v1.3 | April 2022 | TN-2094

TECH NOTE

Flow Network Security

NUTANIX
YOUR ENTERPRISE CLOUD

Copyright

Contents

1. Executive Summary

Flow Network Security offers policy-based network security tightly integrated into Nutanix AHV and Prism Central, and provides rich visualization, automation, and security for VMs running on AHV. Microsegmentation is a component of Flow Network Security that simplifies policy management. Using multiple Prism Central categories (logical groups), you can create a powerful distributed firewall that gives administrators an application-centric policy management tool for securing VM traffic.

Categories allow you to flexibly group VMs based on attributes to define security policies. When you use logical categories, you no longer need to base policy definitions on network addresses or manually update policies to handle network changes. Your security policy can automatically apply to a VM independent of its network configuration. With Flow Network Security, administrators can visualize traffic between groups of VMs to create plain-language policies based on application behavior, rather than defining allowed traffic in the language of IP addresses and subnets.

2. Introduction

Audience

This tech note is part of the Nutanix Solutions Library. We wrote it for architects and administrators responsible for VM networking and security. Readers of this document should already be familiar with Nutanix AHV and Prism Central.

Purpose

In this document, we cover the following topics:

- Flow Network Security
- Flow Network Security microsegmentation
- Flow Network Security visualization
- Creating and managing categories
- Creating and managing security policies
- Identity-based security with Active Directory
- Logging and auditing

Document Version History

Version Number	Published	Notes
1.0	April 2018	Original publication.
1.1	May 2018	Added Service Chain KB article.
1.2	June 2020	Updated for 5.17.

Version Number	Published	Notes
1.3	April 2022	Updated for PNP 2.0. Updated Flow Network Security for 6.1 and naming.

3. Flow Network Security

Flow Network Security is an application-centric microsegmentation solution that provides protection of east-west traffic to your Nutanix environments based on stateful, distributed firewall policies.

Flow Network Security (previously Nutanix Flow) is one component of three major related product areas:

1. Flow Virtual Networking: Provides self-service network provisioning and overlapping IP addresses for true multitenant networking.
2. Security Central: Provides security planning, threat detection, and compliance auditing for both on-premises and cloud environments.
3. Flow Network Security: Provides category- and policy-based microsegmentation.

Flow Network Security, which is fully integrated into AHV, provides application-centric policies that enable complete visibility and traffic control. This policy model allows administrators to implement fine-grained rules regarding traffic sources and destinations, or microsegmentation. These same policies make it possible to visualize traffic flowing within and between application VMs. This granular level of control is an important part of a defense-in-depth strategy against modern datacenter threats.

Architecture

The Nutanix management plane, Prism Central, provides policy administration for Flow Network Security. Prism Central and the registered Nutanix clusters combine to form the control plane for Flow Network Security. The data plane is the Nutanix AHV host Open vSwitch, which processes VM network traffic.

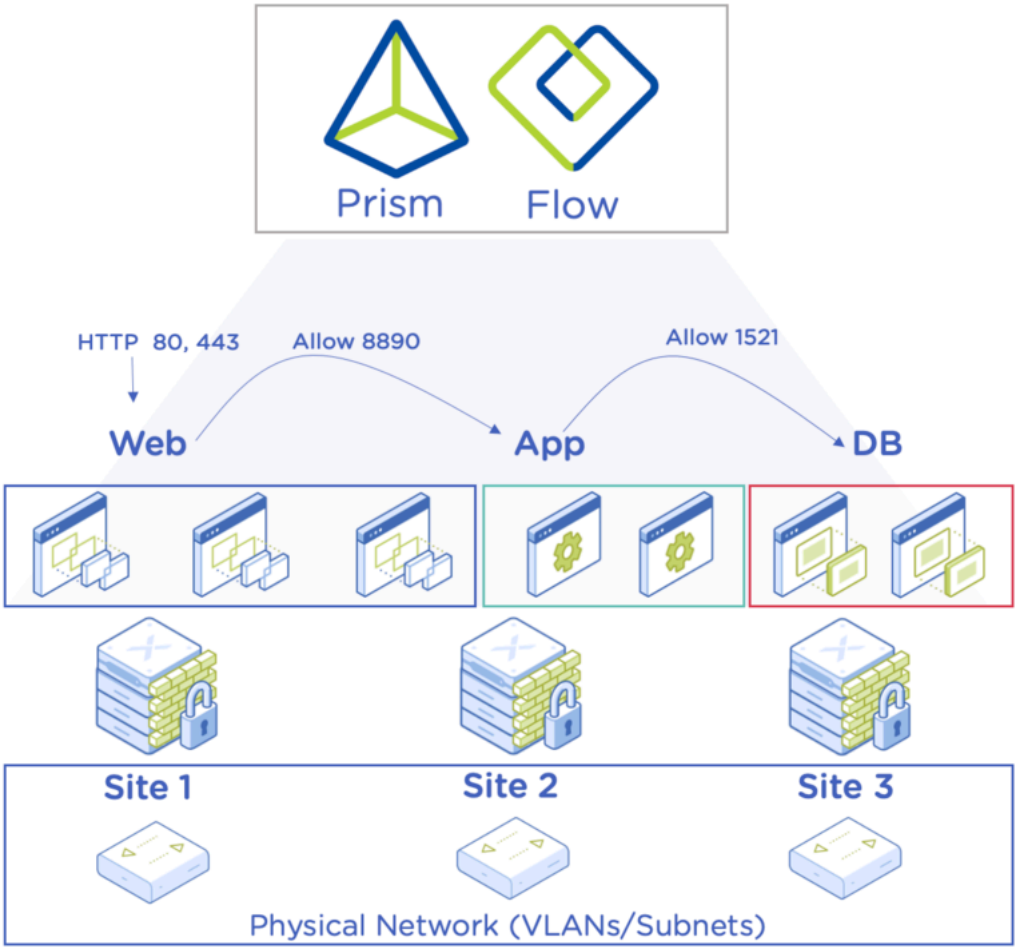

Figure 1: Flow Network Security Architecture

Enabling Microsegmentation and Visualization

Flow Network Security is disabled by default, but it's easy to enable: in Prism Central, click the question mark, select New in Prism Central, then click Microsegmentation. Refer to the Security Policies section of the Flow Microsegmentation Guide before you enable Flow Network Security. Nutanix offers a free 60-day trial period to get started without any licenses required.

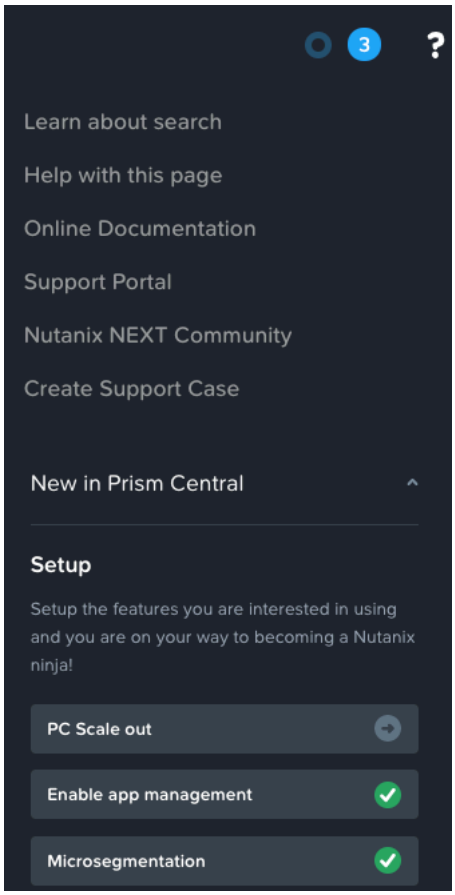

Figure 2: Enable Microsegmentation

Categories

Categories can group entities like VMs together logically. Each category consists of two parts: a key and a value. For example, the Environment key may have several values, such as Production, Development, Staging, or Test. When assigned to a VM, each category consists of a text-based key-value pair. For example, you assign Environment: Production—a key of Environment and a value of Production—to a VM in the production environment.

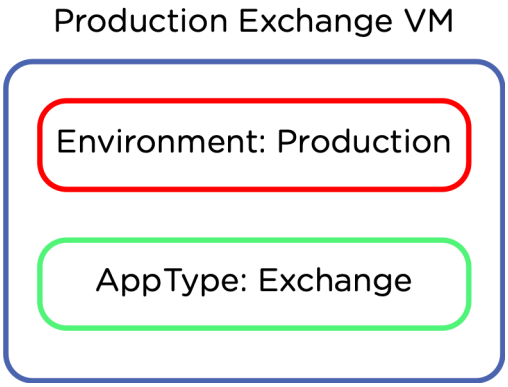

Figure 3: Categories Assigned to a VM

Categories are extremely flexible, and you can create new key-value pairs to group VMs based on application requirements. Categories are the building blocks of security policies, so think carefully about what labels you need to segment network traffic along existing logical boundaries. For example, production may be isolated from development, or a proxy application may require special access from the internet.

There are a number of system categories built into Prism Central that have a special meaning when used in Flow Network Security. Use these system categories to take advantage of security policies, starting with the existing categories and adding your own if needed. The following sections describe several important system categories that Flow Network Security requires for building policies.

AppType System Category

The system category AppType defines a group of VMs that belong to the same application. You can use AppType: Exchange, for example, to define all the entities that belong to the Exchange application. Create new values inside the AppType category to match the applications running in your AHV environment instead of creating new top-level categories.

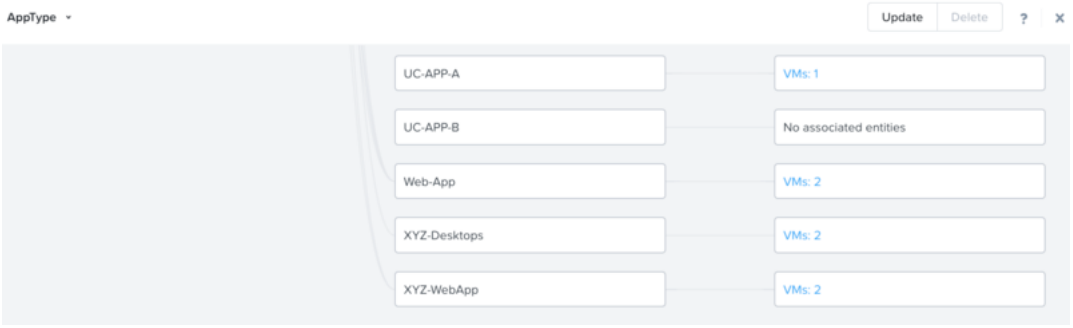

Figure 4: AppType Category List

When creating Application Security Policies, Flow Network Security uses AppType categories to manage traffic to and from an application, so define your applications with the appropriate AppType values. Select AppType, then select Update to add your own custom application names.

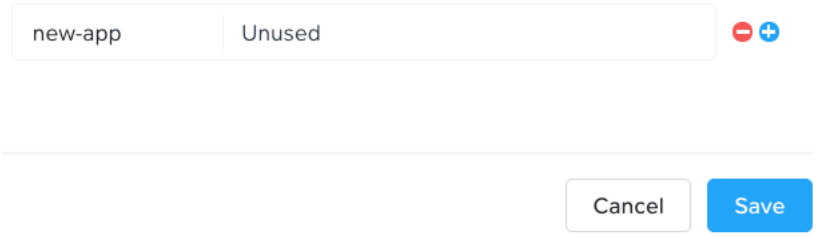

Figure 5: AppType Custom Application Name

AppTier System Category

The category AppTier defines a set of entities that serve the same function in an application. For example, AppTier: Exchange_Mbox and AppTier: Exchange_Edge are two tiers of VMs in the Exchange application that each serve a different function.

The combination of AppType and AppTier is used to uniquely identify an application. In the previous example, an Exchange mailbox VM is assigned both AppType: Exchange and AppTier: Exchange_Mbox.

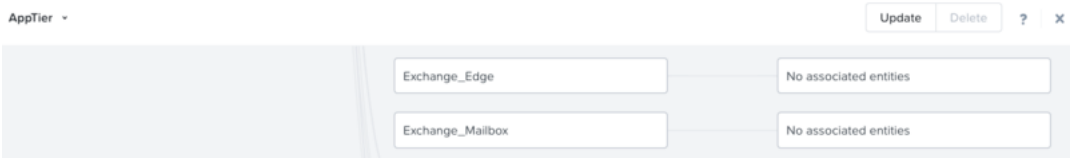

Figure 6: AppTier Categories

If the VM is part of the production environment, it may also have an environment category assigned, as shown here.

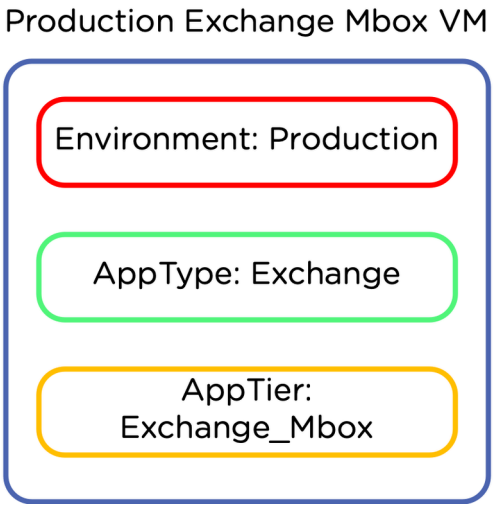

Figure 7: Production Exchange Mailbox VM

When creating application security policies, Flow Network Security uses the AppTier values to secure communication to, from, and between application tiers. Create custom tiers for your application by selecting AppTier, then Update from the Administration, Categories menu in Prism Central.

Creating Custom Categories

Nutanix recommends updating the existing system categories where possible and expanding these with values that suit your needs. For example, expand the AppType and AppTier categories by adding your own applications and tiers as values. AppType and AppTier are the only categories you can select as the target of a security policy.

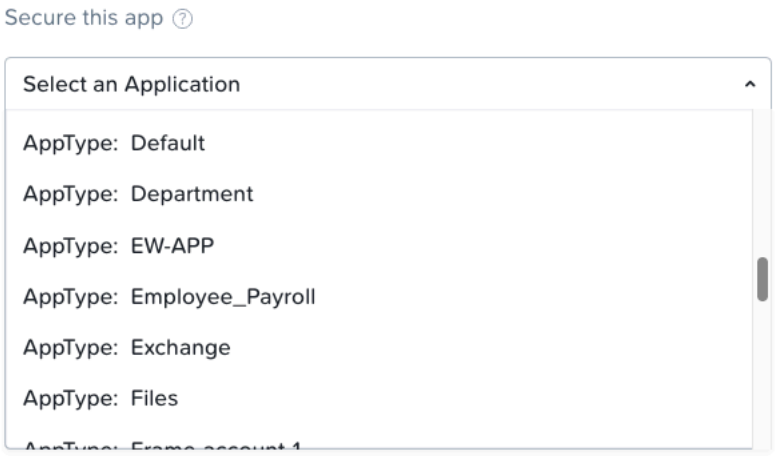

Figure 8: Application Policy Must Use AppType

You can create new category keys as needed for inbound and outbound endpoints and isolation policies, but Nutanix recommends keeping things as simple as possible while still achieving the desired result. Creating the smallest possible number of categories helps ensure that policy creation is quick and easy for the administrator.

The creation of Exchange policies for branch offices offers an example of the tradeoff between the number of categories and the number of policies. Consider the categories and policies required based on the two following approaches: one category per site type and one category per site. The result of each policy is the same, but the effort required is greater in the second policy.

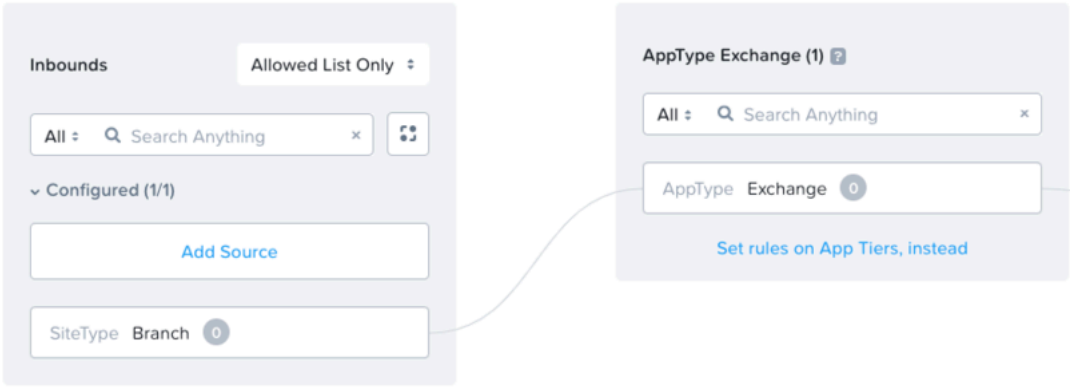

Figure 9: Approach 1: One Category per Site Type

The previous policy creates a single category called SiteType: Branch and applies this category to all VMs at all branches. That means the security policy to connect to Exchange requires just a single entry for all branch VMs. In the next policy, a category is created for Site: Branch-001 through Site: Branch-004 and now four inbound rules are required in this policy.

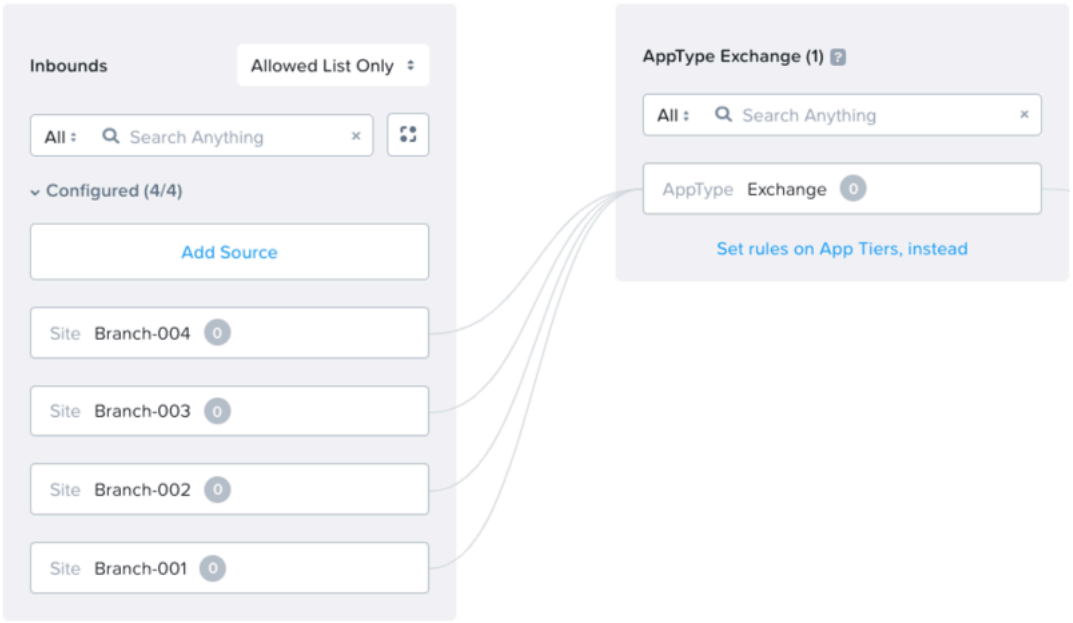

Figure 10: Approach 2: One Category per Site

With the first approach, it's easy to create a security policy as long as all branches have the same relationship to the HQ application. The second approach gives more control per site, but we must define the relationship between each site and HQ. If we have many sites, creating this definition for every site could be time consuming.

Quarantine System Category

By default, quarantine categories isolate a VM from all other network devices. In addition to the standard quarantine category—default, which denies all traffic—a forensic category allows quarantined VMs access to and from a policy-defined set of destinations and sources to assist in analyzing the quarantined VM. You can't modify the Quarantine: Default and Quarantine: Forensic categories or

manually assign them to a VM. Flow Network Security uses these quarantine categories when you select the quarantine action for a VM.

Assigning Categories to VMs

To assign categories to VMs, navigate to the Prism Central menu, then click Compute & Storage, then VMs. Select the VMs you want, then click Manage Categories in the Actions dropdown menu.

You can assign more than one category to each VM, so it's possible to assign multiple category types, such as Site and SiteType described in the Creating Custom Categories section. We could allow connectivity between branches based on one attribute (such as Site: Branch-001) and connectivity from branches to HQ based on another attribute (such as SiteType: Branch).

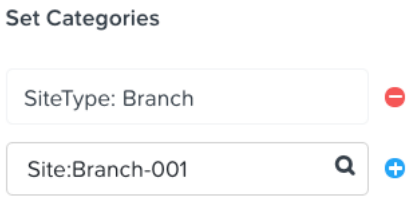

Figure 11: Site and SiteType Categories on a VM

With the Site category, you can define complex relationships between sites by the individual site number. For simpler policy creation, you can use the SiteType category to define relationships between all sites and HQ. This flexibility demonstrates the power of categories and shows that thinking about the relationships in an application can help you create more useful categories.

Security Policies

Security policies managed in Prism Central use categories to define the network traffic allowed or denied between VMs. Different security policy types describe the relationships between categories and hence between applications. Flow Network Security offers the following policy types: quarantine, isolation, application, and VDI.

The quarantine policy has two available tiers: Strict and Forensic. Strict quarantine blocks all traffic to and from VMs. Forensic quarantine allows specific

predefined sources and destinations for quarantined VMs. Use Strict quarantine to completely block all inbound and outbound traffic to a VM. Use Forensic quarantine to allow a defined set of administrators to troubleshoot and analyze a VM.

Isolation policies define specific categories of VMs that shouldn't be allowed to communicate with each other. For example, you could create an isolation policy to ensure that the development environment can never communicate with the production environment.

Next, application security policies allow you to define a multitiered application and control traffic to and from as well as within the application tiers. For example, you could create an application policy to define a mailbox and an edge tier inside the Exchange application. Furthermore, you can then set this policy to allow sources such as Marketing and Sales access to the mailbox tier, with both Exchange tiers allowed access to AD and Windows Update servers.

Finally, a VDI policy uses Microsoft Active Directory integration to categorize VMs based on the AD group of the signed-in user. The VDI policy looks just like an application policy, with the protected desktops at the center and sources and destinations on the side. Each AD group becomes a tier of the policy. The big difference between VDI policies and application policies is that when a user belongs to more than one AD group, the tiers of a VDI policy matching those groups are combined as a union. The corresponding combined allowlist is used to determine allowed traffic.

Nutanix Flow Network Security evaluates security policies in the order shown in the next figure to build traffic rules. If traffic encounters a matching rule, Flow Network Security applies the action in that policy and further policy processing stops. If no quarantine or isolation rule matches the traffic, then the application rule evaluates it for a match. If no application rules match the traffic, it's evaluated for a VDI policy match. If the traffic source or destination doesn't match any rules, Flow Network Security allows the traffic.

Traffic is matched both as it leaves VMs and as it is sent to VMs. Matches are made using both source and destination IP addresses as well as protocol and destination port. Even VMs on the same AHV host that send traffic to each other are protected by these rules.

Forensic quarantine policies, application policies, and VDI policies are all matched based on the combination of a source or destination allowlist and target group, which is the item at the center of the security policy. Strict quarantine policies contain only a target group with no allowlist. Isolation policies contain only a source and destination category and are matched by the source and destination IP addresses of traffic.

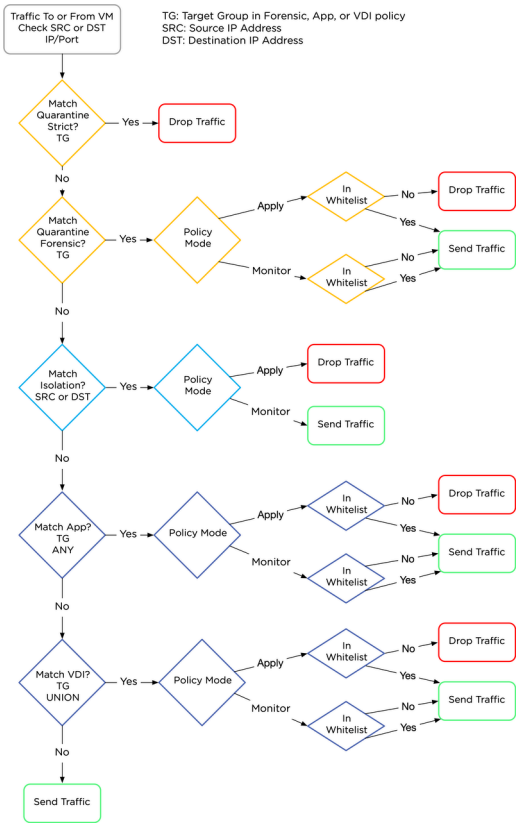

Figure 12: Policy Evaluation Order

Combining Policies

The evaluation order allows you to combine multiple policies and policy types to build a complete security solution for a set of applications. First, use quarantine policies to temporarily isolate a specific VM from all other VMs. Next, use isolation policies to define which groups of VMs must never be allowed to talk with other groups. Use application policies to control traffic allowed to, from, and within applications.

Application policies can allow traffic between two or more applications as well but remember to ensure that the policy of each application allows this traffic exactly. One application policy's outbound connection may be another policy's inbound connection. When this is the case, the outbound of the originating policy must exactly match the inbound of the destination policy as shown in the following example. Mismatch between the ports in these policies may lead to undetermined visualization behavior.

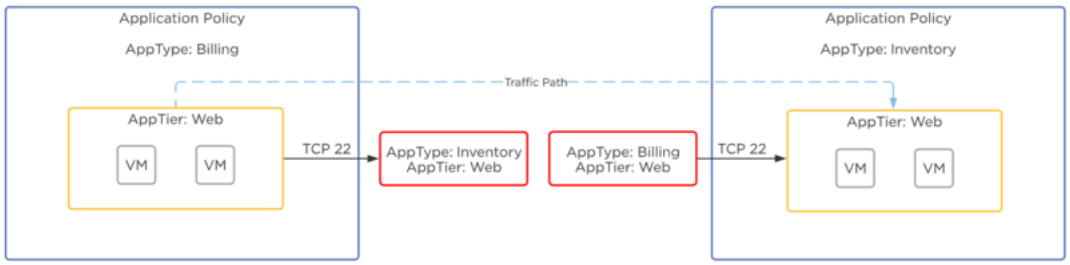

Figure 13: Combining Application Policies

Finally, VDI policies control traffic to and from VMs based on the Active Directory user group of the signed-in user.

Policy Mode

The two selectable modes for security policies are Enforce and Monitor. A policy in Monitor mode allows traffic, even if the policy doesn't specifically define that traffic as allowed. In contrast, Enforce mode only allows specifically defined traffic.

Monitor mode is the default state for a newly created policy. Use Monitor mode in combination with flow visualization to track flows on a newly created policy and ensure that it contains the expected traffic. Once you have confirmed policy accuracy, you can move policies from Monitor to Enforce mode. Use flow visualization in Enforce mode to see blocked traffic.

All policies except strict quarantine have a monitor mode. Pay careful attention to the policy evaluation order, since traffic that matches a policy in monitor mode is allowed and all further policy processing stops. For example, a matching Isolation policy in Monitor mode allows traffic that an Application policy further down the processing order may have otherwise blocked.

Using Quarantine Policies

You have the option to place a VM into a quarantine category in the VM action menu.

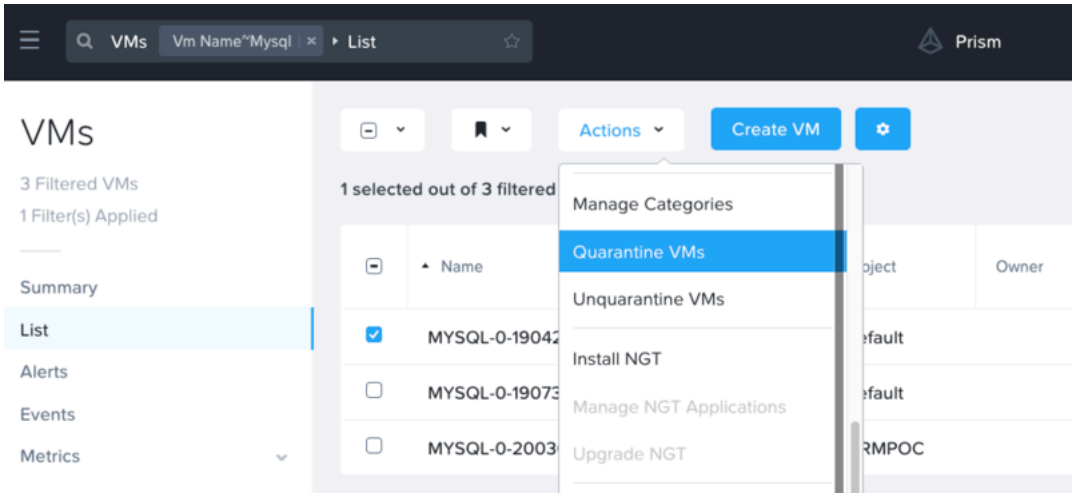

Figure 14: Quarantine VM Action

You can place VMs into either Strict or Forensic quarantine.

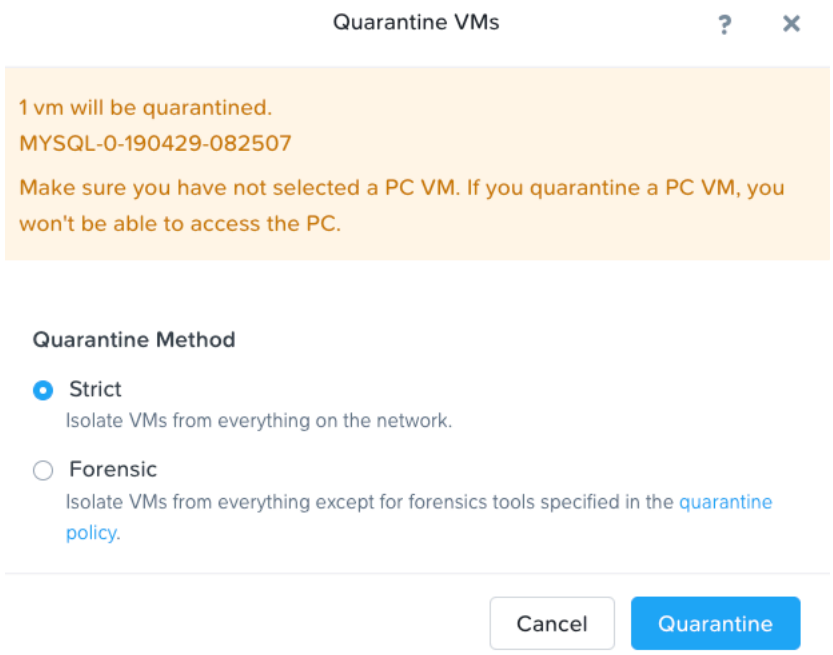

Figure 15: Quarantine Methods

Select Strict to put the VM in the default quarantine category, which restricts all traffic to and from the VM. Select Forensic to put the VM in the forensic category, which allows a defined set of sources and destinations.

Modify the quarantine policy to add your own tools as needed. In this example, the security team is allowed access to the forensic quarantined VMs on specific ports. You can add inbound sources and even outbound destinations to the forensic category; however, you can't add any sources or destinations to the default strict category.

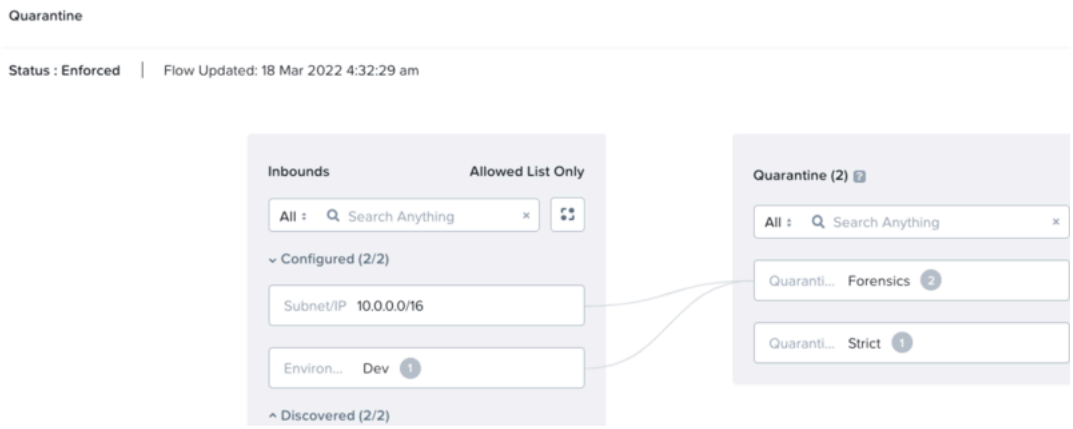

Figure 16: Forensic Quarantine Policy Definition

Quarantine Flow Visualization

Use the quarantine policy view to visualize the traffic to and from quarantined VMs. Traffic blocked by the quarantine policy displays a red blocked symbol.

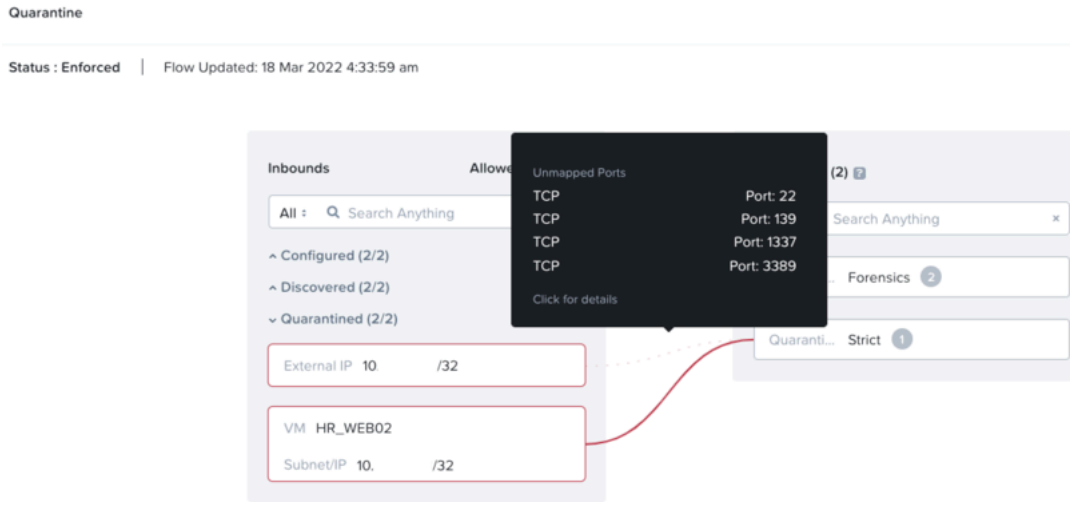

Figure 17: Quarantine Flow Visualization

Using Isolation Policies

Isolation policies prevent two defined groups of entities (VMs) from communicating with each other. This kind of restriction is useful in instances with a limited number of groups, where a group requires absolute isolation from

one other defined group. Isolation policies still allow traffic to undefined groups and traffic within a group. To restrict traffic to undefined groups and to control traffic within the group, use application policies instead of isolation policies.

Environment isolation is a great example of how to use isolation policies, where VMs in category Environment: Production shouldn't communicate with VMs in Environment: Dev.

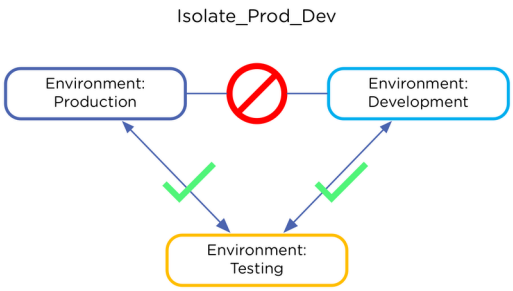

Figure 18: Isolating Prod from Dev

Create isolation policies by navigating to Networking and Security and Security Policies. Click Create Security Policy, then select Isolate Environments (Isolation Policy).

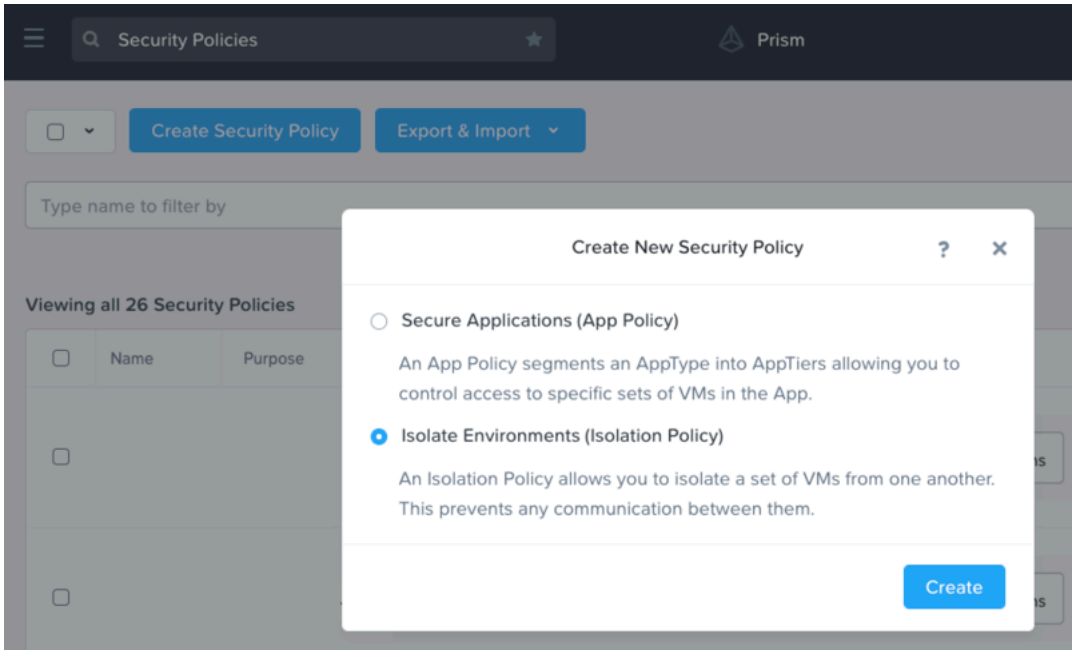

Figure 19: Isolation Policy Menu

Give the isolation policy a name and select the two categories that should be separated from each other. You can use any two categories to create an isolation policy.

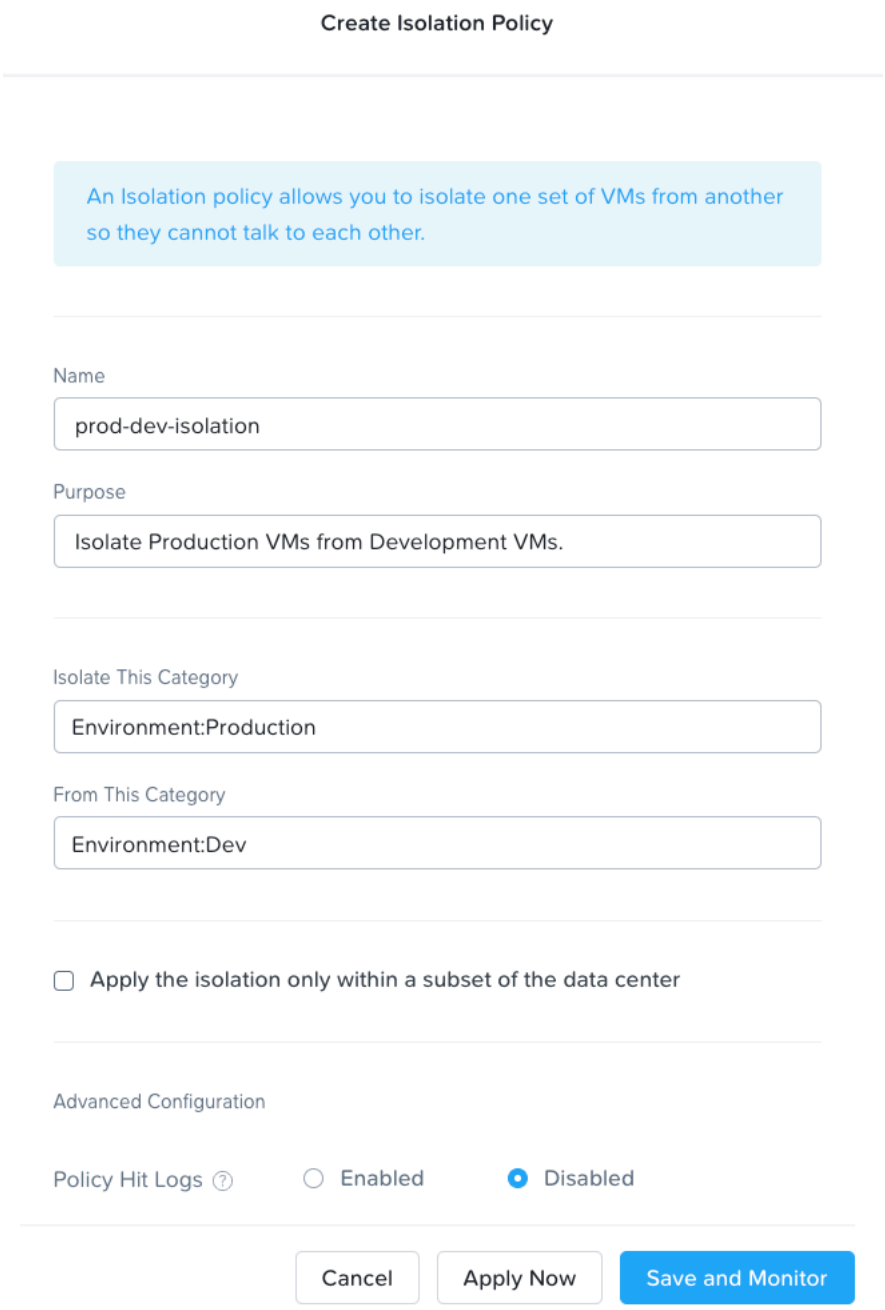

Figure 20: Isolation Policy Creation

If you have more than two groups that require separation (for example, Production, Dev, Staging, and Testing), create an isolation policy for each

unique pair of groups. For a large number of groups, this list of isolation policies can grow long, so you may find application policies to be more effective than isolation policies in this instance. For example, to separate four groups, you would need the following six policies.

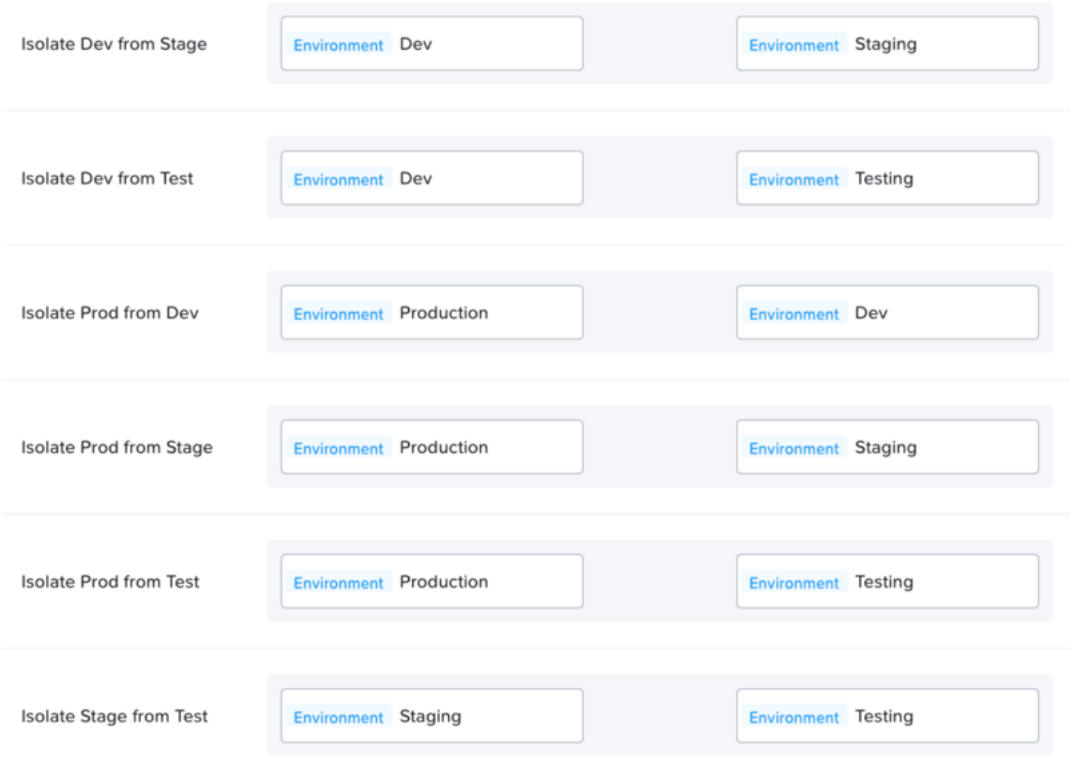

Figure 21: Isolation Policies for Four Categories

Adding one more group that requires isolation, such as Environment: Backup, requires four new isolation policies—one to isolate each of the previously existing groups from the new group.

Isolation Flow Visualization

Isolation policies show the number of connection attempts between isolated groups, along with the discovered ports. For details on connection attempts between isolated categories, such as the source or destination VMs, use Flow Network Security policy hit logs collected on a syslog server.

Isolate Within a Subset

In addition to simple isolation between two groups, Flow Network Security also offers the option for the isolation policy to function within only a subset of VMs in the matching pair. For example, our previous isolation policy separated all Environment: Production from all Environment: Development. If we only wanted to separate Production from Development when the VMs also had the category Site: Branch-001 applied, we add Site: Branch-001 as a subset to the isolation policy.

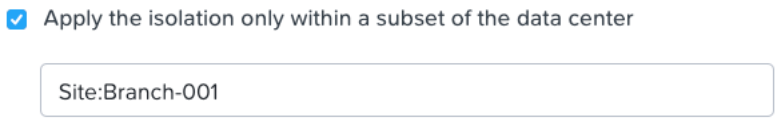

Figure 22: Isolate Within a Subset

Using Application Policies

Application policies are the most flexible type of security policy, defining both inbound traffic sources and outbound destinations for a single- or multitiered application. An application policy can also control traffic to, from, and within individual application tiers. You must configure your application policy with a single AppType category that acts as the target group, because the mapping between application policies and AppType categories is 1:1. For example, you can use the category AppType: Exchange as the AppType in only one application security policy.

Use application policies when you need to allow traffic on specific ports and protocols between sources and destinations. Application policies can also isolate one group of AppType VMs from all other AppType VMs without needing to use isolation policies. The key difference between these two policy types is that application policies allow configurable sources and destinations between apps, while isolation policies enforce strict separation with no exceptions. The following diagram shows the allowed inbound sources on the left and outbound destinations on the right, with a single-tiered application in the center. The central application is labeled Your App in the diagram, but in some contexts Nutanix also calls it the target group.

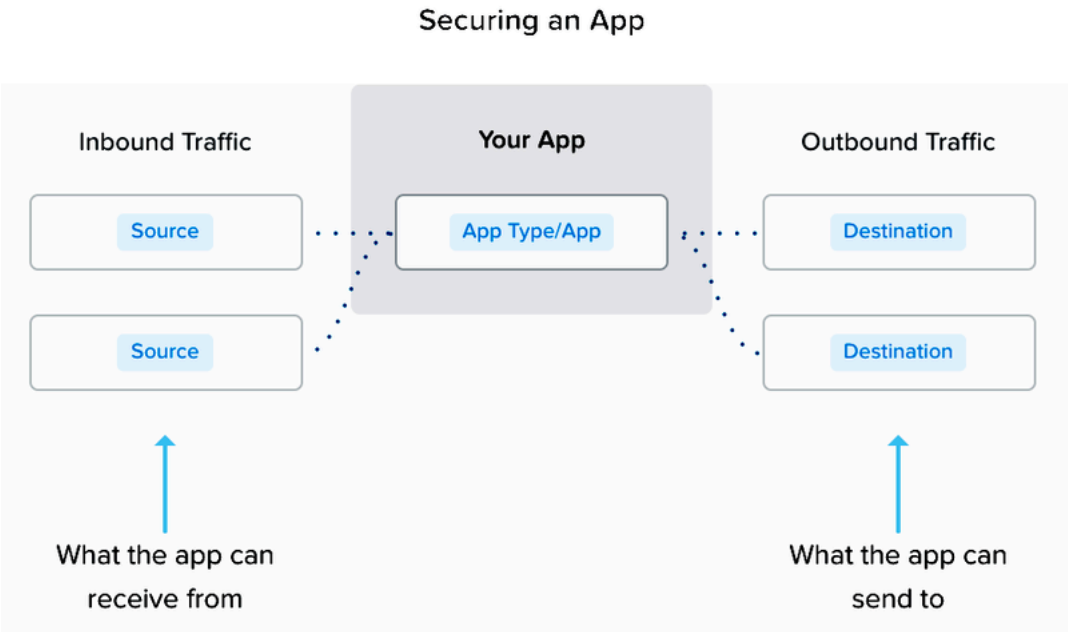

Figure 23: Application Policy Source, Application, and Destination

To create application policies in Prism Central, select Network & Security, Security Policies, click Create Security Policy, then select Secure Applications (App Policy).

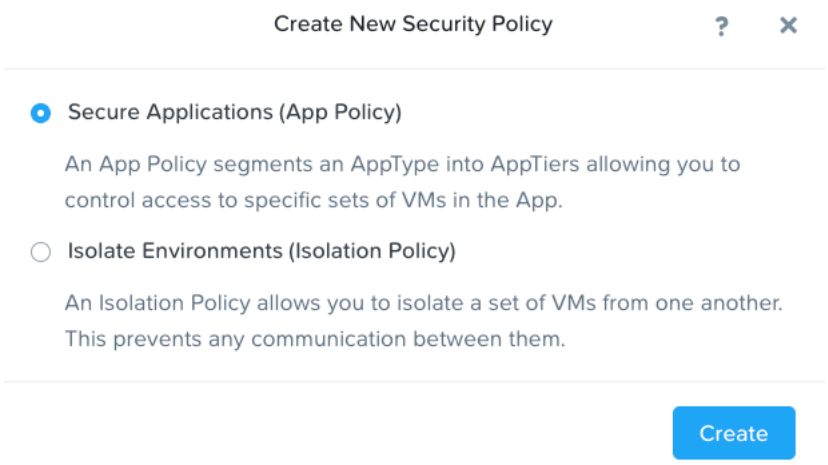

Figure 24: Create Application Policy

Enter a name for the policy and a helpful text description. The application dropdown menu shows the list of available AppType categories.

Note: You can only use an AppType category to create an application policy.

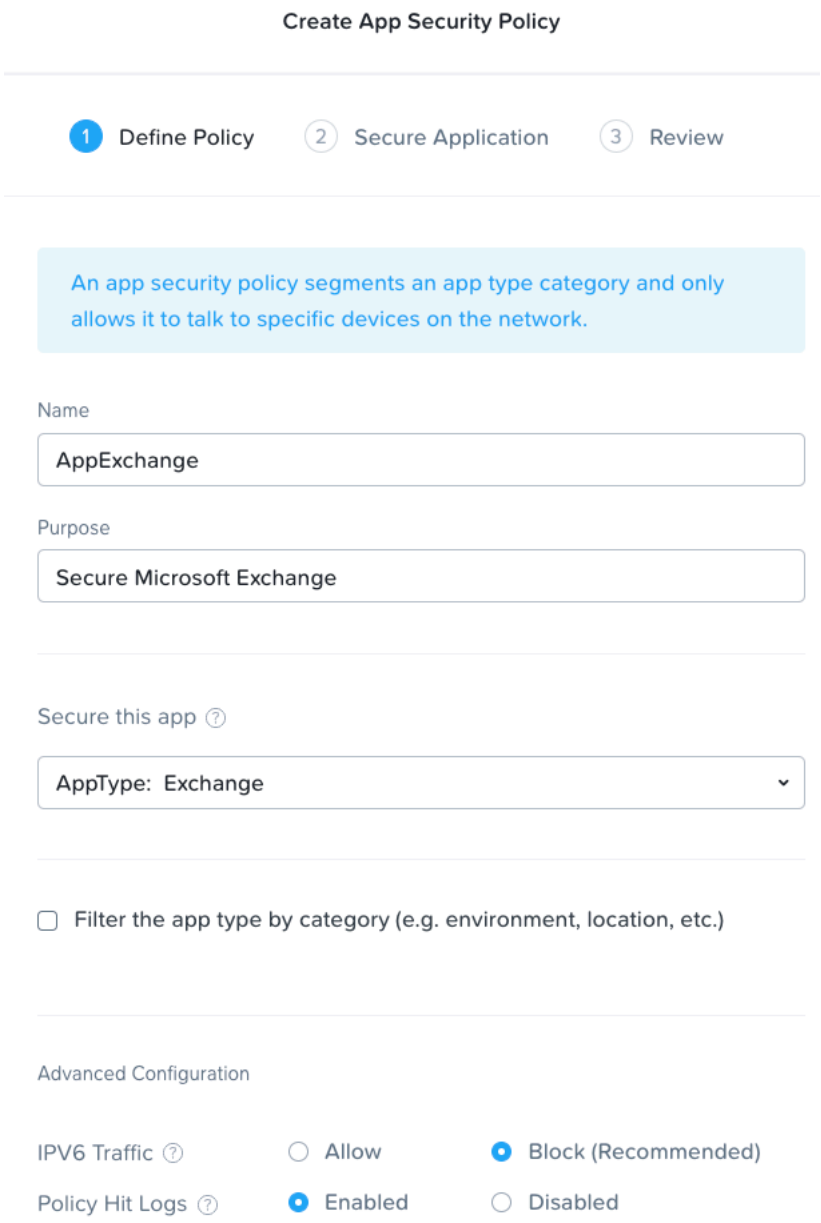

Figure 25: Application Policy Creation

The application policy definition page also controls whether IPv6 traffic is allowed for this policy. Nutanix recommends blocking IPv6 traffic in security policies; otherwise, all IPv6 traffic is allowed to and from the target group, with no restrictions.

You can also enable policy hit log creation for the specified policy. When you select this option, traffic sessions to and from the protected application generate syslog messages to the configured syslog server.

Sources and Destinations

After you select the application using the AppType category, choose whether to allow or safelist inbound sources and outbound destinations. With the Allow All option, all sources or destinations can communicate with the application. To control which sources or destinations are allowed, use the Allowed List Only option to allow explicit categories or IP subnets and block any unspecified traffic.

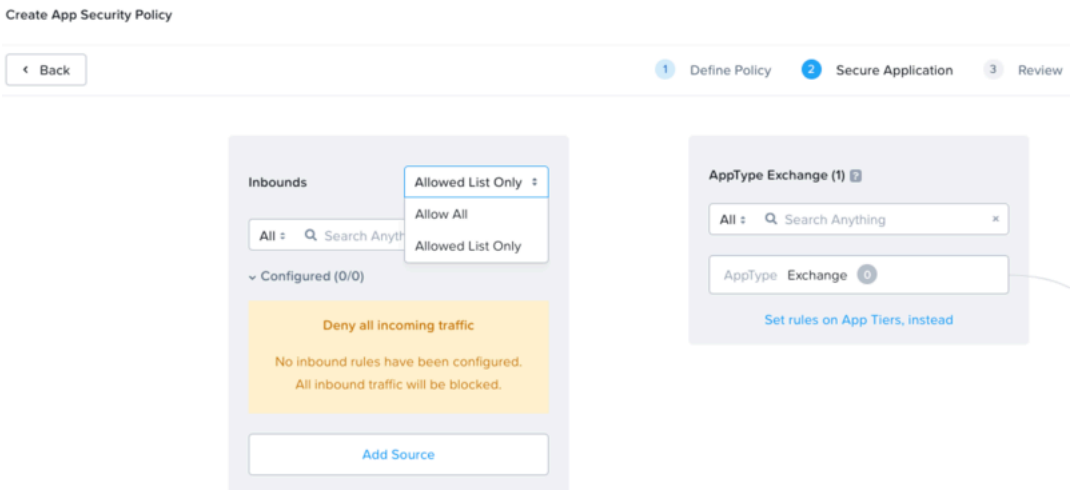

Figure 26: Allow All vs. Allowlist Only

Add inbound sources on the left side of the application policy by category or IP subnet. Select Category from the dropdown menu and use the text box to search for and select the desired category. All VMs tagged with this category are added as an allowed source.

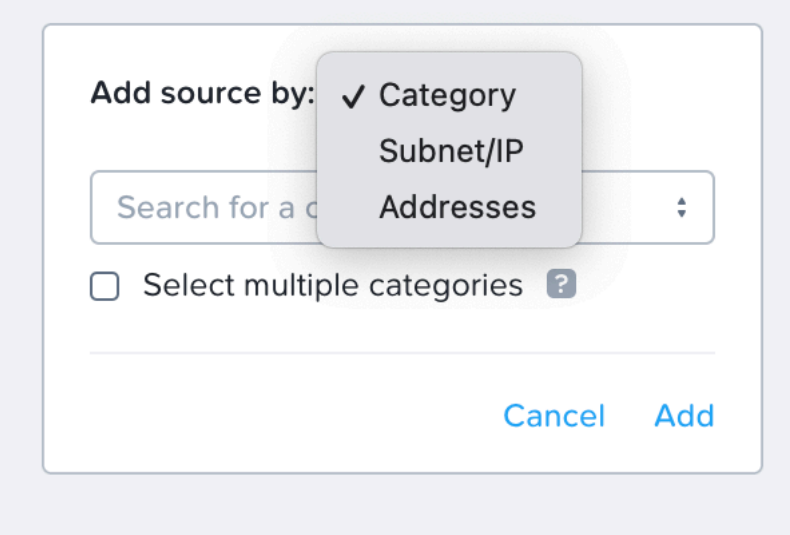

Figure 27: Add Source by Category

Select Subnet/IP and enter the IPv4 network address in CIDR notation (IP/mask length). For example, use a mask length of 24 to allow an entire Class C subnet, or a mask length of 32 to allow only a single host IP address. Using IP subnets is helpful when the source or destination you're adding is external to the Nutanix cluster.

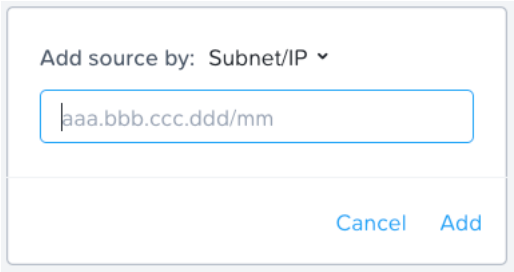

Figure 28: Add Source by Subnet/IP

After clicking Add to add a source, click the plus sign to select the traffic's destination, allowing traffic to a specific tier in the application.

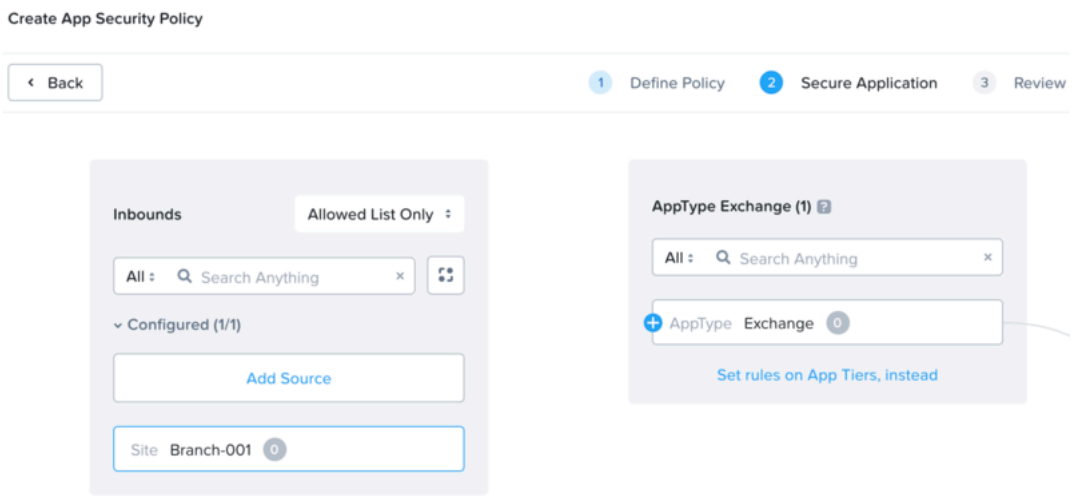

Figure 29: Connecting a Source to AppType

Add allowed protocols and ports such as TCP, UDP, and ICMP to the flow.

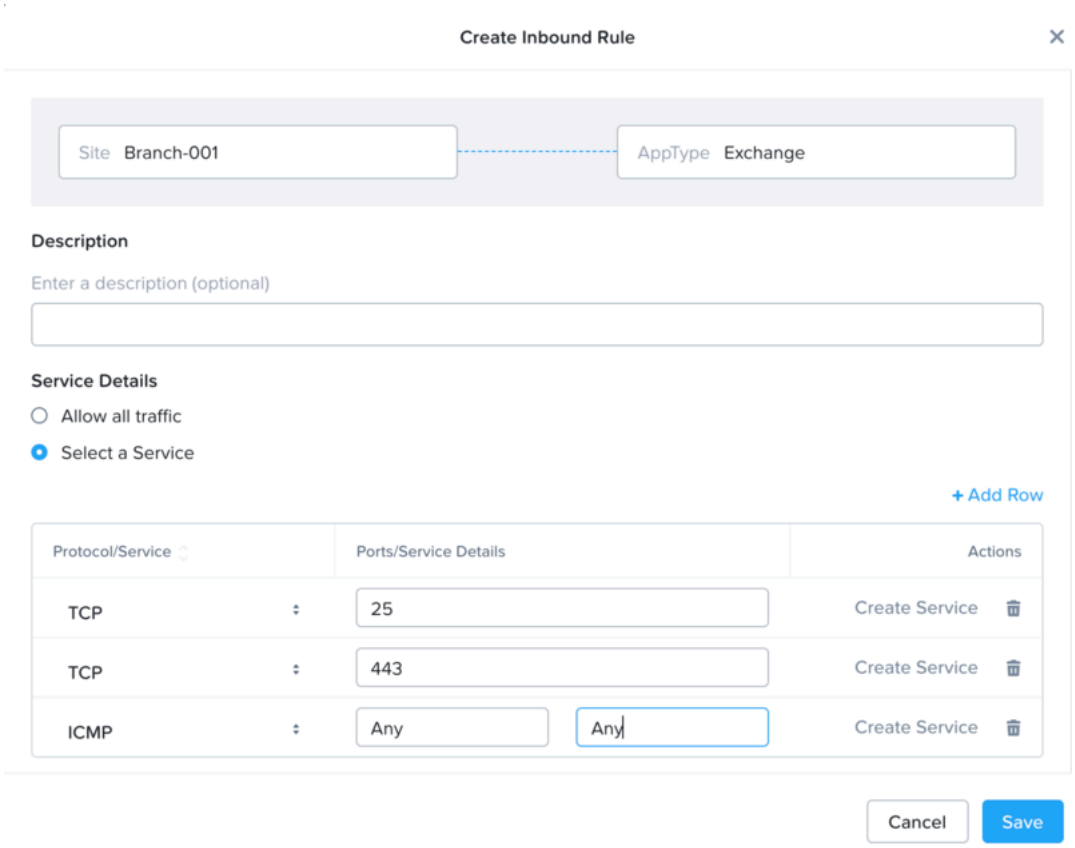

Figure 30: Allowed Flow Specification

Adding categories and IP subnets as destinations in an application policy works in a similar way, but you select the plus sign on the source tier of the application after adding a destination on the right.

Application Tiers

Add tiers to an application that requires granular traffic control for the different VMs in the application. In the previous example, we combined the Exchange AppType category with two AppTier categories: Exchange_Mailbox and Exchange_Edge_Transport. To add these AppTier categories in the application policy, select Set rules on App Tiers, instead to allow traffic to and from individual tiers of the application.

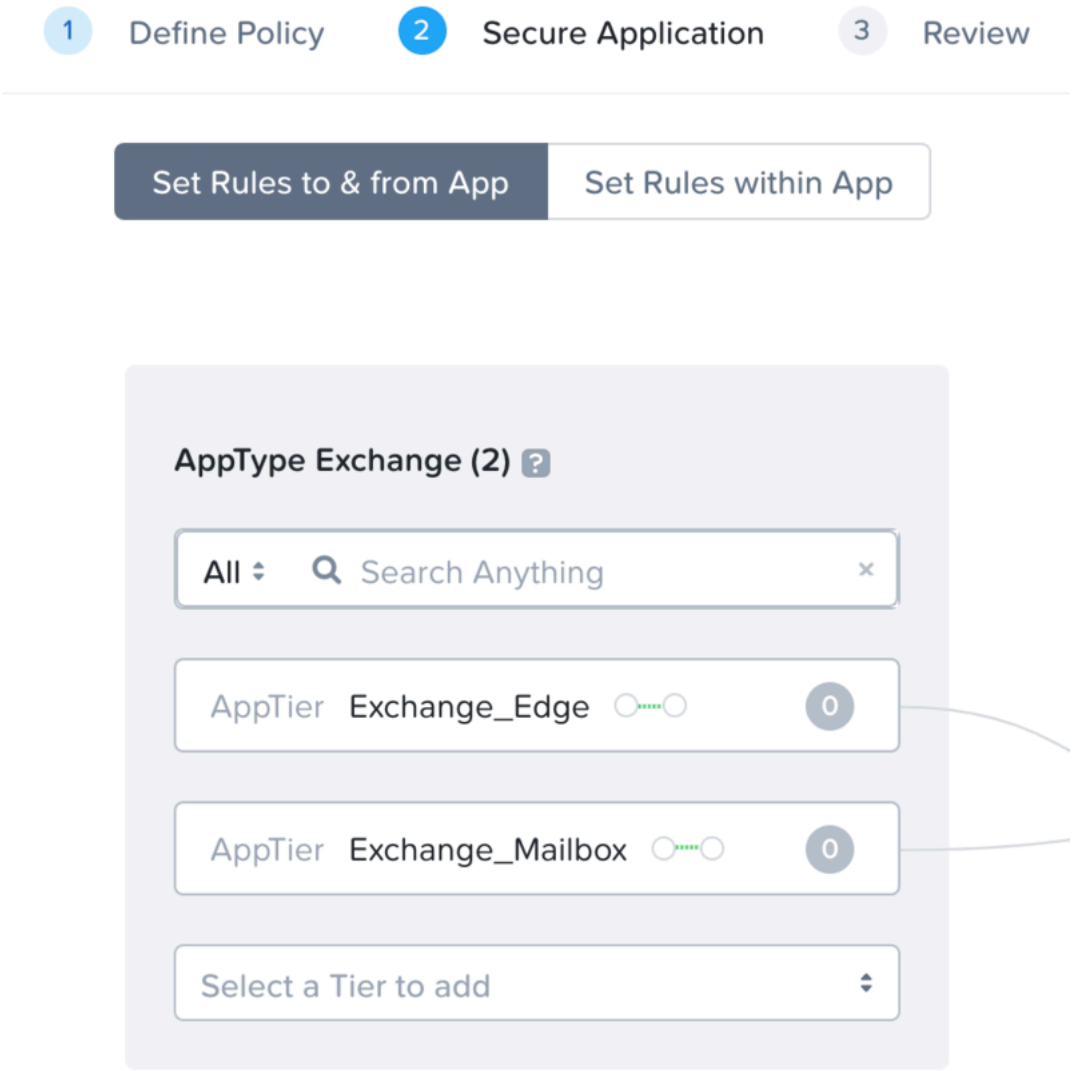

Figure 31: Add Application Tiers

Now you can control sources and destinations at the individual tier level.

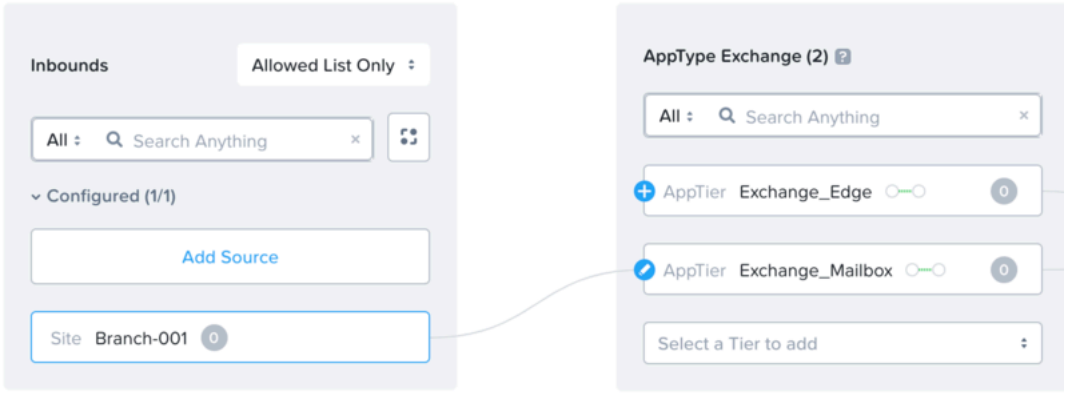

Figure 32: Source to AppTier

To control traffic between application tiers, click Set Rules within App at the top of the screen and select the tiers you want to allow traffic from. Use the plus signs that appear on the other tiers to define the allowed traffic.

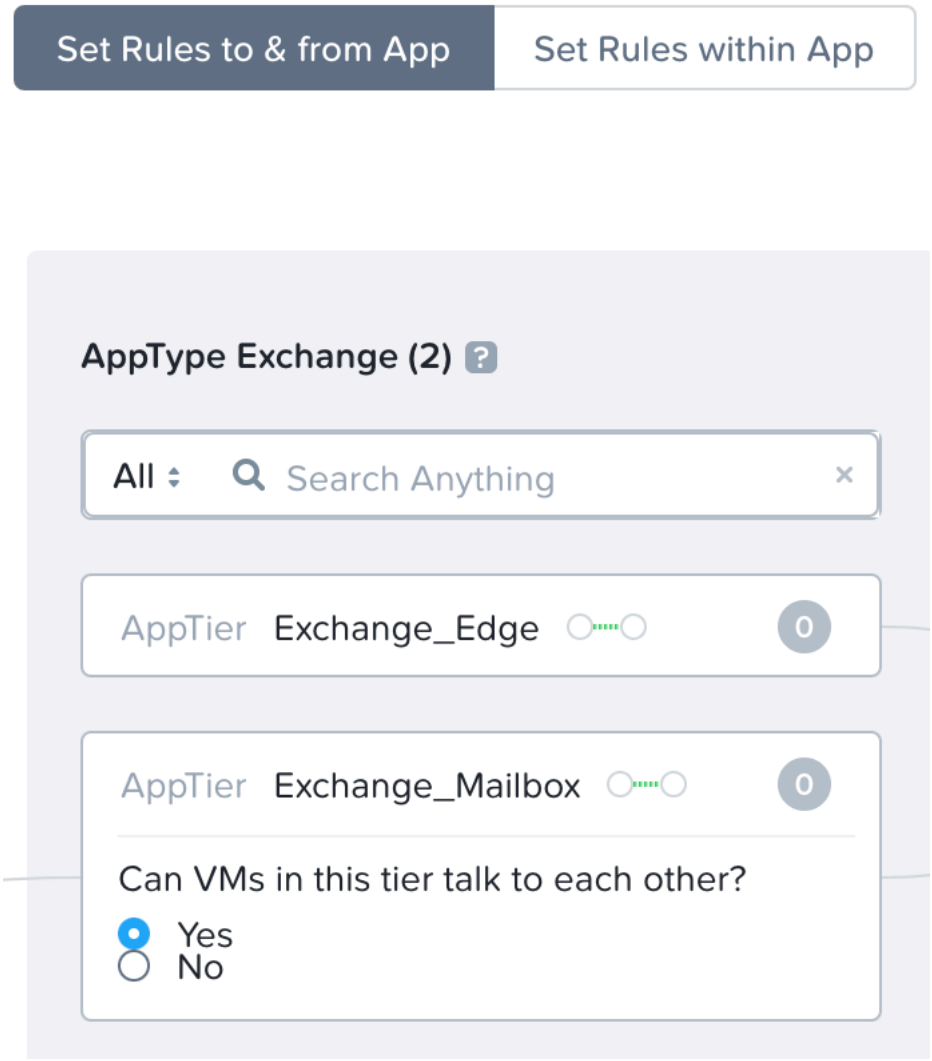

Figure 33: Set Rules Within App

To control traffic between VMs in a single application tier, select whether traffic is allowed within the same tier. For Exchange, it's desirable to allow traffic between servers in the same tier, but an administrator may want to disallow traffic between web VMs in the same front-end web application tier for additional security. This additional specification is also very useful between desktop VMs to prevent the spread of malware.

Service Chain Insertion

You can direct each defined flow in an application policy through a service chain when a chain exists. Service chains define a set of network function VMs (NFVs) for advanced traffic processing. For example, the service chain can direct network traffic on a specific port to a VM for antivirus scanning, deep packet inspection, or packet capture. You can combine multiple NFVs in a chain to apply multiple functions to guest VM traffic.

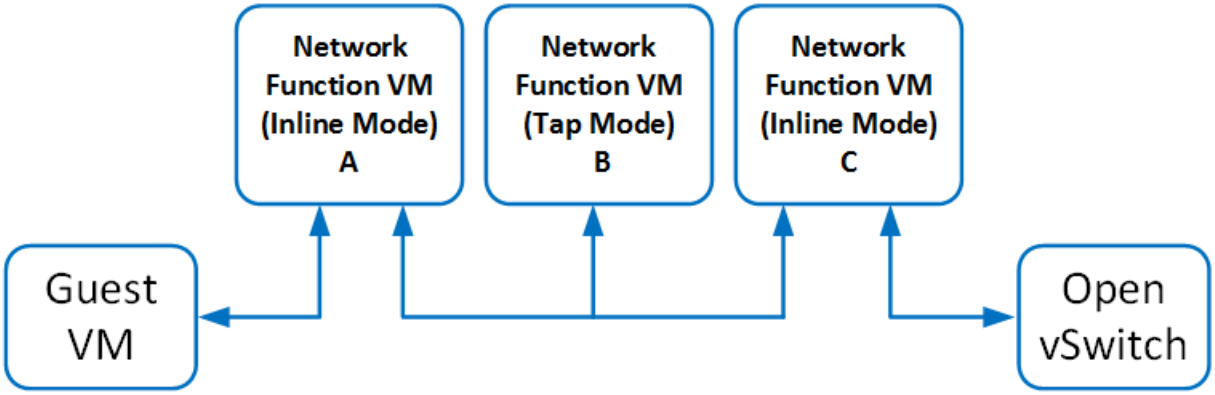

Figure 34: Network Function VM Chain

Nutanix partners with companies that provide NFVs. The deployment workflow of these partner applications can automate the creation of a service chain directly from a Calm blueprint or other orchestration tool. For assistance with manually implementing service chains, see the recommendations in KB 12833.

Once your partner or Nutanix Professional Services creates the service chain, it's available to use in Flow Network Security right away. In Prism Central, use Flow Network Security to create the allowed inbound or outbound rule as usual, then select the desired service chain from the dropdown menu.

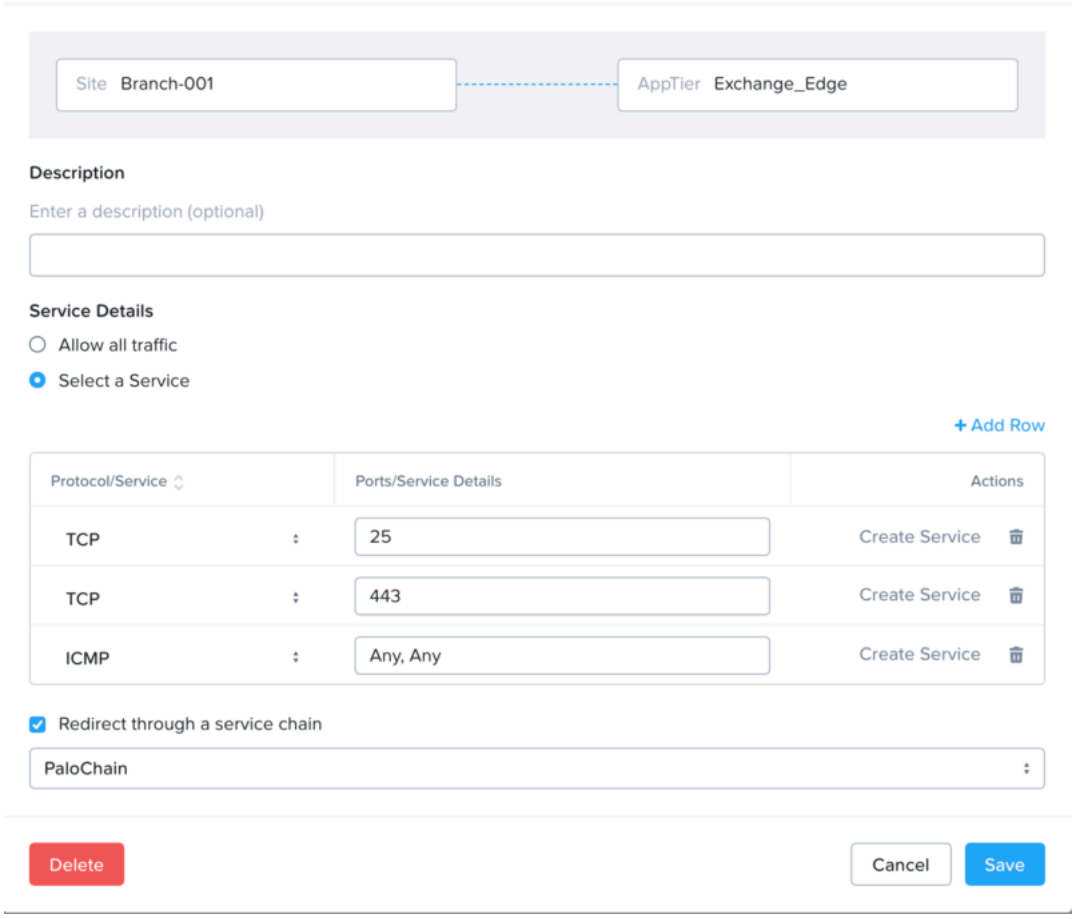

Figure 35: Redirect Through a Service Chain

Flow Network Security Monitoring

Application policies provide flow visualization based on 24 hours of collected traffic statistics to monitor blocked and allowed traffic flows. Nutanix AHV hosts collect traffic information and send it to Prism Central, which builds a flow visualization view inside the policy. The traffic display may take a few minutes for processing and presentation. Flow Network Security monitoring and visualization is meant to simplify policy creation and shouldn't be used as a traffic auditing tool. Instead use syslog as the source of truth for audit purposes.

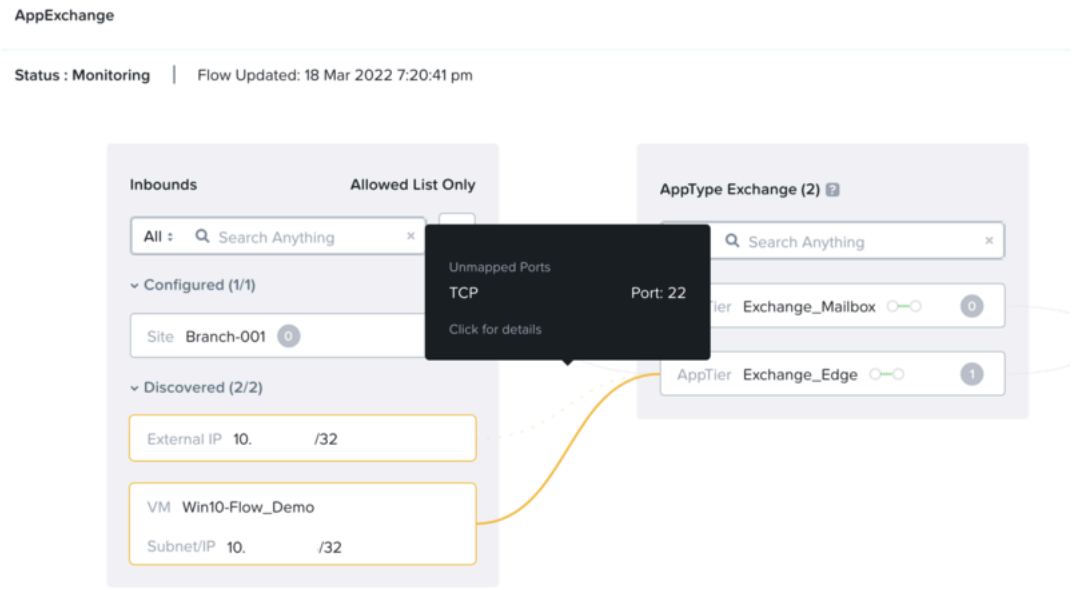

Figure 36: Flows Detected in Monitor Mode

When an application policy is in Monitor mode, traffic to the application that isn't allowed in the policy is shown in yellow. Hovering over the traffic flow shows details about port and protocol. Clicking on the detected flow shows a list of all the ports and protocols. When you edit the policy, hovering over the traffic inbound source or outbound destination allows you to add a specific flow to the policy. This flexibility helps you create an accurate policy, ensure that no traffic is missed, and see what activity is taking place.

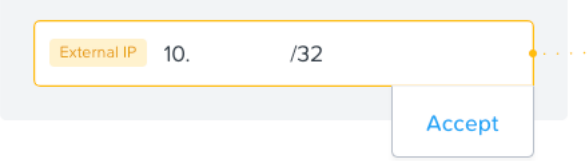

Figure 37: Accept Detected Traffic and Add to Policy

When an application policy is in Enforce mode, denied traffic displays with a red block symbol, denoting connection attempts that the policy blocked.

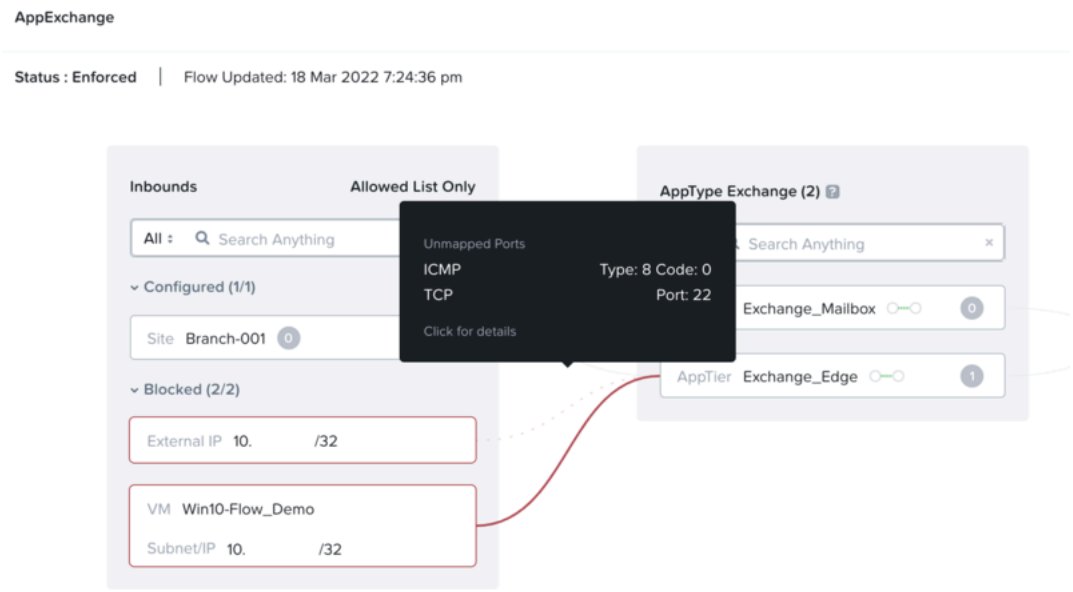

Figure 38: Flows Blocked in Enforce Mode

Each of these visual displays of traffic also generates a syslog entry sent to the configured syslog server if you enable hit logs for the policy.

Application Policies Filter

When you create an application policy, you can also select a specific subset of categories where this application rule applies. Using the AppType: Exchange example and selecting a filter of Environment: Production, the application policy only applies to AppType: Exchange VMs that also share the category Environment: Production.

To avoid ambiguity, when the AppType policy uses the Environment filter, all sources and destinations must have an Environment category explicitly defined.

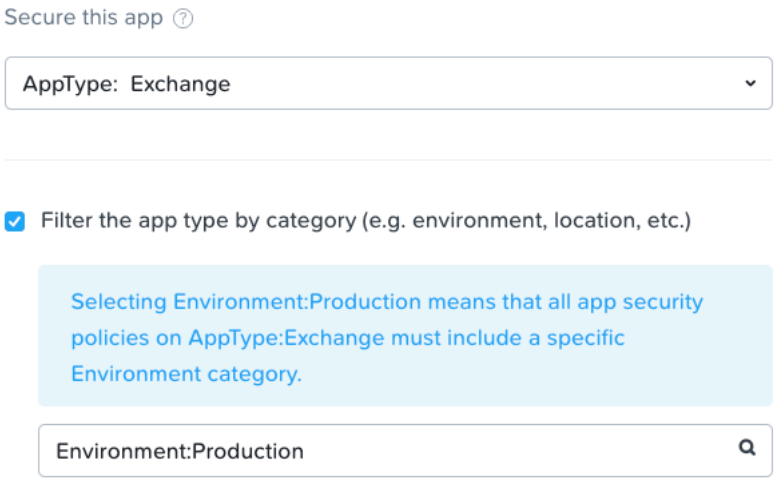

Figure 39: Application Policies Within a Subset

Using the filter allows you to protect applications in different environments with different security policies.

Using VDI Policies for Identity-Based Security

Flow Network Security 5.17 introduced identity-based security to allow administrators to categorize VMs based on the Active Directory group of the signed-in user. Identity-based security protects VDI and adds a new Flow Network Security policy called the VDI policy. You can find more information on using and configuring the VDI policy in the Nutanix Flow Network Security Guide.

There is only one VDI policy in any system using ID-based security, and it's evaluated after all other policies. The VDI policy is similar to the Application policy with inbound sources, outbound destinations, and multiple tiers. User groups are imported from Active Directory and mapped to ADGroup categories in Prism Central, such as ADGroup: Marketing. Each mapped ADGroup becomes a tier in the VDI policy. When a user with a mapped group signs in to an AHV VM that matches the inclusion criteria and isn't already assigned an AppType, ADGroup categories are applied to the VM.

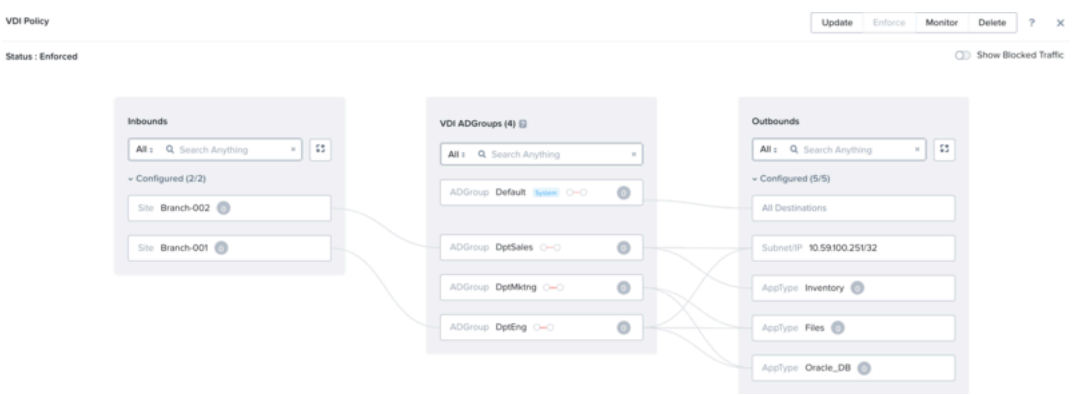

Figure 40: VDI Policy

Flow Network Security evaluates the tiers of matching ADGroups together and applies the resulting combined allowlist to the VM. For example, if a user belongs to ADGroup: Marketing and ADGroup: Engineering, the rules for both Marketing and Engineering are applied to the VM. A user who signs in to a VM that belongs to the engineering and marketing groups has access to Oracle and Files in this example because of the combination of policies.

Use the default intratier settings of the VDI policy to block traffic between VMs in the same tier. This setting helps prevent the spread of malware from one desktop to another.

Logging and Auditing with Syslog

Flow Network Security provides two types of logs: audit logs and policy hit logs. Audit logs track changes to security policy configuration and VM-to-category mappings. Use audit logs to determine when a policy was changed or applied and who changed it.

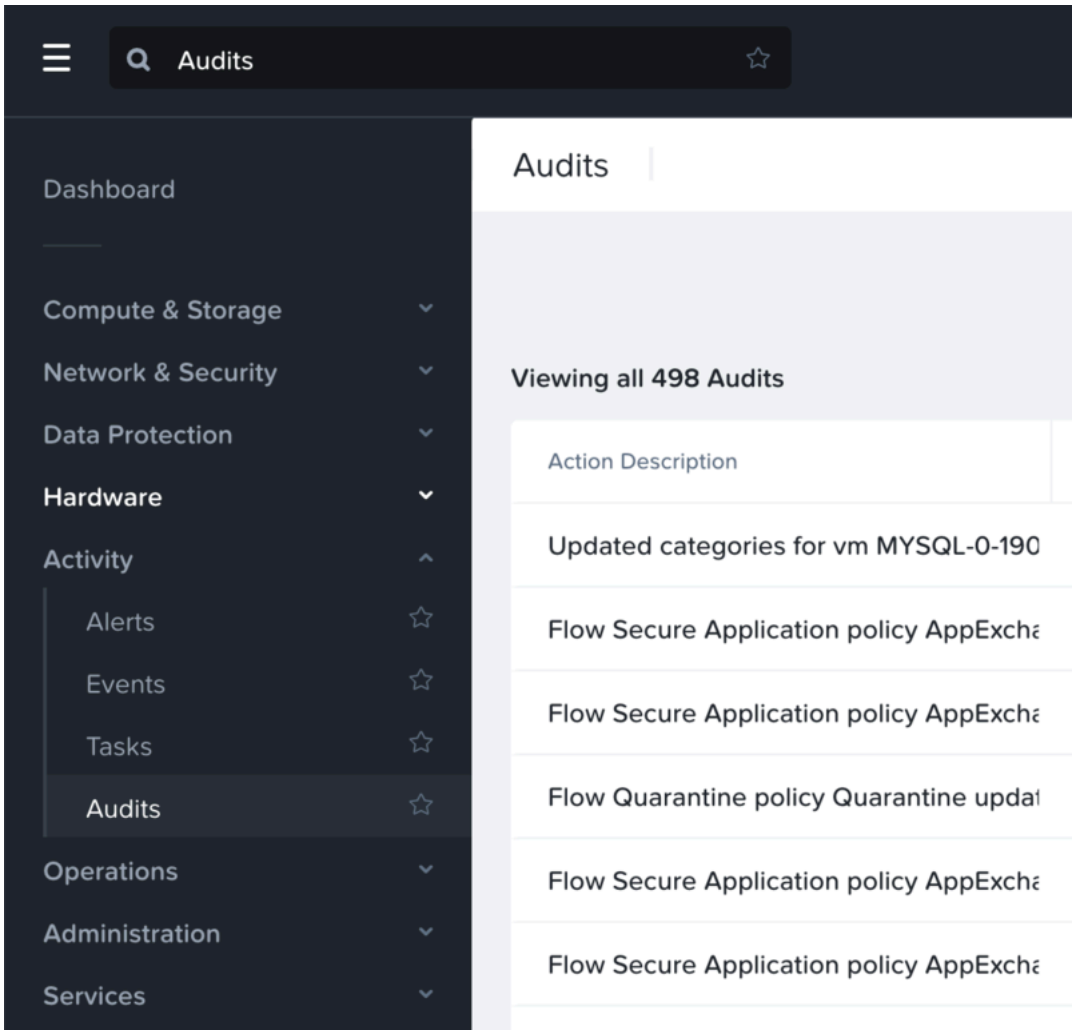

Figure 41: View Audit Events for Policy Changes

Policy hit logs track network flows and whether they were allowed or denied by a specific policy. Use policy hit logs to determine if specific traffic is present on the network and how a security policy affects the traffic. Use policy hit logs as the definitive tool for tracking connections to secured VMs.

Audit logs are enabled by default and capture all changes related to Flow Network Security made in Prism Central. Policy hit logs are enabled for each policy and are disabled by default. Policy hit logs may generate a large amount of data. To analyze the data from policy hit logs, use an external remote syslog server or SIEM system to collect these events. Prism Central sends audit logs to

the remote syslog server, and you can also view them in Prism Central. Policy hit logs are sent directly from each AHV host to the syslog server but generate too much data to consume inside Prism, so you must do policy hit log analysis on the external SIEM.

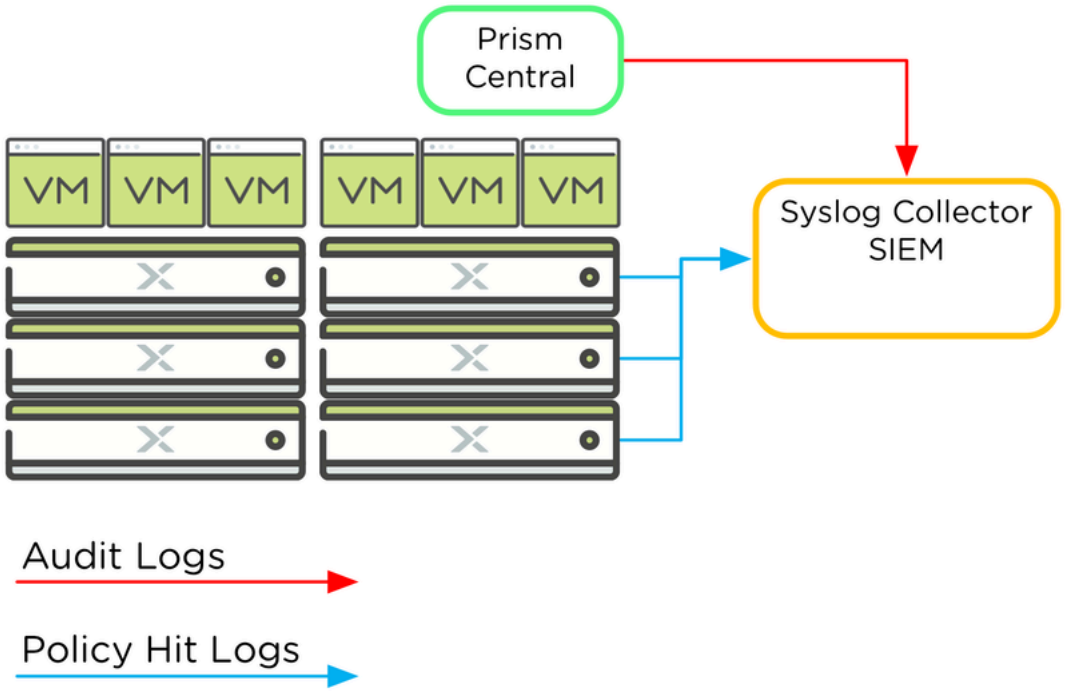

Figure 42: External Logging Architecture

Ensure that the remote syslog server or SIEM expects traffic sourced from both Prism Central and each individual AHV host. Configure a remote syslog server in Prism Central and select the desired port and protocol.

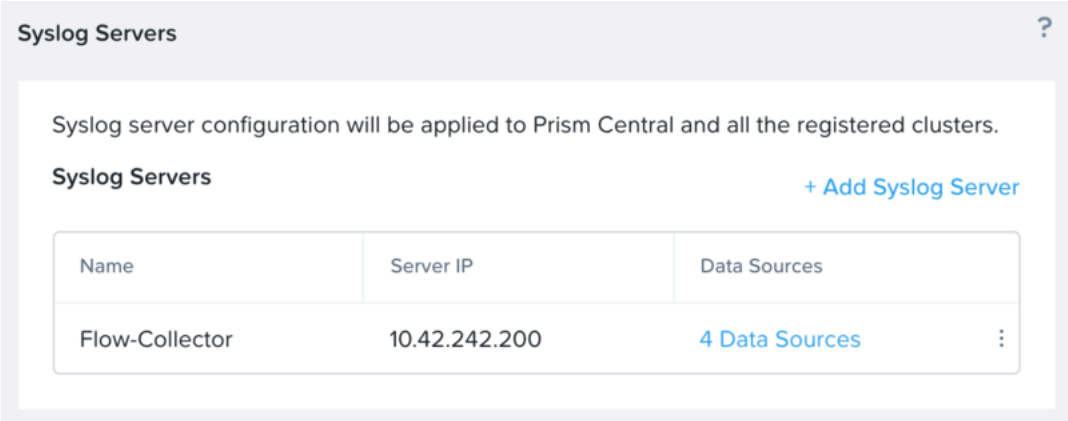

Figure 43: Configure External Logging Destination

Select the desired modules such as Audit and Security Policy Hit Logs to collect audit and policy hit logs and set the severity level to 6 Informational.

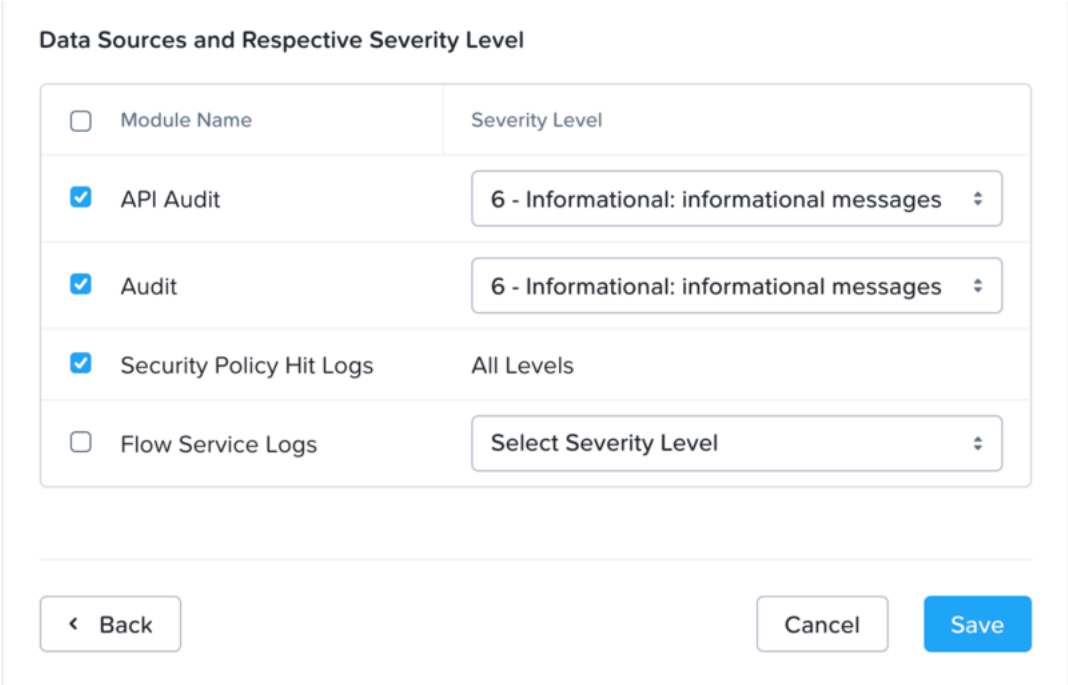

Figure 44: Select Logging Data to Export

The API Audit and Audit modules capture changes to the policy made through REST API and Prism, respectively. The Security Policy Hit Logs module corresponds to a policy hit log and the severity cannot be modified. The

API Audit module captures any changes to policy or category made directly through the REST API endpoint in Prism Central. Only use Flow Network Security service logs at the direction of Nutanix support.

Using Export and Import to Backup Policies

Starting in Prism Central 5.11, use the export file to back up security policies to restore them at a later point or to transfer existing policies to a new system.

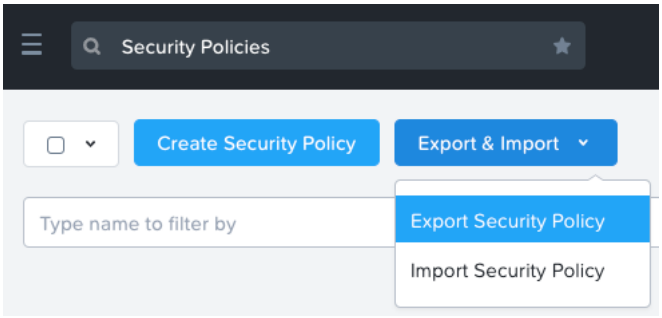

Figure 45: Policy Import and Export

An export captures all security policies and the categories used to build them. An import replaces all existing security policies with the policies from the imported file. Any missing categories are also created at the time of import.

Special Considerations

There are several considerations when building a secure system using Flow Network Security. You must understand how security policies are evaluated and applied to know how VM traffic is going to be processed.

Mapping Categories to VM IP Addresses

Flow Network Security uses categories assigned to VMs and learns a list of IP addresses associated with these VMs. Each VM may have more than one IP address for multiple interfaces or shared virtual IP addresses. Security rules in the hypervisor virtual switch are created based on these lists of learned IP addresses. A security policy that blocks Environment: Dev from reaching Environment: Prod evaluates a list of source and destination IP addresses to do so.

When AHV IP Address Management (IPAM) assigns the VM IP address with a managed network, AHV immediately knows the address of the VM before the VM turns on. When the VM uses static IP addresses or an external DHCP address, the hypervisor must learn the address through DHCP and ARP snooping. There may be a short delay while the hypervisor learns the new IP address of a VM, and during this time security policies based on this new IP address don't yet protect the VM.

Nutanix recommends using AHV IPAM when using security policies to ensure strict policy enforcement.

IPv6 Addresses

Security policies are based on IPv4 addresses and don't support rules based on IPv6 addresses. Nutanix recommends configuring Block IPv6 for all configured security policies so that all IPv6 traffic to and from a protected VM is dropped. If you allow IPv6 in a security policy, all possible IPv6 sources and destinations are allowed to and from this VM on all ports.

Layer 2 Broadcast Traffic

Application security policies don't block ARP and other layer 2 broadcast traffic to and from protected VMs because this traffic is critical to VM operation. Isolation policies do block all layer 2 broadcast traffic between isolated categories.

4. Conclusion

Flow Network Security is a software-defined networking solution for AHV that provides visualization, automation, and security, and uses Prism Central categories and security policies to protect VMs. By using categories to create flexible groups of VMs, Flow Network Security simplifies security policy definition. With security policies, administrators can protect or isolate applications or environments and quarantine infected or rogue VMs.

Through policy- and application-centric traffic monitoring, Flow Network Security visualization answers the vital question of what VM traffic is sent and received in the virtual system.

The ease of creating and assigning categories and policies opens the door for automation. Security is no longer tied directly to VMs or IP addresses, and administrators can respond quickly to datacenter changes.

Because Flow Network Security is built into Prism Central and works natively in AHV, there are no additional components to install or manage. Network security is now just one click away.

For feedback or questions, please contact us using the Nutanix NEXT Community forums.

5. Appendix

References

1. Flow Network Security - Service Chain Integration (KB 12833)
2. Nutanix Prism Central Guide: Category Management
3. Nutanix Prism Central Guide: Security Policies

About Nutanix

Nutanix is a global leader in cloud software and a pioneer in hyperconverged infrastructure solutions, making clouds invisible and freeing customers to focus on their business outcomes. Organizations around the world use Nutanix software to leverage a single platform to manage any app at any location for their hybrid multicloud environments. Learn more at www.nutanix.com or follow us on Twitter @nutanix.

List of Figures