Cx6 NICs) are available in your setup. For more information about the NICs compatibility for ZTR feature, see [NIC Compatibility Matrix for RDMA Features](#) on page 141.

NIC Compatibility Matrix for RDMA Features

The following table provides the information about NIC compatibility with RDMA features; RDMA Port Pass-through and ZTR, and the required workaround:

Table 8: NIC Compatibility Matrix for RDMA Features

NIC Available at Site	RDMA Feature	Compatibility Information	Workaround
NVIDIA Mellanox ConnectX-4 Ethernet Adapter (CX-4)	ZTR	Not supported Nutanix recommends you use ZTR only if CX-5 NICs are available in your setup.	None

NIC Available at Site	RDMA Feature	Compatibility Information	Workaround
	RDMA port pass-through	RDMA port pass-through after cluster creation is not supported.	Perform RDMA port pass-through during Foundation setup, and then upgrade to AOS 6.6 release.
NVIDIA Mellanox ConnectX-5 Ethernet Adapter (CX-5)	ZTR	Supported	None
	RDMA port pass-through	Supported The RDMA port pass-through is supported for both of the following scenarios: <ul style="list-style-type: none"> During Foundation setup where CVM reserves the entire NIC for RDMA port pass-through. In this case, the host cannot use the empty NIC port for other operations nor can it change the selected port. After cluster creation in case of AHV only. 	None
NVIDIA Mellanox ConnectX-6 Ethernet Adapter (CX-6)	ZTR	Supported	None
	RDMA port pass-through	Supported The RDMA port pass-through is supported for both of the following scenarios: <ul style="list-style-type: none"> During Foundation setup where CVM reserves the entire NIC for RDMA port pass-through. In this case, the host cannot use the empty NIC port for other operations nor can it change the selected port. After cluster creation in case of AHV only. 	None

Important:

- A mix of RDMA NIC families on the same node is not supported. For example, a combination of NVIDIA Mellanox ConnectX-4 Ethernet Adapter (CX-4 NIC) and NVIDIA Mellanox ConnectX-5 Ethernet Adapter (CX-5 NIC) or a combination of CX-5 NICs with different speeds such as CX-5 10Gb NIC and CX-5 25Gb NIC, on the same node is not supported. If you have any of these combination types of RDMA NIC families on the same node, the system never allows you to enable RDMA on the cluster.
- Mixing NIC cards across nodes in a cluster is supported. In this case, the oldest family of cards dictates the feature supported for the cluster. For example, if any CX-4 NIC is present, the system blocks the RDMA live port pass-through (RDMA port pass-through after cluster creation) and ZTR functions, and allows only RDMA data replication.

Isolating the Backplane Traffic on an Existing RDMA Cluster

This section describes how to configure the RDMA network segmentation settings using either PFC or ZTR from Prism Central.

About this task

The network segmentation process creates a separate network for RDMA communications on the existing default virtual switch and places the `rdma0` interface (created on the CVMs during upgrade) and the host interfaces on the newly created network. From the specified subnet, IP addresses are assigned to each new interface. Two IP addresses are therefore required per node. If you specify the optional VLAN ID, the newly created interfaces are placed on the VLAN. A separate VLAN is highly recommended for the RDMA network to achieve true segmentation.

Before you begin

Ensure that the following prerequisites are met before you proceed with RDMA network segmentation:

- Specify a non-routable subnet. The interfaces on the backplane network are automatically assigned IP addresses from the subnet. You reserve the entire subnet for the backplane network alone.
- If you plan to specify a VLAN for the RDMA network, ensure that the VLAN is configured on the physical switch ports to which the nodes are connected.
- Configure the switch interface as a trunk port.
- Observe the NICs compatibility information specified in [NIC Compatibility Matrix for RDMA Features](#) on page 141 .
- Mixed configuration in a cluster is not supported where some nodes have RDMA port pass-through or ZTR enabled while other nodes have it disabled. All nodes in the cluster need to be uniformly configured with the RDMA functionality.

The following prerequisites are applicable only for ZTR :

- NVIDIA Mellanox Connect X-5 Ethernet Adapters (Cx5 NICs) or NVIDIA Mellanox Connect X-6 Ethernet Adapters (Cx6 NICs) are available. For details, see [NIC Compatibility Matrix for RDMA Features](#) on page 141.
- NICs on all nodes are running the same NIC firmware. Perform a NCC health check to verify the minimum driver version recommended and supported by Mellanox NIC running on the Nutanix platforms. For information about how to perform NCC Health check, see [KB-4289](#)
- AHV or ESXi hypervisor and AOS 6.6 or later (CVM 6.6) are deployed for the cluster. For information about supported AHV or ESXi hypervisor versions, see [KB-4289](#).

Procedure

To isolate the backplane traffic on an existing RDMA cluster, perform the following steps:

1. Log on to the Prism Element web console, click the gear icon in the top-right corner, and click **Network Configuration** in the **Settings** page.
The **Network Configuration** dialog box displays.

2. Click the **Internal Interfaces** tab.
The **Internal Interfaces** tab displays.

Note: If the system detects that NVIDIA Mellanox Cx5 NIC or NVIDIA Mellanox Cx6 NIC is used, it provides you the option to configure RDMA.

3. Click **Configure** in the **RDMA** row, and set the following attributes in **RDMA** dialog box:

Warning: The **Configure** option is disabled if any of the following conditions exist in your setup:

- All the nodes in the cluster do not contain at least two RDMA-enabled NIC cards.
- A mix of RDMA NIC families on the same node is present. For example, a combination of NVIDIA Mellanox ConnectX-4 Ethernet Adapter (CX-4 NIC) and NVIDIA Mellanox ConnectX-5 Ethernet Adapter (CX-5 NIC) or a combination of CX-5 NICs with different speeds such as CX-5 10Gb NIC and CX-5 25Gb NIC, on the same node is present. If you have any of these combination types of RDMA NIC families on the same node, the system never allows you to enable RDMA on the cluster.

Mixing NIC cards across nodes in a cluster is supported. In this case, the oldest family of cards dictates the feature supported for the cluster. For example, if any CX-4 NIC is

present, the system blocks the RDMA live port pass-through (RDMA port pass-through after cluster creation) and ZTR functions, and allows only RDMA data replication.

For more information, see [NIC Compatibility Matrix for RDMA Features](#) on page 141.

- Under **Subnet IP** , specify a non-routable Subnet IP.

Note: Ensure that the subnet size can accommodate cluster expansion in the future.

- Under **Netmask**, specify a non-routable netmask.
- Under **VLAN**, specify a VLAN ID for the RDMA LAN.

Note: VLAN ID is optional with ZTR but Nutanix recommends you to use it for true network segmentation and enhanced security

- Select either the **Use Zero Touch RoCE** checkbox to enable ZTR or select the **PFC** value configured on the physical switch port.

Note: When the Use Zero Touch RoCE checkbox is selected:

- The PFC field is disabled.
- The system automatically defines the RDMA port if RDMA port pass-through is done during Foundation setup. You cannot change the RDMA port in this case.

RDMA ? >

Descriptive Name	RDMA
Interface	rdma0

Subnet IP

172.16.99.0

Netmask

255.255.255.0

VLAN ID

601

Enable VLAN ID on Physical Switch as well.

Virtual Switch

vs0

☒ Use Zero Touch RoCE ⓘ

PFC

0

Cancel Verify and Save

Figure 68: RDMA Network Segmentation: Pass-through done during Foundation

4. Click **Verify and Save**.

The selected RDMA network segmentation settings are saved.

Note: If the RDMA port pass-through is not done during Foundation setup and only AHV hypervisor is deployed, the system prompts you to define the RDMA port.

Perform the following steps to define RDMA port:

- a. Click **Next** and select the RDMA port in the **Port Selection** tab.

RDMA?×

Network

Port Selection

Descriptive Name	RDMA
Interface	rdma0

Subnet IP

172.200.11.0

Netmask

255.255.255.0

VLAN ID

601

Enable VLAN ID on Physical Switch as well.

Virtual Switch

vs0

☒ Use Zero Touch RoCE ?

PFC

0

Cancel




Next

Figure 69: RDMA Network Segmentation - No Pass-through done during Foundation

- b. Select the port to be reserved for RDMA on all the hosts.

Network · Port Selection

Select port for the RDMA Interface.

Hosts	Ports ?
denahi01-3	3 Ports 
<input checked="" type="radio"/> 1	eth0
<input type="radio"/> 2	eth6 ?
<input type="radio"/> 3	eth7 ?
denahi01-1	3 Ports 
<input checked="" type="radio"/> 1	eth0
<input type="radio"/> 2	eth6 ?
<input type="radio"/> 3	eth7 ?
denahi01-2	3 Ports 
<input checked="" type="radio"/> 1	eth0
<input type="radio"/> 2	eth6 ?

BackSave

Figure 70: RDMA Port Selection

c. Click **Save**.

Support for iSCSI Extensions for RDMA (iSER)

Internet Small Computer Systems Interface (iSCSI) is a Storage Area Networking (SAN) based data transfer model that utilizes TCP/IP to transfer the data between applications to a storage array. The iSER is an extension of iSCSI that utilizes the direct data placement technology of RDMA protocol to eliminate the copies in the data path. The system appropriately utilizes the CPU resources and reduces latency between AHV host and CVM I/O transfer.

iSER provides the following key benefits:

- Enhances the I/O flow performance between AHV host and CVM. iSER bypasses both the AHV host and CVM kernel spaces and enables you to achieve low latency in the I/O data transfer as the data is sent and received between AHV host and CVM without the involvement of the TCP software stack. iSER enables you to achieve a high bandwidth in the I/O data transfer as it removes the unwanted context switches and system calls between AHV host and CVM.
- Saves platform resources due to less CPU utilization as an application can read remote memory without any intervention of the remote processor.

The following figure shows how the I/O traffic flows between AHV host and CVM:




Figure 71: I/O Traffic - AHV Host and CVM

The following figure shows the detailed network architecture for iSER:

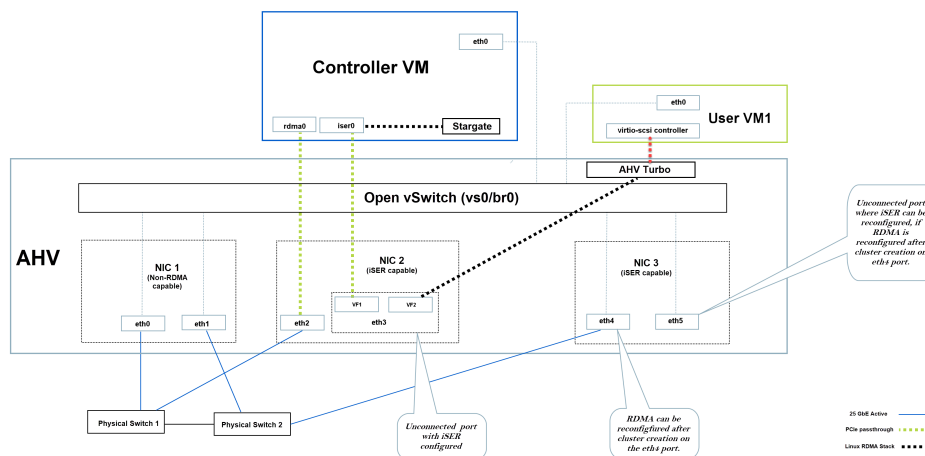


Figure 72: iSER Network Architecture

The following specifications apply to the iSER architecture:

- NIC 1 is a non-RDMA capable NIC for regular network functionality, and NIC 2 and NIC 3 are dual-port qualified iSER NICs for RDMA and iSER. For more information about qualified iSER NICs, see [NIC Compatibility Matrix for iSER](#) on page 153.
- The first port on NIC 2 (eth2) and NIC 3 (eth4) is connected to the uplink switch. The second port; eth3 on NIC 2 and eth5 on NIC 3 is disconnected.
- One of the unconnected ports (eth3 or eth5 in the above diagram) can be used for iSER.
- One of the connected ports on NIC 2 or NIC 3 (eth2 or eth4 in the above diagram) can be passed through to the CVM for RDMA replication.
- VF1 interface on an iSER port is consumed by Stargate using the Linux RDMA stack.
- VF2 interface on an iSER port is consumed by AHV Turbo using the Linux RDMA stack.
- iSER and RDMA replication can either be enabled during cluster creation or dynamically enabled or disabled on the available ports (through the RDMA manageability workflow) after cluster creation.

In the above diagram:

- RDMA can be reconfigured after cluster creation on the eth4 port.
- iSER can be reconfigured on the eth5 port if RDMA is reconfigured after cluster creation on the eth4 port in NIC 3.

Note: Currently, iSER uses fixed internal IPs. There is no supported workflow to change iSER IPs for the CVM or the AHV host.

The following specifications are applicable in the iSER setup:

iSER Specifications	Component	Description
iSER Initiator	AHV (AHV Turbo in AHV)	AHV Turbo providing iSCSI packet processing
iSER Target	CVM (Stargate in CVM)	Stargate in CVM providing RDMA transport and receiving data

Typically, AHV and CVM communicate with each other using the iSCSI protocol over the Open vSwitch TCP datapath. AHV forwards the I/O data that originates from the guest VM to CVM, which processes the I/O data request. After CVM completes the I/O data processing, it returns the response and/or data for the request to AHV.

If iSER is enabled, the following actions occur between AHV and CVM to establish storage connectivity:

1. The system establishes a handshake connection between AHV and CVM over TCP.
2. On the TCP iSCSI login, the CVM responds to AHV with the status of RDMA as enabled or disabled and returns the target IP address for iSER if RDMA is enabled.
3. If the CVM indicates that RDMA connectivity is possible, AHV closes the TCP connection and attempts to establish an RDMA connection to the CVM iSER IP.
4. If the RDMA connection and TCP iSCSI login over the RDMA connection succeed, AHV proceeds to submit the guest VM SCSI requests to CVM using the RDMA connection.

Note: In case of any failure scenarios, the system falls back and re-establishes TCP connections between AHV and CVM through the standard vSwitch path, and the data processing is resumed between AHV and CVM with zero disruption to the running workload.

CVM Memory Requirements for iSER

For iSER, CVM requires a minimum of 64 GiB vRAM for optimum performance. For more information, see [Controller VM \(CVM\) Field Specifications](#) in *Acropolis Advanced Admin Guide*.

iSER Port-Passthrough Mechanism

You can perform iSER NIC port pass-through either during Foundation setup or after the completion of cluster creation using Prism Element.

The following behavior applies to iSER port pass-through:

- If you perform the iSER port pass-through during Foundation setup, the CVM reserves the whole NIC for both iSER and RDMA data replication. The CVM reserves one port on the NIC for RDMA data replication and configures the other port for iSER. For more information about Foundation setup, see [Configuring Foundation VM by Using the Foundation GUI](#) in *Field Installation Guide*.
- If you perform the iSER port pass-through after cluster creation, the system allows you to configure iSER on one of the unconnected NIC ports. You can use the other NIC port to configure RDMA if RDMA port-passthrough is not done during Foundation setup.

Note: Ensure the port that you configure for iSER is not allocated to the uplink vswitch of the host. If the port that you want to configure for iSER is already allocated for uplink vswitch, the system displays it as disabled for selection and you need to remove it from the vswitch manually.

Cluster Expansion Case

When a new node is added to increase the capacity of a cluster, the iSER ports cannot be changed after the cluster is created. In this case, the system checks if iSER is enabled in the cluster. If enabled, the system provisions the iSER for the new node and selects the first port available in the NIC. To change the iSER port, disable the iSER for the whole cluster and reconfigure it with the new port selection.

NIC Compatibility Matrix for iSER

The following table provides the information about NIC compatibility for iSER:

Table 9: NIC Compatibility Matrix for iSER

NIC Available at Site	Number of NICs Required	iSER Port Pass-through Compatibility Information	Workaround
NVIDIA Mellanox ConnectX-4 Ethernet Adapter (CX-4)	Not Applicable	Not supported. Nutanix recommends you use iSER only if CX-5 NICs are available in your setup	None

NIC Available at Site	Number of NICs Required	iSER Port Pass-through Compatibility Information	Workaround
NVIDIA Mellanox ConnectX-5 Ethernet Adapter (CX-5)	Two dual-port NICs for RDMA + iSER	<p>Supported</p> <p>The iSER port pass-through is supported for both the following scenarios:</p> <ul style="list-style-type: none"> During Foundation setup where CVM reserves the whole NIC for both iSER and RDMA data replication. After cluster creation in case of AHV only, where the system provides you an option to configure iSER on one of the unconnected NIC ports. You can use the other NIC port to configure RDMA if RDMA port-passthrough is not done during Foundation setup. 	None
NVIDIA Mellanox ConnectX-6 Ethernet Adapter (CX-6)	Not Applicable	<p>Yet to be qualified.</p> <p>iSER support is yet to be qualified with CX-6 NICs.</p> <p>Nutanix recommends you use iSER only if CX-5 NICs are available in your setup.</p>	None

Important:

- A mix of RDMA NIC families on the same node is not supported. For example, a combination of NVIDIA Mellanox ConnectX-4 Ethernet Adapter (CX-4 NIC) and NVIDIA Mellanox ConnectX-5 Ethernet Adapter (CX-5 NIC) or a combination of CX-5 NICs with different speeds such as CX-5 10Gb NIC and CX-5 25Gb NIC, on the same node is not supported. If you have any of these combination types of RDMA NIC families on the same node, you cannot enable RDMA data replication on the cluster and cannot configure ports for iSER.
- Mixing NIC cards across nodes in a cluster is supported. In this case, the oldest family of cards dictates the iSER port pass-through function. For example, if any CX-4 NIC is present, the system blocks the iSER live port pass-through and allows only RDMA data replication.

iSER Limitations

The following limitations apply to iSER:

- The iSER is supported with AHV only as it involves changes in I/O traffic flow between the AHV host and CVM. iSER support is not available with the ESXi hypervisor in the Nutanix environment.
- The total number of VM disks supported per node using iSER is limited to 200 because of memory limitations, and you cannot specify or configure the VM disks for iSER acceleration.

Configuring iSER Port in an Existing RDMA Cluster

This section describes how to configure the iSER port in an existing RDMA cluster.

About this task

iSER is enabled by default on all nodes with Mellanox CX-5 NIC on which RDMA data replication is successful. The system reserves two IP addresses for iSER. The iSER IP addresses are fixed, and during CVM restart, the system uses these reserved IP addresses for iSER.

Note: During AOS or AHV upgrade from previous versions, iSER gets automatically enabled and AHV and CVM configures an internal IP address for iSER.

Before you begin

Ensure that the following prerequisites are met before you configure the iSER port in an existing cluster:

- AHV is upgraded to AHV-20230302.207 and AOS is upgraded to AOS 6.7 version.
- Only the qualified iSER NICs; NVIDIA Mellanox dual-port CX-5 NICs, are present in your setup, or else iSER functionality is disabled. For more information, see [NIC Compatibility Matrix for iSER](#) on page 153.
- Two dual-port qualified iSER NICs are present in your setup for RDMA and iSER. For more information about qualified iSER NICs, see [NIC Compatibility Matrix for iSER](#) on page 153.
- One non-RDMA capable NIC is present in your setup for regular network functionality.
- The OFED drivers are qualified for iSER. For more information, see *Linux Drivers* information on the [NVIDIA website](#).
- NIC Firmware upgrade is done through LCM.
- RDMA for iSER must be enabled for Virtual Function (VF). For more information, see [Isolating the Backplane Traffic on an Existing RDMA Cluster](#) on page 143.
- An unconnected port is available on NIC for iSER traffic.
- Observe the limitations as specified in [iSER Limitations](#) on page 155.

Procedure

To configure iSER in an existing RDMA cluster, perform the following steps:

1. Log on to the Prism web console, click the gear icon in the top-right corner, and click **Network Configuration** in the **Settings** page.
The **Network Configuration** dialog box displays.

2. Click the **Internal Interfaces** tab.
The **Internal Interfaces** tab displays.

Note: If the system detects that NVIDIA Mellanox Cx5 NIC is used, it provides you the option to configure iSER.

3. Click **Configure** in the **iSER** row, and set the port in following attributes in **Configure iSER** page.

Note: If only AHV is deployed and iSER port pass-through is not done during Foundation setup, the system allows you to define the iSER port.

Perform the following steps to define the iSER port:

- a. Select the port for iSER on all the hosts.

Configure iSER

X

Configure iSER for AOS storage acceleration between AHV and CVM

[Learn more](#)

Descriptive Name	iSER
Interface	iser0

Select port for the iSER interface.

Hosts	Ports	
Calamba01-4	2 Ports	
<input type="radio"/> 1	eth2	
<input type="radio"/> 2	eth3	
Calamba01-1	2 Ports	
Calamba01-2	2 Ports	

Cancel

Save

Figure 73: Configure iSER

- b. Click **Save**.

Physically Isolating the Backplane Traffic on an Existing Cluster

By using the Prism web console, you can configure the eth2 interface on a separate virtual switch if you wish to isolate the backplane traffic to a separate physical network.

If you do not configure as separate virtual switch, the backplane traffic uses another VLAN in the default switch for VLAN-based traffic isolation.

A virtual switch is known as the following in different hypervisors.

Hypervisor	Virtual Switch
AHV	Virtual Switch
ESXi	vSwitch
Hyper-V	Hyper-V Virtual Switch

Network segmentation process creates a separate network for backplane communications on the new virtual switch. The segmentation process places the CVM eth2 interfaces and the host interfaces on the newly created network. Specify a subnet with a network mask and, optionally, a VLAN ID. From the specified subnet or an IP Pool assign IP addresses to each new interface in the new network. You require a minimum of two IP addresses per node.

If you specify the optional VLAN ID, the newly created interfaces are placed on VLAN.

Nutanix highly recommends a separate VLAN for the backplane network to achieve true segmentation.

Requirements and Limitations

- Ensure that physical isolation of backplane traffic is supported by the AOS version deployed.
- Ensure that you configure the network (port groups or bridges) on the hosts and associate the network with the required physical NICs before you enable physical isolation of the backplane traffic.
For AHV, see [Configuring the Network on an AHV Host](#) on page 136. For ESXi and Hyper-V, see VMware and Microsoft documentation respectively.
- Segmenting backplane traffic can involve up to two rolling reboots of the CVMs. The first rolling reboot is done to move the backplane interface (eth2) of the CVM to the selected port group, virtual switch or Hyper-V switch. This is done only for CVM(s) whose backplane interface is not already connected to the selected port group, virtual switch or Hyper-V switch. The second rolling reboot is done to migrate the cluster services to the newly configured backplane interface.

Physically Isolating the Backplane Traffic on an AHV Cluster

Before you begin

On the AHV hosts, do the following:

1. From the default virtual switch vs0, remove the uplinks (physical NICs) that you want to add to a new virtual switch you create for the backplane traffic in the next step.
2. Create a virtual switch for the backplane traffic.

Add the uplinks to the new bond when you create the new virtual switch.

See [Configuring the Network on an AHV Host](#) on page 136 for instructions about how to perform these tasks on a host.

Note: Before you perform the following procedure, ensure that the uplinks you added to the virtual switch are in the UP state.

About this task

Perform the following procedure to physically segment the backplane traffic on an AHV cluster.

Procedure

1. Shut down all the guest VMs in the cluster from within the guest OS or use the Prism Element web console.
2. Place all nodes of a cluster into the maintenance mode (AHV hypervisor and not CVM).

- a. Use SSH to log on to a Controller VM in the cluster
- b. Determine the IP address of all the nodes (AHV hypervisor):

```
nutanix@cvm$ acli host.list
```

Note down all the hypervisor IPs for the cluster.

- c. Put the node into the maintenance mode. Repeat this step for all nodes:

```
nutanix@cvm$ acli host.enter_maintenance_mode hypervisor-IP-address [wait="{ true  
| false }" ] [non_migratable_vm_action="{ acpi_shutdown | block }" ]
```

Note: Never put Controller VM and AHV hosts into maintenance mode on single-node clusters. It is recommended to shutdown user VMs before proceeding with disruptive changes.

Replace *hypervisor-IP-address* with either the AHV hypervisor IP address.

The following are optional parameters for running the `acli host.enter_maintenance_mode` command:

- **wait**
- **non_migratable_vm_action**

Example.

```
nutanix@cvm$ acli host.enter_maintenance_mode 197.116.6.79  
EnterMaintenanceMode: pending  
EnterMaintenanceMode: complete
```

Do not continue if the node has failed to enter the maintenance mode.

- d. Verify if the node is in the maintenance mode:

```
nutanix@cvm$ acli host.get hypervisor-IP-address
```

In the output that is displayed, ensure that **node_state** equals to **EnteredMaintenanceMode** and **schedulable** equals to **False**.

Example.

```
nutanix@cvm$ acli host.get 197.116.6.79  
197.116.6.79 {  
  cpu_usage_ppm: 192941  
  cvm_memory_size_bytes: 21474836480  
  cvm_num_vcpus: 8  
  cvm_num_vnics: 3  
  cvm_uuid: "e96dbef0-d425-4926-9fe8-4d10b1a69902"  
  host_overhead_bytes: 4957721854  
  logical_timestamp: 23  
  max_mem_ha_reserved_bytes: 0  
  mem_assigned_bytes: 0
```

```

mem_usage_bytes: 26654856446
memory_size_bytes: 269842644992
node_state: "EnteredMaintenanceMode"
num_cpus: 32
pool_size_bytes: 0
schedulable: False
uuid: "cc50ea78-49ce-4767-90b3-a0a2ef891446"
}

```

3. Enable backplane network segmentation.

- a. Log on to the Prism web console, click the gear icon in the top-right corner, and then click **Network Configuration** in the **Settings** page.
- b. On the **Internal Interfaces** tab, in the **Backplane LAN** row, click **Configure**.
- c. In the **Backplane LAN** dialog box, do the following:

- In **Subnet IP**, specify a non-routable subnet that is different from the subnet used by the AHV host and CVMs.

The AOS CVM default route uses the CVM eth0 interface, and there is no route on the backplane interface. Therefore, Nutanix recommends only using a non-routable subnet for the backplane network. To avoid split routing, do not use a routable subnet for the backplane network.

Make sure that the backplane subnet has a sufficient number of IP addresses. Two IP addresses are required per node. Reconfiguring the backplane to increase the size of the subnet involves cluster downtime, so you might also want to make sure that the subnet can accommodate new nodes in the future.

- In **Netmask**, specify the network mask.
- If you want to assign the interfaces on the network to a VLAN, specify the VLAN ID in the **VLAN ID** field.

Nutanix strongly recommends configuring a separate VLAN. If you do not specify a VLAN ID, AOS applies the untagged VLAN on the virtual switch.

- In the **Virtual Switch** list, select the virtual switch you created for the backplane traffic.

- d. Click **Verify and Save**.

If the network settings you specified pass validation, the backplane network is created and the CVMs perform a reboot in a rolling fashion (one at a time), after which the services use the new backplane network. The progress of this operation can be tracked on the Prism tasks page.

4. Log on to a CVM in the cluster with SSH and stop Acropolis cluster-wide:

```
nutanix@cvm$ allssh genesis stop acropolis
```

5. Restart Acropolis cluster-wide:

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


6. Remove all nodes from the maintenance mode.

- a. From any CVM in the cluster, run the following command to exit the AHV host from the maintenance mode:

```
nutanix@cvm$ acli host.exit_maintenance_mode hypervisor-IP-address
```

Replace *hypervisor-IP-address* with the IP address of the node.

Example.

```
nutanix@CVM$ acli host.exit_maintenance_mode 197.116.6.79
ExitMaintenanceMode: pending
ExitMaintenanceMode: complete
```

This command migrates (live migration) all the VMs that were previously running on the host back to the host.

- b. Verify if the node has exited the maintenance mode:

```
nutanix@cvm$ acli host.get hypervisor-IP-address
```

Replace *hypervisor-IP-address* with the IP address of the node.

In the output that is displayed, ensure that **node_state** equals to **kAcropolisNormal** or **AcropolisNormal** and **schedulable** equals to **True**.

Example.

```
nutanix@cvm$ acli host.get 197.116.6.79
197.116.6.79 {
  cpu_usage_ppm: 192941
  cvm_memory_size_bytes: 21474836480
  cvm_num_vcpus: 8
  cvm_num_vnics: 3
  cvm_uuid: "e96dbef0-d425-4926-9fe8-4d10b1a69902"
  host_overhead_bytes: 4957721854
  logical_timestamp: 23
  max_mem_ha_reserved_bytes: 0
  mem_assigned_bytes: 0
  mem_usage_bytes: 26654856446
  memory_size_bytes: 269842644992
  node_state: "AcropolisNormal"
  num_cpus: 32
  pool_size_bytes: 0
  schedulable: True
  uuid: "cc50ea78-49ce-4767-90b3-a0a2ef891446"
}
```

7. Power on the guest VMs from the Prism Element web console.

Physically Isolating the Backplane Traffic on an ESXi Cluster

Before you begin

On the ESXi hosts, do the following:

1. Create a vSwitch for the backplane traffic.
2. From vSwitch0, remove the uplinks (physical NICs) that you want to add to the vSwitch you created for the backplane traffic.
3. On the backplane vSwitch, create one port group for the CVM and another for the host. Ensure that at least one uplink is present in the Active Adaptors list for each port group if you have overridden the failover order.

Note: You don't need to manually create the VMkernel adapter for backplane segmentation since the workflow takes care of creating it on the port group selected as the host port group.

See the ESXi documentation for instructions about how to perform these tasks.

Note: Before you perform the following procedure, ensure that the uplinks you added to the vSwitch are in the UP state.

About this task

Perform the following procedure to physically segment the backplane traffic.

Procedure

1. Log on to the Prism web console, click the gear icon in the top-right corner, and then click **Network Configuration** in the **Settings** page.
2. On the **Internal Interfaces** tab, in the **Backplane LAN** row, click **Configure**.
3. In the **Backplane LAN** dialog box, do the following:
 - a. In **Subnet IP**, specify a non-routable subnet that is different from the subnet used by the ESXi host and CVMs.

The AOS CVM default route uses the CVM eth0 interface, and there is no route on the backplane interface. Therefore, Nutanix recommends only using a non-routable subnet for the backplane network. To avoid split routing, do not use a routable subnet for the backplane network.

Make sure that the subnet has a sufficient number of IP addresses. Two IP addresses are required per node. Reconfiguring the backplane to increase the size of the subnet involves cluster downtime, so you might also want to make sure that the subnet can accommodate new nodes in the future.
 - b. In **Netmask**, specify the network mask.
 - c. If you want to assign the interfaces on the network to a VLAN, specify the VLAN ID in the **VLAN ID** field.

Nutanix strongly recommends configuring a separate VLAN. If you do not specify a VLAN ID, AOS applies the default VLAN on the virtual switch.
 - d. In the **Host Port Group** list, select the port group you created for the host.

A port group can be a standard vSwitch port group, a distributed vSwitch port group or a NSX segment reflected as a distributed port group in the vCenter.
 - e. In the **CVM Port Group** list, select the port group you created for the CVM.

A port group can be a standard vSwitch port group, a distributed vSwitch port group or a NSX segment reflected as distributed port group in the vCenter.

Note:

Nutanix clusters support both vSphere Standard Switches and vSphere Distributed Switches with either normal portgroups or NSX segment backed portgroups. However, you must mandatorily configure only one type of virtual switches in one cluster. Configure all the backplane and management traffic in one cluster on either vSphere Standard Switches or vSphere Distributed Switches. Do not mix Standard and Distributed vSwitches on a single cluster. Ensure to have the same port group type (normal or NSX segment) on all nodes of the cluster.

4. Click **Verify and Save**.

If the network settings you specified pass validation, the backplane network is created and the CVMs perform a reboot in a rolling fashion (one at a time), after which the services use the new backplane network. The progress of this operation can be tracked on the Prism tasks page.

Physically Isolating the Backplane Traffic on a Hyper-V Cluster

Before you begin

On the Hyper-V hosts, do the following:

1. Create a Hyper-V Virtual Switch for the backplane traffic.
2. From the default External Switch, remove the uplinks (physical NICs) that you want to add to the backplane Virtual Switch you created for the backplane traffic.
3. On the backplane Virtual Switch, create a subnet and, optionally, assign a VLAN.

See the Hyper-V documentation on the Microsoft portal for instructions about how to perform these tasks.

Note: Before you perform the following procedure, ensure that the uplinks you added to the backplane Virtual Switch are in the UP state.

About this task

Perform the following procedure to physically segment the backplane traffic.

Procedure

1. Log on to the Prism web console, click the gear icon in the top-right corner, and then click **Network Configuration** in the **Settings** page.
2. On the **Internal Interfaces** tab, in the **Backplane LAN** row, click **Configure**.
3. In the **Backplane LAN** dialog box, do the following:
 - a. In **Subnet IP**, specify a non-routable subnet that is different from the subnet used by the Hyper-V host and CVMs.

The AOS CVM default route uses the CVM eth0 interface, and there is no route on the backplane interface. Therefore, Nutanix recommends only using a non-routable subnet for the backplane network. To avoid split routing, do not use a routable subnet for the backplane network.

Make sure that the subnet has a sufficient number of IP addresses. Two IP addresses are required per node. Reconfiguring the backplane to increase the size of the subnet involves cluster downtime, so you might also want to make sure that the subnet can accommodate new nodes in the future.
 - b. In **Netmask**, specify the network mask.
 - c. If you want to assign the interfaces on the network to a VLAN, specify the VLAN ID in the **VLAN ID** field.

Nutanix strongly recommends configuring a separate VLAN. If you do not specify a VLAN ID, AOS applies the default VLAN on the virtual switch.
 - d. In the **Bridge** list, select the Hyper-V switch you created for the backplane traffic.

4. Click **Verify and Save**.

If the network settings you specified pass validation, the backplane network is created and the CVMs perform a reboot in a rolling fashion (one at a time), after which the services use the new backplane network. The progress of this operation can be tracked on the Prism tasks page.

Note: Segmenting backplane traffic can involve up to two rolling reboots of the CVMs. The first rolling reboot is done to move the backplane interface (eth2) of the CVM to the selected port group or virtual switch. This is done only for CVM(s) whose backplane interface is not already connected to the selected port group or bridge virtual switch. The second rolling reboot is done to migrate the cluster services to the newly configured backplane interface.

Reconfiguring the Backplane Network

Backplane network reconfiguration is a CLI-driven procedure that you perform on any one of the CVMs in the cluster. The change is propagated to the remaining CVMs.

About this task

Caution: At the end of this procedure, the cluster stops and restarts, even if only the VLAN is changed, and therefore involves cluster downtime.

To reconfigure the cluster, do the following:

Procedure

1. Log on to any CVM in the cluster using SSH.
2. Reconfigure the backplane network.

```
nutanix@cvm$ backplane_ip_reconfig [--backplane_vlan=vlan-id] \  
[--backplane_ip_pool=ip_pool_name]
```

Replace `vlan-id` with the new VLAN ID, and `ip_pool_name` with the newly created backplane IP pool.

See [Configuring Backplane IP Pool](#) on page 176 to create a backplane IP pool.

For example, reconfigure the backplane network to use VLAN ID 10 and newly created backplane IP pool.

```
nutanix@cvm$ backplane_ip_reconfig --backplane_vlan=10 \  
--backplane_ip_pool=NewBackplanePool
```

Output similar to the following is displayed:

```
This operation will do a 'cluster stop', resulting in disruption of  
cluster services. Do you still want to continue? (Type "yes" (without quotes)  
to continue)  
Type yes to confirm that you want to reconfigure the backplane network.
```

Caution: During the reconfiguration process, you might receive an error message similar to the following.

Failed to reach a node.

You can safely ignore this error message and therefore do not stop the script manually.

Note: The `backplane_ip_reconfig` command is not supported on ESXi clusters with vSphere Distributed Switches. To reconfigure the backplane network on a vSphere Distributed Switch setup, disable the backplane network (see [Disabling Network Segmentation on an ESXi and Hyper-V Clusters](#) on page 165) and enable again with a different subnet or VLAN.

3. Type `yes` to confirm that you want to reconfigure the backplane network.
The reconfiguration procedure takes a few minutes and includes a cluster restart. If you type anything other than `yes`, network reconfiguration is aborted.
4. After the process completes, verify that the backplane was reconfigured.
 - a. Verify that the IP addresses of the `eth2` interfaces on the CVM are set correctly.

```
nutanix@cvm$ svmips -b
```

Output similar to the following is displayed:

```
172.30.25.1 172.30.25.3 172.30.25.5
```

- b. Verify that the IP addresses of the backplane interfaces of the hosts are set correctly.

```
nutanix@cvm$ hostips -b
```

Output similar to the following is displayed:

```
172.30.25.2 172.30.25.4 172.30.25.6
```

The `svmips` and `hostips` commands, when used with the option `-b`, display the IP addresses assigned to the interfaces on the backplane.

Disabling Network Segmentation on an ESXi and Hyper-V Clusters

Backplane network reconfiguration is a CLI-driven procedure that you perform on any one of the CVMs in the cluster. The change is propagated to the remaining CVMs.

About this task

Procedure

1. Log on to any CVM in the cluster using SSH.
2. Disable the backplane network.
 - Use this CLI to disable network segmentation on an ESXi and Hyper-V cluster:

```
nutanix@cvm$ network_segmentation --backplane_network --disable
```

Output similar to the following appears:

```
Operation type : Disable
Network type : kBackplane
Params : {}
Please enter [Y/y] to confirm or any other key to cancel the operation
```

Type `Y` or `y` to confirm that you want to reconfigure the backplane network.

If you type `Y` or `y`, network segmentation is disabled and the controller VMs (CVMs) restart in a rolling manner, one CVM at a time. If you type anything other than `Y` or `y`, network segmentation is not disabled.

This method does not involve cluster downtime.

3. Verify that network segmentation was successfully disabled. You can verify this in one of two ways:

- » Verify that the backplane is disabled.

```
nutanix@cvm$ network_segment_status
```

Output similar to the following is displayed:

```
2017-11-23 06:18:23 INFO zookeeper_session.py:110 network_segment_status is attempting to connect to Zookeeper
```

Network segmentation is disabled

- » Verify that the commands to show the backplane IP addresses of the CVMs and hosts list the same management IP addresses. Run the `svmips` and `hostips` command with and without the `-b` option, and then compare the IP addresses shown in the output.

Important:

```
nutanix@cvm$ svmips
192.127.3.2 192.127.3.3 192.127.3.4
nutanix@cvm$ svmips -b
192.127.3.2 192.127.3.3 192.127.3.4
nutanix@cvm$ hostips
192.127.3.5 192.127.3.6 192.127.3.7
nutanix@cvm$ hostips -b
192.127.3.5 192.127.3.6 192.127.3.7
```

In the above example, the outputs of the `svmips` and `hostips` commands with and without the `-b` option are the same, indicating that the backplane network segmentation is disabled.

Disabling Network Segmentation on an AHV Cluster

About this task

You perform backplane network reconfiguration procedure on any one of the CVMs in the cluster. The change propagates to the remaining CVMs.

Procedure

1. Shut down all the guest VMs in the cluster from within the guest OS or use the Prism Element web console.

2. Place all nodes of a cluster into the maintenance mode.

- a. Use SSH to log on to a Controller VM in the cluster
- b. Determine the IP address of the node you want to put into the maintenance mode:

```
nutanix@cvm$ acli host.list
```

Note the value of Hypervisor IP for the node you want to put in the maintenance mode.

- c. Put the node into the maintenance mode:

```
nutanix@cvm$ acli host.enter_maintenance_mode hypervisor-IP-address [wait="{ true  
| false }" ] [non_migratable_vm_action="{ acpi_shutdown | block }" ]
```

Note: Never put Controller VM and AHV hosts into maintenance mode on single-node clusters. It is recommended to shutdown user VMs before proceeding with disruptive changes.

Replace *host-IP-address* with either the IP address or host name of the AHV host you want to shut down.

The following are optional parameters for running the `acli host.enter_maintenance_mode` command:

- **wait**
- **non_migratable_vm_action**

Do not continue if the host has failed to enter the maintenance mode.

- d. Verify if the host is in the maintenance mode:

```
nutanix@cvm$ acli host.get host-ip
```

In the output that is displayed, ensure that **node_state** equals to **EnteredMaintenanceMode** and **schedulable** equals to **False**.

3. Disable backplane network segmentation from the Prism Web Console.

- a. Log on to the Prism web console, click the gear icon in the top-right corner, and then click **Network Configuration** under the **Settings**.
- b. In the **Internal Interfaces** tab, in the **Backplane LAN** row, click **Disable**.

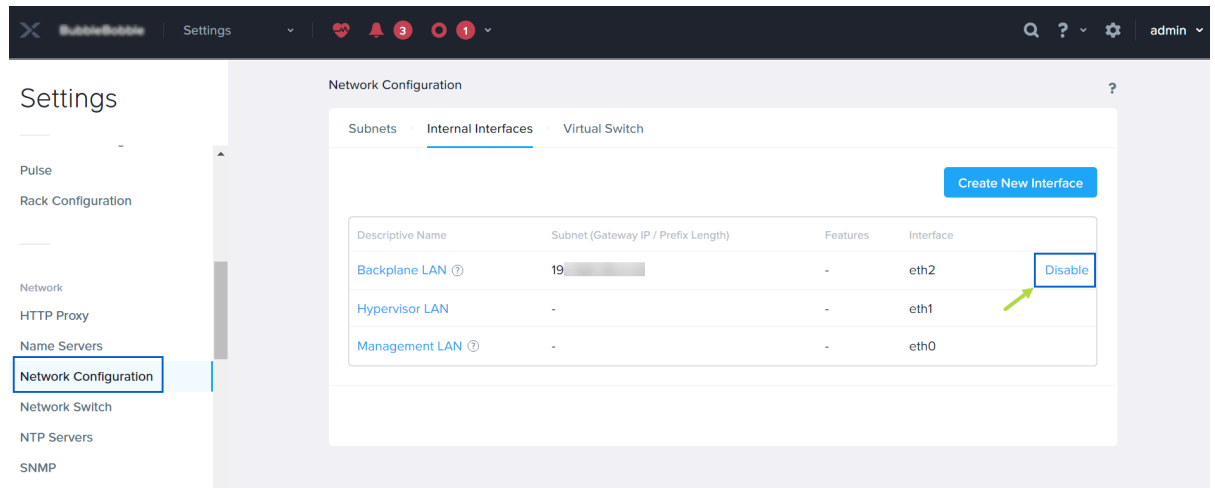


Figure 74: Disable Network Configuration

- c. Click **Yes** to disable Backplane LAN.

This involves a rolling reboot of CVMs to migrate the cluster services back to the external interface.

4. Log on to a CVM in the cluster with SSH and stop Acropolis cluster-wide:

```
nutanix@cvm$ allssh genesis stop acropolis
```

5. Restart Acropolis cluster-wide:

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

6. Remove all nodes from the maintenance mode.

- a. From any CVM in the cluster, run the following command to exit the AHV host from the maintenance mode:

```
nutanix@cvm$ acli host.exit_maintenance_mode host-ip
```

Replace *host-ip* with the new IP address of the host.

This command migrates (live migration) all the VMs that were previously running on the host back to the host.

- b. Verify if the host has exited the maintenance mode:

```
nutanix@cvm$ acli host.get host-ip
```

In the output that is displayed, ensure that **node_state** equals to **kAcropolisNormal** or **AcropolisNormal** and **schedulable** equals to **True**.

7. Power on the guest VMs from the Prism Element web console.

Service-Specific Traffic Isolation

Isolating the traffic associated with a specific service is a two-step process. The process is as follows:

- Configure the networks and uplinks on each host manually. Prism only creates the VNIC that the service requires, and it places that VNIC on the bridge or port group that you specify. Therefore, you must manually create the bridge or /port group on each host and add the required physical NICs as uplinks to that bridge or port group.
- Configure network segmentation for the service by using Prism. Create an extra VNIC for the service, specify any additional parameters that are required (for example, IP address pools), and the bridge or port group that you want to dedicate to the service.

Isolating Service-Specific Traffic

Before you begin

- Ensure to configure each host as described in [Configuring the Network on an AHV Host](#) on page 136.
- Review [Prerequisites](#) on page 135.

About this task

To isolate a service to a separate virtual network, do the following:

Procedure

1. Log on to the Prism web console and click the gear icon at the top-right corner of the page.
2. In the left pane, click **Network Configuration**.
3. In the details pane, on the **Internal Interfaces** tab, click **Create New Interface**.
The **Create New Interface** dialog box is displayed.

4. On the **Interface Details** tab, do the following:

- a. Specify a descriptive name for the network segment.
- b. (On AHV) Optionally, in **VLAN ID**, specify a VLAN ID.
Make sure that the VLAN ID is configured on the physical switch.
- c. In **Bridge** (on AHV) or **CVM Port Group** (on ESXi), select the bridge or port group that you created for the network segment.
- d. To specify an IP address pool for the network segment, click **Create New IP Pool**, and then, in the **IP Pool** dialog box, do the following:
 - In **Name**, specify a name for the pool.
 - In **Netmask**, specify the network mask for the pool.
 - Click **Add an IP Range**, specify the start and end IP addresses in the **IP Range** dialog box that is displayed.
 - Use **Add an IP Range** to add as many IP address ranges as you need.

Note: Add at least $n+1$ IP addresses in an IP range considering n is the number of nodes in the cluster.

- Click **Save**.
- Use **Add an IP Pool** to add more IP address pools. You can use only one IP address pool at any given time.
- Select the IP address pool that you want to use, and then click **Next**.

Note: You can also use an existing unused IP address pool.

5. On the **Feature Selection** tab, do the following:

You cannot enable network segmentation for multiple services at the same time. Complete the configuration for one service before you enable network segmentation for another service.

- a. Select the service whose traffic you want to isolate.
- b. Configure the settings for the selected service.
The settings on this page depend on the service you select. For information about service-specific settings, see [Service-Specific Settings and Configurations](#) on page 172.
- c. Click **Save**.

6. In the **Create Interface** dialog box, click **Save**.

The CVMs are rebooted multiple times, one after another. This procedure might trigger more tasks on the cluster. For example, if you configure network segmentation for disaster recovery, the firewall rules are added on the CVM to allow traffic on the specified ports through the new CVM interface and updated when a new recovery cluster is added or an existing cluster is modified.

What to do next

See [Service-Specific Settings and Configurations](#) on page 172 for any additional tasks that are required after you segment the network for a service.

Modifying Network Segmentation Configured for a Service

To modify network segmentation configured for a service, you must first disable network segmentation for that service and then create the network interface again for that service with the new IP address pool and VLAN.

About this task

For example, if the interface of the service you want to modify is ntnx0, after the reconfiguration, the same interface (ntnx0) is assigned to that service if that interface is not assigned to any other service. If ntnx0 is assigned to another service, a new interface (for example ntnx1) is created and assigned to that service.

Perform the following to reconfigure network segmentation configured for a service.

Procedure

1. Disable the network segmentation configured for a service by following the instructions in [Disabling Network Segmentation Configured for a Service](#) on page 171.
2. Create the network again by following the instructions in [Isolating Service-Specific Traffic](#) on page 169.

Disabling Network Segmentation Configured for a Service

To disable network segmentation configured for a service, you must disable the dedicated vNIC. Disabling network segmentation frees up the name of the vNIC. Disabling network segmentation frees up the vNIC's name. The free name is reused in a subsequent network segmentation configuration.

About this task

At the end of this procedure, the cluster performs a rolling restart. Disabling network segmentation might also disrupt the functioning of the associated service. To restore normal operations, you might have to perform other tasks immediately after the cluster has completed the rolling restart. For information about the follow-up tasks, see [Service-Specific Settings and Configurations](#) on page 172.

To disable the network segmentation configured for a service, do the following:

Procedure

1. Log on to the Prism web console and click the gear icon at the top-right corner of the page.
2. In the left pane, click **Network Configuration**.
3. On the **Internal Interfaces** tab, for the interface that you want to disable, click **Disable**.

Note: The defined IP address pool is available even after disabling the network segmentation.

Deleting a vNIC Configured for a Service

If you disable network segmentation for a service, the vNIC for that service is not deleted. AOS reuses the vNIC if you enable network segmentation again. However, you can manually delete a vNIC by logging into any CVM in the cluster with SSH.

Before you begin

Ensure that the following prerequisites are met before you delete the vNIC configured for a Service:

- Disable the network segmentation configured for a service by following the instructions in [Disabling Network Segmentation Configured for a Service](#) on page 171.
- Observe the Limitation specified in [Limitation for vNIC Hot-Unplugging](#) topic in *AHV Admin Guide*.

you

About this task

Perform the following to delete a vNIC.

Procedure

1. Log on to any CVM in the cluster with SSH.
2. Delete the vNIC.

```
nutanix@cvm$ network_segmentation --service_network --interface="interface-name" --delete
```

Replace `interface-name` with the name of the interface you want to delete. For example, ntnx0.

Service-Specific Settings and Configurations

The following sections describe the settings required by the services that support network segmentation.

Nutanix Volumes

Network segmentation for Volumes also requires you to migrate iSCSI client connections to the new segmented network. If you no longer require segmentation for Volumes traffic, you must also migrate connections back to eth0 after disabling the vNIC used for Volumes traffic.

You can create up to two different networks for Nutanix Volumes with different IP pools, VLANs, and data services IP addresses. For example, you can create two iSCSI networks, one for production and one for non-production traffic, on the same Nutanix cluster.

Follow the instructions in [Isolating Service-Specific Traffic](#) on page 169 again to create the second network for Volumes after you create the first network.

Table 10: Settings to be Specified When Configuring Traffic Isolation

Parameter or Setting	Description
Virtual IP	(Optional) Virtual IP address for the service. If specified, the IP address must be picked from the specified IP address pool. If not specified, an IP address from the specified IP address pool is selected for you.

Parameter or Setting	Description
Client Subnet	<p>The network (in CIDR notation) that hosts the iSCSI clients. Required If the vNIC created for the service on the CVM is not on the same network as the clients.</p> <p>You can specify multiple client subnets while configuring network segmentation for Volumes in the CLI from a CVM.</p> <p>For example: -- <code>client_subnets="10.2.2.0/24,10.2.3.0/24,10.2.4.0/24"</code></p> <p>For more information on the usage of the CLI, run the <code>nutanix@cvm\$ network_segmentation --help</code> command.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>Note:</p> <p>You can specify only one client subnet while configuring network segmentation for Volumes in the Prism Element Web Console.</p> </div>
Gateway	Gateway to the subnetwork that hosts the iSCSI clients. Required If you specify the client subnet.

Migrating iSCSI Connections to the Segmented Network

After you enable network segmentation for Volumes, you must manually migrate connections from existing iSCSI clients to the newly segmented network.

Before you begin

Make sure that the task for enabling network segmentation for the service succeeds.

About this task

Note: Even though support is available to run iSCSI traffic on both the segmented and management networks at the same time, Nutanix recommends that you move the iSCSI traffic for guest VMs to the segmented network to achieve true isolation.

To migrate iSCSI connections to the segmented network, do the following:

Procedure

1. Log out from all the clients connected to iSCSI targets that are using CVM eth0 or the Data Service IP address.
2. Optionally, remove all the discovery records for the Data Services IP address (DSIP) on eth0.
3. If the clients are allowlisted by their IP address, remove the client IP address that is on the management network from the allowlist, and then add the client IP address on the new network to the allowlist.

```
nutanix@cvm$ acli vg.detach_external vg_name initiator_network_id=old_vm_IP
nutanix@cvm$ acli vg.attach_external vg_name initiator_network_id=new_vm_IP
```

Replace `vg_name` with the name of the volume group and `old_vm_IP` and `new_vm_IP` with the old and new client IP addresses, respectively.

4. Discover the virtual IP address specified for Volumes.
5. Connect to the iSCSI targets from the client.

Migrating Existing iSCSI Connections to the Management Network (Controller VM eth0)

About this task

To migrate existing iSCSI connections to eth0, do the following:

Procedure

1. Log out from all the clients connected to iSCSI targets using the CVM vNIC dedicated to Volumes.
2. Remove all the discovery records for the DSIP on the new interface.
3. Discover the DSIP for eth0.
4. Connect the clients to the iSCSI targets.

Disaster Recovery with Protection Domains

The settings for configuring network segmentation for disaster recovery apply to all Asynchronous, NearSync, and Metro Availability replication schedules. You can use disaster recovery with Asynchronous, NearSync, and Metro Availability replications only if both the primary site and the recovery site is configured with Network Segmentation. Before enabling or disabling the network segmentation on a host, disable all the disaster recovery replication schedules running on that host.

Table 11: Settings to be Specified When Configuring Traffic Isolation

Parameter or Setting	Description
Virtual IP	(Optional) Virtual IP address for the service. If specified, the IP address must be picked from the specified IP address pool. If not specified, an IP address from the specified IP address pool is selected for you.
<div> Note: Virtual IP address is different from the external IP address and the data services IP address of the cluster. </div>	
Gateway	Gateway to the subnetwork (subnet).

Remote Site Configuration

After configuring network segmentation for disaster recovery, configure remote sites at both locations. You also need to reconfigure remote sites if you disable network segmentation.

For information about configuring remote sites, see [Remote Site Configuration](#) in the *Data Protection and Recovery with Prism Element Guide*.

Segmenting a Stretched Layer 2 Network for Disaster Recovery

A stretched Layer 2 network configuration allows the source and remote metro clusters to be in the same broadcast domain and communicate without a gateway.

About this task

You can enable network segmentation for disaster recovery on a stretched Layer 2 network that does not have a gateway. A stretched Layer 2 network is usually configured across physically remote clusters, such as a metro availability cluster deployment. A stretched Layer 2 network allows the source and remote clusters to be configured in the same broadcast domain without the usual gateway.

See *AOS Release Notes* for minimum AOS version required to configure a stretched Layer 2 network.

To configure a network segment as a stretched L2 network, do the following.

Procedure

1. Log on to any CVM in the cluster using SSH.

2. Run the following command:

- **On a cluster running ESXi hypervisor**

```
nutanix@cvm$ network_segmentation --service_network --service_name=kDR --  
ip_pool=DR-ip-pool-name  
--service_vlan=DR-vlan-id --desc_name=Description --  
host_physical_network=portgroup  
--stretched_metro
```

- **On a cluster running AHV hypervisor**

```
nutanix@cvm$ network_segmentation --service_network --service_name=kDR --  
ip_pool=DR-ip-pool-name  
--service_vlan=DR-vlan-id --desc_name=Description --host_virtual_switch=virtual-  
switch  
--stretched_metro
```

Replace the following: (See [Isolating Service-Specific Traffic](#) on page 169 for the information)

- `DR-ip-pool-name` with the name of the **IP Pool** created for the DR service or any existing unused IP address pool.
- `DR-vlan-id` with the **VLAN ID** used for the DR service.
- `Description` with a suitable description of this stretched L2 network segment.
- `portgroup` with the details of the **CVM Port Group** used for the DR service in the ESXi hypervisor.
- `virtual-switch` with the details of the **Virtual Switch** used for the DR service in the AHV hypervisor.

This example shows how to configure network segment as a stretched L2 network on a cluster running ESXi hypervisor:

```
nutanix@cvm$ network_segmentation --service_network --service_name=kDR --  
ip_pool=DR_pool  
--service_vlan=124 --desc_name="L2 Strech for ESXi" --  
host_physical_network=portgroup0
```

```
--stretched_metro
```

For example, configure network segment as a stretched L2 network on a cluster running AHV hypervisor:

```
nutanix@cvm$ network_segmentation --service_network --service_name=kDR --  
ip_pool=DR_pool  
--service_vlan=124 --desc_name="L2 Strech for AHV" --host_virtual_switch=vs0  
--stretched_metro
```

For more information about the `network_segmentation` command, see the [Command Reference](#) guide.

Nutanix Disaster Recovery

The settings for configuring network segmentation for Nutanix Disaster Recovery apply to all Asynchronous, NearSync, and Synchronous replication schedules. For detailed information about network segmentation for Nutanix Disaster Recovery, see *Network Segmentation* in the *Nutanix Disaster Recovery Guide*.

Table 12: Settings to be Specified When Configuring Traffic Isolation

Parameter or Setting	Description
Virtual IP	(Optional) Virtual IP address for the service. If specified, the IP address must be picked from the specified IP address pool. If not specified, an IP address from the specified IP address pool is selected for you. Note: Virtual IP address is different from the external IP address and the data services IP address of the cluster.
Gateway	Gateway to the subnetwork.

Configuring Backplane IP Pool

This procedure shows how to create an IP pool for backplane interfaces using the new CLI.

About this task

Network Segmentation for backplane traffic previously required an entire subnet even if a cluster has a small number of nodes. This resulted in inefficient use of IP addresses. The backplane IP pool feature enables you to provide a small IP pool instead of an entire subnet.

You can create an IP address pool using the new `network_segmentation ip_pool` command. The named IP pool includes one or more IP ranges. For example, an IP address from 172.16.1.100 to 172.16.1.105 can be one IP range and 172.16.1.120 to 172.16.1.125 can be another IP range within the same named IP pool and same IP subnet.

At present in the Prism interface, there is no option to create an IP address pool for backplane segmentation. However, the Prism interface allows creating small IP address pools for service-specific traffic such as Volumes and DR. You can use the new `network_segmentation ip_pool` CLI to create IP address pools for Backplane, Volumes, and DR as well. You can also manage (edit, delete, and update) IP address pools that are created for Backplane, Volumes and DR using the new CLI.

Procedure

1. Log on to any CVM in the cluster using SSH

2. Create a new IP pool and define the IP ranges

```
nutanix@cvm$ network_segmentation --ip_pool_name=IP-Pool-name --  
ip_pool_netmask=netmask  
--ip_ranges="['First-IP-Address', 'Last-IP-Address'], ('First-IP-Address', 'Last-IP-  
Address')]"  
ip_pool create
```

Replace:

- *IP-Pool-name* with a user defined IP pool name
- *Netmask* with a network mask in dotted decimal mask notation
- *First-IP-Address* with the first IP Address in the range
- *Last-IP-Address* with the last IP Address in the same range

For example:

```
nutanix@cvm$ network_segmentation --ip_pool_name=BackplanePool --  
ip_pool_netmask=255.255.255.0  
--ip_ranges="['172.16.1.100', '172.16.1.105'], ('172.16.1.120', '172.16.1.125')]"  
ip_pool create
```

Enabling Backplane Network Segmentation on a Mixed Hypervisor Cluster

This procedure shows how to enable Backplane Network Segmentation on a mixed hypervisor cluster.

About this task

You can enable Backplane Network Segmentation on a mixed hypervisor cluster containing:

- ESXi and AHV storage only nodes.
- Hyper-V and AHV storage only nodes.

Procedure

1. Log on to any CVM in the cluster using SSH.
2. Enable network segmentation for backplane traffic

On a cluster containing ESXi and AHV storage only nodes:

```
nutanix@cvm$ network_segmentation --backplane_network  
--ip_pool=IP-pool-name  
--backplane_vlan=VLAN-ID  
[--esx_host_physical_network=ESXi-host-portgroup-name ]  
[--esx_cvm_physical_network=ESXi-cvm-portgroup-name ]  
[--ahv_host_physical_network=AHV-network-name ]
```

On a cluster containing Hyper-V and AHV storage only nodes:

```
nutanix@cvm$ network_segmentation --backplane_network  
--ip_pool=IP-pool-name  
--backplane_vlan=VLAN-ID  
[--hyperv_host_physical_network=HyperV-host-network-name ]  
[--ahv_host_physical_network=AHV-network-name ]
```

In the above command replace:

- *IP-Pool-name* with a user defined IP pool name
- *VLAN-ID* with a backplane VLAN ID

- `ESXi-host-portgroup-name` with the ESXi host network name
- `ESXi-cvm-portgroup-name` with the ESXi CVM network name
- `AHV-network-name` with the AHV storage only node bridge name
- `HyperV-host-network-name` with the Hyper-V switch name

For example, enable network segmentation on a mixed hypervisor containing ESXi and AHV storage only nodes:

```
nutanix@cvm$ network_segmentation --backplane_network
--ip_pool=BackplanePool
--backplane_vlan=1234
--esx_host_physical_network=host-pg
--esx_cvm_physical_network=cvm-pg
--ahv_host_physical_network=br1
```

Updating Backplane Portgroup

You can update the backplane portgroups that are assigned to CVM and host nodes. Earlier, to change a portgroup that is assigned to a CVM and host, you had to disable Network Segmentation and re-enable with new portgroups.

This feature is only supported on a cluster running ESXi hypervisor.

Updating backplane portgroups helps you to:

- Move from one vSphere Standard Switch (VSS) portgroup to another VSS portgroup within the same virtual standard switch
- Move from one VSS portgroup to another VSS portgroup in a different Virtual Standard Switch
- Move from a VSS portgroup to a vSphere Distributed Switch (VDS) portgroup
- Move from a VDS portgroup to a VSS

Note:

For renaming existing VSS or VDS portgroups, you must manually perform the rename operation from the vCenter application or use the ESXi CLI and run the update operation with the new portgroup. This is to ensure the configuration stored in the Nutanix internal database is up to date.

Limitations of Updating Backplane Portgroup

Consider the following limitations before updating backplane portgroups:

- This feature is not supported on clusters running on the AHV and Hyper-V hypervisors.
- This feature does not support updating any other configuration such as VLAN ID, and IP address.

Note: This feature does not perform any network validation on the new portgroups. Hence the user must ensure the portgroup settings are accurate before proceeding with the portgroup update operation. If the settings are not accurate, the CVM on that node may not be able to communicate with its peers and this results in a stuck rolling reboot.

Updating Backplane Portgroup

This procedure shows how to update backplane portgroups:

About this task

You can update the backplane portgroups that are assigned to CVM and host nodes

Procedure

1. Log on to any CVM in the cluster using SSH
2. Update the CVM and host portgroups

```
nutanix@cvm$ network_segmentation --backplane_network
--host_physical_network=new-host-portgroup-name
--cvm_physical_network=new-cvm-portgroup-name
--update
```

In the above command replace:

- `new-host-portgroup-name` with the new host portgroup name
- `new-cvm-portgroup-name` with the new CVM portgroup name

For example:

```
nutanix@cvm$ network_segmentation --backplane_network
--host_physical_network=new-bp-host-pgroup
--cvm_physical_network=new-bp-cvm-pgroup
--update
```

For creating a port group, see [Creating Port Groups on the Distributed Switch](#) in *vSphere Administration Guide for Acropolis*.

IP Address Customization for each CVM and Host

This procedure shows how to create custom IP addresses for each CVM and host for network segmentation.

About this task

IP Address customization for each CVM and host feature enables you to allocate an IP address manually to a CVM and host. This helps in maintaining similarity between external and segmented IP addresses. This feature is supported while configuring backplane segmentation, service specific traffic isolation for Volumes, and service specific traffic isolation for Disaster Recovery.

Procedure

1. Create a JSON file with a mapping of IP Addresses

You must manually define the mapping of the CVM external IP address to the new segmented IP address in a JSON file. The segmented IP addresses should belong to the same IP address pool that is created from the Prism UI or the CLI command before starting the Network Segmentation operation. You can create and save the JSON file in any CVM in the cluster.

Here is the example of JSON file format:

- JSON file format for Backplane Segmentation:

```
{
  "svmips": {
    `cvm_external_ip1`: `cvm_backplane_ip1`,
    `cvm_external_ip2`: `cvm_backplane_ip2`,
    `cvm_external_ip3`: `cvm_backplane_ip3`
  },
  "hostips": {
```

```

    `host_external_ip1`: `host_backplane_ip1`,
    `host_external_ip2`: `host_backplane_ip2`,
    `host_external_ip3`: `host_backplane_ip3`
  }
}

```

For example:

```

{
  "svnmips": {
    "10.47.240.141": "172.16.10.141",
    "10.47.240.142": "172.16.10.142",
    "10.47.240.143": "172.16.10.143"
  },
  "hostips": {
    "10.47.240.137": "172.16.10.137",
    "10.47.240.138": "172.16.10.138",
    "10.47.240.139": "172.16.10.139"
  }
}

```

- JSON file format for Service Segmentation:

```

{
  "svnmips": {
    `cvm_external_ip1`: `cvm_service_ip1`,
    `cvm_external_ip2`: `cvm_service_ip2`,
    `cvm_external_ip3`: `cvm_service_ip3`
  }
}

```

For example:

```

{
  "svnmips": {
    "10.47.240.141": "10.47.6.141",
    "10.47.240.142": "10.47.6.142",
    "10.47.240.143": "10.47.6.143"
  }
}

```

2. Log on to the CVM in the cluster where the JSON file exists using SSH
3. Enable Service Specific Traffic Isolation for Volumes using the JSON file

```

nutanix@CVM:~$ network_segmentation --service_network
--ip_pool=pool1 --desc_name="Volumes Seg 1"
--service_name=kVolumes
--host_physical_network=dv-volumes-network-1
--service_vlan=151
--ip_map_filepath=/home/nutanix/ip_map.json

```

Enabling Physical Backplane Segmentation on Hyper-V Using CLI

This procedure shows how to enable physical backplane segmentation on a cluster containing Hyper-V node using CLI.

About this task

Physical backplane segmentation support is now available on a cluster containing Hyper-V nodes. Earlier, the support was available on AHV and ESXi nodes.

Procedure

1. Log on to any CVM in the cluster using SSH
2. Enable backplane segmentation on a Hyper-V node

```
nutanix@CVM:~$ network_segmentation --backplane_network  
--ip_pool=IP-Pool-name  
--backplane_vlan=VLAN-ID  
--host_physical_network=hyperv_host_physical_network
```

In the above command replace:

- `IP-Pool-name` with a user defined IP pool name
- `VLAN-ID` with a backplane VLAN ID
- `hyperv_host_physical_network` with the Hyper-V switch name

For example:

```
nutanix@CVM:~$ network_segmentation --backplane_network  
--ip_pool=BackplanePool  
--backplane_vlan=1234  
--host_physical_network=BackplaneSwitch
```

Network Segmentation during Cluster Expansion

When expanding a cluster:

- If you enable the backplane network segmentation, Prism allocates two IP addresses for every new node from the backplane IP Pool.
- If you enable service-specific traffic isolation, Prism allocates one IP address for every new node from the respective (Volumes or DR) IP pools.
- If enough IP addresses are not available in the specified network, the Prism Element web console displays a failure message in the tasks page. To add more IP ranges to the IP pool, see [Configuring Backplane IP Pool](#) on page 176.
- If you cannot add more IPs to the IP pool, then reconfigure that specific network segmentation. For more information about how to reconfigure the network, see [Reconfiguring the Backplane Network](#) on page 164.
- The network settings on the physical switch to which the new nodes are connected must be identical to the other nodes in the cluster. New nodes communicate with current nodes using the same VLAN ID for segmented networks. Otherwise, the expand cluster task will fail in the network validation stage.
- After fulfilling the earlier points, you can add nodes to the cluster. For instructions about how to add nodes to your Nutanix cluster, see [Expanding a Cluster](#) in the *Prism Web Console Guide*.

Network Segmentation–Related Changes During an AOS Upgrade

When you upgrade from an AOS version which does not support network segmentation to an AOS version that does, the eth2 interface (used to segregate backplane traffic) is automatically created on each CVM. However, the network remains unsegmented, and the cluster services on the CVM continue to use eth0 until you configure network segmentation.

The vNICs ntnx0, ntnx1, and so on, are not created during an upgrade to a release that supports service-specific traffic isolation. They are created when you configure traffic isolation for a service.

Note:

Do not delete the eth2 interface that is created on the Controller VMs, even if you are not using the network segmentation feature.

Firewall Requirements

The [Ports and Protocols](#) describes detailed port information (like protocol, service description, source, destination, and associated service) for Nutanix products and services. It includes port and protocol information for 1-click upgrades and LCM updates.

Log management

This chapter describes how to configure cluster-wide setting for log-forwarding and documenting the log fingerprint.

Log Forwarding

The Nutanix Controller VM provides a method for log integrity by using a cluster-wide setting to forward all the logs to a central log host. Due to the appliance form factor of the Controller VM, system and audit logs does not support local log retention periods as a significant increase in log traffic can be used to orchestrate a distributed denial of service attack (DDoS).

Nutanix recommends deploying a central log host in the management enclave to adhere to any compliance or internal policy requirement for log retention. In case of any system compromise, a central log host serves as a defense mechanism to preserve log integrity.

Note: The audit in the Controller VM uses the audisp plugin by default to ship all the audit logs to the rsyslog daemon (stored in /home/log/messages). Searching for audispd in the central log host provides the entire content of the audit logs from the Controller VM. The audit daemon is configured with a rules engine that adheres to the auditing requirements of the Operating System Security Requirements Guide (OS SRG), and is embedded as part of the Controller VM STIG.

Use the nCLI to enable forwarding of system, audit, aide, and SCMA logs of all the Controller nodes in a cluster at the required log level. For more information, see [Send Logs to Remote Syslog Server](#) in the *Acropolis Advanced Administration Guide*

Documenting the Log Fingerprint

For forensic analysis, non-repudiation is established by verifying the fingerprint of the public key for the log file entry.

Procedure

1. Login to the CVM.
2. Run the following command to document the fingerprint for each public key assigned to an individual admin.

```
nutanix@cvm$ ssh-keygen -lf /<location of>/id_rsa.pub
```

The fingerprint is then compared to the SSH daemon log entries and forwarded to the central log host (/home/log/secure in the Controller VM).

Note: After completion of the ssh public key inclusion in Prism and verification of connectivity, disable the password authentication for all the Controller VMs and AHV hosts. From the Prism main menu, de-select **Cluster Lockdown configuration > Enable Remote Login with password** check box from the gear icon drop-down list.

SECURITY MANAGEMENT USING NUTANIX COMMAND LINE INTERFACE (NCLI)

Hardening Instructions (nCLI)

This chapter describes how to implement security hardening features for Nutanix AHV and Controller VM.

Hardening AHV

You can use Nutanix Command Line Interface (nCLI) in order to customize the various configuration settings related to AHV as described below.

Table 13: Configuration Settings to Harden the AHV

Description	Command or Settings	Output
Getting the cluster-wide configuration of the SCMA policy.	Run the following command: <pre>nutanix@cvm\$ ncli cluster get-hypervisor-security- config</pre>	<pre>Enable Aide : false Enable Core : false Enable High Strength P... : false Enable Banner : false Schedule : DAILY</pre>
Enabling the Advanced Intrusion Detection Environment (AIDE) to run on a weekly basis.	Run the following command: <pre>nutanix@cvm\$ ncli cluster edit-hypervisor-security- params enable-aide=true</pre>	<pre>Enable Aide : true Enable Core : false Enable High Strength P... : false Enable Banner : false Schedule : DAILY</pre>
Enabling the high-strength password policies (minlen=15, difok=8, maxclassrepeat=4).	Run the following command: <pre>nutanix@cvm\$ ncli cluster edit-hypervisor-security- params \ enable-high-strength- password=true</pre>	<pre>Enable Aide : true Enable Core : false Enable High Strength P... : true Enable Banner : false Schedule : DAILY</pre>
Enabling the defense knowledge consent banner of the US department.	Run the following command: <pre>nutanix@cvm\$ ncli cluster edit-hypervisor-security- params enable-banner=true</pre>	<pre>Enable Aide : true Enable Core : false Enable High Strength P... : true Enable Banner : true Schedule : DAILY</pre>
Changing the default schedule of running the SCMA. The schedule can be hourly, daily, weekly, and monthly.	Run the following command: <pre>nutanix@cvm\$ ncli cluster edit-hypervisor-security- params schedule=hourly</pre>	<pre>Enable Aide : true Enable Core : false Enable High Strength P... : true Enable Banner : true Schedule : HOURLY</pre>
Enabling the settings so that AHV can generate stack traces for any cluster issue.	Run the following command: <pre>nutanix@cvm\$ ncli cluster edit-hypervisor-security- params enable-core=true</pre> <div>Note: Nutanix recommends that Core should not be set to true unless instructed by the Nutanix support team.</div>	<pre>Enable Aide : true Enable Core : true Enable High Strength P... : true Enable Banner : true Schedule : HOURLY</pre>

Description	Command or Settings	Output
When a high governance official needs to run the hardened configuration.	<p>The settings should be as follows:</p> <pre> Enable Aide : true Enable Core : false Enable High Strength P... : true Enable Banner : false Enable SNMPv3 Only : true Schedule : HOURLY Enable Kernel Mitigations : false SSH Security Level : LIMITED Enable Lock Status : true Enable Kernel Core : true </pre>	
When a federal official needs to run the hardened configuration.	<p>The settings should be as follows:</p> <pre> Enable Aide : true Enable Core : false Enable High Strength P... : true Enable Banner : true Enable SNMPv3 Only : true Schedule : HOURLY Enable Kernel Mitigations : false SSH Security Level : LIMITED Enable Lock Status : true Enable Kernel Core : true </pre> <p>Note: A banner file can be modified to support non-DoD customer banners.</p>	

Description	Command or Settings	Output
Backing up the DoD banner file.	<p>Run the following command on the AHV host:</p> <pre>[root@AHV-host ~]# cp -a /etc/puppet/modules/kvm/files/issue.DoD \ /etc/puppet/modules/kvm/files/issue.DoD.bak</pre> <p>Important: Any changes in the banner file are not preserved across upgrades.</p>	
Modifying the DoD banner file.	<p>Run the following command on the AHV host:</p> <pre>[root@AHV-host ~]# vi /etc/issue.DoD</pre> <p>Note: Repeat the above step on each AHV host in a cluster.</p>	
Setting the banner for all nodes through nCLI.	<p>Run the following command:</p> <pre>nutanix@cvm\$ ncli cluster edit-hypervisor-security-params enable-banner=true</pre>	

The following options are configured or customized to harden the AHV:

- **Enable AIDE:** Advanced Intrusion Detection Environment (AIDE) is a Linux utility that monitors a given node. After you enable AIDE in ncli, an executable file named aide gets copied to `/etc/cron.weekly/` directory. When the weekly cron job runs, the executable aide file generates aide database. You can move the aide database to a secure location in a read-only media or on other machines. After aide database is created, you can use the `aide --check` command for the system to check the integrity of the files and directories by comparing the files and directories on your system with the snapshot in the database. In case there are unexpected changes, a report gets generated, which you can review. During AOS version upgrades and changes made to existing files or files added are valid, then the aide database must be updated manually using the `aide --update` command.
- **Enable high strength password:** You can run the command as shown in the table in this section to enable high-strength password policies (`minlen=15`, `difok=8`, `maxclassrepeat=4`).

Note:

- **minlen** is the minimum required length for a password.
 - **difok** is the minimum number of characters that must be different from the old password.
 - **maxclassrepeat** is the number of consecutive characters of same class that you can use in a password.
- **Enable Core:** A core dump consists of the recorded state of the working memory of a computer program at a specific time, generally when the program gets crashed or terminated abnormally. Core dumps are used to assist in diagnosing or debugging errors in computer programs. You can enable the core for troubleshooting purposes.

- **Enable Banner:** You can set a banner to display a specific message. For example, set a banner to display a warning message that the system is available to authorized users only.

Hardening Controller VM

You can use Nutanix Command Line Interface (nCLI) in order to customize the various configuration settings related to the Controller VM as described below.

For the complete list of cluster security parameters, see *Edit the security params of a Cluster* in the [Command Reference](#) guide.

- Run the following command to support cluster-wide configuration of the SCMA policy.

```
nutanix@cvm$ ncli cluster get-cvm-security-config
```

The current cluster configuration is displayed.

```
Enable Aide           : false
Enable Core           : false
Enable High Strength P... : false
Enable Banner         : false
Enable SNMPv3 Only    : false
Schedule              : DAILY
Enable Kernel Mitigations : false
SSH Security Level     : DEFAULT
Enable Lock Status     : false
Enable Kernel Core     : false
```

Important:

SSH Security Level is applied to the nutanix user for SSH login to the Nutanix Cluster. The security levels can be configured with one of the following options:

- **default** - The nutanix user can start an SSH session with either a password or an SSH key without any change to the account privileges.
- **limited** - For SSH sessions started with a password, the nutanix user is switched to the admin user that has lower operational privileges. SSH key-based logins do not change the nutanix user's privileges.
- **restricted** - For all SSH sessions, whether started with a password or an SSH key, the nutanix user is switched to the admin user that has lower operational privileges.

Tip: The admin user, while having lower privileges than the nutanix user, is designed to have sufficient privileges for administrative tasks, see [Controller VM Access](#) in the AHV Administration Guide for details.

Tip: In addition to configuring the SSH Security Level, you can also consider cluster lockdown to disable password-based SSH authentication by adding SSH keys, see [Controlling Cluster Access](#) on page 71.

- Run the following command to schedule weekly execution of Advanced Intrusion Detection Environment (AIDE).

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-aide=true
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : false
```

```

Enable Banner           : false
Enable SNMPv3 Only     : false
Schedule                : DAILY
Enable Kernel Mitigations : false
SSH Security Level      : DEFAULT
Enable Lock Status      : false
Enable Kernel Core      : false

```

- Run the following command to enable the strong password policy.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-high-strength-password=true
```

The following output is displayed.

```

Enable Aide             : true
Enable Core             : false
Enable High Strength P... : true
Enable Banner           : false
Enable SNMPv3 Only     : false
Schedule                : DAILY
Enable Kernel Mitigations : false
SSH Security Level      : DEFAULT
Enable Lock Status      : false
Enable Kernel Core      : false

```

- Run the following command to enable the defense knowledge consent banner of the US department.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-banner=true
```

The following output is displayed.

```

Enable Aide             : true
Enable Core             : false
Enable High Strength P... : true
Enable Banner           : true
Enable SNMPv3 Only     : false
Schedule                : DAILY
Enable Kernel Mitigations : false
SSH Security Level      : DEFAULT
Enable Lock Status      : false
Enable Kernel Core      : false

```

- Run the following command to enable the settings to allow only SNMP version 3.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-snmpv3-only=true
```

The following output is displayed.

```

Enable Aide             : true
Enable Core             : false
Enable High Strength P... : true
Enable Banner           : true
Enable SNMPv3 Only     : true
Schedule                : DAILY
Enable Kernel Mitigations : false
SSH Security Level      : DEFAULT
Enable Lock Status      : false
Enable Kernel Core      : false

```

- Run the following command to change the default schedule of running the SCMA. The schedule can be hourly, daily, weekly, and monthly.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params schedule=hourly
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : true
Enable Banner         : true
Enable SNMPv3 Only    : true
Schedule              : HOURLY
Enable Kernel Mitigations : false
SSH Security Level     : DEFAULT
Enable Lock Status     : false
Enable Kernel Core     : false
```

- Run the following command to enable the settings so that Controller VM can generate stack traces for any cluster issue.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-core=true
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : true
Enable Banner         : true
Enable SNMPv3 Only    : true
Schedule              : HOURLY
Enable Kernel Mitigations : false
SSH Security Level     : DEFAULT
Enable Lock Status     : false
Enable Kernel Core     : true
```

Note: Nutanix recommends that Core should not be set to true unless instructed by the Nutanix support team.

- Run the following command to configure security levels for the nutanix user for ssh login to the Nutanix Cluster.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params ssh-security-level=limited
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : true
Enable Banner         : true
Enable SNMPv3 Only    : true
Schedule              : HOURLY
Enable Kernel Mitigations : false
SSH Security Level     : LIMITED
Enable Lock Status     : true
Enable Kernel Core     : true
```

- Run the following command to enable to locking of the security configuration.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-lock-status=true
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : true
Enable Banner         : true
Enable SNMPv3 Only    : true
Schedule              : HOURLY
Enable Kernel Mitigations : false
SSH Security Level     : LIMITED
Enable Lock Status    : true
Enable Kernel Core    : true
```

Note: Enabling the security configuration lock to true locks the configuration settings, preventing you from making changes through nCLI or API calls. You need to contact Nutanix Support to unlock this configuration.

- Run the following command to enable iTLB Multihit Mitigation (CVE-2018-12207) for all AHV nodes.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-itlb-multihit-mitigation=true
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : true
Enable High Strength P... : true
Enable Banner         : true
Schedule              : HOURLY
Enable iTLB Multihit M... : true
```

Note:

- This settings is disabled by default
- Enabling this setting may have performance impact on the running workloads

•

Scenario-Based Hardening

- When a high governance official needs to run the hardened configuration then the settings should be as follows.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : true
Enable Banner         : false
Enable SNMPv3 Only    : true
Schedule              : HOURLY
Enable Kernel Mitigations : false
SSH Security Level     : LIMITED
Enable Lock Status    : true
Enable Kernel Core    : true
```

- When a federal official needs to run the hardened configuration then the settings should be as follows.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : true
Enable Banner         : true
Enable SNMPv3 Only    : true
Schedule              : HOURLY
Enable Kernel Mitigations : false
SSH Security Level     : LIMITED
Enable Lock Status     : true
Enable Kernel Core     : true
```

DoD Banner Configuration

- **Note:** A banner file can be modified to support non-DoD customer banners.

- Run the following command to backup the DoD banner file.

```
nutanix@cvm$ sudo cp -a /srv/salt/security/CVM/sshd/DODbanner \
/srv/salt/security/CVM/sshd/DODbannerbak
```

- Run the following command to modify DoD banner file.

```
nutanix@cvm$ sudo vi /srv/salt/security/CVM/sshd/DODbanner
```

Note: Repeat all the above steps on every Controller VM in a cluster.

- Run the following command to backup the DoD banner file of the Prism Central VM.

```
nutanix@pcvm$ sudo cp -a /srv/salt/security/PC/sshd/DODbanner \
/srv/salt/security/PC/sshd/DODbannerbak
```

- Run the following command to modify DoD banner file of the Prism Central VM.

```
nutanix@pcvm$ sudo vi /srv/salt/security/PC/sshd/DODbanner
```

- Run the following command to set the banner for all nodes through nCLI.

```
nutanix@cvm$ ncli cluster edit-cvm-security-params enable-banner=true
```

TCP Wrapper Integration

TCP Wrapper is not supported in AOS version 6.8.

IP Set Based Firewall

The IP set-based firewall feature implements strict firewall rules for intra cluster communication. This feature enables individual IP address filtering for intra cluster communication for various internal services. Prism considers only individual Controller VM and hypervisor IP addresses as trusted sources to receive RPC and API request from other CVMs in the same cluster.

The IP set-based firewall feature is enabled by default in AOS version 6.8.

Hardening PCVM

You can use Nutanix Command Line Interface (nCLI) in order to customize the various configuration settings related to the PCVM as described below.

For the complete list of cluster security parameters, see *Edit the security params of a Cluster* in the [Command Reference](#) guide.

- Run the following command to support cluster-wide configuration of the SCMA policy.

```
nutanix@pcvm$ ncli cluster get-pcvm-security-config
```

The current cluster configuration is displayed.

```
Enable Aide           : false
Enable Core           : false
Enable High Strength P... : false
Enable Banner         : false
Enable SNMPv3 Only    : false
Schedule              : DAILY
Enable Kernel Core    : true
Enable Page Poison    : false
Enable Slub Debug     : false
SSH Security Level    : DEFAULT
Enable Lock Status    : false
IP Restriction State  : NORMAL
Enable Dodin Mode     : false
Enable Fapolicy       : false
Enable Processor Mitig... : false
SSH whitelisted addres... :
```

- Run the following command to schedule weekly execution of Advanced Intrusion Detection Environment (AIDE).

```
nutanix@pcvm$ ncli cluster edit-pcvm-security-params enable-aide=true
```

The following output is displayed.

```
Enable Aide           : true
Enable Core           : false
Enable High Strength P... : false
Enable Banner         : false
Enable SNMPv3 Only    : false
Schedule              : DAILY
Enable Kernel Core    : true
Enable Page Poison    : false
Enable Slub Debug     : false
SSH Security Level    : DEFAULT
Enable Lock Status    : false
IP Restriction State  : NORMAL
Enable Dodin Mode     : false
Enable Fapolicy       : false
Enable Processor Mitig... : false
SSH whitelisted addres... :
```

Common Criteria

Common Criteria is an international security certification that is recognized by many countries around the world. Nutanix AOS and AHV are Common Criteria certified by default and no additional configuration is required to enable the Common Criteria mode. For more information, see the [Nutanix Trust](#) website.

Note: Nutanix uses FIPS-validated cryptography by default.

Certificate Revocation Checking (nCLI)

Enabling Certificate Revocation Checking using Online Certificate Status Protocol (nCLI)

About this task

OCSP is the recommended method for checking certificate revocation in client authentication. You can enable certificate revocation checking using the OSCP method through the command line interface (nCLI).

To enable certificate revocation checking using OCSP for client authentication, do the following.

Procedure

1. Set the OCSP responder URL.

```
ncli authconfig set-certificate-revocation set-ocsp-responder=<ocsp url><ocsp url>
```

indicates the location of the OCSP responder.

2. Verify if OCSP checking is enabled.

```
ncli authconfig get-client-authentication-config
```

The expected output if certificate revocation checking is enabled successfully is as follows.

```
Auth Config Status: true
File Name: ca.cert.pem
OCSP Responder URI: http://<ocsp-responder-url>
```

Enabling Certificate Revocation Checking using Certificate Revocation Lists (nCLI)

About this task

Note: OSCP is the recommended method for checking certificate revocation in client authentication.

You can use the CRL certificate revocation checking method if required, as described in this section.

To enable certificate revocation checking using CRL for client authentication, do the following.

Procedure

Specify all the CRLs that are required for certificate validation.

```
ncli authconfig set-certificate-revocation set-crl-uri=<uri 1>,<uri 2> set-crl-refresh-interval=<refresh interval in seconds> set-crl-expiration-interval=<expiration interval in seconds>
```

- The above command resets any previous OCSP or CRL configurations.
- The URIs must be percent-encoded and comma separated.
- The CRLs are updated periodically as specified by the `crl-refresh-interval` value. This interval is common for the entire list of CRL distribution points. The default value for this is 86400 seconds (1 day).
- The periodically updated CRLs are cached in-memory for the duration specified by value of `set-crl-expiration-interval` and expired after the duration, in case a particular CRL distribution point is not reachable. This duration is configured for the entire list of CRL distribution points. The default value for this is 604800 seconds (7 days).

Eliminate Default Passwords during Cluster Creation

Default passwords are a well-known security vulnerability, often exploited in hacking attempts. By removing the usage of default passwords and providing mechanisms for more secure SSH access configurations, you can improve the overall security posture of your Nutanix clusters.

You can eliminate default passwords by configuring the following security settings during the cluster creation operations using the Controller VM CLI. See [Eliminating Default Passwords during Cluster Creation \(CVM-only\)](#) on page 194 for details.

- **SSH Access Configuration**

`Password Lockdown Mode` can be applied to the cluster to restrict password-based SSH access by the default local user account ("nutanix" and "admin").

`Lockdown Mode` can be applied to the cluster to restrict both password-based and public SSH key based authentication.

- **Adding an external Public SSH Keys**

You can add public SSH keys from trusted systems, further securing remote access without relying on password-based authentication.

You can apply the above configuration to eliminate the default password during the cluster creation workflow. The new password and/or publish SSH keys are maintained during cluster management workflow, such as adding a node to a cluster, removing a node from the cluster, or performing a break-fix operation.

For more information on cluster management operations using nCLI, see the [Acropolis Advanced Administration Guide](#).

Note: Elimination of default password is supported only for "nutanix" user and Controller VM.

Eliminating Default Passwords during Cluster Creation (CVM-only)

About this task

You can eliminate default password for SSH access to the Controller VM by the "nutanix" user during cluster creation. To change the default password and SSH access configuration for the Controller VM, do the following.

Procedure

- Method 1: Change the default password and configure SSH access based on the password
This method enables SSH access to Controller VMs using the passwords.
 - a. Log on to any Controller VM in the cluster with SSH.
 - b. Create a cluster and set the SSH access mode.

```
nutanix@cvm$ cluster --cluster_create_password_enforcement=true --  
cluster_create_security_opts=true --encrypted_password="<encrypted_string>" create
```

Note: Replace <encrypted_string> with the encrypted password to be applied on the cluster. Also, see [Controller VM Password Complexity Requirements](#).

- Method 2: Change the default password and configure SSH access based on a public key
This method enables SSH access to the Controller VMs using the public keys you provide.
 - a. Log on to any Controller VM in the cluster with SSH.
 - b. Create a cluster and set the SSH access mode (public SSH key based access only).

```
nutanix@cvm$ cluster --password_lockdown_mode=true --
cluster_create_password_enforcement=true --cluster_create_security_opts=true --
encrypted_password="<encrypted_string>" --external_access_keys="<public_key>"
create
```

Note:

- Replace <encrypted_string> with the encrypted password to be applied on the cluster. Also, see [Controller VM Password Complexity Requirements](#).
- Replace <public_key> with the actual public SSH key of the trusted system from which you want to allow SSH access to the system.
- You can enter multiple public SSH keys for multiple trusted systems in the following way:

```
nutanix@cvm$ cluster --password_lockdown_mode=true
--cluster_create_password_enforcement=true
--cluster_create_security_opts=true --
external_access_keys="<public_key_1>", "<public_key_2>" create
```

The cluster creation starts, and the following settings are applied:

Note:

- The default password for the Controller VM is replaced by the password that you provide.
- SSH access to the Controller VM using password is restricted.
- The new password and/or public SSH key is maintained during cluster management workflows like, adding a node to a cluster, removing a node from the cluster, or breakfix operation.

- Method 3: Disable SSH access completely.
This method disables SSH access to the Controller VM, both password-based and public SSH key based authentication are restricted. This enables the highest level of security for the Controller VM, along with eliminating default password usage.
 - a. Log on to any Controller VM in the cluster with SSH.
 - b. Create a cluster with lockdown mode enabled.

```
nutanix@cvm$ cluster --lockdown_mode=true --
cluster_create_password_enforcement=true --cluster_create_security_opts=true
create
```

The cluster creation starts, and password-based and SSH key based authentication is restricted.

ACCESSING A LIST OF OPEN SOURCE SOFTWARE RUNNING ON A CLUSTER

As an admin user, you can access a text file that lists all of the open source software running on a cluster.

About this task

Perform the following procedure to access a list of the open source software running on a cluster.

Procedure

1. Log on to any Controller VM in the cluster as the admin user by using SSH.
2. Access the text file by using the following command.

```
less /usr/local/nutanix/license/blackduck_version_license.txt
```

COPYRIGHT

Copyright 2024 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

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