



# AMD Pensando Policy and Services Manager for Aruba CX 10000: User Guide

February 2023

## Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD, the AMD Arrow logo, Pensando and combinations thereof are trademarks of Advanced Micro Devices, Inc. VMware ESXi™, VMware vSphere® vMotion® and VMware vCenter® are trademarks of VMware. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

© 2022 – 2023 Advanced Micro Devices, Inc. All Rights Reserved.

[amd.com/pensando](https://amd.com/pensando)

PPD22002

## Revision History

Version	Description	Date
1.0	First release	February 9, 2022
1.1	Miscellaneous updates	February 25, 2022
1.2	1.49.1-T beta: hierarchical security policies, policy scaling, ALG enhancements	May 13, 2022
1.3	1.49.1-T final	June 13, 2022
1.4	1.49.2-T beta	July 2022
1.5	1.49.2-T	August 11, 2022
1.6	1.49.3-T: minor corrections, caveat on disabling vSphere DRS	September 2022
1.7	1.54.1-T: new features (see Release Notes)	December 2022
1.7.1	1.54.2-T: minor errata fixes to guide; no new functionality	February 2023

# Contents

<b>Revision History</b>	<b>3</b>
<b>Introduction</b>	<b>9</b>
Key Features	9
Related Documentation	9
<b>Glossary</b>	<b>10</b>
<b>PSM Overview</b>	<b>12</b>
<b>Initial Deployment Workflow: High-Level Overview</b>	<b>15</b>
<b>PSM Object Model</b>	<b>15</b>
Firewall Objects	15
Apps, Network Security Policy	16
Key PSM Objects	17
Labels	18
<b>PSM Installation</b>	<b>19</b>
Storage Considerations	19
Data Retention	20
PSM Installation on ESXi	21
PSM Installation on KVM	22
Bootstrap the PSM Cluster	25
<b>The PSM Graphical User Interface</b>	<b>28</b>
Online Help	28
Searching	29
Global Icons	31
Server Certificate	32
API Capture	32



<b>Create PSM User Authentication Policy and Users</b>	<b>33</b>
User Authentication Policy	33
Local User Lockout Policy	36
Role-Based Access Control (RBAC)	36
Roles	37
Role Binding	38
<b>System Upgrade</b>	<b>39</b>
AFC Upgrade	39
PSM Upgrade	40
Upload PSM Upgrade Bundle	40
Create Rollout	41
AOS-CX Upgrade	43
<b>Configuration Snapshots</b>	<b>44</b>
<b>Associating a DSS with the PSM</b>	<b>45</b>
AOS-CX CLI	45
ZTP	46
AFC	47
Verification	47
<b>Decommissioning a DSS</b>	<b>48</b>
<b>Associating the PSM to Aruba Fabric Composer</b>	<b>51</b>
<b>Firewall Policy Functionality and Configuration</b>	<b>52</b>
Considerations	52
VSX and Firewall High Availability	54
Supported Topologies	55
Connection Tracking	56
Firewall Policy Configuration	57

Create a VRF	57
Create a Network Security Policy	58
Rule Overlap Detection	59
Attach Policy to a Network	61
Create VRF on AOS-CX	62
Create VLAN on AOS-CX	62
Understanding Firewall Policy Scaling Profiles	63
Switch Policy Scaling Profile	63
Verifying the Number of Rules Consumed in the Data Plane	67
Understanding Hierarchical Security Policy	71
Policy Enforcement	73
Configuration and Verification	76
Enable/Disable Individual Firewall Rules	78
Configuring Firewall Log Export	80
Firewall Log Export Policy Configuration	80
Bind Export Policy to DSS	81
Firewall Log Record Format	82
Example of Firewall Syslog Message	84
Deduplication for Firewall Logs	86
Flow Export (IPFIX)	87
Configuring Apps	88
Protocol And Ports	89
ALG	90
DDoS Detection and Alerting: Maximum Sessions and CPS Limits	92
Maximum Session Limit	92
Maximum CPS	92

Min and Max Values	94
Configuring the Maximum Sessions / CPS on a VRF via the PSM UI	94
Configuring the Maximum Sessions/CPS on a Network via the PSM UI	94
API Examples	95
Behavior on Reaching the Maximum Session Limit	99
Behavior on Reaching the Maximum CPS Limit	99
Implication of Configuring session-limit on an Active System	100
VSX Implications	100
Multiple ALG Types/Apps/Protocols in Firewall Policy Rules	102
Steps To Configure Rules With Multiple Proto-Ports Via the PSM UI	102
Steps to configure rules with multiple ALG types via the PSM UI:	103
API Examples:	103
API Example: Policy with Rule Referencing Multiple ALG Types:	105
IP Protocols Support for Firewall Policy	106
UI Examples	106
API Examples	107
<b>Stateful Firewall Flow Migration with vMotion</b>	<b>109</b>
Considerations for Multi-Homed ESX Servers	110
Behavior with Flow Logs and Flow Statistics	110
Configuration	111
DSS Required AOS-CX CLI Configuration	112
Caveats	113
<b>Monitoring the DSM via the PSM UI</b>	<b>113</b>
Metrics Charts	117
<b>Alerts and Events</b>	<b>120</b>
<b>PSM Automation</b>	<b>121</b>

Python Language Bindings	121
Ansible Modules	121
REST API	123
<b>Tech Support Collection</b>	<b>124</b>
<b>Appendix A: PSM Quorum High Availability</b>	<b>126</b>
<b>Appendix B: PSM Operational Network Ports</b>	<b>127</b>
<b>Appendix C: Configuring Microsegmentation in Non-AFC Environments</b>	<b>130</b>
Topology	130
Configuration on the DSS	131
Global Config Mapping the Primary and Secondary VLAN	131
Host-Facing Interface Configured as Regular Trunk, Allowing Both Primary and Secondary VLAN	131
SVI Config on Primary with Local Proxy ARP	132
Configuring VMware (ESXi)	132
Configuration on the PSM	132
<b>Appendix D: Saving the PSM Recovery Key</b>	<b>133</b>
Saving the Key	133
Recovering the Cluster	134
<b>Appendix E: Using the PSM Network Graph to Create Security Policies</b>	<b>135</b>
First Method	135
Second Method	145

## Introduction

This guide describes how to install and operate the *AMD Pensando Policy and Services Manager (PSM)* to manage the stateful services of the *Aruba CX 10000 with AMD Pensando* distributed services switch (abbreviated as either *CX 10000* or *DSS*).

The PSM can be accessed via the IP address or host name of any of the PSM cluster nodes or, if a load balancer is being used, the IP address or host name presented by the load balancer. In this document, the PSM address will be referenced as either `$PSMaddr` when used in the context of shell commands or scripts, or as *PSMaddr* in other examples.

The PSM is managed through either its browser-based GUI or its secure RESTful API. Most examples in this document show the GUI, which is accessible at the URI

`https://PSMaddr` .

## Key Features

Core functionality supported includes:

- Distributed stateful firewall
- Microsegmentation (using PVLAN)
- Flow logging and metrics
- Full AOS-CX routing and switching feature set (see AOS-CX documentation for further details)
- Fabric and services orchestration with AFC and PSM
- DDoS detection and alerting

## Related Documentation

- *Aruba CX 10000 Switch Series Installation and Getting Started Guide*
- *PSM Release Notes*
- Release notes for AOS-CX and Aruba Fabric Composer (AFC)
- *Aruba Fabric Composer User Guide*
- AOS-CX 10.00 Feature Guides

Aruba documentation can be found at the [Aruba 10000 Switch Series documentation portal](#).

Review the PSM release notes for details and information about new features, known issues, fixed bugs, and supported servers, cables, and switches.

# Glossary

Name	Description
AFC	Aruba Fabric Composer
AOS-CX	The Aruba switch operating system, providing network services functions and management
CoPP	Control Plane Policing
Data traffic	(aka <i>data plane traffic</i> ) the actual network data being processed by the DSS environment
DSE	Distributed Services Entity: collectively describes the services and monitorability provided by the two DSMs in a DSS
DSM	AMD Pensando Distributed Services Module (two per DSS): the stateful services execution engine of the DSS
DSS	Aruba Distributed Services Switch with AMD Pensando
Egress	Traffic leaving a host to fabric, in reference to security policies
Ingress	Traffic entering a host from fabric, in reference to security policies
ISL	Inter-Switch Link, a layer 2 interface between two VSX peer switches
Management and control traffic	(processed by the management and <i>control plane</i> ) network communication related to the interoperability, reporting, and policy management of the DSS environment

Table 1: Glossary of terms (1/2)

Name	Description
Persona	A configuration type that can be set for a port, determining if it is connected to workloads or to the network. Can be either <code>access</code> for host-facing ports, or <code>uplink</code> for fabric-facing ports.
PVLAN	Private virtual LAN
PSM	AMD Pensando Policy and Services Manager
VRF	Virtual Routing and Forwarding instance
VSX	Aruba Virtual Switching Extension, providing high availability and redundancy capabilities
ZTP	Zero-Touch Provisioning: automated network configuration and deployment of managed devices

*Table 1: Glossary of terms (2/2)*

## PSM Overview

The AMD Pensando Policy and Services Manager is a programmable, secure, highly available, centralized system for managing infrastructure policy, with capabilities for:

- Deploying and controlling distributed firewall security
- Telemetry and analytics
- Troubleshooting
- Operations and maintenance: events, alerts, technical support
- Authentication, authorization, and accounting (AAA)

The PSM is designed to establish and manage consistent policies for a number of Distributed Services Switches. (Refer to the *Aruba CX 10000 Release Notes* for current support limits.)

The PSM operates as a 3-node quorum-based cluster running on virtual machines (VMs) hosted on multiple servers for fault tolerance. A PSM cluster can tolerate the loss of one controller node and continue to maintain full service. The PSM cluster is not involved in datapath operations; if it becomes unreachable or multiple nodes fail, there will be no impact on data traffic and stateful services on the DSSes it manages.



Figure 1 is a diagram of the interconnection between the PSM and the switches it manages; interactions take place through an IP network.


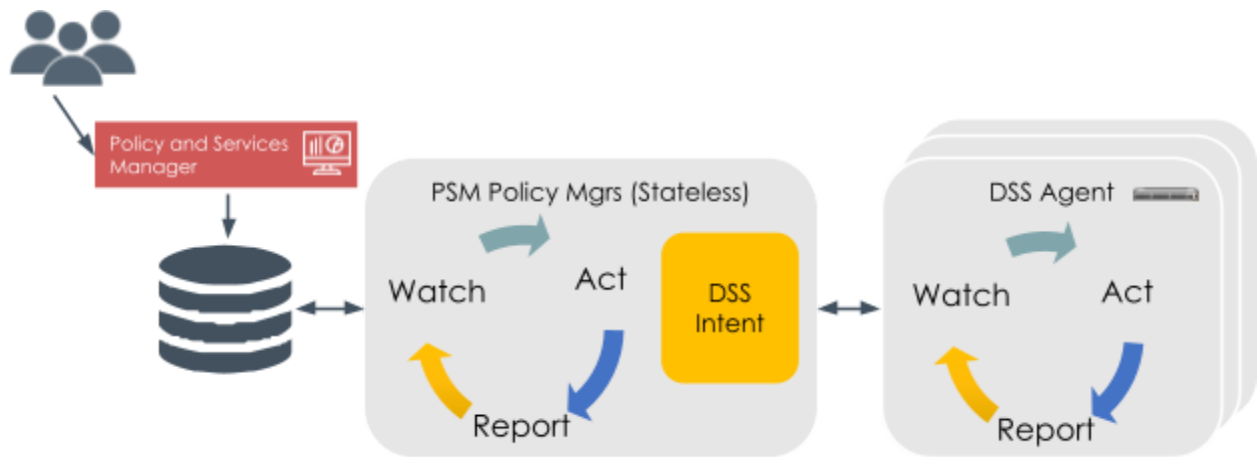


Figure 1. PSM/DSS management plane

Each DSS is configured with an IP address that is used for communication with its associated PSM over any IP network. This is referred to as its *management address*.

Each DSS runs an agent which constantly watches for incoming configuration changes upon which it must take action.

The PSM employs an *intent-based* configuration management structure, similar to Kubernetes. Any configuration changes are continuously monitored within the PSM until it has been confirmed that the changes have been propagated to all DSSes. The PSM resends configuration requests until the desired state is reported back from each DSS, as shown in Figure 2:



*Figure 2. Working principles of intent-based configuration*

Intent is expressed in terms of *policies* established for firewall and flow telemetry.

## Initial Deployment Workflow: High-Level Overview

This is an outline of the steps necessary for initial deployment of a PSM cluster and its associated DSSes. Detailed steps are provided further below in this document.

- **Install the PSM**

- Install the PSM software on either an ESX-based or KVM-based 3-node cluster.
- Configure the PSM using the `bootstrap_PSM.py` utility.
- Save a copy of the PSM recovery key on a different server from PSM, in case the PSM needs to be rebuilt later as part of disaster recovery. (Refer to Appendix D for more details)
- Set the PSM user authentication policy, and create PSM users with appropriate roles.

- **DSS Configuration**

For each DSS:

- Plan for one additional IP address allocated to each DSS as a management interface, configured either from its host or via DHCP.
- Associate each DSS to the PSM
- Admit the DSS into its PSM cluster. The PSM can be configured to do this automatically.

Installation of the PSM cluster is a one-time activity; other procedures may be performed during initial installation, but will also be part of the standard operation of the PSM, performed as more DSSes are added.

## PSM Object Model

The PSM's intent-based paradigm relies on the PSM object model described in this section.

### Firewall Objects

The primary firewall objects are illustrated in Figure 3. The *NSPRule* (Network Security Policy Rule) specifies the firewall behavior, but is not a managed object itself. Instead, the *NetworkSecurityPolicy*<sup>1</sup> is the managed object that contains an array of NSPRule specifications.

---

<sup>1</sup> Refer to the Release Notes for the number of NetworkSecurityPolicy objects supported.

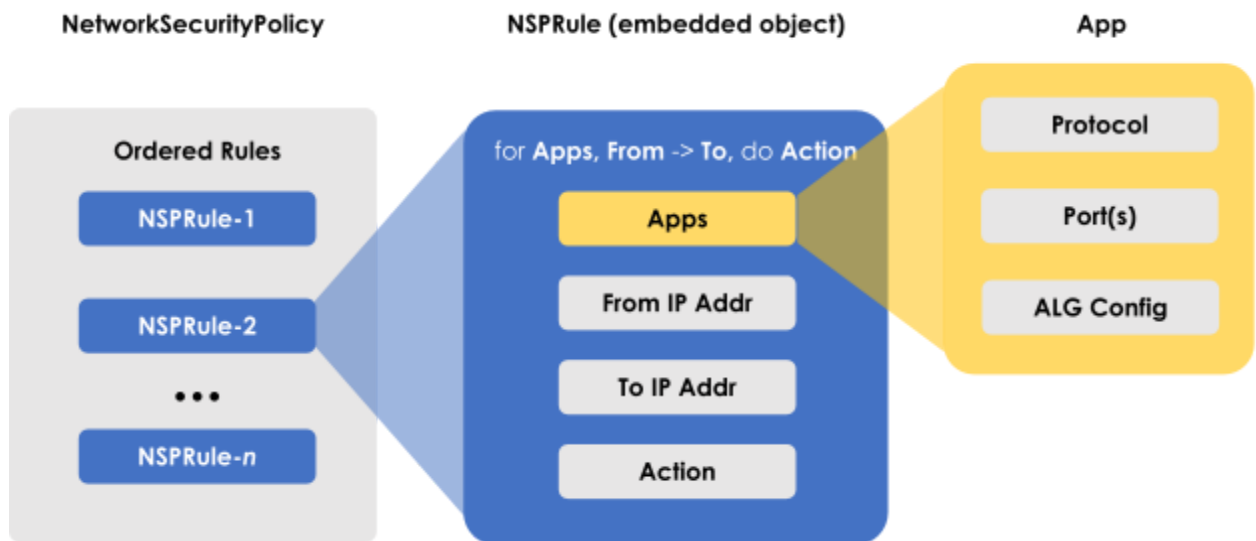


Figure 3. PSM primary firewall objects: NetworkSecurityPolicy, NSPRule, and App

## Apps, Network Security Policy

In PSM terminology, an *App* is a service defined either by a protocol/port pair, or by an application level gateway (ALG) for any of several predefined apps. A *Network Security Policy* is a collection of firewall rules governing App connectivity.

## Key PSM Objects

Table 3 contains sample key PSM objects. For a complete list please refer to the REST API online help available through the PSM GUI.

Object	Description
Distributed Services Entity	<ul style="list-style-type: none"> <li>Entity is a synonym for switch identified by switchname, automatically assigned when a DSS is admitted. Each DSS contains two Distributed Services Modules (DSM).</li> </ul>
Tenant	<ul style="list-style-type: none"> <li>Individual tenant, supporting future multi-tenancy. Currently only “default” is supported.</li> </ul>
VRF	<ul style="list-style-type: none"> <li>Virtual Routing and Forwarding. Collection of subnets.</li> <li>Identified by name</li> </ul>
Network	<ul style="list-style-type: none"> <li>A Network object represents a subnet for which security policy is defined and enforced</li> <li>Identified by a name</li> <li>Contains a VRF, a VLAN ID and the associated ingress/egress security policies to be enforced.</li> </ul>
Security Policy	<ul style="list-style-type: none"> <li>Stateful firewall security policies, defined between network endpoints (IP:port:protocol or apps)</li> <li>Identified by policy name</li> <li>Contains one or more firewall rules</li> </ul>
App	<ul style="list-style-type: none"> <li>Describes the networking specification of an application, service or traffic</li> <li>Identified by name</li> <li>Contains either: <ul style="list-style-type: none"> <li>List of Ports/Protocols</li> <li>Application Layer Gateway (ALG)</li> <li>Both</li> </ul> </li> <li>An App object is subsequently referred to by a SecurityPolicy object</li> </ul>

*Table 3. PSM objects (part 1/2)*

Object	Description
Firewall Export Policy	<ul style="list-style-type: none"> <li>• Syslog destination for firewall logs</li> </ul>
Firewall Profiles	<ul style="list-style-type: none"> <li>• Provides the ability to modify session idle timeout (stateless) and TCP timeout (stateful)</li> </ul>
Alert Policies	<ul style="list-style-type: none"> <li>• A collection of conditions that trigger operator alerts of a given severity type. Triggered alerts are then sent to designated syslog destinations</li> </ul>
Event Policies	<ul style="list-style-type: none"> <li>• The Event stream captures all configuration changes as well as system state changes. All Events can be streamed to designated syslog destinations.</li> </ul>

Table 3. PSM objects (part 2/2)

## Labels

Each object can be associated with one or more *labels* that can be used to refer to a group of objects, which is a very effective way to enable “administration at scale”.



**Note:** Labels that begin with "io.pensando." are reserved for system use, and cannot be created or modified by the user. If the user attempts to create or modify an object's labels with a system label, the label will be silently removed from the user configuration.

# PSM Installation



**Note:** Before installation, see the Release Notes for the minimum resource requirements needed to operate the PSM, as well as the minimum supported versions of virtualization platforms, AOS-CX and Aruba Fabric Composer in conjunction with a given PSM release.

The AMD Pensando PSM software is installed on a virtualized compute infrastructure based on VMware ESXi™ or KVM QEMU emulator. The PSM is delivered as an OVA package for ESXi, or a QCOW2 image for KVM. The recommended configuration to ensure high availability is to install three instances of the PSM software on three physical servers; a single-node PSM cluster configuration is also supported for smaller deployments. See the *Release Notes* for the maximum number of DSSes supported per PSM configuration.

## Storage Considerations

The storage requirement when firewall logs are exported to the PSM shown in Tables 3 and 3a is based on the assumption that the PSM is receiving 1k logs per second.

With increased firewall logs ingestion rates, it is recommended to take into consideration future requirements when initially sizing necessary disk space. However, extra space can be added to the PSM VMs by adding an additional disk. Follow the specific instructions for the hypervisor the PSM is deployed on.

To estimate current and future storage requirements over time for maximum flow log ingestion rates for either a 3-node or 1-node PSM cluster, use Table 4:

Cluster Size	Ingestion Rate	Total Storage			
		3 Days	7 Days	14 Days	30 Days
3 nodes	Sustained 10k LPS	1.04 TB	1.66 TB	2.76 TB	5.28 TB
1 node	Sustained 2k LPS	206.16 GB	320 GB	516.4 GB	967.6 GB

Table 4. Flow log storage estimator

## Data Retention

The PSM has the following retention policy, which is currently not configurable:

- **Events:** retained for 10 days
- **Audit Logs:** retained for 30 days
- **Metrics:** retained for a variable period of time based on roll-up
  - 1 day with 30 second granularity
  - 5 days with 5 minute aggregation
  - 30 days with 1 hour aggregation
  - 1 year with 1 day aggregation
- **Firewall Logs:** retained for up to 30 days (subject to disk space availability, otherwise oldest logs are overwritten).



## PSM Installation on ESXi

### Notes:

- Reserve the amount of memory indicated in the “Installation Requirements” section of the Release Notes for the VM.
- Dedicate a VMFS partition to the VM.
- Live migration (VMware vSphere® vMotion®) is supported.
- VMware HA is supported.

The PSM installs as a virtual appliance (OVA format file), deployed through VMware Virtual Center (vCenter). The PSM deployment depends on vApp and requires vCenter for installation.

1. Log in to vCenter. Locate the ESXi host to install a PSM node on and select “Deploy OVF Template” from the Action button.
2. Specify the URI or Local File name of the PSM OVA file `psm.ova`.
3. Specify the PSM VM name.
4. Under the storage section, select Thick Provision.
5. Specify the OVA properties: hostname, IP address, etc.



**Note:** changing PSM cluster node IP addresses after bootstrapping is not supported.

- a. If using DHCP, leave the IP address blank, and configure a static MAC address-to-IP binding (reservation) for this host in the DHCP server.
  - b. It is strongly recommended that static IP addresses be used.
  - c. Under Password, specify the SSH/console password.
6. Review details. Click “Next” to accept the warnings about advanced configuration and lack of Publisher certificate<sup>2</sup>.
  7. Start the VM in vCenter once the OVA deployment status shows “Completed”. The boot process will untar and install the PSM distribution from a read-only partition.
  8. When the VM comes up, verify that the hostname has changed to what was specified in the OVA properties above and is not “localhost”. If “localhost” appears, then contact Technical Support, as this indicates that the initialization did not complete successfully.
  9. Login to the PSM as user `root`, with the password specified in the OVA properties above (if one was not defined in the OVA properties, the default password is `centos`). If a non-default root password is configured, it may take 1 to 2 minutes for the password

<sup>2</sup> In this release, certificates are self-signed, triggering a warning. This will be changed to authority-signed in a future release.

to take effect after the login prompt becomes available. If this is a concern, make sure network access to the VM is disabled until the password has been reset.

**Note:** Deploying a 3-node cluster involves importing the `psm.ova` file once for each VM instance. The number of imports can be reduced by cloning the first VM as a template, and then deploying subsequent VMs from the template. If taking this approach, follow these steps:

10. Create the first VM from the `psm.ova` file.
11. In vCenter, choose “Clone as Template to Library” to save the VM as a Template (.vmtx) file. Be sure to give the VM a unique name.
12. Select the new VM in vCenter. Select the “Configure” Menu item. Expand “Settings” and select “vApp Options”. Scroll down to “Properties”. Click the radio button for “hostname” and the “Set Value” action to change the hostname to a unique value (typically corresponding to the VM unique name).
13. If applicable, apply any network-specific settings that may have been used in deploying the original VM from the `psm.ova` file.
14. Start the new VM and verify that the VM name and network setting are as intended.
15. Although single-node clusters are supported for smaller production deployments, it is recommended to configure a three node PSM cluster—see the *Release Notes* for the number of DSSes supported in each configuration.

Repeat the deployment of the OVA for the second and third nodes and assign the IP addresses before proceeding further.

By default, the PSM’s `autoadmit` parameter is set to `True`. This means that once a DSS is pointed to a PSM to register, it will be immediately admitted into the cluster. Setting its value to `False` provides additional security, as it requires an operator to manually admit each DSS.

## PSM Installation on KVM

### Notes:

- Ensure no memory oversubscription of VMs running on the same KVM host.
- Dedicate a disk or a Logical Volume to the VM.

1. Verify the available vCPU, memory, and disk resources:

```
$ virsh nodecpumap
CPUs present:    72
CPUs online:     72
$ virsh nodememstats
```

```
total : 131520292 KiB
free : 40785484 KiB
$ df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/centos-root    50G       1.8G   49G   4% /
devtmpfs                  32G         0   32G   0% /dev
tmpfs                     32G         0   32G   0% /dev/shm
tmpfs                     32G       8.9M   32G   1% /run
tmpfs                     32G         0   32G   0% /sys/fs/cgroup
/dev/sda1                 1014M     145M   870M  15% /boot
/dev/mapper/centos-home    57G       33M    57G   1% /home
tmpfs                     6.3G         0   6.3G   0% /run/user/1000
```

2. Verify that the necessary system packages are already installed:

```
$ sudo yum install qemu-kvm qemu-img qemu-system virt-manager
virt-install libvirt libvirt-python libvirt-client bridge-utils -y
```

3. Obtain the PSM qcow2 image.

**Note:** the qcow2 image file will be modified by KVM; make a backup copy of the original image if you need to have a master copy.

4. Use the following command to deploy the PSM node onto KVM, substituting the actual path to the qcow2 image for the highlighted portion:

```
$ virt-install --import --name PSM1 --virt-type kvm --cpu
host-passthrough --os-variant rhel7.6 --ram 32768 --vcpus 12
--network=bridge:br0,model=virtio --disk
path=/home/user/PSM.qcow2,format=qcow2,bus=scsi --controller
scsi,model=virtio-scsi --nographics
Starting install...
Connected to domain PSM1
Escape character is ^]
[ 0.000000] Initializing cgroup subsys cpuset
[ 0.000000] Initializing cgroup subsys cpu
[ 0.000000] Initializing cgroup subsys cpuacct
```

—snip—

5. A console connection will activate while the VM is booting. Wait for the login prompt.
6. Log in with user ID `root` and password `centos`.
7. Run the `config_PSM_networking.py` command to set the hostname and password, and configure the static IP for this PSM VM. (Use `config_PSM_networking.py -h` to see all available parameters.)

Example: set the hostname to `PSM1`, root password to `pensando123`, and configure static IP, gateway, and DNS addresses:

```
[root@localhost ~]# config_PSM_networking.py -m PSM1 -p pensando123
-a static -i 203.0.113.49 -n 255.255.252.0 -g 203.0.113.1 -d
203.0.113,8.8.8.8
```

8. Use `ip addr` to verify the newly configured static IP address. Test logging out and logging back in to verify that the new hostname and password have taken effect.
9. Once everything has been verified, power off the VM:

```
$ poweroff
```

10. To deploy a PSM with more than one interface (as in this example), use the `virsh` command to add the second interface. The below example assumes the VM name is `PSM1` and the second interface is attached to bridge `br54`.

```
# virsh attach-interface PSM1 bridge br54 --model virtio --config
```

11. Power on the VM and log in with the default username `root` and password `centos`.

```
# virsh start PSM1
```

12. The second interface will be detected but unconfigured. The first interface should be configured already as part of step 1. Proceed to configure the second interface (if

required by your topology). Follow the steps in the next section to finish bootstrapping the PSM.

## Bootstrap the PSM Cluster

Before the PSM cluster can be administered, it must be initialized via the `bootstrap_PSM.py` utility. (This utility is in `/usr/local/bin`, which is already in the root user's command path.) Below are some usage examples. The command `bootstrap_PSM.py -h` will show all parameters that can be specified.

Determine the IP address assigned to each PSM VM that has been deployed. In ESXi deployments this can be obtained from vCenter, or from within a CentOS VM with the command `ip addr`. This address is required when launching the `bootstrap_PSM.py` utility and is provided through the `-v` option.

**Note:** A PSM VM should have a single L3 interface that it uses to communicate to other PSM VMs as well as its DSSes. The IP address of this interface should be used in bootstrapping the cluster. The default IP route should point to this interface.

The `bootstrap_PSM.py` utility can be used to provide configuration information to the PSM. The following example provides a cluster name, a domain name, and the address of an NTP server, and activates automatic DSS admission. (If a PSM node has multiple interfaces, use the IP address of the PSM interface that will be used to communicate to DSSes during the bootstrap process.)

**Example:** Bootstrap a 3-node PSM cluster for production, specifying the three nodes' IP addresses. Before executing this command, make sure that all three PSM VMs are already running. The PSM VMs can be deployed using the same OVA, but must have unique IP addresses. The bootstrap script only needs to be executed on one of the nodes.

```
# -distributed_services_switch -clustername Demo -domain training.local
-ntpserver 10.29.5.5 -autoadmit True 192.168.71.134 192.168.68.179
192.168.71.49

2021-12-03 00:22:42.786803: * all messages printed to the console will
also be logged to the file: /var/log/pensando/pen-bootstrap-psm.log
2021-12-03 00:22:42.786834: * start PSM bootstrapping process
2021-12-03 00:22:42.786851: * - list of PSM ips: ['192.168.71.134',
'192.168.68.179', '192.168.71.49']
2021-12-03 00:22:42.786860: * - list of ntp servers: ['10.29.5.5']
2021-12-03 00:22:42.786867: * - using domain name: training.local
2021-12-03 00:22:42.786874: * - auto-admit dse: True
2021-12-03 00:22:42.786881: * - cluster-recovery-keys file path: None
2021-12-03 00:22:42.786888: * - accept EULA: False

2021-12-03 00:22:42.786897: * checking for mandatory
distributed_services_switch option
2021-12-03 00:22:42.786904: * checking for reachability
2021-12-03 00:22:46.790891: * connectivity check to 192.168.71.134
passed
2021-12-03 00:22:50.794810: * connectivity check to 192.168.68.179
passed
2021-12-03 00:22:54.798926: * connectivity check to 192.168.71.49 passed
—snip—
```

### Notes:

1. The cluster name must start and end with an alphanumeric character, and can only contain alphanumeric characters, dash, underscore and period.
2. If the PSM is not connected to the Internet or can't resolve IP names for other reasons, specify the IP addresses of the NTP servers (as shown in this example) instead of their FQDNs.

If everything completes successfully, the message below will be seen in the log:

```
—snip—
2021-12-03 00:26:17.707248: * PSM bootstrap completed successfully
2021-12-03 00:26:17.707280: * you may access PSM at
https://192.168.71.134
```

The PSM browser GUI and REST API should now be available at any of the PSM addresses. Note that there is no virtual IP address for the 3-node cluster. If this is desired, a load balancer should be installed in front of the cluster.

If an `admin` password is not specified when bootstrapping the PSM cluster, the default is `Pensando0$`.

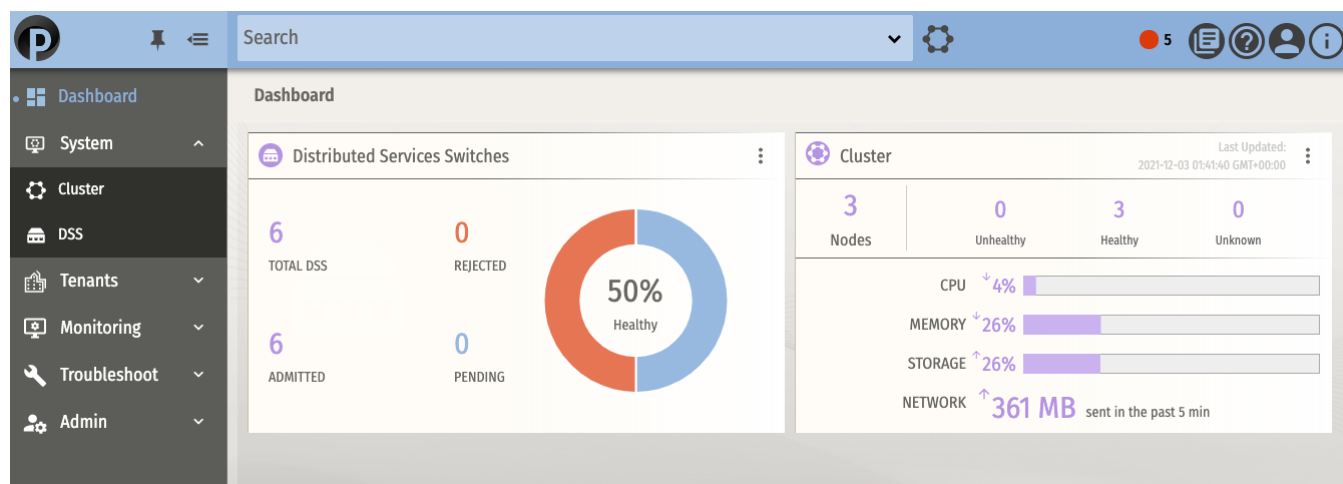
Once this process completes (which may take a few minutes), open a browser connection to the PSM cluster as described in the next section, allowing you to verify that the PSM is healthy.



**Note:** After the PSM is bootstrapped, a copy of its recovery key should be generated and stored in a safe location, in case the PSM needs to be rebuilt later as part of disaster recovery. See [Appendix D](#).


## The PSM Graphical User Interface

The PSM Graphical User Interface (GUI) is accessible via a web browser at `https://PSMaddr`, where *PSMaddr* corresponds to the IP address of any of the PSM cluster nodes. Figure 4 shows a configurable dashboard that offers an overview of the status of the AMD Pensando Distributed Services Platform; the main menu on the left side provides access to the configuration of all supported features.



*Figure 4. AMD Pensando PSM interface, showing the dashboard, with the “System” section of the main menu open; its two subsections, “Cluster” and “DSS”, are visible and selectable.*

## Online Help

Detailed and comprehensive online help is offered in a context-sensitive manner for each page in the PSM GUI, accessible through the help icon in the upper right-hand corner: .

The help icon is context sensitive, showing information related to the currently displayed GUI elements. For example, clicking the help icon while in the Networks overview will display descriptive help and examples on how to create a Network object, as shown in Figure 5. Similarly, clicking the help icon while in the Monitoring -> Alerts & Events view will show help on configuring Alert Policies.



The screenshot shows the 'Networks Overview' section of the PSM interface. It features a table with 7 columns: VRF, Name, VLAN, Ingress Policy, and Egress Policy. The table lists several network configurations. To the right, the 'PSM Help' window is open, displaying information about the 'Network' section and how to 'Add Network'.

VRF	Name	VLAN	Ingress Policy	Egress Policy
vrf600	601	601	SGpolicy	SGpolicy
vpc1	net30	30	SGpolicy	SGpolicy
VRF401	411	411	SGpolicy	SGpolicy
vpc1	kapil-vlan11	11	kapil-ing	kapil-egr
vpc1	505	505	SGpolicy	SGpolicy
vpc1	kapil-vlan20	20	kapil-ing	kapil-egr
vpc1	720	720	SGpolicy	SGpolicy

**Network**  
The Network section provides two functions:  
1. Add a network.  
2. Network overview: Provides a detailed table view of the existing networks created along with the VLAN ID and the Orchestrator information.

**Add Network**  
The Network refers to a DSS network that the PSM creates for network traffic. A network is created with the following attributes:  
name | Name of the object, unique within a Namespace for scoped objects. Must start and end with alpha numeric and can have alphanumeric, -, \_ . Length of string should be between 2 and 64.

Figure 5. Example of PSM help

Online Help windows can be easily undocked, redocked, resized, or closed.

Online Help has its own presentation context, so it does not need to be closed prior to subsequent operations; selecting different items from the left-hand-side Navigation pane will automatically display the corresponding Online Help information.

## Searching

The easy-to-use search facility is accessed from the search bar at the top of the screen.

The screenshot shows the search bar at the top of the PSM interface. The search bar contains the text 'ae-s7'. Below the search bar, a list of search results is displayed, including 'Alert', 'AuditEvent', 'DistributedServiceCard', 'Endpoint', 'Event', 'Host', 'TechSupportRequest', 'UserPreference', and 'Workload'. The results are sorted by relevance, with 'Event' having the highest count (298).

ae-s7

Keyword: ae-s7

Search

Alert	57
AuditEvent	15
DistributedServiceCard	2
Endpoint	4
Event	298
Host	2
TechSupportRequest	1
UserPreference	1
Workload	4

Figure 6. PSM search facility

In the example in Figure 6 above, doing a free form text search for the string “ae-s7” shows a summary of the various objects where that string appears, along with a count of the number of occurrences for each object type.



Figure 7. Accessing Advanced Search

Clicking on the downward arrow on the right hand side of the text box (shown in Figure 7 above) gives access to the Advanced Search capability shown in Figure 8, where users can search based on object Category, Kind or Tag (arbitrary labels associated to objects):

Search ×

In: Category

☐ Auth
☐ Cluster
☐ Monitoring
☐ Network

☐ Objstore
☐ Preferences
☐ Rollout
☐ Security

☐ Staging

Is: Kind

☐ Alert
☐ AlertDestination
☐ AlertPolicy
☐ ArchiveRequest

☐ AuditEvent
☐ AuditPolicy
☐ AuthenticationPolicy
☐ Buffer

☐ Cluster
☐ ClusterProfile
☐ ConfigurationSnapshot
☐ Credentials

☐ DistributedServiceSwitch
☐ Event
☐ EventPolicy
☐ License

☐ Network
☐ NetworkInterface
☐ NetworkSecurityPolicy
☐ Node

☐ Role
☐ RoleBinding
☐ Rollout
☐ RolloutAction

☐ SnapshotRestore
☐ TechSupportRequest
☐ Tenant
☐ UIGlobalSettings

☐ User
☐ Version
☐ VirtualRouter

Tag: Label +


CLEAR

SEARCH

Figure 8. Advanced Search

All the keywords used in Advanced Search can also be typed directly into the search bar to avoid having to bring up the Advanced Search tab.

## Global Icons

The GUI makes use of common/global icons for many actions, regardless of context, such as “Edit” or “Delete”, which can be used to edit fields, add labels, or delete objects such as “VRF”, “Networks”, and “Security Policies”, as shown in Figure 9 : 




Figure 9. Edit and Delete icons

Many of the tables displayed in the GUI can be exported as CSV or JSON text files, as shown in Figure 10:

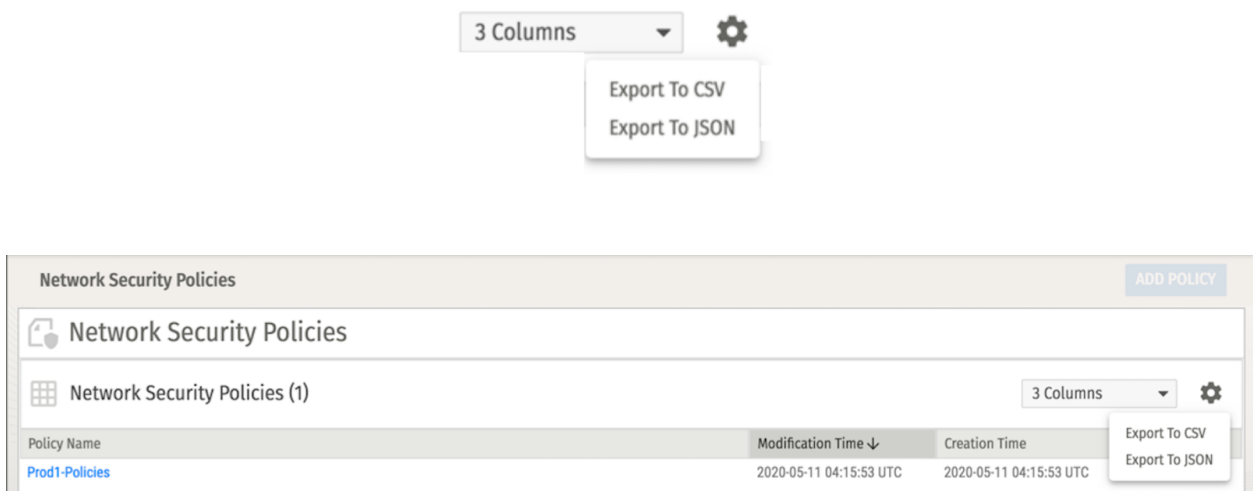


Figure 10. Example of how a table can be exported in CSV or JSON format

## Server Certificate

By default, the PSM uses a self-signed certificate to authenticate itself to browser-based GUI or REST clients.

Sites may instead provide a custom key and certificate for the PSM to use. If the root certificate authority (CA) of the custom certificate is included in either the browser or the client hosts' trusted root CA certificate list, warning messages related to certificate validity will no longer be shown when accessing the PSM cluster login page.

The two supported encoded key formats are RSA and ECDSA. To change the PSM certificate, click "Admin" --> "Server Certificate". On the top right hand side, click "UPDATE". Enter the key and certificate in Privacy Enhanced Mail (PEM) format and then click "Upload" to apply the change, as shown in Figure 11.

**Note:** *this action will not disrupt existing connections, even if they were established with a previous certificate.*



Update Server Certificate

UPDATE CLEAR FORM

⚠ Updating the server certificate changes the HTTPS certificate used by the server. This action can not be reversed!

Key

Certs

Figure 11. Changing the PSM server certificate

## API Capture

Users of the PSM's REST API for external integration can take advantage of the fact that the PSM GUI itself uses this API, to see examples of it in use. REST API calls sent from the GUI to the PSM as it implements the configurations created by the user can be examined using the API Capture feature.

When the API Capture menu item is selected, a view as shown in Figure 12 appears. Use this screen to browse sample API calls (in the API Capture tab) or a live capture of APIs generated while navigating the GUI (in the Live API Capture tab).

The scope of the live capture tool results is per GUI session; large responses are trimmed down to two records to present the look and feel of the response, rather than the entire response.

The screenshot shows the 'API Capture' interface. At the top right is a 'Disable Live Capture' button. Below the header, there are tabs for 'API Capture' and 'Live API Capture'. The 'Live API Capture' tab is active, displaying a table of captured API calls. The table has columns: Category, Kind, Method, and URL. A search bar is located above the table. The first entry in the table is for a GET request to the URL `./configs/cluster/v1/hosts`. Below the table, the details of the selected API call are shown, including the category, kind, method, url, request, and response.

Category	Kind	Method	URL
cluster	hosts	GET	<code>./configs/cluster/v1/hosts</code>

**category:** cluster  
**kind:** hosts  
**method:** GET  
**url:** `./configs/cluster/v1/hosts`  
**Request:**  
**Response:**

```
{
  "kind": "HostList",
  "api-version": "v1",
  "list-meta": {},
  "items": null
}
```

Figure 12. Example of captured REST API calls

## Create PSM User Authentication Policy and Users

Each PSM will have one or more users defined. Users are assigned roles granting them privileges depending on the tasks they need to perform; during the installation process, an initial user, named `admin` with the password `Pensando0$`, is created with full administrator privileges.

A different name for the initial user as well as a custom password can be provided as parameters to the `bootstrap_PSM.py` utility used to initialize the cluster, as described in the section [“Bootstrap The PSM Cluster”](#).

### User Authentication Policy

The PSM supports Local (i.e., username and password), LDAP and RADIUS authenticators, as shown in Figure 13. Creation of authenticators should be done early in the system setup

process. Once two or more authenticators are created, they can be re-ordered dynamically to specify the priority with which they should be applied.

See the *PSM LDAP Configuration Guide* for specific configuration details.

The screenshot displays the 'Auth Policy' configuration page. At the top, there are 'Refresh' and 'New Authorization Token' buttons. The main section is titled 'Authentication Policy' and contains three authenticators, each with a large blue number indicating its priority:

- Local (Priority 1):** The first authenticator in the list.
- Radius (Priority 2):** The second authenticator. It has a 'CONFIG' tab and a 'SERVERS' tab. Under 'SERVERS', there is a table with columns for 'NAS ID', 'psm-1', and a server address '10.29.65.25:1812 pap'.
- LDAP (Priority 3):** The third authenticator. It has three sub-sections:
  - CREDENTIALS:**
    - \* Bind DN: PENSANDO\Administrator
    - \* Bind Password: \*\*\*\*\*
  - SCOPE:**
    - \* Base DN: DC=pensando, DC=local
    - \* User Object Class: user
    - \* Group Object Class: group
  - ATTRIBUTE MAPPING:**
    - \* User: sAMAccountName (Full name:)
    - \* Group: memberOf (Email:)
    - Tenant:

*Figure 13. Managing user authentication policy*

To create an LDAP Authenticator, click the “CREATE LDAP AUTHENTICATOR” button. Active Directory (AD) and OpenLDAP providers are supported.

Configure *Credentials*, *Scope* (which controls user and group entry search) and *Attribute* mapping (which contains the name of the attributes that should be extracted from the LDAP entry of the user, such as full name and email address) as appropriate, ensuring all required (\*) fields are properly filled, as in Figure 14:

Auth Policy

Refresh New Authorization Token

### Authentication Policy

Authenticators

User credentials will be tested against the following authentication methods until one of them returns successfully.

#### LDAP

SAVE CANCEL

CREDENTIALS	SCOPE	ATTRIBUTE MAPPING
* Bind DN: <input type="text" value="cn=admin user,ou=users"/> <input type="button" value="Test Bind"/>	* Base DN: <input type="text" value="dc=pensando,dc=io"/>	* User: <input type="text" value="uid"/> Full name: <input type="text" value="cn"/>
Bind Password: <input type="password" value="password"/>	* User Object Class: <input type="text" value="inetOrgPerson"/>	* Group: <input type="text" value="ou"/> Email: <input type="text" value="mail"/>
	* Group Object Class: <input type="text" value="posixGroup"/>	Tenant: <input type="text"/>

#### SERVER CONFIGURATION

+ ADD

Server Port:  ☒ Use TLS

☒ Verify Server Certificate

Server Name:

```
-----BEGIN CERTIFICATE-----
MIIDHTCCAm2gAwIBAgIJAMnL/d3yxAhMA0GCSqGSIb3DQEBCwUAMFkxCzAJBgNV
BAYTAihyMRUwEwYDVQHDxhZWNpdHkxHDAaBgNVBAoMEORIZmF1bHQg
Q29tcGFueSBMdGQxFTATBgNVBAMMDG9wZW5sZGFwdm0wMTAeFw0xOTcwMDMwOTMy
MDdaFw0yMDcwMDMwOTMyMDdaFkxCzAJBgNVBAYTAihyMRUwEwYDVQHDxhZWNpdHkx
HDAaBgNVBAoMEORIZmF1bHQgQ29tcGFueSBMdGQxFTATBgNVBAMM
```

Figure 14. LDAP configuration

Once saved, the values should be visible, as shown in Figure 15. The order of the various authenticators can be changed (using the small arrows on the right hand side).

Auth Policy

Refresh New Authorization Token

### Authentication Policy

Authenticators

User credentials will be tested against the following authentication methods until one of them returns successfully.

#### LDAP

CREDENTIALS	SCOPE	ATTRIBUTE MAPPING
* Bind DN: <input type="text" value="cn=admin user,ou=users,dc=pensando,dc=io"/>	* Base DN: <input type="text" value="dc=pensando,dc=io"/>	* User: <input type="text" value="uid"/> Full name: <input type="text" value="cn"/>
* Bind Password: <input type="password" value="*****"/>	* User Object Class: <input type="text" value="inetOrgPerson"/>	* Group: <input type="text" value="ou"/> Email: <input type="text" value="mail"/>
	* Group Object Class: <input type="text" value="posixGroup"/>	Tenant: <input type="text"/>

SERVERS: ☒ 192.168.71.153:389

#### Local

1 ↓

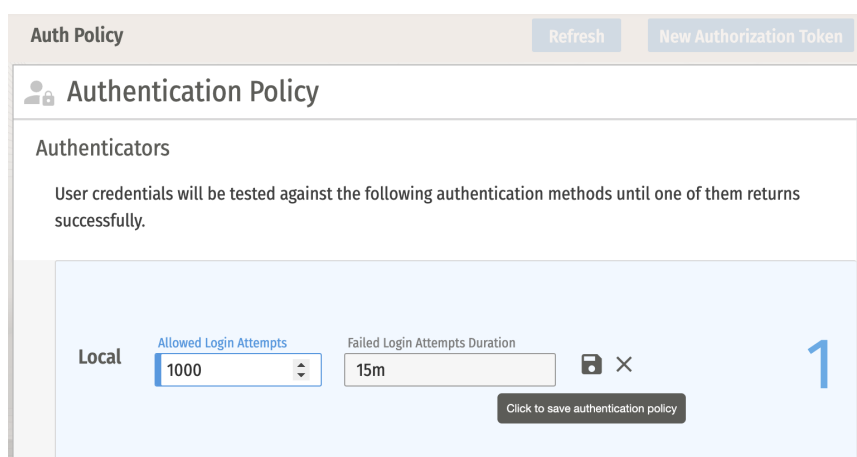
2 ↑

Figure 15. Changing authentication order

## Local User Lockout Policy

By default, any local account on the PSM will be locked out after 10 incorrect login attempts within a 15 minute interval. This policy applies to all PSM local accounts, including admin accounts. Locked accounts can be unlocked by an admin account.

To change the values for the maximum number of attempts, navigate to *Admin → Auth Policy*. Click the Edit icon (✎) on the Local Authenticators line and set a new value. Click on the Save (💾) icon when done.



The screenshot shows the 'Auth Policy' configuration page. At the top, there are buttons for 'Refresh' and 'New Authorization Token'. Below this is the 'Authentication Policy' section, which includes a description: 'User credentials will be tested against the following authentication methods until one of them returns successfully.' The 'Authenticators' section is highlighted in blue. It contains a table with one row for 'Local' authentication. The 'Allowed Login Attempts' field is set to '1000' and the 'Failed Login Attempts Duration' field is set to '15m'. There is a 'Click to save authentication policy' button at the bottom right of the table. A large blue number '1' is visible on the right side of the table.

Figure 16. Changing lockout policy

It is strongly recommended that sites create a backup admin local account on the PSM in case all other admin accounts become locked.

## Role-Based Access Control (RBAC)

The User Management menu gives access to the RBAC Management screen, shown in Figure 17, which allows management of users, roles, and the association of users to roles. Each action can be selected with the drop-down menu on the top right corner. It is recommended to first create one or more Roles, followed by one or more Users, and then create the corresponding associations between Roles and Users (rolebinding).



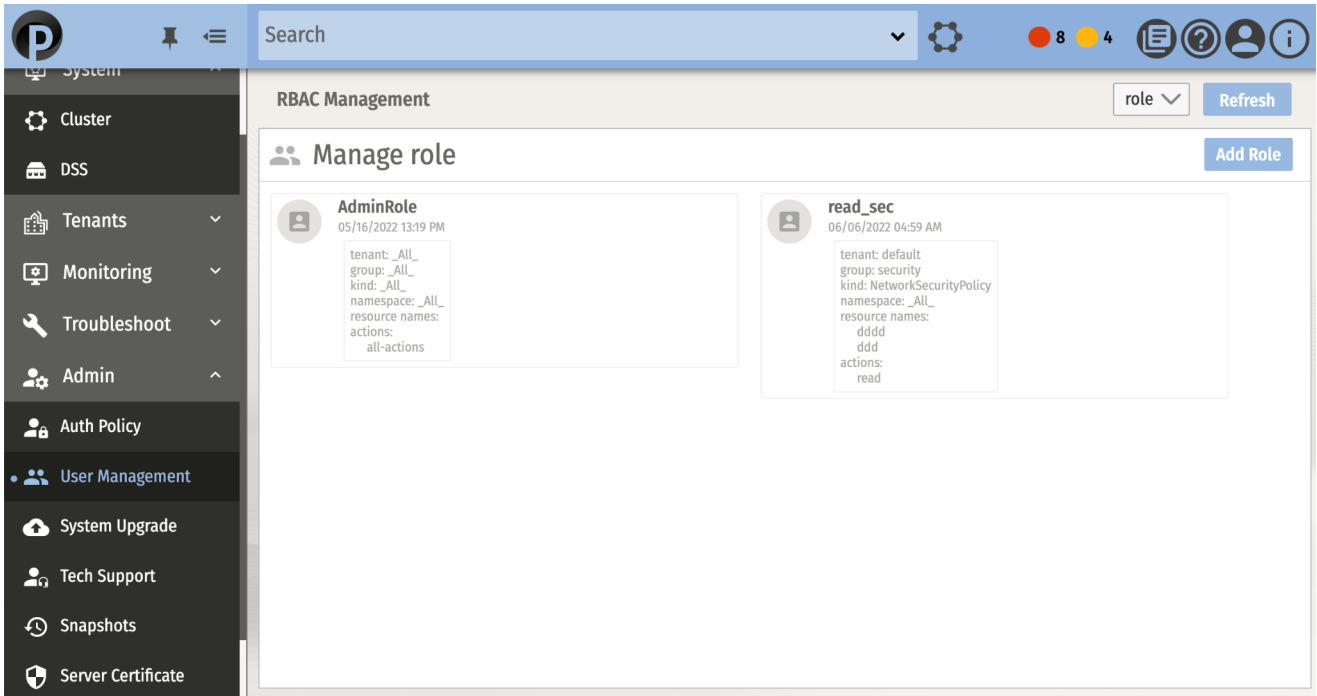


Figure 17. RBAC management

## Roles

*Roles* are created to control access to classes of features by sets of users. Roles can have scope over various objects, which are grouped by the PSM in the categories:

- Auth
- Cluster
- Diagnostics
- Monitoring
- Network
- Objstore
- Preferences
- Rollout
- Security
- Staging

**RBAC Management**

**Manage role**

Name:

Permissions:

Configurations:

Group:

Kind:

+ AND

*Figure 18. RBAC roles*

As shown in Figure 18 above, various kinds of management aspects are available for a given Group. Once one is selected, access to actions can be added or removed, as shown in Figure 19:

- Create
- Read
- Update
- Delete

**RBAC Management**

**Manage role**

Name:

Permissions:

Configurations:

Group:

Kind:

Actions:

☐ All Actions

☐ Create

☐ Delete

☐ Read

☐ Update

+ AND

*Figure 19. Assigning actions to a group*

## Role Binding

Once a Role is created, a corresponding rolebinding is automatically created. Rolebindings allow users to be flexibly mapped to various sets of roles. Figure 20 shows the view to modify

a “rolebinding” that allows to associate any of the users defined in the system (in the left list titled Available) with the Role specified in the form. Users successfully associated with the Role appear in the right list titled Selected.

The rolebinding can be also specified using the “Group” attribute value configured in the LDAP authenticator “Attribute Mapping” section and retrieved from the LDAP user entry. This is the distinguished name of the LDAP group entry to which a user belongs.

**Manage rolebinding**

Name: AdminRoleBinding

Role: AdminRole

User Groups: cn=admin,ou=groups,dc=pensando,dc=io

Users:

Available	Selected
Search by name	Search by name
	admin

Navigation buttons: >, >>, <, <<

Buttons: Cancel, Save

Figure 20. Rolebinding

## System Upgrade

The PSM platform nodes can be upgraded using the “System Upgrade” option, available in the Admin menu. The process is broken into two tasks:

- Download the upgrade bundle from your vendor’s support site.
- Upgrade the PSM cluster (approximately 10 minutes per cluster node).

When upgrading a fabric, make sure that AFC (if present) is upgraded first, then the PSM cluster nodes, and finally the CX 10000 switches.

See the PSM release notes for supported upgrade paths for AOS-CX, PSM and AFC.

## AFC Upgrade

Upgrading AFC is documented in the *Aruba Fabric Composer Installation Guide*.

## PSM Upgrade

### Upload PSM Upgrade Bundle

1. Download the PSM upgrade bundle. This bundle is part of the software release package, and will have a file name like `psm_upgrade_bundle_release_number.tar`. Once this file is downloaded, upload it to the PSM. Click on the **ROLLOUT IMAGES** button in the Admin > System Upgrade view.

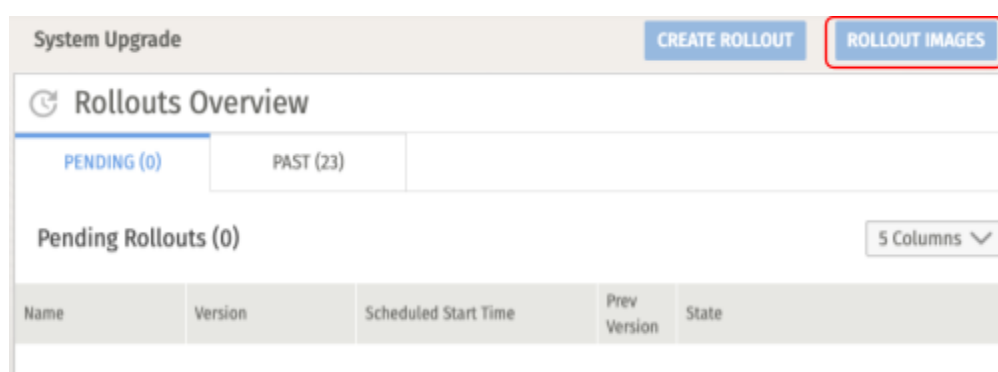


Figure 21. Rollouts Overview screen

2. Click on “Upload Image File”:


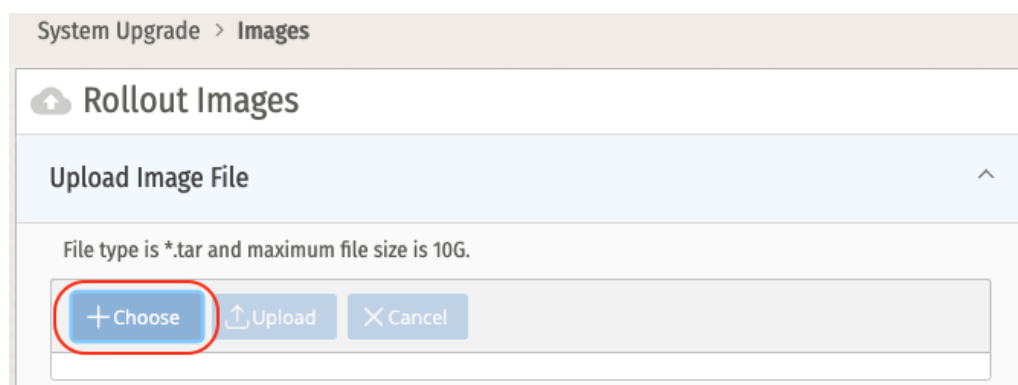


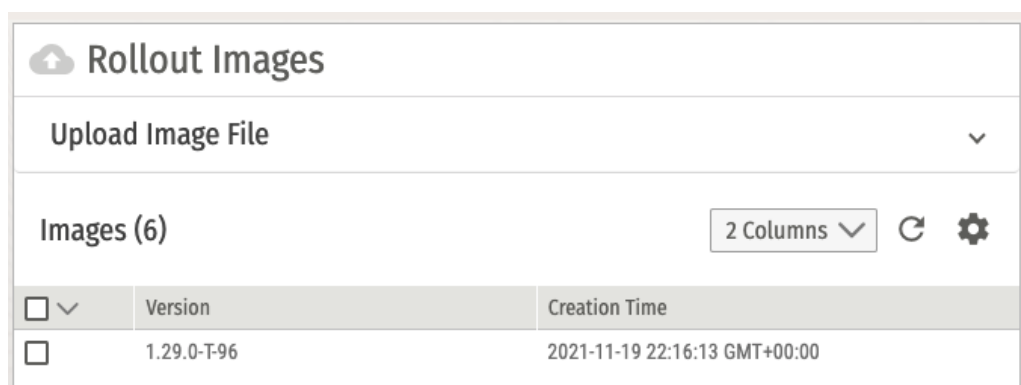
Figure 22. Bring up image file selector

- Click on “Choose” and select the PSM upgrade image to be uploaded.



*Figure 23. Choose rollout image file*

Once uploaded, the new bundle will show in the Images repository, and will now be available to be included in a rollout configuration.



*Figure 24. New upgrade image is available to select.*

## Create Rollout

Once the upgrade bundle is uploaded to PSM, create a rollout. The rollout can be scheduled to start at a particular date/time or immediately.

- Return to the Admin -> System Upgrade screen (as in Figure 21), and click “Create Rollout”. The rollout creation form will appear, as shown in Figure 25.

**System Upgrade** SAVE ROLLOUT CANCEL

### Rollouts Overview

**PENDING (0)** **PAST (4)**

Unique Rollout Name  
test\_rollout

version  
1.29.0-T-96

Start Time  
11/19/2021 22:30 ☒ Schedule Now

End Time ☐ Specify rollout duration

*Figure 25. Specify name, version, and time to begin upgrade.*

- **Unique Rollout Name:** Name of the rollout; can be any arbitrary name
  - **Version:** Choose the version to upgrade to. The choices are auto-populated from the image(s) uploaded to the PSM.
  - **Start time:** Time when the rollout should start.
    - **Schedule Now:** Select this to start the upgrade as soon as “Save Rollout” is applied.
2. Click on **SAVE ROLLOUT** to start the upgrade. If “Schedule Now” was selected, the PSM upgrade will start immediately.

## AOS-CX Upgrade

For upgrading AOS-CX on the switches being managed, consult Aruba documentation at <https://asp.arubanetworks.com/>. High level instructions are provided here for quick reference.

1. Copy image to non-current boot bank on the DSS using the command

```
switch# copy scp://server/path/switch_image.swi  
(primary|secondary) (vrf mgmt)
```

- a. The existing image in the primary or secondary location will be overwritten.
- b. Use the appropriate VRF.

2. Verify that the image has been copied to the DSS:

```
switch# show images
```

3. Confirm if any device component updates are required:

```
switch# boot system (primary|secondary)
```

Enter **n** when prompted to continue.

- a. If any non-failsafe device updates are required, enter configuration mode and type in “allow-unsafe-updates 30” before proceeding to the next step.

4. Start the upgrade using the command

```
switch# boot system (primary|secondary)
```

If AFC is used to upgrade the switches, this step can be skipped.

Consult Aruba AOS-CX documentation for information on related topics such as VSX Live Upgrade.

# Configuration Snapshots

The PSM configuration can be saved and later restored via *Snapshots*. A PSM configuration Snapshot can be created by going to Admin->Snapshots, and clicking on **SAVE A CONFIG SNAPSHOT**, as shown in Figure 26.

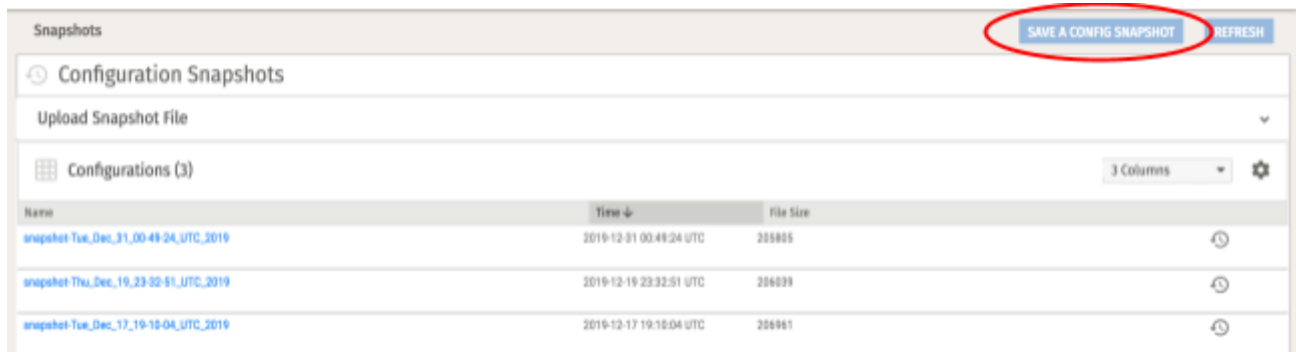


Figure 26. Configuration Snapshots screen

You may specify a snapshot name, in which case the snapshot will be saved in the local file `name.gz`; otherwise the snapshot will be saved with a filename based on the current day and time.

To restore the state of the PSM to an earlier configuration, click the “Restore config” icon corresponding to the snapshot, as shown in Figure 27.

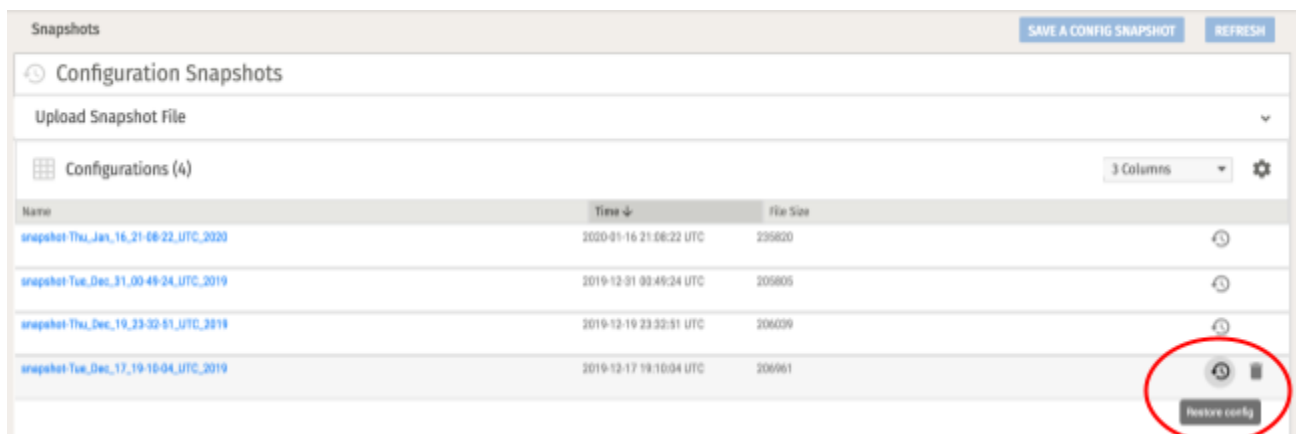


Figure 27. Accessing the Restore config icon

The PSM will be unavailable during the configuration restore process.

Configuration backups are backward compatible: a configuration snapshot taken with an older version of the PSM software can be restored on a later version of the PSM. Features that were not supported when the snapshot was taken will not be configured upon restore.



## Associating a DSS with the PSM

A DSS can be associated with a PSM using:

- The Aruba AOS-CX CLI
- ZTP (Zero Touch Provisioning), by passing PSM network addresses
- Aruba Fabric Composer (AFC) for day 0 provisioning

**Note:** Before associating a DSS with a PSM:

1. Ensure that there will be Layer 3 connectivity between the DSS and the PSM through the management or inband interface. Make sure that the PSM service ports are open for any firewall in-between; see the list of PSM-to-DSS ports in [Appendix B](#).
2. Make sure that the DSS is not currently associated with another PSM. If you need to move a DSS from one PSM cluster to another PSM cluster, follow the steps in the [“Decommissioning a DSS”](#) section.

### AOS-CX CLI



**Note:** Verify that the time is set correctly on the switch.

1. Enter configuration mode on the switch and type `psm` to get into the PSM configuration section. To show the commands that you can execute, enter the `?` symbol (the `?` will not display on the screen when you enter it).

```
switch# conf t
switch(config)# psm
switch (config-psm)# ?
end      End current mode and change to enable mode.
exit     Exit current mode and change to previous mode
host     Configure the Policy and Services Manager address
list     Print command list
no       Negate a command or set its defaults
show     Show running system information
```

2. Use the `host` command to provide the IP addresses for the three nodes of the PSM to associate the DSS with.

```
switch(config-psm) # host 203.0.113.49 203.0.113.50 203.0.113.51 vrf default
```

3. If connecting to the PSM over the in-band network in the default VRF, it is required to either have the management IP of the switch configured or use the below commands:

```
switch# conf t  
switch(config) # ip source-interface psm ip address vrf vrf name
```

## ZTP

Add the below options to `dhcp.conf`; see the Aruba [Using ZTP to Provision a Managed Device](#) documentation for further information.

```
// Global options:  
  
# Need this for CX 10000 ZTP  
option local-proxy-config code 252 = text;  
option space Aruba code width 1 length width 1 hash size 4;  
option Aruba.image-file code 145 = text;  
option Aruba.conf-file code 144 = text;  
option Aruba.cop-location code 146 = text;  
option Aruba.http-proxy-location code 148 = text;  
  
host cx10000-xyz {  
    ##### DSS  
    hardware ethernet 04:90:81:00:0a:ab;  
    fixed-address 10.30.2.49;  
    option tftp-server-name "10.30.4.188";  
    vendor-option-space Aruba;  
    option Aruba.conf-file "p0-29-running-config.cfg";  
}
```

Add the lines highlighted in the example above to specify the TFTP server IP address and configuration file to use (10.30.4.188 and p0-29-running-config.cfg in the example).

The configuration file is a regular ASCII file containing the PSM IP addresses.

## AFC

Refer to the “PSM Integration” section of the *Aruba Fabric Composer User Guide* for details.

## Verification

Once the association is completed, the DSS should show up in the PSM GUI.<sup>3</sup>

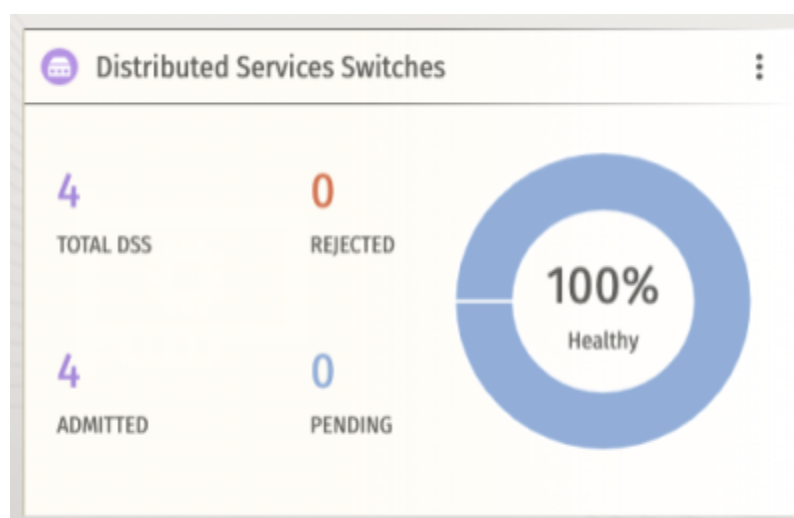


Figure 28. The PSM dashboard DSS status card, showing the number of admitted and pending DSSes. See also the System->DSS view, for a list of DSSes associated with this PSM.

<sup>3</sup> The DSS may need to be manually admitted, depending on the `autoadmit` admission policy selected when bootstrapping the PSM quorum, as discussed [above](#).

The AOS-CX CLI can also be used to verify the existing configuration:

```
switch# show psm

Policy and Services Manager Information

Operational Status: admitted
Host Addresses: 203.0.113.49 203.0.113.50 203.0.113.51
VRF: default
switch#
```

## Decommissioning a DSS



**Note:** Once a DSS is decommissioned, all stateful policies on the DSS will be purged and there will be no stateful inspection/services performed until it is readmitted to a PSM.

This procedure can be used to move a DSS from one PSM cluster to another, by first removing the DSS from the PSM, and then removing the configuration from AOS-CX.


- On the PSM, Navigate to **System** → **DSS**
- For the DSS to be decommissioned, move your cursor over to the right end of the row corresponding to the DSS. You will see a mouseover option that reads “Decommission DSS”:

Distributed Services Switches											
Distributed Services Switches Overview											
Distributed Services Switches (3)											
11 Columns											
<input type="checkbox"/>	Switch Name	DSS-ID	DSM	Switch OS Version	Switch Serial #	Management	Health	Labels	FW Logging Export	Policy Scale	Modification Time
<input type="checkbox"/>	leaf2a	DSS-0490.8100.37dc	DSM 1/1: 0490.8100.383c DSM 1/2: 0490.8100.384c	DL.10.10.0001BG...	US17L6B03E	Mode: Inband IP address: 10.30.22.2				Max 24K R...	2022-05-24 05:47:33 GMT+00:00
<input type="checkbox"/>	leaf2b	DSS-0490.8100.13ce	DSM 1/1: 0490.8100.142e DSM 1/2: 0490.8100.143e	DL.10.10.0001BG...	US14L96031	Mode: Inband IP address: 10.30.2.52				Max 24K R...	2022-05-24 05:55:01 GMT+00:00
<input type="checkbox"/>	leaf1	DSS-0490.8100.0596	DSM 1/1: 0490.8100.05f6 DSM 1/2: 0490.8100.0606	DL.10.10.0001BG...	US12L96017	Mode: Inband IP address: 10.30.2.53				Max 24K R...	2022-05-24 05:33:36 GMT+00:00
Decommission DSS											

Figure 29. The Decommission icon

- Click on this icon, and confirm that you want to decommission the DSS:





**Are you sure that you want to decommission the DSS?**

 Once you decommission DSS leaf1, it can not be undone. DSS reboot is required if you want to re-admit a decommissioned DSS back to same/different PSM.

✓Decommission
✗No




*Figure 30. Confirm decommissioning; stateful services will no longer be active on this switch until it is readmitted to a PSM.*

- Once the DSS is decommissioned, the Health icons will be updated as shown in Figure 31:

Distributed Services Switches											
Distributed Services Switches Overview											
Distributed Services Switches (3)											
Switch Name	DSS-ID	DSM	Switch OS Version	Switch Serial #	Management	Health	Labels	FW Logging Export	Policy Scale	Modification Time	
leaf2a	DSS-0490.8100.37dc	DSM 1/1: 0490.8100.383c DSM 1/2: 0490.8100.384c	DL.10.10.0001BG...	US17L6B03E	Mode: Inband IP address: 10.30.22.2				Max 24K R...	2022-05-24 05:47:33 GMT+00:00	
leaf2b	DSS-0490.8100.13ce	DSM 1/1: 0490.8100.142e DSM 1/2: 0490.8100.143e	DL.10.10.0001BG...	US14L96031	Mode: Inband IP address: 10.30.2.52				Max 24K R...	2022-05-24 05:55:01 GMT+00:00	
leaf1	DSS-0490.8100.0596	DSM 1/1: 0490.8100.05f6 DSM 1/2: 0490.8100.0606	DL.10.10.0001BG...	US12L96017	Mode: Inband IP address: 10.30.2.53				Max 24K R...	2022-05-24 18:13:58 GMT+00:00	

*Figure 31. Health icons show that the DSS is no longer admitted.*

- Click on the Delete DSS (  ) icon to delete the DSS object:

Distributed Services Switches Overview											
Distributed Services Switches (2)											
Switch Name	DSS-ID	DSM	Switch Serial #	Management	Health	FW Logging Export					
CX10K-Support-1	DSS-0490.8100.2cb0	DSM 1/1: 0490.8100.2d10 DSM 1/2: 0490.8100.2d20	FSJ21280015	Mode: Out of Band IP address: 192.168.68.128		test					
CX10K-Support-2	DSS-0490.8100.2080	DSM 1/1: 0490.8100.20e0 DSM 1/2: 0490.8100.20f0	FSJ21280021	Mode: Out of Band IP address: 192.168.69.237		test					

*Figure 32. The Delete option is now available.*

- Confirm that you want to delete the DSS object:

Are you sure you want to delete DSS0490.8100.2080?

This action cannot be reversed

✓ Delete
✗ No

*Figure 33. Confirm removing the DSS from this PSM.*

Once this is done, the DSS will be removed from the PSM cluster. Once you have confirmed the DSS does not show up in the PSM GUI, proceed to the next step.

Distributed Services Switches											
Distributed Services Switches Overview											
Distributed Services Switches (2) <span style="float: right;">11 Columns </span> <span style="float: right;">Search <span style="font-size: 0.8em;">Aa Ab R! </span></span>											
<input type="checkbox"/>	Switch Name	DSS-ID	DSM	Switch OS Version	Switch Serial #	Management	Health	Labels	FW Logging Export	Policy Scale	Modification Time
<input type="checkbox"/>	leaf2a	DSS-0490.8100.37dc	DSM 1/1: 0490.8100.383c DSM 1/2: 0490.8100.384c	DL.10.10.0001BG...	US17L6B03E	Mode: Inband IP address: 10.30.22.2				Max 24K R...	2022-05-24 05:47:33 GMT+00:00
<input type="checkbox"/>	leaf2b	DSS-0490.8100.13ce	DSM 1/1: 0490.8100.142e DSM 1/2: 0490.8100.143e	DL.10.10.0001BG...	US14L96031	Mode: Inband IP address: 10.30.2.52				Max 24K R...	2022-05-24 05:55:01 GMT+00:00

*Figure 34. The removed DSS is no longer listed.*

- After removing the DSS from PSM, log in to AOS-CX and remove the PSM configuration. (If you want to switch the DSS to a new PSM, you can also enter the new PSM addresses at this point.) Save the configuration and reload the switch.

```

switch# show running-config psm
psm
    host 20.3.0.37 20.3.0.38 20.3.0.39 vrf default

switch (config)# no psm
Warning: If PSM is removed, REBOOT will be required

Continue (y/n)? Y

switch (config)# show running-config psm
switch (config)#

switch (config)# copy running-config startup-config
Copying configuration: [Success]
  
```

```
switch# boot system primary
```



**Note:** *If the switch being decommissioned is also being managed as part of an AFC fabric, it may also need to be removed from AFC management.*

## Associating the PSM to Aruba Fabric Composer

Aruba Fabric Composer (AFC) automates CX switch provisioning and management, including coordination with stateful services via the PSM. It can be used to create and provision leaf-spine networks, simplify troubleshooting, and perform life cycle management.

To synchronize these capabilities with the PSM's stateful services, AFC must be configured to be associated with the PSM by providing AFC with the PSM instance's network address. Refer to the AFC documentation for details on this process.

## Firewall Policy Functionality and Configuration

The CX 10000 extends the capabilities of the leaf-spine fabric to natively provide a distributed stateful firewall for east-west traffic, zero trust segmentation, and pervasive telemetry. These built-in security services help minimize attack surface within the data center and block lateral movement in attack attempts.

Functionality supported by the stateful L4 firewall includes:

- Connection tracking:
  - Initial TCP handshake validation
  - Session closing sequence validation
  - TCP reset checks
- Application-level gateways (TFTP, DNS, FTP, MSRPC, SUNRPC, RTSP)
  - The DNS ALG, for example, tracks and correlates DNS queries with corresponding responses; once a response is received and forwarded, the ALG will accelerate aging/closing of the DNS UDP session, resulting in better DNS session scale.
  - Other ALGs follow standard protocol implementation.
- Firewall Log Export to the PSM
- Firewall Log Export to external collectors, in the default VRF context.

Network Security Policies are created, then attached to either networks, VRFs or both. When a policy is attached to a VRF it is internally applied to all networks in that VRF. *Ingress* security policies are applied to traffic entering a host and *egress* security policies are applied to traffic leaving a host. The same policy can be used for both ingress and egress, or different policies for ingress and egress can be used for more granular control.

Policy direction is from a workload/host perspective.

## Considerations

- All traffic leaving the workload/host and entering the switch in host-facing ports (“access” persona) is subject to the egress policy on the switch.
- All traffic destined for the workload/host, entering the switch from the fabric side (“uplink” persona) ports, is subject to ingress policy on the switch.
- The `persona` CLI config option on a port is used to set the appropriate policy (ingress/egress) on the DSM.
- By default, all 10G/25G ports are configured as workload/host facing ports (“access” persona) and the 100G/40G ports are configured as fabric ports (“uplink” persona).
  - If a 40G/100G port needs to be used as a workload/host facing port, then the “access” persona needs to be explicitly configured on it.



- If a 10G/25G port needs to be used as a fabric port, then the “uplink” persona needs to be explicitly configured on it.
  - Inter-switch links (ISLs) should be configured to “No Persona”
- For locally-switched flows on the switch, the traffic from the host is subject to policy processing only once and only egress policy is enforced.
- For locally-routed flows on the switch, the traffic from the host is subject to egress policy in source VLAN (pre-routing) first and then subject to ingress policy in destination VLAN (post-routing). In this case, the firewall log is generated for the ingress policy applied on the post-routed destination VLAN. No firewall log will be generated for the egress policy applied on the source VLAN.
- For PVLAN flows that are locally switched on the switch, the traffic from the host is subject to egress policy in the source VLAN (which is the secondary PVLAN) first and then subject to ingress policy in destination VLAN (which is the primary PVLAN). A firewall log record will be generated for the ingress policy applied to the destination VLAN.
- Users can choose to export Firewall Log records, which are generated for flows subject to policy evaluation.
  - If PSM is chosen as an export target for Firewall Log records, the supported sustained ingestion rate is 10K logs per second (LPS), with a burst rate of 20K LPS.
- If a security policy attached to a network has no rules, then there is an implicit default deny rule to drop all the traffic. In this scenario, there are no rule statistics available for denied flows. Flow logs generated in this scenario will not contain the policy and rule information.
- If a VLAN or VNI belonging to a given VRF needs a security policy, users *must* define network objects on the PSM for all VLANs in that VRF.
- If rules are edited within an existing in-use Firewall Policy, all existing rule statistics within the Firewall Policy will be reset.
- If a Firewall Policy is updated, existing active flows will be re-evaluated and the new updated policy will be enforced. The new policy will be enforced only after the initiator of the flow sends traffic or at least 1 packet. This dynamic flow refresh feature is not supported for data-flows of ALG type of traffic including TFTP, RTSP, SUNRPC and MSRPC ALG. The policy for these types of active ALG data flows will not be updated dynamically. For FTP, the policy for active data flows will be updated only after there is a control packet exchange from the client.


Firewall profiles				UPDATE
 Firewall Profile				
Details				
Name:	default	Tenant:	default	
Modification Time:	2022-11-02 05:26:14 GMT+00:00	Creation Time:	2022-05-16 20:19:47 GMT+00:00	
Specs				
Session Idle Timeout:	1000s	TCP Timeout:	1000s	
Status				
Propagation:	Propagation Complete. Updated on 1 DSSs.			

Figure 35. Default Firewall Profile and Timeouts

- With flow log export, there are no periodic updates for long-lived flows.
- If route leaking is enabled on the DSS and if the VRFs involved in the route leaking are pinned to different DSMs, then stateful services on traffic between VRFs is not supported. It is recommended to use the VRF pinning CLI for statically pinning both the VRFs to the same DSM instance.
- When a DSS is acting as an L2 access switch, having different actions (PERMIT/DENY) for the same flow across ingress and egress policies on different VLANs is not permitted.
- Broadcast (ARP), L2 multicast, IP multicast and IPv6 traffic are not redirected to the DSE and therefore are not subject to policy evaluation.
- Firewall policy enforcement on VXLAN centralized gateways is not supported; this is only supported on VXLAN distributed anycast gateway designs.
- On non-AFC controlled setups, for a given VRF, the same VRF name *must* be used on the PSM and the DSS.

## VSX and Firewall High Availability

Aruba VSX can be used to pair switches, increasing firewall availability by preventing application connections from being dropped when a switch fails or is upgraded.



**NOTE:** In order for VSX firewall high availability to function effectively, `linkup-delay-timer` and `inter-switch-link peer-detect-interval` **must be configured**.

- Configure the required `linkup-delay-timer` and `inter-switch-link peer-detect-interval` to 600 seconds under VSX in AOS-CX

```
switch# conf t
switch(config)# vsx
switch(config-vsx)# linkup-delay-timer 600
switch(config-vsx)# inter-switch-link
peer-detect-interval 600
```

- Configure OSPF `max-metric` using the following commands:

```
switch# conf t
switch (config)# router ospf <>
switch (config-ospf-1)# max-metric router-lsa include-stub
on-startup
```

This command advertises a high metric for 600s (default) post reload.

- When adding ports as members of an ISL LAG for VSX, make sure to configure `no persona` for each one.
- Connection tracking with VSX is fully supported.
- Using ALG protocols in security policies with VSX deployments is fully supported.
- Logging considerations:
  - Only the VSX node that sees the first packet in a flow will increment rule statistics.
  - With Firewall Log export, an open record will be generated by either the primary or secondary switch that receives the first flow miss, and a “close” record will be generated by both primary or secondary switches.
- The maximum session limit on VRF/network will apply to both local learn flows and to flows that get flow-synced from the VSX peer.

## Supported Topologies

**Note:** See the CX 10000 Stateful Services Deployment Guide for more detailed information.

CX 10000 firewall can be supported in standard switching deployment architectures, including traditional 3-tier architecture with VSX (core, aggregation, and access), and Clos-based spine-leaf fabrics with VXLAN overlays.




Figure 36. Deployment architectures

## Connection Tracking

TCP connection tracking is always enabled on the CX 10000 platform; there is no need to explicitly enable this on PSM as part of the security profile. This feature is supported on both the standalone single-switch CX 10000 as well as the VSX pair. TCP connection tracking is a key feature as part of the stateful firewall implementation, providing detailed per-flow validation of TCP control packets during the initial 3-way handshake sequence (SYN, SYN-ACK, ACK) and also the closing sequence involving the FIN/RST packets.

As part of connection-tracking, invalid TCP control packet sequences, invalid TCP packets with incorrect flag combinations, incorrect sequence numbers in TCP control packets are detected and dropped providing a more robust and security stateful firewall.

Connection tracking on VSX provides at high level several categories of checks:

- Checks based on existing state of flow and direction of flow (i-flow/r-flow)
- Checks for invalid TCP packets based on current state of flow
- Checks for invalid sequences in each direction of flow
- Sequence number based checks for different types of TCP packets, including SYN/ACK/RST/FIN during the initial 3-way handshake and session closures due to FIN/RST

Packets are dropped for these invalid conditions and will not result in changes of the per-flow state machine.

## Firewall Policy Configuration

To enable stateful L4 firewalling, follow these steps:

- Create a VRF on the PSM
- Create a Security Policy on the PSM
- Attach the Policy to a Network on the PSM
- Create VRF and VLAN on AOS-CX

The VRF and VLAN must also be defined on the switch using AOS-CX, either before or after configuring security policy for those networks in PSM. If this is not done, traffic will not be redirected within the switch to the DSM, and there will be no stateful security evaluation of flows for the specific VRF or VLAN/network.

### Create a VRF

On the PSM, navigate to `Tenants` → `VRF` and click [ADD VRF](#).



*Figure 37. PSM VRF screen detail showing ADD VRF button*

Provide a name for this VRF. To save, click on **ADD VRF** on this screen as well.

Figure 38. VRF creation screen

**Note:** The VRF names in AOS-CX and the PSM must match.

## Create a Network Security Policy

There are two ways to create a network security policy: manually entering a predetermined policy, or by using the PSM UI network graph feature to help construct policies.

If the intended policy has already been determined, use the *table view* within the UI to define it. (An API call is available as well.) The *network graph* feature can be used instead, which aids in discovering what flows should be allowed and designing a zero trust network policy to only allow those flows. The network graph method is covered in detail in Appendix E.

To use table view, navigate to `Tenants → Security Policy` and click on **ADD SECURITY POLICY**:

Figure 39. PSM Security Policy screen detail showing ADD SECURITY POLICY button

Enter the relevant information and click **CREATE POLICY** to save the new policy.

The screenshot shows the 'Network Security Policies' interface. At the top right, there are buttons for 'Table View', 'CREATE POLICY' (circled in red), and 'CANCEL'. The main form area is titled 'Network Security Policies' and contains a 'Policy Name' field with the value 'demo-policy'. Below this is a 'RULES: (specify one or more)' section. Inside this section, there is a 'Rule Name' field with 'r1', a 'Description' field, and a 'Source IP Addresses' field with '1111/32'. The 'Action' dropdown is set to 'Permit', and the 'Destination IP Addresses' field has '1.2.2.2/32'. Under the 'PROTO-PORTS' section, the 'Protocol' is set to 'tcp'. Under the 'APPS' section, the 'Ports' field has '80'. There is a '+ ADD' button at the bottom of the rules section.

Figure 40. Security Policy creation screen

The Network Graph UI feature helps discover the relationships between endpoints within a VRF, within and between VLANs, and within or between endpoint groups. The relationships help better understand common connection occurrences within the data center that should be permitted. See [Appendix E](#) for a description of two methods for determining policies.

### Rule Overlap Detection

During policy creation, if there are rules in a policy that overlap, the PSM will flag them to prevent a possible misconfiguration, as shown in Figure 41:

The warning dialog box has a title bar that says 'Warning: Rule overlaps have been found'. The main content area contains a warning icon and the text: 'Click **SUBMIT** to continue or **Cancel** to go back to the Policy Editing Form.' Below this, it lists two overlapping rules: 'Rule 1 overlaps with rule 2 on source IP address(es) "10.1.1.1" and destination IP address(es) "10.1.1.2" on one or more ports or applications.' and 'Rule 2 overlaps with rule 1 on source IP address(es) "10.1.1.1" and destination IP address(es) "10.1.1.2" on one or more ports or applications.' At the bottom right, there are two buttons: 'Submit' and 'Cancel'.

Figure 41. Rule overlap warning

Rule Overlap Summary						
Rule	Source IP	Action	Destination IP	Protocol Port	Applications	
Rule 2	10.11.1/32	permit	10.11.2/32	tcp/80		
Impacted by Rule	Source IP	Action	Destination IP	Protocol Port	Application	
Rule 1	10.10.10.0/8	permit	10.10.10.0/8	tcp/80		

Figure 42. Rule overlap summary

Network Security Policies > 1-ui

SAVE POLICY

CANCEL ALL

Security Policy

Security Policy Details

Policy Name:

1-ui

Tenant:

default

Created on:

2022-05-20 17:58:04 GMT+00:00

Ingress Policy for VRFs:

Ingress Policy for Networks:

Propagation:

Attach-tenant:

true

Last Modified:

2022-05-20 21:29:19 GMT+00:00

Egress Policy for VRFs:

Egress Policy for Networks:

Rules Overlap Detected:

1

Policy Rules (3)

Search:

Rule Name

Source IP

Destination IP

<Protocol>/<Port>

App

Action

8 Columns

	<div><div></div></div>	Number	Source IPs	Destination IPs	Action	Protocol Port	Applications	Status	Total Connection Hits
<div></div>	<div></div>	1	10.10.10.0/8	10.10.10.0/8	Permit	tcp/80		Enabled	N/A - policy is not saved
<div></div>	<div></div>	2	10.1.1.1/32	10.1.1.2/32	Permit	tcp/80		Enabled	N/A - policy is not saved
<div></div>	<div></div>	3	10.2.1.1/32	10.3.1.1/32	Permit	tcp/80		Enabled	N/A - policy is not saved

Figure 43. Display of policy showing a detected rule overlap

In the example in Figure 43, there is an overlap between rule1 and rule2/rule3, so one overlapping rule is reported as being detected. In the list a warning is indicated against rule 1.

**Note:** It is not recommended to use overlapping rule detection for policies with large rule scale.



## Attach Policy to a Network

Navigate to **Tenants** → **Networks** and click on **ADD NETWORK**.




Figure 44. PSM Network screen detail showing ADD button

Enter the relevant information and click **CREATE NETWORK**.

- Select the VRF to be used.
- Select the ingress and egress security policies.
- Specify the VLAN where this policy is to be applied. This does not create the AOS-CX VLAN.

 A screenshot of the 'Networks Overview' form. At the top, there are 'CREATE NETWORK' and 'CANCEL' buttons. The form contains several fields: 'Name' (required), 'VLAN ID' (required), 'Ingress Security Policy' (dropdown), 'Egress Security Policy' (dropdown), 'VRF' (dropdown), 'Maximum CPS' (with options 'Inherit from VRF', 'Unlimited', 'Set Value'), and 'Maximum Sessions' (with options 'Inherit from VRF', 'Unlimited', 'Set Value').

Figure 45. Network creation screen

Once the network is created, you will be prompted to create the VLAN in AOS-CX. Creating the policy on the PSM does not automatically create the VLAN on AOS-CX.

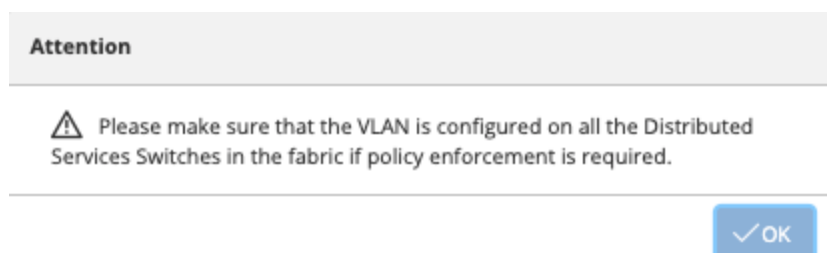


Figure 46. Warning notice for required AOS-CX action

## Create VRF on AOS-CX

If the VRF does not already exist in AOS-CX, create the VRF in AOS-CX using the AOS-CX CLI, AFC UI, or via the AOS-CX REST API. The example shown below uses the CLI.

```
switch# configure
switch (config)# vrf test
switch (config-vrf)# exit
switch (config)#
```

## Create VLAN on AOS-CX

Create the VLAN via the AOS-CX CLI, AFC UI, or programmatically. The example shown below uses the CLI.

```
switch# configure
switch (config)# vlan 199
switch (config-vlan-199)# name demo-vrf
switch (config-vlan-199)# exit
switch (config)# exit
```

## Understanding Firewall Policy Scaling Profiles

By default, the CX 10000 firewall can scale up to 6K rules per policy. This can be seen in the PSM UI; navigate to the cluster object and the policy scale is shown as in Figure 47:




Figure 47. PSM cluster settings, including policy scale limit.

While this scale should satisfy the majority of enterprise site needs, it is not uncommon for some sites to require firewall policy scaling beyond the default. CX 10000 firewall offers the ability to increase policy scaling by switching to a different policy scaling profile, which will allow up to 24K rules per policy.

### Switch Policy Scaling Profile

To change the profile, use the following workflow:

- Navigate to Tenant | Networks and click on the network name. Detach all the attached policies from the Network or VRF. If one or more policies remain attached, the policy profile change will fail.

Name: Subnet-1336

VRF: Test-VRF-84

Ingress Security Policy: ingress-network-4

Egress Security Policy: egress-network-4

VLAN ID: 1336

Click to detach policies

Figure 48. Detach policies

- There must be no attached policies from networks or VRF. Verify this by navigating to **Tenants | Networks** or **VRF** to confirm there is no policy attached.

**Networks**

**Networks Overview**

Networks (1,000)

**No Policy Attached**

<input type="checkbox"/>	Name	VRF	VLAN	Ingress Policy	Egress Policy
<input type="checkbox"/>	Subnet-1001	Test-VRF-1	1001		
<input type="checkbox"/>	Subnet-1002	Test-VRF-1	1002		
<input type="checkbox"/>	Subnet-1003	Test-VRF-1	1003		
<input type="checkbox"/>	Subnet-1004	Test-VRF-1	1004		

Figure 49. No policies are listed for Ingress or Egress.

- Switch the profile to “24K” under **System | Cluster** in the PSM UI:

Edit Cluster Security Policy Rule Scale

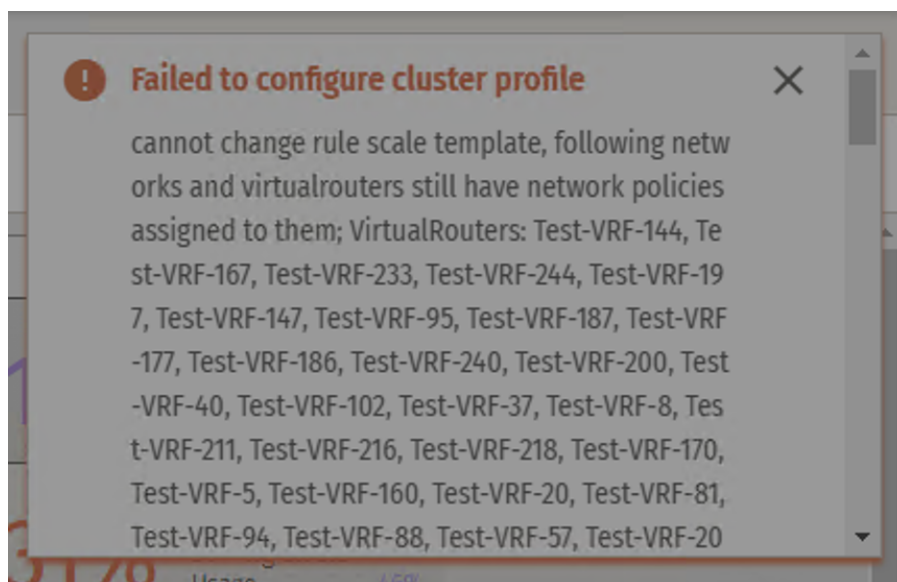
Security Policy Rule Scale: 6K

6K

24K

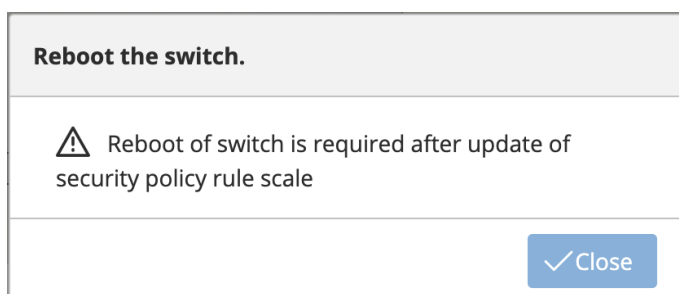
Figure 50. Changing the policy scale limit

If any policies are attached to the VRF/network, the cluster-wide scale profile change is not allowed and a pop-up message indicating the problem will appear at the top right corner:



*Figure 51. Failure notification due to attached policies, which must be detached in order to proceed.*

- After the change is saved, a pop-up message will appear, noting that a reboot is required to activate the profile on each switch.



*Figure 52. Alert that rebooting the switch is needed*

- Before the reboot, under `System | DSS`, “Distributed Services Switches Overview” will show the previous scale setting of “6k” and the “Health” column will indicate a reboot is required:










Distributed Services Switches Overview						
<div> <div>Distributed Services Switches (3)</div> <div>6 Columns</div> <div>Search</div> <div>Aa Ab R*</div> </div>						
<input type="checkbox"/>	Switch Name	Switch OS Version	↑	Management	Health	Policy Scale
<input type="checkbox"/>	P1-27	DL.10.10.0001AJ-45-g8cd099a9cbd77		Mode: Inband IP address: 10.30.22.2	  	Max 6K Rules
<input type="checkbox"/>	P0-36	DL.10.10.0001AJ-45-g8cd099a9cbd77		Mode: Inband IP address: 10.30.2.52	  	Max 6K Rules
<input type="checkbox"/>	P0-11	DL.10.10.0001AJ-45-g8cd099a9cbd77		Mode: Inband IP address: 10.30.2.53	  	Max 6K Rules

Figure 53. The power icon's presence indicates which switches need rebooting.

- After the reboot, System | DSS now shows the new scale setting of “24k”







Distributed Services Switches Overview						
<div> <div>Distributed Services Switches (3)</div> <div>6 Columns</div> <div>Search</div> <div>Aa Ab R*</div> </div>						
<input type="checkbox"/>	Switch Name	Switch OS Version	↑	Management	Health	Policy Scale
<input type="checkbox"/>	P0-11	DL.10.10.0001AJ-45-g8cd099a9cbd77		Mode: Inband IP address: 10.30.2.53	 	Max 24K Rules
<input type="checkbox"/>	P0-36	DL.10.10.0001AJ-45-g8cd099a9cbd77		Mode: Inband IP address: 10.30.2.52	 	Max 24K Rules
<input type="checkbox"/>	P1-27	DL.10.10.0001AJ-45-g8cd099a9cbd77		Mode: Inband IP address: 10.30.22.2	 	Max 24K Rules

Figure 54. Switch status after the switch is rebooted and new profile is applied

**NOTE:** Policy Rule Scale is a cluster-wide setting; all switches managed by the same PSM cluster will have the same scale setting. Mixing different rule scale settings for different switches within the same cluster is not supported.

When a new switch joins a PSM cluster, if the default profile value on the switch does not match the value configured in the PSM, then the switch will need another reboot to get adjusted automatically.

**NOTE:** *With a 6K policy profile:*

- *The max-rules per policy is set to 6138.*
- *Up to 80 policies can be defined if each policy is of size 6k.*
- *Up to 500 policies can be defined if each policy is of size  $\leq 1k$ .*

*With a 24K policy profile:*

- *The max-rules per policy is set to 24,570.*
- *Up to 17 unique policies can be defined if each policy is of size 24k.*
- *Up to 102 unique policies can be defined if each policy is of size  $\leq 4k$*

### Verifying the Number of Rules Consumed in the Data Plane

A given rule in the PSM gets expanded internally into multiple rules in the data plane. This expansion is based on the number of unique source prefixes, destination prefixes, port numbers/ranges, protocol types, and other data.

With the 24K policy profile, if the number of expanded rules is  $\geq 4095$ , then the DSM internally combines multiple policies to accommodate them. For example, if a 24K policy is created, then the DSM uses 6 internal policies of 4K rules each to support it.

In order to avoid exceeding the switch's available internal resources, it is important to be aware of how many internal policies and internal expanded rules are actually consumed in the data plane.

To verify internal policies and rules, navigate to `Tenants | Security Policies`, then click on the name of the policy you wish to verify. Under Security Policy Details, click on “View Details” .

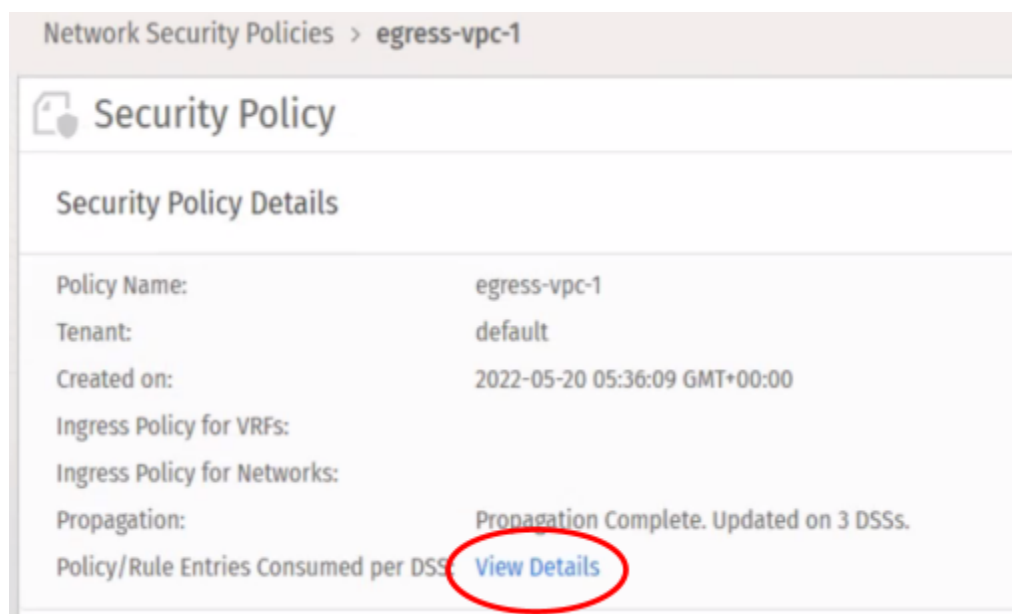


Figure 55. Select View Details to see Policy/Rule Entries Consumed per DSS.

1. A new window, as seen in figure Figure 56, shows the actual number of internal policies and actual number of rule entries consumed in hardware.

Policy/Rule Entries Consumed per DSS			
Policy/Rule Entries Consumed per DSS (3)			
Host Name	DSS ID	Policy Entries Consumed	Rule Entries Consumed
leaf1	DSS-0490.8100.0596	6	24570
leaf2a	DSS-0490.8100.37dc	6	24570
leaf2b	DSS-0490.8100.13ce	6	24570

Figure 56. Detail view: Policy/Rule Entries Consumed per DSS

Alerts can be set to notify if the number of policies and rules in the data plane have reached the resource limit:

2. Navigate to `Monitoring | Alerts & Events` , then add a new Alert Policy
3. Click on the “OBJECT BASED ALERT POLICIES” tab:






Figure 57. Object Based Alert Policies tab selector

- Type in the policy name and choose “NetworkSecurityPolicy”

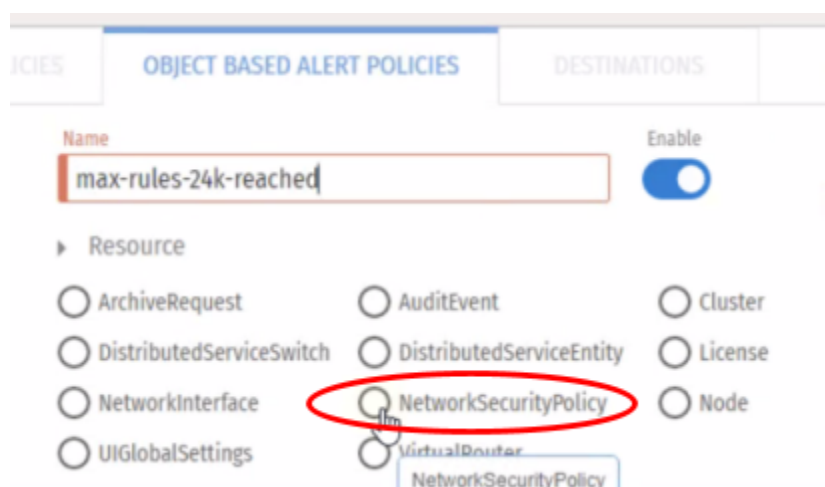


Figure 58. Object Based Alert Policy for Network Security Policy

- In the next screen, from the “Key” dropdown menu select “status.rule-metrics-status.rule-entries-consumed”

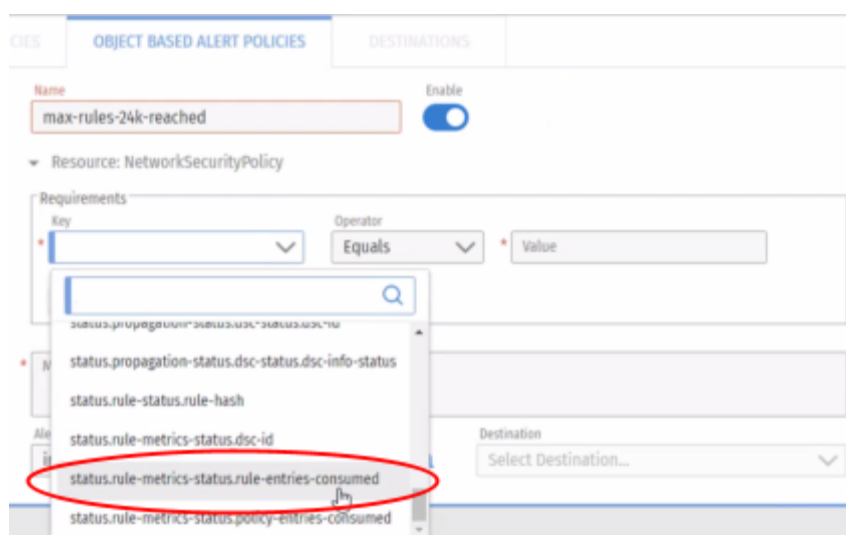


Figure 59. Object Based Alert Policy for Rule-Entries Consumed in Policy

- Put the max allowed value of 24570 in the “Value” field and change the severity level to “warn”

The screenshot shows the configuration for an object-based alert policy. The policy is named 'max-rules-24k-reached' and is currently enabled. It is associated with the resource 'NetworkSecurityPolicy'. Under the 'Requirements' section, a key-value pair is defined: the key is 'status.rule-metrics-sta...' (partially visible), the operator is 'Equals', and the value is '24570'. Below this, a message template is provided: 'Max 24k rules have been reached'. At the bottom, the 'Alert Severity Level' is set to 'warn', and the 'Destination' is set to 'Select Destination...'.

Figure 60. Selecting values and Alert Security Level

- Similarly, create another alert policy once the internal policies number reaches 6

This screenshot shows the configuration for a second object-based alert policy named 'policy-max', which is also enabled. It is linked to the 'NetworkSecurityPolicy' resource. In the 'Requirements' section, the key 'status.rule-metrics-sta...' is selected, and the operator is 'Equals'. A dropdown menu is open, displaying a list of available keys. The key 'status.rule-metrics-status.policy-entries-consumed' is highlighted at the bottom of the list. The 'Value' field is currently empty. The message field is also empty, and the destination is set to 'Select Destination...'.

Figure 61. Object Based Alert Policy for Policy Entries Consumed

Figure 62. Selecting values and Alert Security Level for Policy Entries Consumed

With these two alert policies configured, once the internal policy number or the internal expanded rule number reaches the configured limit, alert messages will be raised to warn the administrator.

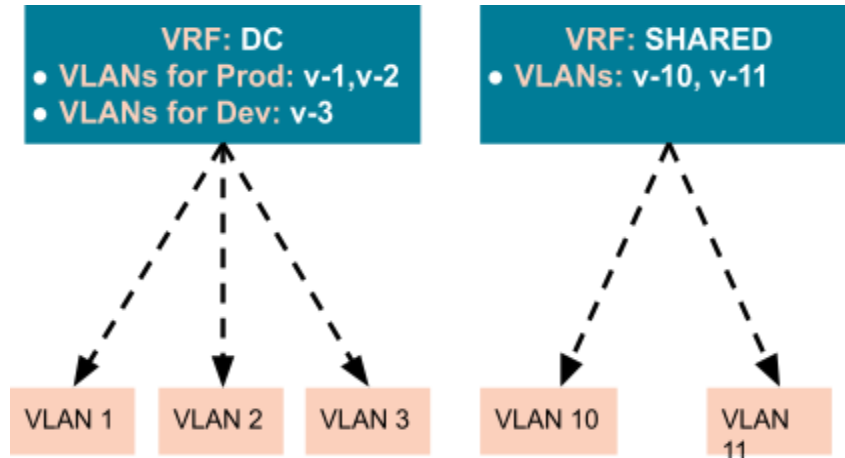
Severity	Message	State	Alert Type	Source Node & Component	Total Hits	Creation Time	Modification Time
warn	[ ingress-vpc-1 ] Max policy limit has been reached	Open	object	Ref: Reporter: psm-alertmgr	1	2022-06-01 23:32:37 GMT+00:00	2022-06-01 23:32:46 GMT+00:00
warn	[ ingress-vpc-1 ] Max 24k rules have been reached	Open	object	Ref: Reporter: psm-alertmgr	1	2022-06-01 23:32:37 GMT+00:00	2022-06-01 23:32:46 GMT+00:00

Figure 63. Warnings that the policy limit has been reached

**NOTE:** Profile change is a significantly disruptive operation, as it involves policy detachment and restarting switches. It is highly recommended to plan and implement the appropriate scale profile for your site during the initial design to be established when the PSM cluster is bootstrapped.

## Understanding Hierarchical Security Policy

Data center networking is usually implemented as a hierarchical design: a VRF represents the higher level network, and each VRF network consists of multiple VLAN networks which represent the lower level network, as in figure 64:



*Figure 64. A typical data center hierarchical network design*

Attaching and enforcing security policies at only the VLAN level may not be sufficient for some enterprise sites for any of several reasons:

- For simplicity purposes, sites may want to have a single unified security policy enforced across all VLANs within the VRF. Attaching the same policy explicitly to each and every VLAN can be cumbersome.
- In a large enterprise data center deployment, common scenarios include:
  - A set of commonly shared services are provided to all workloads in the data center, such as DNS, DHCP, or NTP.
  - Different VLAN networks provide different services such as production vs development, depending on what apps are running on them.
  - The commonly shared services are usually stable or fixed over time. Administrators commonly would prefer not to change security policies on a regular basis. However, app-specific services may change frequently. Applications can be brought up and down regularly in a dynamic environment. Their related policies would need to be modified in turn.

The DSS hierarchical security policy model is designed to handle the above scenarios by supporting the security policy to be attached and enforced at the VRF level:




Figure 65. Attaching Security Policy

## Policy Enforcement

- Two levels of policy enforcement will happen in both directions
- Traffic going out of the host will be subject to egress policy enforcement, first at Network level, followed by VRF level. Traffic coming into the host will be subject to ingress policy enforcement first at VRF level followed by and Network level.




Figure 66. Egress/Ingress Policy Enforcement at VRF/Network levels

- For *local switched traffic* (between hosts connected to the same DSS), only egress policy is enforced.




Figure 67. Hierarchical design

- If there is a “Deny” at any level (VRF/Network) the traffic will be denied (both ingress and egress)
- If no policy is configured, that will be considered as “Allow” and the traffic will pass
- Tables 5 and 6 list the final result for egress/ingress policy evaluation:

Network→	VRF	Final result
Allow	Allow*	<b>Allow</b>
Allow	Deny*	<b>Deny</b>
Deny*	Allow	<b>Deny</b>
Deny*	Deny	<b>Deny</b>
No policy	Allow*	<b>Allow</b>
No policy	Deny*	<b>Deny</b>
Allow*	No policy	<b>Allow</b>
Deny*	No policy	<b>Deny</b>
No policy	No policy	<b>Allow</b>

*Table 6. Egress policy enforcement rules. For Egress, network policy is evaluated first followed by VRF policy.*

*\* indicates rule-stats / FW log hits*

VRF→	Network	Final result
Allow	Allow*	<b>Allow</b>
Allow	Deny*	<b>Deny</b>
Deny*	Allow	<b>Deny</b>
Deny*	Deny	<b>Deny</b>
No policy	Allow*	<b>Allow</b>
No policy	Deny*	<b>Deny</b>
Allow*	No policy	<b>Allow</b>
Deny*	No policy	<b>Deny</b>
No policy	No policy	<b>Allow</b>

*Table 7. Ingress policy enforcement rules. For Ingress, VRF policy is evaluated first followed by network policy.*

*\* indicates rule-stats / FW log hits*

**Rule-stats:** Statistics that are maintained by DSM at a per-rule level for each policy accounting for number of packets hitting a rule during flow-create. These statistics are also exported to PSM and are available to view as “Connection Hits” in the policy UI page against each rule.

**FW Log hits:** A FW log is created and exported to configured destination (PSM/external) for each rule-hit within policy at the time of flow-create / flow-delete

### Configuration and Verification

To apply Stateful L4 Security Policy to a VRF, navigate to `Tenants` → `VRF` in the PSM GUI, and click on [ADD VRF](#).



Name

Ingress Security Policy  
 X ▾

Egress Security Policy  
 X ▾

Select flow export policies... ▾

Maximum CPS

Maximum Sessions

Maximum CPS

Maximum Sessions

Figure 68. Hierarchical design

- **Name** of the VRF (string).
- **Ingress Security Policy**: choose a previously configured security policy.
- **Egress Security Policy**: choose a previously configured security policy.
- **Flow export policies**: choose a previously configured policy.
- The Maximum CPS and Maximum Sessions for a VRF can be set to Unlimited or to a value.
- **Maximum CPS per DSE**, which can be up to the maximum supported connections per Second (range 1000-1000,000) for the entire system.
- **Maximum Sessions per DSE**, which can be up to the maximum supported sessions (range 10,000 - 5000,000) for the entire system.

The VRF policies can be verified in the VRF Overview under `Tenants → VRF`

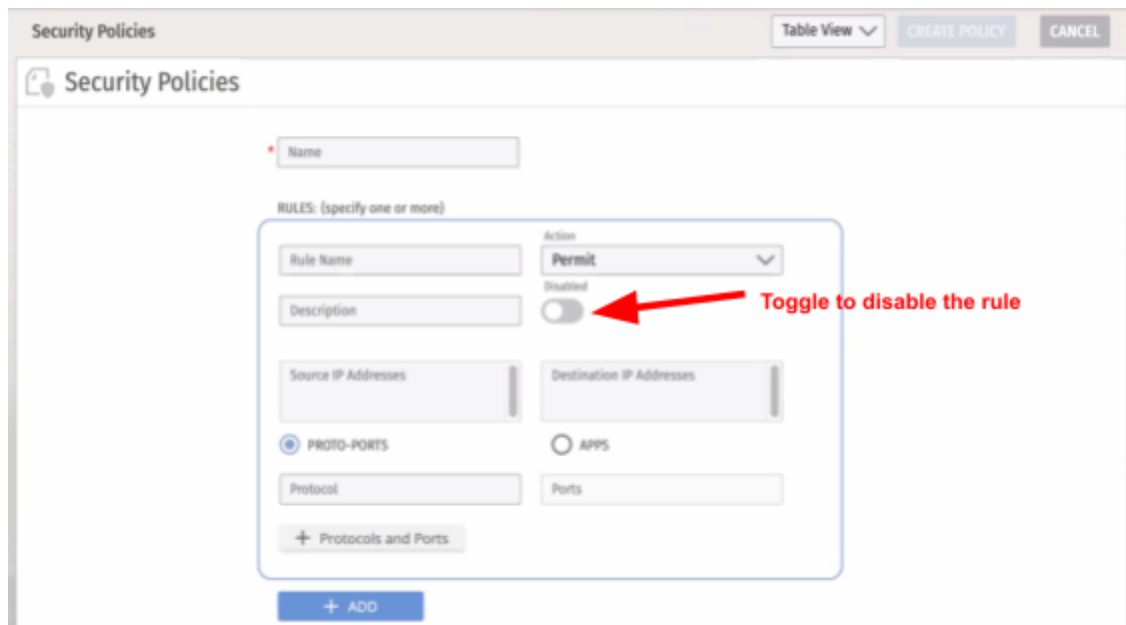
VRF						
VRF Overview						
VRF (252) <span style="float: right;">8 Columns</span>						
<input type="checkbox"/> ▾	Name	Associated VLANs	Ingress Policy	Egress Policy	Maximum CPS	Maximum Sessions
<input type="checkbox"/>	demo-vrf		ingress-vpc-1	egress-vpc-1	1000	10000
<input type="checkbox"/>	default				Unlimited	Unlimited
<input type="checkbox"/>	Test-VRF-206	1821, 1822, 1823, 1824	ingress-vpc-1	egress-vpc-1	400000	4000000

Figure 69. Hierarchical design

## Enable/Disable Individual Firewall Rules

Firewall rules can be enabled or disabled individually. There are two ways to use this feature:

1. When creating a new security policy, each rule can be added with the option to disable it (Rules are enabled by default):



The screenshot shows the 'Security Policies' configuration window. At the top, there's a 'Table View' dropdown and 'CREATE POLICY' and 'CANCEL' buttons. Below the title bar, there's a 'Name' field. The main section is titled 'RULES: (specify one or more)'. Inside this section, there's a form for adding a new rule. The form includes fields for 'Rule Name', 'Description', 'Source IP Addresses', and 'Destination IP Addresses'. There are also radio buttons for 'PROTO-PORTS' (selected) and 'APPS'. Below these are fields for 'Protocol' and 'Ports'. A red arrow points to the 'Disabled' toggle switch, which is currently turned off. A red text label 'Toggle to disable the rule' is next to the arrow. At the bottom of the rule form, there's a '+ ADD' button.

Figure 70. Specify a new rule should be disabled

- For existing security policies, go to the Policy Editing page, where each rule can be enabled or disabled individually:

Security Policies > East-West\_segmentation CLONE POLICY

### Security Policy

Policy Editing

Policy Name: East-West\_segmentation SAVE CANCEL

Rule Order: 1 Action: Permit Rule Name: Disabled Description:

Source IP Addresses: 10.1.0/24 Destination IP Addresses: 10.2.0/24

☒ PROTO-PORTS ☐ APPS

Protocol: TCP Ports: 443

+ Protocols and Ports

Preview

Number	Rule Name	Sources	Destinations	Action	Protocol Port	Applications	Description	Status
1		IPs: 10.1.0/24	IPs: 10.2.0/24	Permit	TCP/443			Enabled

Figure 71. Disable an existing rule

A disabled rule is indicated in the “Status” field on the Security Policy page:

Security Policies > East-West\_segmentation CLONE POLICY COMPARE CLONED POLICY

### Security Policy

#### Security Policy Details

Policy Name: East-West\_segmentation

Tenant: default Attach-tenant: true

Created on: 2022-10-10 17:52:32 GMT+00:00 Last Modified: 2022-10-10 17:54:35 GMT+00:00

Ingress Policy for VRFs: Egress Policy for VRFs:

Ingress Policy for Networks: Egress Policy for Networks:

Propagation:

Policy Rules (1) Search: Rule Name Source IP/Workload Group Destination IP/Workload Group <Protocol>/<Port> App

	Number	Rule Name	Sources	Destinations	Action	Protocol Port	Applications	Description	Status	Total Connection Hits
	1		IPs: 10.1.0	IPs: 10.2.0	Permit	TCP/443			Disabled	

Figure 72. Identify disabled rules

## Configuring Firewall Log Export

Firewall logs from the DSS can be exported to an external destination/collector, or to the PSM. Export to external collector destinations is only supported in the default VRF with inband IP.

Enabling firewall log export requires two steps:

- Defining Firewall Log export policy configuration
- Binding policy to the DSS

### Firewall Log Export Policy Configuration

Navigate to **Tenants** → **Firewall Export Policies** and click on **ADD FIREWALL LOG EXPORT POLICY**:




Figure 73. Firewall Log Policies screen detail, showing ADD FIREWALL button

Enter the relevant information and click **CREATE FIREWALL LOG POLICY**.

- Fill out the mandatory fields: name, destination, transport.
- Select the format and facility override.

 A screenshot of the 'Firewall Log Export Policies' form for creating a new policy. The form has a header bar with the title 'Firewall Log Export Policies' and two buttons: 'CREATE FIREWALL LOG POLICY' (highlighted with a red box) and 'CANCEL'. Below the header, the form contains several sections:
 

- Name:** A text input field containing 'demo-log-export'.
- Export:** A dropdown menu currently showing 'All Logs'.
- Enable PSM Target:** A toggle switch that is currently turned off.
- SYSLOG EXPORT CONFIGURATION:** A section containing:
  - Facility Override:** A dropdown menu showing 'User'.
  - Format:** Two radio button options: 'BSD' (selected) and 'RFC5424'.
- Target:** A section containing:
  - destination:** A text input field containing '192.168.100.1'.
  - Transport:** A text input field containing 'udp/511'.

Figure 74. Firewall log policy creation screen

Figure 75 shows an example of defining a firewall log policy to export to the PSM:

Figure 75. Firewall Log Export to PSM

## Bind Export Policy to DSS

Navigate to “System → DSS” and click on the DSS-ID of the switch to apply the policy to.

Switch Name	DSS-ID	DSM	Switch OS Version	Switch Serial #	Management	Health	Labels	FW Logging Export	Policy Scale	Modification Time
leaf2a	DSS-0490.8100.37dc	DSM 1/1: 0490.8100.383c DSM 1/2: 0490.8100.384c	DL.10.10.0001BH...	US17L6B03E	Mode: Inband IP address: 10.30.22.2	✓			Max 6K Rul...	2022-05-25 21:14:14 GMT+00:00
leaf2b	DSS-0490.8100.13ce	DSM 1/1: 0490.8100.142e DSM 1/2: 0490.8100.143e	DL.10.10.0001BH...	US14L96031	Mode: Inband IP address: 10.30.2.52	✓			Max 6K Rul...	2022-05-25 21:14:05 GMT+00:00
leaf1	DSS-0490.8100.0596	DSM 1/1: 0490.8100.05f6 DSM 1/2: 0490.8100.0606	DL.10.10.0001BH...	US12L96017	Mode: Inband IP address: 10.30.2.53	✓			Max 6K Rul...	2022-05-25 21:14:22 GMT+00:00

Figure 76. DSS list screen

Click on the icon to apply Firewall Log Policy:




Figure 77. Select the “assign” action.

Select the previously-defined policy, then click **ASSIGN**:




Figure 78. Policy assignment window

Distributed Services Switches Overview										
Distributed Services Switches (3)										
Switch Name	DSS-ID	DSM	Switch OS Version	Switch Serial #	Management	Health	Labels	FW Logging Export	Policy Scale	Modification Time
leaf2a	DSS-0490.8100.37dc	DSM 1/1: 0490.8100.383c DSM 1/2: 0490.8100.384c	DL 10.10.0001BH...	US17L6803E	Mode: Inband IP address: 10.30.22.2	✓			Max ...	2022-05-25 21:14:14 GMT+00:00
leaf2b	DSS-0490.8100.13ce	DSM 1/1: 0490.8100.142e DSM 1/2: 0490.8100.143e	DL 10.10.0001BH...	US14L96031	Mode: Inband IP address: 10.30.2.52	✓			Max ...	2022-05-25 21:14:05 GMT+00:00
leaf1	DSS-0490.8100.0596	DSM 1/1: 0490.8100.0516 DSM 1/2: 0490.8100.0506	DL 10.10.0001BH...	US12L96017	Mode: Inband IP address: 10.30.2.53	✓		demo-external-osplog	Max ...	2022-05-26 04:00:01 GMT+00:00

Figure 79. List of switches, showing one with an attached firewall log policy

Once a firewall log policy is attached to a DSS, the list of switches should indicate the attachment, as shown in Figure 80; it will also be listed on the switch's detail page.

## Firewall Log Record Format

Field Name	Description
ts	Flow record timestamp, RFC3339
sessionstate	Can be flow_create or flow_delete
act	Allow or Deny action applied for the session
vpcid	Source VRF UUID for the flow record
sip	Source IPv4 address of the flow
sport	Source TCP/UDP port of the flow

Table 8. Firewall record fields (1/2)

Field Name	Description
dip	Destination IPv4 address of the flow
dport	Destination TCP/UDP port of the flow
proto	Network protocol
sessionid	Flow session identifier
securitypolicyid	UUID of the security policy
ruleid	Hash of the rule against which the flow was evaluated
rulename	User assigned name of the rule within the policy
iflowpkts	Packet count sent from Source to Destination
iflowbytes	Bytes count sent from Source to Destination
rflowpkts	Packet count sent from Destination to Source
rflowbytes	Bytes count sent from Destination to Source
vlan	VLAN configured in Network Object
polycyname	Name of security policy used to evaluate this flow
producttype	DSS
softwareversion	AOS-CX version
serialnumber	Serial Number of DSS
devicename	MAC address of DSS
unitid	Unit ID for DSM. Can be 1 or 2, since there are two DSMs per DSS.
version	V3

*Table 8. Firewall record fields (2/2)*

## Example of Firewall Syslog Message

Example of security policy deny rule hit:

```
Apr 5 05:54:15 P0-29 P0-29[718]:
2022-04-05T05:54:15Z,flow_create,deny,11d576b5-2a27-400a-af00-d10ce033bbf8,63.126.3.
2,39110,63.126.4.2,21,6,4338336,56205fb5-7a49-4341-b9ec-909995faced7,698156806178603
2682,deny-tcp,0,0,0,0,2503,
deny-tcp-policy,DSS,DL.10.09.1010,US14L9601F,0490.8100.0aaa,2,v3
```

Example of security policy permit rule hit:

```
Apr 5 05:54:45 P0-29 P0-29[718]:
2022-04-05T05:54:45Z,flow_create,allow,11d576b5-2a27-400a-af00-d10ce033bbf8,63.126.3
.2,15608,63.126.4.2,2048,1,14765445,dc53b47d-6574-4a41-b89e-c18a3a2f1aec,95549616758
68087418,,0,0,0,0,2504,allow-icmp-policy,DSS,DL.10.09.1010,US14L9601F,0490.8100.0aaa
,1,v3
```



Field	Value from Syslog
ts	2022-04-05T05:54:15Z
sessionstate	flow_create
act	deny
vpcid	11d576b5-2a27-400a-af00-d10ce033bbf8
sip	63.126.3.2
sport	39110
dip	63.126.4.2
dport	21
protocol	6
sessionid	4338336
securitypolicyid	56205fb5-7a49-4341-b9ec-909995faced7
ruleid	6981568061786032682
rulename	deny-tcp
iflowpkts	0
iflowbytes	0
rflowpkts	0
rflowbytes	0
vlan	2503
policyname	deny-tcp-policy
producttype	DSS
AOS-CX version	DL.10.09.1010
serialnumber	US14L9601F
devicename	0490.8100.0aaa
unitid	2
version	3

*Table 9. Values shown in deny rule example*

## Deduplication for Firewall Logs



**NOTE:** At the time of this writing, this feature is included for evaluation purposes only, and is not supported for production use. See the release notes for the release you are using for the most current information about its support status.

The PSM backend supports a deduplication option for the API flowlog query option. This option can be enabled via the PSM, under the `Tenants | Firewall Logs` page, using the `Dedup` toggle:

The screenshot shows the AMD Pensando Firewall Logs management interface. On the left is a navigation sidebar. The main panel is titled 'Firewall Logs' and contains a search bar and various filters. A red circle highlights the 'Dedup' toggle switch, which is currently turned off. Below the filters, there is a table header for 'Firewall Logs (0)' with columns for various log details. The bottom of the interface shows 'AMD Pensando Systems' and the 'Server Time: 2022-10-18 16:24:57 UTC'.

Figure 80. Enabling deduplication

When `Dedup` is enabled, duplicate flow records for a given 5-tuple are removed. When the same 5-tuple flow exists on a given switch over different time instances, only 1 instance, which is the latest instance of the flow record, is returned. Also, when the same 5-tuple flow record gets exported from multiple switches in the fabric, only 1 instance is maintained.

API users can enable flow log de-duplication with a specific type of aggregation.

```
message FwLogQuery {
    "start-time": xxx
    "end-time": xxx
    "source-ips": xxx
    "destination-ips": xxx
    —snip—
    filter-out-duplicates: "do-not-filter" |
    "srcip-destip-destport-proto"
}
```

When submitting a JSON format request, users can either enable 4-tuple based deduplication by specifying `srcip-destip-destport-proto` or disable the deduplication by specifying `do-not-filter`.

#### Notes:

- By default the query will be with `do-not-filter` and all the logs in response will not be deduplicated.
- `srcip-destip-destport-proto` will make each combination of the four tuples only have one log entry in the response.
- The `filter-out-duplicates` option only supports the latest 24 hours of entries, including queries from any to any IP or queries with a specific IP address.
- The `filter-out-duplicates` option will not work with the `countOnly` type of `FwLogQuery`.

## Flow Export (IPFIX)



**NOTE:** At the time of this writing, this feature is included for evaluation purposes only, and is not supported for production use. See the release notes for the release you are using for the most current information about its support status.

IP Flow export, implementing the IETF IPFIX standard, is used to gain visibility into network traffic usage patterns within data center networks, which can then be used as a basis for creating firewall policies.

The following AOS-CX CLI commands are needed to enable Flow Exports:

```
switch(config)# dsm
switch (config-dsm)# ipfix workload-migration ip source-interface
ipfix 1.1.1.1
```

If the `source-interface` CLI -command is absent, then the DSS DSM is used as the source IP address of IPFIX packets (169.254.101.2 and 169.254.101.3 for the two DSMs).

Flow Export can be found under the “Troubleshoot” menu.

Flow Export policies can be configured at a DSS or VRF level to export flows directly to an external collector. Flow records with flow information (including source and destination IP addresses, source and destination ports, permitted and dropped packets) are exported to the external collector(s), based on the configured timers. Flows that are cleaned up due to aging or due to session close are immediately exported to the collector.

Up to 16 million flows can be exported.

The screenshot displays the 'Flow Export Policies' configuration window. At the top right, there are buttons for 'CREATE FLOW EXPORT POLICY' and 'CANCEL'. The main section is titled 'Flow Export' and contains the following fields and controls:

- Name:** Elastic
- Disabled:** A toggle switch is currently turned off.
- Interval:** 10s
- Template Refresh Rate:** 5m
- EXPORT CONFIGURATION:**
  - Format:** Ipfix (selected with a radio button)
- Target:**
  - destination:** 192.168.89.240
  - Transport:** udp/9995
  - + TARGET:** A button to add a new target.

At the bottom right of the configuration area, there are buttons for 'SAVE' and 'CANCEL'.

Figure 81. Defining a Flow Export policy

## Configuring Apps

An App is a service defined either by a protocol/port pair, or by an application-level gateway (ALG). It allows a user to define commonly-deployed applications' protocol/ports, allowing the app's name to be used in policies instead of repeatedly entering the protocol/port information.

App objects also allow for the specification of ALG configurations that define application-specific firewall behavior.

Supported ALGs include FTP, TFTP, RTSP, DNS, SUNRPC, and MSRPC. A set of 41 pre-defined Apps are available.

Apps <span>ADD APP</span>				
Apps Overview				
<div> <div>Apps (41)</div> <div>4 Columns</div> <div>Search</div> <div>Aa Ab R </div> <div>Q</div> <div>▼</div> <div></div> </div>				
<input type="checkbox"/> Name	Configuration	Creation Time	Modification Time	↓
<input type="checkbox"/> RADIUS	Protocol: udp Ports: 1812, 1813	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> DHCP6	Protocol: udp Ports: 546, 547	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> RDP	Protocol: tcp Ports: 3389	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> MYSQL	Protocol: tcp Ports: 3306	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> TFTP	Protocol: udp Ports: 69	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> TELNET	Protocol: tcp Ports: 23	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> SYSLOG	Protocol: udp Ports: 514	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> SSH	Protocol: tcp Ports: 22	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> SNMP	Protocol: tcp Ports: 161, 162, Protocol: udp Ports: 161, 162	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> SIP	Protocol: tcp Ports: 5060, Protocol: udp Ports: 5060	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> PING	Protocol: ICMP Ports: any	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	
<input type="checkbox"/> NTP	Protocol: tcp Ports: 123, Protocol: udp Ports: 123	2022-05-16 20:19:47 GMT+00:00	2022-05-16 20:19:47 GMT+00:00	

Figure 82. The Apps Overview screen

Figure 82 shows some of the predefined Apps which can be used in policy configuration. Additional apps can be created manually, by navigating to **Tenants** → **Apps** and selecting **ADD APP**, as shown in Figure 83.




Figure 83. Apps screen detail, showing ADD button

## Protocol And Ports

In this example, an App is created for the protocol/port pair of **TCP** and **443**. Click on **CREATE APP** when complete.

Figure 84. App definition window

This App can now be used with a security policy. When creating a rule within a policy, select the APPS radio button and choose the name of an App that was created previously.

Figure 85. Adding an App to a security policy

## ALG

In this example, an App is created for DNS ALG.

The screenshot shows the 'Apps' configuration window. At the top right, there are 'CREATE APP' and 'CANCEL' buttons. The main section is titled 'Apps Overview'. It contains a 'Name' field with the value 'demo-app-alg'. Below this are three radio buttons: 'ALG Only', 'Protocols And Ports Only', and 'Both', with 'Both' selected. The 'ALG' section is expanded, showing a 'Type' dropdown set to 'dns'. It includes several toggle switches: 'Drop multi-question packets' (off), 'Drop long label packets' (off), 'Drop large domain name packets' (off), and 'Max Message Length' (512). The 'Max Response Timeout' is set to '60s'. The 'Protocols and Ports' section is also expanded, showing a 'Protocol' dropdown set to 'udp' and a 'Ports' field with the value '53'. A '+ Protocols and Ports' button is visible below these fields.

Figure 86. Creating an ALG App (compare to Figure 84)

This App can now be used with a security policy. When creating a rule within a policy, select the APPS radio button and select the name of the App that was created previously.

The screenshot shows the 'Network Security Policies' configuration window. At the top right, there are 'CREATE POLICY' and 'CANCEL' buttons. The main section is titled 'Network Security Policies'. It contains a 'Policy Name' field with the value 'demo-app-algdns'. Below this is a section for 'RULES: (specify one or more)'. A rule is defined with the following details: 'Rule Name' is 'algdns', 'Action' is 'Permit', 'Description' is empty, 'Source IP Addresses' is '192.168.100.0/24', and 'Destination IP Addresses' is '192.168.101.0/24'. Below the IP address fields are two radio buttons: 'PROTO-PORTS' and 'APPS', with 'APPS' selected. Under 'APPS', a dropdown menu shows 'demo-app-alg'. A '+ ADD' button is located at the bottom of the rule configuration area.

Figure 87. Adding the DNS ALG app to the policy

## DDoS Detection and Alerting: Maximum Sessions and CPS Limits

Two limits can be configured for distributed denial-of-service (DDoS) protection: *maximum session limit* and *maximum connections per second (CPS)*. The PSM is used to configure these limits, and the DSM enforces these limits in the data plane.

Configuring these limits can help to detect unusual network activities which might indicate a DDoS attack underway in the network, and mitigate their effect. When the session limit is enforced and reached, the switch will stop learning any new sessions, preventing the switch's flow table from being exhausted by a DDoS attack. When CPS limits are enforced, a high-rate DDoS attack (SYN attack) will be prevented from overwhelming the switch with new connection requests. All new connections will be dropped when a CPS limit is reached.

These are not per-protocol (TCP/UDP/ICMP) limits; rather, they can be either applied at VRF level or at a network level or both. This section explains details about these limits, including how and where they can be enforced and used.

### Maximum Session Limit

In the CX 10000 stateful firewall implementation, a *session* is defined as a pair of flows: an *I-flow* (initiator flow) and an *R-flow* (responder flow). Note that session and CPS limits apply to total session count and not flow count. This limit can be either applied at VRF level, or at network level, or both. When applied at the VRF level, all networks that are part of that VRF will inherit the configuration. The maximum session limit is not enforced based on strict values; there is a leeway of about 512 sessions between the configured value and the enforced value.

### Maximum CPS

CPS (connections per second) is essentially the rate of incoming new connections that can be successfully processed and installed in the switch's flow table. A DDoS attack with a high rate of incoming SYN packets can potentially exhaust the switch's DSMs and cause it to start dropping valid connections, which is highly undesirable.

Configuring maximum CPS enables a limit to throttle the maximum number of new connections that can be set up. From the data plane perspective this limit is enforced by the DSM's P4 engine with the help of a policer that is applied at a per-VLAN level. The PSM user has options to set this limit as per-VRF, per-network or both.

In general, the best practice is to set general VRF-wide (applied to all networks in VRF) max session-limits and max CPS to values slightly greater than the general operating load in steady-state production. Specific network-level CPS and session limits can be applied to only those networks where more specific values (fine-tuned) are required.



Table 10 describes the behavior of max session limit and max CPS with different configurations. (UI) in parentheses refers to what is seen in the UI, and (API) in parentheses refers to the value sent in the REST payload.

VRF-Level	Network-Level	Expected Behavior
Not-configured (API)	(API) - Not configured	No CPS/max-session limits enforced. This is the default behavior.
Not-configured (API)	Inherit from VRF (UI) “-1” (API)	No CPS/max-session limits enforced
Not-configured (API)	Unlimited (UI) “0” (API)	No CPS/max-session limits enforced
Not-configured (API)	Set Value (UI) “x” sessions “y” CPS	Limit to max of “x” sessions and “y” CPS rate for the given network
Configured “x” sessions “y” CPS	Not-configured (API)	Not-configured defaults to inherit. The configured value of “x” sessions and “y” CPS will be enforced.
Configured “x” sessions “y” CPS	Inherit from VRF (UI) “-1” (API)	The network in the VRF inherits the configured value of the max session limit/rate under the VRF. User will get max of “x” sessions and “y” CPS for each network in VRF
Configured “x” sessions “y” CPS	Unlimited (UI) “0” (API)	No CPS/max-session limits enforced.  <b>NOTE:</b> The configured value of “x” sessions and “y” CPS <b>will not be enforced.</b>
Configured “x1” sessions “y1” CPS	Set Value (UI) “x2” sessions “y2” CPS	Maximum session limit of “x2” sessions and “y2” CPS gets enforced for the given network

Table 10. Expected results for different options for VRF and Network-level configs

## Min and Max Values

The minimum value of the maximum session limit field is 10,000. The maximum value of this field is 5M (5,000,000), well beyond the supported limit of max sessions per DSM, which is 2M sessions.

The minimum value for the max CPS limit field is 1,000. The maximum value of this field is 1,000,000 (1M), well beyond the supported limit of 800K CPS (single-switch / non-VSX).

The following section describes and gives examples of how the maximum sessions and CPS features can be configured via the PSM UI.

## Configuring the Maximum Sessions / CPS on a VRF via the PSM UI

### **For new VRFs:**

1. Select Tenants -> VRF from side menu
2. Select **ADD VRF**
3. Enter a name for the VRF
4. Select Set Value (under Maximum CPS), and enter value
5. Select Set Value (under Maximum Sessions), and enter a value

The screenshot shows the 'VRF Overview' configuration page. At the top right are 'ADD VRF' and 'CANCEL' buttons. The main form contains the following fields:

- Name:** A text input field with an asterisk indicating it is required.
- Ingress Security Policy:** A dropdown menu currently showing 'Ingress Security Policy' with a close (X) and expand (v) icon.
- Egress Security Policy:** A dropdown menu currently showing 'Egress Security Policy' with a close (X) and expand (v) icon.
- Maximum CPS:** A section with a 'Set Value' button and a text input field containing '2000'.
- Maximum Sessions:** A section with a 'Set Value' button and a text input field containing '20000'.

*Figure 88. Adding a VRF*

### **For existing VRFs:**

1. Select Tenants -> VRF from side menu
2. Hover over the VRF you wish to modify, and click on the Update VRF (pencil) icon
3. Update the Maximum CPS value
4. Update the Maximum Sessions value

## Configuring the Maximum Sessions/CPS on a Network via the PSM UI

### **For New Networks:**

1. Select Network from panel (left)
2. Add Network
3. Select from one of the options (“Inherit from VRF” or “Unlimited” or “Set Value”)
4. Create Network

#### For Existing Networks:

1. Select Network from panel (left)
2. Select Network
3. Update Network
4. Select from one of the options (“Inherit from VRF” or “Unlimited” or “Set Value”)
5. Save

When “Inherit from VRF” is used, the Maximum CPS and Maximum Sessions will inherit the value from the value configured in the corresponding VRF. When Unlimited is used, no limits are enforced for Maximum sessions/ Max CPS. When “Set value” is used, the configured values will be applied and will also override the values configured at the corresponding VRF level.

The screenshot shows a configuration page for a VRF named 'Test-VRF-1'. The page has a header with a navigation bar containing 'Subnet-1001', 'Test-VRF-1', and '1001'. The main content area is divided into two columns. The left column contains fields for 'Name' (Subnet-1001), 'Ingress Security Policy' (Ingress Security Policy), 'VLAN ID' (1001), 'Maximum CPS' (1000), and 'Maximum Sessions' (2500). The right column contains fields for 'VRF' (Test-VRF-1), 'Egress Security Policy' (Egress Security Policy), and 'Maximum Sessions' (2500). Below the 'Maximum CPS' and 'Maximum Sessions' fields, there are three buttons: 'Inherit from VRF', 'Unlimited', and 'Set Value'. The 'Set Value' button is highlighted in blue. At the top right of the form, there are 'SAVE' and 'CANCEL' buttons.

Figure 89. Configuring session and CPS limits

## API Examples

An example REST API call to set maximum session limit/CPS for a given VRF:

PUT: [https://\\$PSMaddr/configs/network/v1/tenant/default/virtualrouters/Test-VRF-1](https://$PSMaddr/configs/network/v1/tenant/default/virtualrouters/Test-VRF-1)

Body:

```
{
  "kind": "VirtualRouter",
  "api-version": "v1",
  "meta": {
    "name": "Test-VRF-1",
```

```

    "tenant": "default"
  },
  "spec": {
    "maximum-sessions-per-network-per-distributed-services-entity": 400000,
    "maximum-cps-per-network-per-distributed-services-entity": 40000
  }
}

```

In the above example, a max session-limit of 400000 and max-CPS of 40000 is enforced on VRF Test-VRF-1.

An example REST API call to set maximum session limit/CPS for a given network:

PUT: [https://\\$PSMaddr/configs/network/v1/tenant/default/networks/Test-Subnet-1](https://$PSMaddr/configs/network/v1/tenant/default/networks/Test-Subnet-1)

Body:

```

{
  "kind": "Network",
  "api-version": "v1",
  "meta": {
    "name": "Test-Subnet-1",
    "tenant": "default"
  },
  "spec": {
    "vlan-id": 5,
    "virtual-router": "Test-VRF-1",
    "firewall-profile": {
      "maximum-cps-per-distributed-services-entity": 10000,
      "maximum-sessions-per-distributed-services-entity": 25000
    }
  }
}

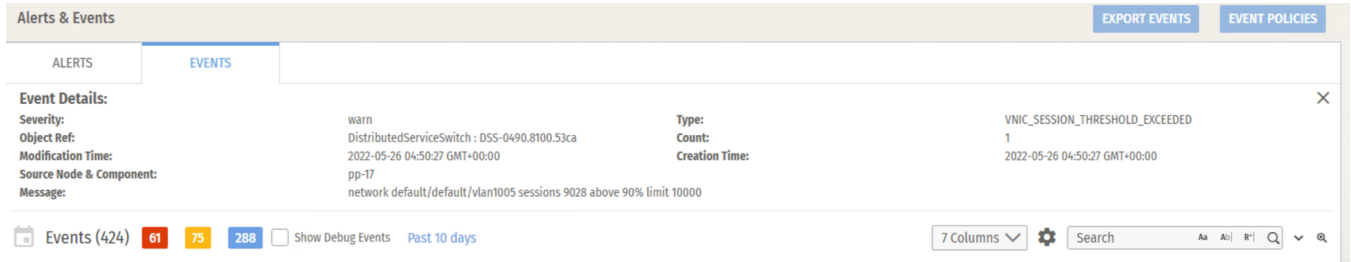
```

In the above example, a per-network limit of 25000 sessions and 10000 CPS is set. Events/Alerts Related To Maximum Configured Session-Limit

- As the flow table is populated, a warning event will be generated when 90% of the configured max limit is reached.
- A critical Alert will be generated for a VRF/network when it reaches 100% of its configured limit.
- When the session count drops below 80% of the configured limit, an INFO event is generated.



## Examples:

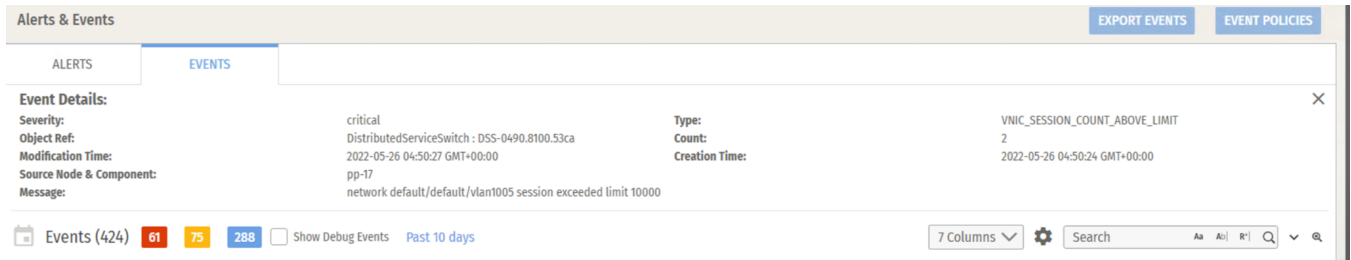


The screenshot shows the 'Alerts & Events' section with the 'EVENTS' tab selected. The event details are as follows:

Event Details:		Type:	
Severity:	warn	Type:	VNIC_SESSION_THRESHOLD_EXCEEDED
Object Ref:	DistributedServiceSwitch : DSS-0490.8100.53ca	Count:	1
Modification Time:	2022-05-26 04:50:27 GMT+00:00	Creation Time:	2022-05-26 04:50:27 GMT+00:00
Source Node & Component:	pp-17		
Message:	network default/default/vlan1005 sessions 9028 above 90% limit 10000		

At the bottom, the status bar shows 'Events (424)' with counts for 61 (red), 75 (yellow), and 288 (blue) events. There are also buttons for 'Show Debug Events' and 'Past 10 days'.

*Figure 90. A warning event generated when 90% of the configured limit is reached*

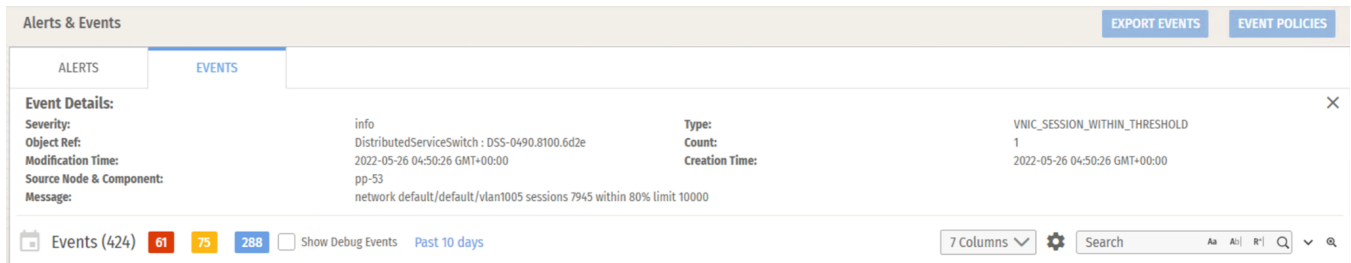


The screenshot shows the 'Alerts & Events' section with the 'EVENTS' tab selected. The event details are as follows:

Event Details:		Type:	
Severity:	critical	Type:	VNIC_SESSION_COUNT_ABOVE_LIMIT
Object Ref:	DistributedServiceSwitch : DSS-0490.8100.53ca	Count:	2
Modification Time:	2022-05-26 04:50:27 GMT+00:00	Creation Time:	2022-05-26 04:50:24 GMT+00:00
Source Node & Component:	pp-17		
Message:	network default/default/vlan1005 session exceeded limit 10000		

The status bar at the bottom is identical to Figure 90, showing 'Events (424)' with 61, 75, and 288 event counts.

*Figure 91. A critical event generated when a limit is reached*



The screenshot shows the 'Alerts & Events' section with the 'EVENTS' tab selected. The event details are as follows:

Event Details:		Type:	
Severity:	info	Type:	VNIC_SESSION_WITHIN_THRESHOLD
Object Ref:	DistributedServiceSwitch : DSS-0490.8100.6d2e	Count:	1
Modification Time:	2022-05-26 04:50:26 GMT+00:00	Creation Time:	2022-05-26 04:50:26 GMT+00:00
Source Node & Component:	pp-53		
Message:	network default/default/vlan1005 sessions 7945 within 80% limit 10000		

The status bar at the bottom is identical to the previous figures, showing 'Events (424)' with 61, 75, and 288 event counts.

*Figure 92. An INFO event generated when the count returns to below 80%.*

## Behavior on Reaching the Maximum Session Limit

- Flow learning (learning of new flows) will not be stopped at exactly 100% of the configured limit. An excess of 5% is allowed, so flow learning is completely stopped at 105% of the configured limit value and a critical alert will be reported to the PSM.

## Behavior on Reaching the Maximum CPS Limit

When the incoming CPS/flow-create/SYN rate is higher than the configured rate, then incoming connections will start getting dropped (or policed) at the rate defined by the configuration. This policing action is done by the P4 engine itself. However, the event that gets generated will get generated only on exceeding the configured limit. No warning events for 90% of configured CPS are reported to the PSM. Also, note that it could take up to 2 minutes for the MAX CPS exceed event to get reported to the PSM.

Example of event reported at PSM:

The screenshot shows the Pensando PSM web interface. The left sidebar contains navigation links: Dashboard, System, Tenants, Overview, VRF, Networks, Security Policies, Apps, Firewall Export Policies, Firewall Profiles, Firewall Logs, Monitoring, Alerts & Events (selected), Archive Logs, Metrics, Audit Events, Troubleshoot, and Admin. The main content area is titled 'Alerts & Events' and has tabs for 'ALERTS' and 'EVENTS'. An 'Alert Details' pop-up is visible, showing the following information:


- Severity:** critical
- Modification Time:** 2022-06-01 13:27:53 GMT+00:00
- Alert Type:** stats
- Source Node & Component:** Ref: 0490.8100.53ca Reporter: 0490.8100.53ca - pen-turret
- Message:** difference value of counter [VnicDrops:CoPPDrops] is above the threshold of 1 on 0490.8100.53ca/vlan1003 [DSS-0490.8100.53ca] host: pp-17
- State:** open
- Creation Time:** 2022-06-01 13:27:53 GMT+00:00
- Total Hits:** 1

Below the alert details, there is a table of alerts. The table has columns: Severity, Message, State, Alert Type, Source Node & Component, Total Hits, Creation Time, and Modification Time. The first row shows a critical alert with the message 'difference value of counter [VnicDrops:CoPPDrops] is above the threshold of ...', state 'Open', alert type 'stats', and a total hit of 1.

At the bottom of the interface, the status bar shows 'Pensando Systems' on the left and 'Activate Windows' and 'Server Time: 2022-06-01 13:55:28 UTC' on the right. The Windows taskbar at the very bottom shows the time as 6:28 AM on 6/1/2022.

Figure 93. Example of an event generated upon reaching MAX CPS limit

In addition to checking for above alert, PSM users can create a metric chart and look for CoPP drops (see Figure 94):



*Figure 94. Example of Flow Miss CoPP Drops graph;  
Max CPS exceeded*

## Implication of Configuring `session-limit` on an Active System

Configuring the `session-limit` value on a live system (a switch with active flows) for any given VRF/network does not have any traffic impact or disruption of existing active flows.

Active flows are not purged if the new configured value is less than the total number of active sessions for that VRF/network. The change or the new limit will only apply to new flows created after the last config change.

## VSX Implications

Important considerations in VSX topology with regard to this configuration:

- Each switch in VSX independently applies the configuration at a per-VRF /network level. The PSM will push identical VRF/network configurations to both switches in VSX, so



configuration consistency is expected and independently enforced without the need to manually coordinate between peers.

- Both flows locally learned by VSX switch and flows learned through the process of flow-sync from VSX peer are accounted towards the calculation of max sessions before enforcing the max session-limit.

Consider the following examples:

**CASE 1:** A and B in VSX pair and switch A gets all the new flows, no flow hashes to switch B.

Configured limit is 10000 sessions on network/VRF.

Switch A can only learn a max of 10000 sessions. It will sync 10000 sessions to B. Both A and B are at max session limit of 10000 and cannot accept any more flows.

**CASE 2:** A and B are in VSX pair and there is equal distribution of flows from host to switch

Configured limit is 10,000 sessions on network/VRF.

In this case, each switch Switch A and Switch B can learn only up to 5K new flows from hosts since the remaining 5k flows are learnt from flow-sync from the respective VSX peer.

- The max-CPS limit is independently applied on each switch in a VSX pair. This is illustrated in the following examples:

**CASE 1:** A and B in VSX pair and switch A gets all the new flows, no flow hashes to switch B.

Application sends traffic at a rate of 20,000 CPS. A limit of 5,000 CPS is applied on both switches A and B.

All traffic hashes to switch A and gets throttled to 5KCPS. This 5KCPS will be flow-synced to switch B without further throttling on switch B.

Effective CPS achieved: 5K CPS.

**CASE 2:** A and B are in VSX pair, flows hash evenly to switch A and switch B.

Application sends traffic at a rate of 20K CPS. A limit of 5K CPS is applied to both switches A and B.

The 10K CPS traffic that hashes to A will get throttled to 5K CPS and the 10K CPS traffic that gets hashed to switch B will also get throttled to 5K CPS. Each switch will

flow sync its 5K CPS to the other peer. No further throttling takes place when it is synced to peer.

Effective CPS achieved: 10K CPS.

## Multiple ALG Types/Apps/Protocols in Firewall Policy Rules

It is possible to match on multiple ALG types / multiple apps / multiple IP protocol types (TCP/UDP/ICMP/GRE/AH/ESP) within a single rule inside a PSM policy. One or more of the defined ALG types/apps/IP protocols can be used. This section provides examples of defining such a rule via the PSM UI, as well as API examples.

### Steps To Configure Rules With Multiple Proto-Ports Via the PSM UI

1. Security Policies (left panel)
2. Add Security Policy
3. Select Proto-Ports

The screenshot displays the 'Network Security Policies' configuration window. At the top, there's a 'Table View' dropdown and buttons for 'CREATE POLICY' and 'CANCEL'. The main area is titled 'Network Security Policies' and contains a 'Policy Name' field with the value 'test-multiple-protocols'. Below this, a 'RULES: (specify one or more)' section is highlighted with a blue border. Inside this section, a 'Rule Name' field contains 'r1'. The 'Action' dropdown is set to 'Permit'. A 'Description' field is empty. The 'Source IP Addresses' field contains '1.1.1/32', and the 'Destination IP Addresses' field contains '2.2.2/32'. Under the 'PROTO-PORTS' section, several protocols are listed: 'tcp', 'udp', 'icmp', 'gre', 'ah', and 'esp'. Each protocol has a corresponding 'Ports' field. For 'tcp', the ports are '1001-1002, 1004, 1005'. For 'udp', the ports are '2001-2002, 2012, 2034'. The other protocols have empty 'Ports' fields. At the bottom of the rule configuration, there is a '+ Protocols and Ports' button. The bottom right corner of the window shows 'Activate Windows'.

Figure 95. UI example of rules with multiple proto-ports

As shown above, a mix of TCP ports, ID ports and protocol types of ICMP, GRE, IPsec (AH/ESP) can be referenced within a single rule.

A rule can also reference multiple ALG types, as shown below:

### Steps to configure rules with multiple ALG types via the PSM UI:

1. Security Policies (left panel)
2. Add Security Policy
3. Select Apps
4. From drop-down menu choose multiple defined ALG types from list

The screenshot shows the 'Network Security Policies' configuration page. At the top, there's a 'Table View' dropdown, 'CREATE POLICY', and 'CANCEL' buttons. Below, the 'Policy Name' is 'multi-apps'. The 'RULES: (specify one or more)' section contains a rule named 'r1' with a 'Permit' action and a 'Disabled' toggle. The 'Source IP Addresses' field contains '1.1.1.1/32' and the 'Destination IP Addresses' field contains '2.2.2.2/32'. Under the 'APPS' section, '5 items selected' is shown in a dropdown menu. A tooltip lists the selected apps: FTP, RTSP, SUNRPC, and TFTP. At the bottom, a table lists 'Network Security Policies (5)' with columns for Policy Name, Policy/Rule Entries Consumed per DSS, and Associated Network. To the right, there's a search bar and a table with columns: Networks (Egress Policy), Associated VRFs (Ingress Policy), Associated VRFs (Egress Policy), and Propagation Status.

Figure 96. UI example of rules with multiple ALG types

**NOTE:** A policy rule can either reference multiple proto-ports (list of TCP and list of UDP ports) or multiple apps/ALG types but NOT a mix of both proto-ports and app/ALG types

### API Examples:

This section provides API examples to create a policy with rules referencing either multiple protocol ports or multiple apps.

Example:

POST: [https://\\$PSMADDR/configs/security/v1/tenant/default/networksecuritypolicies](https://$PSMADDR/configs/security/v1/tenant/default/networksecuritypolicies)

Request body:

```
{
  "kind": "NetworkSecurityPolicy",
```

```

"api-version": "v1",
"meta": {
  "name": "rule-with-multiple-protocols",
  "tenant": "default"
},
"spec": {
  "attach-tenant": true,
  "rules": [
    {
      "proto-ports": [
        {
          "protocol": "tcp",
          "ports": "1001-1002,2001,3001"
        },
        {
          "protocol": "udp",
          "ports": "2001-2002,3003,4004"
        },
        {
          "protocol": "icmp"
        },
        {
          "protocol": "gre"
        },
        {
          "protocol": "ah"
        },
        {
          "protocol": "esp"
        }
      ],
      "action": "permit",
      "from-ip-addresses": [
        "1.1.1.1/32"
      ],
      "to-ip-addresses": [
        "2.2.2.2/32"
      ]
    }
  ]
}

```

```
    ]
  }
}
```

## API Example: Policy with Rule Referencing Multiple ALG Types:

```
{
  "kind": "NetworkSecurityPolicy",
  "api-version": "v1",
  "meta": {
    "name": "rule-with-multiple-alg-types",
    "tenant": "default"
  },
  "spec": {
    "attach-tenant": true,
    "rules": [
      {
        "apps": [
          "FTP",
          "DNS"
        ],
        "action": "permit",
        "from-ip-addresses": [
          "1.1.1.1/32"
        ],
        "to-ip-addresses": [
          "2.2.2.2/32"
        ]
      }
    ]
  }
}
```

## IP Protocols Support for Firewall Policy

Any IP protocol can be matched inside of a firewall policy rule. The rule can either match on reserved keywords like “gre”, “ah” and “esp” for GRE, IPsec-AH and IPsec-ESP protocols respectively or it can match on any valid numeric protocol number in the range of 0-254.

The following section gives examples for doing this using the PSM UI as well as the REST API.

### UI Examples

The screenshot displays the 'Network Security Policies' management interface. At the top, there's a header bar with 'Table View' and buttons for 'CREATE POLICY' and 'CANCEL'. Below the header, the main area is titled 'Network Security Policies'. A 'Policy Name' field contains 'policy-match-on-ip-protocols'. Under the 'RULES: (specify one or more)' section, a rule named 'r1' is configured. The rule's action is set to 'Permit'. The 'Source IP Addresses' field contains '1.1.1.1/32' and the 'Destination IP Addresses' field contains '2.2.2.2/32'. The 'PROTO-PORTS' tab is selected, showing a list of protocols: 1, 2, and 89. Each protocol has a corresponding 'Ports' field. A '+ ADD' button is located at the bottom of the rule configuration area.

*Figure 97. UI example of rules with multiple IP protocol numbers*

The above figure shows an example of a policy created with a rule referencing a combination of multiple IP protocol numbers.

Network Security Policies Table View CREATE POLICY CANCEL

Network Security Policies

Policy Name  
policy-match-on-ip-protocols

RULES: (specify one or more)

Rule Name  
r1

Action  
Permit

Description

Source IP Addresses  
1.1.1/32

Destination IP Addresses  
2.2.2/32

☒ PROTO-PORTS

☐ APPS

Protocol  
gre

Ports

Protocol  
ah

Ports

Protocol  
esp

Ports

+ Protocols and Ports

Figure 98. UI example of rules with multiple named IP protocols

The above figure shows an example where policy is created with a rule referencing a combination of named IP protocols such as GRE, AH or ESP (IPsec protocols).

## API Examples

The following creates a policy with a rule referencing multiple IP protocol numbers:

```
{
  "kind": "NetworkSecurityPolicy",
  "api-version": "v1",
  "meta": {
    "name": "rule-with-proto-numbers",
    "tenant": "default"
  },
  "spec": {
    "attach-tenant": true,
    "rules": [
      {
        "proto-ports": [
          {
            "protocol": "10"
          },
          {
            "protocol": "254"
          }
        ]
      }
    ]
  }
}
```

```

        },
        {
            "protocol": "67"
        },
        {
            "protocol": "202"
        },
        {
            "protocol": "143"
        },
        {
            "protocol": "199"
        }
    ],
    "action": "permit",
    "from-ip-addresses": [
        "1.1.1.1/32"
    ],
    "to-ip-addresses": [
        "2.2.2.2/32"
    ]
}
]
}
}
}

```



## Stateful Firewall Flow Migration with vMotion

In the event of a vMotion migration of VMs between two ESX hosts directly connected to a DSS (or a VSX pair of DSSes), PSM integration with vCenter allows firewall flows to be statefully migrated from the original connected DSS to the new connected DSS.

The vMotion migration itself is transparent to the hosted application, as the stateful nature of flows are preserved, with applications experiencing very minimal loss during the event. The intent for flow migration is to allow connection tracking to be enabled and to preserve associated features for flows, such as flow statistics and flow logs, following a vMotion event. To support vMotion of VMs between ESX hosts connected to DSSes in the cluster at steady state, some additional information is collected from vCenter and AOS-CX.

Prerequisites:

1. All ESX hosts **must** be directly connected to a DSS, as this feature relies on LLDP for locality resolution.
2. vMotion will be supported only within the same PSM cluster: The source DSS and the destination DSS must be controlled by the same PSM cluster.
3. All ESX hosts **must** have LLDP enabled and LLDP **must** be enabled on all DSS ports.

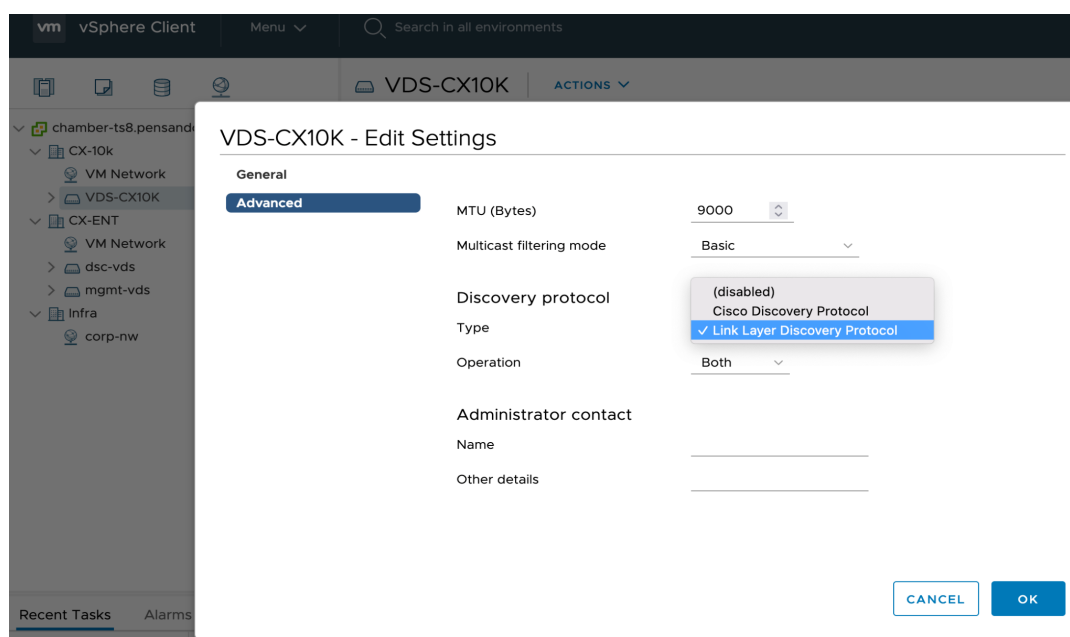


Figure 99. Enabling LLDP on a vDS

4. All DSSes **must** be reachable by each other via the default VRF in-band.
5. Each DSM is available on a different TCP port. They share the same common inband IP of the DSS.

- a. NAT rules are programmed in the DSS for each DSM using the reserved L4 ports (11400-11450 for DSM1 and 11451-11502 for DSM2)
6. LLDP is not supported on standard vSwitches. A workload's network interfaces **must** be attached to distributed virtual switch (DVS) if stateful firewall vMotion support is needed.
7. vCenter versions 6.7 and 7.0 are supported.

## Considerations for Multi-Homed ESX Servers

The DSS can be deployed as VSX pairs where the flows are synchronized between the pair. Hosts can then be dual-homed to the pair with the multi-chassis LAG. These hosts can be running ESX. VMs can then move from any combination of VSX pair and single homed host; for example, a single homed host to a VSX pair and back, or a VSX pair to a VSX pair. All such scenarios are supported. If the source is a VSX pair, then the flows are pulled from any one of the source DSS switches and then locally deleted on both. If the destination is a VSX pair, then the primary switch is used for the flow synch; the flows will then be synched to the secondary switch. vMotion for VMs in a PVLAN is now supported.

## Behavior with Flow Logs and Flow Statistics

When a flow is migrated from the source DSS to the destination DSS:

1. A flow log OPEN packet is sent from the destination DSS once flow migration completes and the flow is installed. The statistics of the migrated flow are reset to 0—statistics are not carried over during flow migration.
2. A flow log CLOSE message is sent from the source DSS once flow migration completes, and will include the flow's I-flow and R-flow packet/byte statistics.

When the flow eventually is removed/aged out on its new location on the destination DSS, another CLOSE packet is sent with I-flow and R-flow statistics that correspond to statistics on the destination DSS. The flow log CLOSE packet that is sent from the destination DSS post-vMotion will not have the original policy-name and the rule-name contained in the logs.

## Configuration

1. The PSM connects to vCenter via user-provided credentials. There is an option to monitor either all or specific datacenters on the vCenter and receive notifications as appropriate.

The screenshot shows the 'vCenter Overview' configuration window. It includes fields for Name (TPM-vCenter), URI (10.29.192.51), and a Credential section with User Name (administrator@penlabs.net) and Password (masked with asterisks). There is a 'Disable Server Authentication' toggle switch which is currently turned on. At the bottom, there are two radio buttons: 'Monitor All Datacenters' (selected) and 'Monitor Individual Datacenters'.

Figure 100. Add a connection to vCenter

2. The PSM reads the Host, Workloads and LLDP information from vCenter via the VMware VIM API, and is notified by vCenter of vMotion-related events.

The screenshot shows the 'vCenter Overview' window with a table titled 'vCenters (1)'. The table has 8 columns: Name, FQDN, Datacenters, Workloads, Orchestrator ID, Status, Connection Transition Time, and Modification Time. The first row shows 'TPM-Vcenter' with FQDN '10.29.192.51', Datacenters 'All Datacenters', and a list of workloads including demo-app02, demo-app01, pod1-nginx-opencart, demo-web01, pod1-web2-opencart, demo-db, and Tower. The Status is 'Connection: Success' and 'Discovered Datacenters: Datacenter'. The Connection Transition Time is '2022-07-27 14:53:16 GMT+00:00' and the Modification Time is '2022-07-27 14:53:21 GMT+00:00'.

Name	FQDN	Datacenters	Workloads	Orchestrator ID	Status	Connection Transition Time	Modification Time
TPM-Vcenter	10.29.192.51	All Datacenters	demo-app02(orch1-vm-15) demo-app01(orch1-vm-16) pod1-nginx-opencart(orch1-vm-17) demo-web01(orch1-vm-18) pod1-web2-opencart(orch1-vm-19) demo-db(orch1-vm-20) Tower(orch1-vm-21) ...> 20 More	1	Connection: Success Discovered Datacenters: Datacenter	2022-07-27 14:53:16 GMT+00:00	2022-07-27 14:53:21 GMT+00:00

Figure 101. All the workloads imported by the PSM from vCenter

- Information received from vCenter is used to create the ESXi Host to local DSS mapping.

Workloads Overview						
Workloads						
Workloads (28)		6 Columns		Search		
<input type="checkbox"/>	Workload Name	Host Name	Switch & Port	Interfaces	Labels	Modification Time
<input type="checkbox"/>	Systest-Vlan2063-VM1	10.29.192.113	tpm-leaf3 1/1/1 tpm-leaf4 1/1/1	IP: 51.0.12.10 VLAN: 2063 Network: DP-Systest-2063	DataCenterName: TPM-DC, VcenterName: TPM-Vcenter	2022-08-01 15:16:49 GMT+00:00
<input type="checkbox"/>	Systest-Vlan2063-VM2	10.29.192.112	tpm-leaf2 1/1/2 tpm-leaf1 1/1/2	IP: 51.0.34.10 VLAN: 2063 Network: DP-Systest-2063	DataCenterName: TPM-DC, VcenterName: TPM-Vcenter, vcenter.Environment: Production	2022-08-01 15:16:05 GMT+00:00

Figure 102. Two workloads and their local mappings and tags/labels

- This database of ESC and workload/VM locality is maintained in the PSM by correlating the Host Object (ESXi Node) information and the LLDP information populated in the DSS Object.

Workloads Overview > Systest-Vlan2063-VM2		
WORKLOAD DETAILS		
Name: Systest-Vlan2063-VM2 orch1--vm-59	Last Modified: 2022-08-01 15:16:05 GMT+00:00 Created: 2022-07-29 08:36:49 GMT+00:00	Labels
Host: 10.29.192.112 orch1--host-27		Key Value
Switch: tpm-leaf2 1/1/2 tpm-leaf1 1/1/2		DataCenterName TPM-DC
		VcenterName TPM-Vcenter
		vcenter.Environment Production
INTERFACES		
MAC Address	IP Addresses	External VLAN
0050.568e.3b36	51.0.34.10	2063

Figure 103. A workload and its local specific mappings

## DSS Required AOS-CX CLI Configuration

Neighbor resolution must be enabled using the following AOS-CX commands on each DSS:

- MAC-IP Bindings
  - New CLI to trigger the ARP snooping on DSS.

```
switch(config)# dsm
switch(config-dsm)# workload-migration
```

- DSM-DSM communication for vMotion flow migration
  - DSM will communicate with DSS via bond interfaces created during init in vlan 4093.
  - **Note:** The following CLI configuration is *required* on each DSS to create the NAT rules to allow sessions establishment between the source and destination DSMs:



```
switch(config)# ip source-interface workload_migration
interface vlxanX (or routeable loopback X)
```

## Caveats

- Even though vCenter can be configured to support up to 8 parallel vMotion migrations, a DSS will process one flow migration at a time; all the vMotion move-in and move-out requests received are queued while processing the existing request, and then processed sequentially.
- Port Group VLAN changes between the source DVS and the destination DVS are not supported for the Migrating VM, even though vCenter allows it. In other words, the DPG on the source and destination DVS must have the same VLAN-ID.
- When a vMotion migration occurs for a workload from behind a non-DSS connected host to a DSS connected host, existing TCP applications may experience packet loss and re-transmissions as all mid-stream non-SYN data packets will be dropped. The existing TCP connections will time out and connections need to be re-established.
- The PSM can be configured with up to four vCenters. However, flow migration is only supported when vMotion occurs between hosts managed by the same vCenter.

## Monitoring the DSM via the PSM UI

The PSM dashboard page provides an overview of the total number of CX 10000 switches in the fabric, the number pending admission, number rejected and a health readout.

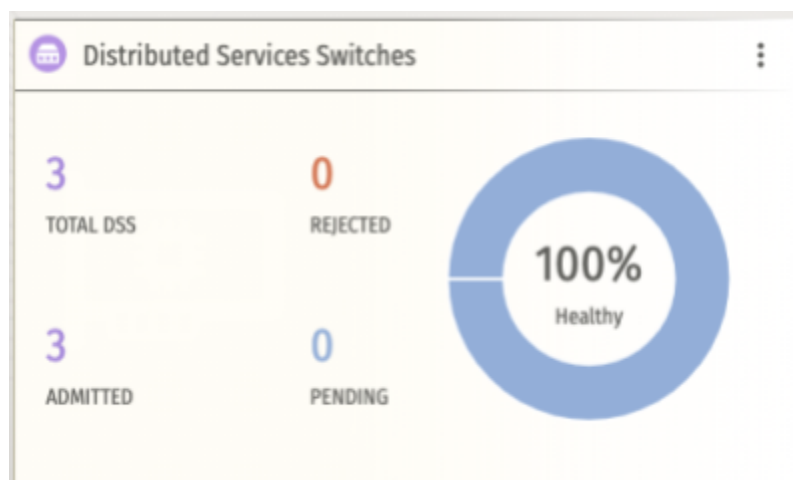


Figure 104. Summary card for CX 10000 switches from the PSM dashboard:

TOTAL DSS: Number of DSSes attached to this PSM

ADMITTED: DSSes admitted to this PSM

REJECTED: DSSes that have been rejected from joining this PSM quorum

PENDING: DSSes pending admission

The Distributed Services Switches Overview page, which can be reached by clicking on the menu symbol (⋮) and choosing `Navigate to DSSes`, lists top-level information, including the health of CX 10000 switches admitted to the cluster. Health status reflects the health and reachability of the DSM present in the DSS.

## Distributed Services Switches Overview

Distributed Services Switches (6)							
<input type="checkbox"/>	Switch Name	DSS-ID	DSM	Switch OS Version	Management	Health	Labels
<input type="checkbox"/>	P2-10	DSS-0490.8100.4e80	DSM 1/1: 0490.8100.4ee0 DSM 1/2: 0490.8100.4ef0	DL.10.09.0002F	Mode: Inband IP address: 10.30.22.14		Belong: Vijay, Board: P2-10
<input type="checkbox"/>	P2-11	DSS-0490.8100.4a00	DSM 1/1: 0490.8100.4a60 DSM 1/2: 0490.8100.4a70	DL.10.09.0002F	Mode: Inband IP address: 10.30.22.15		Belong: Vijay, Board: P2-11
<input type="checkbox"/>	P2-09	DSS-0490.8100.4800	DSM 1/1: 0490.8100.4860 DSM 1/2: 0490.8100.4870	DL.10.09.0002F	Mode: Inband IP address: 10.30.22.13		Belong: Vijay, Board: P2-09

*Figure 105. Distributed Services Switches Overview list*

Switch Name: *Name name of the DSS. Clicking on this will open a link to that DSS's Web UI.*

DSS-ID: *ID of the DSS, used in metrics and logs; derived from switch MAC address. Clicking on this will switch to this DSS's detail page.*

DSM: *Unique MAC address IDs of the 2 DSMs in the DSS*

Switch OS Version: *AOS-CX version the switch is running*

Management: *Mode used to communicate between DSM and PSM*

Health: *Health and admission status of the DSS*

Labels: *Admin-definable label field*

Clicking on the DSS-ID of any admitted CX 10000 in this table will switch to its Distributed Services Switch detail page, as shown in Figure 106, which provides an overview of the two DSMs in the CX 10000, and any specific Events/Alerts related to the DSMs.

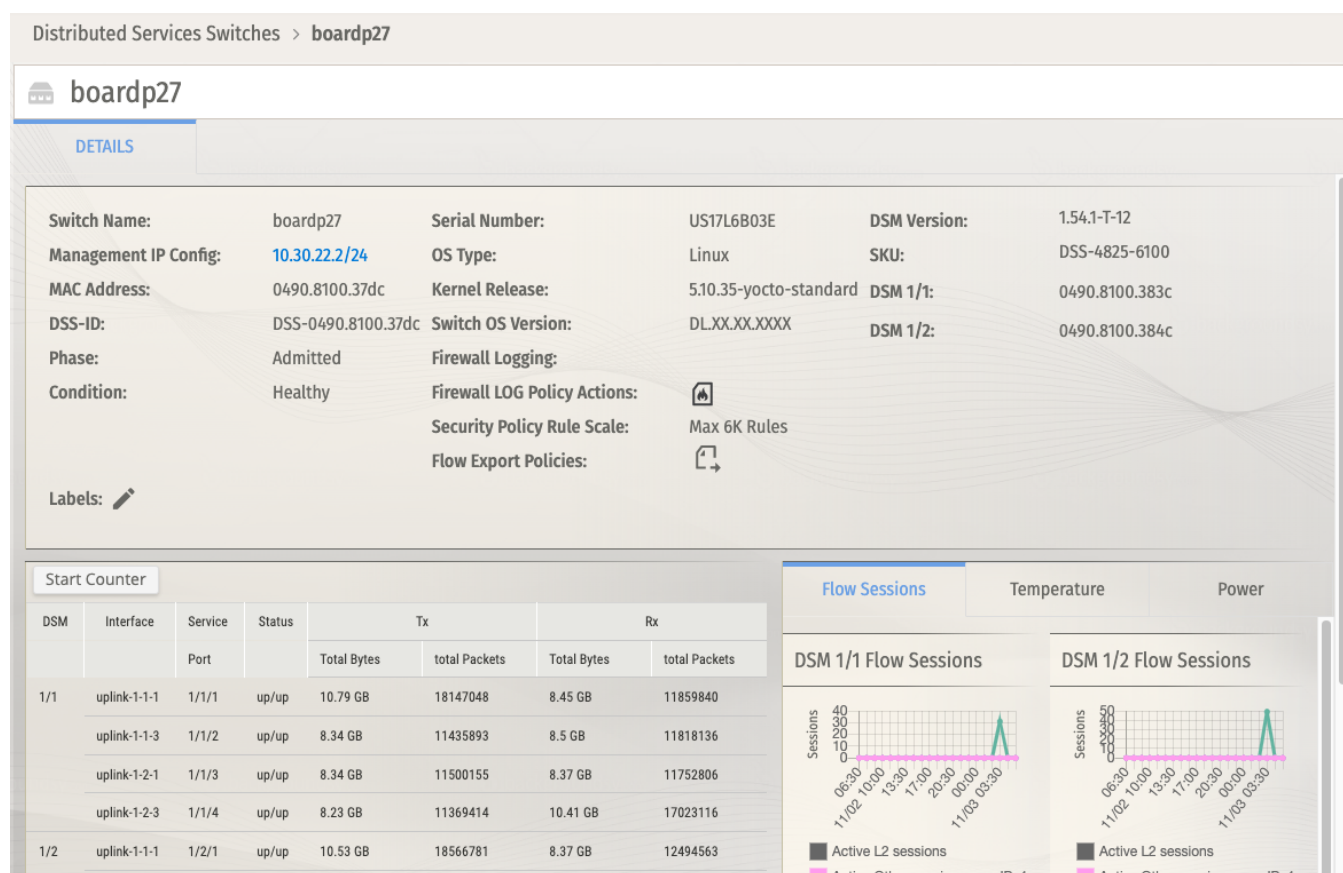


Figure 106. CX 10000 detail page

Alerts and events for all DSSes admitted to the PSM can be viewed from the Monitoring tab.

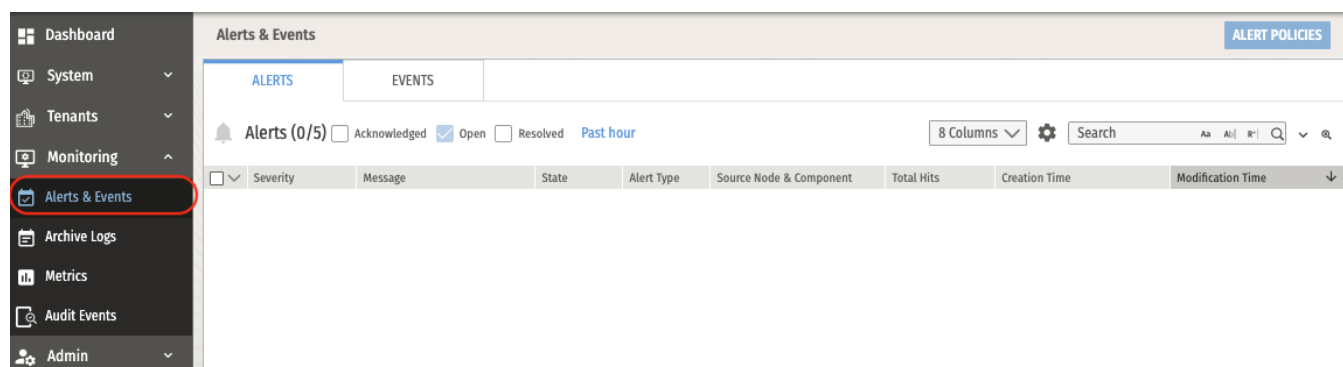


Figure 107. Monitoring -> Alerts & Events screen



There is also an option for viewing audit logs related to the DSMs and the PSM. Audit logs are a history of configuration changes related to the PSM and the DSM.

Who	Action	Resource Tenant	Act On (kind)	Act On (name)	Outcome	Client	Service Node	Service Name	Modification Time
admin	Update	default	Alert	4ceaa12a-07b4-...	success	192.168.32.3	dss_venice_ova1	pen-apiserver	2021-09-24 21:38:34 GMT+00:00
admin	Update	default	Alert	16e163fc-c720-4...	success	192.168.32.3	dss_venice_ova1	pen-apiserver	2021-09-24 21:37:59 GMT+00:00
admin	Update	default	Alert	aebe4f47-b3b0-4...	success	192.168.32.3	dss_venice_ova1	pen-apiserver	2021-09-24 21:37:48 GMT+00:00
admin	Update	default	Alert	f2ab7a39-903a-...	success	192.168.32.3	dss_venice_ova1	pen-apiserver	2021-09-24 21:37:37 GMT+00:00
admin	Update	default	Alert	ae4c5015-e247-...	success	192.168.32.3	dss_venice_ova1	pen-apiserver	2021-09-24 21:37:31 GMT+00:00
admin	Update	default	FirewallProfile	default	success	192.168.66.9	dss_venice_ova1	pen-apiserver	2021-09-24 21:32:39 GMT+00:00

Figure 108. Monitoring -> Audit Events screen

## Metrics Charts

Metrics charts can be created for various switch statistics. Select Metrics from the left menu. On the Metrics screen, select **CREATE CHART**:

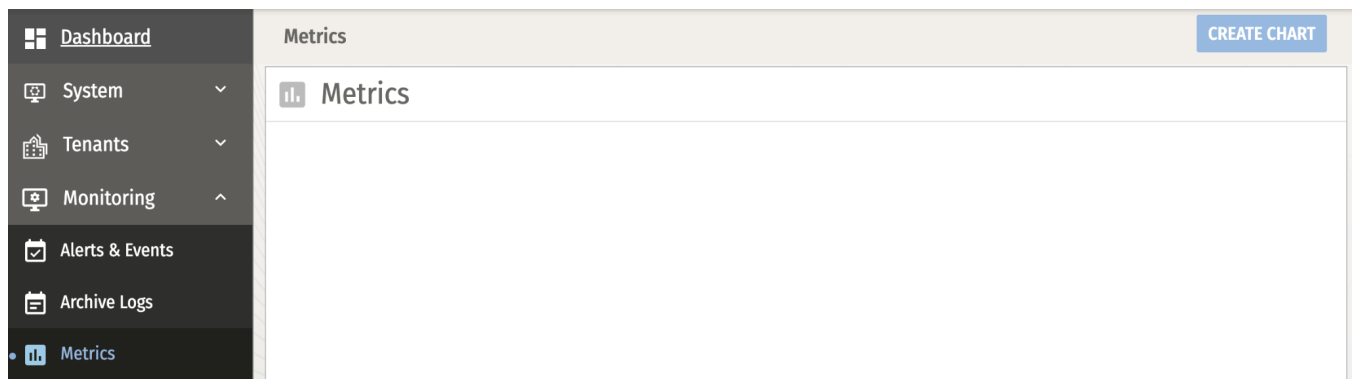
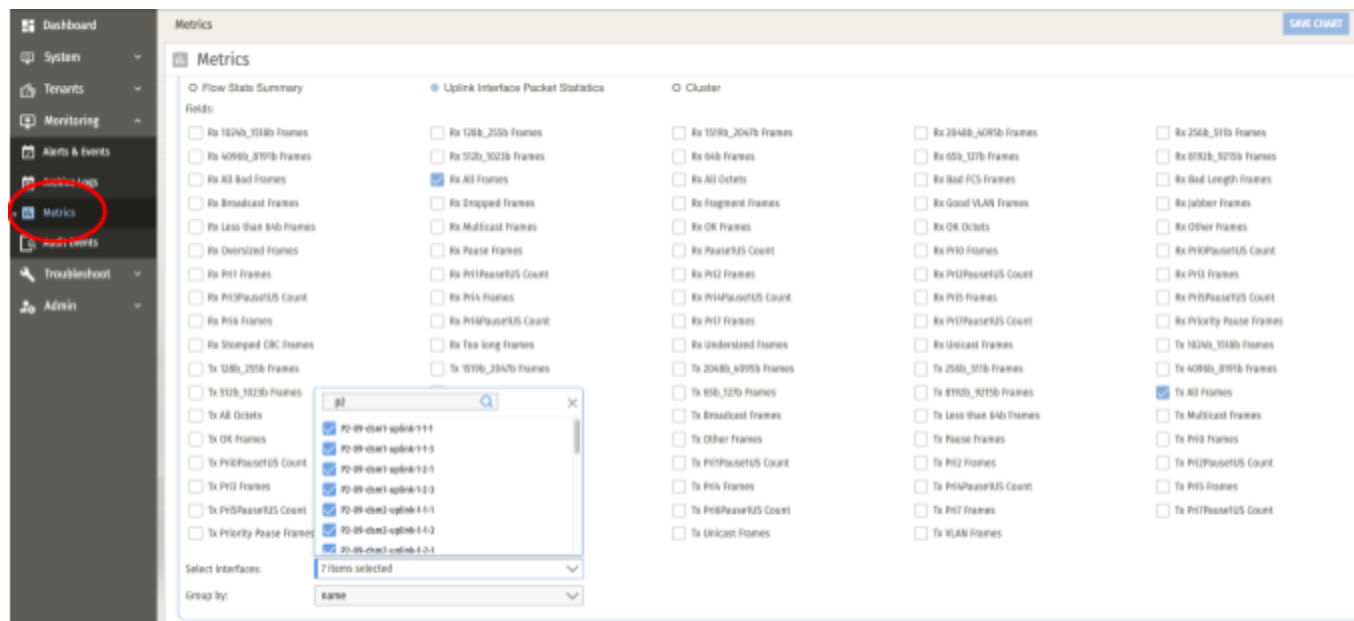



Figure 109. Metrics screen showing CREATE CHART button

Here, the Tx/Rx frames are selected:



PSM for Aruba CX 10000 User Guide

Once the chart definition is saved, a graph is created:



*Figure 112. The ID used for selecting interfaces is the CX 10000 MAC address.  
The ID shown in the result is the two DSMs that are part of the CX 10000.  
First DSM MAC (reporterID)  $\Rightarrow$  CX 10000 system MAC + 96.  
Second DSM MAC (reporterID)  $\Rightarrow$  CX 10000 system MAC + 112.*

Any chart can be “pinned”, which will place it on the PSM dashboard. To pin a chart, hover over it, and a selection of action icons will appear.

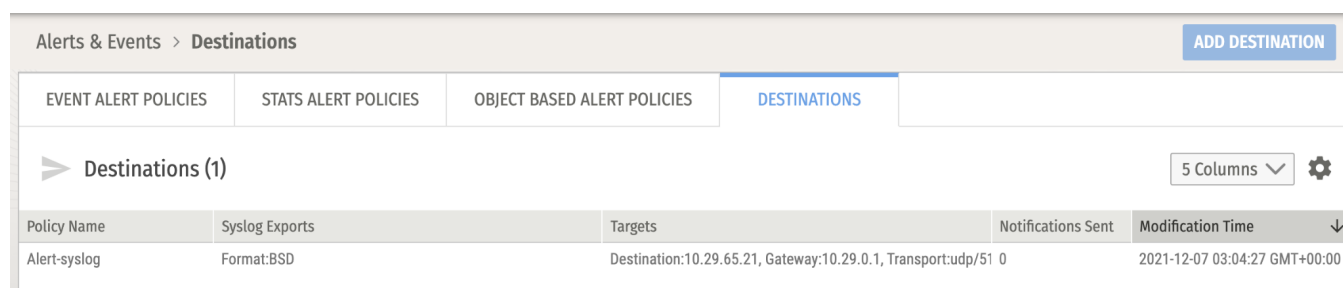


*Figure 113. An example chart, showing the pin, edit, and delete controls*

Select the pin (📌) icon to pin this chart. The icon will change to a white outlined pin (📌) to indicate its pinned status; you can unpin it by selecting it again.

## Alerts and Events

There are two types of alerts that may be managed by policy. Stats Alerts are triggered based on min/max thresholds for given statistics/metrics. Object-Based Alerts are triggered based on attribute values of a given object. Once policies are created for Stats and Object-Based Alerts, the resulting alerts get sent to one of the syslog-based destinations that have been defined.

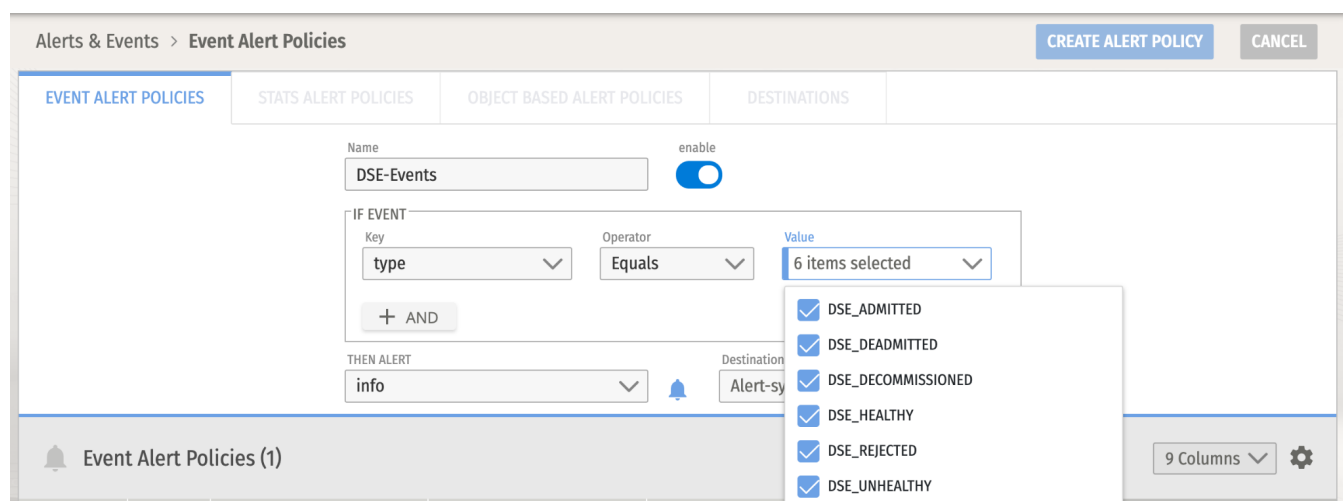


Alerts & Events > Destinations				
<a href="#">EVENT ALERT POLICIES</a> <a href="#">STATS ALERT POLICIES</a> <a href="#">OBJECT BASED ALERT POLICIES</a> <a href="#">DESTINATIONS</a>				
<div>Destinations (1)</div> <div>5 Columns</div>				
Policy Name	Syslog Exports	Targets	Notifications Sent	Modification Time
Alert-syslog	Format:BSD	Destination:10.29.65.21, Gateway:10.29.0.1, Transport:udp/51 0		2021-12-07 03:04:27 GMT+00:00

Figure 114. Viewing the list of defined alert destinations

First, select the “Destinations” tab and click “Add Destination” to configure a syslog collector that should receive the Alerts and Events.

Next, select the “Event Alert Policies” tab and click “Add Alert Policy” to provide criteria for an event to be captured. In the example below, an event is sent to the destination for any of the DSE actions. Event alert conditions can be combined into a single Alert Policy (“+AND”), or be created as individual Alert Policies.



Alerts & Events > Event Alert Policies

[CREATE ALERT POLICY](#)
[CANCEL](#)

[EVENT ALERT POLICIES](#)
[STATS ALERT POLICIES](#)
[OBJECT BASED ALERT POLICIES](#)
[DESTINATIONS](#)

Name: DSE-Events
enable: ☒

IF EVENT

Key: type
Operator: Equals
Value: 6 items selected

+ AND

THEN ALERT

info
Destination: Alert-syslog

Event Alert Policies (1)

9 Columns

☒ DSE\_ADMITTED  
☒ DSE\_DEADMITTED  
☒ DSE\_DECOMMISSIONED  
☒ DSE\_HEALTHY  
☒ DSE\_REJECTED  
☒ DSE\_UNHEALTHY

Figure 115. Adding an Alert Policy

## PSM Automation

PSM automation and programmability can be controlled via the following methods:

- Python language bindings
- Ansible modules
- REST API

### Python Language Bindings

Python language bindings are available through `pypi.org` and can be installed easily via the `pip` utility.

The pypi.org link is : <https://pypi.org/project/pensando-dss/>

The Python bindings can be easily installed via:

```
# pip install pensando-dss
```

where “dss” indicates the bindings specific to distributed services switches.

Documentation for the Python bindings is provided in the GitHub repo:

<https://github.com/pensando/pypi>

Please refer to the top-level [README](#) file.

Documentation for the objects are documented at :

[https://github.com/pensando/pypi/tree/main/src\\_dss/pensando\\_dss](https://github.com/pensando/pypi/tree/main/src_dss/pensando_dss) and include working code examples.

All the code in the AMD Pensando PyPi repo is automatically code-generated from the PSM’s Swagger specification.

### Ansible Modules

A number of Ansible modules are available at :

- <https://gitlab.com/pensando/tbd/ansible>
- <https://galaxy.ansible.com/pensando>

Unlike the Python bindings, the Ansible modules are not auto-generated.

Examples for the corresponding playbooks are included in the module sources. The following example corresponds to creating a network security policy:

```
---
- hosts: localhost
  connection: local
  gather_facts: no

  vars:
    psm_ip: 203.0.113.42
    psm_username: '{{ lookup("env", "PSM_USER") }}'
    psm_passwd: '{{ lookup("env", "PSM_PASSWORD") }}'

  tasks:

    - name: Test Network Security Policy
      pnso_network_security_policy:
        state: present
        policy_name: test_policy_new
        tenant: default
        psm_user: '{{ psm_username }}'
        psm_password: '{{ psm_passwd }}'
        psm_host: '{{ psm_ip }}'
        attach_tenant: true
        rules:
          - proto_ports:
              - protocol: udp
                ports: '111-222'
            action: permit
            from_ip_addresses: [ 1.1.1.1, 2.2.2.2, 3.3.3.3 ]
            to_ip_addresses: [ 5.5.5.5, 6.6.6.6, 7.7.7.7, 8.8.8.8 ]
          - proto_ports:
              - protocol: tcp
                ports: '22'
            action: permit
            from_ip_addresses:
              - 10.10.10.10
            to_ip_addresses:
              - 20.20.20.20
          - proto_ports:
              - protocol: udp
                ports: '169'
              - protocol: tcp
                ports: '80-92,1010-1100'
              - protocol: tcp
```

```

        ports: '9000,9100,9200'
    action: deny
    from_ip_addresses:
        - 12.12.12.12
        - 13.13.13.13
        - 14.14.14.14
    to_ip_addresses:
        - 42.24.42.24

```

## REST API

The PSM REST API documentation can be accessed at `https://$PSMaddr/docs` , where `$PSMaddr` corresponds to the PSM cluster address at your site.

A sample Postman collection is available at :

`https://$PSMaddr/docs/examples/Sample.postman_collection.json`

The PSM Swagger specification is available at:

`https://$PSMaddr/docs/generated/swaggeruri.html`

## Tech Support Collection

The Tech Support feature collects various logs and troubleshooting information needed by technical support teams in case of an issue.

Tech support files are collected for both the PSM and the specified CX 10000 DSSes (including their DSMs).

Individual PSM nodes and DSSes can be selected. Technical Support will provide information on what components of the PSM/DSS fabric (including selecting specific DSSes) may need to be collected.

Tech Support Requests can be created in the PSM views shown below. First, navigate to Admin -> Tech Support, and click on **ADD TECH-SUPPORT REQUEST**.

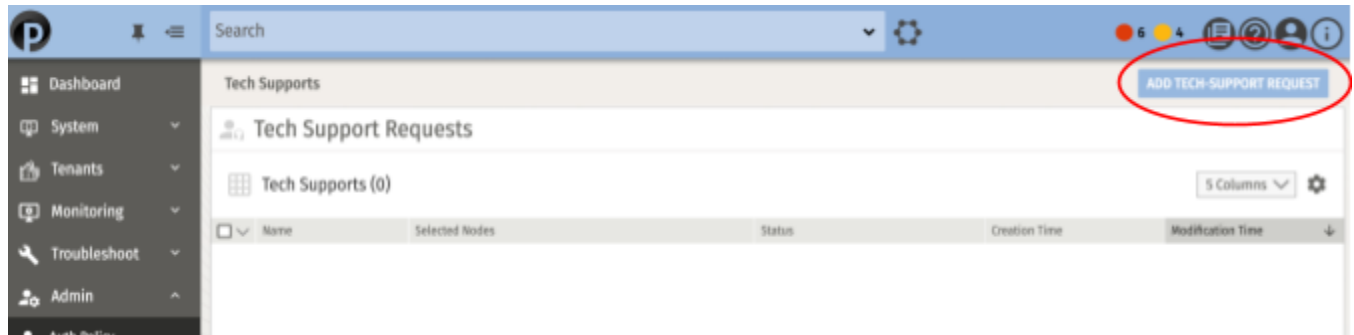


Figure 116. Tech Support screen

Specify a name for this request. A list of available PSM nodes and a list of available DSSes are shown; a search mechanism is provided to quickly find the specific switches to be included.

Once you have selected the nodes and switches to be included, click **CREATE TECH SUPPORT REQUEST**. It may take some time for the creation process to complete.



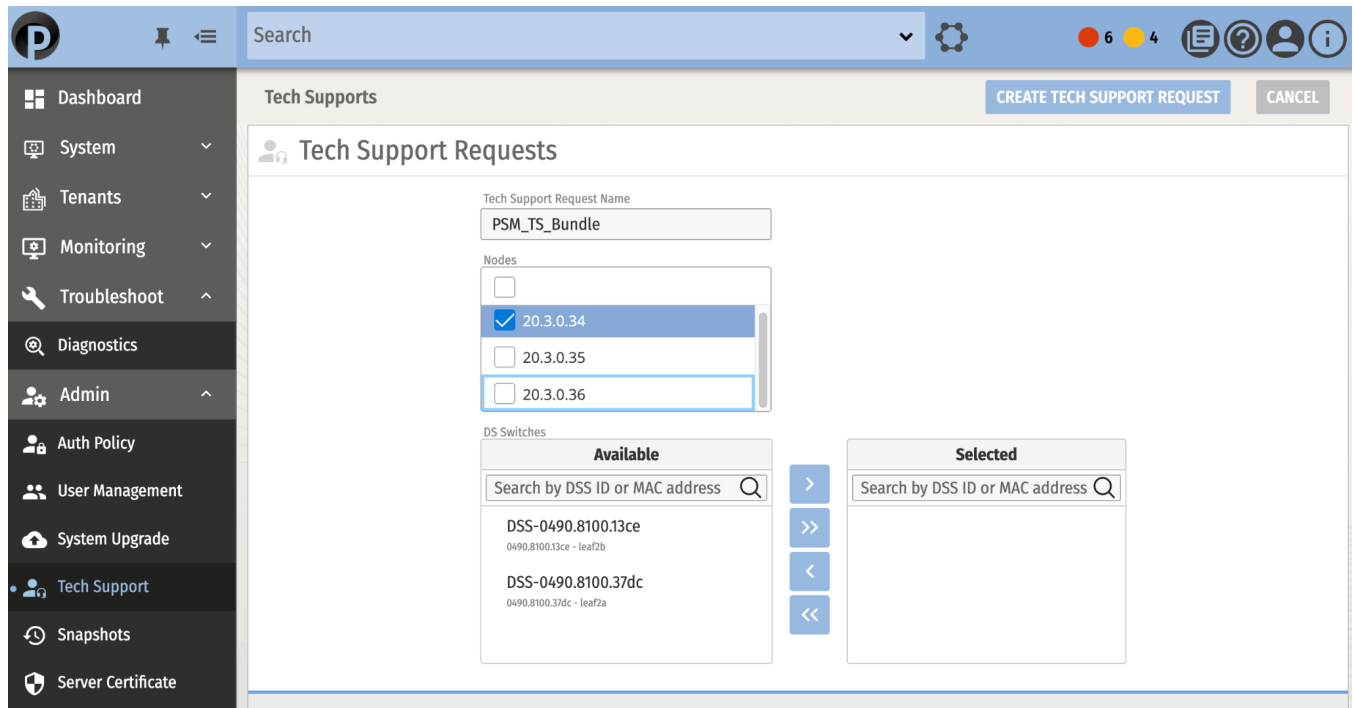


Figure 117. PSM Tech Support Requests screen

With the tech-support request completed, the new request should appear on the list of tech supports. Click on "Download all files". The zip file will be downloaded to the default location of your browser.

## Appendix A: PSM Quorum High Availability

The management plane is the communication channel between DSSes and the PSM cluster (gRPC over HTTPS). This traffic includes configuration, metrics, and logs. It is functionally distinguished from the data plane: the actual network traffic being managed by the DSS environment.

The data plane, including all active stateful services, is never affected by any full or partial PSM outage. In addition, since there are three PSM VMs in a cluster, the management plane remains fully operational unless two or more PSM VMs are down or unreachable.

Table 11 describes the various PSM failure conditions and how the PSM and DSSes can tolerate different failure types. It explains the various failure conditions and the expected impact to the management (configuration, logs) and data plane (traffic forwarding) from the DSS perspective.

Scenario	Management Plane Impact	Data Plane Impact
Management plane is down between a DSS and the PSM	Configuration changes are accepted by the PSM, but will not take effect until the DSS and PSM can communicate. Metrics and logs will be buffered at the DSS as local memory availability permits	No impact to stateful services Configuration changes are not possible
One PSM node fails or becomes unreachable	Configuration changes are accepted by the PSM and applied to DSSes Metrics and logs are regularly collected	No impact to stateful services No impact on configuration changes
Two PSM nodes fail or become unreachable	The PSM does not allow new user logins and does not accept configuration changes, to avoid a “split brain” scenario. Metrics and logs are not collected by the PSM, but will be buffered at the DSS subject to local storage availability.	No impact to stateful services Configuration changes are not possible

*Table 11. PSM failure conditions and impacts*

## Appendix B: PSM Operational Network Ports

These ports must be opened in each direction in order for the PSM cluster to function correctly.

TCP Port	Service
<b>From user station to PSM node</b>	
22	sshd (for node management)
80	redirects to 443
443	ApiGw HTTPS
9001	Initial POST for cluster bootstrap
<b>From PSM node to PSM node</b>	
5001	etcd (peer)
5002	etcd (client)
6443	Kubernetes APIServer
7000	Citadel
7087	Citadel Query
9002	Cluster management
9003	ApiServer
9004	Orchestrator Hub
9009	Resolver

*Table 12. Required port availability (part 1/3)*

TCP Port	Service
<b>From PSM node to PSM node (cont'd)</b>	
9010	EventsManager
9011	Spyglass
9012	EventsProxy
9014	CMD Leader Services
9015	Rollout
9020	TPM
9030	TSM
9051	VOS
9200	Elastic (client)
9300	Elastic (peer)
10250	Kubelet
10257	Kubernetes Controller Manager
10259	Kubernetes Scheduler
10777	Citadel Collector
19001	Minio

*Table 12. Required port availability (part 2/3)*

TCP Port	Service
<b>From PSM to DSS and from management station to PSM</b>	
8888	agent reverse proxy for <code>pentl</code> and diagnostics
<b>From DSS to PSM</b>	
9005	NPM
9009	Resolver
9010	EventsManager
9012	EventsProxy
9014	CMD Health Updates
9015	Rollout
9019	NIC Registration
9020	TPM
9030	TSM
9051	VOS
10777	Citadel Collector

*Table 12. Required port availability (part 3/3)*

## Appendix C: Configuring Microsegmentation in Non-AFC Environments

### Topology

Figure 118 shows an example topology for PVLAN setup.

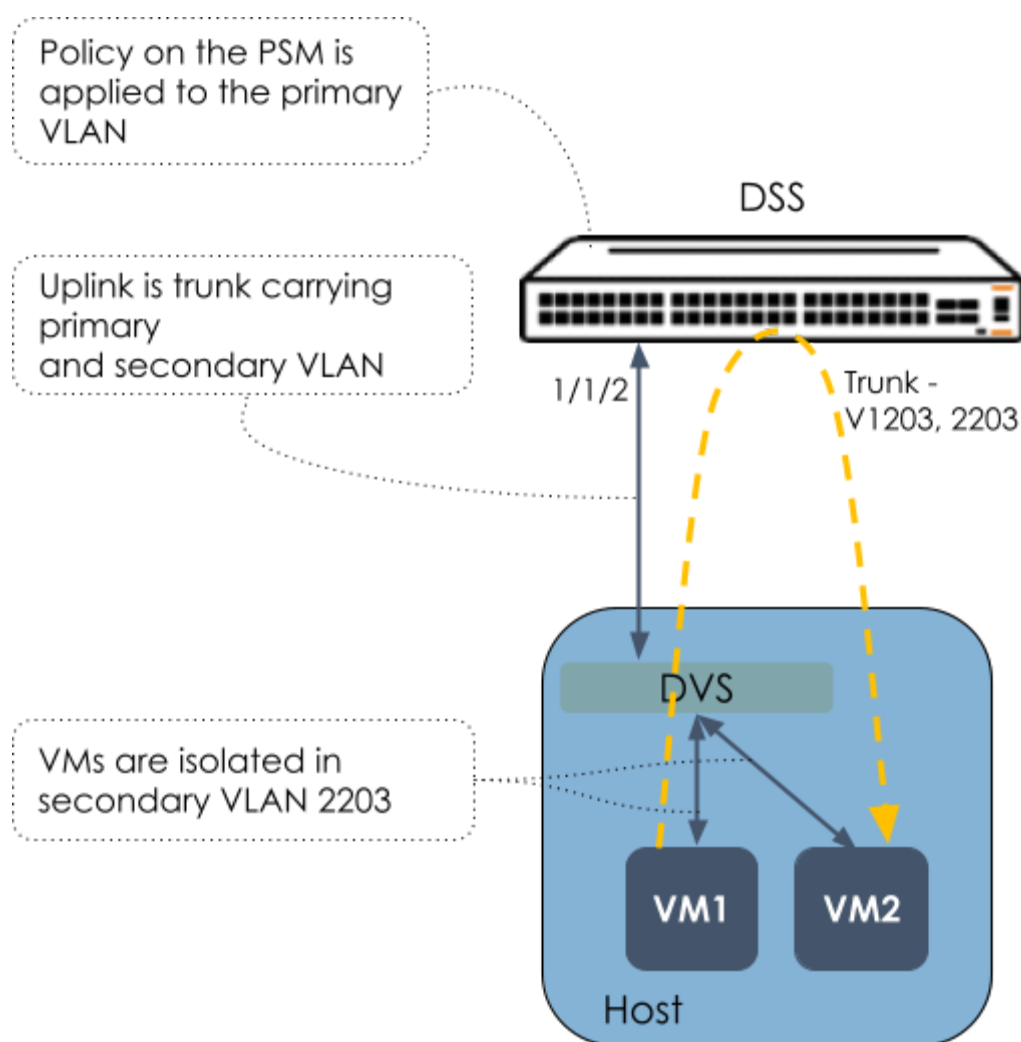


Figure 118. PVLAN setup and configuration between VMs

## Configuration on the DSS

In the topology in Figure 118, traffic between VM1 and VM2 would normally be switched using the DVS. By using PVLANS, policy can be enforced on the traffic between VM1 and VM2. This stateful policy enforcement is done on the DSS.

Consider 2 VMs, VM1 and VM2 which are part of the same DVS. To enable stateful policy enforcement for traffic between VM1 and VM2:

1. Configure PVLAN on the ESXi server
2. Configure PVLAN on AOS-CX
3. Configure Security Policy for the primary VLAN on the PSM

The example below shows VLAN 1203 as primary VLAN and VLAN 2203 as secondary/isolated VLAN in AOS-CX.

### Global Config Mapping the Primary and Secondary VLAN

```
vlan 1203
private-vlan primary
vlan 2203
private-vlan isolated primary-vlan 1203
```

### Host-Facing Interface Configured as Regular Trunk, Allowing Both Primary and Secondary VLAN

```
interface 1/1/2
no shutdown
persona access
mtu 9198
no routing
vlan trunk native 1
vlan trunk allowed 1,1203,2203
```

## SVI Config on Primary with Local Proxy ARP

```
# show running-config interface vlan 1203
interface vlan1203
ip address 10.6.203.2/24
active-gateway ip mac 00:00:00:00:01:00
active-gateway ip 10.6.203.1
ip mtu 9198
ip local-proxy-arp
Exit
```

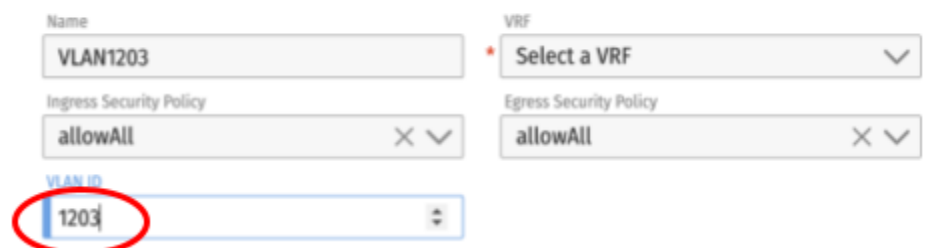
## Configuring VMware (ESXi)

Refer to VMware documentation for the appropriate procedures for configuring a PVLAN.

## Configuration on the PSM

On the PSM, create the primary private VLAN and apply the appropriate policy to the particular object (ingress and/or egress):

### Networks Overview



The screenshot shows the 'Networks Overview' configuration page. It includes the following fields:

- Name:** VLAN1203
- VRF:** Select a VRF (dropdown menu)
- Ingress Security Policy:** allowAll (dropdown menu)
- Egress Security Policy:** allowAll (dropdown menu)
- VLAN ID:** 1203 (dropdown menu, circled in red)

Figure 119. Setting up the VLAN



## Appendix D: Saving the PSM Recovery Key

When a PSM cluster is created, a private/public key pair and corresponding self-signed certificate is generated (the cluster credentials). These credentials are used to secure communication for different PSM and DSS functions. Whenever a failed PSM node is replaced or a new PSM node is added to the cluster, these credentials are automatically passed along to the new node.

In a catastrophic situation where all of the PSM nodes are lost and a new cluster must be created, it will have new credentials, and DSSes that were admitted to the old cluster will not be able to connect to it, requiring all DSSes to be decommissioned and readmitted, even if the cluster is reloaded with a snapshot of the old cluster.

To simplify recovery from this type of loss, the PSM allows for a one-time download of its credentials (its recovery key), which can then be used with `bootstrap_PSM.py` when the replacement cluster is created.

After a first-time installation of a cluster, a privileged user may use either the PSM browser interface or a REST request to save the credentials to a file. Once this is done, the new cluster can be created using the original credentials, allowing the existing DSSes to remain admitted.

### Saving the Key

From the System menu, go to the Cluster view. The “Download Cluster Recovery Key” icon will be present if the key has not already been downloaded; select this option, and save the file in a secure location.

**NOTE:** Once the key is obtained, the icon will no longer appear in the UI.

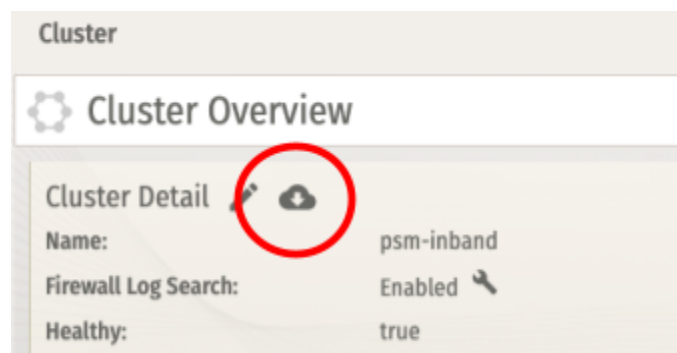


Figure 120. System/Cluster screen, showing the position of the download button

Alternatively, the key may be obtained from a REST GET call to the `sysruntime/v1/cluster/recoverykeys` endpoint. After either method is used, the icon will no longer appear in the UI, and a REST call to the endpoint will fail.

**Note:** *The recovery key should be stored on a system other than those hosting the PSM.*

## Recovering the Cluster

To restore the cluster, three components are needed:

1. The recovery key file
2. The IP addresses previously used for the cluster
3. A snapshot of the cluster configuration

Follow the same procedure used to initially install the cluster, as described above, with one additional element: when running `bootstrap_PSM.py`, include the option

```
-recovery_keys RECOVERY_KEYS
```

where `RECOVERY_KEYS` is replaced by the path to the file containing the saved credentials.

Note that the new cluster must use the same IP addresses as the old cluster. Restore the cluster snapshot to finish the recovery.

# Appendix E: Using the PSM Network Graph to Create Security Policies

In addition to manually adding predetermined policies via table view in the UI or via the equivalent API call, the PSM UI *network graph* feature can be used instead, which aids in discovering what flows should be allowed and designing a zero trust network policy to only allow those flows. This Appendix describes two approaches to using the network graph.

## First Method

The Network Graph method helps discover the relationships between endpoints within a VRF, within and between VLANs, and within or between endpoint groups. The relationships help better understand common connection occurrences within the data center that should be permitted. Through an iterative process of analyzing the behavior and connections of your endpoints and applications over a time domain, the network graph method will help discover a zero trust network policy.

Navigate to `Tenants → Security Policy`; click on the drop-down menu, which by default has `Table View` selected, and change it to `Network Graph`.



*Figure 121. Network Security Policy screen, showing the policy selection method drop-down menu; change to `Network Graph`.*

Use the VRF drop-down menu to select the VRF whose firewall logs you wish to evaluate. Select `Past hour` or `Past day` as the range of data to evaluate.

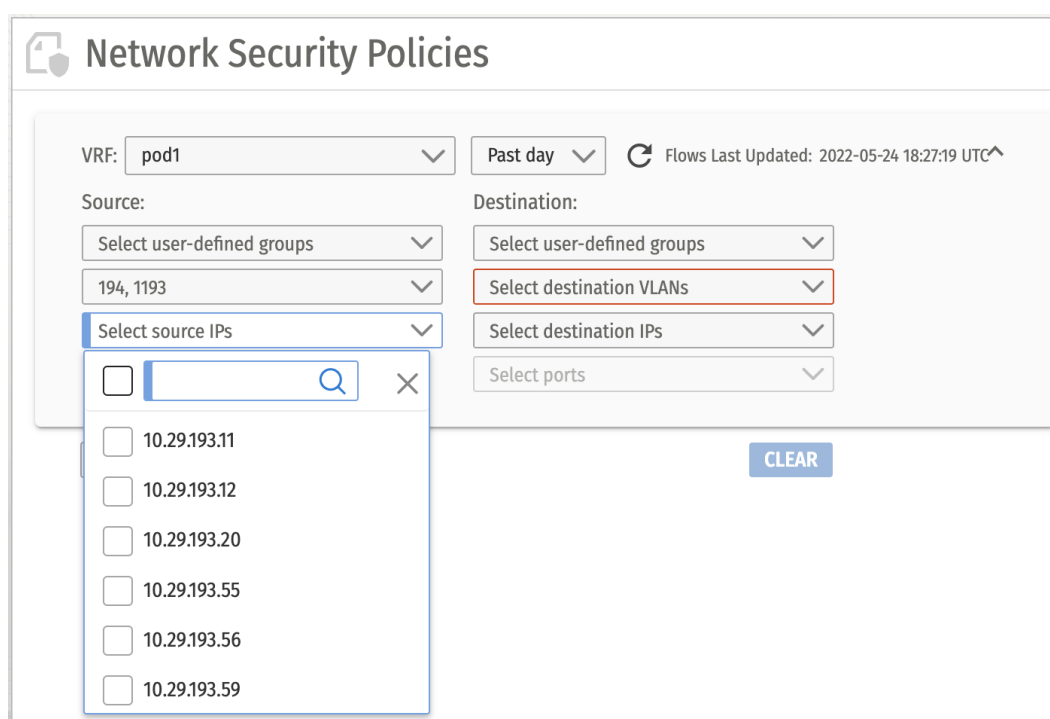
## Network Security Policies



VRF: `pod1` ▼ `Past day` ▼ ↻

*Figure 122. Choose the VRF and the range of data to evaluate.*

The firewall logs can be filtered even more narrowly by selecting VLAN, IP, port and policy, as shown in Figure 123:



**Network Security Policies**

VRF: `pod1` ▼ `Past day` ▼ ↻ Flows Last Updated: 2022-05-24 18:27:19 UTC ^

Source: Select user-defined groups ▼ 194, 1193 ▼ Select source IPs ▼

Destination: Select user-defined groups ▼ Select destination VLANs ▼ Select destination IPs ▼ Select ports ▼

CLEAR

☐  Q X

- ☐ 10.29.193.11
- ☐ 10.29.193.12
- ☐ 10.29.193.20
- ☐ 10.29.193.55
- ☐ 10.29.193.56
- ☐ 10.29.193.59

*Figure 123. Restrict the flows evaluated by user group, VLAN, IP, port and policy*

Select `policy/rules to filter` to see that the flows for the VRF selected currently have no flows protected by any policy/rules.

Select policy/rules to filter

CLEAR

☐

☐ Flows not covered by any rules 1669 flows

☐ Pol1 - 1: Pol1\_Protoicmp\_from\_s1 0 flows

☐ Pol1 - 2: Pol1\_Protoudp\_from\_s1 0 flows

☐ Pol1 - 3: Pol1\_Prototcp\_from\_s1 0 flows

☐ Pol1 - 4: Pol1\_Protoicmp\_prefixRoute 0 flows

☐ Pol1 - 5: Pol1\_Protoudp\_prefixRoute 0 flows

Figure 124. Check for any flows protected by policy/rules

On the right of the screen, NetSec Summary will display flows captured for the VRF, time frame, and any other filters selected:

NetSec Summary

Unique Flows

Flows

Policies

User-Defined Groups

Unique Flows (27)

Search

Aa Ab R+ Q

<input type="checkbox"/> <div></div>	Source IP	Destination IP	Destination Port	Protocol	#
<input type="checkbox"/>	10.29.193.20	10.29.5.7	53	udp	1328
<input type="checkbox"/>	10.29.193.20	8.8.8.8	53	udp	1354
<input type="checkbox"/>	10.29.193.170	10.29.5.7	53	tcp	449
<input type="checkbox"/>	10.29.193.59	10.29.5.7	53	tcp	451
<input type="checkbox"/>	10.29.193.55	10.29.5.7	53	tcp	444
<input type="checkbox"/>	10.29.193.56	10.29.5.7	53	tcp	444
<input type="checkbox"/>	10.29.193.170	10.29.193.55	2048	icmp	38
<input type="checkbox"/>	10.29.193.170	10.29.193.55	5666	tcp	75

25

1

2

Showing 1 to 25 of 27 records

Figure 125. NetSec Summary shows the flows captured.

Sort by unique flows by clicking on the # column to filter the most commonly occurring flows from the top down; next, sort by clicking on the destination port column to filter by service and include the number of unique occurrences. The approach allows a focus on building policy based on a destination service and frequency of connections.

NetSec Summary

Unique Flows

Flows

Policies

User-Defined Groups

Unique Flows (22)

Search

Aa Ab R\* Q

<input type="checkbox"/> <input type="checkbox"/>	Source IP	Destination IP	Destination Port ↑	Protocol	#
<input type="checkbox"/>	10.29.193.20	10.29.5.7	53	udp	1389
<input type="checkbox"/>	10.29.193.20	8.8.8.8	53	udp	1434
<input type="checkbox"/>	10.29.193.55	10.29.5.7	53	tcp	432
<input type="checkbox"/>	10.29.193.170	10.29.5.7	53	tcp	432
<input type="checkbox"/>	10.29.193.59	10.29.5.7	53	tcp	435
<input type="checkbox"/>	10.29.193.56	10.29.5.7	53	tcp	431
<input type="checkbox"/>	10.29.193.56	10.29.5.7	53	udp	18
<input type="checkbox"/>	10.29.193.12	10.29.5.7	53	udp	12

25

⏮

⏪

1

⏩

⏭

Showing 1 to 22 of 22 records

Figure 126. The summary list can be sorted by selecting columns.

Optionally, instead of sorting by column you use the search bar to filter for a specific service or IP. The search bar can be manually entered or auto-populated by right-clicking a row of interest from the unique flow table and selecting your search criteria.

NetSec Summary

Unique Flows

Flows

Policies

User-Defined Groups

Unique Flows (22)

Search

Aa Ab R+ Q

<input type="checkbox"/>	Source IP	Destination IP	Destination Port ↑	Protocol	#
<input type="checkbox"/>	10.29.193.20	8.8.8.8	53	udp	1270
<input type="checkbox"/>	10.29.193.20	10.29.5.7	53	udp	1238
<input type="checkbox"/>	10.29.193.59	10.29.5.7	53	tcp	493
<input type="checkbox"/>	10.29.193.55	10.29			493
<input type="checkbox"/>	10.29.193.170	10.29			492
<input type="checkbox"/>	10.29.193.56	10.29.5.7	53		488
<input type="checkbox"/>	10.29.193.56	10.29.5.7	53		
<input type="checkbox"/>	10.29.193.170	10.29.5.7	53		

25

⏪

⏴

1

⏵

⏩

Showing 1 to 22 of 22 records

Figure 127. Use the Search bar to filter flows



Select the filtered flows to build rules for by clicking the left column box:

**NetSec Summary**

Unique Flows | Flows | Policies | User-Defined Groups

Unique Flows (22) **+**

Search Aa Ab R+ Q

<input type="checkbox"/>	Source IP	Destination IP	Destination Port ↑	Protocol	#
<input checked="" type="checkbox"/>	10.29.193.20	8.8.8.8	53	udp	1485
<input checked="" type="checkbox"/>	10.29.193.20	10.29.5.7	53	udp	1432
<input checked="" type="checkbox"/>	10.29.193.170	10.29.5.7	53	tcp	419
<input checked="" type="checkbox"/>	10.29.193.55	10.29.5.7	53	tcp	416
<input checked="" type="checkbox"/>	10.29.193.59	10.29.5.7	53	tcp	414
<input checked="" type="checkbox"/>	10.29.193.56	10.29.5.7	53	tcp	410
<input checked="" type="checkbox"/>	10.29.193.12	10.29.5.7	53	udp	12
<input checked="" type="checkbox"/>	10.29.193.12	8.8.8.8	53	udp	12

25 1 Showing 1 to 22 of 22 records

Figure 128. Check the filtered flows to build rules for, and click “Add to policy”.

Click “Add to policy” (the + sign, as shown in Figure 128). Based on the flows selected, the predefined source, destination, ports and protocols are pre-populated to form a rule, which can be added to an existing or new policy. Provide a policy and rule name. If a new policy, the policy is not attached to a network or VRF automatically. Remember to click save.

NetSec Policy Editing

Policy Name

pod1

SAVE

CANCEL

Rule Order

1

Action

Permit

Rule Name

pod1\_DNS\_Acce

Disabled

Description

Source IP Addresses

10.29.193.20

10.29.193.55

10.29.193.170

10.29.193.59

10.29.193.56

10.29.193.12

10.29.193.11

Destination IP Addresses

10.29.5.7

8.8.8.8

☒ PROTO-PORTS
 ☐ APPS

Protocol

udp

Ports

53

Protocol

tcp

Ports

53

+ Protocols and Ports

Preview

Number	Rule Name	Source IPs	Destination IPs	Action	Protocol Port	Applications	Description	Status
1	pod1_DNS_Access	10.29.193.20, 10.29.193.55, 10.29.193.170, 10.29.193.59, 10.29.193.56, 10.29.193.12, 10.29.193.11	10.29.5.7, 8.8.8.8	Permit	udp/53, tcp/53			Enabled

Figure 129. Rule table populated with objects from the flow table selected

To use this method for policy discovery it is essential to add a “permit any any” rule, to collect which flows would have been denied. Using a “permit any any” rule at the end of the policy assists with policy evaluation while iterating through the discovery of new flows and evaluating new time dimensions of the firewall logs within the VRF and other selected filters.

Select the Policies tab, search for your policy name, and click `Insert Rule After`, as shown in Figure 130. Enter a rule as shown in Figure 131, to catch all traffic that did not match rule 1.

PSM for Aruba CX 10000 User Guide

NetSec Policy Editing

Policy Name

pod1

SAVE

CANCEL

Rule Order

2

Action

Permit

Rule Name

pod1-catch-all

Disabled

☐

Description

policy-discovery-delete

Source IP Addresses

any

Destination IP Addresses

any

☒ PROTO-PORTS
 ☐ APPS

Protocol

any

Ports

+ Protocols and Ports

Preview

Number	Rule Name	Source IPs	Destination IPs	Action	Protocol Port	Applications	Description	Status
1	pod1_DNS_Access	10.29.193.20, 10.29.193.5	10.29.5.7, 8.8.8.8	Permit	udp/53			Enabled
2	pod1-catch-all	any	any	Permit	any		policy-discovery-delete	Enabled

Figure 131. Creation of the catch-all rule for policy discovery. Remember to click Save.

Remember to click **SUBMIT CHANGES** in the top right of the screen.

Network Graph

CANCEL CHANGES

SUBMIT CHANGES

Figure 132. Click SUBMIT CHANGES

At the next time domain of collected firewall logs for the VRF policy discovery, follow the steps below to filter flows to analyze. Click on `Select policy/rules` to filter the policy/rules to analyze what flows match the specified permit rules and what flows match the catch-all “permit any” rule. Any such flows that did not match a specific permit rule and matched the catch-all rule require further analysis to determine if these flows should be permitted with a specific rule. Through several iterations this process will help create a network zero trust policy for a micro, macro, or zone based segmentation strategy.

The example in Figure 133 shows that 24 hours later, 564 new flows are net new and matched the catch-all rule. These flows require evaluation to determine whether they should be considered permitted flows. The NetSec Summary, on the right of Figure 133, displays the unique flows that would have been implicitly denied if not for the catch-all permit. Iterate

through this process until you are satisfied that a zero trust implicit deny at the end of the policy will not block any needed flows.

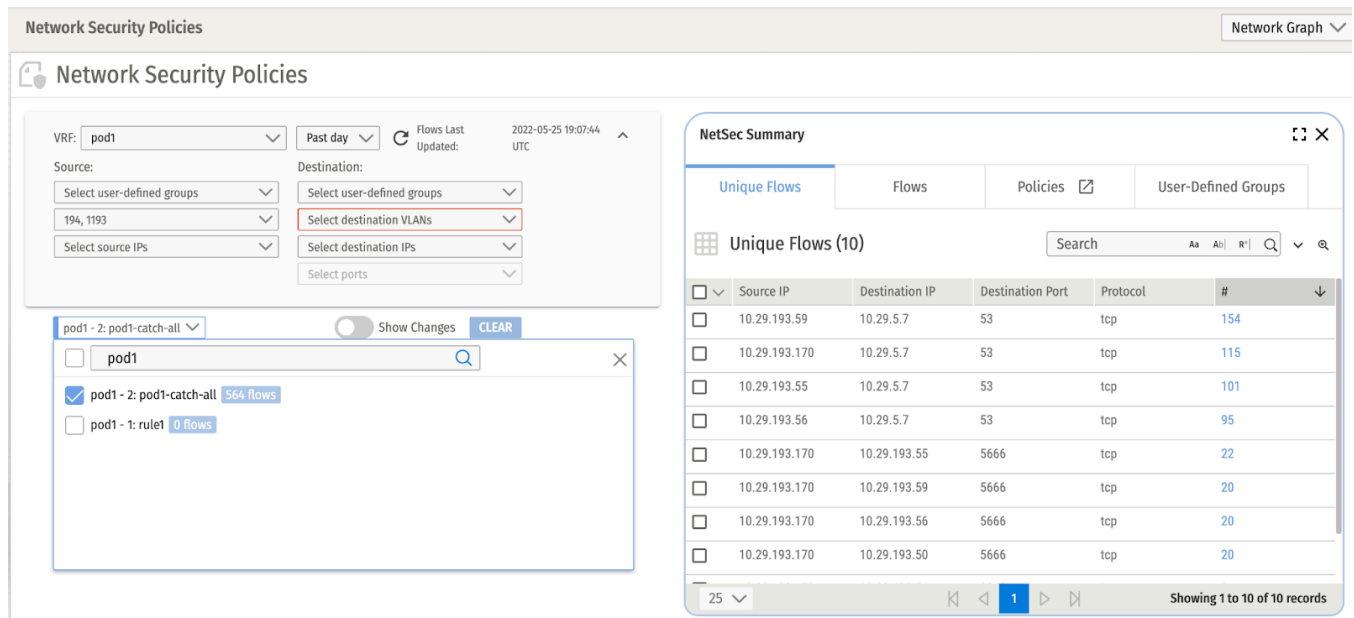


Figure 133. Analysis of catch-all rule and related flows and occurrences

## Second Method

This Network Graph method helps discover the relationships between endpoints within a VRF, within and between VLANs, and within or between user-defined groups. End users who prefer to evaluate connections and relationships rendered as a graph and build policy will choose this approach. For example, you can manually create user-defined groups to analyze their application's flows (Eg: Web->App, App->DB, etc.)

Navigate to `Tenants → Security Policy`; click on the drop-down menu, which by default has `Table View` selected, and change it to `Network Graph`.



Figure 134. Network Security Policy screen, showing the policy selection method drop-down menu; change to Network Graph.

Use the VRF drop-down menu to select the VRF whose firewall logs you wish to evaluate. Select `Past hour` or `Past day` as the range of data to evaluate.

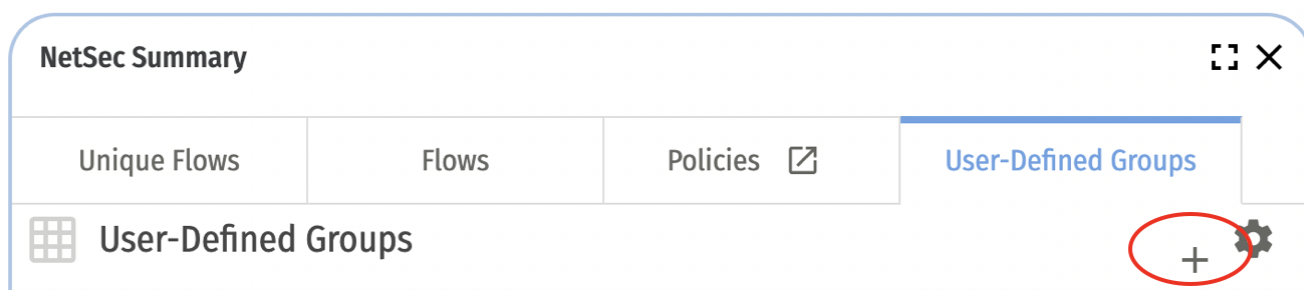
## Network Security Policies



VRF:     

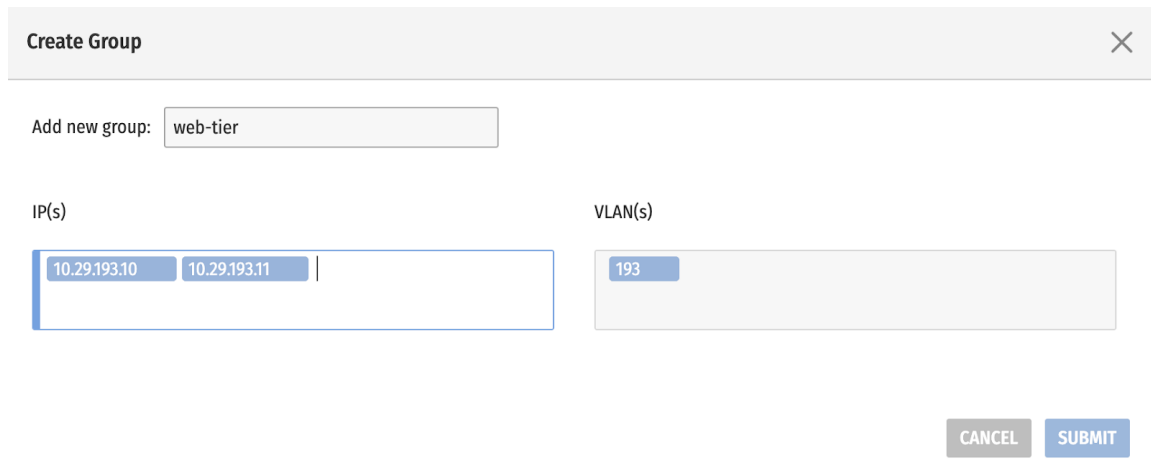
*Figure 135. Choose the VRF and the range of data to evaluate.*

*User-defined groups* (UDGs) can be manually created by clicking on the `User-Defined Groups` tab and manually defining each group based on known IPs and/or VLANs. Click the “Add to group” button (the “+” sign highlighted in Figure 136); the Create Group popup will appear as seen in Figure 137.



*Figure 136. The User-Defined Groups tab, as shown on the right side of the Network Security Policies screen*

Example of a UDG based on IP addresses and/or VLAN:



The 'Create Group' popup window contains the following elements:

- Title Bar:** 'Create Group' with a close button (X).
- Add new group:** A text input field containing 'web-tier'.
- IP(s):** A list of IP addresses: 10.29.193.10, 10.29.193.11.
- VLAN(s):** A list of VLAN IDs: 193.
- Buttons:** 'CANCEL' and 'SUBMIT' buttons at the bottom right.

Figure 137. Create Group popup

Optionally, known groups can be uploaded from a JSON file. Select the gear icon, and choose “Upload Groups”:

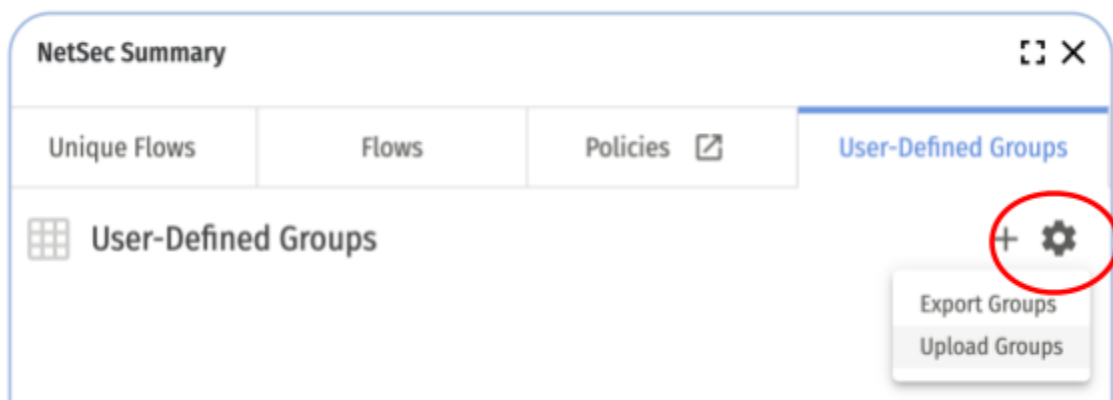


Figure 138. The gear icon accesses the Export/Upload Groups menu.

Exam`ple file:

```
{
  "App-Tier": [
    {
      "type": "VLAN",
      "value": "193"
    },
    {
      "type": "IP",
      "value": "10.29.193.56"
    }
  ]
}
```

A final option for creating a UDG is to select the `Flows` tab, right-click on a row, and select `Add to Group` (which can be new or existing):



NetSec Summary

Unique Flows

Flows

Policies

User-Defined Groups

Unique Flows (24)

Search

Aa Ab R\* Q

<input type="checkbox"/> <span>▼</span>	Source IP	Destination IP	Destination Port	Protocol	#	<span>↓</span>
<input type="checkbox"/>	10.29.193.20	8.8.8.8	53	udp	1245	
<input type="checkbox"/>	10.29.193.20	10.29.5.7	53	udp	1235	
<input type="checkbox"/>	10.29.193.170	10.29.5.7	53	tcp	497	
<input type="checkbox"/>	10.29.193.56	10.29.5.7	53	tcp	492	
<input type="checkbox"/>	10.29.193.55		53	tcp	492	
<input type="checkbox"/>	10.29.193.59		53	tcp	490	
<input type="checkbox"/>	10.29.193.170	10.29.193.55	5666	tcp	90	
<input type="checkbox"/>	10.29.193.170	10.29.193.56	5666	tcp	84	

25 ▼

⏮
⏪
1
⏩
⏭

Showing 1 to 24 of 24 records

*Figure 139. Creating/Updating a user-defined group by right-clicking on a row and selecting Add to Group*

After creating user-defined groups, filter the flows under analysis by clicking on **User-Defined Groups** for the desired source and destination. The result is a rendered graph of the relationships between the groups of endpoints, as shown in Figure 140:

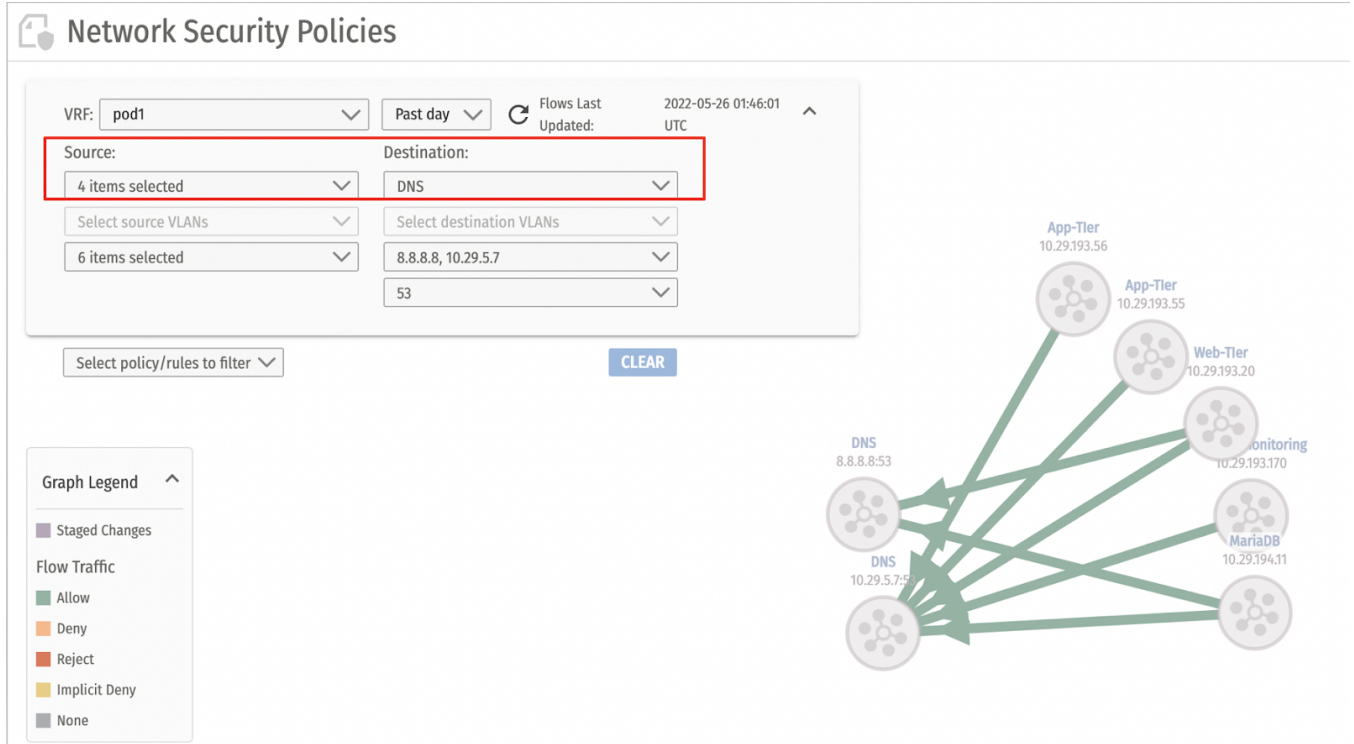
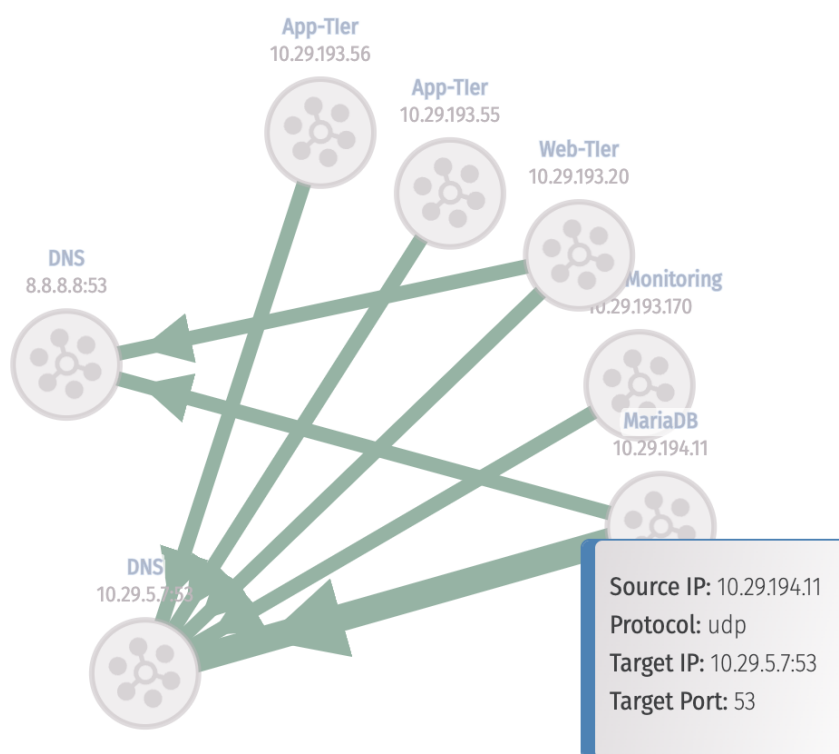


Figure 140. Selecting user-defined groups to render a graph representing UDG's relationships

Hover over a directional line to highlight the 4-tuple information of the connection as shown in Figure 141:



*Figure 141. Hovering over directional line between MariaDB and DNS displays the 4-tuple information for the connection*

To create rules based on the connections rendered in the graph, click on the directional lines of interest and right click on any of the directional lines selected as shown in Figure 142. The result is the connections are added as a new rule, as shown in Figure 143. The predefined source, destination, ports and protocols are pre-populated to form a rule, which can be added to an existing or new policy. Provide a policy and rule name. If a new policy, the policy is not attached to a network or VRF automatically. Remember to click save.




Figure 142. Click on the green directional lines to highlight and right click on a highlighted line to add to policy

**Add Flows to NetSec Policy**
✖

☐ Add to existing policy  
☒ Add to new policy

Policy Name

pod1

Rule Order

1

Action

Permit ▼

Rule Name

app-1

Disabled

☐

Description

application-1

Source IP Addresses

10.29.193.56

10.29.193.20

10.29.194.11

10.29.193.55

10.29.193.170

Destination IP Addresses

10.29.5.7

8.8.8.8

☒ PROTO-PORTS

☐ APPS

Protocol

tcp

Ports

53

Protocol

udp

Ports

53

+
Protocols and Ports

Preview

Number	Rule Name	Source IPs	Destination IPs	Action	Protocol Port	Applications	Description	Status
1	app-1	10.29.193.56, 10.29.193.20, 10.29.194.11, 10.29.193.55, 10.29.193.170	10.29.5.7, 8.8.8.8	Permit	tcp/53, udp/53		application-1	Enabled

*Figure 143. The highlighted connections (Figure 142) are pre-populated as a rule. Provide a name, description, select a policy and click SAVE.*

To continue with policy discovery, it's also essential to add a “permit any any” to identify what flows would have been denied. Using a “permit any any” rule at the end of the policy assists with policy evaluation while iterating through the discovery of new flows and evaluating new time dimensions of the firewall logs within the VRF and other selected filters.

Select the Policies tab, search for the policy name, and click `Insert Rule After` as shown in figure 144. Enter a rule as shown below to catch-all traffic that didn't match rule 1.

PSM for Aruba CX 10000 User Guide

153

NetSec Summary

Unique Flows

Flows

Policies

User-Defined Groups

pod1b

Pod1B

<input type="checkbox"/>	#	Name	Source IPs	Destination IPs	Proto/Ports/Apps	Action
<input type="checkbox"/>	1	app-1	10.29.193.20,10.29.193.11,10.29.193.56,10.29.193.12,10.29.193.170,10.29.194.11,10.29.193.55	10.29.5.7,8.8.8.8	Protocol / Ports: udp/53, tcp/53	permit

Insert Rule After

Figure 144. Filter for policy name and Insert a catch-all rule after specific permit rules

Enter a rule as shown in Figure 145 to catch all traffic that didn't match rule 1. Remember to click SAVE.

NetSec Policy Editing

Policy Name

Pod1B

SAVE

CANCEL

Rule Order

2

Action

Permit

Rule Name

Pod1b-Catch-All

Disabled

☐

Description

Policy-Discovery-Delete

Ready to submit

Source IP Addresses

any

Destination IP Addresses

any

☒ PROTO-PORTS
 ☐ APPS

Protocol

any

Ports

+ Protocols and Ports

Preview

Number	Rule Name	Source IPs	Destination IPs	Action	Protocol Port	Applications	Description	Status
1	app-1	10.29.193.20, 10.29.193.1	10.29.5.7, 8.8.8.8	Permit	udp/53, tcp/53		application-1	Enabled
2	Pod1b-Catch-All	any	any	Permit	any		Policy-Discovery-Delete	Enabled

Figure 145. Create a “permit any any” rule as a catch-all rule to assist with policy discovery.

Remember to click **SUBMIT CHANGES** in the top right of the screen.

Network Graph

CANCEL CHANGES

SUBMIT CHANGES

Figure 146. SUBMIT CHANGES button

At the next time domain of collected firewall logs for the VRF discovering policy, follow the steps below to filter flows to analyze. Click on Select policy/rules to filter the policy/rules to analyze what flows match the specified permit rules and what flows match the catch-all “permit any” rule. Any such flows that did not match a specific permit rule and matched the catch-all rule require further analysis to determine if these flows should be permitted with a specific rule. Through several iterations this process will help create a network zero trust policy for a micro, macro, or zone based segmentation strategy. In the example in Figure 147, one hour later, two new unique flows are net new and matched on the catch-all. The catch-all flows must be evaluated whether they should be permitted flows or not. The rendered graph displays the unique flows that would have been implicitly denied if not for the catch-all permit. Iterate through this process until you are satisfied with a zero trust implicit deny at the end of the policy.

**Network Security Policies**

**Network Security Policies**

VRP: **pod1** **Past hour** **Flows Last Updated:** 2022-05-26 03:38:12 UTC

Source: **4 items selected** **Destination:** **DNS**

Select source VLANs: **5 items selected** Select destination VLANs: **8.8.8.8, 10.29.5.7**

4 items selected

**Pod1B - 2: Pod1b-Catch-All** **CLEAR**

☐ **pod1b** **pod1b**

☒ **Pod1B - 2: Pod1b-Catch-All** **2 flows**

☐ **Pod1B - 1: app-1** **0 flows**

**Gr**

**S**

**Flw**

**A**

**C**

**R**

**Implicit Deny**

**Web-Tier** 10.29.193.11

**DNS** 10.29.5.7:53

**Nagios-Monitoring** 10.29.193.170

**DNS** 10.29.5.7:52

**Source IP:** 10.29.193.170  
**Protocol:** tcp  
**Target IP:** 10.29.5.7:53  
**Target Port:** 53

*Figure 147. Select policy/rules to filter and search for a policy name.  
Click on the catch-all rule to render a graph showing new flows discovered.*