

Flow Network Security Next-Gen Release Version 4.2.x Guide

Flow Network Security Next-Gen 4.2.x
March 20, 2025

Contents

Flow Network Security Product Generation and Release Version.....	4
Flow Network Security Release Version 4.2.x Overview.....	5
FNS Next-Gen Deployments.....	5
Concepts and Terminologies.....	6
Installation and Upgrades.....	8
Enabling Microsegmentation.....	9
Microsegmentation Requirements for FNS Next-Gen 4.2.0.....	9
Limitations.....	10
Migration Experience.....	11
Guided Tour.....	11
Simulated Mode.....	12
Configuring Simulated Mode.....	12
Migrate Flow Network Security to Flow Network Security Next-Gen.....	13
Migration Requirements.....	15
Migration Limitation.....	15
Migrating FNS Policies to FNS Next-Gen.....	16
Security Policies.....	19
Security Policy Model.....	20
Types of Policies.....	23
Services.....	25
Creating a Service.....	25
Addresses.....	26
Creating an Address.....	26
Import Address Groups.....	26
Built-In Categories for Security Policies.....	27
Flow Network Security Next-Gen and Disaster Recovery.....	29
Flow Network Security Support for Single-PC Disaster Recovery.....	29
Policy Consumption and Visualization.....	30
Policy Views.....	30
Security Policies View.....	30
Individual Policy View.....	32
Applying Filtering and Grouping.....	33
Allowing Discovered Traffic.....	35

Role-Based Access Control.....	36
Flow Network Security Roles and Permissions.....	36
Creating an Authorization Policy for FNS Next-Gen.....	38
Application Policy Configuration.....	40
Creating an Application Policy with a VLAN Scope.....	40
Creating an Application Policy within a VPC Scope.....	44
Modifying an Application Policy.....	48
Applying an Application Policy.....	48
Cloning a Security Policy.....	49
Deleting an Application Policy.....	49
VDI Policy Configuration.....	50
Creating a VDI Policy.....	50
Default VDI Policy.....	55
Configuring Active Directory Domain Services.....	55
Modifying the VDI Policy.....	58
Applying the VDI Policy.....	58
Monitoring the VDI Policy.....	58
Deleting the VDI Policy.....	59
Isolation Environment Policy.....	60
Creating an Isolation Environment Policy.....	61
Modifying an Isolation Environment Policy.....	64
Applying an Isolation Environment Policy.....	64
Monitoring an Isolation Environment Policy (Visualizing Network Flows).....	65
Deleting an Isolation Environment Policy.....	67
Quarantine Policy Configuration.....	68
Configuring the Quarantine Policy.....	68
Quarantining a VM.....	70
Removing a VM from the Quarantine.....	71
Exporting and Importing Security Policies.....	72
Disabling Microsegmentation.....	73
Copyright.....	74

FLOW NETWORK SECURITY PRODUCT GENERATION AND RELEASE VERSION

The FNS 4.2.0 is the first release in the FNS 4.2.x series and this release is bundled with the AOS 7.0 and Prism Central pc.2024.3 release bundle.

This table lists different FNS product generations:

Product Generation	Description
Flow Network Security (formerly Flow Microsegmentation)	Legacy microsegmentation solution
Flow Network Security Next-Gen	New generation microsegmentation solution

This table lists FNS release versions:

Release Version	Description
Flow Network Security Release Version 4.2.0	FNS NG release with dual stack that has feature payload for VLAN only and VPC only environments.
Flow Network Security Release Version 4.1.0	FNS NG release with dual stack that has feature payload for VLAN only and VPC only environments.
Flow Network Security Release Version 3.1.1	FNS release with a feature payload only for Basic VLAN environments. This release is a patch release for FNS 3.1.0
Flow Network Security Release Version 3.0.0	FNS NG release with a feature payload for VPC environments.
Flow Network Security Release Version 1.0.x	FNS release bundled with AOS

FLOW NETWORK SECURITY RELEASE VERSION 4.2.X OVERVIEW

The FNS release version 4.2.0 is the latest release for the Flow Network Security product generation that supports Network Controller-managed VLAN and VPC environments.

FNS release version 4.2.0 supports Legacy FNS and FNS Next-Gen (Network Controller managed VLAN and VPC) modes.

For more information on Flow Network Security Next-Gen supported features and configurations, see [Flow Network Security Next-Gen Release Version 4.2.x Guide](#).

For information on Flow Network Security (Legacy) supported features and configurations, see [Flow Network Security Release Version 4.2.x Guide](#)

FNS Next-Gen Deployments

This section explains various FNS Next-Gen deployment scenarios based on your current FNS environment setup and also covers information on new deployments. The deployment scenarios are classified as follows:

- VLAN only Environments
- VPC only Environments
- Mixed Environments

Furthermore, each scenario covers information on deploying FNS Next-Gen on on-prem and NC2.

Note: FNS Next-Gen is supported on clusters that have compute-only nodes.

FNS Next-Gen has the following pre-requisites:

- Network controller
- Network controller-enabled VLANs and VPC

For more information on network controllers, see [Flow Virtual Networking Overview](#).

VLAN Only Environments

- **On-Prem**
 - If you are enabling FNS for the first time after deploying FNS 4.2.0, you default to FNS Next-Gen.
 - If you are running Legacy FNS and deploy FNS 4.2.0, then you continue to use Legacy FNS. A migration banner to migrate to FNS Next-Gen is displayed in the Security Policies page.
- **NC2**
 - **AWS:**
 - FNS Next-Gen in FNS VLAN environment is not supported on NC2 on AWS.
 - If an Legacy FNS customer using NC2 on AWS upgrades to FNS 4.2.0, the customer continues to use Legacy FNS.
 - **Azure:** FNS Next Gen in FNS VLAN environment is not supported on NC2 on Azure.

VPC Only Environments

- **On-Prem:**
 - **Feature coverage:** All the features of FNS release version 4.1.0 are supported in FNS release version 4.2.0
 - **Customer coverage:**
 - If you have FNS Next-Gen running release version 4.1.0 on a VPC environment and you upgrade to FNS 4.2.0, the current feature set continues to work as-is.
 - If an existing VPC customer deploys FNS 4.2.0 and then enables FNS, FNS Next-Gen is loaded by default.
 - If an existing FNS customer using Legacy FNS deploys FNS 4.2.0, you must migrate to FNS Next-Gen to be able to use VPC with FNS in the Next-Gen.
- **NC2**
 - **AWS:** FNS Next Gen in FVN environment is supported on NC2 on AWS.
 - **Azure:** FNS Next-Gen in FVN environment is supported on NC2 on Azure.

Mixed Environments

A mixed environment is when a customer has Network controller-managed VLANs and VPC infrastructure within the same Prism Central. In a mixed environment, FNS doesn't support creating or managing FNS overlapping/extended policies across VPC and VLAN environments.

Concepts and Terminologies

This section describes key concepts and terminology you need to understand to use Flow Network Security Next-Gen effectively.

Microsegmentation

The process of breaking down a network into smaller segments to make it more difficult for an attacker to access a whole system. Each segment acts as its own barrier: If an attacker broke into a system, the intruder would only be able to get to a single segment first, rather than the entire system.

Security Policy Model

A schema of policies for specifying and enforcing a desired behaviour. A Policy Model will have one or more policies.

Security Policy

Defines how to protect assets from threats and how to handle situations when they do occur. Security policy is a collection of security rules and assets [entities, endpoints, categories, applications etc] on which the rules have to be enforced together.

Category

Nutanix construct for the well known concept of Tags, are used to define groups of entities which policies and enforcement are applied to. They typically apply, but are not limited to: environment, application type, application tier, etc. Category: Key/Value "Tag". Examples: app | tier | group | location | subnet | etc. These categories are leveraged by policies to determine what rules / actions to apply.

Category Set

A collection of Categories which are evaluated with an AND operation to resolve to a set of VMs where all categories in the Category Set are assigned.

Entity

Nutanix entity is one or more instances of an object type such as a VM, cluster, security policy, project, or report. For the scope of Flow, we shall refer 'Entity' for end-points of traffic which can be a source or target of a protected entity:

- Source Entity: The entity from where the inbound traffic to a Secured Entity is to be controlled by a Flow policy
- Secured Entity: The entity which is being protected by the Flow policy
- Destination Entity: The entity to which the outbound traffic from the Secured Entity needs to be controlled.

Zero Trust

A security concept centered on the belief that organizations should not automatically trust anything inside or outside its perimeters and instead must verify anything and everything trying to connect to its systems before granting access.

Blocklist

Automatically approves everything. The user has to explicitly define what should be rejected.

Allowlist [least privilege model]

Automatically denies everything. The user has to explicitly define what is allowed.

INSTALLATION AND UPGRADES

Before installing or upgrading to FNS release version 4.2.0, ensure that you have AOS 7.0 or later and pc.2024.3 or later versions.

The FNS Next-Gen release version 4.2.0 is bundled with the AOS 7.0 and Prism Central pc.2024.3 release bundle. When you deploy or upgrade to AOS 7.0 release, FNS Next-Gen 4.2.0 bundled release is deployed on PC and PE respectively.

Connected Site

For Prism Central and Prism Element, installing or upgrading FNS 4.2.0 using LCM 1-Click method on a connected site is not available. You can install using the dark site method as described below.

Dark Site

You install or upgrade to FNS Next-Gen 4.2.0 when you install or upgrade to AOS 7.0 and Prism Central pc.2024.3 using LCM at a dark site. For more information, see [LCM Settings for Dark Sites - No Internet Connectivity](#).

Note: If you are running FNS 5.0.0 and upgrade AOS to AOS 7.0 and PC to pc.2024.3 version, then you continue to run on FNS 5.0.0

ENABLING MICROSEGMENTATION

Microsegmentation is disabled by default. Before you can configure and use application security policies, isolation environment policies, and quarantine policies, you must enable the feature.

Before you begin

- The feature requires a Flow license. For information on Flow license, see [License Manager Guide](#). If you have not installed a Flow license, you can try the feature for a period of 60 days. After this period expires, you will be required to install the license to continue using the feature.
- Ensure that you meet [Microsegmentation Requirements for FNS Next-Gen 4.2.0](#) on page 9.

About this task

To enable microsegmentation, do the following:

Procedure

1. Log on to the Prism Central web console.
2. Click the **Prism Central Settings** to display the Settings page.
3. Click **Microsegmentation** from the Settings menu (on the left). The **Enable Microsegmentation** dialog box is displayed.
4. Select the **Enable Microsegmentation** check box.
5. Click **Save**.

Microsegmentation Requirements for FNS Next-Gen 4.2.0

The FNS NG feature has the following requirements:

- Supported only on the AHV clusters running:
 - AOS 7.0 or later version
 - AHV version 10 or later version.
- The Prism Central (pc.2024.3 or later) instance is hosted on one of the AHV clusters registered with it.
- The host must have at least 2 GB of additional memory for each Prism Central VM hosted on it.
- You must enable Advanced Networking if you are enabling Microsegmentation for the first time.

For more information on the memory requirements for Advanced Networking, see [Flow Virtual Networking guide](#).

- If you are running a Prism Central scale-out instance, all the VMs in the Prism Central cluster must be powered on.
- Allow AHV hosts to communicate with the Prism Central VMs over TCP port 9446.

Keeping the port open enables the hosts to send the Prism Central VMs connection tracking data. Prism Central uses that data to show network flows.

Caution:

- When Flow Microsegmentation is enabled, a Kafka container called **flow_data** is automatically created on the cluster that hosts Prism Central. The **flow_data** container stores data that is required and essential for Flow visualization. Do not delete this container.

Limitations

FNS Next-Gen (VPC and VLAN environments) is not supported on clusters that have storage-only nodes.

The following features are not supported for FNS Next-Gen release version 4.2.0:

For VPC Environment

- Multi-PC Disaster Recovery (DR).
- Single PC and two PC Cross Cluster Live Migration (CCLM) with VPCs.
- Overlapping or conflicting policy configuration. These policy configurations cause unintended interruption to network services.
- VDI policy for VPC scope
- Flow Security Central (FSC) is not supported.
- Network Function Chain. This feature allows network traffic to and from VMs running on AHV to be transparently redirected through a service VM such as firewall, intrusion detection, etc.
- Live Migration of Connection Tracking (LMCT).
- Remote Office Branch Office (ROBO)

For VLAN Environment

- Multi-PC Disaster Recovery (DR).
- Single PC and two PC Cross Cluster Live Migration (CCLM) with Network Controller Managed VLANs.
- Overlapping or conflicting policy configuration. These policy configurations cause unintended interruption to network services.
- Network Function Chain. This feature allows network traffic to and from VMs running on AHV to be transparently redirected through a service VM such as firewall, intrusion detection, etc.
- Live Migration of Connection Tracking (LMCT).
- Remote Office Branch Office (ROBO)

MIGRATION EXPERIENCE

Flow Network Security Next-Gen provides an intuitive curated framework to migrate FNS policies to FNS Next-Gen. The Guided Tour and Simulated Mode features helps you understand the new capabilities of FNS NG and also gives you the insight on how the policies work post migration.

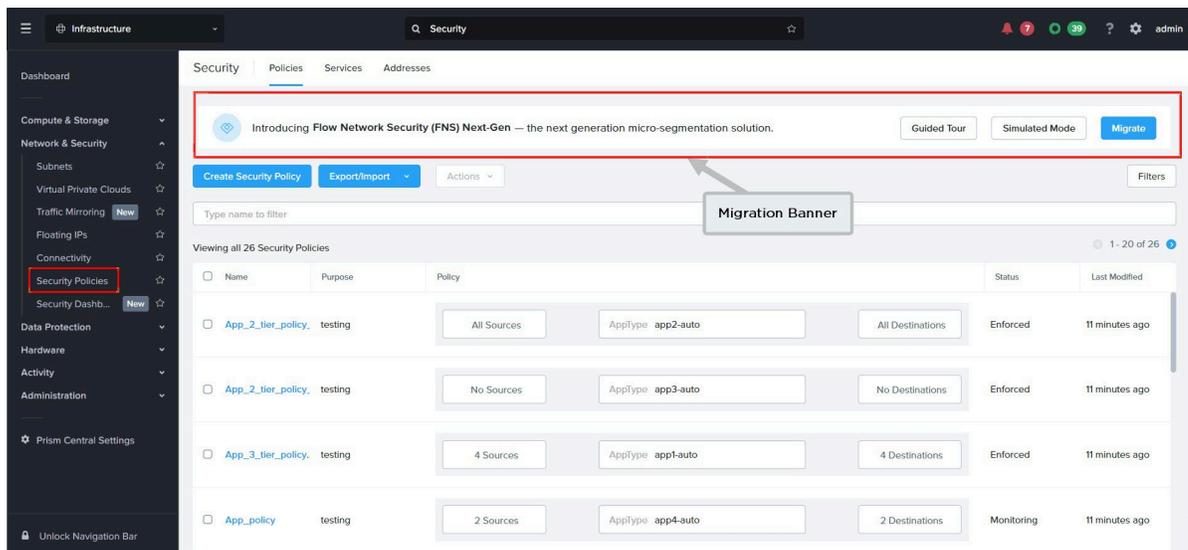


Figure 1: Migration Banner In the Security Policies page

The migration banner appears only for Legacy FNS release version 4.2.0 deployment.

Note:

If you have enabled Remote Office Branch Office (ROBO) feature and want to continue using ROBO, then Nutanix recommends that you do not migrate to FNS Next-Gen.

Guided Tour

About this task

The Guided Tour feature is a walk through that briefly explains the new features and capabilities that FNS NG has to offer.

Procedure

1. In the **Security Policies** page, click **Guided Tour**.

The **Guided Tour** opens in another tab of the browser. The Guided Tour shows information on various FNS Next-Gen features step-by-step.

- In the first step the Guided Tour displays the information on the **Policy Details**. Go through the details and click **Next**.

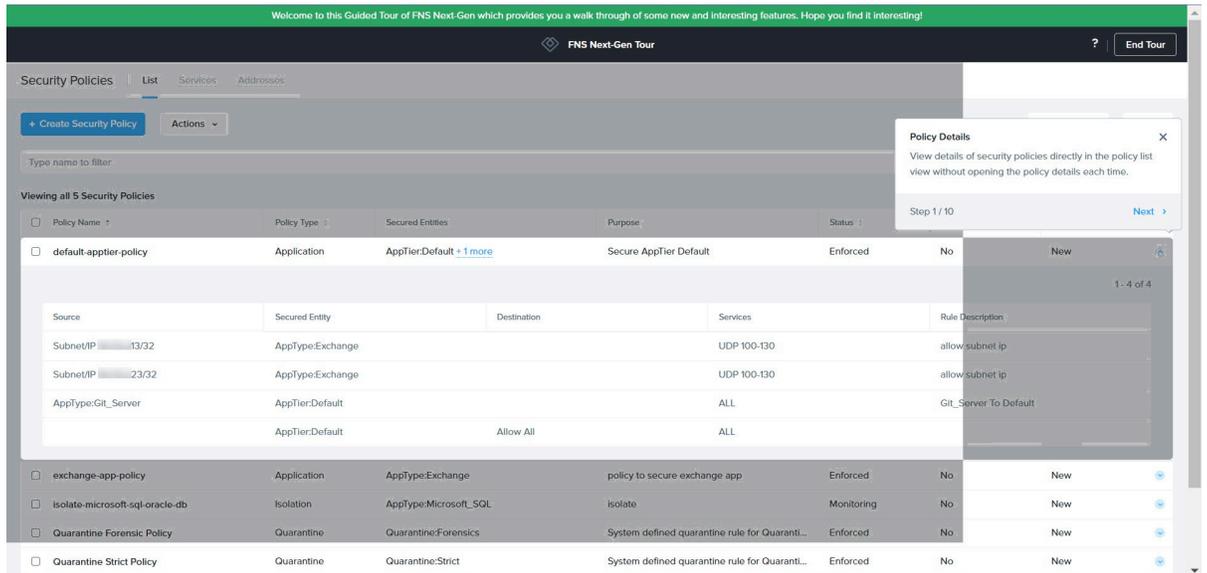


Figure 2: Guided Tour Step 1: Policy Details

Continue reading about all the new features step-by-step.

- In the last step, click **End Tour**.

You can restart the tour by clicking **Back to Tour**, or launch the tour again from the migration banner in the **Security Policies** page.

Simulated Mode

Simulated Mode allows you to experience FNS Next-Gen in a simulated environment.

To experience FNS NG in the simulated mode, import your current policies to a simulated environment; preview and test your policies in FNS NG. Any changes performed during the simulation are not saved. Therefore, your current security policies remain unaffected. You can import up to 20 policies from your current environment and perform the following tasks in the simulated mode:

- Create a security policy.
- View the policy details in visualization and table view.

The simulated mode has limited features and does not support:

- Creating a security policy with VPC scope.
- Updating a security policy.
- Addresses and Services features.

Configuring Simulated Mode

About this task

Simulation mode allows you to peek into the new security features and performances. Any changes performed during the simulation are not saved. Therefore, your current security policies remain unaffected.

Procedure

1. In the **Security Policies** page, click **Simulated Mode**.
2. Select the policies to simulate and click **Start Simulation**.

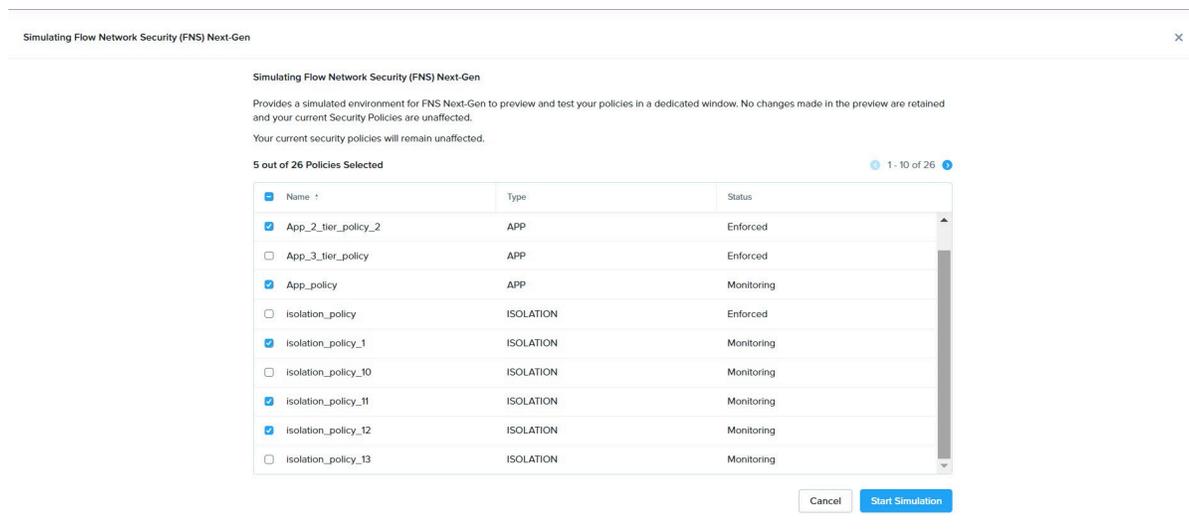


Figure 3: Select Policies

You can select up to 20 policies for simulation.

3. In the **Simulating FNS Next-Gen** screen, click **Get Started**.
4. In the Simulated mode, you can create a policy, and view existing policies.
For information on how to create a policy, see [Creating an Application Security Policy](#) and [Creating an Isolation Environment Policy](#).
5. Click **Exit Simulated Mode** to go back to the Security Policies page.

Migrate Flow Network Security to Flow Network Security Next-Gen

Flow Network Security provides an intuitive curated framework to migrate FNS policies and subnets to FNS Next-Gen. You can migrate your FNS set-up to the Flow Network Security Next-Gen without disrupting your security posture. Although there is no disruption to the security posture during migration, we recommend performing the migration during a maintenance window.

The migration experience is a curated experience. Nutanix plans to enable migration for Legacy FNS users in multiple phases. For more information on the different phases, see the KB article [15412](#).

It is important to note the total number of security policies, App-Tiers, subnets, and vNICs configured in Prism Central before migration.

- **Total Number of Security Policies:** You can view total number of security policies in the Security Policies landing page.

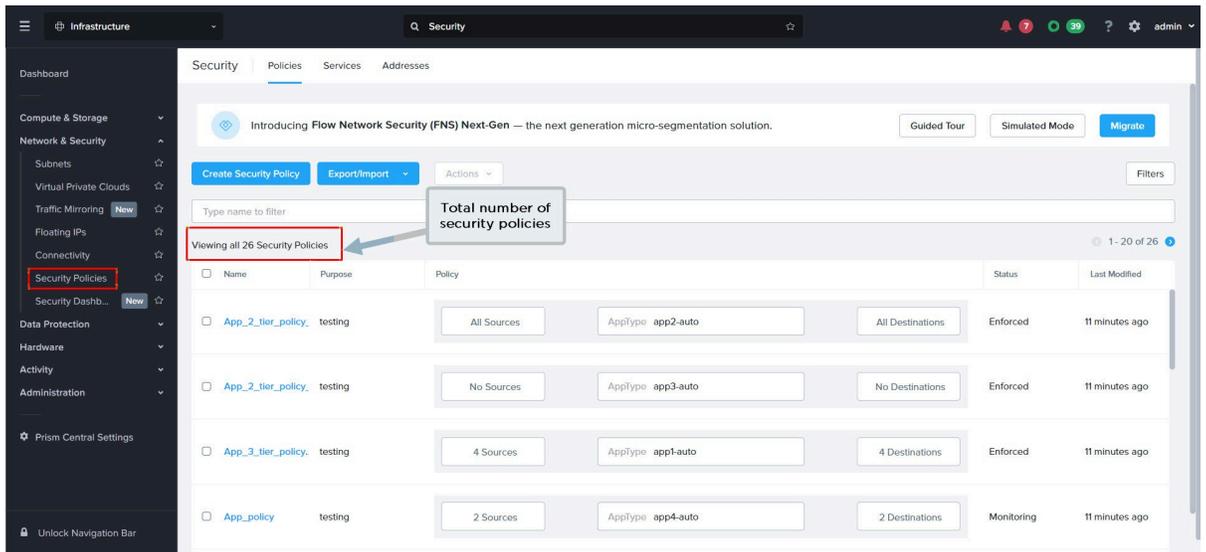


Figure 4: Security Policies page

- **Total Number of App-Tiers:** You can view number of App-Tiers in the **Secured Entities** tile of a security policy.



Figure 5: Secured Entity Tile

Nutanix recommends you to go through the [Guided Tour](#) and [Simulated Mode](#) topics before experiencing the migration.

The following tasks cannot be performed during migration:

- Updates to security policies
- Updates to the selected VLAN basic subnets or associated VM NICs

The following tasks are involved in migration:

- Migration pre-checks
- Security policy migration
- Subnet migration

After the migration is complete, you cannot revert the migration. If migration fails before starting subnet migration, then policies continue to work as earlier. You can retry the migration after sometime. However, if migration fails during or after subnet migration, contact support for assistance.

The following table shows the supported enforced scale during migration.

Table 1: FNS Next-Gen Enforced Scale

Feature	Scale
Application policies	1000
Isolation policies	100
Rules	80 per security policy
Tier	16 per security policy
Services	160 per security policy

Note: The system does not allow you to migrate if you have a feature scale beyond the specified numbers in the above table. Contact support for assistance.

Migration Requirements

The following are the requirements for migrating FNS policies to FNS NG.

- Ensure the cluster is running the following minimum versions:

- AOS 6.7
- pc.2023.3
- AHV 20230302.207
- ANC 3.0.0

- Ensure that Network Controller is enabled on the Prism Central.

For more information about upgrading Network Controller, see [Upgrading Flow Virtual Networking](#) topic in the *Flow Virtual Networking guide*.

- Ensure you have migrated VLAN basic subnets to network controller managed VLAN subnets.

For more information about VLAN basic subnet migration pre-requisites, see [Migration of VLAN Basic Subnets](#) topic in the *Flow Virtual Networking guide*.

Migration Limitation

The system does not allow you to migrate the policies to FNS NG if you have any of the following scenarios:

- Any policy with a network function chain (NFC) attached does not get migrated to FNS NG, as FNS NG does not support NFC.
- If any of the VLAN subnets does not have DVS reference.
- At least one VM has multiple vNICs; some are on subnets being migrated, and the remaining are not.
- If any of the VLAN basic subnets have trunked vNIC.

Migrating FNS Policies to FNS Next-Gen

Before you begin

- You may need a passcode to start the migration. For more information, see the KB article [15412](#)
- Ensure you have noted the number of security policies and App-Tiers configured in the Prism Central. For more information, see [Total Number of Security Policies](#) section.

About this task

This procedure explains how to migrate FNS policies to FNS Next-Gen.

Procedure

1. In the **Security Policies** page, click **Migrate**.
A screen to enter a passcode appears.
2. Enter the passcode and click **Verify and Proceed**.
3. Read the instructions and click **Get Started** to begin migration.

The migration process takes place in two steps.

- Review Security Policies
- Review Subnets

4. In the **Review Security Policies** step, do the following.

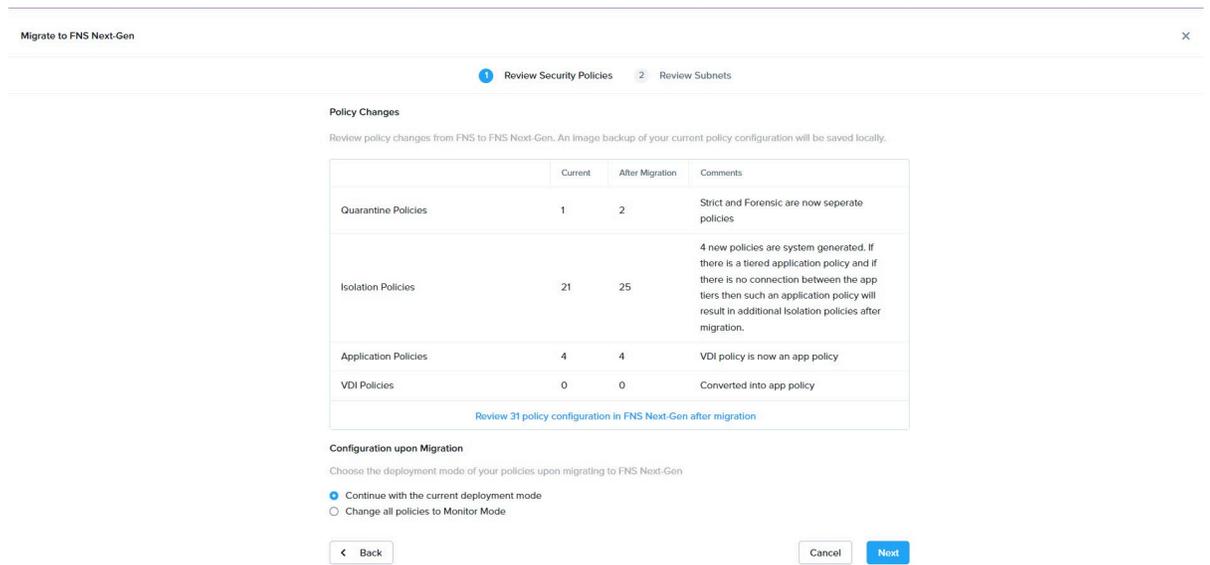


Figure 6: Review Security Policies

- a. The screen shows the number of existing policies and the number of policies after migration with the reason for adding new policies. Click **Review policy configuration in FNS Next-Gen after migration** and review the policy changes.

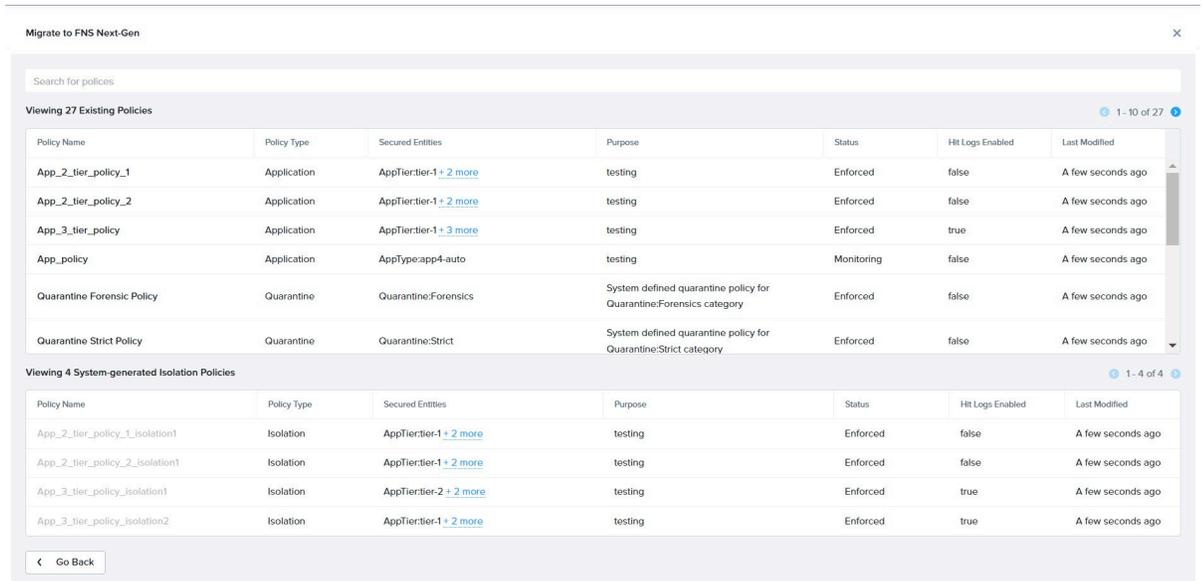


Figure 7: Review Policy Configuration Changes

Policy changes refer to converting:

- few application policies to isolation policies while migration.

For example, if there is a tiered application policy and no connection between the app tiers, then such an application policy results in an additional isolation policy after migration.

- a VDI policy to an application policy.
- a single quarantine policy into two sets of quarantine policies; one set for VLAN (forensic and strict policy) and the other for VPC (forensic and strict policy).

b. Choose the deployment mode of your policies upon migrating to FNS Next-Gen and then click **Next**.

- » **Continue with current deployment mode:** retains the current deployment modes as-is while migrating to FNS Next-Gen
- » **Change all policies to monitor mode:** the system changes modes of all policies to monitor mode after migration to FNS Next-Gen

5. In the **Review Subnets** step, choose from.

- » Migrate all VLAN Basic Subnets to Network Controller managed VLAN subnets
- » Migrate only the VLAN Basic Subnets that are part of security policies

During subnet migration, the VLAN Basic Subnets are migrated to Network Controller (ANC 3.0.0 or later) managed VLAN Subnets. The system creates new Network Controller managed VLAN Subnets with the same configuration. If the migration fails partially, the system reuses already created subnets for migration.

6. Click **Migrate**.

The status of migration is shown in the Security Policies landing page.

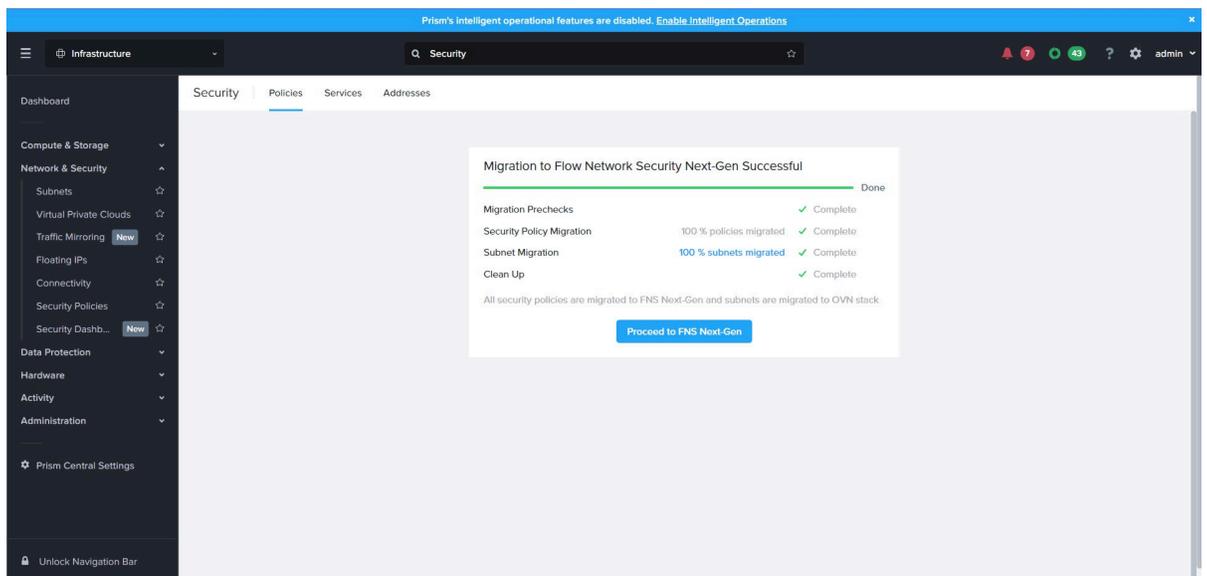


Figure 8: Migration Status

7. Click **Proceed to FNS Next-Gen**.

SECURITY POLICIES

Traditional data centers use firewalls to implement security checks at the perimeter—the points at which traffic enters and leaves the data center network. Such perimeter firewalls are effective at protecting the network from external threats. However, they offer no protection against threats that originate from within the data center and spread laterally, from one compromised machine to another.

The problem is compounded by virtualized workloads changing their network configurations and hosts as they start, stop, and migrate frequently. For example, IP addresses and MAC addresses can change as applications are shut down on one host and started on another. Manual enforcement of security policies through traditional firewalls, which rely on network configurations to inspect traffic, cannot keep up with these frequent changes and are error-prone.

Network-centric security policies also require the involvement of network security teams that have intimate knowledge of network configuration in terms of VLANs, subnets, and other network entities.

Nutanix Flow includes a policy-driven security framework that inspects traffic within the data center. The framework works as follows:

- Security policies inspect traffic that originates and terminates within a data center and help eliminate the need for additional firewalls within the data center.
- The framework uses a workload-centric approach instead of a network-centric approach. Therefore, it can scrutinize traffic to and from VMs no matter how their network configurations change and where they reside in the data center. The workload-centric, network-agnostic approach also enables the virtualization team to implement these security policies without having to rely on network security teams.
- Security policies are applied to categories (a logical grouping of VMs) and not to the VMs themselves. Therefore, it does not matter how many VMs are started up in a given category. Traffic associated with the VMs in a category is secured without administrative intervention, at any scale.
- Prism Central offers a visualization-based approach to configuring policies and monitoring the traffic to which a given policy applies.
- Using Prism Central, you can configure syslog monitoring by forwarding Flow logs to an external syslog server. For more information, see *Configuring Syslog Monitoring* in the [Prism Central Admin Center Guide](#).

Note: Nutanix Flow supports only AHV hypervisor; security policies can not be applied to VMs running on other hypervisors.

SECURITY POLICY MODEL

Application-centricity

The security policy model uses an application-centric policy language instead of the more complex, traditional network-centric policy language. Configuring an application security policy involves specifying which VMs belong to the application you want to protect and then identifying the entities or networks, in the inbound and outbound directions, with which you want to allow communication.

All the entities in an application security policy are identified by the categories to which they belong and not by their IP address, VLAN, or other network attributes. After a VM is associated with a category and the category is specified in a security policy, traffic associated with the VM is monitored even if it migrates to another network or changes its IP address.

The default options for allowing traffic on the inbound and outbound directions are also inherently application centric. For application security policies, the default option for inbound traffic is **Allowed List**, which means that **Allowed List** is usually the recommended option for inbound traffic. The default option can be changed to **Allow All** traffic. The default option in the outbound direction allows the application to send traffic to all destinations, but you can configure a destination **Allowed List** if desired.

For forensic quarantine policies, the default option in both directions is **Allowed List**, but you can **Allow All** traffic in both directions. For strict quarantine policies, no traffic is allowed in either direction.

All the VMs within a category can communicate with each other. For example, in a tiered application, regardless of how you configure tier-to-tier rules, the VMs within a given tier can communicate with each other.

Allowed List-Based Policy Expression

An application security policy is expressed in terms of the categories and subnets with which you want the application to communicate and therefore, by extension, the traffic you want to allow. A more granular policy expression can be achieved by specifying which protocols and ports can be used for communication.

Any category or subnet that is not in the allowed list is blocked. You cannot specify the categories and subnets you want to block because the number of such entities are typically much larger and grow at a much higher rate than the categories and subnets with which an application should be allowed to communicate. Expressing a policy in terms of allowed traffic results in a smaller, tighter policy configuration that can be modified, monitored, and controlled more easily.

Policy Scope

Prism Central associates each policy with either VLAN Subnets or VPC policy scope.

VLAN Subnets scope refers to securing VMs that are part of Network Controller managed VLANs.

VPC scope refers to securing VMs that are part of a VPC.

Policy Modes

All policies, whether associated with securing an application, or isolating environments, can be run in the following modes:

Save Mode

You can keep the policy in a draft state using the Save mode if you do not want to apply the policy at the time of creating it.

Apply (Monitor) Mode

Allows all traffic, including traffic that is not allowed by the policy. This mode enables you to visualize both allowed and disallowed traffic and fine-tune the policy before applying it.

Apply (Enforce) Mode

Blocks all traffic that is not allowed by the policy.

Automated Enforcement

A policy uses categories to identify the VMs to which it must apply. This model allows the automatic enforcement of a policy to VMs regardless of their number and network attributes. Connectivity between Prism Central and a registered AHV cluster is required only when creating and modifying policies, or when changing the mode of operation (applied or monitoring) of a policy. Policies are applied to the VMs in a cluster even if the cluster temporarily loses network connectivity with the Prism Central instance with which it is registered. New policies and changes are applied to the cluster when connectivity is restored.

Priorities Between Policies

Prism Central does not provide a way for you to specify priorities between policies of a single type. For example, you cannot prioritize one security policy over another. There is no limit to the number of inbound and outbound rules that you can add to a security policy, allowing you to define all of an application's security requirements in a single policy. This makes priorities between policies unnecessary.

However, priorities exist between the different policy types and policy modes. In the old priority model quarantine policies have the highest priority followed by isolation policies, and application security policies, in that order.

From FNS Next-Gen release version 4.2.0, the priority model is enhanced. The new priority model gives precedence to the policy in enforce mode over monitor mode. The new priority between policies is:

1. Quarantine policy in enforce mode
2. Isolation policy in enforce mode
3. Application policy in enforce mode
4. Quarantine, Isolation, and Application policy in monitor mode

The following flow chart illustrates the new priority framework. It shows that for incoming traffic or outgoing traffic, based on the rules defined in the policies that match the traffic, Quarantine policies in the enforce mode take precedence over Isolation policies in the enforce mode. Similarly, Isolation policies in the enforce mode take precedence over Application security policies in the enforce mode.

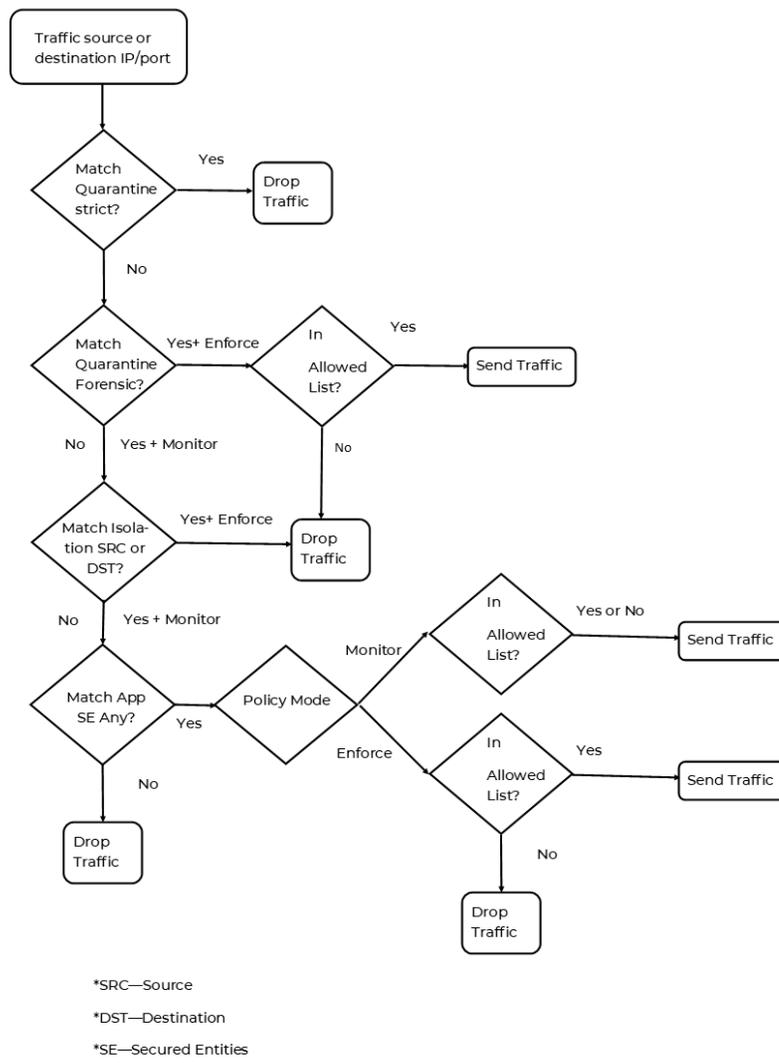


Figure 9: Enhanced Priority Framework Flowchart

- If the isolation environment policy is in the enforce mode, it blocks all traffic to the application, including the traffic that is allowed by the application security policy.
- If the isolation environment policy is in the monitoring mode, it allows all traffic to the application, including any traffic that is disallowed by the application security policy.

Policies with the Same Category as Secured Entity

Flow Network Security Next-Gen allows you to create policies with the same or overlapping categories as secured entities. This feature helps you create micro policies instead of large monolithic single policies.

Note: When the same or overlapping categories are repeated across multiple policies, policy visualization is available only in one policy.

This feature has the following benefits:

- Splitting a large monolithic policy into multiple granular micro policies ensures smooth migration to FNS NG
- Cloning a policy is simple, as the system allows you to save the cloned policy with the same or overlapping categories or secured entities.

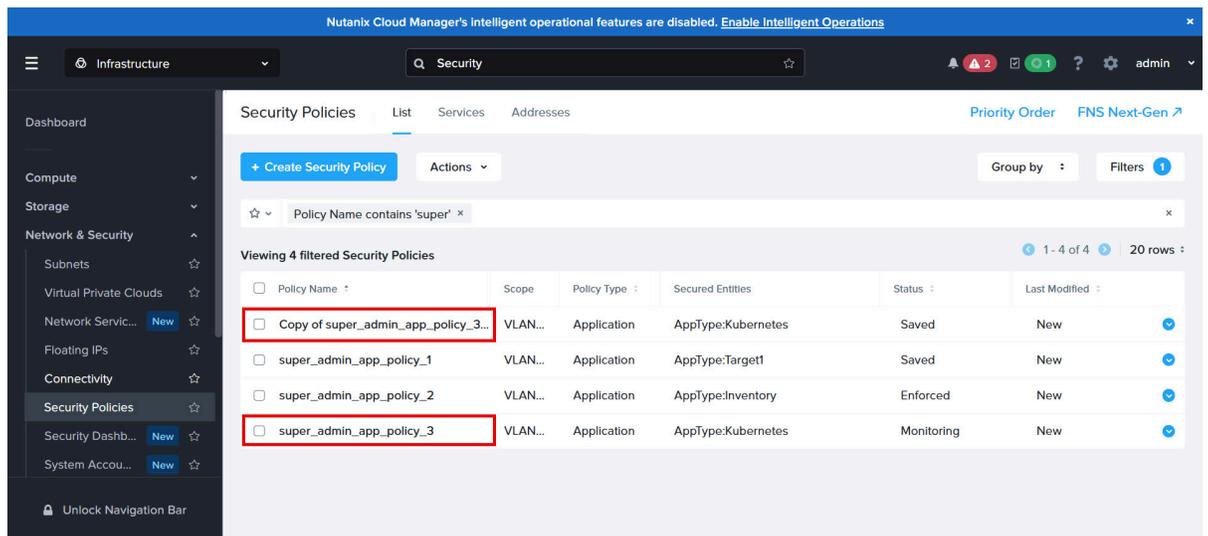


Figure 10: Cloning a User Defined Policy

Types of Policies

The types of policies in Prism Central and their use cases are described here.

Table 2: Types of Policies

Policy Type	Use Case
Application Policy	<p>This is a user defined policy. You can create the following application policy type:</p> <ul style="list-style-type: none">• Generic Policy <p>Use an application policy when you want to secure an application by specifying allowed traffic sources and destinations. This method of securing an application is typically called <i>application ring fencing</i>.</p> <p>For example, use an application security policy when you want to allow only those VMs in the categories <code>department: engineering</code> and <code>department: customersupport</code> (the allowed sources) to communicate with an issue tracking tool in the category <code>AppType: IssueTracker</code> (the secured application), and you want the issue tracking tool to be able to send traffic only to an integrated customer relationship management application in the category <code>AppType: CRM</code>.</p> <p>The secured application itself can be divided into tiers by the use of categories (the built-in <code>AppTier</code> category). For example, you can divide the issue tracking tool into web, application, and database tiers and configure tier-to-tier rules.</p> <div data-bbox="906 1100 1437 1182" style="border: 1px solid #ccc; padding: 5px;"><p>Note: In the new policy model, any category can be used as a secured entity.</p></div> <p>The rules within a policy (application, isolation or quarantine policy) act commonly on a VM. For example: if a VM is in the allowed-list in one policy of an application, then it is allowed-list across all the policies that VM belongs to.</p> <p>For more information, see Application Policy Configuration on page 40.</p>

Policy Type	Use Case
Isolation Policy	<p>This is a user defined policy.</p> <p>Use an isolation environment policy when you want to block all traffic, regardless of direction, between two groups of VMs identified by their category. VMs within a group can communicate with each other.</p> <p>For example, use an isolation environment policy when you want to block all traffic between VMs in the category <code>Environment: sandbox</code> and VMs in the category <code>Environment: production</code>, and you want to allow all the VMs within each of those categories to communicate with each other.</p> <p>For more information, see Isolation Environment Policy on page 60.</p>
VDI Policy	<p>Use a VDI policy when you want to secure your VDI environment.</p>
Quarantine Policy	<p>This is a system defined policy.</p> <p>Use a quarantine policy when you want to isolate a compromised or infected VM and optionally want to subject it to forensics.</p>

Services

Services is a group of protocol-port combination. You can use any of the default services or create a custom service. The ability to use the service entities in the policy creation workflow reduces any manual configuration error and enables reusability of available entities.

- To create or update a custom service, see [Creating a Service](#) on page 25.
- To view the list of available services (built-in and custom services), go to **Network & Security > Security Policies > Services**.

Creating a Service

About this task

To create a custom service, do the following.

Procedure

1. Log on to the Prism Central web console.
2. Select Infrastructure application from the [Application Switcher Function](#), and navigate to **Network & Security > Security Policies > Services**.
3. Click **Create Service Group**.
4. Enter a name and description for the Service.
5. Select the **Protocol** from the drop-down menu and enter the port number or port range in the **Port** field.
You can add multiple protocol-port combinations in a single service. To add more protocol-port combination, click **Add Row** and specify the required values.

6. Click **Save**.

Addresses

Addresses refer to IP addresses. Address groups are sets of contiguous or non-contiguous IP addresses or IP address ranges of IP address pools. You can create an Address Group entity and use this entity for creating security policies. By using address groups in policy creation workflows, you can avoid manual configuration errors. The address groups you create are reusable entities that prevent the repeated creation of IP address sets.

- To create or update a addresses or address group, see [Creating an Address](#) on page 26.
- To import address groups in bulk, see [Importing Address Groups](#) on page 27.

Creating an Address

About this task

To create an Address, do the following.

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#) select Infrastructure application, and navigate to **Network & Security > Security Policies > Addresses**.
3. Click **Create Address**.
4. Enter a name and description for the Address.
5. You can enter IP address or a IP range in the two ways:
 - » Enter or copy-paste the IP address or a IP range in the **Subnet** field or
 - » Click **Import CSV File** and upload the address groups file in CSV format.

You must create an address group file in the CSV format. The following table shows an example of a single address group:

192.127.0.1/32	192.127.0.2/32
192.127.0.3/32	192.127.0.4/32
192.127.0.5/32	192.64.0.1
192.64.0.2	192.64.0.3

6. Click **Create Address**.

What to do next

To view the newly created address group details, click the name of the address group in the Addresses page.

Import Address Groups

Prism Central allows you to import address groups in bulk. This reduces the manual effort required to create a large set of address groups. Create address groups in an spreadsheet and save it in the CSV file format. You can download the CSV template with sample address groups and use it to create your address groups CSV file.

You can import a maximum of 20 address groups at once.

Importing Address Groups

About this task

Prism Central allows you to import address groups in bulk. To import Address Groups in bulk, do the following.

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#) select Infrastructure application, and go to **Network & Security > Security Policies > Addresses**.
3. Click **Import CSV File**.
4. Click **download the template** to download the sample CSV file format.
You can update the template and add multiple address groups.
5. Browse the file containing multiple address groups and click **Import**.
Ensure the file adheres to the standards mentioned in the template.

Note: In case of errors in the CSV file, an appropriate error message is displayed when you import. You can rectify the errors and retry.

After the address groups are imported, they are displayed on the Addresses page.

Built-In Categories for Security Policies

Prism Central includes built-in categories that you can use in application security policies and isolation policies. It also includes a built-in category for quarantining VMs.

Table 3: Built-In Categories

Category	Description
AppTier	Add values for the tiers in your application (such as web, application_logic, and database) to this category and use the values to divide the application into tiers when configuring a security policy.
AppType	Associate the VMs in your application with the appropriate built-in application type such as Exchange and Apache_Spark. You can also update the category to add values for applications not listed in this category.
Environment	Add values for environments that you want to isolate from each other and then associate VMs with the values.

Category	Description
Quarantine	<p>Add a VM to this category when you want to quarantine the VM. You cannot modify this category. The category has the following values:</p> <p>Strict Use this value when you want to block all inbound and outbound traffic.</p> <p>Forensic Use this value when you want to block all inbound and outbound traffic except the traffic to and from categories that contain forensic tools.</p>
ADGroup	<p>This category is managed by ID Based Security (ID Firewall). Each ADGroup value represents an imported group from Active Directory. To add or remove values to use in Flow policies use the ID Based Security configuration page (Prism Central Settings > Flow > ID Based Security).</p>
ADGroup:Default	<p>This category is applied to the VDI VMs of the AD group when the VM inclusion criteria is set and allows you to apply a default set of rules for the VDI VMs (without the requirement of user logons).</p>

Bring in Your Own Category

You can create your own categories and use in application security policies and isolation policies. For information on creating categories, see [Category Management](#) topic in the *Prism Central Infrastructure Guide*.

Avoid using system-defined category names when creating your own categories.

FLOW NETWORK SECURITY NEXT-GEN AND DISASTER RECOVERY

Disaster Recovery (DR) enables you to synchronize and recover the entities during any network outage or planned maintenance activity. Disaster recovery supports synchronizing and recovering Flow Network Security Next-Gen entities at the recovery site in the following scenarios:

- Single PC DR – Network controller managed VLAN and VPC

FNS supports Protection Domain-based DR with OVN VLAN and FNS OVN VLAN policies from pc.2024.3.1 and AOS 7.0.1 version.

Flow Network Security Support for Single-PC Disaster Recovery

FNS Next-Gen supports single PC DR for FNS policies with network controller managed VLAN and VPC.

This feature does not require any new configuration for FNS, as the VM is replicated to another cluster managed by the same Prism Central in which the FNS policy is configured. However, you should configure a single PC DR. For more information, see [Nutanix Disaster Recovery Guide](#). After the VM fails over, the FNS policy is applied automatically to the recovered VM with no impact on the security posture of the secured VM.

This feature ensures VM security posture during:

- planned and unplanned cold migration.
- asynchronous, nearsync, and synchronous replication (Metro).

Note:

IP mapping for recovery plan is supported only when categories are used in FNS policy fields such as Secured Entities, Inbound Rules, and Outbound Rules. IP mapping for recovery plan is not supported when subnets or address groups are used in Inbound or Outbound Rules.

For more information, see [Nutanix Disaster Recovery Guide](#).

Requirements

To support synchronization between the local and remote sites, the two sites should meet the following requirements:

- ANC should be enabled.
- FNS Next-Gen should be enabled.
- FNS Next-Gen version 4.2.0 must be the minimum version on both sites.
- Nutanix clusters must be running AOS version 7.0 or later.
- Nutanix clusters must be running PC version pc.2024.3 or later.
- AHV version 10 or later

Limitations

- FNS Next-Gen currently does not support Cross-Cluster Live Migration (CCLM)
- FNS Next-Gen support for single-PC DR is only for intra-VPC traffic.

POLICY CONSUMPTION AND VISUALIZATION

Policy Views

You can use the following policy views for efficient policy management.

- Security Policies View
- Individual Policy View

Security Policies View

This section describes the functionalities available on the Security Policies landing page.

The Security Policies page has the following tabs:

- List
- Services
- Addresses

The Security Policies List view is also available on the Virtual Private Clouds landing page.

List Tab

The List tab has the following options:

- **Create Security Policy:** Allows you to create an application, VDI or isolation policy. For more information, see [Creating an Application Policy](#), and [Creating Isolation Policy](#).
- **Actions:** You can perform the following actions on a selected security policy:
 - **Update:** Allows you to update a policy
 - **Clone:** Allows you to clone a policy
 - **Delete:** Allows you to delete a policy. You can select multiple policies and perform the bulk delete action
 - **Apply (Monitor):** Allows you to place a policy in a monitor mode.
 - **Apply (Enforce):** Allows you to Enforce a policy that is in Saved mode. You can select multiple application policies and perform the bulk Enforce action
 - **Enable Hit Logs:** Enables policy hit logs
 - **Disable Hit Logs:** Disables policy hit logs
 - **Export All Policies:** Exports all policies and saves them in your local machine
 - **Import Policies:** Imports new policies from your local machine. After the import, existing policies will be overridden with the configuration from the import. And, policies that are not part of this import will be deleted. However, quarantine policies are never deleted.
- **Search:** You can perform a basic search by typing the policy name or using filters

- **Filters:** You can refine your search to a specific policy using the below filters:
 - Policy Name
 - Scope - available only for VPC environments
 - Policy Type
 - Secured Entity
 - Status

For more information, see [Filtering and Grouping](#).

- **Security policies list:** The page displays the list of security policies that are grouped based on policy types: Quarantine, VDI, Isolation, and Application. This page provides the following information on each policy:
 - Policy Name: Displays the name of the policy
 - Scope: Displays the VPC scope of the policy
 - Policy Type: Displays the policy type name
 - Secured Entities: Displays the secured entities that are protected in this policy
 - Purpose: Displays the brief description of the policy
 - Last Modifies: Displays the policy as new or modified
 - In-line Policy Configuration View (down icon): Displays more information on the policy such as Source, Secured Entity, Destination, Services, and Rule Description

Service Tab

The Services tab has the following options:

- **Create Service Group:** You can create a custom service group. For more information, see [Creating a Service](#).
- **Actions:** You can perform the following actions on a service:
 - Update: Allows you to update a service
 - Clone: Allows you to clone a service
 - Delete: Allows you to delete a service
- **Filters:** You can filter the service by clicking the Filter option and selecting the filtering parameters such as name and description.

Addresses Tab

The Addresses tab has the following options:

- **Create Address:** You can create an address entity. For more information, see [Creating an Address](#).
- **Actions:** You can perform the following actions on an address
 - Update: Allows you to update an address
 - Delete: Allows you to delete an address
- **Filters:** You can filter the addresses by clicking the Filter option and selecting the filtering parameters such as name and description.

Individual Policy View

To view the details of a policy, click the name of the policy. The policy opens in a visualization view.

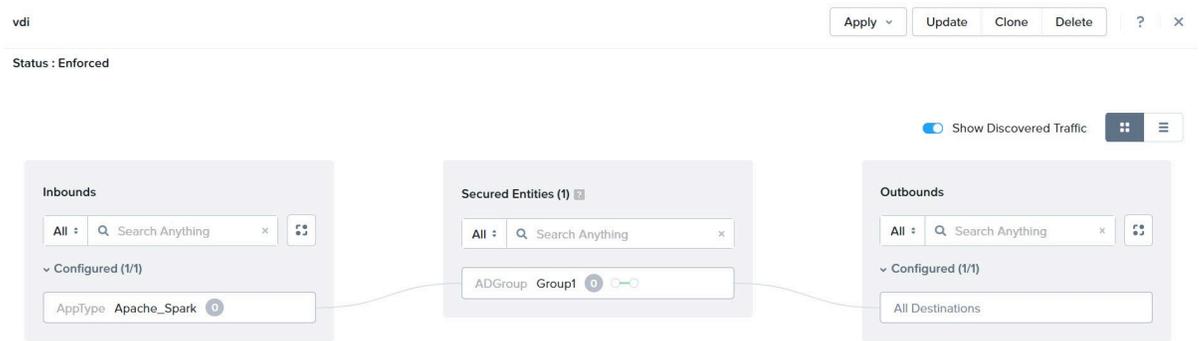


Figure 11: Policy in Visualization View

Click the hamburger icon on the top-right of the page to open the policy in the table view.

Source	Secured Entity	Services	Description
AppType Apache_Spark	ADGroup Group1	All	

Figure 12: Policy in Table View

Both the policy views show the status of the policy and allow you to perform the following actions:

- **Apply** : Allows you to change the policy mode.
- **Update**: Allows you to update an application, isolation or quarantine policy.
- **Clone**: Allows you to clone an application policy.
- **Delete**: Allows you to delete an application or isolation policy.

The visualization view has the following panes:

- **Inbounds**: Displays the allowlist for the source traffic.

- **Secured Entities:** Displays the secured entities protected by the policy. This pane also shows the scope of the VPC to which the policy belongs. To know the details on the protected VMs and the VMs that are not in the policy scope, click the fraction number within the category.

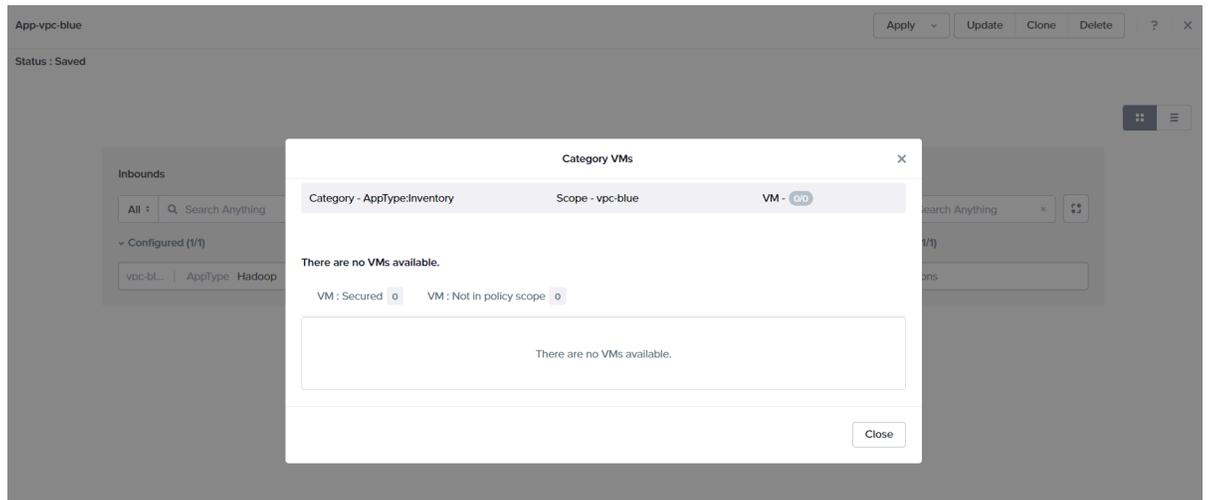


Figure 13: Category VMs

- **Outbounds:** Displays the allowlist for the destination traffic.

The table view displays the following information:

- **Inbound Rules:** Displays more information on the policy such as Source, Secured Entity, Services and Description.
- **Outbound Rules:** Displays more information on the policy such as Secured Entity, Destination, Services and Description.
- **Search:** Allows you to search an entity by following options:
 - All
 - Source
 - Secured Entity
 - Subnets
 - Categories
 - Services
 - Addresses
 - Rule Name
 - TCP/UDP
 - ICMP
 - VM

Applying Filtering and Grouping

You can apply different types of filters to view results based on properties like policy name, policy type, scope, and more. Using the right filter options you can view targeted results and also perform bulk

actions on the filtered result. You can also group related rule attributes together for easier visualization of connection flows. Grouping and filtering work together to provide an intuitive view for the security policy. Also, grouping and filtering work seamlessly when applied to the inbound and outbound rules in the Table view.

About this task

To apply filtering and grouping to a security policy, do the following.

Procedure

1. Log on to the Prism Central web console.
2. Click the hamburger menu and go to **Network & Security > Security Policies**
The **Security Policies** page appears.
3. On the **Security Policies** page, click **Filters** to open the filtering drop-down menu.

Tip:

- You can also use the filtering functionality for the inbound and outbound rules in the **Table** view.
- Filtering works with grouped results as well, see Step 4 for applying grouping.

4. To view specific rule properties, do one of the following.
 - » In the **Type name to filter** search box, enter the policy name that you want to view.
 - » Select any filter from the available filter types. You can refine your search further by entering a particular value for each filter.

Policy Name

filter results by policy name

Scope

filter results by VPC name

Policy Type

filter results by policy type (application, quarantine, and isolation)

Secured Entity

filter results for a secured entity using the category name

Status

filter results based on the policy enforcement status (enforced, monitoring, and saved)

5. To group the list of policies, go to **Security Policies > List** page, then click **Group by** and choose from **Policy Type** or **Status**.

The group option organizes related rule attributes like subnet IP, categories, and service in distinct boxes. To view all the entities belonging to a group, click the down-arrow icon to expand the group.

Tip: You can also use the grouping functionality for the inbound and outbound rules in the **Table** view. The group option organizes related rule attributes like subnet IP, categories, and service in distinct boxes. To view all the

entities belonging to a group, click the down-arrow icon to expand the group. To apply grouping on the inbound and outbound rules of a policy, do the following:

1. Open any security policy.
2. Switch to the **Table** view.
3. Click **Group**.

Allowing Discovered Traffic

About this task

When a policy is set in the enforce mode, the policy engine discovers the traffic denied by the policy. The discovered traffic can be viewed in the policy details page of the UI. To allow the discovered traffic, manually update the policy. Perform the following steps to allow the discovered traffic.

Procedure

1. Log on to the Prism Central web console.
2. Select Infrastructure application from the [Application Switcher Function](#), and navigate to **Network & Security > Security Policies**.
3. In the **Security Policies** page, click and open any policy.
The policy page shows the discovered traffic in the **Inbounds** and **Outbounds**.
4. Click **Update**.
The **Update Security Policy** page appears.
5. Click **Next**. You can perform this step in the List view or Visual view.
 - **List view:** Click **Inbound Rules** or **Outbound Rules** and under **Discovered Traffic** section, select the traffic source and click **Allow**.
 - **Visual view:** In the **Inbounds > Discovered** or **Outbounds > Discovered** section, hover over the traffic source and click **Allow Traffic**.
6. In the **Allow Traffic** page, select **Show Tiers** or **Show All Traffic** to display the discovered traffic.
Show Tiers: Shows a source trying to reach different AppTypes within an AppTier.
Show All Traffic: Shows a consolidated list of sources trying to reach all AppTiers (secured VMs).
7. In the Discovered traffic section, select the source you want to allow and choose a service.
You can select more than one traffic source. Also, you can create a new service if the required service is not available. For more information, see [Creating a Service](#) on page 25.
8. Click **Allow Discovered Traffic** and click **Next**.
9. In the **Review** page, select a policy mode and click **Confirm** to allow discovered traffic.

ROLE-BASED ACCESS CONTROL

Flow Network Security Next-Gen supports role-based access control (RBAC) and granular 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 and the users assigned to those roles.

Granular RBAC enables you to create custom roles with finer permissions, like allowing a user to update a specific address group as compared to the broader permissions to update all address groups. For example, you can create a role with permissions to create flow policies in the monitor and save mode. The user has access to create, update, and view flow policies in monitor and save mode only.

Support for categories RBAC is available.

For more information, see [Controlling User Access \(RBAC\)](#) in the *Nutanix Security Guide*.

Limitations

- Support for bulk operations on policies for a custom role is not available.
- A user without permission to access the address group, service group, and flow policy can import the address group, service group, and flow policy with only import permission.
- The system does not support selective policy import or export. A user can import or export all policies.
- The system revokes any RBAC permissions to a user on the newly imported policies, as after importing new policies, the existing policies are overridden with new policies. However, this does not affect users with Full Access or All Entities access for flow policies.
- A user can view policies in the policy list even if access to VPC is not provided.
- The system does not allow you to add an AD server config entity to a role when a user has access to only self-owned entities. You must provide access to all entities as well.
- The system does not list the categories on the policy details page when you filter the category as Category_Key.

Flow Network Security Roles and Permissions

Prism Central provides the following roles for Flow Network Security management:

- Flow Admin
- Flow Policy Author
- Flow Viewer

The table provides the list of permissions that are available for the two roles.

Note: Even though the roles have pre-configured permissions, they are not effective. You must define the scope for the entities at the time of creating authorization policy.

Table 4: FNS Roles and Permissions

Entity	Flow Admin	Flow Policy Author	Flow Viewer
Address Group			
Create Address Group	Yes	Yes	No

Entity	Flow Admin	Flow Policy Author	Flow Viewer
Delete Address Group	Yes	Yes	No
Update Address Group	Yes	Yes	No
View Address Group	Yes	Yes	Yes
AHV VM			
View AHV VM	Yes	Yes	Yes
Category			
Create Category	Yes	No	No
Delete Category	Yes	No	No
Update Category	Yes	No	No
View Category	Yes	Yes	Yes
Category Mapping			
Create Category Mapping	Yes	Yes	No
Delete Category Mapping	Yes	Yes	No
Update Category Mapping	Yes	Yes	No
View Category Mapping	Yes	Yes	Yes
Directory Service			
Search Directory Service	Yes	Yes	Yes
View Directory Service	Yes	Yes	Yes
Directory Server Config			
Create Directory Server Config	Yes	Yes	No
Delete Directory Server Config	Yes	Yes	No
Update Directory Server Config	Yes	Yes	No
View Directory Server Config	Yes	Yes	Yes
Flow Policy-			
Create Flow Policy	Yes	Yes	No
Delete Flow Policy	Yes	Yes	No
Update Flow Policy	Yes	Yes	No
View Flow Policy	Yes	Yes	Yes

Entity	Flow Admin	Flow Policy Author	Flow Viewer
View Rules Flow Policy	Yes	Yes	Yes
Export Flow Policy	Yes	Yes	No
Import Flow Policy	Yes	Yes	No
Identity Categorization Config			
Update Identity Categorization Config	Yes	Yes	No
View Identity Categorization Config	Yes	Yes	Yes
Service Group			
Create Service Group	Yes	Yes	No
Delete Service Group	Yes	Yes	No
Update Service Group	Yes	Yes	No
View Service Group	Yes	Yes	Yes

Creating an Authorization Policy for FNS Next-Gen

About this task

An authorization policy helps you map a role to a user or group and allows you to configure more precise role assignments (AHV only) such as providing access to a specific entity, or an individual entity.

To assign a role to selected users or groups and to provide access to a specified set of FNS Next-Gen entities to the role, do the following:

Procedure

1. Log in to Prism Central as an administrator or any user with super admin access.
2. (Optional) Select **Admin Center** in the Application Switcher. Go to **IAM > Settings > Authentication**, click **+ New Directory** and add your preferred Active Directory.

This step configures Active Directory settings. You can skip this step if an Active Directory is already configured

3. Select **Admin Center** in the Application Switcher.
4. Select **IAM** and go to **Roles**.
The page displays system defined and custom roles.
5. Select a role in the roles dashboard, then click **Actions > Add Authorization Policy**.

You can select any of the system-defined roles or a custom role.

The **Create New Authorization Policy** page appears.

6. (Optional) Click the edit icon to edit the name of the authorization policy.
7. In the **Choose Role** tab, ensure that the role that you selected in Step 5 appears and click **Next**.
You have the option to change the role in this step.

8. In the **Define Scope** tab, do any of the following:

- Select **Full access** to provide full access to all entities and instances.
- Select **Configure access** to provide access to selected entity types and instances.

Note: FNS roles do not support cluster and category-based scoping for All Entities.

You can configure granular RBAC for these entities:

- **Address Group:** Filter using **Individual Entity**, search for an address group, or provide access to all address groups.
- **Service Group:** Filter using **Individual Entity**, search for a service group, or provide access to all service groups.
- **Directory Service:** Filter by Individual Entity
- **Flow Policy:** Filter based on **Individual Entity** or **Mode**. If you filter using Individual Entity, then search for a specific flow policy or select All Flow Policies. If you filter using Mode, select ENFORCE, MONITOR, or SAVE mode. The user has access to view, update, and create flow policies in the specified mode depending on the role.

Note: To export policies, a user must have permission for all the flow policies.

The following are the other entities that you can configure access to the role:

- **AHV VM:** Filter by Individual Entity, In Category, In Project, or On Cluster
- **Category:** Filter by Individual Entity or Category Key
- **Category Mapping:** Users has access to All Category Mapping
- **Directory Server Config:** Users get access to All Directory Server Config
Directory Server Config is used for v4 APIs.
- **Identity Categorization Config:** Users get access to All Identity Categorization Config
Identity Categorization Config is used for v3 API calls. If you provide the Identity Categorization Config entity access to a role and try to access the ID Based Security page in Prism settings, then the system does not load the page as the UI uses v4 APIs. Therefore, you must use the View Directory Server Config entity permission for a role.
- Check **Allow users access to entities created by them** field to allow users or groups to access the entities that they create.

9. Click **Next**.

10. In the **Assign Users** tab, do the following:

- a. Select the configured AD or IDP from the dropdown menu.

Typing few letters in the search field displays a list of users from which you can select, and you can add multiple user names in this field.

- b. Select **Local User** to add a local user or group.

Typing few letters in the search field displays a list of users from which you can select, and you can add multiple user names in this field.

11. Click **Save**.

APPLICATION POLICY CONFIGURATION

You can create the following type of application policy:

- **Generic Policy:** An application policy that includes user-defined and system-defined categories apart from AD groups.
- **VDI Policy:** The VDI Policy is based on identity-based categorization of the VDI VMs using Active Directory group membership. You should add an Active Directory Domain to create VDI policies. To add an AD, go to **Prism Central Settings > ID Based Security**. For more information, see [VDI Policy Configuration](#) on page 50.

Creating an Application Policy with a VLAN Scope

Perform the procedure to create an application policy to secure entities with a VLAN scope.

Before you begin

- Ensure you have created Network Controller managed VLANs.
- Create the categories you need and associate the VMs that you want to protect with those categories. You can create categories for the following purposes. Some categories or category values are required while others are optional:
 - Every security policy must be associated with a value in the AppType category, so make sure that you update the AppType category with appropriate values if the built-in values do not work for you. For information about this category and its values, see *Category Details View* in the [Prism Central Infrastructure Guide](#).
 - If you need to apply the policy to an application in a specific environment (for example, development, test, or production) or an application at a specific location, create the category you need and apply it to the application. Prism Central includes a built-in Environment category that you can use or update with values of your own. You can also create your own categories.
 - If you want to specify categories for traffic sources and destinations instead of allowing all inbound and outbound traffic, create those categories and apply them to the traffic sources and destinations.
 - The AppTier category has a built-in `default` value, but you can update the category to add values of your choice.

For information about categories and their values, see *Category Management* in the [Prism Central Infrastructure Guide](#).

- Security policy configuration might require more time than the default session timeout allows you. You might want to increase the session timeout so that you do not lose a configuration that is left unattended while you perform associated tasks such as referring to this documentation. For information about changing session timeout, see *Configuring Prism Central UI Settings* in the [Prism Central Admin Center Guide](#).

About this task

To secure an application, do the following:

Procedure

1. Log on to Prism Central.

2. Select Infrastructure application from the [Application Switcher Function](#), and navigate to **Network & Security > Security Policies**.

The **Security Policies** page is displayed.

3. On the **Security Policies** page, click **Create Security Policy** and then **Secure an Application**.

The **Create App Security Policy** page is displayed.

4. On the **Define Policy** tab, do the following in the indicated fields:

- a. **Policy Name**: Enter a name for the security policy.

- b. **Purpose**: Describe the purpose of the security policy.

- c. **Secure Entities (Application Policy)**: Select Generic Policy to create an application policy that includes user-defined and system-defined categories apart from AD groups.

- d. In the **Scope of Secured Entities** section, do the following:

- Select the **VLAN Subnets** option to secure VMs that span across one or more Network Controller managed VLANs.

Note: In FNS NG, policies will not have any effect on VMs that are part of VLAN Basic subnets.

- e. In the **Advanced Configuration** section:

- Click the **Allow** option to allow **IPv6 traffic**. The policy rules apply to IPv4 traffic only and all IPv6 traffic is blocked by default.

Note: If you choose to block IPv6 traffic, the IPv6 traffic remains blocked even in the monitoring mode.

- Click the **Enabled** option (in the **Policy Hit Logs** field) to collect flow policy hit logs on a syslog server. Policy hit logs track network flows; whether the flows were allowed or denied by a specific policy. Policy hit logs are useful to determine if specific traffic is present on the network and how a security policy affects the traffic.

You can configure syslog monitoring for the policy hit logs for Flow. For more information, see *Configuring Syslog Monitoring* in the [Prism Central Admin Center Guide](#).

Note: Policy hit logs are not generated if both source and destination are in inbound or outbound category.

5. Click **Next**.

The **Secure Application** tab is displayed.

The schematic on the Secure Application tab is divided into three areas of configuration:

- **Inbounds** : to configure allowlist for the source traffic. The VMs that are part of Network Controller managed VLANs as well as VLAN Basic subnets can be added in the inbounds.
- **Secured Entities** : to add an application as a secured entity. FNS policies work only on VMs part of Network Controller managed VLANs.
- **Outbounds** : to configure allowlist for the destination traffic. The VMs that are part of Network Controller managed VLANs as well as VLAN Basic subnets can be added in the outbounds.

6. On the **Secure Application** tab, do the following, and then click **Next**:
- In the **Secured Entities** section, to add an application as a secured entity, click **Add Secured Entity** and start typing the name of the default or custom category for the application that you want to protect. Then, click **Add**. You can add multiple categories in a single policy.

You can create different application policies using categories assigned to a same VM.

The Secure Entities tab displays the scope of the policy at the top-right corner. For example, the current policy belongs to the **Scope: VLAN Subnets**. To know the details on the protected VMs and the VMs that are not in the policy scope, click the fraction number within the category.

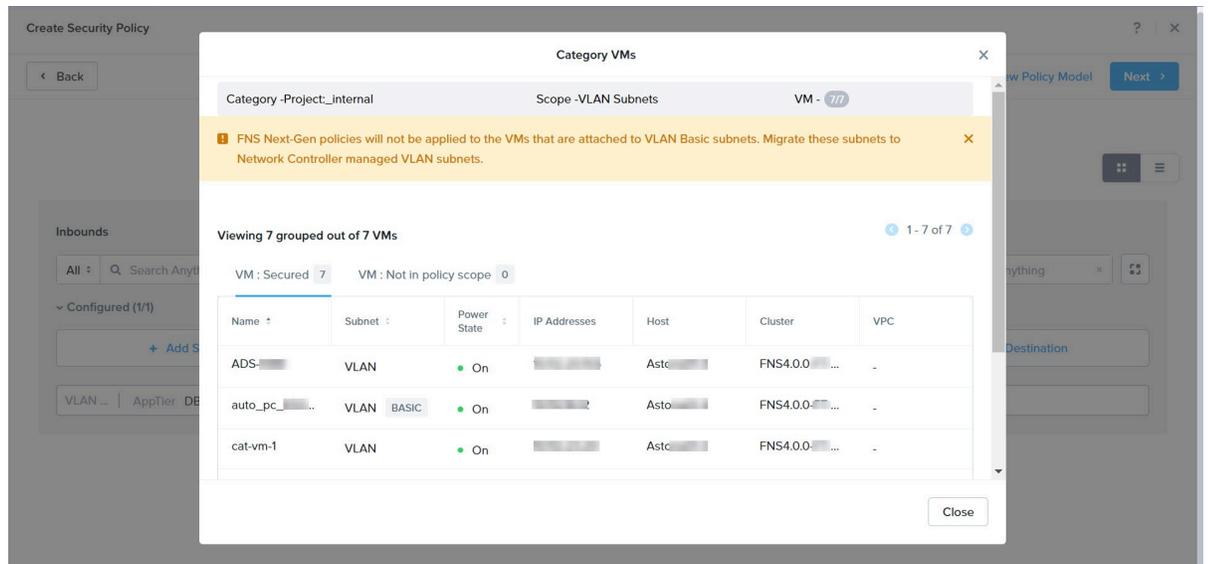


Figure 14: Category VMs

If any of the VMs are in VLAN Basic subnet, a message to migrate VLAN Basic subnets to Network Controller managed VLAN subnets appears.

- In the **Inbound** section to add traffic sources, do the following:

Note: You can create your own category and add as inbound apart from the built-in categories such as AppType, and AppTier.

- Click **Add Source**, and then do the following:
 - Select one of the following options from the **Add source by** drop-down list:
 - Category:** Allows traffic only if that traffic originates from entities that are in the selected category.
 - Subnet/IP:** Allows traffic only if that traffic originates from entities that are in the selected subnet.
 - Addresses:** Allows traffic only if the traffic originates from the entities that are in the selected address.
 - Allow All:** Allows traffic to all destinations.
 - Enter the value in the text box, and then click **Add**. To edit any previously entered value or destination option, click **Edit**.

When entering the name of a category, a list of matching names appears, and select the name you want to specify. The subnet mask must be specified in the CIDR format.

When entering the address, a list of available address group names appears. Select the address group or you can create a new address group.

3. To add another category, subnet or address, click **Add Source**. Add as many categories, subnets or address as you want to allow.

- Each entry in this list represents a stream of inbound traffic.

- c. To add traffic destinations, on the **Outbound** side, do the following:

Note: You can create your own category and add as outbound apart from the built-in categories such as AppType, and AppTier.

- Click **Add Destination**, and then do the following:

1. Select one of the following options from the **Add destination by** drop-down list:

- **Category:** Allows traffic only if that traffic is destined for entities in the selected category.

If you have selected **Inside a VPC** option in Step 2d, then only the categories that are available within the VPC are listed here.

- **Subnet/IP:** Allows traffic only if that traffic is destined for entities in the selected subnet.

- **Addresses:** Allows traffic only if the traffic originates from the entities that are in the selected address.

- **Allow All:** Allows traffic to all destinations.

2. Enter the value in the text box, and then click **Add**. To edit any previously entered value or destination option, click **Edit**.

When entering the name of a category, a list of matching names appears, and select the name you want to specify. The subnet mask must be specified in the CIDR format.

When entering the address, a list of available address group names appears. Select the address group or you can create a new address group.

3. To add another category, subnet or address, click **Add Destination**. Add as many categories, subnets or address, as you want to allow.

Each entry in this list represents a stream of outbound traffic.

- To specify the protocols that you want to allow from each stream of inbound and outbound traffic, do the following:

1. Click the traffic source or traffic destination (a category or subnet if you have configured a allowlist or **All Sources** if you have chosen to allow all sources) for which you want to create a rule.

2. Click the plus icon that appears on the application (if you are treating the application as a single entity) or application tier (if you have divided the application into tiers). The **Create Inbound Rule** or **Create Outbound Rule** dialog box appears.

3. Enter a description for the rule.

4. In **Service Details**, click **Allow all traffic** to allow all types of traffic or click **Select a service** to choose any default or custom service.

5. Click **Save**.

After you configure a rule, a dotted line appears between the two endpoints of the rule. Point to the dotted line to show the list of ports that the rule allows.

7. On the **Review** tab, select a policy mode from the following options:

- **Save:** Saves the policy without applying it. Saving a policy allows you to retain the policy in a draft stage without having the need to apply (enforce) it at the time of creation.
- **Apply (Monitor):** Monitors how the security policy works.

When a policy is in the Monitor state, the application continues to receive all traffic even from the disallowed source, but disallowed traffic is highlighted on the monitoring page. Traffic is not blocked until the policy is applied.

- **Apply (Enforce):** Applies the security policy on the application. The application receives and transmits traffic only from/to the allowed sources and destinations. The policy blocks the traffic from entities that are not defined as sources.

The **Review** tab also displays the policy configuration summary.

8. Click **Confirm**.

Creating an Application Policy within a VPC Scope

Perform the procedure to create an application policy to secure entities within a VPC scope.

Before you begin

- You must create at least one VPC before creating an application policy to secure entities within a VPC.
- Create the categories you need and associate the VMs that you want to protect with those categories. You can create categories for the following purposes. Some categories or category values are required while others are optional:
 - Every security policy must be associated with a value in the AppType category, so make sure that you update the AppType category with appropriate values if the built-in values do not work for you. For information about this category and its values, see *Category Details View* in the [Prism Central Infrastructure Guide](#).
 - If you need to apply the policy to an application in a specific environment (for example, development, test, or production) or an application at a specific location, create the category you need and apply it to the application. Prism Central includes a built-in Environment category that you can use or update with values of your own. You can also create your own categories.
 - If you want to specify categories for traffic sources and destinations instead of allowing all inbound and outbound traffic, create those categories and apply them to the traffic sources and destinations.
 - The AppTier category has a built-in `default` value, but you can update the category to add values of your choice.

For information about categories and their values, see *Category Management* in the [Prism Central Infrastructure Guide](#).

- Security policy configuration might require more time than the default session timeout allows you. You might want to increase the session timeout so that you do not lose a configuration that is left unattended while you perform associated tasks such as referring to this documentation. For information about changing session timeout, see *Configuring Prism Central UI Settings* in the [Prism Central Admin Center Guide](#).
- You can create up to 1000 user-defined application policies.

About this task

To secure an application, do the following:

Procedure

1. You can create an application policy at a global context or a VPC context.

- » To create an application policy at a global context, go to **Network & Security > Security Policies**, click **+ Create Security Policy**
- » To create an application policy at a VPC context, go to **Network & Security > Virtual Private Clouds**, select a VPC. From the **Security Policies** tab click **+ Create Security Policy**.

The **Create Security Policy** page is displayed.

2. On the **Define Policy** tab, do the following in the indicated fields:

- a. **Policy Name:** Enter a name for the security policy.
- b. **Purpose:** Describe the purpose of the security policy.
- c. **Secure Entities (Application Policy):** Select Generic Policy to create an application policy that includes user-defined and system-defined categories apart from AD groups.
- d. In the **Scope of Secured Entities** section, do the following:
 - Select the **Subnets Inside a VPC** option to secure VMs that are part of a VPC. Select a VPC from the drop-down.

Note: In FNS NG, policies will not have any effect on VMs that are part of VLAN Basic subnets.

e. In the **Advanced Configuration** section, by default **Allow IPv6 traffic** is blocked and **Policy Hit Logs** is disabled.

Note: The system does not allow you to change Advanced Configuration as FNS Next-Gen does not support:

- IPv6 traffic.
- Policy Hit Logs.

3. Click **Next**.

The **Secure Application** tab is displayed.

The schematic on this tab can be divided into three areas of configuration:

- **Inbounds:** to configure allowlist for source traffic. The VMs that are part of Network Controller managed VLANs as well as VLAN Basic subnets can be added in the inbounds.
- **Secured Entities:** to add an application as a secured entity. FNS policies work only on VMs part of Network Controller managed VLANs.
- **Outbounds:** to configure allowlist for destination traffic. The VMs that are part of Network Controller managed VLANs as well as VLAN Basic subnets can be added in the outbounds.

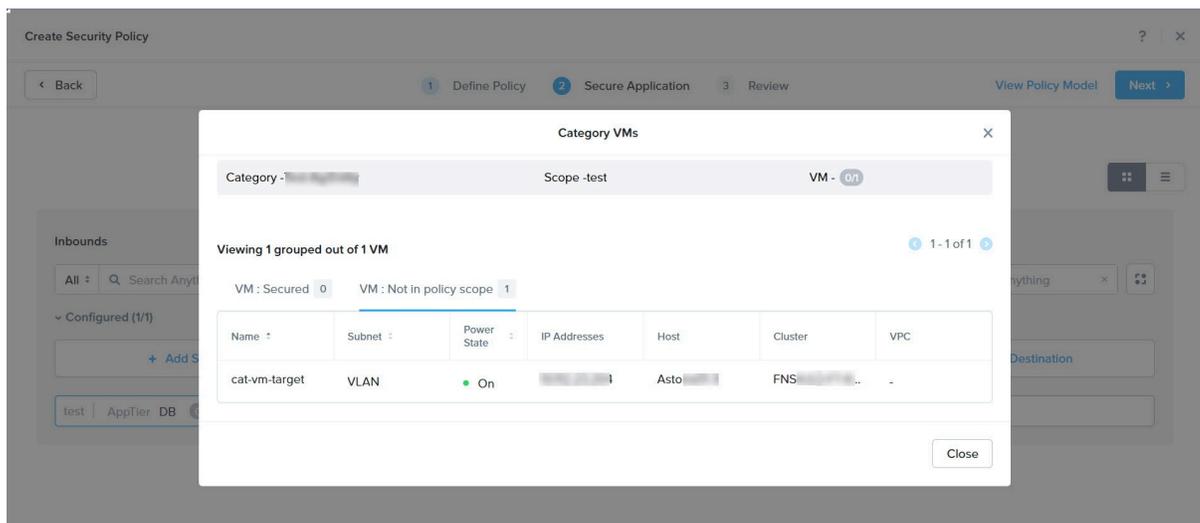
4. On the **Secure Application** tab, do the following, and then click **Next**:

- a. In the **Secured Entities** section, to add an application as a secured entity, click **Add Secured Entity** and start typing the name of the default or custom category for the application that you want to protect. Then, click **Add**. You can add multiple categories in a single policy.

You can create different Application policies using categories assigned to a same VM.

The Secure Entities tab displays the scope of the policy at the top-right corner. For example, the current policy belongs to the VPC **Scope: test**. To know the details on the protected VMs and the VMs that are not in the policy scope, click the fraction number within the category.

Figure 15: Category VMs



- b. In the **Inbound** section to add traffic sources, do the following:

Note: You can create your own category and add as inbound apart from the built-in categories such as AppType, and AppTier.

- Click **Add Source**, and then do the following:
 1. Select one of the following options from the **Add source by** drop-down list:
 - **Category:** Allows traffic only if that traffic originates from entities that are in the selected category.
 - **Subnet/IP:** Allows traffic only if that traffic originates from entities that are in the selected subnet.
 - **Addresses:** Allows traffic only if the traffic originates from the entities that are in the selected address.
 - **Allow All:** Allows traffic to all destinations.
 2. Enter the value in the text box, and then click **Add**. To edit any previously entered value or destination option, click **Edit**.

When entering the name of a category, a list of matching names appears, and select the name you want to specify. The subnet mask must be specified in the CIDR format.

When entering the address, a list of available address group names appears. Select the address group or you can create a new address group.

3. To add another category, subnet or address, click **Add Source**. Add as many categories, subnets or address as you want to allow.
- Each entry in this list represents a stream of inbound traffic.
- c. To add traffic destinations, on the **Outbound** side, do the following:

Note: You can create your own category and add as outbound apart from the built-in categories such as AppType, and AppTier.

- Click **Add Destination**, and then do the following:
 1. Select one of the following options from the **Add destination by** drop-down list:
 - **Category:** Allows traffic only if that traffic is destined for entities in the selected category.
If you have selected **Inside a VPC** option in Step 2d, then only the categories that are available within the VPC are listed here.
 - **Subnet/IP:** Allows traffic only if that traffic is destined for entities in the selected subnet.
 - **Addresses:** Allows traffic only if the traffic originates from the entities that are in the selected address.
 - **Allow All:** Allows traffic to all destinations.
 2. Enter the value in the text box, and then click **Add**. To edit any previously entered value or destination option, click **Edit**.

When entering the name of a category, a list of matching names appears, and select the name you want to specify. The subnet mask must be specified in the CIDR format.

When entering the address, a list of available address group names appears. Select the address group or you can create a new address group.
 3. To add another category, subnet or address, click **Add Destination**. Add as many categories, subnets or address, as you want to allow.

Each entry in this list represents a stream of outbound traffic.

- To specify the protocols that you want to allow from each stream of inbound and outbound traffic, do the following:
 1. Click the traffic source or traffic destination (a category or subnet if you have configured a allowlist or **All Sources** if you have chosen to allow all sources) for which you want to create a rule.
 2. Click the plus icon that appears on the application (if you are treating the application as a single entity) or application tier (if you have divided the application into tiers). The **Create Inbound Rule** or **Create Outbound Rule** dialog box appears.
 3. Enter a description for the rule.
 4. In **Service Details**, click **Allow all traffic** to allow all types of traffic or click **Select a service** to choose any default or custom service.
 5. Click **Save**.

After you configure a rule, a dotted line appears between the two endpoints of the rule. Point to the dotted line to show the list of ports that the rule allows.

5. On the **Review** tab, select a policy mode from the following options:

- **Save:** Saves the policy without applying it. Saving a policy allows you to retain the policy in a draft stage without having the need to apply (enforce) it at the time of creation.
- **Apply (Monitor):** Monitors how the security policy works.

When a policy is in the Monitor state, the application continues to receive all traffic even from the disallowed source, but disallowed traffic is highlighted on the monitoring page. Traffic is not blocked until the policy is applied.

Note: Monitor mode is not supported for a policy with VPC scope.

- **Apply (Enforce):** Applies the security policy on the application. The application receives and transmits traffic only from/to the allowed sources and destinations. The policy blocks the traffic from entities that are not defined as sources.

The **Review** tab also displays the policy configuration summary.

6. Click **Confirm**.

Modifying an Application Policy

About this task

To modify an application policy, do the following:

Procedure

1. Log on to the Prism Central web console.
2. Select Infrastructure application from the [Application Switcher Function](#), and navigate to **Network & Security > Security Policies**.
3. In the **Security Policies** page, select the policy to modify.
4. From the **Actions** drop-down click **Update**.
5. Make the changes, select the policy mode and then click **Confirm**.

Applying an Application Policy

Applying a security policy enforces the security policy on the application, and any traffic from sources that are not allowed is blocked.

About this task

To apply an application policy, do the following:

Procedure

1. In the **Security Policies** page, select the policy.

2. From the **Actions** drop-down select any of the actions:

- **Apply (Monitor)**: Select this mode when you want to monitor how the security policy works.

When a policy is in the Monitor state, the application continues to receive all traffic even from the disallowed source, but disallowed traffic is highlighted on the monitoring page. Traffic is not blocked until the policy is applied.

- **Apply (Enforce)**: Select this mode when you want to apply the security policy on the application. When the policy is Enforce state, the application receives and transmits traffic only from/to the allowed sources and destinations. The policy blocks the traffic from entities that are not defined as sources.

3. Confirm by typing **ENFORCE** or **MONITOR** as applicable in the dialog box and then click **Confirm**.

Cloning a Security Policy

About this task

To clone an existing security policy, do the following.

Procedure

1. In the **Security Policies** page, select an existing application policy to clone.

Note: You can clone only one policy at a time.

2. From the **Actions** drop-down, click **Clone**.

3. Edit the name and purpose of the cloned policy in the **Define Policy** tab. The default name of the cloned policy is **Copy of <Policy Name>**.

Note: You can not change the policy type while cloning a policy.

4. On the **Secure Application** tab, you can edit the policy definition including inbound rules, outbound rules, or secured application configurations.

The system does not allow you to save the cloned policy that has the same secured entities as that of the original policy. Therefore, you must update the secured entities.

Deleting an Application Policy

About this task

To delete an application policy, do the following:

Procedure

1. In the **Security Policies** page, select the policy to delete.

You can select multiple policies and delete them at once.

2. From the **Actions** drop-down click **Delete** .

3. Confirm by typing **DELETE** in the dialog box, and then click **Confirm**.

VDI POLICY CONFIGURATION

The **VDI Policy** is based on identity-based categorization of the VDI VMs using Active Directory group membership. Configuring VDI policy includes adding an Active Directory domain that is used for the ID firewall (**ID Based Security**) and configuring a service account for the domain.

ID Based Security

ID firewall is an extension to Flow that allows you to write security policies based on users and groups in an Active Directory domain in which your VDI VMs are attached. When using ID firewall, you can import groups from Active Directory into Prism Central as categories (in the category key ADGroup), and then write policies around these categories, just as you would for any other category. A new type of policy has been added for this purpose - the **VDI Policy**. ID firewall takes care of automatically placing VDI VMs in the appropriate categories on detecting user logons into the VM hosted on Nutanix infrastructure associated with Prism Central, thus allowing user and group based enforcement of Flow policies.

- See [Configuring Active Directory Domain Services](#) on page 55 to import user groups for identity-based security policies.
- See [Creating a VDI Policy](#) on page 50 to create a VDI policy.
- See [Default VDI Policy](#) on page 55 configuration to define a default VDI policy.

Note:

- It is recommended to disable credential caching on VDI VMs for Flow ID Firewall. The Flow ID Firewall checks the domain controller events for logon attempts. If the VM connection to the domain controller is not available, a user is able to logon (if credential caching enabled) but no event is generated on the domain controller inhibiting the ID Firewall to detect the logon.
- To disable credential caching, see *Interactive logon: Number of previous logons to cache (in case domain controller is not available)* on *Microsoft* documentation website.
- A basic assumption of VDI Policies is that a single end-user is logged on to each desktop VM at a point in time. As a result, if multiple users log into a single desktop VM at once, the security posture of the VM may change in unpredictable ways. Please ensure that for predictable behavior, only one user is logged into the desktop VMs at a time.
- FNS does not support VDI policy for a VPC scope.

Creating a VDI Policy

ID firewall integrates Nutanix Flow with Microsoft Active Directory (AD), such that the groups in the AD can be imported into Prism Central as categories. These imported categories can then be used in the VDI policy as target groups, inbound traffic, and outbound traffic. Prism Central automatically places VMs inside the imported AD group categories when user logons are detected on VMs that are part of the Active Directory domain and also present on Nutanix managed clusters, thus applying security policies based on user group membership.

Before you begin

Note:

- Flow ID firewall is supported only for AHV host compatible with AOS version 5.17 and above and Prism Central version 5.17 and above.

- Flow ID firewall does not detect user logoffs. The policy applied to a VM is kept applied until next user logon on the same VM.
- VMs with an **AppType** category assignment do not get categorized by **ID Based Security**.
- You can use the [Default VDI Policy](#) on page 55 to apply a default set of rules for the VDI VMs (without the requirement of user logons).
- Since a VM user can be a member of multiple ADGroups that are mapped into Prism Central from Active Directory, when a user logs on, a VM may be placed in multiple ADGroups at once. This is the correct behavior, and the policy applied to the VM will be a union of the respective combination of inbounds and outbounds across all ADGroups the VM is placed into.

- If not already available, configure an Active Directory domain that is used for ID firewall, see [Configuring Active Directory Domain Services](#) on page 55.
- Configure a service account with required configuration for the Active Directory domain, see [Configure Service Account for ID Firewall](#) on page 57.

About this task

To secure a VDI environment, do the following:

Procedure

1. In the **Security Policies** page, click **Create Security Policy**.
The **Create Security Policy** page is displayed.

2. On the **Define Policy** tab, do the following in the indicated fields, and then click **Next**:
 - a. **Policy Name**: Enter a name for the security policy.
 - b. **Purpose**: Describe the purpose of the security policy.
 - c. **Secure Entities (Application Policy)**: Select **VDI Policy**.
 - d. **Scope of Secured Entities**: VDI policy does not support VPC scope., therefore, by default **VLAN Subnets** option is selected. The VLAN scope secures VMs that span across one or more Network Controller managed VLANs.

Note: In FNS NG, VDI policies will not have any effect on VMs that are part of VLAN Basic subnets.

- e. In the **Advanced Configuration** section, do the following:

- Click the **Allow** option to allow **IPv6 traffic**. The policy rules apply to IPv4 traffic only and all IPv6 traffic is blocked by default.

Note: If you choose to block IPv6 traffic, the IPv6 traffic remains blocked even in the monitoring mode.

- Click the **Enabled** option (in the **Policy Hit Logs** field) to collect flow policy hit logs on a syslog server. Policy hit logs track network flows; whether the flows were allowed or denied by a specific policy. Policy hit logs are useful to determine if specific traffic is present on the network and how a security policy affects the traffic.

You can configure syslog monitoring for the policy hit logs for Flow. For more information, see *Configuring Syslog Monitoring* in the [Prism Central Admin Center Guide](#).

Note: Policy hit logs are not generated if both source and destination are in inbound or outbound category.

3. In the **Secure Application** tab, do the following in the indicated fields and click **Next**.

- a. **Inbound Traffic:** Click **+ Add Source** and enter the category or subnets that the VDI group can receive the traffic from, as the source.

The VMs that are part of Network Controller managed VLANs as well as VLAN Basic subnets can be added in the inbounds.

- b. **Secured Entities:** Click **+Add AD Group** to select the AD groups (categorized VDI VMs) that you want to secure.

FNS policies work only on VMs part of Network Controller managed VLANs.

The Secure Entities tab displays the scope of the policy at the top-right corner. For example, the current policy belongs to the **Scope: VLAN Subnets**. To know the details on the protected VMs and the VMs that are not in the policy scope, click the fraction number within the category.

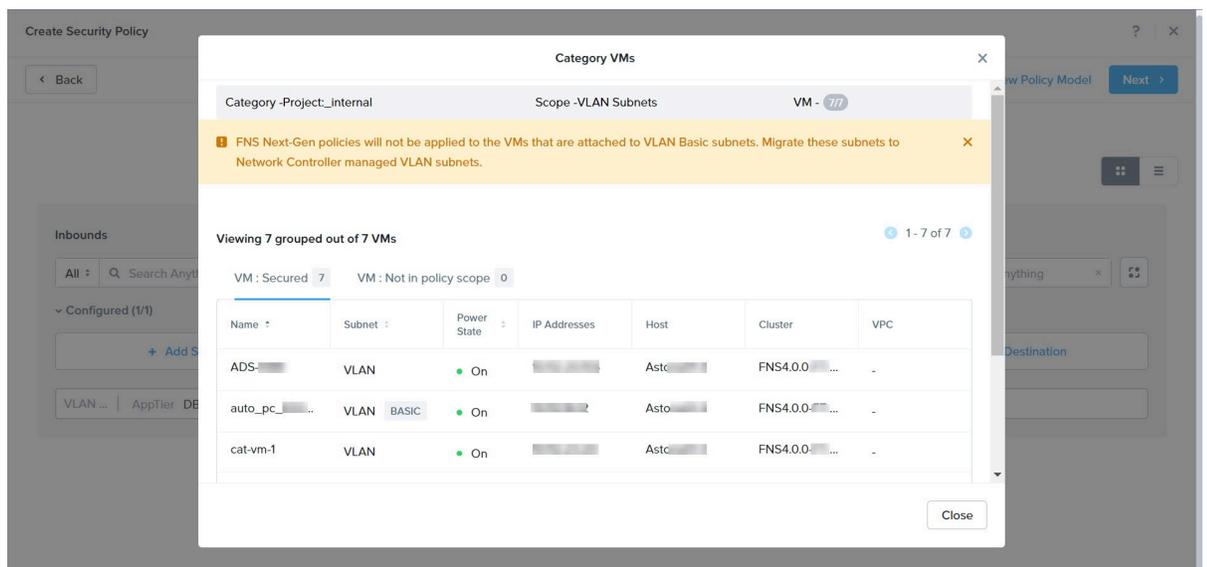


Figure 16: Category VMs

- c. Hover over the AD Group (VDI secured category) and click **Add VDI VM Filter** to filter the VMs.

-
-
- 1 Define Policy 2 Secure Application 3 Review
-

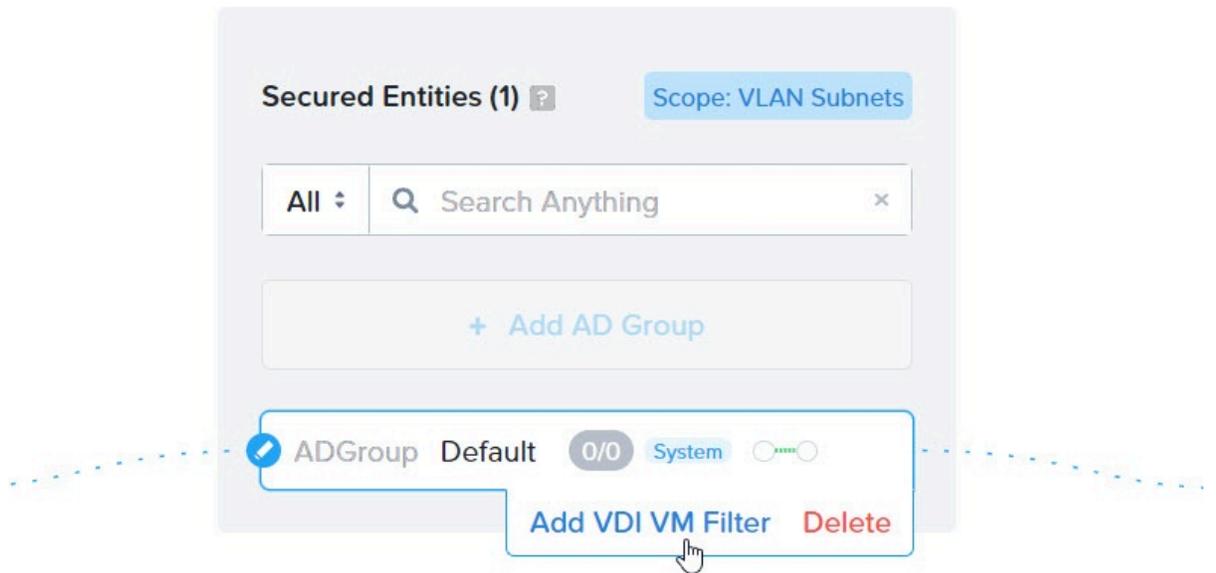


Figure 17: Filter the VMs

- d. In the VDI VM Filter page that appears, do anyone of the following.
- Select **Include VMs by name** and enter the matching criteria in the **VM Name Contains** field. Select the **Add these VMs to a default policy (assign category ADGroup:Default)** check-box to apply a default posture to the VMs. Optionally, select the **Keep the default category upon user logon** check-box to preserve the default category even after user logon.

- Select **Include all VMs** to include all the VMs in the AD group in the policy. Note that non-VDI VMs will also be included in the policy if Include all VMs option is selected.
- e. **Outbound Traffic:** Click **+ Add Destination** and enter the category or subnets that the VDI group can send the traffic to, as the destination.

The VMs that are part of Network Controller managed VLANs as well as VLAN Basic subnets can be added in the outbounds.

Note: If you have not used the default VDI option in *Step 2b*, ensure that you add all of your Active Directory domain controllers as part of this step, using either categories or subnets, for each ADGroup.

4. On the **Review** tab, select a policy mode from the following options:

- **Save:** Saves the policy without applying it. Saving a policy allows you to retain the policy in a draft stage without having the need to apply (enforce or monitor) it at the time of creation.
- **Apply (Monitor):** Monitors how the security policy works.
When a policy is in the Monitor state, the application continues to receive all traffic even from the disallowed source, but disallowed traffic is highlighted on the monitoring page. Traffic is not blocked until the policy is applied.
- **Apply (Enforce):** Applies the security policy on the application. The application receives and transmits traffic only from/to the allowed sources and destinations. The policy blocks the traffic from entities that are not defined as sources.

Note: VDI Policy does not support visualization.

The **Review** tab also displays the policy configuration summary.

5. Click **Confirm**.

Default VDI Policy

The Default VDI policy feature allows you to apply a default set of rules as defined by the desktop administrator for VDI VMs and users. There are two primary use cases for Default VDI Policy (**ADGroup:Default**).

- To ensure that a VDI VM is secure even before a user logs on to the VDI VM.
- To enable access to common network resources without the need to add the resources to every tier of a VDI policy.

You can define a default VDI policy at the time of creating a new VDI policy, or by updating any existing VDI policy. For more information, see *Step 2b* of the [VDI Policy Configuration](#) topic.

Configuring Active Directory Domain Services

Active Directory Domain Services configuration is used to import user groups for identity based security policies.

Before you begin

- Microsegmentation must be enabled to be able to use the ID Firewall feature. For more information, see [Enabling Microsegmentation](#) on page 9.
- You must allow WMI access from Prism Central to all the Active Directory Domain Controllers in your network firewall and Active Directory firewall.

- Active Directory Requirements:
 - Minimum supported domain functional level in Active Directory is Windows Server 2008 R2.
 - ID Firewall checks the membership of Security Groups only, Distribution Groups are not supported.
 - NTP must be configured on Active Directory and Prism Central.
 - DNS must be configured on Prism Central if you want to use host name for domain controllers.

About this task

To configure an Active Directory domain, do the following.

Procedure

1. Log on to the Prism Central web console.
2. Click the hamburger button on the left of the main menu and then select **Prism Central Settings** to display the Settings page.
3. Click **ID Based Security** from the Settings menu (on the left).
The **ID Based Security** page is displayed. This page allows you to **Add New Domain** or use an **Existing AD**.
4. If you select **Use Existing AD** in step 3, do the following in the indicated fields:
 - a. Click the **Manually Add Domain Controller** button, then click **+ Domain Controller**.
 - b. Enter the **IP Address** or **Host Name** of the domain controllers that you want to monitor for user logons events. You must add all the domain controllers associated with your Active Directory manually.
Click **+** and add each domain controller individually, then click the blue check mark icon to save.

Note: DNS must be configured on Prism Central for the host name option to work.

5. If you select **Add New Domain** in step 3, a set of fields is displayed. Do the following in the indicated fields:
 - 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 LDAP address of the directory, including the port number.
 - d. **Service Account Username:** Enter the service account user name in the `user_name@domain.com` format that you want Prism Central to use to detect logons and query user and group information from Active Directory.

Caution: Do not use the Domain Admin account as the service account considering the security best practices. Create a new domain user and grant it required permissions as described in [Configure Service Account for ID Firewall](#) on page 57.

A service account is a special user account that an application or service uses to interact with the Active Directory. Enter your Active Directory service account credentials in this (username) and the following (password) field.

Note: Ensure that you update the service account credentials here whenever the service account password changes or when a different service account is used.

- e. **Service Account Password:** Enter the service account password.
 - f. When all the fields are correct, click the **Save** button (lower right).
ID Firewall uses the service account for ID based security with additional requirements. For more information, see [Configure Service Account for ID Firewall](#) on page 57.
Once saved, the **Referenced AD Groups** section is displayed. You can add a new user group by clicking **+ Add User Group** and edit the auto-generated **Category Value**. After the active directory configuration is complete, you can create the VDI Policy. For more information, see [Creating a VDI Policy](#) on page 50
6. Click **Add Inclusion Criteria** under **Manage the VM Inclusion Criteria** to specify which VMs are assigned to AD Group categories upon user logon based on VM name.

Note: It is recommend that users add inclusion criteria if at all possible to prevent any unintended categorizations.

Note: The VMs with AppType category assigned cannot be categorized by ID Based Security.

Configure Service Account for ID Firewall

Active Directory service account in Prism Central is used for connectivity with the Active Directory domain services. ID Firewall also uses the same service account for ID based security.

To configure a service account for ID firewall, do the following.

1. Create a new user in the Active Directory.
2. Add the user to the **Distributed COM Users** group and the **Event Log Readers** domain groups.
3. Start the **dcomcnfg.exe** utility and go to **Component Services > Computers > My Computer > DCOM Config**.
4. Right-click on **Windows Management and Instrumentation** and select **Properties** from the menu.
5. Switch to **Security** tab, select **Customize** option in the **Access Permissions** section and then click **Edit**.
6. Add the user and grant **Local Access** and **Remote Access** permissions to the user. Click **OK** to confirm changes.
7. Run the `WMIMGMT.msc` command to start **Windows Management Instrumentation** snap-in.

8. Right-click on **WMI control (local)** and select **Properties** from the menu.
9. Switch to **Security** tab and expand **Root** tree.
10. Select **CIMV2** in the expanded tree and click **Security**.
11. Go to **Advanced > Add > Principal** and enter the user name.
12. Change scope by selecting **This namespace and subnamespaces** in the **Applies to** drop-down menu.
13. Click the check-box to grant the **Enable Account** and **Remote Enable** permissions. Click **OK** to confirm changes.
14. Restart the `winmgmt` service.

```
C:\> net stop winmgmt
C:\> net start winmgmt
```

Alternatively, reboot the domain controller.

15. Repeat step 3 to step 14 on every domain controller.

Modifying the VDI Policy

About this task

To modify the VDI policy, do the following:

Procedure

1. In the **Security Policies** page, select the policy to modify, click **Actions**, and then click **Update**.
2. Make the changes and then apply or save and monitor the policy.

The update options are the same as those for creating a policy. For information about the options, see [Creating a VDI Policy](#) on page 50.

Applying the VDI Policy

Applying the VDI policy enforces the policy on the specified categories (VDI AD groups), and any traffic between the categories is blocked.

About this task

To apply the VDI policy, do the following:

Procedure

1. In the **Security Policies** page, select the policy to apply, click **Actions**, and then click **Apply**.
2. Confirm by typing **Apply** in the dialog box, and then click **OK**.

Monitoring the VDI Policy

About this task

The VMs in VDI AD Groups in the VDI policy are allowed to communicate with each other when the policy is in the monitoring state. Traffic is blocked only during the time the policy is applied.

To monitor a security policy, do the following:

Note: VDI Policy does not support visualization.

Procedure

1. In the **Security Policies** page, select the policy to monitor, click **Actions**, and then click **Monitor**.
2. Confirm by typing **Monitor** in the dialog box, and then click **OK**.

Deleting the VDI Policy

About this task

To delete the VDI policy, do the following:

Procedure

1. In the **Security Policies** page, select the VDI policy.
2. Click **Delete** in the **Actions** menu.

ISOLATION ENVIRONMENT POLICY

An isolation policy identifies a minimum of two and a maximum of 32 entities by category and blocks communications between them. Each entity can have up to 10 categories.

For example, consider that you have an `application` category with values `app1` and `app2` and that you have associated some VMs with `application: app1` and some VMs with `application: app2`. Also, consider that these same VMs are distributed between two sites, and have accordingly been associated with values `site1` and `site2` in a category named `location` (`location: site1` and `location: site2`).

In this example, you might want to block communications between the VMs in the two locations. Additionally, you might want to restrict the scope of the policy to VMs in category `application: app1`. In other words, `app1` VMs in `site1` cannot communicate with `app1` VMs in `site2`. The following diagram illustrates the desired outcome. The red connectors illustrate blocked traffic. The green connectors illustrate allowed traffic.

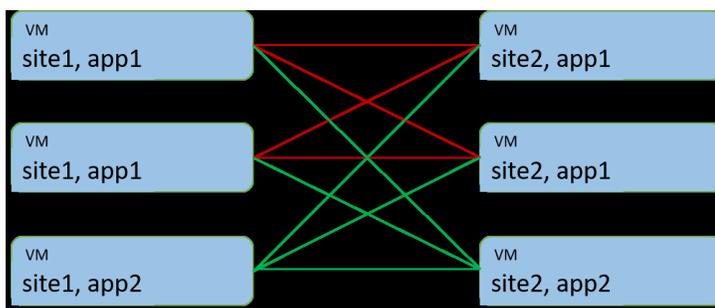


Figure 18: Applications Across Sites

You can configure an isolation policy for this by creating the following categories and isolation policy in Prism Central:

Table 5: Sample Configurations For Categories and the Isolation Policy

Entity	Values
Categories	<ul style="list-style-type: none"> • Name: <code>application</code> • Values: <code>app1</code> and <code>app2</code> • Name: <code>location</code> • Values: <code>site1</code> and <code>site2</code>
Isolation Policy	<ul style="list-style-type: none"> • Name: <code>eng_isolation_policy_across_sites</code> • Description: Isolate engineering VMs across sites • Isolate This Category: <code>location: site1</code> • From This Category: <code>location: site2</code> • Apply the isolation only within a subset of the data center: <code>application: app1</code>

Layer 2 Isolation

Isolation is done at Layer L2 which filters Layer 2 packets across all isolated entities. When an isolation policy is applied between two category-based VM groups, all ingress and egress traffic (broadcast, unknown-unicast, and multicast traffic) is dropped at the destination VM group.

IPv6 traffic between isolated VMs is blocked by default with the introduction of layer 2 isolation.

Note:

- If VMs are part of both isolation policy (in enforced mode) and quarantine policy (in enforced mode), the quarantine policy takes priority of processing over the isolation policy. For example, if VMs with category app1 are isolated from VMs with category app2 using an isolation policy, the traffic between these VM groups are not dropped if the VM groups are also part of a quarantine forensic policy that allows communication between these VMs. Therefore, quarantine policy takes precedence and this policy allows the traffic, even when the isolation policy is enforced.

Creating an Isolation Environment Policy

An isolation policy identifies a minimum of two and a maximum of 32 entities and blocks communications between the entities.

Before you begin

- Ensure you have created Network Controller managed VLANs.

About this task

To create an isolation policy, do the following:

Note:

You can create up to 100 isolation policies.

Procedure

1. In the **Security Policies** page, click **Create Security Policy**.
The **Create Security Policy** page is displayed.

2. On the **Define Policy** tab, do the following in the indicated fields and click **Next**:
 - a. **Policy Name**: Enter a name for the security policy.
 - b. **Purpose**: Describe the purpose of the security policy.
 - c. Select the **Isolation Environments (Isolation Policy)** policy type.
 - d. In the Scope of Secured Entities, do the following:
 - » Select **VLAN Subnets** to isolate VMs that span across one or more Network Controller managed VLANs.
 - » Select **Subnets Inside a VPC** to isolate VMs within a VPC scope. Select the VPC from the drop-down.
 - e. In the **Advanced Configuration** section:
 - Click **Enabled** option for the **Policy Hit Logs** to collect the flow policy hit logs on a syslog server.

Policy hit logs track network flows; whether the flows are allowed or denied by any of the policy. Policy hit logs are useful to determine if specific traffic is present on the network and how a security policy affects the traffic.

For multi-isolation policy, the policy hit logs shows **inbound** and **outbound** as directions instead of source and destination.

You can configure syslog monitoring for the policy hit logs for Flow. For more information, see *Configuring Syslog Monitoring* in the [Prism Central Admin Center Guide](#).

Note: Policy hit logs are not generated if both source and destination are in inbound or outbound category.

3. On the **Secure Application** page do the following:, the system allows you to isolate two sets of categories. Search for the category you want to isolate and press Enter. You can select multiple categories in one set. After adding the categories, click **Next**.
 - a. Click **Add Entity** to add more entities.
 - The system allows you to isolate up to 32 entities.
 - b. Search for the category to isolate and press **Enter**.
 You can add up to 10 categories per entity. If you select multiple categories, the system displays the number of common VMs across multiple categories. The isolation policy protects only the common VMs.
 - c. Click **Next**.

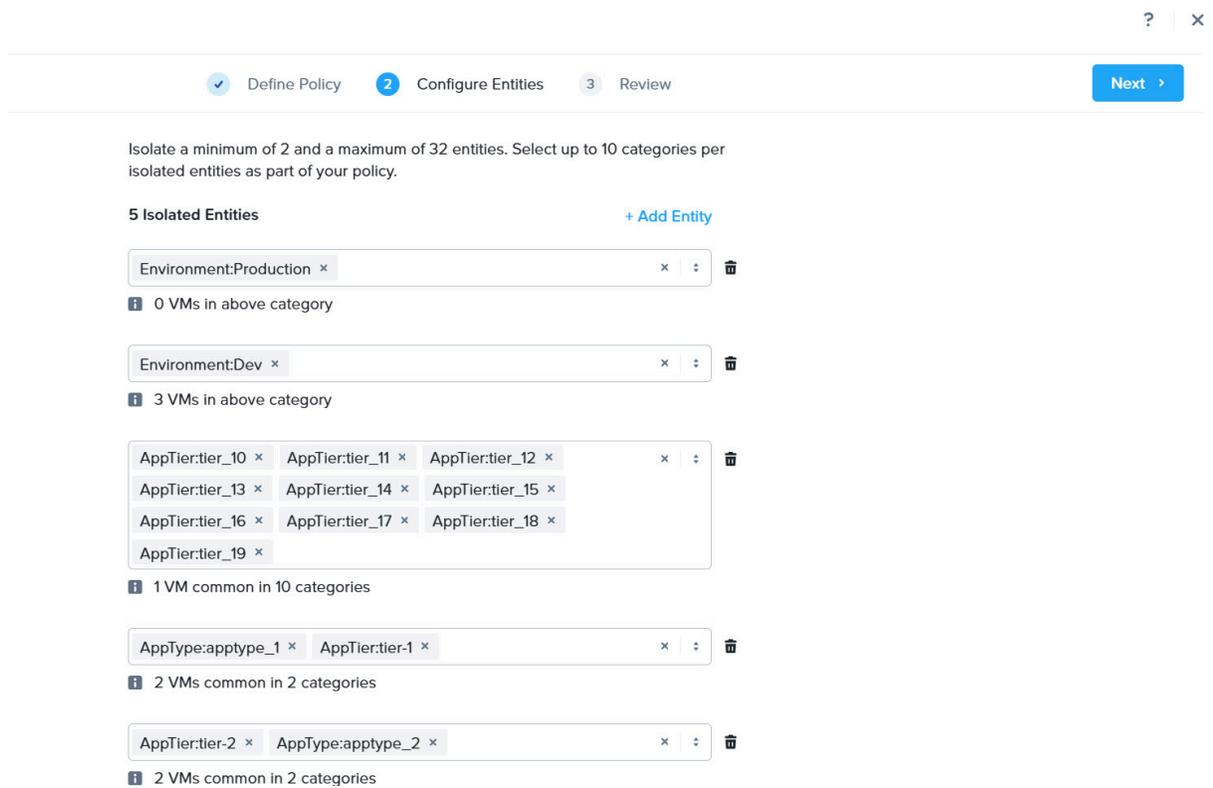


Figure 19: Configuring Entities for Isolation Policy

You can create your own category apart from using built-in categories such as AppType, and AppTier.

Note: You can create your own category apart from using built-in categories such as AppType, and AppTier. For more information, see [Creating a Category](#)

4. On the **Review** tab, select a policy mode from the following options:

- **Save:** Saves the policy without applying it. Saving a policy allows you to retain the policy in a draft stage without having the need to apply (enforce or monitor) it at the time of creation.
- **Apply (Monitor):** Monitors how the security policy works.

When a policy is in the Monitor state, the entities continues to receive all traffic even from the isolated source, but traffic to be isolated is highlighted on the visualization. Traffic is not blocked until the policy is applied.

The traffic flow visualization is available in the monitor mode.

- **Apply (Enforce):** Applies the isolation policy on the entities. The traffic is blocked between the entities defined in the isolation policy.

Between the isolated source and destination entities, traffic visualization is not shown in the enforce mode. However, for the given source and destination if ARP is resolved for the destination in the source before creating the policy or enforcing the policy then visualization is shown in the enforce mode.

5. Click **Confirm**.

Modifying an Isolation Environment Policy

About this task

To modify an isolation environment, do the following:

Procedure

1. In the **Security Policies** page, select the isolation policy.
2. From the **Actions** drop-down click **Update**.
3. Make the changes, select the policy mode and then click **Confirm**.

Applying an Isolation Environment Policy

Applying an isolation environment policy enforces the policy on the specified categories, and any traffic between the categories is blocked.

About this task

Note: Changing the state of an isolation environment policy affects the functioning of any conflicting application security policies. For more information, see [Priorities Between Policies](#) on page 21.

To apply an isolation environment policy, do the following:

Procedure

1. In the **Security Policies** dashboard, select the policy that you want to apply.

2. From the **Actions** drop-down select any of the actions:

- **Apply (Monitor)**: Select this mode when you want to monitor how the security policy works.

When a policy is in the Monitor state, the application continues to receive all traffic even from the disallowed source, but disallowed traffic is highlighted on the monitoring page. Traffic is not blocked until the policy is applied.

- **Apply (Enforce)**: Select this mode when you want to apply the security policy on the application. When the policy is in the Enforce state, the application receives and transmits traffic only from/to the allowed sources and destinations. The policy blocks the traffic from/to entities that are not defined as sources/destinations.

3. Confirm by typing **ENFORCE** or **MONITOR** as applicable in the dialog box and then click **Confirm**.

Monitoring an Isolation Environment Policy (Visualizing Network Flows)

About this task

The VMs in the categories in an isolation environment policy are allowed to communicate with each other when the policy is in the monitoring state. When a policy is in monitor mode, the application continues to receive all traffic, even from the disallowed source. Disallowed traffic is highlighted on the monitoring page. Traffic is not blocked until the policy is applied.

Note: Changing the state of an isolation environment policy affects the functioning of any conflicting application security policies. For more information, see [Priorities Between Policies](#) on page 21.

To monitor a security policy, do the following:

Procedure

1. In the **Security Policies** page, select the policy that you want to monitor.
2. From the **Actions** drop-down click **Apply (Monitor)**.
3. Confirm by typing **MONITOR** in the dialog box, and then click **Confirm**.

The following figure shows a multi-isolation policy in monitor mode. In the figure, a circle represents an isolated entity. The categories within the entity are mentioned inside the circle. The discovered traffic is highlighted in the dotted lines. The arrows in the dotted lines for the discovered traffic indicate the direction of the traffic flow. A single arrow represents unidirectional traffic flow, whereas two arrows represent bidirectional traffic flow. The traffic between the **AppTier tier-2** to **Environment Dev** is unidirectional. The traffic between the **AppTier tier_10** to **AppTier tier-1** is bidirectional. For an entity with multiple categories, the system displays the number of common VMs in a small circle on top of the entity circle. The isolation policy protects only the common VMs.

For example, in the figure, the **AppTier tier_10** entity has 10 categories and only one VM is common across these 10 categories.



Figure 20: Multi-Isolation Policy Visual View

- (Optional) To view the details of any entity, click the entity circle.

A side window appears. The side window displays the list of categories, the total number of common (protected) VMs, and the discovered traffic.

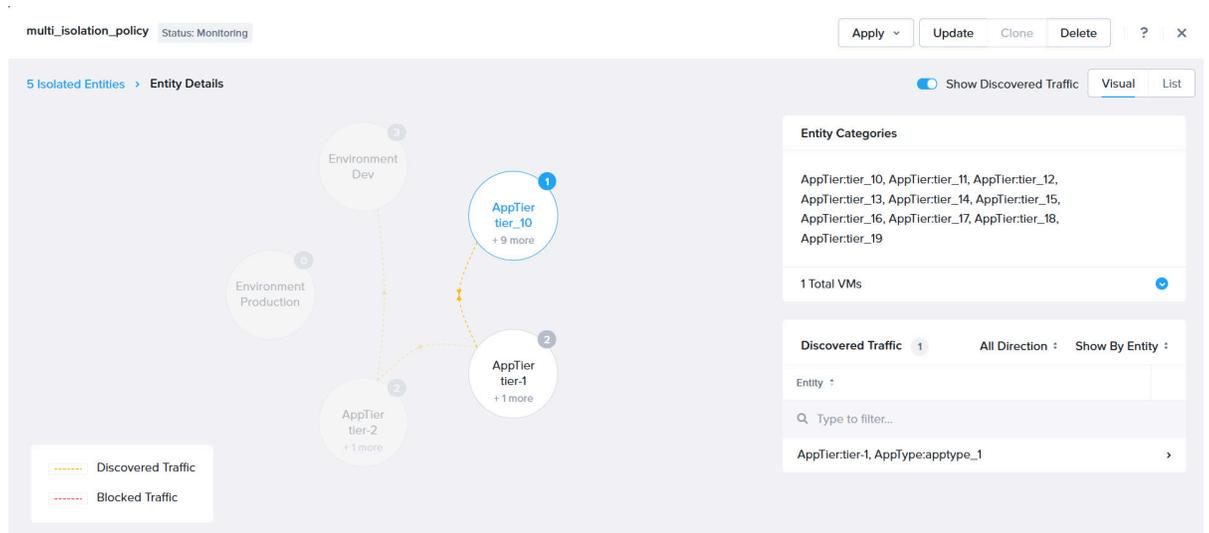


Figure 21: Multi-Isolation Policy - Category Details

5. (Optional) To view the details of the discovered traffic, click the dotted lines.

A side window appears. The following figure shows the discovered traffic details between **AppTier tier_10** and **AppTier tier-1** entity.

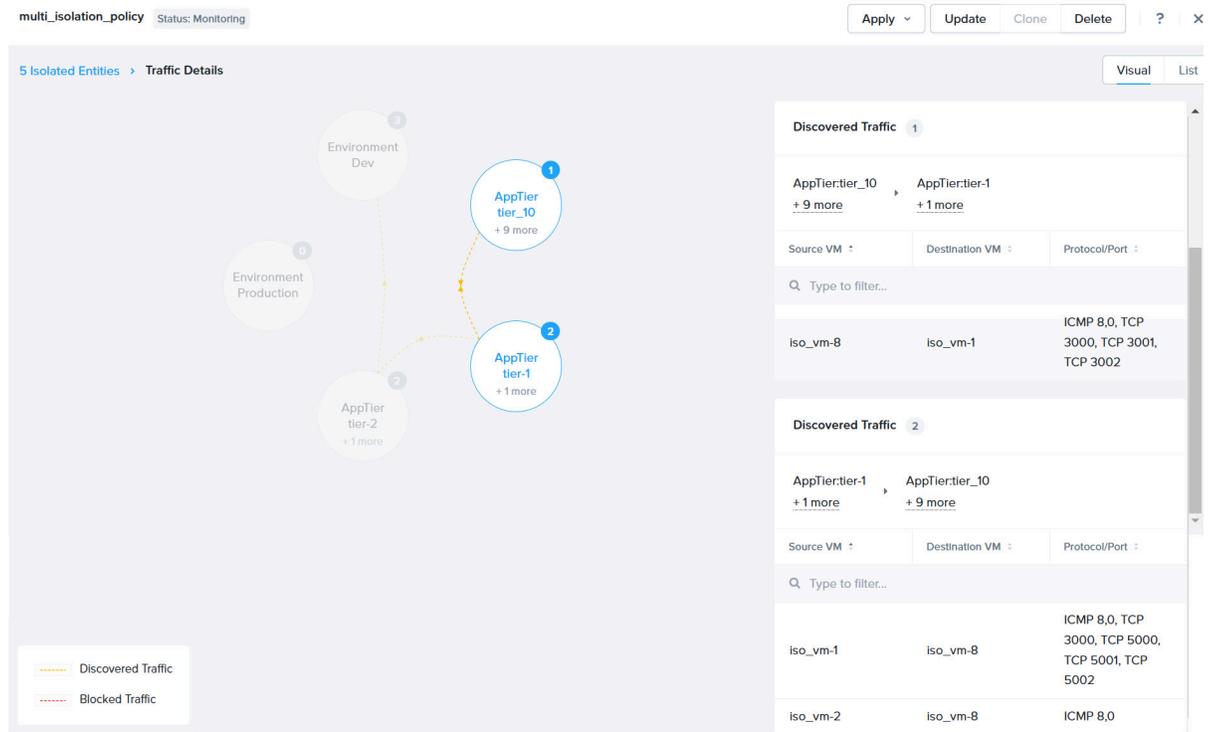


Figure 22: Multi-Isolation Policy - Discovered Traffic Details

Deleting an Isolation Environment Policy

About this task

To delete an isolation environment policy, do the following:

Procedure

1. In the **Security Policies** page, select the policy.
You can select multiple policies to delete them at once.
2. From the **Actions** drop-down click **Delete**.
3. Confirm by typing **DELETE** in the dialog box, and then click **Confirm**.

QUARANTINE POLICY CONFIGURATION

Prism Central has system-defined quarantine policies that enable you to perform the following tasks:

- Completely isolate an infected VM that must not have any traffic associated with it.
- Isolate an infected VM but specify a set of forensic tools that can communicate with the VM.

The system-defined quarantine policies are created for All VLANs and VPCs.

The system-defined VLAN specific quarantine policies are:

- Quarantine Forensic Policy - VLAN Subnets (Scope)
- Quarantine Strict Policy - VLAN Subnets (Scope)

The system-defined VPC specific quarantine policies are:

- Quarantine Forensic Policy - VPC (Scope)
- Quarantine Strict Policy - VPC (Scope)

For these use cases, Prism Central includes built-in categories that are included in the system-defined quarantine policy.

Note: You cannot create or delete a quarantine policy. However, you can modify existing (system-defined) quarantine forensic policy.

Configuring the Quarantine Policy

In the built-in quarantine policy, you specify categories that can communicate with VMs that have been added to the **Quarantine Forensics Policy** or **Quarantine Strict Policy** category.

About this task

To configure the quarantine policy, do the following:

Procedure

1. Log on to Prism Central.
2. Select **Infrastructure** application from the [Application Switcher Function](#), and navigate to **Network & Security > Security Policies**
The **Security Policies** page appears.
3. In the **Security Policies** page, select any of the following built-in quarantine policy and then click **Update** in the **Actions** menu.
 - » Quarantine Forensic Policy for VLAN Subnets (Scope)
 - » Quarantine Forensic Policy for VPC (Scope)

4. (Optional) In the **Advanced Configuration** under the **Define Policy** tab, do the following.
 - a. Select the **Allow** radio button to allow **IPv6 traffic**. The policy rules apply to IPv4 traffic only and all IPv6 traffic is blocked by default. You can configure the allow option for both **Forensic** and **Strict** modes.
 - b. Optionally, click the toggle button against **Policy Hit Logs** to log traffic flow hits on the policy rules.

You can configure syslog monitoring for the policy hit logs for Flow. For more information, see *Configuring Syslog Monitoring* in the *Prism Central Admin Center Guide*.

You can enable the policy hit log option for both **Forensic** and **Strict** modes.

Note: Policy hit logs are not generated if both source and destination are in inbound or outbound category.

5. On the **Secure Application** tab, do the following, and then click **Next**:
 - a. To specify the categories that contain forensic tools, on the **Inbound** and **Outbound** sides of the policy diagram, do the following:
 - Click **Add Source** or **Add Destination**, and then do the following:
 1. Select one of the following options from the drop-down list:
 - **Category**: Allows traffic to or from the specified category.
 - **Subnet/IP**: Allows traffic to or from the specified subnet.
 - **Addresses**: Allows traffic only if the traffic originates from the entities that are in the selected address.
 - **Allow All**: Allows traffic associated with all sources or destinations.
 2. Enter the value in the text box, and then click **Add**.

When entering the name of a category, a list of matching names is displayed, and you can select the name you want to specify. The subnet mask must be specified in the CIDR format.

When entering the address, a list of available address group names appears. Select the address group or you can create a new address group.
 3. To add another category, subnet or address, click **Add Source** or **Add Destination**. Add as many category, subnet or address as you want to allow.
 - b. To specify the protocols and ports over which the forensic tools can communicate with the VMs in the forensic category, do the following:
 - 1. On the **Inbound** and **Outbound** sides of the policy diagram, click a category or subnet (if you have configured a allowed list) or **All Sources** (if you have chosen to allow all sources) for which you want to create a rule.
 - 2. Click the plus icon that appears on the **Quarantine: Forensic** category. The **Create Inbound Rule** or **Create Outbound Rule** dialog box

3. Enter a description for the rule.

Note: The policy rule description is captured in the policy hitlog data.

- Policy hitlog must be enabled
- Rule description is added to the hitlog only for allowed traffic

4. In **Service Details**, click **Allow all traffic** to allow all types of traffic or click **Select a service** to choose any default or custom service.
5. Click **Save**.

After you configure a rule, a line appears between the two endpoints of the rule. Point to the line to show the list of ports that the rule allows.

- c. The **Quarantine** tab displays the scope of the policy at the top-right corner.

For example, the current policy belongs to the VPC **Scope: test**. To know the details on the protected VMs and the VMs that are not in the policy scope, click the fraction number within the category.

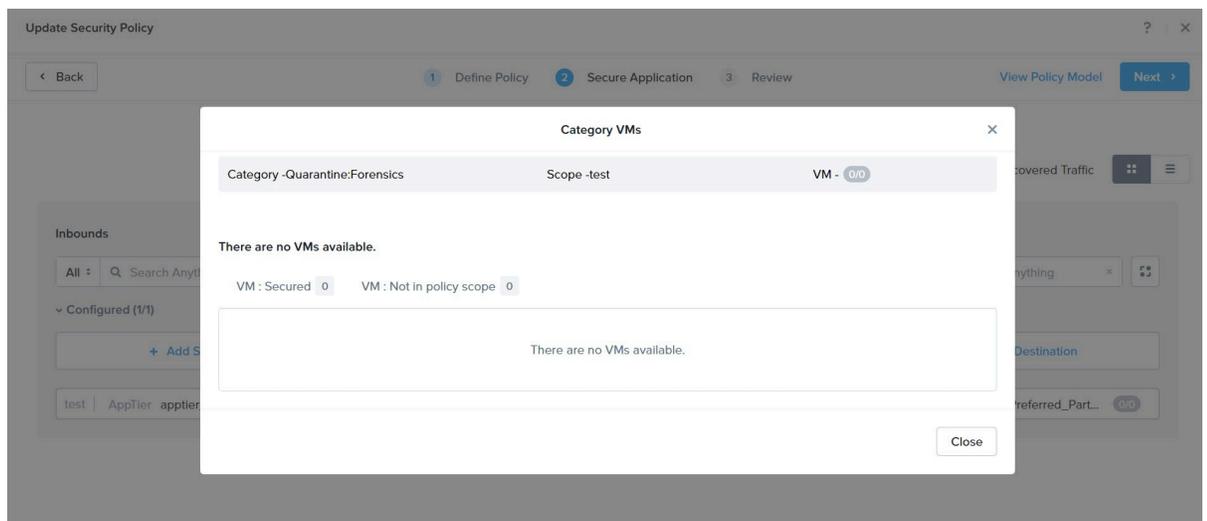


Figure 23: Category VMs

6. On the **Review** tab, do one of the following:
 - » Click **Save** to save the policy without applying it. Saving a policy allows you to retain the policy in a draft stage without having the need to apply (enforce or monitor) it at the time of creation.
 - » Click **Apply (Monitor)** to save the configuration and place the quarantine policy in the monitoring mode.
 - » Click **Apply (Enforce)** to apply the quarantine policy.

You can switch between the monitoring and enforce states by selecting quarantine policy on the **Security Policies** page and clicking the appropriate option in the **Actions** menu.

7. Click **Confirm**.

Quarantining a VM

You quarantine a VM by adding the VM to a quarantine category.

About this task

To add an infected VM to a quarantine category, do the following:

Procedure

1. In the VMs dashboard **List** tab, select the infected VM, click **Actions**, and then click **Quarantine VMs**.
2. Under Quarantine Method, click one of the following options:
 - » **Strict**. Isolates the VM from all traffic. No exceptions can be made for forensics.
 - » **Forensic**. Isolates the VM from all traffic except traffic from categories specified in the built-in quarantine policy. The allowed categories contain forensic tools that enable you to perform forensics on the VM.For VMs added to the strict quarantine, a red icon is displayed in the name column.
3. Click **Quarantine**.

Removing a VM from the Quarantine

About this task

To remove a VM from the quarantine, do the following:

Procedure

1. In the VMs dashboard **List** tab, select the VM that you want to remove from the quarantine, click **Actions**, and then click **Unquarantine VMs**.
You can select multiple VMs and remove them from the quarantine in a single step.
2. In the **Unquarantine VMs** dialog box, click **Unquarantine**.

EXPORTING AND IMPORTING SECURITY POLICIES

Prism Central allows you to export and import security policies for the following security administration aspects.

- Have a snapshot of a working security configuration so that system can be restored to the desired state when needed.
- Ability to apply security policies as templates. This scenario is useful in ROBO environments (disaster recovery deployments) where the datacenters are being managed by multiple Prism Central instances.

Exporting Security Policies

To export security policy, do the following in the **Security Policies** page.

- Select **Actions > Export All Policies**. The security policies binary file is downloaded.

Importing Security Policies

To import any previously exported security policies binary file, do the following in the **Security Policies** page.

1. Select **Actions > Import Policies** , then click **Browse** to select the binary file.

The Import Summary is displayed. This summary shows the details of the import operation. The details includes Overridden, Deleted, and New Addition count. You can click the number and the page displays the names of overridden policies, newly created policies, and deleted policies in respective columns. Also, the New Addition column shows categories, address groups, policies and service groups information.

2. Click **Import**. The security policies are imported.

Note: Existing policies are overridden with new policies. Policies that are not part of this import are deleted.

DISABLING MICROSEGMENTATION

Prism Central web console provides you the ability to disable the microsegmentation feature.

About this task

To disable microsegmentation, do the following:

Procedure

1. Log on to the Prism Central web console.
2. Click the **Prism Central Settings** gear icon in the main menu and then select **Microsegmentation**.
3. Click **Disable Microsegmentation**.

A confirmation message is displayed.

4. Click **Confirm** to disable the microsegmentation feature.

After disabling microsegmentation all security policies applied in monitor or enforce mode are deactivated. Therefore, disabling protection for VMs defined in these policies.

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.